



Strategic Manipulation in Social Choice Theory

Patrick Lederer

Vollständiger Abdruck der von der TUM School of Computation, Information, and Technology der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz:

Prof. Dr. Jens Großklags

Prüfende der Dissertation:

1. Prof. Dr. Felix Brandt
2. Prof. Dr. Ulle Endriss
3. Prof. Dr. Arunava Sen

Die Dissertation wurde am 25.09.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information, and Technology am 31.01.2024 angenommen.

STRATEGIC MANIPULATION IN SOCIAL CHOICE THEORY

PATRICK LEDERER

Patrick Lederer: *Strategic Manipulation in Social Choice Theory* © September, 2023.

E-MAIL:

ledererp@in.tum.de

This thesis was created based on a template by Hans Georg Seedig. It was typeset using \LaTeX and the ClassicThesis style by André Miede, combined with the ArsClassica package by Lorenzo Pantieri and some minor own modifications. The text is set in Palatino with math in Euler, both due to Hermann Zapf. Headlines are set in Iwona by Janusz M. Nowacki, the monospace font is Bera Mono designed by Bitstream, Inc. Most of the graphics were created using TikZ by Till Tantau.

ABSTRACT

Social choice theory studies mechanisms, so-called voting rules, which aggregate the preferences of a group of agents over some alternatives into a group decision. One of the most fundamental properties of such voting rules is *strategyproofness*, which requires that agents should not be able to benefit from lying about their true preferences. Unfortunately, this property is known to be extremely restrictive in social choice theory as Gibbard (1973) and Satterthwaite (1975) have shown that only unattractive voting rules are strategyproof. However, their result only applies to voting rules that always choose a single winner deterministically. In this thesis, we thus analyze whether it is possible to escape this impossibility theorem by modifying its underlying assumptions. In more detail, we study strategyproofness for *social decision schemes* (SDSs), which return lotteries over the alternatives, *social choice correspondences* (SCCs), which return non-empty subsets of the alternatives, and *party-approval committee* (PAC) *voting rules*, which return multisets of the alternatives of fixed size. The idea of both SDSs and SCCs is to ultimately select a single winner: for SDSs, the final winner will be selected by chance according to the chosen lottery, and for SCCs, the final winner will be picked from the choice set by some tie-breaking mechanism. By contrast, PAC voting rules model committee elections, where the seats of a committee are assigned to parties and each party can have multiple seats in the committee.

For all of these models, various notions of strategyproofness can be defined and we derive both positive and negative results. For instance, we show for SDSs that even rather weak strategyproofness notions conflict with a decisiveness criteria called Condorcet-consistency. By contrast, we also prove that there are attractive strategyproof SDSs when restricting the domain of feasible preference profiles by characterizing the set of strategyproof and non-imposing SDSs on the Condorcet domain. For SCCs, we derive similar results and, e.g., characterize an attractive SCC called the top cycle based on a mild strategyproofness notion. Finally, for PAC elections, we demonstrate a variant of the Gibbard-Satterthwaite theorem by proving that no anonymous PAC voting rule satisfies strategyproofness and simultaneously guarantees that the selected committee proportionally represents the voters' preferences. All of these results are unified by the idea of precisely pinpointing when strategyproof social choice is possible.

ZUSAMMENFASSUNG

Sozialwahltheorie untersucht Mechanismen, sogenannte Wahlverfahren, die die Präferenzen einer Gruppe von Agenten über einer Menge von Alternativen zu einer Gruppenentscheidung aggregieren. Eine der grundlegendsten Eigenschaften solcher Wahlverfahren ist *Nicht-manipulierbarkeit*, welche besagt, dass Agenten nicht durch eine strategische Fehldarstellung ihrer wahren Präferenzen profitieren können. Leider ist bekannt, dass diese Eigenschaft in der Sozialwahltheorie sehr restriktiv ist, da Gibbard (1973) und Satterthwaite (1975) gezeigt haben, dass nur unattraktive Wahlverfahren nicht-manipulierbar sind. Dieses Ergebnis gilt jedoch nur für Wahlverfahren, die immer einen einzelnen Gewinner deterministisch auswählen. In dieser Arbeit analysieren wir daher, ob wir dieses Unmöglichkeitstheorem durch die Modifikation der zu Grunde liegenden Annahmen umgehen können. Insbesondere untersuchen wir die Nicht-manipulierbarkeit für *Social Decision Schemes* (SDSs), welche Lotterien über den Alternativen zurückgeben, *Social Choice Correspondences* (SCCs), welche nicht-leere Teilmengen der Alternativen zurückgeben, und *Party-Approval Committee* (PAC) *Voting Rules*, welche Multimengen der Alternativen mit einer bestimmten Größe zurückgeben. Sowohl SDSs als auch SCCs zielen darauf ab, final einen einzelnen Wahlsieger auszuwählen: für SDSs wird der endgültige Gewinner zufällig durch die ausgewählte Lotterie bestimmt, und für SCCs wird der endgültige Gewinner aus der gewählten Menge durch einen zweiten Mechanismus ermittelt. Im Gegensatz dazu modellieren PAC Voting Rules Wahlverfahren, bei denen den Parteien die Sitze eines Komitees zugewiesen werden und jede Partei mehrere Sitze im Komitee erhalten kann.

Für alle diese Modelle können unterschiedliche Varianten von Nicht-manipulierbarkeit definiert werden und wir zeigen sowohl positive als auch negative Ergebnisse. Beispielsweise beweisen wir für SDSs, dass selbst schwache Konzepte der Nicht-manipulierbarkeit im Widerspruch zu einem Kriterium namens Condorcet-Konsistenz stehen. Im Gegensatz dazu beweisen wir auch, dass es attraktive und nicht-manipulierbare SDSs gibt, wenn die Menge der erlaubten Präferenzprofile eingeschränkt ist, indem wir die Menge der nicht-manipulierbaren SDSs, die jede Alternative mit Wahrscheinlichkeit 1 wählen können, auf der Condorcet Domäne charakterisieren. Für SCCs leiten wir ähnliche Ergebnisse her und charakterisieren beispielsweise eine attraktive SCC namens Top Cycle mittels einer schwachen Variante von Nicht-manipulierbarkeit. Schließlich beweisen wir für PAC Voting Rules eine Variante des Gibbard-Satterthwaite-Theorems, indem wir zeigen, dass keine anonyme PAC Voting Rule gleichzeitig nicht-manipulierbar ist und garantiert, dass das gewählte Komitee die Präferenzen der Wähler proportional repräsentiert. Alle diese Ergebnisse sind durch die Idee vereint, genau zu bestimmen unter welchen Annahmen es nicht-manipulierbare Wahlverfahren gibt.

CORE PUBLICATIONS

This thesis is based on the following core publications in their original, published form.

RANDOMIZED SOCIAL CHOICE

- [1] Relaxed notions of Condorcet-consistency and efficiency for strategyproof social decision schemes. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 181–189, 2022 (with F. Brandt and R. Romen).^{*}
- [2] Strategyproof social decision schemes on super Condorcet domains. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1734–1742, 2023 (with F. Brandt and S. Tausch).[†]
- [3] Strategyproof randomized social choice for restricted sets of utility functions. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 306–312, 2021.[‡]
- [4] Incentives in social decision schemes with pairwise comparison preferences. In *Games and Economic Behavior*, 142:266–291, 2023 (with F. Brandt and W. Suksompong).[§]

^{*} This paper has also been presented at the *8th International Workshop on Computational Social Choice (COMSOC)*, 2021, and the *16th Meeting of the Society for Social Choice and Welfare (SSCW)*, 2022.

[†] This paper has also been presented at the *4th AAMAS Workshop on Games, Incentives, and Algorithms*, 2022.

[‡] This paper has also been presented at the *3rd AAMAS Workshop on Games, Incentives, and Algorithms*, 2021.

[§] An earlier version of this paper appeared in the *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 130–136, 2022.

SET-VALUED SOCIAL CHOICE

- [5] Characterizing the top cycle via strategyproofness. *Theoretical Economics*, 18(2):837–883, 2023 (with F. Brandt).[¶]
- [6] On the indecisiveness of Kelly-strategyproof social choice functions. *Journal of Artificial Intelligence Research*, 73:1093–1130, 2022 (with F. Brandt and M. Bullinger).^{||}

COMMITTEE ELECTIONS

- [7] Strategyproofness and proportionality in party-approval multi-winner elections. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5591–5599, 2023. (with T. Delemazure, T. Demeulemeester, M. Eberl, and J. Israel).^{**}

[¶] This paper has also been presented at the *COMSOC Video Seminar*, 2021, and the *16th Meeting of the Society for Social Choice and Welfare (SSCW)*, 2022.

^{||} This paper has also been presented at the *8th International Workshop on Computational Social Choice (COMSOC)*, 2021. An earlier version of this paper appeared in the *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 251–259, 2021.

^{**} This paper has also been presented at the *Topological and Computational Social Choice Workshop*, 2022, seminars at *Université Paris-Dauphine* and *Université Paris-Saclay*, and the *9th International Workshop on Computational Social Choice (COMSOC)*, 2023.

FURTHER PUBLICATIONS

This is a list of further publications by the author, which are not part of this thesis.

- [8] Characterizations of Sequential Valuation Rules. In Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 1697–1705, 2023 (with C. Dong).^{††}

^{††} This paper has also been presented at the *Online Social Choice and Welfare Seminar*, 2023, and the *9th International Workshop on Computational Social Choice (COMSOC)*, 2023.

ACKNOWLEDGMENTS

I want to express my sincere gratitude to a number of people without whom this thesis would not have been possible in this form. In particular, I want to thank

- my supervisor Felix Brandt for his scientific guidance throughout the last four years. You always had an open door for me, showed me that research can be more than just work, and helped me to grow as a person. I believe you are a great role model for every aspiring researcher and I consider myself lucky to have had you as my supervisor.
- my colleagues in the PAMAS group who always created a relaxed and enjoyable atmosphere. Thank you Christian, Christian, Martin, Anaëlle, René, Matthias, and Chris for the great time we had.
- my co-authors Martin, René, Chris, Sascha, Warut, Tom, Théo, Jonas, Manuel, and of course Felix, with whom I had the pleasure to work on amazing problems.
- Ulle Endriss and Arunava Sen for agreeing to examine my thesis.
- Dominik Peters and Jérôme Lang who hosted me at their group at Université Paris Dauphine.
- my amazing proofreaders Armin, Jonas, Martin, Chris, Matthias, and René.
- my family and friends who always had a word of encouragement for me when academia was showing me its rough sides. There are too many to name all of you, but I want to give special thanks to my brother Armin, without whom I would not have ended up in academia.
- my girlfriend Hayoung for bearing with me even when I (once again) ended up working past 7pm.

You are amazing and I hope that our paths cross again!

CONTENTS

I	SYNTHESIS OF CONTRIBUTIONS	1
1	INTRODUCTION	3
1.1	Strategyproofness in Social Choice: an Overview	4
1.2	Contribution and Outline	6
2	SOCIAL CHOICE – A MATHEMATICAL VIEW	9
2.1	Fundamentals of Single Winner Elections	9
2.2	Three Types of Voting Rules	10
2.3	Strategyproofness for Single Winner Elections	15
2.4	Electing Multiple Winners	22
3	SUMMARY OF PUBLICATIONS	27
3.1	Results for Social Decision Schemes	27
3.2	Results for Social Choice Correspondences	40
3.3	Results for Committee Voting Rules	44
4	METHODOLOGY	47
4.1	Decisiveness versus Strategyproofness	47
4.2	SAT Solving	51
5	CONCLUSION AND FURTHER DIRECTIONS	57
II	ORIGINAL PUBLICATIONS	61
6	RELAXED NOTIONS OF CONDORCET-CONSISTENCY AND EFFICIENCY FOR STRATEGYPROOF SOCIAL DECISION SCHEMES	63
7	STRATEGYPROOF SOCIAL DECISION SCHEMES ON SUPER CONDORCET DOMAINS	77
8	STRATEGYPROOF SOCIAL CHOICE FOR RESTRICTED SETS OF UTILITY FUNCTIONS	91
9	INCENTIVES IN SOCIAL DECISION SCHEMES WITH PAIRWISE COMPARISON PREFERENCES	101
10	CHARACTERIZING THE TOP CYCLE VIA STRATEGYPROOFNESS	133
11	ON THE INDECISIVENESS OF KELLY-STRATEGYPROOF SOCIAL CHOICE FUNCTIONS	183
12	STRATEGYPROOFNESS AND PROPORTIONALITY IN PARTY-APPROVAL MULTIWINNER ELECTIONS	225
13	BIBLIOGRAPHY	237

Part I

SYNTHESIS OF CONTRIBUTIONS

INTRODUCTION

Voting is one of the oldest tools of human society to reach joint decisions. Indeed, already in antique civilizations such as the Greek Poleis and the Roman Republic, elections were used to pass bills, sentence criminals, and determine the holders of political positions (Staveley, 1972). Very likely, as old as voting is also the insight that voters may act strategically by, e.g., voting for the second-most preferred alternative when their most preferred alternative has no chance of winning. While there is no data for verifying this claim for antique democracies, empirical research finds that voters act strategically in modern elections, regardless of the considered country, decade, or voting rule in use (Stephenson et al., 2018). For instance, for the 2019 United Kingdom general election, Mellon (2022) reports that at least 10% of the participants strategically voted for a party that they did not prefer the most but which they believed to have a higher chance of winning the election.

Such phenomena in elections, and more generally voting rules themselves, are studied in the field of social choice theory from a theoretical perspective. In more detail, social choice theorists formalize elections and voting rules based on a rigorous mathematical model: a voting rule is a function that maps the voters' preferences to a group decision. This model then allows to define desirable properties of voting rules, so-called axioms, and to reason for or against specific voting rules by showing that they satisfy or fail desirable properties. A particularly important axiom is strategyproofness, which requires that voters can never benefit by misreporting their true preferences. Hence, investigating voting rules with respect to strategyproofness allows us to answer why strategic voting happens in practice and to check whether there are voting rules that completely resolve this issue.

Unfortunately, it turns out that strategyproofness is an extremely restrictive axiom in social choice theory: Gibbard (1973) and Satterthwaite (1975) have independently shown that only very unattractive voting rules satisfy strategyproofness if there are three or more alternatives. In more detail, these authors have proven that the only strategyproof voting rules are dictatorships, which always select the most preferred alternative of a specific voter, or imposing voting rules, which never select some alternatives. Since such voting rules are unacceptable in practice, the question of how we can circumvent or mitigate this negative result arises.

In this thesis, we will analyze whether we can escape the Gibbard-Satterthwaite theorem by relaxing one of its implicit assumptions, namely that voting rules always choose a single winner deterministically. This condition is frequently criticized as it inherently conflicts with basic fairness conditions.¹ For instance, if there are two alternatives and each is favored by half of the voters, both alterna-

¹ For example, Taylor (2005) writes that “if there is a weakness to the Gibbard-Satterthwaite theorem, it is the assumption that winners are unique” and Kelly (1977) that “the Gibbard-Satterthwaite theorem [...] uses an assumption of singlevaluedness which is unreasonable” (see also Barberà, 1977b; Gärdenfors, 1976; Duggan and Schwartz, 2000; Nehring, 2000; Brandt et al., 2022c). Moreover, Moulin (1983, theorem 1) has even formally proven that every voting rule that always picks a single winner deterministically fails basic fairness axioms.

tives are equally acceptable and there is no fair way to choose a single winner deterministically. In this thesis, we will therefore analyze voting rules that need not always return a single winner. In more detail, we will study strategyproofness for set-valued voting rules (which return subsets of alternatives from which the final winner will be chosen), randomized voting rules (which return lotteries over the alternatives that will be used to determine the final winner), and committee voting rules (which return a fixed number of alternatives that are all winning). The goal of this thesis is hence to explore whether we can circumvent the Gibbard-Satterthwaite theorem by focusing on more flexible models of voting.

1.1 STRATEGYPROOFNESS IN SOCIAL CHOICE: AN OVERVIEW

The possibility that voters may act strategically is deeply rooted in the field of social choice theory. This is perhaps best illustrated by Borda's statement (1784) that his "scheme is only intended for the honest men" (see Black, 1958, p. 215), but also the founding works of modern social choice theory by Arrow (1951) and Black (1948) make explicit reference to the possibility of strategic voting. For instance, Arrow writes that he will not consider the "game aspects" of voting introduced by the possibility that voters may benefit by misreporting their preferences (Arrow, 1951, pp. 6–8). Similar remarks can be found in various papers (e.g., Majumdar, 1956, footnote 5; Sen, 1966, footnote 1; Fishburn, 1970, p. 122), thus demonstrating the awareness of the problem of strategic misrepresentation.

Despite this general awareness, there is little early work that specifically analyzes voting rules with respect to strategyproofness. One of the first scholars who explicitly studied this axiom in social choice was Farquharson (1956a; 1956b; 1961) who, motivated by the works of Nash (1950) and Arrow (1951), analyzed equilibria in voting. Moreover, Vickrey (1960) observed that strategyproofness is closely connected to Arrow's condition of independence of irrelevant alternatives, thus indicating that this axiom may be difficult to satisfy. However, it took another decade until the study of strategyproofness became popular: at the end of the 1960s, the textbooks by Farquharson (1969) and Murakami (1968) dealt in depth with strategyproofness and showed first impossibility results for restricted classes of voting rules. These books led to numerous follow-up works (e.g., Grofman, 1969; Wilson, 1969; Pattanaik, 1973; Zeckhauser, 1973). Finally, the influential textbooks by Sen (1970, Section 11.3) and Fishburn (1973, pp. 97–99) included detailed discussions of strategyproofness in voting, thus paving the way for the Gibbard-Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975).

With the publication of the Gibbard-Satterthwaite theorem, strategyproofness soon became one of the central axioms in social choice theory as researchers tried to escape the negative consequences of this result (e.g., Brams and Fishburn, 1978; Moulin, 1980) or to transfer it to related settings (e.g., Gibbard, 1977; Hylland, 1980). In what follows, we will give an overview of the most significant attempts to circumvent the Gibbard-Satterthwaite theorem. For more detailed discussions, we refer to the surveys by Taylor (2005) and Barberà (2010).

SET-VALUED VOTING RULES. As already noted, the Gibbard-Satterthwaite theorem is frequently criticized for its assumption that voting rules always choose a

single winner deterministically. A large body of research thus focuses on so-called social choice correspondences, which return non-empty subsets of the alternatives instead of single winners (e.g., Barberà, 1977b; Kelly, 1977; Feldman, 1979b; Duggan and Schwartz, 2000; Nehring, 2000; Brandt, 2015). The idea of social choice correspondences is to choose a set of possible winners from which a tie-breaking mechanism will choose the final winner. While there are different ways to define strategyproofness in this context, most lead to rather negative results, as, e.g., demonstrated by the impossibility theorem of Duggan and Schwartz (2000). By contrast, weak notions of strategyproofness are satisfied by attractive set-valued voting rules (e.g., Brandt, 2015). We refer to Section 3.2 and the work by Brandt et al. (2022c) for more details on this line of work.

RANDOMIZED VOTING RULES. Another way to escape the unfairness caused by deterministically picking a single winner is to allow for randomized voting rules. This idea leads to social decision schemes, which are voting rules that return lotteries over the alternatives instead of single winners. The final winner will then be selected by chance according to the chosen lottery. In one of the first results on strategyproofness for randomized voting rules, Gibbard (1977) has shown that every non-imposing and strategyproof social decision scheme is a random dictatorship, a result known as the random dictatorship theorem. This result has caused a significant amount of follow-up works which explore various aspects of strategyproof social decision schemes in more detail (e.g., Barberà, 1979a,b; Hylland, 1980; Brandt et al., 2018; Aziz et al., 2018). For instance, Hylland (1980) has shown a variant of the random dictatorship theorem for cardinal preferences (see also Nandeibam, 2013), and Aziz et al. (2018) investigate several weaker strategyproofness notions. More details on strategyproof social decision schemes can be found in Section 3.1 and in the survey by Brandt (2017).

RESTRICTED DOMAINS. A classical approach to finding attractive strategyproof voting rules is to restrict the domain of feasible preference profiles. First, if there are only two alternatives, it is well-known that there are strategyproof voting rules (May, 1952; Picot and Sen, 2012). Moreover, Brams and Fishburn (1978) have shown that there are strategyproof voting rules when voters have dichotomous preferences over the alternatives, and Moulin (1980) has characterized the set of strategyproof voting rules on the domain of single-peaked preferences. Departing from these early results, more and more domains have been considered (e.g., Barberà et al., 1993; Aswal et al., 2003; Nehring and Puppe, 2007), and modern results pinpoint the boundary on when a domain allows for attractive strategyproof voting rules (e.g., Roy and Storcken, 2019; Chatterji and Zeng, 2023). Similar results have also been explored for the randomized case (e.g., Ehlers et al., 2002; Chatterji et al., 2014; Roy and Sadhukhan, 2020). Domain restrictions will be explained in more detail in Section 3.1.3 and we refer to the surveys by Barberà et al. (2020) and Roy et al. (2022) for an exhaustive overview of the literature.

QUANTITATIVE ANALYSIS. A natural follow-up question to the Gibbard-Satterthwaite theorem is how frequently voters can beneficially manipulate a given voting rule by lying about their preferences. For instance, Pazner and Wesley (1978) and Peleg (1979) show for several voting rules that it is unlikely that a single voter can manipulate the election outcome in her favor if the number of voters is large.

More recent results extend this approach to essentially all known voting rules and give tight bounds on the number of manipulable profiles (Favardin and Lepelley, 2006; Xia, 2023). Moreover, Maus et al. (2007) study voting rules that minimize the number of manipulable preference profiles. These results are complemented by experimental analyses, which quantify the manipulability of voting rules through computer simulations (Chamberlin, 1985; Smith, 1999; Aleskerov and Kurbanov, 1999; Aleskerov et al., 2012). We refer to Xia (2023) and Pritchard and Wilson (2007) for an overview of this line of work.

WEAKENINGS OF STRATEGYPROOFNESS. Another natural venue in light of the Gibbard-Satterthwaite theorem is to consider weaker variants of strategyproofness. For single-valued voting rules, this typically means to impose some restrictions on what we count as a manipulation (e.g., Muller and Satterthwaite, 1977; Campbell and Kelly, 2009; Sanver and Zwicker, 2009; Sato, 2013; Kumar et al., 2021). For instance, Sato (2013) studies local strategyproofness, where voters can only misreport their preferences by swapping two adjacent alternatives in their preference relation. In the context of randomized or set-valued voting rules, there are multiple natural strategyproofness notions, which allows for a more fine-grained analysis (e.g., Cho, 2016; Aziz et al., 2018; Brandt et al., 2022c). For randomized voting rules, this approach has been surveyed by Brandt (2017).

ITERATIVE VOTING. Finally, a relatively new escape route to the Gibbard-Satterthwaite theorem is iterative voting. The idea of this approach is to accept that voting rules are manipulable and to study whether the resulting voting games have an equilibrium and, if so, which properties such an equilibrium satisfies. For instance, Meir et al. (2010) study under which conditions the plurality rule converges to an equilibrium. Similar research is carried out by, e.g., Dhillon and Lockwood (2004), Reinjoud and Endriss (2012), Rabinovich et al. (2015), and Kavner and Xia (2021). We refer to Meir (2017) for a survey of this topic.

1.2 CONTRIBUTION AND OUTLINE

In this thesis, we will study whether it is possible to circumvent the Gibbard-Satterthwaite theorem by dropping the assumption that voting rules always choose a single winner deterministically. In more detail, we will analyze strategyproofness for randomized voting rules, set-valued voting rules, and committee voting rules. Randomized voting rules are typically called social decision schemes (SDSs) and return lotteries over the alternatives instead of single winners. The final winner of the election will then be selected by chance according to the chosen lottery. Similarly, set-valued voting rules, which are formally called social choice correspondences (SCCs), return non-empty sets of alternatives with the understanding that the final winner will be picked from the chosen set by some tie-breaking mechanism. Thus, even though SDSs and SCCs return lotteries over the alternatives and sets of alternatives, they ultimately aim to select a single winner. By contrast, the idea of committee voting rules is to select a fixed number of winning alternatives. In this thesis, we study a particular type of committee voting rules called party-approval committee (PAC) voting rules, which assign the seats of a committee to parties based on the voters' approval preferences over these parties.

As hinted at in Section 1.1, there is already a significant amount of work on strategyproofness for SCCs and SDSs. In more detail, in the context of randomized social choice, Gibbard (1977) has proven that random dictatorships are the only SDSs that jointly satisfy a strategyproofness notion called strong \succsim^{SD} -strategyproofness and *ex post* efficiency. Random dictatorships pick each voter with a fixed probability and implement the favorite alternative of the chosen voter as the winner of the election. This result, which is today known as the random dictatorship theorem, is more positive than the Gibbard-Satterthwaite theorem since random dictatorships can be fair towards the voters. We nevertheless interpret the random dictatorship theorem as a negative result because random dictatorships do not allow for compromise and often use a lot of randomization to select the winner. For instance, if all voters have different favorite alternatives but agree on the second-best alternative, we find it desirable to compromise by selecting the common second-best alternative. However, random dictatorships do not allow for this outcome since they only randomize over the top-ranked alternatives. In Section 3.1, we thus explore the boundaries of the random dictatorship theorem with the aim of finding more attractive strategyproof SDSs. A few highlights of this section are:

- an approximate strengthening of the random dictatorship theorem: we show that strongly \succsim^{SD} -strategyproof SDSs that are almost *ex post* efficient are almost random dictatorships (cf. Section 3.1.2).
- a study of the random dictatorship theorem on restricted domains: we characterize the strongly \succsim^{SD} -strategyproof and non-imposing SDSs on the Condorcet domain as mixtures of random dictatorships and the Condorcet rule (cf. Section 3.1.3).
- a tight analysis on the strategyproofness notions that are compatible with Condorcet-consistency for SDSs on the full domain (cf. Section 3.1.4).

Similar to SDSs, it has also been shown for SCCs that strong strategyproofness notions lead to results similar to the Gibbard-Satterthwaite theorem (e.g., Duggan and Schwartz, 2000; Barberà et al., 2001; Benoît, 2002). Moreover, even for weak strategyproofness notions, only few SCCs are known to be strategyproof (Feldman, 1979b; Brandt, 2015). In Section 3.2, we thus try to better understand when SCCs are strategyproof in order to narrow the gap between the strategyproofness notions that result in possibility results and those that lead to impossibility results. In particular, we discuss in Section 3.2:

- a characterization of an SCC known as the top cycle based on a mild strategyproofness notion due to Gärdenfors (1979) (cf. Section 3.2.1).
- far-reaching impossibilities demonstrating that even very weak strategyproofness notions cannot be satisfied by attractive SCCs when voters may be indifferent between alternatives (cf. Section 3.2.2).

Moreover, our theorems and proof techniques give deep insights into why certain axioms are incompatible with strategyproofness. In more detail, our results suggest a strong correlation between the decisiveness of voting rules (in the sense of how often they choose a single winner without tie-breaking) and their strategyproofness: when voting rules are too decisive, strategyproofness becomes impossible. This can also be observed in many results in the literature (e.g., Barberà,

1977b; Duggan and Schwartz, 2000; Benoît, 2002), but it has never been explicitly mentioned. We thus discuss in Section 4.1 the tradeoff between strategyproofness and decisiveness.

Finally, we will also analyze strategyproofness for PAC voting rules, which assign the seats of a committee to the parties based on the voters' approval preferences over these parties (cf. Section 3.3). This model has only recently been introduced by Brill et al. (2022) and our work is thus the first one on strategyproof PAC voting rules. In particular, we show a far-reaching impossibility theorem for this setting by proving that no anonymous PAC voting rule satisfies strategyproofness and always chooses a committee that proportionally represents the voters' preferences. This theorem precludes the existence of attractive strategyproof PAC voting rules because the proportional representation of the voters' preferences is one of the central goals of committee elections. Thus, our impossibility result transfers the negative consequences of the Gibbard-Satterthwaite theorem to PAC elections. Since we prove this result based on a computer-aided technique called SAT solving, we will also explain this approach in detail in Section 4.2.

In summary, our results enhance the understanding of strategyproof voting rules when dropping the assumption of single-valuedness in three ways. Firstly, we show numerous impossibility results based on rather mild strategyproofness notions and common decisiveness conditions for all our settings. While these results are negative, they are important to improve our understanding of when strategyproof social choice is possible. Secondly, we characterize several attractive voting rules based on strategyproofness, thus singling out the most desirable strategyproof voting rules. Finally, our results significantly narrow the gap between the strategyproofness notions used for impossibility results and those used for possibility results. Hence, we also derive insights about the assumptions on the voters' preferences over, e.g., lotteries that lead to strategyproof social choice.

In this chapter, we introduce our mathematical framework for voting. To this end, we recall that the goal of every election ultimately is to aggregate the voters' preferences over some alternatives into a group decision. From this high-level description, we can already derive some basic insights. In particular, in every election, there is a set of voters N who report preferences over a set of alternatives A . A voting rule then maps these preferences to a group decision. We denote by $\mathcal{O}(A)$ the set of all possible outcomes of the election. Moreover, we want to use voting rules for more than a single set of preferences and thus define $\mathcal{D}(A, N)$ as the set of admissible preference profiles. Equivalently, $\mathcal{D}(A, N)$ can be viewed as the domain of our voting rule. Finally, a voting rule is a function of the type $\mathcal{D}(A, N) \rightarrow \mathcal{O}(A)$.

Clearly, this definition is rather generic as the domain $\mathcal{D}(A, N)$ and the set of possible outcomes $\mathcal{O}(A)$ have not been specified. We thus discuss in the following four sections various types of voting rules. In more detail, in Sections 2.1 to 2.3, we introduce three models for choosing a single winner based on the voters' preferences, define strategyproofness for this setting, and formalize further desirable properties. Finally, in Section 2.4, we discuss committee voting rules, which select a fixed number of winners based on the voters' preferences.

2.1 FUNDAMENTALS OF SINGLE-WINNER ELECTIONS

In the majority of this thesis, we will consider the problem of selecting a single winner based on the voters' preferences. We will formalize this setting as follows: we assume that there is a finite set of voters $N = \{1, \dots, n\}$ and a finite set of alternatives $A = \{a_1, \dots, a_m\}$. Following the standard in the literature, we denote by n the number of voters and by m the number of alternatives. Moreover, we suppose that $m \geq 2$ throughout this thesis. Every voter $i \in N$ is assumed to report a *preference relation* \succsim_i over the alternatives A , which is a complete and transitive binary relation on A . As for all binary relations, we will write \succ_i for the strict part of \succsim_i (i.e., $x \succ_i y$ if and only if $x \succsim_i y$ and not $y \succsim_i x$) and \sim_i for the indifference part of \succsim_i (i.e., $x \sim_i y$ if and only if $x \succsim_i y$ and $y \succsim_i x$). We call a preference relation \succsim_i *strict* if it is additionally antisymmetric, i.e., if it contains no ties between alternatives. When we want to stress that a preference relation is strict, we write \succ_i instead of \succsim_i . Conversely, we discuss *weak* preference relations to emphasize that the considered preference relations need not be strict. The set of all strict preference relations is denoted by \mathcal{L} and the set of all weak preference relations is \mathcal{R} . Preference relations will be written as comma-separated lists, where brackets indicate ties. For instance, $\succsim_i = a, \{b, c\}$ means that voter i strictly prefers a to both b and c and is indifferent between the latter two alternatives.

Next, a *preference profile* R assigns a preference relation to every voter $i \in N$, i.e., it is an element of \mathcal{R}^N . We call a preference profile strict if the preference




Figure 2.1: Example of a preference profile and the corresponding majority relation.

relations of all voters are strict and define \mathcal{L}^N as the set of all strict preference profiles. Even though \mathcal{R}^N is more general than \mathcal{L}^N as $\mathcal{L}^N \subseteq \mathcal{R}^N$, we will follow the literature and mainly focus on the domain of strict preference profiles in our analysis. When multiple voters share a preference relation in a profile, we write the set of voters that report the same preference relation before it. To this end, we define $[i \dots j] = \{k \in N : i \leq k \leq j\}$ and note that we omit set brackets for singleton sets. For example, $[1 \dots 3] : a, b, c$ means that voters 1, 2, and 3 prefer a to b to c .

In the literature, there are two antipodal views on preference profiles. The first option is to evaluate alternatives by comparing them to each other. To this end, we define the *support* of an alternative x against another alternative y in a profile R as $n_{xy}(R) = |\{i \in N : x \succ_i y\}|$, i.e., as the number of voters who strictly prefer x to y . For strict preference profiles $R \in \mathcal{L}^N$, it holds that $n_{xy}(R) + n_{yx}(R) = n$ for all distinct $x, y \in A$. The support between alternatives gives rise to the *majority relation* \succsim_M , which is defined by $x \succsim_M y$ if and only if $n_{xy}(R) \geq n_{yx}(R)$ for all distinct alternatives $x, y \in A$. In words, $x \succsim_M y$ holds if at least as many voters prefer x to y as vice versa. As for preference relations, \succ_M denotes the strict part of \succsim_M and \sim_M the indifference part. An example of a preference profile and its majority relation is shown in Figure 2.1.

The second perspective on preference profiles evaluates the quality of an alternative only based on its position in the voters' preference relations. This idea leads to the notion of the *rank*, which we define for strict preference relations $\succ_i \in \mathcal{L}$ as $r(\succ_i, x) = 1 + |\{y \in A \setminus \{x\} : y \succ_i x\}|$. That is, an alternative has rank k in the preference relation of voter i if it is her k -th best alternative. For weak preferences, it is less straightforward to define the rank as it is unclear how to handle ties. We thus define the *rank tuple* by $\bar{r}(\succsim_i, x) = (1 + |\{y \in A \setminus \{x\} : y \succ_i x\}|, |\{y \in A \setminus \{x\} : y \sim_i x\}|)$ as a rather general extension of the rank to weak preferences.

Finally, we want to stress again that our standard domain for single winner elections is \mathcal{L}^N . We will nevertheless define numerous concepts for weak preference relations as this allows us to easily transfer them to different settings. Moreover, in our results, we sometimes modify the domain by, e.g., allowing a variable number of voters or restricting the set of feasible preference profiles.

2.2 THREE TYPES OF VOTING RULES

We will next introduce the types of single-winner voting rules considered in this thesis. In particular, we will discuss social choice functions (which return single alternatives), social decision schemes (which return lotteries over the alternatives), and social choice correspondences (which return non-empty subsets of the

alternatives). Despite the different output formats, all these types of voting rules ultimately aim to select a single winner. We will thus also introduce common axioms for these voting rules. Most concepts in this section are defined for strict preferences, but it is usually easy to adapt them to weak preferences.

2.2.1 Social Choice Functions

Social choice functions are perhaps the simplest type of voting rules: given the voters' preferences, these functions return a single winner. When assuming strict preferences, this corresponds to the following definition.

DEFINITION 2.1 (*Social Choice Functions*)

A *social choice function* (SCF) is a function of the type $\mathcal{L}^N \rightarrow A$.

We note that there are many prominent examples of SCFs. However, all of the subsequent rules except dictatorships need tie-breaking to ensure that there is a single winner in the case that multiple alternatives are tied for the win.

- Maybe the most prominent SCF is the *plurality rule* f_{PL} , which chooses the alternative that is top-ranked by the most voters. To formalize this idea, we define the scoring function s_{PL} by $s_{PL}(\succ_i, x) = 1$ if $r(\succ_i, x) = 1$ and $s_{PL}(\succ_i, x) = 0$ otherwise. The plurality score of an alternative x in a profile R is then defined by $s_{PL}(R, x) = \sum_{i \in N} s_{PL}(\succ_i, x)$ and the plurality rule chooses an alternative with maximal plurality score.
- Another prominent SCF is the *Borda rule* f_{Borda} , which is named after the Chevalier de Borda who promoted this rule already in the 18th century (1784). The idea of this rule is that each voter gives $m - k$ points to her k -th best alternative. To formalize this, we define the Borda score of an alternative x in a preference relation \succ_i as $s_{Borda}(\succ_i, x) = m - r(\succ_i, x)$ and in a profile R as $s_{Borda}(R, x) = \sum_{i \in N} s_{Borda}(\succ_i, x)$. The Borda rule then chooses an alternative with maximal Borda score.
- A conceptually rather different SCF is the *Copeland rule* $f_{Copeland}$, which chooses the alternative that beats the maximal number of alternatives in a pairwise majority comparison. Formally, the Copeland rule selects in every preference profile R an alternative x that maximizes the Copeland score $s_{Copeland}(R, x) = |\{y \in A \setminus \{x\} : x \succ_M y\}| + \frac{1}{2} |\{y \in A \setminus \{x\} : x \sim_M y\}|$ (Copeland, 1951).
- A rather unattractive class of SCFs are dictatorships: the *dictatorship* of voter i , denoted by d_i , always chooses the most preferred alternative of voter i .

For an example of these rules, we consider the profile R shown in Figure 2.1 and assume lexicographic tie-breaking. For this profile, it holds that $f_{PL}(R) = b$ as both b and e are top-ranked by two voters and the lexicographic tie-breaking picks b . Moreover, $f_{Borda}(R) = b$, too, as b has a maximal score of $s_{Borda}(R, b) = 13$. Finally, $f_{Copeland}(R) = a$ since both a and b have a Copeland score of 3 and the lexicographic tie-breaking chooses a .

2.2.2 Social Decision Schemes

As explained in the introduction, SCFs are inherently unfair as they need to break ties in situations where multiple alternatives are equally acceptable. One way to

deal with this problem is to allow for randomization in choosing the winner of the election. This idea leads to social decision schemes, which are voting rules that map the voters' preferences to lotteries over the alternatives. The final winner will then be chosen by chance according to the probabilities assigned by the social decision scheme. To formalize this, we define lotteries as probability distributions over alternatives, i.e., a lottery p is a function of the type $A \rightarrow [0, 1]$ such that $\sum_{x \in A} p(x) = 1$. We denote the set of all lotteries by $\Delta(A)$ and define social decision schemes as follows.

DEFINITION 2.2 (*Social Decision Schemes*)

A social decision scheme (SDS) is a function of the type $\mathcal{L}^N \rightarrow \Delta(A)$.

To simplify notation, we define $f(R, x)$ as the probability that the SDS f assigns to alternative x in the profile R . Moreover, we extend this notation to sets $X \subseteq A$ by $f(R, X) = \sum_{x \in X} f(R, x)$.

We note that SDSs are a strict generalization of SCFs: every SCF can be represented as an SDS that puts probability 1 on some alternative for every preference profile. Perhaps more surprisingly, SDSs can also be interpreted as SCFs where the new alternatives are the lotteries over the old alternatives. Since preferences over lotteries are usually structured, the SCFs are then defined for some restricted domain of preferences.

We conclude this section by introducing several important SDSs.

- First, we note that every SCF that assigns scores to the alternatives can be turned into an SDS by randomizing proportional to these scores. For example, the *randomized Borda rule* is defined by $f_{RB}(R, x) = \frac{2}{n \cdot m(m-1)} s_{Borda}(R, x)$ and the *randomized Copeland rule* by $f_{RC}(R, x) = \frac{2}{m(m-1)} s_{Copeland}(R, x)$.
- Another important class of SDSs are random dictatorships. These rules pick every voter with a fixed probability and return the most preferred alternative of the chosen voter. More formally, let d_i denote the SDS that always puts probability 1 on voter i 's most preferred alternative. Then, an SDS f is a *random dictatorship* if it is a convex combination of the SDSs d_i , i.e., if there are values $\gamma_i \geq 0$ for all $i \in N$ such that $\sum_{i \in N} \gamma_i = 1$ and $f(R) = \sum_{i \in N} \gamma_i d_i(R)$ for all preference profiles R . A particularly interesting SDS within this class is the uniform random dictatorship f_{RD} , which chooses every voter with probability $1/n$ and thus enjoys a high degree of fairness.
- As the last class of SDSs, we introduce maximal lottery rules, which have been suggested by Fishburn (1984) and recently promoted by Brandl et al. (2016). For defining these SDSs, we denote $ML(R) = \{p \in \Delta(A) : \forall q \in \Delta(A) : \sum_{x,y \in A} n_{xy}(R)p(x)q(y) \geq \sum_{x,y \in A} n_{xy}(R)q(x)p(y)\}$ by the set of maximal lotteries in the profile R . The set of maximal lotteries is always non-empty by the minimax theorem and almost always a singleton. In particular, if the number of voters is odd, there is always a unique maximal lottery (Lafond et al., 1997; Le Breton, 2005). Finally, an SDS f_{ML} is a *maximal lottery rule* if $f_{ML}(R) \in ML(R)$ for all preference profiles R .

For an example of these SDSs, we consider again the profile R in Figure 2.1. The randomized Copeland rule returns $f_{RC}(R) = [a : 3/10, b : 3/10, c : 2/10, d : 2/10, e : 0]$ for this profile, which can be verified by counting the outgoing edges of each alternative in the majority relation. Next, $f_{RD}(R) = [a : 1/5, b : 2/5, c : 0, d : 0, e : 2/5]$

which follows by counting how often each alternative is top-ranked. Finally, the unique maximal lottery in R is $f_{ML}(R) = [a : 1/3, b : 1/3, c : 0, d : 1/3, e : 0]$.

2.2.3 Social Choice Correspondences

Another option to deal with the unfairness caused by SCFs is to completely separate voting rules from the tie-breaking and thus allow voting rules to choose sets of possible winners. We formalize this idea with social choice correspondences which are voting rules that return non-empty sets of alternatives.²

DEFINITION 2.3 (*Social Choice Correspondences*)

A social choice correspondence (SCC) is a function of the type $\mathcal{L}^N \rightarrow 2^A \setminus \{\emptyset\}$.

The intuition of an SCC is to choose a set of possible winners from which a tie-breaking mechanism will ultimately select the final winner. For example, the final winner may be chosen by a lottery or by a chairperson. Thus, even though SCCs return sets of alternatives, a single winner will eventually be chosen from this set. When combining an SCC with a deterministic tie-breaking mechanism, we derive an SCF. Furthermore, we derive an SDS when combining an SCC with a randomized tie-breaking mechanism. Conversely, we can also turn an SDS into an SCC by returning the support of the chosen lottery (i.e., the set of alternatives with a positive probability). Perhaps surprisingly, this means that SCCs introduce much more uncertainty on the final winner of the election than SDSs: whereas SCCs may use any tie-breaking mechanism to select the final winner, the tie-breaking mechanism of SDSs is already specified and the chosen lottery can be broadcasted to the voters. For the analysis of SCCs, we will largely ignore the tie-breaking.

We note that the plurality rule, the Borda rule, and the Copeland rule can also be seen as SCCs if we return the sets of alternatives that maximize the respective scores. Moreover, there are also SCCs which tend to choose large choice sets and thus make no sense as SCFs, e.g.:

- The *omnination rule* chooses all alternatives that are top-ranked by at least one voter, i.e., $f_{OMNI}(R) = \{x \in A : \exists i \in N : r(\succ_i, x) = 1\}$. Alternatively, this SCC can be defined as the support of the uniform random dictatorship f_{RD} (i.e., the set of alternatives x with $f_{RD}(R, x) > 0$).
- The *top cycle* f_{TC} chooses all alternatives that reach all other alternatives on some path in the majority relation.³ To make this more formal, we define a path in the majority relation \succsim_M as a sequence of alternatives (a_1, \dots, a_k) such that $a_i \succsim_M a_{i+1}$ for all $i \in \{1, \dots, k-1\}$. Moreover, we write $a \succsim_M^* b$ if there is a path in the majority relation that starts at a and ends at b . The relation \succsim_M^* is the transitive closure of the majority relation, and we can define the top cycle as the SCC that chooses the maximal elements of \succsim_M^* , i.e., $f_{TC}(R) = \{x \in A : \forall y \in A \setminus \{x\} : x \succsim_M^* y\}$.

For the profile R in Figure 2.1, it is easy to compute that $f_{OMNI}(R) = \{a, b, e\}$. Moreover, $f_{TC}(R) = \{a, b, c, d\}$ because e loses all pairwise majority comparisons and the alternatives in $\{a, b, c, d\}$ form a cycle in the majority relation.

² Social choice correspondences are sometimes called set-valued social choice functions in the literature. In particular, this is the case in Publication 6.

³ The top cycle is also known as Good set (Good, 1971), Smith set (Smith, 1973), weak closure maximality (Sen, 1977), and GETCHA (Schwartz, 1986).

2.2.4 Basic Axioms

In the last part of this section, we introduce basic axioms for our voting rules. In particular, since all three types of voting rules considered in this section aim to eventually select a single winner, they share common desiderata. To avoid repetition when formalizing these desiderata, we say a function f is a *single-winner voting rule* if it is an SCF, an SDS, or an SCC.

ANONYMITY. As the first axiom, we introduce a basic fairness notion for voters called anonymity. Informally, this axiom states that all voters should be treated equally. More formally, we call a single-winner voting rule f *anonymous* if $f(R) = f(\pi(R))$ for all preference profiles R and permutations $\pi : N \rightarrow N$ of the voters. Here, $R' = \pi(R)$ denotes the profile defined by $\succ'_i = \succ_{\pi(i)}$ for all $i \in N$. Put differently, anonymity requires that the outcome is invariant under exchanging the voters' identities. We note that this axiom is independent of the output type of a voting rule, so its definition is the same for SCFs, SCCs, and SDSs. Anonymity is a very mild condition that is satisfied by all single-winner voting rules in this section except dictatorships and non-uniform random dictatorships.

NEUTRALITY. Analogous to anonymity, we define next a fairness notion for alternatives. To this end, we say that a single-winner voting rule f is *neutral* if $f(\tau(R)) = \tau(f(R))$ for all preference profiles R and permutations $\tau : A \rightarrow A$. This time, the profile $R' = \tau(R)$ is defined by $\tau(x) \succ'_i \tau(y)$ if and only if $x \succ_i y$ for all voters $i \in N$ and alternatives $x, y \in A$. In words, neutrality requires that when renaming alternatives in the profile, we need to rename them accordingly in the outcome. The definition of $\tau(f(R))$ depends on the type of f : if f is an SCF, then $\tau(f(R))$ is already well-defined; if f is an SCC, then $\tau(f(R)) = \{\tau(x) : x \in f(R)\}$; finally, if f is an SDS, then $p = \tau(f(R))$ is the lottery defined by $p(\tau(x)) = f(R, x)$ for all alternatives $x \in A$. Just as anonymity, neutrality is a mild condition that is satisfied by all SCCs and SDSs defined in this section. By contrast, SCFs that rely on lexicographic tie-breaking violate this axiom.

NON-IMPOSITION. Another common fairness notion is non-imposition, which requires that every alternative should be the unique winner in some preference profile. For instance, if all voters agree that x is the best option, then x should intuitively be the winner of the election. To formalize this, we say a single-winner voting rule f is *non-imposing* if for every alternative $x \in A$ there is a profile R such that $f(R) = x$ if f is an SCF, $f(R) = \{x\}$ if f is an SCC, and $f(R, x) = 1$ if f is an SDS. Non-imposition can be seen both as a fairness condition because it guarantees that all alternatives are the unique winner in some profile, and as a minimal decisiveness notion because it rules out that SCCs and SDSs always return multiple possible winners. All voting rules in this section except the randomized Borda rule and the randomized Copeland rule satisfy non-imposition.

PERFECT-PARETO-OPTIMALITY. As our fourth condition, we introduce Pareto-optimality which formalizes a mild efficiency criterion: an alternative should not be chosen if there is another alternative that makes some voters better off without making any voter worse off. We formalize this idea subsequently for weak preference profiles and thus say that an alternative x *Pareto-dominates* another alternative y in a profile R if $x \succsim_i y$ for all voters $i \in N$ and $x \succ_i y$ for some voter $i \in N$. If

all voters have strict preferences, x Pareto-dominates y if all voters strictly prefer the former alternative to the latter one. Next, an alternative is *Pareto-optimal* if it is not Pareto-dominated by any other alternative. We denote the set of Pareto-optimal alternatives by $f_{PO}(R) = \{x \in A: x \text{ is Pareto-optimal in } R\}$. This function can be interpreted as an SCC and is typically called the Pareto rule. For example, it can be checked that $f_{PO}(R) = \{a, b, d, e\}$ for the profile R in Figure 2.1 as alternative b Pareto-dominates alternative c . Finally, a single-winner voting rule is Pareto-optimal if only Pareto-optimal alternatives have a chance to win the election. Formally, an SCF f is *Pareto-optimal* if $f(R) \in f_{PO}(R)$ for all profiles R , an SCC f if $f(R) \subseteq f_{PO}(R)$ for all profiles R , and an SDS f if $f(R, f_{PO}(R)) = 1$ for all profiles R . We note that for SDSs, Pareto-optimality is typically called *ex post efficiency* and can be equivalently defined by $f(R, x) = 0$ for all $x \notin f_{PO}(R)$. Finally, it is easy to see that Pareto-optimality implies non-imposition for single-winner voting rules. Even though Pareto-optimality and *ex post* efficiency are rather basic conditions, some rules violate these axioms. For instance, the top cycle, the randomized Borda rule, and the randomized Copeland rule fail this condition as they choose the Pareto-dominated alternative c with positive probability in the profile R in Figure 2.1. All other rules in this section are Pareto-optimal.

CONDORCET-CONSISTENCY. The last axiom that we introduce here is Condorcet-consistency. To this end, we define a *Condorcet winner* in a profile R as an alternative x that beats all other alternatives in a pairwise majority comparison, i.e., $x \succ_M y$ for all $y \in A \setminus \{x\}$. While a Condorcet winner is not guaranteed to exist, it is always unique if it does. For instance, the profile R in Figure 2.1 does not admit a Condorcet winner as every alternative loses a majority comparison. Many social choice theorists agree that the Condorcet winner should be uniquely chosen whenever there is one. This property is known as *Condorcet-consistency* and formally requires for all profiles R with Condorcet winner x that $f(R) = x$ if f is an SCF, $f(R) = \{x\}$ if f is an SCC, and $f(R, x) = 1$ if f is an SDS. Clearly, Condorcet-consistency implies non-imposition. For SCCs and SDSs, Condorcet-consistency can be viewed as a decisiveness criterion as it requires that a single winner is chosen for many profiles. We note that numerous voting rules, such as the plurality rule, the Borda rule, dictatorships, the omninomination rule, and the randomized variants of these rules, fail Condorcet-consistency. By contrast, the Copeland rule, the top cycle, and maximal lottery rules satisfy this axiom.

2.3 STRATEGYPROOFNESS FOR SINGLE WINNER ELECTIONS

We next turn to the central axioms of this thesis: strategyproofness and manipulability. Intuitively, strategyproofness requires that voters cannot be better off by lying about their true preferences. Conversely, a voting rule is manipulable if it is not strategyproof. Strategyproofness is important for several reasons: firstly, if a voting rule is not strategyproof, we cannot expect the voters to report their true preferences, and we may choose a socially non-optimal alternative due to wrong information. Moreover, all desirable properties of a manipulable voting rule are in question as these are typically shown under the assumption that voters act truthfully. For example, we may not be able to identify the Condorcet winner correctly

if the voters do not report their true preferences, and Condorcet-consistency may thus be violated with respect to the voters' true preferences. Finally, in the more general field of mechanism design, strategyproofness is frequently motivated by the revelation principle which states that any property that can be obtained in an equilibrium can be obtained in the truthful equilibrium of a direct strategyproof mechanism. Thus, it suffices to restrict the attention to strategyproof mechanisms when studying properties of equilibrium outcomes.

Since strategyproofness depends on the output type of the voting rule, we separately define this axiom for each of our three types of single-winner voting rules. First, in the case of SCFs, it is straightforward to define strategyproofness and manipulability: voters can simply compare the outcomes chosen by an SCF according to their preference relations. Formally, an SCF f is *strategyproof* if $f(R') \not\succ_i f(R)$ for all preference profiles R, R' and voters $i \in N$ such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$. Note here that $f(R') \not\succ_i f(R)$ is equivalent to $f(R) \succeq_i f(R')$ as preference relations are complete. Conversely, we say that an SCF is *manipulable* if it is not strategyproof, i.e., if there are profiles R, R' and a voter $i \in N$ such that $f(R') \succ_i f(R)$ and $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$. For example, it is easy to check that the plurality rule with lexicographic tie-breaking is manipulable: in the profile R in Figure 2.1, it holds that $f_{PL}(R) = b$, but if voter 5 reports a as her favorite alternative instead of e , then $f_{PL}(R) = a$. Because voter 5 prefers a to b according to her true preference relation, this deviation constitutes a successful manipulation.

By contrast, it is unclear how to define strategyproofness for SDSs and SCCs since voters only report preferences over the alternatives and not over lotteries or sets of alternatives. For example, if a voter prefers a to b to c , it is unclear whether she prefers the set $\{a, c\}$ to the set $\{b\}$. For our study of SDSs and SCCs, we will thus rely on lottery and set extensions, which lift the voters' preferences over alternatives to preferences over lotteries and sets of alternatives, respectively. In more detail, we discuss four lottery extensions in Section 2.3.1 and two set extensions in Section 2.3.2. Since some of these extensions have originally been defined for weak preference relations, we introduce these concepts always for weak preferences. Finally, we define strategyproofness for SCCs and SDSs in Section 2.3.3.

2.3.1 Lottery Extensions

To define strategyproofness for SDSs, voters need to compare lotteries and we thus introduce several lottery extensions, which lift the voters' preferences over alternatives to preferences over lotteries. To this end, we note that several lottery extensions have been suggested in the literature (see, e.g., Cho, 2016; Brandt, 2017; Aziz et al., 2018). Subsequently, we discuss four of these extensions.

SD EXTENSION. Maybe the most prominent lottery extension in the literature is the *stochastic dominance (SD) extension* (see, e.g., Gibbard, 1977; Barberà, 1979b; Bogomolnaia and Moulin, 2001; Ehlers et al., 2002; Brandl et al., 2018). According to this notion, a voter $i \in N$ (weakly) prefers a lottery p to a lottery q if p stochastically dominates q according to \succeq_i . More formally, it holds for all lotteries $p, q \in \Delta(A)$ and preference relations \succeq_i that

$$p \succeq_i^{SD} q \iff \forall x \in A: \sum_{y \in A: y \succeq_i x} p(y) \geq \sum_{y \in A: y \succeq_i x} q(y).$$

The appeal of the *SD* extension stems from the fact that it can also be formalized based on von Neuman-Morgenstern (vNM) utilities. To this end, we define vNM utility functions u as mappings from the set of alternatives A to the set of real numbers \mathbb{R} , i.e., a vNM utility function u states for every alternative $x \in A$ a value $u(x)$ that measures the subjective quality of x . We say a vNM utility function u is consistent with a preference relation \succsim if $x \succsim y$ if and only if $u(x) \geq u(y)$ for all alternatives $x, y \in A$. Moreover, we define \mathcal{U}^{\succsim} as the set of all vNM utility functions that are consistent with the preference relation \succsim and let $\mathcal{U} = \bigcup_{\succsim \in \mathcal{R}} \mathcal{U}^{\succsim}$ (resp. $\hat{\mathcal{U}} = \bigcup_{\succ \in \mathcal{L}} \mathcal{U}^{\succ}$) denote the set of vNM utility functions that are consistent with an arbitrary (resp. strict) preference relation. The idea of the *SD* extension is now that each voter i uses a vNM utility function u_i to compare lotteries by their expected utility $\mathbb{E}_p[u_i] = \sum_{x \in A} p(x)u_i(x)$, i.e., voter i prefers a lottery p to a lottery q if $\mathbb{E}_p[u_i] \geq \mathbb{E}_q[u_i]$. However, the vNM utility functions of the voters are not known by the mechanism designer and the *SD* extension thus quantifies over all vNM utility functions in \mathcal{U}^{\succsim_i} : for all preference relations \succsim_i and lotteries $p, q \in \Delta(A)$, it holds that $p \succsim_i^{SD} q$ if and only if $\mathbb{E}_p[u_i] \geq \mathbb{E}_q[u_i]$ for all $u_i \in \mathcal{U}^{\succsim_i}$ (Sen, 2011; Brandl et al., 2018).

U EXTENSION. In Publication 3, we suggest a new class of lottery extensions called *U* extensions, which are closely related to the *SD* extension. The basic idea of this concept is again that voters use vNM utility functions to compare lotteries but, in contrast to the *SD* extension, it is not necessary to consider all vNM utility functions as some may not be plausible for a given scenario. We hence specify a set of vNM utility functions $U \subseteq \mathcal{U}$ that will be used to compare lotteries by the *U* extension. More formally, we define the *U* extension for all preference relations \succsim_i , sets of vNM utility functions $U \subseteq \mathcal{U}$, and lotteries $p, q \in \Delta(A)$ by

$$p \succsim_i^U q \iff \forall u_i \in U \cap \mathcal{U}^{\succsim_i}: \mathbb{E}_p[u_i] \geq \mathbb{E}_q[u_i].$$

It follows from this definition that the *SD* extension is equivalent to the \mathcal{U} extension. Furthermore, the smaller the considered set of vNM utility functions U , the more lotteries are comparable by the *U* extension. For strict preferences, an important special case of this extension arises when the set U contains a single utility function u and its permutations. We thus define the set u^Π for a vNM utility function $u \in \hat{\mathcal{U}}$ by $u' \in u^\Pi$ if and only if there is a permutation $\pi: A \rightarrow A$ such that $u'(x) = u(\pi(x))$ for all $x \in A$. In particular, the u^Π extension associates every strict preference relation with a single canonical vNM utility function. We note that similar but less general concepts than the *U* extension have been considered by, e.g., Sen (2011) and Mennle and Seuken (2021).

PC EXTENSION. Another approach for comparing lotteries over alternatives is the concept of *pairwise comparison* (*PC*): a voter prefers a lottery p to a lottery q if it is more likely that she prefers the alternative drawn from p to the alternative drawn from q than vice versa. More formally, the *PC* extension is defined for all lotteries $p, q \in \Delta(A)$ and preference relations \succsim_i as follows:

$$p \succsim_i^{PC} q \iff \sum_{x, y \in A: x \succ_i y} p(x)q(y) \geq \sum_{x, y \in A: y \succ_i x} p(x)q(y).$$

The *PC* extension has been suggested by Aziz et al. (2015a) in the context of randomized social choice, but it has been considered before in decision theory (Blyth,

1972; Packard, 1982; Blavatsky, 2006). Moreover, this extension has recently attracted attention in social choice theory as it allows for strong positive results (Brandl et al., 2019; Brandl and Brandt, 2020). Finally, we note that the *PC* extension can also be interpreted in the context of utility functions. To this end, we introduce Fishburn's skew-symmetric bilinear (SSB) utility functions u , which assign values to all pairs of alternatives $x, y \in A$ such that $u(x, y) = -u(y, x)$ (Fishburn, 1982). A voter i with SSB utility function u_i then prefers a lottery p to another lottery q if $\sum_{x, y \in A: x \succ_i y} u_i(x, y)p(x)q(y) \geq \sum_{x, y \in A: y \succ_i x} u_i(y, x)p(x)q(y)$. The *PC* extension then arises when each voter $i \in N$ uses the SSB utility function u_i given by $u_i(x, y) = 1$ if $x \succ_i y$ and $u_i(x, y) = 0$ if $x \sim_i y$ to compare lotteries and it is thus conceptually related to the u^Π extension.

DD EXTENSION. As the last lottery extension in this section, we consider the *deterministic dominance (DD) extension* suggested by Brandt (2017). The idea of this extension is that a voter prefers a lottery p to a lottery q if she weakly prefers every outcome that can be chosen by p to every outcome that can be chosen by q . Formally, this extension is defined as follows:

$$p \succsim_i^{DD} q \iff \forall x, y \in A : p(x)q(y) > 0 \implies x \succsim_i y.$$

The *DD* extension only allows for rather uncontroversial comparisons between lotteries and thus formalizes highly risk-averse voters. Notably, this extension can be computed based only on the supports of the considered lotteries (i.e., the sets of alternatives with positive probabilities) and it is hence related to a set extension due to Kelly (1977).

For each of our lottery extensions \succsim^X , we define by \succ^X its strict part and by \sim^X its indifference part. Furthermore, we note that all our lottery extensions but the u^Π extension and the *PC* extension are incomplete, and all lottery extensions but the *PC* extension are transitive. Finally, it holds for all preference relations $\succsim \in \mathcal{R}$ and sets of vNM utility functions U, U' with $U \subseteq U'$ that $\succsim^{DD} \subseteq \succsim^{SD} \subseteq \succsim^{U'} \subseteq \succsim^U$ and $\succsim^{DD} \subseteq \succsim^{SD} \subseteq \succsim^{PC}$. In contrast, the *U* extension and the *PC* extension are not related by subset inclusion. Analogous inclusions hold for the strict parts of the lottery extensions, too, when requiring that $U \cap \mathcal{W}^\succ \neq \emptyset$ for the *U* extension.

EXAMPLE 2.4

We next give an example to illustrate the various lottery extensions. To this end, let $\succsim_i = a, b, c, d$ and define the lotteries p, q, r, s by

$$\begin{aligned} p(c) &= 1, & q(b) &= 1/3, q(c) = 2/3 \\ r(b) &= r(c) = 1/2, & s(a) &= 5/11, s(d) = 6/11. \end{aligned}$$

It is easy to see that $q \succ_i^{DD} p$ and $r \succ_i^{DD} p$ and that all other pairs of lotteries are incomparable by the *DD* extension. Next, $r \succ_i^{SD} q$ since $r(a) = q(a) = 0$, $r(a) + r(b) = 1/2 > 1/3 = q(a) + q(b)$, and $r(a) + r(b) + r(c) = 1 = q(a) + q(b) + q(c)$. Moreover, $q \succ_i^{DD} p$ and $r \succ_i^{DD} p$ imply that $q \succ_i^{SD} p$ and $r \succ_i^{SD} p$. By contrast, the *SD* extension cannot compare lottery s to any of the other lotteries. The *PC* extension states that $x \succ_i^{PC} s$ for all lotteries $x \in \{p, q, r\}$ because, e.g., $q(b)s(d) + q(c)s(d) = 6/11 > 5/11 = s(a)q(b) + s(a)q(c)$. For all remaining pairs, the *PC* extension agrees with the *SD* extension. Finally, it holds for the utility function u defined by $u(a) = 3$, $u(b) = 2$, $u(c) = 1$, and $u(d) = 0$ that $r \succ_i^{u^\Pi} s \succ_i^{u^\Pi} q \succ_i^{u^\Pi} p$ because r has an expected utility of $3/2$, s of $15/11$, q of $4/3$, and p of 1 .

2.3.2 Set Extensions

We next turn to set extensions, which lift the voters' preferences over alternatives to preferences over sets of alternatives. While there are numerous set extensions considered in the literature (we refer to Gärdenfors (1979), Taylor (2005), and Jimeno et al. (2009) for an overview), we will only consider Kelly's and Fishburn's extension in this thesis. These extensions only allow for rather uncontroversial comparisons between sets of alternatives and we refer to Erdamar and Sanver (2009) and Brandt et al. (2022c) for detailed discussions about their interpretation.

KELLY'S EXTENSION. One of the first set extensions suggested in the literature is due to Kelly (1977) and requires that a voter prefers a set X to a set Y if she weakly prefers every alternative $x \in X$ to every alternative $y \in Y$. Formally, given a preference relation \succsim_i and two non-empty sets $X, Y \subseteq A$, *Kelly's extension* is defined by

$$X \succsim_i^K Y \iff \forall x \in X, y \in Y: x \succsim_i y.$$

When \succsim_i is a strict preference relation, $X \succsim_i^K Y$ requires that $|X \cap Y| \leq 1$. Kelly's extension can be motivated by the assumption that the final winner will be selected from the choice set X by some tie-breaking mechanism unknown to the voters. Then, $X \succsim_i^K Y$ if, regardless of the tie-breaking mechanism, voter i prefers the choice from X at least as much as the choice from Y . Alternatively, Kelly's extension can be motivated by the assumption of randomized tie-breaking: a voter prefers the set X to the set Y if the former guarantees her at least as much expected utility as the latter for every lottery over X , every lottery over Y , and every vNM utility function that is consistent with her preference relation. Kelly's extension coincides with the *DD* extension when applied to the support of the lotteries.

FISHBURN'S EXTENSION. The second set extension we consider in this thesis is due to Gärdenfors (1979), who attributes it to Fishburn (1972) as it is the smallest set extension that satisfies five axioms proposed by Fishburn. Following Gärdenfors' suggestion, we thus call it Fishburn's extension. This set extension requires that a voter prefers a set X to a set Y if she prefers all alternatives in $X \setminus Y$ to all alternatives in Y and all alternatives in X to all alternatives in $Y \setminus X$. More formally, *Fishburn's extension* is defined as follows for all preference relations \succsim_i and all non-empty sets of alternatives $X, Y \subseteq A$:

$$X \succsim_i^F Y \iff \forall x \in X \setminus Y, y \in Y: x \succsim_i y \wedge \forall x \in X, y \in Y \setminus X: x \succsim_i y.$$

This extension can be motivated by the assumption that a chairperson picks the final winner according to her own strict preference relation. Then, $X \succsim_i^F Y$ if and only if voter i weakly prefers the outcome selected from X to the outcome selected from Y for every possible preference relation of the chairperson. Alternatively, Fishburn's extension can be motivated by the assumption of *a priori* weights w_x of the alternatives such that the likelihood for an alternative x to be selected from a set Y with $x \in Y$ is $\frac{w_x}{\sum_{y \in Y} w_y}$. Then, $X \succsim_i^F Y$ if and only if the expected utility of the alternative chosen from X is at least as high as the expected utility of the alternative chosen from Y for all *a priori* weights and all vNM utility functions that are consistent with voter i 's preference relation.

We define \succ^K and \sim^K (resp. \succ^F and \sim^F) as the strict part and indifference part of \succsim^K (resp. \succsim^F). Moreover, for all preference relations \succsim , it holds that \succsim^K and \succsim^F are transitive but incomplete binary relations and that $\succsim^K \subseteq \succsim^F$ and $\succ^K \subseteq \succ^F$.

EXAMPLE 2.5

To illustrate Kelly's and Fishburn's extensions, let $\succsim_i = a, b, c$. It is now easy to verify that, e.g., $\{a\} \succ_i^K \{a, b\}$ and $\{a, b, c\} \succ_i^K \{c\}$. By contrast, Kelly's extension does not allow to compare the sets $X = \{a, b\}$ and $Y = \{a, b, c\}$ as $a \succ_i b$ and $X \cap Y = \{a, b\}$. In turn, it holds that $X \succsim_i^F Y$ since $X \setminus Y = \emptyset$ and $x \succ_i y$ for all $x \in X, y \in Y \setminus X$. Finally, the sets $\{a, b\}$ and $\{a, c\}$ are incomparable with respect to Fishburn's extension as a is contained in both sets and preferred to both b and c .

2.3.3 Strategyproofness for SDSs and SCCs

Finally, we are ready to define strategyproofness for SDSs and SCCs. Since many set and lottery extensions are incomplete, there are two ways to define strategyproofness, which differ in how incomparable outcomes are handled.

DEFINITION 2.6 (Strong and weak \succsim^X -strategyproofness)

Given a lottery extension (resp. set extension) \succsim^X , an SDS (resp. SCC) f satisfies

- *strong \succsim^X -strategyproofness* if $f(R) \succsim_i^X f(R')$ for all profiles R, R' and voters $i \in N$ such that $\succsim_j = \succsim_j'$ for all $j \in N \setminus \{i\}$.
- *weak \succsim^X -strategyproofness* if $f(R') \not\succ_i^X f(R)$ for all profiles R, R' and voters $i \in N$ such that $\succsim_j = \succsim_j'$ for all $j \in N \setminus \{i\}$.

Conversely, we call an SDS or SCC *strongly \succsim^X -manipulable* if it is not weakly \succsim^X -strategyproof and *weakly \succsim^X -manipulable* if it is not strongly \succsim^X -strategyproof.

The difference between weak and strong \succsim^X -strategyproofness lies in the fact how incomparable sets or lotteries are handled. The strong notion requires that a voter always \succsim^X -prefers the outcome when voting honestly to any outcome she could obtain by lying. Hence, deviating to an outcome that is incomparable with respect to \succsim^X is a manipulation. By contrast, the weak notion only prohibits that voters can deviate to a strictly \succsim^X -preferred outcome and a deviation to an incomparable outcome is thus no successful manipulation. Consequently, strong \succsim^X -strategyproofness implies weak \succsim^X -strategyproofness for all set and lottery extensions \succsim^X . Notably, when a set or lottery extension is complete (such as \succsim^{PC} and \succsim^{u^Π}), weak and strong strategyproofness coincide. We thus write \succsim^{PC} -strategyproofness and \succsim^{u^Π} -strategyproofness without the prefix strong or weak. Moreover, for two lottery extensions (resp. set extensions) \succsim^X, \succsim^Y with $\succsim^X \subseteq \succsim^Y$, it follows that strong \succsim^X -strategyproofness implies strong \succsim^Y -strategyproofness and weak \succsim^Y -strategyproofness implies weak \succsim^X -strategyproofness. We thus infer the relations depicted in Figure 2.2 for our strategyproofness notions.

EXAMPLE 2.7

For an example of weak and strong strategyproofness, we consider the following three preference profiles.

R^1 :	1: a, b, c	2: a, b, c	3: b, c, a	4: b, c, a	5: c, a, b
R^2 :	1: a, b, c	2: a, b, c	3: b, c, a	4: b, c, a	5: a, c, b
R^3 :	1: a, b, c	2: a, b, c	3: b, c, a	4: c, b, a	5: c, a, b

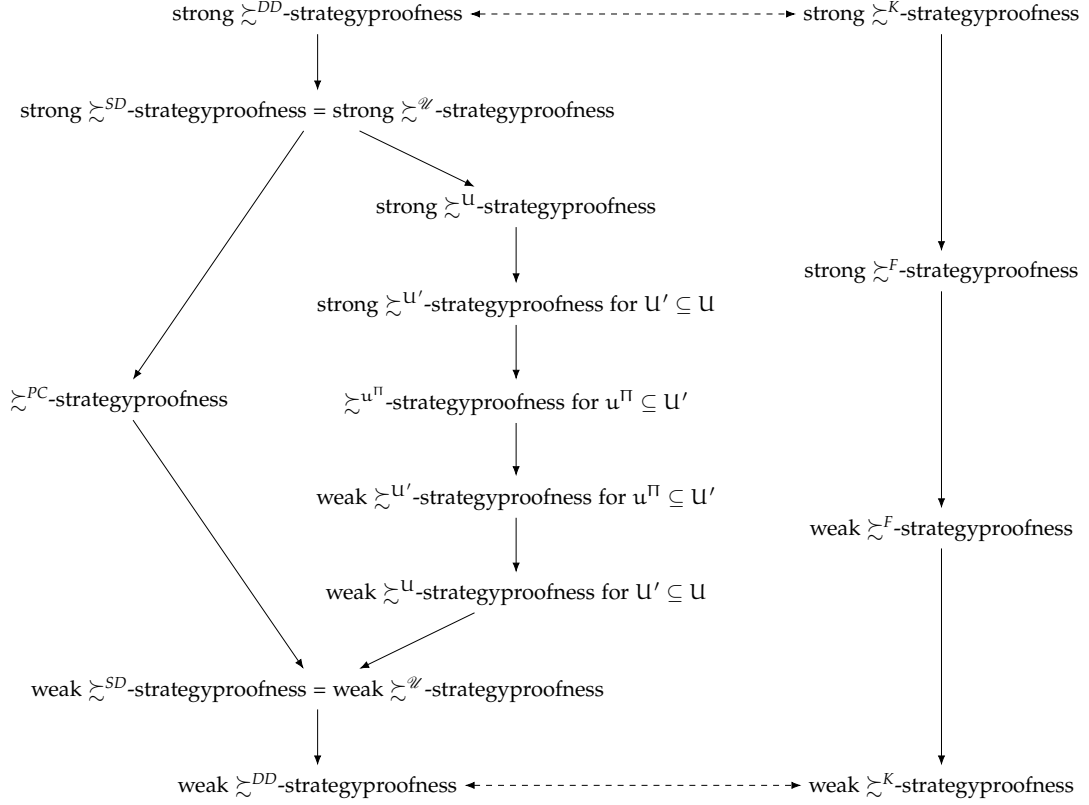


Figure 2.2: Overview of the relations between our strategyproofness notions for SDSs and SCCs. An arrow from X to Y means that X implies Y . The implications involving weak \succsim^U -strategyproofness require that $U \cap \mathcal{U} \neq \emptyset$ for all preference relations \succsim . The dotted line between strong (resp. weak) \succsim^{DD} -strategyproofness and strong (resp. weak) \succsim^K -strategyproofness indicates that these two notions coincide when considering the support of SDSs.

It can be computed that $f_{ML}(R^1) = [a : 3/5, b : 1/5, c : 1/5]$, $f_{ML}(R^2) = [a : 1, b : 0, c : 0]$, and $f_{ML}(R^3) = [a : 1/3, b : 1/3, c : 1/3]$ for every maximal lottery rule f_{ML} . Now, $f_{ML}(R^1)$ is incomparable to $f_{ML}(R^2)$ according to the SD extension of \succsim_5^1 . Hence, deviating from R^1 to R^2 is a weak \succsim^{SD} -manipulation for voter 5 but no strong \succsim^{SD} -manipulation. By contrast, it holds that $f(R^3) \succ_4^{SD} f(R^1)$ for \succsim_4^1 , so voter 4 can strongly \succsim^{SD} -manipulate by deviating from R^1 to R^3 .

For SDSs, it is common to consider the strong variant of strategyproofness. In particular, strong \succsim^{SD} -strategyproofness has attracted significant attention (e.g., Gibbard, 1977; Barberà, 1979b; Ehlers et al., 2002; Roy et al., 2022). Moreover, several SDSs mentioned in Section 2.2.2 satisfy strong \succsim^{SD} -strategyproofness when preferences are strict: results by Gibbard (1977) and Barberà (1979b) imply that all random dictatorships, the randomized Borda rule, and the randomized Copeland rule are strongly \succsim^{SD} -strategyproof. In Section 3.1.4, we discuss SDSs that satisfy strong \succsim^U -strategyproofness for a large set of utility functions U but fail strong \succsim^{SD} -strategyproofness. An example of a weakly \succsim^{SD} -strategyproof SDS is the randomized Condorcet rule, which chooses the Condorcet winner with probability 1 if it exists and otherwise returns the uniform lottery over all alternatives. Moreover, if $m = 3$, this SDS is even \succsim^{PC} -strategyproof. Finally, maximal lottery

rules that assign probability 1 to an alternative only if $|ML(R)| = 1$ are weakly \succsim^{DD} -strategyproof but not weakly \succsim^{SD} -strategyproof (Brandt, 2017).

By contrast, when studying SCCs, we will focus on weak strategyproofness. The reason for this is that \succsim^K and \succsim^F allow to compare too few sets: for instance, not even the SCC that always returns all alternatives is strongly \succsim^K -strategyproof as \succsim^K does not allow to compare the set of all alternatives to itself. Moreover, although weak \succsim^K -strategyproofness and weak \succsim^F -strategyproofness are rather mild strategyproofness notions, only few SCCs satisfy these axioms. For instance, the plurality rule, the Borda rule, and the Copeland rule all fail weak \succsim^K -strategyproofness.⁴ On the other hand, the top cycle, the omninomination rule, and the Pareto rule satisfy weak \succsim^F -strategyproofness for strict preferences. Finally, there are also SCCs that are weakly \succsim^K -strategyproof but not weakly \succsim^F -strategyproof, e.g., the uncovered set and the bipartisan set (see Brandt et al., 2016, for details).

2.4 ELECTING MULTIPLE WINNERS

We will also consider elections that aim to elect multiple alternatives instead of a single winner in this thesis. The prime example of such elections are parliamentary elections, where multiple seats of a parliament have to be distributed to the parties based on the voters' preferences. Moreover, this type of election can also be used to model technical applications such as recommender systems (e.g., Gawron and Faliszewski, 2022) or medical diagnostic support systems (Gangl et al., 2019).

While it is possible to study committee elections based on strict or weak preference relations (see, e.g., Elkind et al., 2017; Faliszewski et al., 2017), it is more common in the recent literature to analyze committee elections based on dichotomous preferences (Lackner and Skowron, 2023). In these preference relations, voters only distinguish between acceptable and unacceptable alternatives. More formally, we call a preference relation $\succsim \in \mathcal{R}$ *dichotomous* if it has at most two equivalence classes, i.e., for all triples of alternatives $x, y, z \in A$ with $x \succ y$, it holds that $x \sim z$ or $y \sim z$. Furthermore, we say that a voter i *approves* alternative x if $x \succsim_i y$ for all other alternatives $y \in A$. Dichotomous preferences are typically motivated by their simplicity (Laslier and Sanver, 2010). Since all alternatives that are not acceptable are by definition unacceptable in dichotomous preferences, we only write the set of approved alternatives to indicate a voter's preference relation and we call this set *approval ballot*. For instance, $1 : \{a, b\}$ means that voter 1 approves alternatives a and b . The set of all approval ballots, or equivalently the set of all dichotomous preferences, is denoted by \mathcal{A} and is the set of all non-empty subsets of A . Moreover, we define \mathcal{A}^N as the set of all *approval profiles*.

Given an approval profile, our goal is to select a committee of predefined size k . In this thesis, we aim to formalize the elections of city councils or parliaments and thus interpret our alternatives as parties. When studying committee voting rules, we will even refer to the elements $x \in A$ as parties instead of alternatives. Since every party can have multiple seats in a city council, we define committees as multisets of size k . More formally, a *committee* W is a function of the type

⁴ The Copeland rule and the Borda rule can even be manipulated in situations where single alternatives are chosen before and after the manipulation. This is the most severe type of manipulation as it is independent of the tie-breaking assumptions.

$A \rightarrow \mathbb{N}_0$ such that $\sum_{x \in A} W(x) = k$. Furthermore, \mathcal{W}_k denotes the set of all size k committees. Then, we define party-approval committee voting rules as follows.

DEFINITION 2.8 (*Party-approval Committee Voting Rules*)

A party-approval committee (PAC) voting rule is a function of the type $\mathcal{A}^N \rightarrow \mathcal{W}_k$.

Given a party x , an approval profile R , and a PAC voting rule f , we define $f(R, x)$ as the number of seats that f assigns to party x . We extend this notion also to sets of parties $X \subseteq A$ by defining $f(R, X) = \sum_{x \in X} f(R, x)$ and to dichotomous preference relations \succsim_i by $f(R, \succsim_i) = f(R, A_i)$, where A_i is the approval ballot of voter i . Similarly, $W(X)$ denotes the number of seats assigned to the parties $x \in X$ by W and $W(\succsim_i)$ is the number of seats assigned to voter i 's approved parties.

We note that, in contrast to most of the literature, we define committee voting rules for a fixed committee size k . The reason for this is that we will study PAC voting rules with respect to strategyproofness, which does not allow us to change the committee size. Thus, the above definition avoids an additional input parameter. Moreover, PAC voting rules are by definition resolute, i.e., they always return a single winning committee. Just as for SCFs, this will require tie-breaking whenever there are multiple committees tied for the win.

The model of PAC voting has only recently been introduced by Brill et al. (2022). However, it is closely related to apportionment (where voters can only approve a single party; see, e.g., Balinski and Young (2001)), approval-based committee (ABC) voting (where voters approve individual candidates rather than parties and the output is thus a size k subset of the alternatives instead of a multiset; see, e.g., Lackner and Skowron (2023)), and fair mixing (where the output is a probability distribution over the alternatives rather than a multiset; see, e.g., Aziz et al. (2020)). Firstly, the relation to apportionment is obvious because PAC voting rules relax the restriction that voters only vote for a single party. Next, ABC voting rules are a special case of PAC voting rules as every set is also a multiset. Conversely, one can also model PAC voting rules as ABC voting rules by replacing every party with k alternatives representing its members. Then, PAC elections constitute a domain restriction for ABC voting rules as voters either have to approve all members of a party or none. Finally, we can turn every PAC voting rule into a fair mixing rule by returning $\frac{f(R, x)}{k}$ for every party $x \in A$, which shows that fair mixing is an even more general model than PAC voting.

In the rest of this section, we define a class of PAC voting rules called Thiele rules (cf. Section 2.4.1) and discuss desirable axioms for PAC elections (cf. Section 2.4.2).

2.4.1 Thiele Rules

Thiele rules, which have first been suggested by Thiele (1895), are one of the most prominent class of committee voting rules. The idea of these rules is similar to that of single-winner scoring rules: voters give points to the committees based on the number of seats allocated to their approved parties and the winning committee is the one that maximizes the total score. To formalize this, we define Thiele scoring functions s as mappings of the type $\mathbb{N}_0 \rightarrow \mathbb{R}$ such that $s(\ell + 1) \geq s(\ell)$ and $s(\ell + 2) - s(\ell + 1) \leq s(\ell + 1) - s(\ell)$ for all $\ell \in \mathbb{N}_0$. The first assumption is typically motivated by the idea that voters prefer committees that contain more of their approved parties, and the second one by the observation that the gain of each

additionally approved committee member is decreasing. Next, we define the score of a committee W in an approval profile R as $s(R, W) = \sum_{i \in N} s(W(\succ_i))$. Finally, a PAC voting rule f is a *Thiele rule* if there is a Thiele scoring function s such that f selects a committee W with maximal total score $s(R, W)$ for each approval profile R . We note that Thiele rules need tie-breaking to ensure that a single winner is chosen when there are multiple committees with the same maximal total score.

Next, we introduce three important Thiele rules. We refer to Example 2.9 for an illustration of these rules.

- *Multiwinner approval voting* (f_{AV}) is the Thiele rule defined by the Thiele scoring function $s_{AV}(x) = x$ for all $x \in \mathbb{N}_0$.
- *Proportional approval voting* (f_{PAV}) is the Thiele rule defined by the Thiele scoring function $s_{PAV}(x) = \sum_{y=1}^x \frac{1}{y}$ for all $x \in \mathbb{N}$ and $s(0) = 0$.
- *Chamberlin-Courant approval voting* (f_{CCAV}) is the Thiele rule defined by the Thiele scoring function $s_{CCAV}(x) = 1$ for all $x \in \mathbb{N}$ and $s(0) = 0$.

2.4.2 Axioms

We now turn to the desirable properties for PAC elections. At the end of this section, we also present an example to illustrate all introduced axioms.

ANONYMITY. We first note that anonymity can be defined just as for single-winner elections: a PAC voting rule f is *anonymous* if $f(R) = f(\pi(R))$ for all approval profiles R and permutations $\pi : N \rightarrow N$, i.e., if it treats all voters equally. All typically considered PAC voting rules satisfy this axiom.

WEAK REPRESENTATION. A central desideratum in committee elections is to select committees that fairly represent the voters' preferences. In particular, if a large group of voters approves a common party, this group should be represented by some members in the elected committee. A mild axiom motivated by this idea is *weak representation*, which requires of a PAC voting rule f that $f(R, x) \geq 1$ for all approval profiles R and parties $x \in A$ such that there are at least n/k voters that uniquely approve party x in R . Less formally, this axiom postulates that a party should have at least one seat in the chosen committee if there is a group of voters with size at least n/k that unanimously and uniquely approves this party. The reason for this is that $1/k$ -th of the voters should determine $1/k$ -th of the seats in the committee. Weak representation is a weakening of justified representation, which is a well-known proportionality notion for approval-based committee elections (Aziz et al., 2017).

WEAK PROPORTIONAL REPRESENTATION. While weak representation guarantees large cohesive groups of voters some representation, this may not be strong enough in practice. For example, if half of the voters uniquely approve a single party, this party is only ensured a single seat in the committee rather than half of the seats. We thus introduce a stronger representation axiom: a PAC voting rule f satisfies *weak proportional representation* if $f(R, x) \geq \ell$ for all approval profiles R and parties $x \in A$ such that there are at least $\ell \frac{n}{k}$ voters that uniquely approve party x in R . We note that this axiom is implied by a property known as *proportional justified representation*, which is frequently considered in ABC elections (Aziz et al., 2017;

Sánchez-Fernández et al., 2017). For instance, f_{PAV} satisfies weak proportional representation,⁵ f_{CCAV} satisfies weak representation but not weak proportional representation, and f_{AV} fails both proportionality notions.

STRATEGYPROOFNESS. As the last condition for PAC voting rules, we introduce strategyproofness. To this end, we say a PAC voting rule f is *strategyproof* if $f(R, \succsim_i) \geq f(R', \succsim_i)$ for all approval profiles R, R' and voters $i \in N$ such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$. Put differently, a PAC voting rule is strategyproof if a voter approves at least as many members of the committee selected when voting honestly as of any committee that she could obtain by lying about her preferences. Similar to strategyproofness for SCCs and SDSs, we thus use a committee extension to lift the voters' dichotomous preferences over the parties to weak preference relations over the committees. We can therefore interpret results based on our strategyproofness notion also in the context of single-winner voting, where the alternatives are the committees and the set of feasible preferences is strongly restricted. It should be mentioned that this strategyproofness notion has already been studied for approval-based committee elections under the name cardinality-strategyproofness (e.g., Peters, 2018; Lackner and Skowron, 2018; Aziz et al., 2015b). For instance, it is known that f_{AV} satisfies strategyproofness and that all other Thiele rules fail this condition.

EXAMPLE 2.9

To illustrate the concepts of this section, we consider an example. To this end, we suppose there are five parties $A = \{a, b, c, d, e\}$ and that our target committee size is $k = 4$. Moreover, we consider the approval profiles R^1 and R^2 shown below.

R^1 :	[1 ... 6]: {a}	7: {a, b}	[8 ... 10]: {c}	11: {d}	12: {e}
R^2 :	[1 ... 6]: {a}	7: {b}	[8 ... 10]: {c}	11: {d}	12: {e}

We start by considering the profile R^1 . First, it is easy to see that multiwinner approval voting (f_{AV}) chooses the committee W^1 that assigns all 4 seats to a . The reason for this is that a is approved by 7 voters, whereas every other party is approved by less voters, so W^1 achieves a maximal score of $s_{AV}(R^1, W^1) = 7 \cdot 4 + 5 \cdot 0 = 28$. This also shows that AV fails weak representation: 3 of the 12 voters uniquely approve party c , but it does not get a seat in the committee.

Next, proportional approval voting (f_{PAV}) selects for the profile R^1 the committee W^2 that assigns three seats to party a and one to party c . This committee achieves a score of $s_{PAV}(R^1, W^2) = 7 \cdot 11/6 + 3 \cdot 1 + 2 \cdot 0 = 95/6$, which is maximal. It can moreover be checked that this committee satisfies weak proportional representation: this axiom postulates that a gets at least two seats as 6 of the 12 voters uniquely approve it and that c gets at least one seat as 3 out of the 12 voters uniquely approve it.

Finally, Chamberlin-Courant approval voting (f_{CCAV}) chooses for R^1 the committee W^3 that gives one seat to each of a, c, d and e . The reason for this is that f_{CCAV} chooses the committee that contains at least one approved member for as many voters as possible. Since the committee W^3 contains for every voter an approved party, it achieves the maximal score of $s_{CCAV}(R^1, W^3) = 12$. Consequently, CCAV satisfies weak representation for R^1 but fails weak proportional representation.

⁵ In fact, it is known that f_{PAV} satisfies a much stronger proportionality notion called the core in PAC elections (Brill et al., 2022).

Finally, we turn to the profile R^2 . First, it is easy to see that $f_{AV}(R^2) = W^1$ and $f_{PAV}(R^2) = W^2$, i.e., these rules choose the same committees as for R^1 . By contrast, assuming suitable tie-breaking, f_{CCAV} will choose the committee W^4 that assigns one seat to each of a , b , c , and d . This shows that voter 7 can manipulate f_{CCAV} by deviating from R^1 to R^2 since $f_{CCAV}(R^2, \{a, b\}) = 2 > 1 = f_{CCAV}(R^1, \{a, b\})$.

SUMMARY OF PUBLICATIONS

In this chapter, we discuss the main results of our publications and relate our work to the literature. In particular, all of our results have in common that they aim to circumvent or to strengthen the Gibbard-Satterthwaite theorem. We formally state this result next.

THEOREM 3.1 (Gibbard, 1973; Satterthwaite, 1975)

If $m \geq 3$, every strategyproof and non-imposing SCF on \mathcal{L}^N is a dictatorship.

An analogous result also holds if we allow voters to express weak preferences: in this case, every strategyproof and non-imposing SCF always chooses one of the most preferred alternatives of a specific voter. On the other hand, if $m = 2$, the majority rule, which chooses alternative a if $a \succsim_M b$ and alternative b otherwise, is strategyproof. Consequently, we will focus for all results on the case that $m \geq 3$ as the problem of strategyproof social choice is not interesting otherwise.

Since the Gibbard-Satterthwaite theorem is an impossibility result that rules out the existence of attractive strategyproof SCFs, we try to better understand when strategyproof social choice is possible in related settings. In more detail, we study strategyproofness for SDSs in Section 3.1, for SCCs in Section 3.2, and for PAC voting rules in Section 3.3. We note that for each of the publications in Part II, we only give the main theorems; the publications typically state more results and discuss various additional aspects in the form of remarks.

3.1 RESULTS FOR SOCIAL DECISION SCHEMES

As our first escape route to the Gibbard-Satterthwaite theorem, we analyze strategyproofness for SDSs. We note that this is also one of the oldest approaches to circumvent the Gibbard-Satterthwaite theorem as, e.g., Gibbard (1977) and Barberà (1979b) studied the notion of strong \succsim^{SD} -strategyproofness shortly after the publication of the impossibility result for SCFs. In particular, these authors characterize the set of strongly \succsim^{SD} -strategyproof SDSs for strict preferences.⁶ A prominent consequence of the work by Gibbard (1977) is the so-called random dictatorship theorem: if there are at least three alternatives, an SDS is *ex post* efficient and strongly \succsim^{SD} -strategyproof if and only if it is a random dictatorship.

While this result may seem like a straightforward extension of the Gibbard-Satterthwaite theorem, it is much more positive because the uniform random dictatorship enjoys a high degree of fairness. On the other hand, random dictatorships do not allow for compromise: for instance, if all voters have different favorite alternatives but a common second-best alternative, it seems reasonable to select the second-best alternative. However, random dictatorships can only randomize

⁶ While Barberà (1979b) restricts his attention to anonymous and neutral SDSs, Gibbard (1977) indeed characterizes the set of all strongly \succsim^{SD} -strategyproof SDSs.

over the top-ranked alternatives and therefore do not allow for this compromise. On a formal level, this observation relates to the fact that random dictatorships fail Condorcet-consistency. Since random dictatorships also tend to use a large amount of randomization to determine the winner (see Brandl et al., 2022), we interpret the random dictatorship as an impossibility theorem and thus consider several approaches for deriving more positive results for SDSs in this section.

To this end, we first discuss the random dictatorship theorem in more detail in Section 3.1.1. After that, we explore several escape routes: we investigate the set of strongly \succsim^{SD} -strategyproof SDSs with respect to relaxations of *ex post* efficiency and Condorcet-consistency in Section 3.1.2, restrict the domain of preference profiles in Section 3.1.3, and study weaker strategyproofness notions in Section 3.1.4.

3.1.1 The Random Dictatorship Theorem

Before presenting our results, we will discuss and prove the random dictatorship theorem because of its central role in the literature on strategyproof SDSs. Indeed, this theorem, which has first been shown by Gibbard (1977),⁷ has caused a large amount of follow-up works: there are numerous alternative proofs of this result (e.g., Duggan, 1996; Nandeibam, 1997; Tanaka, 2003) as well as extensions to cardinal preferences (e.g., Hylland, 1980; Dutta et al., 2007; Nandeibam, 2013), weaker notions of strategyproofness (e.g., Benoît, 2002; Sen, 2011; Aziz et al., 2018; Brandl et al., 2018), and manipulations by groups (Barberà, 1979a). Most of these results show that the negative consequences of the random dictatorship theorem prevail when modifying the underlying assumptions. By contrast, the most successful escape routes to this theorem are domain restrictions (see Roy et al., 2022) and very weak strategyproofness notions (see Brandt, 2017).

We will next formally state and prove the random dictatorship theorem. We note that our proof is new and not part of any of the publications in this thesis. The main observation for our proof is that every strongly \succsim^{SD} -strategyproof and *ex post* efficient SDS for $n + 1$ voters induces an SDS for n voters that satisfies the same axioms by fixing the preference relation of a voter. Based on this insight, we then show the random dictatorship theorem by an induction over the number of voters. While there are already inductive proofs of the random dictatorship theorem (e.g., Sen, 2011), these use the inductive argument to generalize the result to larger n after establishing it for two voters. By contrast, our proof is mainly carried by the inductive argument and gives a natural way to turn a random dictatorship for n voters into one for $n + 1$ voters. Moreover, the subsequent reasoning can be seen as an example of our proof techniques as most of our proofs use similar arguments. Finally, since non-imposition and strong \succsim^{SD} -strategyproofness imply *ex post* efficiency, it is easy to derive the Gibbard-Satterthwaite theorem from our proof as dictatorships are the only deterministic random dictatorships.

THEOREM 3.2 (Gibbard, 1977)

If $m \geq 3$, an SDS on \mathcal{L}^N is strongly \succsim^{SD} -strategyproof and *ex post* efficient if and only if it is a random dictatorship.

Proof. First, we note that random dictatorships are *ex post* efficient as they only randomize over the top-ranked alternatives of the voters. Moreover, these SDSs

⁷ Gibbard (1977) attributes the random dictatorship theorem to Hugo Sonnenschein.

are strongly \succsim^{SD} -strategyproof since each voter assigns a fixed probability to her favorite alternative and she can only shift this probability to a worse alternative by deviating. Hence, we focus on the inverse direction and suppose that f is a strongly \succsim^{SD} -strategyproof and *ex post* efficient SDS. Our goal is to show that f is a random dictatorship. To this end, we first derive several auxiliary claims on the behavior of f and finally prove the theorem in Step 4.

Step 1: As the first step, we show that strong \succsim^{SD} -strategyproofness implies two properties known as localizedness and non-perversity (Gibbard, 1977).⁸ For defining these axioms, we let the upper contour set $U(\succsim_i, x) = \{y \in A: y \succsim_i x\}$ denote the set of alternatives that voter i prefers to x . Then, an SDS f is *localized* if $f(R, x) = f(R', x)$ for all profiles $R, R' \in \mathcal{L}^N$ such that $U(\succsim_i, x) = U(\succsim'_i, x)$ for all $i \in N$. Moreover, an SDS f is *non-perverse* if $f(R, x) \geq f(R', x)$ for all profiles $R, R' \in \mathcal{L}^N$ such that $U(\succsim_i, x) \subseteq U(\succsim'_i, x)$ for all $i \in N$. Now, let f denote a strongly \succsim^{SD} -strategyproof SDS; we will show that f is localized and non-perverse.

For localizedness, consider two profiles R and R' and an alternative x such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$ and $U(\succsim_i, x) = U(\succsim'_i, x)$. By strong \succsim^{SD} -strategyproofness from R to R' and vice versa, we infer that $f(R, U(\succsim_i, x)) = f(R', U(\succsim_i, x))$ and $f(R, U(\succsim_i, x) \cup \{x\}) = f(R', U(\succsim_i, x) \cup \{x\})$. So, $f(R, x) = f(R', x)$ and repeating this argument for all voters shows that f is localized. For non-perversity, consider two preference profiles R, R' and an alternative x such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$ and $U(\succsim_i, x) \subseteq U(\succsim'_i, x)$. Moreover, let \hat{R}' denote the profile derived from R' by reordering the alternatives in $U(\succsim'_i, x)$ in the preference relation of voter i according to \succsim_i . In particular, voter i prefers all alternatives in $U(\succsim_i, x)$ to those in $U(\succsim'_i, x) \setminus U(\succsim_i, x)$ in \hat{R}' . Next, let \hat{R} denote the profile derived from \hat{R}' by moving x up in the preference of voter i such that $U(\succsim_i, x) = U(\hat{\succsim}_i, x)$. By localizedness, we deduce that $f(\hat{R}, x) = f(R, x)$ and $f(\hat{R}', x) = f(R', x)$. Moreover, strong \succsim^{SD} -strategyproofness shows that $f(\hat{R}, U(\succsim_i, x)) = f(\hat{R}', U(\succsim_i, x))$ as $U(\succsim_i, x)$ is an upper contour set in both $\hat{\succsim}_i$ and $\hat{\succsim}'_i$. Finally, this axiom also entails that $f(\hat{R}, U(\succsim_i, x) \cup \{x\}) \geq f(\hat{R}', U(\succsim_i, x) \cup \{x\})$ and we thus derive that $f(R, x) = f(\hat{R}, x) \geq f(\hat{R}', x) = f(R', x)$. By repeating this argument for one voter after another, it follows that f is non-perverse.

Step 2: Next, we will investigate the outcomes for specific profiles. In particular, consider two distinct alternatives $x, y \in A$, two profiles R and R' , and a partition (N_x, N_y) of the voters such that all voters in N_x top-rank x in both R and R' , and all voters in N_y top-rank y in both profiles. We claim that $f(R) = f(R')$ and $f(R, x) + f(R, y) = f(R', x) + f(R', y) = 1$. To see this, let R^* denote the profile in which all voters in N_x top-rank x and second-rank y , and all voters in N_y top-rank y and second-rank x ; the remaining alternatives can be placed arbitrarily. By *ex post* efficiency, it holds that $f(R^*, x) + f(R^*, y) = 1$. Next, let R denote an arbitrary profile in which the voters in N_x top-rank x and the voters in N_y top-rank y . Moreover, we define R^1 as the profile with $\succsim_i^1 = \succsim_i^*$ for all $i \in N_x$ and $\succsim_i^1 = \succsim_i$ for all $i \in N_y$, and R^2 as the profile with $\succsim_i^2 = \succsim_i$ for all $i \in N_x$ and $\succsim_i^2 = \succsim_i^*$ for all $i \in N_y$. We will next show that $f(R^*, x) = f(R^1, x) = f(R, x)$ and $f(R^*, y) = f(R^2, y) = f(R, y)$, which entails that $f(R^*) = f(R)$. For the claim on R^1 , we first note that R^* and R^1 only differ in the preferences of the

⁸ Our definitions of localizedness and non-perversity differ from those given by Gibbard (1977), but they are logically equivalent. We find our variants easier to use.

voters in N_y and that these voters still top-rank y . Hence, $f(R^*, y) = f(R^1, y)$ by localizedness. Moreover, *ex post* efficiency requires that $f(R^1, z) = 0$ for all $z \in A \setminus \{x, y\}$ because the voters in N_x name y as their second-best alternative. This implies that $f(R^1, x) = 1 - f(R^1, y) = 1 - f(R^*, y) = f(R^*, x)$. Finally, only the voters in N_x deviate when moving from R^1 to R and all these voters top-rank x in both profiles. Using again localizedness, we infer that $f(R, x) = f(R^1, x) = f(R^*, x)$. An analogous argument works for R^2 , so $f(R) = f(R^*)$. This implies that $f(R) = f(R^*) = f(R')$ and $f(R, x) + f(R, y) = f(R', x) + f(R', y) = 1$ for all profiles R and R' in which the voters in N_x top-rank x and the voters in N_y top-rank y .

Step 3: For the third step, we consider two profiles R^1, R^2 , two pairs of alternatives x_1, y_1 and x_2, y_2 , and a partition of the voters (N_x, N_y) such that all voters in N_x top-rank x_1 in R^1 and x_2 in R^2 , and all voters in N_y top-rank y_1 in R^1 and y_2 in R^2 . We aim to show that $f(R^1, x_1) = f(R^2, x_2)$ and $f(R^1, y_1) = f(R^2, y_2)$ for all such profiles and alternatives. To this end, consider three distinct alternatives x, y, z and let \hat{R}^1 denote a profile in which the voters in N_x top-rank x and the voters in N_y top-rank y and second-rank z . All other alternatives can be placed arbitrarily in \hat{R}^1 . Moreover, we define \hat{R}^2 as the profile derived from \hat{R}^1 by letting all voters $i \in N_y$ swap y and z . By localizedness, we have that $f(\hat{R}^1, x) = f(\hat{R}^2, x)$, and Step 2 then entails that $f(\hat{R}^1, y) = 1 - f(\hat{R}^1, x) = 1 - f(\hat{R}^2, x) = f(\hat{R}^2, z)$. Even more, Step 2 shows that this holds for all profiles R^1 and R^2 where the voters have the same favorite alternatives as in \hat{R}^1 and \hat{R}^2 . Since a symmetric construction also works for N_x , we can now infer the claim of this step by repeating this argument.

Step 4: We are finally ready to prove the theorem. For this, we will use an induction over the number of voters n . The base case $n = 1$ is trivial as *ex post* efficiency requires that the favorite alternative of the single voter gets probability 1. We therefore assume that the theorem holds for the electorate $N = \{1, \dots, n\}$ and suppose that f is defined for the electorate $N' = \{1, \dots, n+1\}$. In the following, we will derive strongly \succsim^{SD} -strategyproof and *ex post* efficient SDSs g_\succ for n voters from f by fixing the preference relation of voter $n+1$. By the induction hypothesis, the SDSs g_\succ hence are random dictatorships. Since these SDSs also describe f , we infer as the last step that f is a random dictatorship, too.

To formalize this idea, we denote by t_\succ the top-ranked alternative of the preference relation \succ and by R^\succ a profile in which voter $n+1$ reports \succ and every other voter bottom-ranks t_\succ . Moreover, we let $\Delta_\succ = f(R^\succ, t_\succ)$ and note that Δ_\succ does not depend on the preferences of the voters in N on the alternatives in $A \setminus \{t_\succ\}$ because of localizedness. We will next show that, if $\Delta_\succ = 1$ for some $\succ \in \mathcal{L}$, then f is the dictatorship of voter $n+1$. To this end, we let x denote an alternative in $A \setminus \{t_\succ\}$ and consider the profile \hat{R} derived from R^\succ by turning x into the best alternative of the voters $i \in N$. Localizedness shows that $f(\hat{R}, t_\succ) = f(R^\succ, t_\succ) = 1$, so $f(\hat{R}, x) = 0$. Next, we deduce from Steps 2 and 3 that $f(R, y) = 1$ for all profiles R and alternatives $y, z \in A$ such that voter $n+1$ top-ranks y , and all other voters top-rank z and bottom-rank y . Finally, non-perversity entails that $f(R, y) = 1$ for all profiles R in which voter $n+1$ top-ranks alternative y , so f is the dictatorship of voter $n+1$ if $\Delta_\succ = 1$ for some preference relation $\succ \in \mathcal{L}$.

We thus suppose that $\Delta_\succ < 1$ for all $\succ \in \mathcal{L}$ and define the SDSs g_\succ for the electorate N and each $\succ \in \mathcal{L}$ by $g_\succ(R, t_\succ) = \frac{1}{1-\Delta_\succ} (f((R, \succ), t_\succ) - \Delta_\succ)$ and $g_\succ(R, x) = \frac{1}{1-\Delta_\succ} f((R, \succ), x)$ for all $x \in A \setminus \{t_\succ\}$. Our goal is to show that the

SDSs g_{\succ} are a strongly \succ^{SD} -strategyproof and *ex post* efficient because the induction hypothesis then implies that they are random dictatorships. We thus fix one of these SDSs and first show that it is a well-defined SDS. To this end, we note that $g_{\succ}(R, x) \geq 0$ for all $x \in A \setminus \{t_{\succ}\}$ and $R \in \mathcal{L}^N$ since $f((R, \succ), x) \geq 0$. Moreover, $g_{\succ}(R, t_{\succ}) \geq 0$ for all $R \in \mathcal{L}^N$ because $f((R, \succ), t_{\succ}) \geq \Delta_{\succ}$ due to non-perversity. Finally, $\sum_{x \in A} g_{\succ}(R, x) = \frac{1}{1-\Delta_{\succ}} \sum_{x \in A} f((R, \succ), x) - \frac{\Delta_{\succ}}{1-\Delta_{\succ}} = \frac{1}{1-\Delta_{\succ}} - \frac{\Delta_{\succ}}{1-\Delta_{\succ}} = 1$ for all profiles $R \in \mathcal{L}^N$, which shows that g_{\succ} is indeed well-defined.

Next, we prove that g_{\succ} is strongly \succ^{SD} -strategyproof. If this was not true, there are two profiles $R^1, R^2 \in \mathcal{L}^N$, a voter $i \in N$, and an alternative $x \in A$ such that $g_{\succ}(R^2, U(\succ_i^1, x)) > g_{\succ}(R^1, U(\succ_i^1, x))$ and $\succ_j^1 = \succ_j^2$ for all $j \in N \setminus \{i\}$. Hence,

$$\begin{aligned} f((R^2, \succ), U(\succ_i^1, x)) &= \mathbb{1}_{t_{\succ} \in U(\succ_i^1, x)} \Delta_{\succ} + (1 - \Delta_{\succ}) g_{\succ}(R^2, U(\succ_i^1, x)) \\ &> \mathbb{1}_{t_{\succ} \in U(\succ_i^1, x)} \Delta_{\succ} + (1 - \Delta_{\succ}) g_{\succ}(R^1, U(\succ_i^1, x)) \\ &= f((R^1, \succ), U(\succ_i^1, x)), \end{aligned}$$

where $\mathbb{1}_{t_{\succ} \in U(\succ_i^1, x)} = 1$ if $t_{\succ} \in U(\succ_i^1, x)$ and 0 otherwise. This contradicts the strong \succ^{SD} -strategyproofness of f , so g_{\succ} must satisfy this condition, too.

Finally, we prove that g_{\succ} is *ex post* efficient. Towards a contradiction, we suppose that there is a profile R^1 and two alternatives $x, y \in A$ such that all voters $i \in N$ prefer x to y but $g_{\succ}(R^1, y) > 0$. This implies that $f((R^1, \succ), y) > 0$, too. Next, consider the profile R^2 derived from R^1 by making x into the favorite alternative of every voter $i \in N$. Localizedness shows that $f((R^2, \succ), y) = f((R^1, \succ), y)$. Now, if $x = t_{\succ}$, this contradicts the *ex post* efficiency of f as every voter top-ranks t_{\succ} in (R^2, \succ) . Furthermore, if $t_{\succ} \notin \{x, y\}$, this contradicts Step 2 as $f((R^2, \succ), t_{\succ}) + f((R^2, \succ), x) \neq 1$. Finally, if $y = t_{\succ}$, we consider the profile R^3 in which the voters in N top-rank x and bottom-rank t_{\succ} . By Step 2, it holds that $f((R^2, \succ)) = f((R^3, \succ))$ as we did not change the first-ranked alternatives of the voters. Moreover, localizedness shows that $f((R^3, \succ), t_{\succ}) = f(R^{\succ}, t_{\succ}) = \Delta_{\succ}$ because the voters $i \in N$ bottom-rank t_{\succ} in both R^3 and R^{\succ} . Chaining our equations shows that $f(R^1, t_{\succ}) = f(R^{\succ}, t_{\succ}) = \Delta_{\succ}$ since $y = t_{\succ}$. Hence, $g_{\succ}(R^1, t_{\succ}) = 0$ by definition, which contradicts the assumption that $g_{\succ}(R^1, y) > 0$. Since we exhausted all cases, g_{\succ} satisfies *ex post* efficiency.

By the induction hypothesis, we now infer that the SDSs g_{\succ} are random dictatorships, i.e., there are values $\gamma_i^{\succ} \geq 0$ for all voters $i \in N$ and $\succ \in \mathcal{L}$ such that $\sum_{i \in N} \gamma_i^{\succ} = 1$ and $g_{\succ}(R) = \sum_{i \in N} \gamma_i^{\succ} d_i(R)$ for all profiles R . It thus follows for all profiles $R \in \mathcal{L}^N$ and preference relations $\succ \in \mathcal{L}$ that $f((R, \succ), t_{\succ}) = \Delta_{\succ} + (1 - \Delta_{\succ}) g_{\succ}(R, t_{\succ}) = \Delta_{\succ} + (1 - \Delta_{\succ}) \sum_{i \in N} \gamma_i^{\succ} d_i(R, t_{\succ})$ and $f((R, \succ), x) = (1 - \Delta_{\succ}) g_{\succ}(R, x) = (1 - \Delta_{\succ}) \sum_{i \in N} \gamma_i^{\succ} d_i(R, x)$ for all $x \in A \setminus \{t_{\succ}\}$. So, we only need to prove that $\Delta_{\succ} = \Delta_{\succ'}$ and $\gamma_i^{\succ} = \gamma_i^{\succ'}$ for all $i \in N$ and $\succ, \succ' \in \mathcal{L}$ to show that f is a random dictatorship. To this end, consider two preference relations \succ^1 and \succ^2 and let $R^k \in \mathcal{L}^N$ for $k \in \{1, 2\}$ denote profiles in which all voters $i \in N$ bottom-rank t_{\succ^k} and top-rank another alternative x_k . By localizedness, we derive that $f((R^k, \succ^k), t_{\succ^k}) = f(R^{\succ^k}, t_{\succ^k}) = \Delta_{\succ^k}$. Consequently, Step 3 shows that $f((R^1, \succ^1), t_{\succ^1}) = f((R^2, \succ^2), t_{\succ^2})$, so $\Delta_{\succ^1} = \Delta_{\succ^2}$. Next, suppose for contradiction that there is a voter $i \in N$ such that $\gamma_i^{\succ^1} \neq \gamma_i^{\succ^2}$. This means that there are two preference relations \succ^3 and \succ^4 such that $\gamma_i^{\succ^3} \neq \gamma_i^{\succ^4}$ and \succ^3 differs from \succ^4 only by swapping two alternatives x and y . Now, con-

sider the profile $R \in \mathcal{L}^N$ in which all voters $j \in N \setminus \{i\}$ top-rank alternative x and voter i top-ranks an alternative $z \in A \setminus \{x, y\}$. Localizedness entails that $(1 - \Delta_{\succ^3})\gamma_i^{\succ^3} = f((R, \succ^3), z) = f((R, \succ^4), z) = (1 - \Delta_{\succ^4})\gamma_i^{\succ^4}$. Since $\Delta_{\succ^3} = \Delta_{\succ^4} < 1$, it follows that $\gamma_i^{\succ^3} = \gamma_i^{\succ^4}$. Hence, we conclude that f is the random dictatorship defined by $f(R) = \Delta d_{n+1}(R) + (1 - \Delta) \sum_{i \in N} \gamma_i d_i$, where $\Delta = \Delta_{\succ}$ and $\gamma_i = \gamma_i^{\succ}(1 - \Delta)$ for an arbitrary preference relation $\succ \in \mathcal{L}$. \square

3.1.2 Analysis of the Set of Strongly \succsim^{SD} -Strategyproof SDSs

We next turn to the escape routes to the random dictatorship theorem, and more generally to the Gibbard-Satterthwaite theorem, and investigate as the first approach whether there are attractive and strongly \succsim^{SD} -strategyproof SDSs other than random dictatorships. Surprisingly, this question has not attracted much attention: after the characterizations of the set of strongly \succsim^{SD} -strategyproof SDSs by Gibbard (1977) and Barberà (1979b), only few authors have investigated these SDSs in more detail (e.g., Procaccia, 2010; Filos-Ratsikas and Miltersen, 2014; Ebadian et al., 2022). Moreover, these papers typically analyze modern properties, such as the distortion of randomized voting rules, instead of classical axioms. In Publication 1, we thus analyze strongly \succsim^{SD} -strategyproof SDSs with respect to relaxations of *ex post* efficiency and Condorcet-consistency.

In more detail, we first note that every strongly \succsim^{SD} -strategyproof SDS other than a random dictatorship has to fail *ex post* efficiency because of Theorem 3.2. However, this does not exclude the possibility that there are strongly \succsim^{SD} -strategyproof SDSs that assign negligibly small probabilities to Pareto-dominated alternatives and that are otherwise axiomatically attractive. For instance, if an SDS only chooses a Pareto-dominated alternative with a probability of 10^{-100} , then this rule will effectively never select a Pareto-dominated alternative and it is thus as efficient as an *ex post* efficient SDS for all practical matters. As the first question, we hence ask whether there are appealing strongly \succsim^{SD} -strategyproof SDSs that only assign small probabilities to Pareto-dominated alternatives. To formalize this idea, we introduce the following relaxation of *ex post* efficiency: an SDS f is β -*ex post* efficient if $f(R, x) \leq \beta$ for all profiles R and alternatives $x \in A$ that are Pareto-dominated in R . That is, an SDS is β -*ex post* efficient if it assigns a probability of at most β to each Pareto-dominated alternative. For example, random dictatorships are 0-*ex post* efficient, and the uniform lottery rule f_U , which always chooses every alternative with probability $\frac{1}{m}$, is $\frac{1}{m}$ -*ex post* efficient.

Unfortunately, it turns out that the random dictatorship theorem is robust with respect to relaxing *ex post* efficiency: if an SDS is almost *ex post* efficient, it is almost a random dictatorship. To formalize this observation, we say an SDS f is γ -randomly dictatorial if $\gamma \in [0, 1]$ is the maximal value for which there is a random dictatorship d and another strongly \succsim^{SD} -strategyproof SDS g such that $f(R) = \gamma d(R) + (1 - \gamma)g(R)$ for all profiles R . We note that the requirement that g is strongly \succsim^{SD} -strategyproof is necessary as otherwise, SDSs that intuitively are completely non-randomly dictatorial are not 0-randomly dictatorial. For example, the uniform lottery rule can be represented as $f_U(R) = \frac{1}{m}d_1(R) + \frac{m-1}{m}g(R)$, where d_1 is the dictatorship of voter 1 and g is the manipulable SDS that randomizes uniformly over all alternatives but the favorite one of voter 1. By contrast, the

uniform lottery rule is 0-randomly dictatorial according to our definition. Moreover, random dictatorships are 1-randomly dictatorial, and maximal lottery rules are not γ -randomly dictatorial for any $\gamma \in [0, 1]$ since these SDSs fail strong \succsim^{SD} -strategyproofness but γ -random dictatorships satisfy this axiom by definition. As a less trivial example, we note that the randomized Borda rule f_{RB} is $\frac{2}{m(m-1)}$ -randomly dictatorial. This follows as $f_{RB}(R) = \frac{2}{m(m-1)}f_{RD}(R) + (1 - \frac{2}{m(m-1)})g(R)$ for all profiles R , where f_{RD} is the uniform random dictatorship and g is the SDS that randomizes proportional to the scoring function $s(R, x) = s_{Borda}(R, x) - s_{PL}(R, x)$. This representation works as f_{RD} randomizes proportional to the plurality scores of the alternatives. It can moreover be shown that f_{RB} cannot be γ -randomly dictatorial for $\gamma > \frac{2}{m(m-1)}$.

Based on β -*ex post* efficiency and γ -random dictatorships, we prove the following continuous strengthening of the random dictatorship theorem in Publication 1.

THEOREM 3.3 (Brandt et al., 2022b)

For every $\epsilon \in [0, 1]$, every strongly \succsim^{SD} -strategyproof and $\frac{1-\epsilon}{m}$ -*ex post* efficient SDS on \mathcal{L}^N is γ -randomly dictatorial for $\gamma \geq \epsilon$ if $m \geq 3$.

When $\epsilon = 1$, Theorem 3.3 implies the random dictatorship theorem. Moreover, if ϵ is close to 1, the considered SDS is close to a random dictatorship. By contrast, the further away an SDS is from a random dictatorship, the less *ex post* efficient it is. In particular, it follows from Theorem 3.3 that every 0-randomly dictatorial and strategyproof SDS is at least $\frac{1}{m}$ -*ex post* efficient and thus as inefficient as the uniform lottery rule. In summary, this means that relaxing *ex post* efficiency does not lead to more attractive strongly \succsim^{SD} -strategyproof SDSs than random dictatorships as all these rules are either very close to random dictatorships or fail β -*ex post* efficiency for a large β .

In light of Theorem 3.3, it is necessary to give up *ex post* efficiency when trying to find appealing strongly \succsim^{SD} -strategyproof SDSs other than random dictatorships. We therefore turn to another objective: we next aim to find strongly \succsim^{SD} -strategyproof SDSs that guarantee Condorcet winners as much probability as possible. Unfortunately, the results by Gibbard (1977) entail that no strongly \succsim^{SD} -strategyproof SDS satisfies Condorcet-consistency. We thus relax this axiom as follows: an SDS f is α -Condorcet-consistent if $f(R, x) \geq \alpha$ for all profiles R in which alternative x is the Condorcet winner. Less formally, α -Condorcet-consistency guarantees that Condorcet winners are chosen with probability at least α . We observe that random dictatorships may assign probability 0 to the Condorcet winner, so these SDSs are 0-Condorcet-consistent and not sensible for our new objective. As another example, the uniform lottery rule is $\frac{1}{m}$ -Condorcet-consistent.

As the second main contribution of Publication 1, we show that the randomized Copeland rule f_{RC} maximizes the α -Condorcet-consistency among all strongly \succsim^{SD} -strategyproof SDSs. Recall that this SDS randomizes proportional to the Copeland scores, i.e., $f_{RC}(R, x) = \frac{2}{m(m-1)}s_{Copeland}(R, x)$ for all profiles R and alternatives $x \in A$. It follows from this definition that $f_{RC}(R, x) = \frac{2}{m}$ whenever x is the Condorcet winner in R as its Copeland score is then $m-1$. This shows that the randomized Copeland rule is $\frac{2}{m}$ -Condorcet-consistent and we characterize this rule based on this condition, strong \succsim^{SD} -strategyproofness, anonymity, and neutrality. Moreover, we prove that no strongly \succsim^{SD} -strategyproof SDS is α -Condorcet-consistent for $\alpha > \frac{2}{m}$, even when dropping anonymity and neutrality.

THEOREM 3.4 (Brandt et al., 2022b)

The randomized Copeland rule is the only strongly \succsim^{SD} -strategyproof SDS on \mathcal{L}^N that satisfies anonymity, neutrality, and $\frac{2}{m}$ -Condorcet-consistency if $m \geq 3$ and $n \geq 3$. Moreover, no strongly \succsim^{SD} -strategyproof SDS on \mathcal{L}^N is α -Condorcet-consistent for $\alpha > \frac{2}{m}$ if $n \geq 3$.

This theorem shows that the randomized Copeland rule can be seen as counterpart to random dictatorships when optimizing for α -Condorcet-consistency rather than β -*ex post* efficiency. Or, put differently, if the main criterion for evaluating SDSs is how likely they choose the Condorcet winner, the randomized Copeland rule is more desirable than random dictatorships. However, we also have to acknowledge that the randomized Copeland rule is only “twice as Condorcet-consistent” as the uniform lottery rule and severely fails *ex post*-efficiency. Hence, one may also interpret Theorem 3.4 as a negative result: strong \succsim^{SD} -strategyproofness does not allow for a reasonable amount of Condorcet-consistency.

3.1.3 Strongly \succsim^{SD} -Strategyproof SDSs on the Condorcet Domain

As our second approach to circumvent the random dictatorship theorem, we weaken the assumption that SDSs are defined for all preference profiles in \mathcal{L}^N and study SDSs for a restricted domain of preference profiles $\mathcal{D} \subseteq \mathcal{L}^N$. The motivation for this approach is that often not all preference profiles are likely or plausible in practice, and we can thus omit some of them from our analysis. Moreover, since it tends to be easier to satisfy desirable axioms for smaller domains, this approach is a promising escape route from impossibility theorems. All the results in this section are taken from Publication 2.

The study of restricted domains has attracted significant attention in the realm of social choice. In particular, Moulin (1980) has shown in a seminal paper that there are attractive strategyproof SCFs on the domain of single-peaked preference profiles, thus giving an appealing escape route to the Gibbard-Satterthwaite theorem. Similar positive results have been shown for a multitude of other domains (e.g., Barberà et al., 1993; Nehring and Puppe, 2007; Saporiti, 2009; Barberà et al., 2012; Chatterji et al., 2013). On the other hand, restricted domains have also been used to strengthen the Gibbard-Satterthwaite theorem by showing that this result also holds for smaller domains (e.g., Aswal et al., 2003; Sato, 2010; Gopakumar and Roy, 2018). In more recent works, the positive and negative results converge by giving exact criteria for deciding whether strategyproof and non-dictatorial SCFs exist on a domain (Chatterji and Sen, 2011; Chatterji et al., 2013; Roy and Storcken, 2019; Chatterji and Zeng, 2023). While similar results have been shown for SDSs, this setting is less understood. For instance, Ehlers et al. (2002) characterize the set of strongly \succsim^{SD} -strategyproof and non-imposing SDSs on the domain of single-peaked preference profiles (see also Peters et al., 2014; Pycia and Ünver, 2015). Moreover, several other domains have been discussed (Chatterji et al., 2014; Peters et al., 2017; Chatterji and Zeng, 2018; Roy and Sadhukhan, 2020; Peters et al., 2021), showing either that the random dictatorship theorem holds on these domains or that there are other attractive and strongly \succsim^{SD} -strategyproof SDSs. We refer to Roy et al. (2022) for an overview of recent work on this topic.

In this section, we aim to identify maximal domains that still allow for attractive strongly \succsim^{SD} -strategyproof SDSs. To this end, we will study the *Condorcet domain* $\mathcal{D}_C = \{R \in \mathcal{L}^N : \text{there is a Condorcet winner in } R\}$, which is precisely the set of preference profiles that have a Condorcet winner. This domain is significant as there is strong empirical evidence that most real-world elections admit a Condorcet winner (Regenwetter et al., 2006; Laslier, 2010; Gehrlein and Lepelley, 2011; Boehmer and Schaar, 2023). Moreover, it is well-known that the Condorcet rule, which always assigns probability 1 to the Condorcet winner, is strongly \succsim^{SD} -strategyproof on this domain. This demonstrates that the Condorcet domain allows for positive results. Even more, Campbell and Kelly (2003) have shown that, if the number of voters n is odd, the Condorcet rule is the only strategyproof, non-imposing, and non-dictatorial SCF on the Condorcet domain (see also Merrill, 2011; Campbell and Kelly, 2015, 2016).

As our first theorem in Publication 2, we revisit the result of Campbell and Kelly (2003) for randomized voting rules and characterize the set of non-imposing and strongly \succsim^{SD} -strategyproof SDSs on the Condorcet domain \mathcal{D}_C for odd n .

THEOREM 3.5 (Brandt et al., 2023c)

Assume n is odd and $m \geq 3$. An SDS on \mathcal{D}_C is non-imposing and strongly \succsim^{SD} -strategyproof if and only if it is a convex combination of the Condorcet rule and a random dictatorship.

Note that Theorem 3.5 directly generalizes the result of Campbell and Kelly (2003) to SDSs: the set of strongly \succsim^{SD} -strategyproof and non-imposing SDSs on \mathcal{D}_C is simply the convex combination of all strategyproof and non-imposing SCFs. This means that the Condorcet domain satisfies the deterministic extreme point property (see, e.g., Pycia and Ünver, 2015; Roy and Sadhukhan, 2020) if n is odd as every strongly \succsim^{SD} -strategyproof and non-imposing SDS can be represented as a mixture of deterministic SCFs that satisfy these properties. Moreover, Theorem 3.5 shows that the deterministic Condorcet rule is the most attractive strongly \succsim^{SD} -strategyproof SDS on the Condorcet domain, even if we allow for randomization. In particular, we can easily characterize this SDS based on Theorem 3.5 and the notion of γ -random dictatorships introduced in Section 3.1.2: the Condorcet rule is the only non-imposing, strongly \succsim^{SD} -strategyproof, and 0-randomly dictatorial SDS on the Condorcet domain \mathcal{D}_C if n is odd and $m \geq 3$.

As our next question, we investigate whether there is a superset of the Condorcet domain that still allows for attractive and strongly \succsim^{SD} -strategyproof SDSs. To this end, we first note that if we add a profile R to the Condorcet domain that differs in at least two voters from every profile $R' \in \mathcal{D}_C$, the resulting domain allows for strategyproof SDSs other than random dictatorships because no voter can manipulate to R . To avoid such artificial situations, we will focus on connected domains. For defining these, we first need to define ad-paths in a domain \mathcal{D} : an *ad-path* between two profiles $R, R' \in \mathcal{D}$ is a sequence of preference profiles R^1, \dots, R^k such that $R^1 = R$, $R^k = R'$, $R^i \in \mathcal{D}$ for all $i \in \{1, \dots, k\}$, and the profile R^{i+1} is derived from the profile R^i by only swapping two adjacent alternatives in the preference relation of a single voter. Then, a domain \mathcal{D} is *connected* if (i) for all profiles $R, R' \in \mathcal{D}$ there is an ad-path from R to R' and (ii) for all profiles $R, R' \in \mathcal{D}$ and all alternatives $x \in A$ such that $U(\succ_i, x) = U(\succ'_i, x)$ for all $i \in N$, there is an ad-path from R to R' along which x is never moved. We recall here that $U(\succ_i, x)$ denotes

the set of alternatives that voter i prefers to x . Our notion of connectedness is rather mild and, e.g., weaker than Sato's non-restoration property (Sato, 2013) and Nehring's connectedness notion (Nehring, 2000).

We note that the Condorcet domain is connected if n is odd. Even more, we will show next that the Condorcet domain is a maximal connected domain that allows for attractive strongly \succsim^{SD} -strategyproof SDSs if n is odd as only random dictatorships satisfy non-imposition and strong \succsim^{SD} -strategyproofness on connected supersets of the Condorcet domain. We note that the subsequent theorem also strengthens the random dictatorship theorem by showing that this result holds for strongly restricted domains.

THEOREM 3.6 (Brandt et al., 2023c)

Assume n is odd, $m \geq 3$, and let \mathcal{D} denote a connected domain with $\mathcal{D}_C \subsetneq \mathcal{D}$. An SDS on \mathcal{D} is non-imposing and strongly \succsim^{SD} -strategyproof if and only if it is a random dictatorship.

Finally, we observe that both Theorems 3.5 and 3.6 only discuss the case where the number of voters is odd. Indeed, if n is even, there are larger domains than the Condorcet domain that allow for strongly \succsim^{SD} -strategyproof SDSs other than random dictatorships as a single voter cannot change the Condorcet winner. We thus consider a slight extension of the Condorcet domain: for an arbitrary preference relation $\triangleleft \in \mathcal{L}$, we define the tie-breaking Condorcet domain $\mathcal{D}_C^{\triangleleft} = \{R \in \mathcal{L}^N : \text{there is a Condorcet winner in } (R, \triangleleft)\}$ as the set of profiles that have a Condorcet winner after adding an additional voter with the preference relation \triangleleft . If n is even, every tie-breaking Condorcet domain is connected and $\mathcal{D}_C \subseteq \mathcal{D}_C^{\triangleleft}$. Moreover, we define the tie-breaking Condorcet rule as the SDS that assigns probability 1 to the Condorcet winner in (R, \triangleleft) .

By focusing on tie-breaking Condorcet domains, we now transfer Theorems 3.5 and 3.6 to the case of even n .

THEOREM 3.7 (Brandt et al., 2023c)

Assume $n \geq 4$ is even, $m \geq 3$, let $\triangleleft \in \mathcal{L}$ be an arbitrary preference relation, and let \mathcal{D} denote a connected domain.

1. Suppose $\mathcal{D} = \mathcal{D}_C^{\triangleleft}$. An SDS on \mathcal{D} is non-imposing and strongly \succsim^{SD} -strategyproof if and only if it is a convex combination of the tie-breaking Condorcet rule and a random dictatorship.
2. Suppose $\mathcal{D}_C^{\triangleleft} \subsetneq \mathcal{D}$. An SDS on \mathcal{D} is non-imposing and strongly \succsim^{SD} -strategyproof if and only if it is a random dictatorship.

In summary, our results show that always choosing the Condorcet winner is the most appealing strategyproof voting rule when Condorcet winners are guaranteed to exist. In particular, this observation even holds when allowing for randomization. Moreover, our result demonstrate that the Condorcet domain is effectively a maximal domain that allows for attractive strategyproof social choice.

3.1.4 Analysis of Weaker Strategyproofness Notions

As our last escape route to the random dictatorship theorem, we analyze the consequences of replacing strong \succsim^{SD} -strategyproofness with weaker strategyproofness

notions. The motivation for this approach comes from the fact that, while strong \succsim^{SD} -strategyproofness guarantees the mechanism designer that no voter can manipulate if her preferences over lotteries are represented by a vNM utility function, this may not be necessary in practice. For instance, not all vNM utility functions may be plausible for a given election or vNM utility functions may not be the right model to describe the voters' preferences over lotteries (cf. Section 2.3.1). In such situations, strong \succsim^{SD} -strategyproofness is unnecessarily restrictive and we thus investigate weaker strategyproofness notions in Publications 3 and 4.

We note that there is already some work on analyzing weaker notions of strategyproofness for randomized social choice and we refer to Brandt (2017) for an overview. For instance, Aziz et al. (2018) and Brandl et al. (2018, 2021) investigate the compatibility of weak strategyproofness notions with various efficiency notions. However, these authors consider the case of weak preferences and then infer strong negative results. By contrast, it is easy to see that weak strategyproofness notions, such as weak \succsim^{SD} -strategyproofness and weak \succsim^{DD} -strategyproofness, allow for more positive results when the voters' preferences are strict. For instance, maximal lottery rules are weakly \succsim^{DD} -strategyproof (when using suitable tie-breaking for the case that $|ML(R)| > 1$) and the randomized Condorcet rule, which chooses the Condorcet winner with probability 1 if it exists and otherwise returns the uniform lottery, is weakly \succsim^{SD} -strategyproof. By contrast, negative results typically require restrictive strategyproofness notions such as strong \succsim^{SD} -strategyproofness (see, e.g., Gibbard, 1977; Benoît, 2002) when voters report strict preferences. Consequently, there is a gap between the strategyproofness notions that lead to positive results and those that lead to negative ones.

To close this gap, we consider intermediate strategyproofness notions in this section. In more detail, we study the concept of strong \succsim^U -strategyproofness in Publication 3 and \succsim^{PC} -strategyproofness in Publication 4. Unfortunately, we mainly infer negative results based on these notions, which demonstrates how far the incompatibility of strategyproofness with other axioms goes.

Strong \succsim^U -Strategyproofness

The idea of strong \succsim^U -strategyproofness is that some vNM utility functions may not be plausible in practice, and we thus do not need to consider them when analyzing strategyproofness. This leads to the question of the vNM utility functions for which SDSs can be strategyproof while satisfying further desirable axioms.

Given our previous results, perhaps the most apparent question is whether Condorcet-consistent SDSs can be strongly \succsim^U -strategyproof for a large set of vNM utility functions U . In Publication 3, we answer this question negatively by proving a strong impossibility theorem. To state this result, we recall that $\hat{\mathcal{U}}$ is the set of vNM utility functions that are consistent with strict preference relations. Moreover, given a utility function $u \in \hat{\mathcal{U}}$, u^Π is the set of utility functions that can be derived by permuting u , i.e., u^Π assigns to every preference relation a single canonical utility function. We recall that the corresponding set extension is complete and we therefore write \succsim^{u^Π} -strategyproofness without strong or weak.

THEOREM 3.8 (Lederer, 2021)

No Condorcet-consistent SDS on \mathcal{L}^N satisfies \succsim^{u^Π} -strategyproofness for any vNM utility function $u \in \hat{\mathcal{U}}$ if $m \geq 4$ and $n \geq 10$.

Since \succsim^{u^Π} -strategyproofness is weaker than strong \succsim^U -strategyproofness for all sets of vNM utility functions U with $u^\Pi \subseteq U$, this result shows that Condorcet-consistency and strong \succsim^U -strategyproofness are already in conflict if the set of utility functions U satisfies minimal symmetry properties. Hence, these concepts are completely incompatible. Moreover, Theorem 3.8 complements the fact that, e.g., the randomized Condorcet rule is weakly \succsim^{SD} -strategyproof because weak \succsim^{SD} -strategyproofness means that for every possible manipulation, there is a vNM utility function such that the expected utility of the manipulator decreases. This means that, while such a utility function always exists for the randomized Condorcet rule, it cannot be the same for each possible manipulation.

Since Theorem 3.8 proves that there are no Condorcet-consistent and strongly \succsim^U -strategyproof SDSs, we will next consider a weaker decisiveness axiom. To this end, we introduce the concept of k -unanimity: an SDS f is k -unanimous if $f(R, x) = 1$ for all profiles in which at least $n - k$ voters prefer x the most. Less formally, k -unanimity postulates that if all but k voters agree on a favorite alternative, this option should be chosen with probability 1. Thus, 0-unanimity requires that an alternative is selected with probability 1 if it is top-ranked by all voters and is satisfied by random dictatorships. By contrast, it is known that 1-unanimity conflicts with strong \succsim^{SD} -strategyproofness (Benoît, 2002) and we analyze therefore the tradeoff between strong \succsim^U -strategyproofness and k -unanimity.

To this end, we first show that there are k -unanimous SDSs for every $k < n/2$ that are strongly \succsim^U -strategyproof for a large set of vNM utility functions U . In more detail, we consider the following variant of the uniform random dictatorship: the SDS f_{RD}^k chooses an alternative with probability 1 if it is top-ranked by at least $n - k$ voters and returns the outcome of the uniform random dictatorship otherwise. Clearly, this SDS is k -unanimous by definition and we show in Publication 3 that it is strongly \succsim^U -strategyproof for the set $U = \{u \in \hat{\mathcal{U}} : u(1) - u(2) \geq k(u(2) - u(m))\}$ (here, $u(\ell)$ denotes the value assigned to the alternative with the ℓ -th highest utility). Even more, this set is maximal in the sense that f_{RD}^k fails $\succsim^{\{u\}}$ -strategyproofness for every vNM utility function $u \in \hat{\mathcal{U}} \setminus U$.

A natural follow-up question is whether the SDSs f_{RD}^k optimally solve the tradeoff between k -unanimity and strong \succsim^U -strategyproofness or whether there are k -unanimous SDSs that satisfy strong \succsim^U -strategyproofness for larger sets of utility functions. We partially answer this question by focusing on the class of rank-based SDSs. To define these SDSs, we introduce the rank vector $r^*(R, x)$ of an alternative x , which contains the ranks $r(\succ_i, x)$ for every voter $i \in N$ in increasing order. More formally, $r^*(R, x) = (r(\succ_{i_1}, x), \dots, r(\succ_{i_n}, x))$ where the voters are ordered such that $r(\succ_{i_j}, x) \leq r(\succ_{i_{j+1}}, x)$ for all $j \in \{1, \dots, n-1\}$. We then say that an SDS f is *rank-based* if $f(R) = f(R')$ for all profiles R, R' such that $r^*(R, x) = r^*(R', x)$ for all $x \in A$. Clearly, all SDSs f_{RD}^k are rank-based and the randomized Borda rule also satisfies this condition. As the next theorem shows, rank-based SDSs cannot do significantly better than the SDSs f_{RD}^k with respect to the tradeoff between k -unanimity and strong \succsim^U -strategyproofness.

THEOREM 3.9 (Lederer, 2021)

Let $k \in \mathbb{N}$ such that $0 < k < n/2$ and let $u \in \hat{\mathcal{U}}$ denote a vNM utility function such that $u(1) - u(2) < \sum_{i=\max(3, m-k+1)}^m u(2) - u(i)$. No rank-based SDS on \mathcal{L}^N is both k -unanimous and \succsim^{u^Π} -strategyproof if $m \geq 3$ and $n \geq 3$.

We note that the SDSs f_{RD}^k show that the bound in this theorem is tight since these SDSs are \succsim^{u^π} -strategyproof for every vNM utility function u with $u(1) - u(2) \geq k(u(2) - u(m))$. As a consequence, we can find a utility function $u \in \hat{\mathcal{U}}$ for every $k \leq m-2$ and $\epsilon > 0$ such that $k(u(2) - u(m)) \leq u(1) - u(2) < \sum_{i=\max(3, m-k+1)}^m u(2) - u(i) + \epsilon$. Thus, f_{RD}^k satisfies all axioms of Theorem 3.9 for some vNM utility function when relaxing the bound on u . Conversely, there is a gap between the set \mathcal{U} for which f_{RD}^k is strongly \succsim^U -strategyproof and for which the impossibility holds. While we give also another SDS that satisfies k -unanimity for $k < \frac{n}{2}$ and strong \succsim^U -strategyproofness for a large set of vNM utility functions \mathcal{U} in Publication 3, there is still some room to improve the bound.

\succsim^{PC} -Strategyproofness

Our second weakening of strong \succsim^{SD} -strategyproofness is \succsim^{PC} -strategyproofness. Note that, since the PC extension is complete, we again drop the prefix strong for this strategyproofness notion. The study of \succsim^{PC} -strategyproofness is motivated by two observations: firstly, the representation of preferences over lotteries via vNM utility functions has come under scrutiny due to empirical research (e.g., Allais, 1953; Kahneman and Tversky, 1979; Machina, 1989; Anand, 2009), and it is therefore reasonable to study lottery extensions that are founded on a different normative basis. Secondly, the PC extension has led to strong positive results in randomized social choice. In particular, Brandl and Brandt (2020) show that this lottery extension can be used to circumvent Arrow's impossibility, and Brandl et al. (2019) have demonstrated how to escape Moulin's no-show paradox based on the PC extension. A natural follow-up question to these results is whether we can also escape the random dictatorship theorem based on this lottery extension. In this section, we answer this question in the negative by proving two strong impossibilities. All results in this section are taken from Publication 4 and answer open questions of Brandt (2017).

Similar to strong \succsim^U -strategyproofness, we first ask whether there is an SDS that satisfies both \succsim^{PC} -strategyproofness and Condorcet-consistency. Unfortunately, we prove that this is not the case.

THEOREM 3.10 (Brandt et al., 2023b)

No Condorcet-consistent SDS on \mathcal{L}^N satisfies \succsim^{PC} -strategyproofness if $m \geq 4$ and $n \geq 5$ is odd.

Since Condorcet-consistency conflicts with \succsim^{PC} -strategyproofness, we will once again modify our objective. To this end, we follow the approach of Aziz et al. (2018) and Brandl et al. (2018) and study the compatibility of strategyproofness and efficiency. Because we use the PC extension to compare lotteries, it is a natural choice to consider \succsim^{PC} -efficiency. To introduce this axiom, we generalize the idea of Pareto-dominance to lotteries and thus say that a lottery p \succsim^{PC} -dominates a lottery q in a profile R if $p \succsim_i^{PC} q$ for all $i \in N$ and $p \succ_i^{PC} q$ for some voter $i \in N$. Moreover, a lottery p is \succsim^{PC} -efficient in a profile R if it is not \succsim^{PC} -dominated by any other lottery. Less formally, a lottery is \succsim^{PC} -efficient if we cannot make a voter better off without making another voter worse off according to the PC extension. Finally, an SDS f is \succsim^{PC} -efficient if $f(R)$ is \succsim^{PC} -efficient for every profile R . We note that one can analogously define \succsim^X -efficiency for every lottery extension \succsim^X .

It is known that maximal lottery rules satisfy \succsim^{PC} -efficiency but fail \succsim^{PC} -strategy-proofness, and that random dictatorships fail \succsim^{PC} -efficiency but satisfy \succsim^{PC} -strategy-proofness (Brandt, 2017). The question of whether there is an SDS that satisfies both \succsim^{PC} -strategyproofness and \succsim^{PC} -efficiency is thus equivalent to the question of whether there is an SDS that combines the advantages of maximal lottery rules and random dictatorships. Unfortunately, it turns out that no such SDS exists when additionally requiring anonymity and neutrality.

THEOREM 3.11 (Brandt et al., 2023b)

No anonymous and neutral SDS on \mathcal{L}^N satisfies both \succsim^{PC} -strategyproofness and \succsim^{PC} -efficiency if $m \geq 4$ and $n \geq 7$.

We note that Theorem 3.11 strengthens a result of Aziz et al. (2018) who have shown an analogous claim for weak preferences. However, when we allow weak preferences, much stronger impossibilities are known (see, e.g., Section 3.2.2 or the results by Brandl et al. (2018, 2021)).

Finally, to derive more positive results, we consider a weakening of \succsim^{PC} -strategy-proofness in Publication 4. To this end, we say a voter \succsim^{PC1} -prefers a lottery p to a lottery q if $p \succsim^{PC} q$ and p or q are degenerate (i.e., p or q assigns probability 1 to a single alternative). The $PC1$ extension is sparser than the PC extension and only allows for particularly simple PC comparisons between lotteries. Based on this lottery extension, we finally derive positive results: all maximal lottery rules that only return a degenerate lottery if there is a unique maximal lottery satisfy \succsim^{PC} -efficiency, Condorcet-consistency, and weak \succsim^{PC1} -strategyproofness.

3.2 RESULTS FOR SOCIAL CHOICE CORRESPONDENCES

In this section, we turn to the study of social choice correspondences and thus analyze whether it is possible to obtain strategyproof voting rules when choosing sets of alternatives rather than single winners. In more detail, we will first derive a strong possibility result in Section 3.2.1 by giving a characterization of the top cycle based on weak \succsim^F -strategyproofness and several auxiliary conditions for the case of strict preferences. By contrast, when voters are allowed to report weak preferences, we will show in Section 3.2.2 that even the very mild notion of weak \succsim^K -strategyproofness precludes the existence of attractive SCCs.

Before presenting our results, we note that there is already a large body of literature that investigates strategyproofness for social choice correspondences. Roughly, these works can be divided into four categories. Firstly, early works rely on mild strategyproofness conditions but require strong additional conditions (e.g., Barberà, 1977a,b; Kelly, 1977; MacIntyre and Pattanaik, 1981; Bandyopadhyay, 1983). For instance, Barberà (1977b) shows that all positively responsive SCCs violate a variant of weak \succsim^K -strategyproofness. However, of all commonly studied SCCs, only the Borda rule and Black's rule (1958) satisfy positive responsiveness, so this result only affects a narrow class of SCCs. The second line of work derives negative results similar to the Gibbard-Satterthwaite theorem by using strong strategyproofness notions (e.g., Duggan and Schwartz, 2000; Ching and Zhou, 2002; Benoît, 2002; Sato, 2008, 2014). For instance, Ching and Zhou (2002) show that every strongly \succsim^F -strategyproof SCC is either dictatorial or constant. One drawback

of these results is that the considered strategyproofness notions rely on strong assumptions that may be difficult to motivate in practice. The third line of work shows that mild strategyproofness notions allow for positive results (Gärdenfors, 1976; Feldman, 1979a; Nehring, 2000; Brandt, 2015). For instance, Brandt (2015) shows that numerous SCCs that only rely on the majority relation to compute the outcome (such as the top cycle) are weakly \succsim^K -strategyproof. We note that these positive results only hold for strict preferences and often break down when allowing voters to report ties between alternatives (Brandt, 2015; Brandt et al., 2022c). Finally, a last stream of research assumes that voters express preferences over all sets of alternatives but that not all preference relations are valid (e.g., Barberà et al., 2001; Özyurt and Sanver, 2009). For instance, Barberà et al. (2001) show a remarkable variant of the Gibbard-Satterthwaite theorem when the voters' preferences on sets have to obey Fishburn's extension. We refer to Brandt et al. (2022c) for a more detailed discussion of the literature.

3.2.1 Characterization of the Top Cycle based on Weak \succsim^F -Strategyproofness

In this section, we will demonstrate how to escape the Gibbard-Satterthwaite theorem by moving from SCFs to SCCs. To this end, we first observe that even rather mild strategyproofness notions are only satisfied by few commonly studied SCCs. For instance, among all SCCs named in this thesis, only the Pareto rule, the omninomination rule, and the top cycle satisfy weak \succsim^F -strategyproofness. By contrast, most positional scoring rules (e.g., the Plurality rule and the Borda rule) and all SCCs that satisfy Pareto-optimality and only rely on the majority relation to compute the winners fail this condition (Brandt and Geist, 2016). Motivated by these insights, we aim in Publication 5 to better understand the set of weakly \succsim^F -strategyproof SCCs. In more detail, the main contribution of this paper is a characterization of the top cycle based on weak \succsim^F -strategyproofness and other mild axioms. This result essentially turns the Gibbard-Satterthwaite impossibility into a characterization of the top cycle by moving from SCFs to SCCs.

To formally state this characterization, we first have to slightly change our model: while we defined SCCs for a fixed electorate in Chapter 2, we allow in this section a variable electorate. To this end, we let $\mathbb{N} = \{1, 2, 3, \dots\}$ denote an infinite set of voters and $\mathcal{F}(\mathbb{N})$ is the set of non-empty and finite subsets of \mathbb{N} . Intuitively, \mathbb{N} is the set of all possible voters and an element $N \in \mathcal{F}(\mathbb{N})$ is a concrete electorate. For this section, we assume that the domain of SCCs is $\mathcal{L}^* = \bigcup_{N \in \mathcal{F}(\mathbb{N})} \mathcal{L}^N$ instead of \mathcal{L}^N for some fixed electorate N , i.e., SCCs are defined for all possible electorates. By contrast, the set of alternatives A is still fixed.

In addition to weak \succsim^F -strategyproofness, we need three further axioms for our characterization of the top cycle:

- An SCC f is *pairwise* if $f(R) = f(R')$ for all preference profiles $R, R' \in \mathcal{L}^*$ such that $n_{xy}(R) - n_{yx}(R) = n_{xy}(R') - n_{yx}(R')$ for all $x, y \in A$.
- An SCC f is *set non-imposing* if for every non-empty set of alternatives $X \subseteq A$, there is a profile $R \in \mathcal{L}^*$ such that $f(R) = X$.
- An SCC f is *homogeneous* if $f(R) = f(\ell R)$ for all profiles $R \in \mathcal{L}^*$ and integers $\ell \in \mathbb{N}$. Here, ℓR denotes the profile that consists of ℓ copies of R ; the identities of the voters will not matter as pairwise implies anonymity.

The first axiom, pairwise, restricts the information that an SCC can use to determine the winner by defining an equivalence relation on the profiles. While this axiom rules out some SCCs (e.g., the omninomination rule and the Pareto rule), numerous important SCCs are pairwise (e.g., the top cycle, the Borda rule, and the Copeland rule). We furthermore note that pairwise relates profiles defined for different numbers of voters. Next, set non-imposition states that all possible outcomes can be chosen by the considered SCC, i.e., its range coincides with its codomain. Due to the variable electorate assumption, this condition is mild as we can construct profiles where a given set of alternatives is perfectly symmetric. Finally, homogeneity is a weak consistency axiom for variable electorates that states that multiplying the whole profile does not affect the outcome. The idea of this axiom is that multiplying all voters' preferences does not give any new information, so the outcome should not change.

We are now ready to state the main result of this section.

THEOREM 3.12 (*Brandt and Lederer, 2023*)

The top cycle is the only pairwise SCC on \mathcal{L}^* that satisfies weak \succsim^F -strategy-proofness, set non-imposition, and homogeneity.

We first note that this result essentially turns the impossibility theorem of Gibbard (1973) and Satterthwaite (1975) into a complete characterization of the top cycle by moving from SCFs to pairwise SCCs. In particular, we only need to replace strategyproofness with weak \succsim^F -strategyproofness and non-imposition with set non-imposition (and additionally require homogeneity) to arrive at the characterization of the top cycle. Thus, we view Theorem 3.12 as an attractive escape route to the negative consequences of Theorem 3.1. On the other hand, we have to acknowledge that Theorem 3.12 itself can be interpreted as an impossibility theorem as the top cycle fails Pareto-optimality. In more detail, our result implies that no Pareto-optimal and pairwise SCC satisfies weak \succsim^F -strategyproofness, set non-imposition, and homogeneity, which constitutes a strong negative result.

Finally, it should be mentioned that we also present a more general variant of Theorem 3.12 in Publication 5. To explain this result, we introduce the notion of dominant sets: a set X is dominant in the majority relation if $x \succ_M y$ for all $x \in X$, $y \in A \setminus X$. Dominant sets are ordered by set inclusion and it can be shown that the top cycle is the SCC that always returns the smallest dominant set. We then generalize Theorem 3.12 by replacing set non-imposition with neutrality and non-imposition: the given axioms then characterize a class of SCCs whose defining feature is to always return a dominant set in the majority relation.

3.2.2 The Case of Weak Preferences

So far, we have focused on the case that voters report strict preferences in our analysis. However, there are numerous situations where weak preferences arise naturally. For instance, if there are many similar alternatives, it seems reasonable that voters may declare to be indifferent between all of them. In this section, we thus investigate strategyproofness for SCCs for the case that voters report weak preference relations. All the results in this section are taken from Publication 6.

Before discussing our results, we note that many possibility theorems crucially rely on the assumption that voters' preferences are strict: for instance, Theo-

rem 3.12 ceases to hold when allowing for indifferences in the voters' preference relations because all Condorcet-consistent SCCs on \mathcal{R}^N fail even weak \succsim^K -strategyproofness (Brandt, 2015). Moreover, Brandt et al. (2022c) have shown that no anonymous and Pareto-optimal SCC is weakly \succsim^F -strategyproof when the voters' preference relations can contain ties. Similar negative results are also known for SDSs: for instance, Brandl et al. (2018) show that no anonymous and neutral SDS satisfies weak \succsim^{SD} -strategyproofness and \succsim^{SD} -efficiency (which requires that it is not possible to give one voter an \succsim^{SD} -preferred outcome without making another voter worse off). Furthermore, Aziz et al. (2018) and Brandl et al. (2021) show similar results for various other lottery extensions. In summary, these results demonstrate that already very mild forms of strategyproofness lead to far-reaching impossibility theorems when voters can report weak preference relations.

On the other side, it is known that some SCCs are weakly \succsim^K -strategyproof even when preferences are weak. For instance, the Pareto rule and the omninomination rule satisfy weak \succsim^K -strategyproofness even for weak preferences. However, these SCCs tend to choose rather large choice sets and thus violate common decisiveness desiderata. In this section, we therefore analyze weak \succsim^K -strategyproofness for SCCs on the domain of weak preferences, and we will show that all SCCs that satisfy this axiom are undesirable as they tend to choose large choice sets.

In more detail, we will first focus on two important subclasses of SCCs: rank-based and support-based ones. Rank-basedness has already been introduced for SDSs in the case of strict preferences and we will now generalize this definition to weak preferences. To this end, we redefine the rank vector $\bar{r}^*(R, x)$ based on the rank tuple $\bar{r}(\succsim_i, x) = (1 + |\{y \in A \setminus \{x\} : y \succ_i x\}|, |\{y \in A \setminus \{x\} : y \sim_i x\}|)$: $\bar{r}^*(R, x)$ is the vector that contains the rank tuples of all voters in lexicographically increasing order. That is, $\bar{r}^*(R, x) = (\bar{r}(\succsim_{i_1}, x), \dots, \bar{r}(\succsim_{i_n}, x))$, where the voters are ordered such that $\bar{r}(\succsim_{i_j}, x)_1 \leq \bar{r}(\succsim_{i_{j+1}}, x)_1$ and $\bar{r}(\succsim_{i_j}, x)_2 \leq \bar{r}(\succsim_{i_{j+1}}, x)_2$ if $\bar{r}(\succsim_{i_j}, x)_1 = \bar{r}(\succsim_{i_{j+1}}, x)_1$ for all $j \in \{1, \dots, n-1\}$. Then, an SCC f is *rank-based* if $f(R) = f(R')$ for all profiles R, R' such that $\bar{r}^*(R, x) = \bar{r}^*(R', x)$ for all alternatives $x \in A$. For instance, the Borda rule, the omninomination rule, and the plurality rule are in this class of SCCs. In Publication 6, we prove that all weakly \succsim^K -strategyproof and rank-based SCCs fail Pareto-optimality.

THEOREM 3.13 (Brandt et al., 2022a)

No rank-based and weakly \succsim^K -strategyproof SCC on \mathcal{R}^N satisfies Pareto-optimality if $m \geq 4$ and $n \geq 3$.

The second class of SCCs that we consider are support-based ones. To introduce this class, we recall that the support between two alternatives $x, y \in A$ is $n_{xy}(R) = |\{i \in N : x \succ_i y\}|$. Then, an SCC f is *support-based* if $f(R) = f(R')$ for all preference profiles R, R' such that $n_{xy}(R) = n_{xy}(R')$ for all $x, y \in A$. Put differently, support-based SCCs can only rely on the supports between alternatives to compute the winning alternatives. We note that support-basedness is implied by pairwiseness (see Section 3.2), and numerous important SCCs, such as the top cycle, the Pareto rule, and the Borda rule, are thus support-based. As the next result shows, all support-based and weakly \succsim^K -strategyproof SCCs must choose large choice sets. In more detail, we will show that these rules either fail Pareto-optimality or always choose at least one of the most preferred alternatives of every voter $i \in N$.

THEOREM 3.14 (Brandt et al., 2022a)

Every support-based and weakly \succsim^K -strategyproof SCC on \mathcal{R}^N fails Pareto-optimality or always chooses at least one of the most preferred alternatives of every voter $i \in N$ if $m \geq 3$.

We note that this theorem implies an impossibility result of Brandt et al. (2022c) for pairwise and weakly \succsim^K -strategyproof SCCs. By contrast, Theorem 3.14 is no impossibility as the Pareto rule satisfies all axioms, but it also chooses always at least one of the most preferred alternatives of every voter. Finally, we again emphasize that almost all commonly studied SCCs are either rank-based or support-based. Consequently, Theorems 3.13 and 3.14 show that most SCCs in the literature either fail weak \succsim^K -strategyproofness or sometimes select unreasonably large choice sets.

As the last contribution of Publication 6, we prove a third impossibility theorem which demonstrates a far-reaching variant of the Gibbard-Satterthwaite theorem. To formalize this result, we introduce the Condorcet loser property. A *Condorcet loser* x in a profile R is an alternative such that $y \succ_M x$ for all $y \in A \setminus \{x\}$. That is, a Condorcet loser loses all pairwise majority comparisons and it is thus undesirable to choose this alternative. The *Condorcet loser property* requires that $x \notin f(R)$ whenever alternative x is the Condorcet loser in the profile R . As we show next, non-imposition and the Condorcet loser property are jointly incompatible for weakly \succsim^K -strategyproof SCCs. Hence, every SCC that satisfies non-imposition and weak \succsim^K -strategyproofness selects the Condorcet loser in some profile.

THEOREM 3.15 (Brandt et al., 2022a)

No non-imposing and weakly \succsim^K -strategyproof SCC on \mathcal{R}^N satisfies the Condorcet loser property if $m \geq 3$ and $n \geq 4$.

Finally, we note that all our results in this section immediately carry over to SDSs when using *ex post* efficiency instead of Pareto-optimality and weak \succsim^{DD} -strategyproofness instead of weak \succsim^K -strategyproofness. The reason for this is that these axioms for SDSs only rely on the supports of the chosen lotteries (i.e., the sets of alternatives with positive probability), which allows for a joint treatment of SCCs and SDSs. Thus, the results in Publication 6 also rule out that there are attractive strategyproof SDSs when voters have weak preferences. In particular, our results can be understood as a strengthening of a seminal impossibility result by Brandl et al. (2018) for restricted classes of SDSs.

3.3 RESULTS FOR COMMITTEE VOTING RULES

In this section, we move past the classical setting of single-winner elections and instead consider committee elections, where a fixed number of winners is chosen. In more detail, we will study strategyproofness for party-approval committee (PAC) voting rules, which assign the seats of a committee to parties based on the voters' approval ballots over these parties. We note that this setting has only recently been introduced by Brill et al. (2022), and our results are thus the first on strategyproof PAC voting rules. In particular, we are interested in the question of whether there are strategyproof PAC voting rules that guarantee proportional representation of

the voters' preferences (we refer to Section 2.4 for formal definitions). The results in this section are based on Publication 7.

The research question in this section draws significant motivation from related settings. In particular, for approval-based committee (ABC) elections (where voters approve individual candidates rather than parties and a committee is a subset of the candidates instead of a multiset of the parties), Peters (2018) has shown that very mild notions of strategyproofness and proportional representation are incompatible with each other.⁹ Due to the central role of proportional representation in ABC voting, this result can be seen as a counterpart of the Gibbard-Satterthwaite theorem for these elections. Hence, researchers have explored various escape routes, such as restricting the domain of feasible preference profiles (Botan, 2021) or studying strategyproofness for set-valued committee voting rule (Kluiving et al., 2020). Our work fits well in this stream as one can see party-approval elections also as a domain restriction for ABC elections by replacing each party with k alternatives representing its members and requiring that voters either approve all members of a party or none. In this section, we will thus answer whether it is possible to circumvent the negative result of Peters (2018) by moving from ABC elections to PAC elections.

Moreover, there are even some results that indicate that there could be appealing strategyproof PAC voting rules: this setting is related to two models called apportionment and fair mixing, both of which allow for strategyproof voting rules that satisfy proportional representation (Aziz et al., 2020; Pukelsheim, 2014). In more detail, apportionment can be seen as the special case of PAC voting where voters only approve a single party (Balinski and Young, 2001; Pukelsheim, 2014), and it is easy to see that many voting rules guarantee proportional representation and strategyproofness in this setting. Furthermore, in fair mixing, where the output is a probability distribution over the alternatives rather than a multiset (Bogomolnaia et al., 2005; Aziz et al., 2020; Brandl et al., 2021), a rule called the conditional utilitarian rule satisfies strategyproofness and proportional representation.

Unfortunately, it turns out that strategyproofness and proportional representation are incompatible for PAC elections. In more detail, we show in Publication 7 that no anonymous PAC voting rule satisfies both of these conditions.

THEOREM 3.16 (Delemazure et al., 2023)

No anonymous PAC voting rule satisfies weak representation and strategyproofness if there are $k \geq 3$ seats in the committee, $m \geq k + 1$ parties, and $n = 2\ell k$ voters for some $\ell \in \mathbb{N}$.

We note that Theorem 3.16 only holds for specific configurations of the parameters. In particular, we require that the number of voters is a multiple of $2k$, which is mainly due to the fact that we found no induction step that generalizes the impossibility to all n . While we believe that the impossibility also holds for different values of n (though not necessarily all), it remains open to close this gap.

The second assumption that may seem controversial is that $m \geq k + 1$, i.e., there are more parties than seats in the committee. While this assumption is not valid for large parliamentary elections, it is easy to see that it is necessary

⁹ Before this result, it was already known that no commonly considered ABC voting rule satisfies both strategyproofness and proportional representation (Aziz et al., 2015b; Lackner and Skowron, 2018). These results hinted at the general impossibility by Peters (2018).

for our impossibility theorem: otherwise, every rule that always returns a fixed committee W with $W(x) \geq 1$ for all parties x satisfies both weak representation and strategyproofness. However, rather than showing that there are attractive PAC voting rules that satisfy strategyproofness when $m \leq k$, this example hints at the fact that weak representation is unreasonably weak in such situations. Indeed, when replacing this axiom with weak proportional representation, we can extend our impossibility theorem to the case that $m \leq k$.

THEOREM 3.17 (*Delemazure et al., 2023*)

No anonymous PAC voting rule satisfies weak proportional representation and strategyproofness if there are $k \geq 3$ seats in the committee, $m \geq 4$ parties, and $n = 2\ell k$ voters for some $\ell \in \mathbb{N}$.

In light of these negative results, it is a natural follow-up question how to derive more positive theorems. To this end, we consider in Publication 7 a weaker form of strategyproofness. In more detail, we study the notion of strategyproofness for unrepresented voters, which requires that $f(R, \succsim_i) \geq f(R', \succsim_i)$ for all approval profiles $R, R' \in \mathcal{A}^N$ and voters $i \in N$ such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$ and $f(R, \succsim_i) = 0$. Put differently, this axiom only prohibits voters who approve of none of the elected committee members from manipulating. We find this notion reasonable as voters without representation have the most incentive to change the outcome. Unfortunately, it turns out that almost all commonly studied PAC voting rules even fail strategyproofness for unrepresented voters in some approval profile. Indeed, we only found two exceptions to this: multiwinner approval voting (which satisfies even full strategyproofness but fails every notion of proportional representation) and Chamberlin-Courant approval voting. Since Chamberlin-Courant approval voting satisfies weak representation, we thus reconcile at least minimal notions of strategyproofness and proportional representation.

THEOREM 3.18 (*Delemazure et al., 2023*)

Chamberlin-Courant approval voting is the only Thiele rule that satisfies weak representation and strategyproofness for unrepresented voters for all committee sizes k , numbers of parties m , and numbers of voters n .

We note, however, that this result is only attractive when $k \leq m$ because Chamberlin-Courant approval voting does not distinguish between parties once every voter approves one committee member. In more detail, once every party has a seat in the committee, the remaining seats are allocated by the tie-breaking mechanism as the Chamberlin-Courant score of all such committees is maximal. Thus, this theorem is only a first step to more positive results on strategyproofness in committee elections and there is still room to explore numerous other approaches.

In this section, we will explain the techniques and central observations that have been used to derive our results. To this end, we first note that all our proofs use the axioms to directly reason about preference profiles. On the one hand, this means that our proofs do not offer completely new techniques. On the other hand, our proofs can be verified based on rather basic knowledge and are independent of other results. Moreover, although we do not use complicated techniques, our proofs are often quite involved as they reason about numerous preference profiles.

To derive these elaborate proofs, we have often relied on some common observations, which we will explain subsequently in more detail. In particular, the central insight for many of our results is that there is a tradeoff between decisiveness (in the sense of avoiding randomization or large sets of winners whenever possible) and strategyproofness. In Section 4.1, we will explain this aspect in detail with the goal of highlighting why decisiveness and strategyproofness conflict and how to use this insight to prove new theorems. As our second main technique, we rely on SAT solving, a computer-aided approach for theorem proving, and we will explain this method in Section 4.2.

4.1 DECISIVENESS VERSUS STRATEGYPROOFNESS

A crucial property of SCCs and SDSs is their decisiveness: whenever there is a clear winner in a preference profile, this alternative should be uniquely selected by a voting rule. Indeed, if such an alternative exists but is not uniquely chosen, another less desirable alternative has a chance to win. Moreover, this means that we ignore information contained in the preference profile and instead leave the decision on the final winner to a tie-breaking mechanism that is independent of the voters' preferences.

While these reasons show that decisiveness is important for SCCs and SDSs, many strategyproof voting rules are rather indecisive. This is best illustrated by the uniform random dictatorship: while this rule is arguably the most appealing strongly \succsim^{SD} -strategyproof SDS, it only chooses a winner deterministically if all voters agree on the best choice and randomizes in all other cases. Hence, the uniform random dictatorship is rather indecisive and the results by Gibbard (1977) and Barberà (1979b) entail that essentially all strongly \succsim^{SD} -strategyproof SDSs need to use a lot of randomization for selecting the winner. Numerous other impossibility results in the literature also follow this theme: for instance, Barberà (1977b) shows that all SCCs that satisfy positive responsiveness, an axiom that implies that a single winner is chosen for many profiles, fail weak \succsim^K -strategyproofness. Weaker decisiveness criteria such as Condorcet consistency, residual resoluteness, or 1-unanimity, are, e.g., used for impossibility theorems by Gärdenfors (1976), Duggan and Schwartz (2000), and Benoît (2002). Finally, even

results that do not directly use decisiveness criteria are often proven by implicitly showing that the considered axioms entail a strong degree of decisiveness (e.g., Brandl et al., 2018; Aziz et al., 2018; Brandt et al., 2022c).

The tradeoff between decisiveness and strategyproofness is also visible in many of our results: for instance, Condorcet-consistency can naturally be interpreted as a decisiveness notion as it postulates that a single alternative is chosen whenever there is a Condorcet winner, and we show in numerous results (cf. Theorems 3.4, 3.6, 3.8 and 3.10) that Condorcet-consistency conflicts with various notions of strategyproofness for SDSs. Even more explicit examples are Theorem 3.9, where we investigate the tradeoff between strategyproofness (in the form of \succsim^u -strategyproofness) and decisiveness (in the form of k -unanimity) for rank-based SDSs, or the results in Section 3.2.2 which indicate that only rather indecisive SCCs are weakly \succsim^k -strategyproof when the voters' preferences are weak. Finally, even results such as Theorems 3.11 and 3.12, which may seem unrelated to decisiveness at a first glance, reason about decisiveness notions in their proofs.

On an intuitive level, there is a simple reason why so many results in the literature describe the conflict between strategyproofness and decisiveness: the more decisive an SCC or SDS is, the closer it becomes to an SCF. In particular, in the extreme case that an SCC or SDS always chooses a single winner, it turns into an SCF and the Gibbard-Satterthwaite theorem applies again. By contrast, the maximally indecisive voting rules—the SCC that always chooses all alternatives and the SDS that always picks every alternative with probability $1/m$ —are clearly strategyproof as voters cannot affect the outcome. Thus, the research on strategyproof SDSs and SCCs tries to figure out for which preference profiles a single winner can be chosen and for which profiles strategyproofness necessitates returning a set or a lottery over the alternatives. As a consequence of this insight, many impossibility results can be interpreted as showing that strategyproofness (possibly combined with further axioms) requires very indecisive voting rules.

In the next two subsections, we aim to explain how the tradeoff between strategyproofness and decisiveness can be used to prove new results. The reason for this is that many of our theorems do not only demonstrate the tradeoff between these two desiderata, but their proof is directly driven by this conflict.

4.1.1 Decisiveness as the Cause of Manipulability

The most direct way to show that decisiveness conflicts with strategyproofness is to prove that no strategyproof voting rule satisfies a given decisiveness axiom. In particular, this approach typically leads to rather elegant proofs as decisiveness notions determine the outcomes for numerous profiles and we only need to relate these outcomes by strategyproofness to infer an impossibility. This technique has, for example, been used to show Theorems 3.4 and 3.8 to 3.10. Furthermore, the impossibility theorems by, e.g., Benoît (2002) and Brandt (2015) also rely on this approach.

As a simple example of this technique, we discuss next a proposition that summarizes the main idea of the proof of Theorem 3.8. Recall for the subsequent result that $u(k)$ denotes the utility assigned to a voter's k -th best alternative.

PROPOSITION 4.1 (*Lederer, 2021*)

No Condorcet-consistent SDS on \mathcal{L}^N satisfies strong \succsim^{u^Π} -strategyproofness for a vNM utility function $u \in \mathcal{U}$ with $u(1) - u(2) < u(2) - u(3)$ if $m = 3$ and $n = 3$.

Proof. Assume for contradiction that there is an SDS f that satisfies all requirements of the proposition and consider the following four preference profiles.

R^1 :	1: a, b, c	2: c, a, b	3: b, c, a
R^2 :	1: b, a, c	2: c, a, b	3: b, c, a
R^3 :	1: a, b, c	2: a, c, b	3: b, c, a
R^4 :	1: a, b, c	2: c, a, b	3: c, b, a

First, it is easy to verify that b is the Condorcet-winner in R^2 , a in R^3 , and c in R^4 . Hence, Condorcet-consistency entails that $f(R^2, b) = f(R^3, a) = f(R^4, c) = 1$. Now, by strong \succsim^{u^Π} -strategyproofness from R^1 to R^2 , R^3 , and R^4 , we infer the following inequalities.

$$\begin{aligned} f(R^1, a)u(1) + f(R^1, b)u(2) + f(R^1, c)u(3) &\geq u(2) \\ f(R^1, b)u(1) + f(R^1, c)u(2) + f(R^1, a)u(3) &\geq u(2) \\ f(R^1, c)u(1) + f(R^1, a)u(2) + f(R^1, b)u(3) &\geq u(2) \end{aligned}$$

By summing up these inequalities, we derive that $u(1) + u(2) + u(3) \geq 3u(2)$ which conflicts with our assumption that $u(1) - u(2) < u(2) - u(3)$. This contradicts the strong \succsim^{u^Π} -strategyproofness of f and hence proves the proposition. \square

In the proof of Proposition 4.1, we do not directly relate the profiles for which Condorcet-consistency determines the outcomes. This is typical for this type of proof as decisiveness notions only determine the outcome if there is a clear winner and it is usually impossible to directly manipulate in such profiles. The main challenge is thus to find suitable intermediate profiles to connect the outcomes given by the decisiveness notion. In the proof of Proposition 4.1, we use for this a rather classical idea: the profile R^1 appears in numerous impossibility theorems and is, e.g., also the cause for Arrow's impossibility. While many of our proofs rely on variants of this profile, we want to emphasize that it is sometimes rather difficult to find the intermediate profiles because every strategyproofness application typically loses information.

4.1.2 Inferring Decisiveness as a Tool

Perhaps a more surprising insight than the one in Section 4.1.1 is that decisiveness notions can also be used as auxiliary tools to prove new theorems. In particular, in many of our results, we first infer a decisiveness notion from the given axioms and then complete the proof based on this decisiveness axiom. For instance, the proofs of Theorems 3.6, 3.11 and 3.13 to 3.15 follow this pattern. Moreover, the results of Brandl et al. (2018) and Brandt et al. (2022c) also adhere to this scheme by implicitly inferring that the considered voting rules are, e.g., 1-unanimous (which is a strong condition if there are only few voters). This is possible as strategyproofness combined with other axioms often implies decisiveness.

As a concrete example of this approach, we consider next a part of the proof of Theorem 3.11, where we show that the axioms postulated in Theorem 3.11 imply

k-unanimity for $k = \frac{n}{2} - 1$ if $m = 3$ and $n \geq 4$ is even. Recall here that k-unanimity requires that $f(R, x) = 1$ if at least $n - k$ voters top-rank x in R .

PROPOSITION 4.2 (Brandt et al., 2023b)

Every anonymous and neutral SDS that satisfies \succsim^{PC} -efficiency and \succsim^{PC} -strategy-proofness is k-unanimous for $k = \frac{n}{2} - 1$ if $m = 3$ and $n \geq 4$ is even.

Proof. Let f denote an SDS that satisfies anonymity, neutrality, \succsim^{PC} -efficiency, and \succsim^{PC} -strategyproofness, and consider the following profile R :

$$R^1: \quad 1: a, b, c \quad 2: a, c, b \quad [3 \dots \frac{n}{2} + 1]: b, a, c \quad [\frac{n}{2} + 2 \dots n]: c, a, b$$

First, it is easy to see that b and c are symmetric in R^1 , so anonymity and neutrality imply that $f(R^1, b) = f(R^1, c)$. In turn, \succsim^{PC} -efficiency shows that $f(R^1, a) = 1$ as every lottery with $f(R^1, b) = f(R^1, c) > 0$ is \succsim^{PC} -inefficient. Next, we change the preferences of all voters $i \in \{1, \dots, \frac{n}{2} + 1\}$ to a, c, b to derive the profile R^2 .

$$R^2: \quad [1 \dots \frac{n}{2} + 1]: a, c, b \quad [\frac{n}{2} + 2 \dots n]: c, a, b$$

By repeatedly applying \succsim^{PC} -strategyproofness, it follows that $f(R^2, a) = 1$ as all manipulators prefer a the most after they deviate. Moreover, we note that c Pareto-dominates b in R^2 , so we can now let the voters $i \in \{\frac{n}{2} + 2, \dots, n\}$ change their preference relations to c, b, a . By \succsim^{PC} -efficiency, we have that $f(R, b) = 0$ for all intermediate profiles and \succsim^{PC} -strategyproofness then requires that $f(R, a) = 1$. So, we infer that $f(R^3, a) = 1$ for the subsequent profile.

$$R^3: \quad [1 \dots \frac{n}{2} + 1]: a, c, b \quad [\frac{n}{2} + 2 \dots n]: c, b, a$$

From here on, it is easy to show that f satisfies k-unanimity for $k = \frac{n}{2} - 1$. Firstly, when a voter $i \in \{1, \dots, \frac{n}{2} + 1\}$ changes her preferences but a remains the favorite alternative, \succsim^{PC} -strategyproofness implies that a still gets probability 1. Secondly, when a voter $i \in \{\frac{n}{2} + 2, \dots, n\}$ changes her preference relation, a still must be chosen with probability 1 as every other outcome constitutes a manipulation. Hence, $f(R, a) = 1$ for all profiles in which the voters $i \in \{1, \dots, \frac{n}{2} + 1\}$ top-rank a . Finally, we can rename the voters and alternatives by anonymity and neutrality, so f is k-unanimous for $k = \frac{n}{2} - 1$. \square

To complete the proof of Theorem 3.11, we generalize Proposition 4.2 in Publication 4 first to all $n \geq 7$ and $m \geq 3$ and then prove that no \succsim^{PC} -strategyproof and \succsim^{PC} -efficient SDS satisfies k-unanimity for $k = \lfloor \frac{n-1}{2} \rfloor$. So, we reduce the impossibility of Theorem 3.10 to an impossibility result that analyzes the compatibility of strategyproofness and decisiveness.

While there is no universal way to infer decisiveness notions, Proposition 4.2 shows some key steps that are frequently used. In particular, the first and maybe most important step is to find a profile where an alternative x is chosen with probability 1 even though it is not unanimously top-ranked. In Proposition 4.2, this is the profile R^1 . From this profile, we then go to a profile where a group of voters top-rank alternative x and all other voters top-rank another alternative. Strategyproofness usually implies for this step that x still gets probability 1. Finally, we bottom-rank x in the preference relations of the voters that do not top-rank this alternative. Just as in Proposition 4.2, it is often possible to use *ex post* efficiency or similar axioms for this step. Based on this profile, it is then easy to infer that the considered voting rule satisfies, e.g., k-unanimity. These steps are, for instance, also visible in the proofs of Propositions 1 and 2 in Publication 6.

Finally, we note that we directly infer a strong decisiveness notion in Proposition 4.2. An alternative approach is to first show that the considered voting rule only satisfies a mild degree of decisiveness (such as 1-unanimity) and then inductively infer stronger notions of decisiveness. For instance, for Theorems 3.14 and 3.15, we first establish that the considered rules are 1-unanimous and then show inductively that every k -unanimous SCC that satisfies the requirements of these theorems is also $k + 1$ -unanimous. Based on this approach, it is straightforward to infer an impossibility as no voting rule satisfies k -unanimity for $k \geq \frac{n}{2}$, but the induction step typically implies the existence of such a rule.

4.2 SAT SOLVING

Our second main method for proving results is SAT solving, a computer-aided theorem proving technique. The idea of this method is to encode the existence of a voting rule that satisfies some desired axioms as a logical formula and then let a computer program decide whether the formula is satisfiable or not. In particular, if the computer returns unsatisfiable, this means that no voting rule satisfies all of the specified axioms, thus proving an impossibility theorem. Among the results in this thesis, Theorems 3.16 and 3.17 have been fully proven by SAT solving, and the proofs of numerous other results have been supported by such computer-aided techniques. In this section, we will thus explain SAT solving as an example of computer-aided theorem proving.

Before going into the details of SAT solving, we note that computer-aided theorem proving techniques have been used to show a wide range of theorems in social choice theory (we refer to Geist and Peters (2017) for a survey on this topic). In particular, after the pioneering work by Tang and Lin (2009), these methods soon become popular in computation social choice (e.g., Geist and Endriss, 2011; Grandi and Endriss, 2013; Brandt and Geist, 2016; Brandt et al., 2017; Brandl et al., 2018; Endriss, 2020; Kluiving et al., 2020; Brandt et al., 2022c; Brandl et al., 2021; Brandt et al., 2023a). Notably, most of these works rely on SAT solving, which emphasizes the importance of this method.¹⁰ Moreover, it should be noted that most of these results use the computer to prove impossibility theorems. By contrast, there is also significant amount of work on automated mechanism design, where computers are used to find good mechanisms or outcomes (Conitzer and Sandholm, 2002, 2003; Narasimhan et al., 2016; Mittelman et al., 2022), or explainability in social choice, where the goal is to find explanation for why some alternative should be chosen (Cailloux and Endriss, 2016; Boixel and Endriss, 2020; Peters et al., 2020; Schmidtlein and Endriss, 2023). We will, however, ignore these works as they have a significantly different goal.

To explain SAT solving in detail, we will revisit Theorem 3.16 and explain in several steps how this result is proven. We note that our code for this result is publicly available at <https://zenodo.org/record/7356204>.

¹⁰ Noteworthy exceptions are due to Brandl et al. (2018), who use SMT solving to analyze randomized voting rules, and Brandt et al. (2023a), who use search algorithms to tackle enormous problem instances.

4.2.1 Encoding the Formula

As the first step for using SAT solving, we have to encode the considered problem as a logical formula. For Theorem 3.16, this means that we need to encode the existence of an anonymous PAC voting rule that satisfies strategyproofness and weak representation in propositional logic. To this end, we first need to fix the parameters of our voting rule because the corresponding formula will become infinitely large if we, e.g., allow for an infinite number of voters. We thus assume that there are exactly k seats in the committee, m parties, and n voters for some fixed numbers k , m , and n . Moreover, most SAT solvers require the formula to be in a special format called conjunctive normal form, so we will write all our conditions in this form.

Next, we need to define the meaning of the variables. While there are numerous ways to encode voting rules, the most common one is to add a variable $x_{R,O}$ for each profile R and possible outcome O , which should be true if and only if the encoded voting rule chooses the outcome O for the profile R . For Theorem 3.16, this means that we add a variable $x_{R,W}$ for every approval profile $R \in \mathcal{A}^N$ and committee $W \in \mathcal{W}_k$. Furthermore, to enforce that these variables indeed encode a voting rule, we need to add constraints. In particular, we have to ensure that precisely one outcome is chosen for each profile, which corresponds to the following two constraints in the context of PAC voting rules. The first constraint ensures that at least one committee is chosen for every approval profile R , and the second one that not more than one committee is chosen. It is straightforward to adapt these constraints to different types of voting rules.

$$\begin{aligned} \forall R \in \mathcal{A}^N : \bigvee_{W \in \mathcal{W}_k} x_{R,W} \\ \forall R \in \mathcal{A}^N \text{ and } W, W' \in \mathcal{W}_k \text{ with } W \neq W' : \neg x_{R,W} \vee \neg x_{R,W'} \end{aligned}$$

Next, we turn to our axioms, which will be encoded as constraints. Since the encoding depends on the given axiom, we subsequently focus on the conditions in Theorem 3.16 and first explain how to enforce weak representation. To this end, we recall that weak representation requires that, if there is a group of voters of size at least n/k that uniquely approves a party x , then x needs to have at least one seat in the committee. Or, put differently, if there is such a group of voters for a party x , then no committee W with $W(x) = 0$ can be chosen. We model this by requiring that $x_{R,W}$ is false if W violates weak representation for R . More formally, let $WRep(R)$ denote the set of committees that satisfy weak representation in R . Then, weak representation corresponds to the following constraints.

$$\forall R \in \mathcal{A}^N \text{ and } W \in \mathcal{W}_k \setminus WRep(R) : \neg x_{R,W}$$

We next encode strategyproofness. This axiom postulates of a PAC voting rule f that $f(R, \succsim_i) \geq f(R', \succsim_i)$ for all profiles $R, R' \in \mathcal{A}^N$ and voter $i \in N$ with $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$ (recall that $f(R, \succsim_i)$ denotes the number of seats assigned to the approved parties of voter i). Or, in other words, if we choose a committee W for R , then we are not allowed to choose a committee W' for R' with $W(\succsim_i) < W'(\succsim_i)$. This leads to the following constraints, where $R_{-i} = R'_{-i}$ indicates that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$. The formulation of strategyproofness is rather universal.

$$\forall R, R' \in \mathcal{A}^N, i \in N, \text{ and } W, W' \in \mathcal{W}_k \text{ with } R_{-i} = R'_{-i} \\ \text{and } W(\succsim_i) < W'(\succsim_i) : x_{R,W} \implies \neg x_{R',W'}$$

Finally, we consider anonymity and note that it is possible to encode this axiom with further constraints: if R is derived from R' by permuting the voters, then $x_{R,W}$ is true if and only if $x_{R',W}$ is true. However, a more clever way to encode anonymity is to change the domain of PAC voting rules from approval profiles to anonymous approval profiles. Hence, instead of viewing profiles as tuples that state the ballot of every voter, we view profiles as multisets that state how often each approval ballot is reported. Because all profiles that differ only by permuting voters correspond to the same anonymous profile, we can encode this axiom without further constraints by changing the domain of our voting rules to anonymous approval profiles. Moreover, by doing so, we significantly reduce the number of profiles and the formula's size.

Based on this encoding, we can, at least in theory, construct the formula, encode it in a computer-readable format, and check with a computer program, a so-called SAT solver (e.g., Biere, 2008; Audemard and Simon, 2019), whether it is satisfiable. However, this approach will only work in practice if the constructed formula is not too big, which depends crucially on the numbers of voters and alternatives. For instance, for Theorem 3.16, we need to use the parameters $k = 3$, $m = 4$, and $n = 6$ to infer an impossibility and the corresponding formula is then so large that even state-of-the-art SAT solvers cannot decide whether the formula is satisfiable in a reasonable amount of time. Consequently, many authors further optimize the encoding by, e.g., excluding some preference profiles or removing some constraints. In particular, this step requires human knowledge or educated guesses about the constraints and preference profiles that are important for the impossibility. For example, for Theorem 3.16, we exclude numerous profiles from the analysis to speed up the SAT solving. After applying this additional trick, the SAT solver is finally able to solve the formula and we derive the following impossibility as the SAT solver returns unsatisfiable for the constructed formula.

PROPOSITION 4.3 (*Delemazure et al., 2023*)

No anonymous PAC voting rule satisfies strategyproofness and weak representation if there are $k = 3$ seats in the committee, $m = 4$ parties, and $n = 6$ voters.

4.2.2 Verifying the Impossibility

After deriving a result such as Proposition 4.3 based on a computer-aided approach, the next question is how to verify the result. For instance, if there is a mistake in the program that writes the logical formula or in the program that checks whether the formula is satisfiable, the theorem proven by SAT solving could be wrong. Of course, one can eliminate such mistakes by verifying that the underlying programs are correct, but this is tedious and not insightful.

Thus, the standard approach to verify the correctness of results shown by computer-aided methods is to extract a human-readable proof. In more detail, one can typically extract a minimal unsatisfiable subset (MUS) from the original formula. Given an unsatisfiable propositional formula, a MUS is an inclusion-minimal sub-

set of the original constraints that is still unsatisfiable. Thus, a MUS can be considered as the core of why a problem is unsatisfiable and there are again computer programs to efficiently derive MUSes (e.g., Belov and Marques-Silva, 2012; Nadel et al., 2014). In social choice theory, MUSes are often quite small, containing only around 20 profiles, which makes it possible to translate a MUS back into a human-readable proof. Then, one can simply check the human-readable proof to verify the correctness of the theorem proven by SAT solving and is completely independent of the computer proof. However, in more recent results, MUSes became quite large: for instance, the computer proof extracted by Brandl et al. (2018) relies on 47 profiles, and while the authors managed to extract a human-readable proof, it is still very tedious to verify the result based on it. An even more drastic example is an impossibility theorem by Brandl et al. (2021), for which the computer proof uses 386 profiles. Because of this huge number of profiles, just verifying this proof by hand seems error-prone.

As a consequence, it has recently become popular to verify computer proofs by interactive theorem provers, which are computer programs specifically designed for this task. An example of such an interactive theorem prover is Isabelle (Nipkow et al., 2002). The idea of these programs is to take a proof and only verify its correctness. To this end, interactive theorem provers typically support higher-order math (such as universal quantification) and therefore allow to formalize axioms very similar to their original definitions. Moreover, to ensure the prover does not make mistakes, it can only rely on a small and highly trustworthy kernel of basic assertions. Because of all these precautions, the verification of results by interactive theorem provers is considered highly trustworthy and experts in this field even claim that computer-verified proofs are much less likely to be wrong than peer-reviewed proofs (e.g., Hales et al., 2017).

For the verification of Proposition 4.3, we also relied on an interactive theorem prover, namely Isabelle, to verify the correctness. The reason for this is that the smallest MUS that we found contained over 20,000 constraints and 635 profiles. Moreover, an inspection of this MUS revealed that it is not structured in an insightful way. Hence, even if we had extracted a human-readable proof, it would not provide significant benefit. By contrast, we can simply hand these 635 profiles to Isabelle, create the constraints based on verified code, and check whether the resulting logical formula is unsatisfiable. This results in a highly trustworthy proof, thus ensuring that Proposition 4.3 is indeed correct. The Isabelle code for verifying Proposition 4.3 is available in the Archive of Formal Proofs (Delemazure et al., 2022).

4.2.3 Inductive Arguments

Finally, we note that results proven by SAT solving only hold for fixed parameters of k , m , and n , but it is typically desirable to make the results independent of these specific values. To extend the impossibility theorems to more general parameters, it is common to use inductive arguments. Frequently, these inductive arguments are quite straightforward: for instance, if Pareto-optimality is part of the impossibility, one can add more alternatives by bottom-ranking them in the preferences of all voters. Pareto-optimality then ensures that they are not chosen and they hence do not affect the analysis.

To derive Theorem 3.16 from Proposition 4.3, we use three inductive arguments—one for the committee size, one for the number of parties, and one for the number of voters. We exemplarily explain the inductive argument for the number of voters as it is the easiest one. In particular, we show in the following lemma that if there is a PAC voting rule that satisfies all conditions of Theorem 3.16 for a large number of voters, there is also such a PAC voting rule for a small number of voters. Since there is no PAC voting rule that satisfies these axioms for a small number of voters by Proposition 4.3, this means that there is also no such rule for a large number of voters. This type of contrapositive reasoning is typical for inductive arguments in social choice theory.

LEMMA 4.4 (*Delemazure et al., 2023*)

Assume there is an anonymous PAC voting rule for committee size k , m parties, and $n = 2k\ell$ voters (for some $\ell \in \mathbb{N}$) that satisfies strategyproofness and weak representation. There is also such a PAC voting rule for committee size k , m parties, and $n = 2k$ voters.

Proof. Let $\ell \in \mathbb{N}$ and assume that f is a PAC voting rule that satisfies anonymity, strategyproofness, and weak representation for k , m , and $n = 2k\ell$. We define the PAC voting rule g for $n = 2k$ voters as follows: given an approval profile R on $2k$ voters, we clone every voter ℓ time to infer a profile ℓR on $2k\ell$ voters. Then, we set $g(R) = f(\ell R)$. Clearly, g is a well-defined PAC voting rule and always returns a committee of size k as f does so. It hence remains to show that g satisfies anonymity, strategyproofness, and weak representation.

First, we note that g inherits anonymity from f . Indeed, if $R = \pi(R')$ for some profiles R, R' with $2k$ voters and a permutation π , then it also holds that $\ell R = \pi'(\ell R')$ for a permutation π' . Thus, $g(R) = f(\ell R) = f(\ell R') = g(R')$ due to the anonymity of f , which shows that g is also anonymous.

Next, we consider weak representation and assume that R is a profile for $2k$ voters and x an alternative that is uniquely approved by at least $\frac{2k}{k} = 2$ voters in R . As a consequence, there are at least $\frac{2k\ell}{k} = 2\ell$ voters who uniquely approve x in ℓR and the weak representation of f shows that $g(R, x) = f(\ell R, x) \geq 1$.

Finally, for strategyproofness, we assume that g fails this axiom. Thus, there are approval profiles R, R' (for $2k$ voters) and a voter i such that $\tilde{z}_j = \tilde{z}_j'$ for all other voters $j \neq i$ and $g(R', \tilde{z}_i) > g(R, \tilde{z}_i)$. This means that $f(\ell R', \tilde{z}_i) > f(\ell R, \tilde{z}_i)$ by the definition of g . However, we can let the ℓ clones of voter i one after another manipulate starting from ℓR and then infer from the strategyproofness of f that they cannot increase the number of their approved members in the selected committee. Hence, $f(R', \tilde{z}_i) \leq f(R, \tilde{z}_i)$ which contradicts our previous insight. So, g also inherits strategyproofness and satisfies all requirements of the lemma. \square

CONCLUSION AND FURTHER DIRECTIONS

In this thesis, we study the problem of strategic manipulation in social choice theory from a mathematical perspective. Since a seminal theorem by Gibbard (1973) and Satterthwaite (1975) shows that no attractive single-valued voting rule is immune to strategic manipulations by voters, we study in this thesis whether we can circumvent this negative result by studying more flexible models for elections.

In more detail, we first investigate strategyproofness for randomized and set-valued voting rules, which are commonly called social decision schemes (SDSs) and social choice correspondences (SCCs), respectively. SDSs return a lottery over the alternatives instead of a single winner and the final winner will be chosen by chance according to this lottery. Similarly, SCCs return a set of possible winners and some unspecified tie-breaking mechanism will eventually choose a final winner from this set. Both of these approaches are well-known in the literature and there are various ways to define strategyproofness for these models. As a consequence, for both SDSs and SCCs, both impossibility results similar to the Gibbard-Satterthwaite theorem (e.g., Gibbard, 1977; Dutta et al., 2002; Brandl et al., 2018) and possibility results that show the existence of attractive strategyproof voting rules (e.g., Nehring, 2000; Ehlers et al., 2002; Brandt, 2015) have been demonstrated. However, there is typically a large gap between the assumptions that lead to positive results and those that lead to negative ones. In Sections 3.1 and 3.2, we thus try to close this gap by weakening various conditions, such as the considered domain of preference profiles or the underlying strategyproofness notion. In more detail, we provide several positive and negative results, which all have in common that they are tight in the sense that modifying the assumptions turns an impossibility result into a possibility result or vice versa. Our results also indicate a deep tradeoff between the decisiveness and the strategyproofness of voting rules, which we explain in Section 4.1.

We moreover study strategyproofness for committee elections, where the goal is to select $k > 1$ winners instead of one. Such elections have recently attracted significant attention (e.g., Faliszewski et al., 2017; Lackner and Skowron, 2023), but strategyproofness has not been explored in detail for this setting yet. We thus analyze party-approval committee elections, where the seats of a committee are distributed to parties based on the voters' approval ballots. This model has recently been introduced by Brill et al. (2022) and we show a variant of the Gibbard-Satterthwaite theorem for it: no anonymous party-approval committee voting rule satisfies strategyproofness and always chooses a committee that fairly represents the voters' preferences. Since this result is proven based on a computer-aided approach called SAT solving, we also explain this technique in Section 4.2.

While our results settle several challenging and interesting questions regarding strategyproof social choice, we still see several directions that could be explored. We thus list below the, in our opinion, most intriguing open questions concerning strategyproof SDSs, SCCs, and committee voting rules.

OPEN PROBLEM 5.1

Can we quantify the tradeoff between decisiveness and strategyproofness for SDSs?

Our first open problem aims to better understand the tradeoff between decisiveness (in the sense of avoiding randomization whenever possible) and strategyproofness. In particular, our results show that various decisiveness notions, such as Condorcet-consistency and k -unanimity, are incompatible with strategyproofness for SDSs. However, while both Condorcet-consistency and 1-unanimity conflict with strong \succ^{SD} -strategyproofness, it is intuitively clear that a 1-unanimous SDS might be much more strategyproof than a Condorcet-consistent one. Hence, it seems interesting to quantify how severely SDSs fail strategyproofness. For instance, in the context of strong \succ^{SD} -strategyproofness, one could consider the minimal ϵ such that $\sum_{y \in A: y \succ_i x} f(R, y) + \epsilon \geq \sum_{y \in A: y \succ_i x} f(R', y)$ for all profiles R, R' , voters $i \in N$, and alternatives $x \in A$ such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$. Put differently, ϵ measures the maximal violation of strong \succ^{SD} -strategyproofness for a given SDS and can thus be used to define a quantitative variant of strategyproofness. In particular, the value ϵ allows us to quantify the tradeoff between decisiveness and strategyproofness by, e.g., studying the minimal ϵ for Condorcet-consistent SDSs.

OPEN PROBLEM 5.2

Which sets of tie-breaking mechanisms allow for attractive strategyproof SCCs?

Our second problem aims to better understand the assumptions on the tie-breaking mechanisms required for SCCs. In particular, as explained in Section 2.3.2, strategyproofness for SCCs is typically defined by quantifying over a class of tie-breaking mechanisms, which captures a large degree of uncertainty in the outcome. For example, Kelly's extension (and hence also the corresponding strategyproofness notion) is equivalent to saying that a voter prefers a set X to another set Y if she weakly prefers the alternative chosen from X to the alternative chosen from Y , regardless of how these alternatives are selected. Put differently, this set extension quantifies over all tie-breaking mechanisms and thus introduces an enormous degree of uncertainty for the voters because they do not know which tie-breaking mechanism will be chosen. From a practical point of view, this seems undesirable as it is typically impossible to introduce so much uncertainty in the outcome. Therefore, it seems interesting to study strategyproofness notions derived from more restricted classes of tie-breaking mechanisms. For instance, when considering lotteries as tie-breaking mechanisms, one could restrict the set of lotteries that can be used to infer the final winner. In the extreme case that only a single lottery is left for every set, the corresponding SCC turns into an SDS, so this approach also helps to better understand the relation between these two models.

OPEN PROBLEM 5.3

When is strategyproofness possible if voters have weak preferences?

Surprisingly, strategyproofness is not well-understood for weak preferences. In particular, not even for deterministic SCFs, the class of strategyproof voting rules is fully understood: while it is known that these rules must always choose one of the most preferred alternatives of a specific voter, not all SCFs that satisfy this criterion are strategyproof when the voters' preferences contain ties. Note, however,

that we can at least interpret this insight as an impossibility result as no SCF is anonymous, non-imposing, and strategyproof for weak preferences. Another interesting avenue is thus to investigate how far such impossibilities go when considering SDSs and SCCs. For instance, in Publication 6, we prove several far-reaching impossibility results based on weak \succsim^K -strategyproofness for the case that voters report weak preferences and Brandl et al. (2018) show another strong impossibility based on weak \succsim^{SD} -strategyproofness. However, while these results are already quite strong, we conjecture that an even more general impossibility theorem holds: no anonymous (and neutral) SDS satisfies both weak \succsim^{DD} -strategyproofness and \succsim^{SD} -efficiency. It also seems interesting to investigate variants of this claim by, e.g., studying decisiveness axioms for weak preferences. Moreover, domain restrictions could also help to find more positive results when voters have weak preferences.

OPEN PROBLEM 5.4

How can we avoid manipulability in committee elections?

Our last open problem addresses the issue of strategic manipulation in committee elections. In particular, Peters (2018) has shown that no approval-based committee voting rule satisfies both strategyproofness and fairly represents the voters' preferences, and we have extended this result to party-approval committee voting rules in Publication 7. This naturally leads to the question of how we can circumvent these negative results, and many classical escape routes have not been studied yet. For instance, one could investigate domain restrictions for approval profiles (see, e.g., Elkind and Lackner, 2015) and check whether known voting rules satisfy strategyproofness on these domains or design new rules that satisfy both strategyproofness and additional desiderata. Secondly, it seems also interesting to study committee voting rules in the context of iterative voting. In particular, if iterative voting leads to reasonable outcomes, this may mitigate the problems caused by manipulable committee voting rules. Finally, one could also introduce randomization into committee elections to obtain strategyproof committee voting rules.

Part II

ORIGINAL PUBLICATIONS

CORE PUBLICATION [1]: RELAXED NOTIONS OF CONDORCET-CONSISTENCY AND EFFICIENCY FOR STRATEGYPROOF SOCIAL DECISION SCHEMES

SUMMARY

We study social decision schemes (SDSs), which map the preferences of a group of voters over a set of m alternatives to a probability distribution, with respect to strong \succsim^{SD} -strategyproofness (which is called strategyproofness in the following). For this strategyproofness notion, the *random dictatorship theorem* by Gibbard (1977) shows that all strategyproof and *ex post* efficient SDSs are random dictatorships. However, random dictatorships do not allow for any compromise and typically use a lot of randomization for choosing the winner. In this paper, we thus analyze whether there are attractive strategyproof SDSs other than random dictatorships.

To answer this question, we analyze relaxations of two classic axioms, namely Condorcet-consistency and *ex post* efficiency. Condorcet-consistency requires that the Condorcet winner is chosen with probability 1 if it exists, and it follows from the work of Gibbard (1977) that no strategyproof SDS satisfies this axiom. We thus study α -Condorcet-consistency which postulates that a Condorcet winner always gets a probability of at least α . We then show that the randomized Copeland rule (which randomizes proportional to the Copeland scores) is the only anonymous, neutral, and strategyproof SDS that satisfies $\frac{2}{m}$ -Condorcet-consistency. Moreover, we prove that no other strategyproof SDS can exceed this bound, even when dropping anonymity and neutrality.

Secondly, we also study a relaxation of *ex post* efficiency. This axiom ensures that Pareto-dominated alternatives always get probability 0, and we weaken this condition by requiring that each Pareto-dominated alternative gets a probability of at most β . We call this relaxation β -*ex post* efficiency and prove a continuous strengthening of Gibbard's random dictatorship theorem based on this new axiom: the less probability we put on Pareto-dominated alternatives, the closer the resulting SDS is to a random dictatorship. In summary, this demonstrates that we cannot escape the random dictatorship theorem by weakening *ex post* efficiency as SDSs that put a negligible probability on Pareto-dominated alternatives are essentially random dictatorships.

Finally, we also identify a tradeoff between α -Condorcet-consistency and β -*ex post* efficiency for strategyproof SDSs. In particular, we show that every strategyproof and α -Condorcet-consistent SDS fails β -*ex post* efficiency for $\beta < \frac{m-2}{m-1}\alpha$. This demonstrates that we cannot jointly optimize our two objectives. Finally, we identify the anonymous, neutral, and strategyproof SDSs that optimize this tradeoff as mixtures of the randomized Copeland rule and the uniform random dictatorship.

REFERENCE

F. Brandt, P. Lederer, and R. Romen. **Relaxed notions of Condorcet-consistency and efficiency for strategyproof social decision schemes.** In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 181–189, 2022.
DOI: <https://dl.acm.org/doi/10.5555/3535850.3535872>

INDIVIDUAL CONTRIBUTION

My coauthor René Romen and I, Patrick Lederer, are the joint main authors of this paper. In particular, the results and the conceptual design have been developed in close collaboration. Moreover, we are mainly responsible for the write-up of the manuscript.

COPYRIGHT AGREEMENT

The right to present this paper in a doctoral thesis has been granted by the publisher, the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), in the copyright form presented below. There, IFAAMAS grants permission to “*personal reuse of all or portions of the above article/paper in other works of their own authorship.*” This form can also be found at <https://www.ifaamas.org/AAMAS/aamas07/IFAAMASCopyrightForm.pdf> (accessed August 24, 2023).

TERMINOLOGY

In this paper, strategyproofness refers to strong \succsim^{SD} -strategyproofness and preferences are always strict.

**INTERNATIONAL FOUNDATION FOR AUTONOMOUS AGENTS AND
MULTIAGENT SYSTEMS (IFAAMAS)
COPYRIGHT FORM**

Title of Article/Paper: _____

Publication in Which Article Is to Appear: _____

Author's Name(s): _____

Please type or print your name as you wish it to appear in print
(Please read and sign Part A only, unless you are a government employee and created your article/paper as part of your employment. If your work was performed under Government contract, but you are not a Government employee, sign Part A and see item 5 under returned rights.)

PART A – Copyright Transfer Form

The undersigned, desiring to publish the above article/paper in a publication of the International Foundation for Autonomous Agents and MultiAgent Systems (IFAAMAS), hereby transfer their copyrights in the above article/paper to the International Foundation for Autonomous Agents and MultiAgent Systems (IFAAMAS), in order to deal with future requests for reprints, translations, anthologies, reproductions, excerpts, and other publications.

This grant will include, without limitation, the entire copyright in the article/paper in all countries of the world, including all renewals, extensions, and reversions thereof, whether such rights current exist or hereafter come into effect, and also the exclusive right to create electronic versions of the article/paper, to the extent that such right is not subsumed under copyright.

The undersigned warrants that they are the sole author and owner of the copyright in the above article/paper, except for those portions shown to be in quotations; that the article/paper is original throughout; and that the undersigned right to make the grants set forth above is complete and unencumbered.

If anyone brings any claim or action alleging facts that, if true, constitute a breach of any of the foregoing warranties, the undersigned will hold harmless and indemnify IFAAMAS, their grantees, their licensees, and their distributors against any liability, whether under judgment, decree, or compromise, and any legal fees and expenses arising out of that claim or actions, and the undersigned will cooperate fully in any defense IFAAMAS may make to such claim or action. Moreover, the undersigned agrees to cooperate in any claim or other action seeking to protect or enforce any right the undersigned has granted to IFAAMAS in the article/paper. If any such claim or action fails because of facts that constitute a breach of any of the foregoing warranties, the undersigned agrees to reimburse whomever brings such claim or action for expenses and attorneys' fees incurred therein.

Returned Rights

In return for these rights, IFAAMAS hereby grants to the above authors, and the employers for whom the work was performed, royalty-free permission to:

1. Retain all proprietary rights other than copyright (such as patent rights).
2. Personal reuse of all or portions of the above article/paper in other works of their own authorship.
3. Reproduce, or have reproduced, the above article/paper for the author's personal use, or for company use provided that IFAAMAS copyright and the source are indicated, and that the

copies are not used in a way that implies IFAAMAS endorsement of a product or service of an employer, and that the copies per se are not offered for sale. The foregoing right shall not permit the posting of the article/paper in electronic or digital form on any computer network, except by the author or the author's employer, and then only on the author's or the employer's own web page or ftp site. Such web page or ftp site, in addition to the aforementioned requirements of this Paragraph, must provide an electronic reference or link back to the IFAAMAS electronic server, and shall not post other IFAAMAS copyrighted materials not of the author's or the employer's creation (including tables of contents with links to other papers) without IFAAMAS's written permission.

4. Make limited distribution of all or portions of the above article/paper prior to publication.

5. In the case of work performed under U.S. Government contract, IFAAMAS grants the U.S. Government royalty-free permission to reproduce all or portions of the above article/paper, and to authorize others to do so, for U.S. Government purposes.

In the event the above article/paper is not accepted and published by IFAAMAS, or is withdrawn by the author(s) before acceptance by IFAAMAS, this agreement becomes null and void.

Author's Signature

Date

Employer for whom work was performed

Title (if not author)

PLEASE FAX THIS SIGNED COPYRIGHT FORM TO:

Jay Modi

FAX: 215-895-0545

Department of Computer Science

Drexel University

Philadelphia, PA, 19104

INTERNATIONAL FOUNDATION FOR AUTONOMOUS AGENTS AND
MULTIAGENT SYSTEMS

c/o Professor Ed Durfee

University of Michigan

EECS Department

2260 Hayward Street

Ann Arbor, MI 48109 USA

PART B – U.S. Government Employee Certification

This will certify that all authors of the above article/paper are employees of the U.S. Government and performed this work as part of their employment, and that the article/paper is therefore not subject to U.S. copyright protection. The undersigned warrants that they are the sole author/translator of the above article/paper, and that the article/paper is original throughout, except for those portions shown to be in quotations.

U.S. Government Employee Authorized Signature

Date

Name of Government Organization

Title (if not author)

Relaxed Notions of Condorcet-Consistency and Efficiency for Strategyproof Social Decision Schemes

Felix Brandt
Technical University of Munich
Munich, Germany
brandtf@in.tum.de

Patrick Lederer
Technical University of Munich
Munich, Germany
ledererp@in.tum.de

René Romen
Technical University of Munich
Munich, Germany
rene.romen@tum.de

ABSTRACT

Social decision schemes (SDSs) map the preferences of a group of voters over some set of m alternatives to a probability distribution over the alternatives. A seminal characterization of strategyproof SDSs by Gibbard implies that there are no strategyproof Condorcet extensions and that only random dictatorships satisfy *ex post* efficiency and strategyproofness. The latter is known as the *random dictatorship theorem*. We relax Condorcet-consistency and *ex post* efficiency by introducing a lower bound on the probability of Condorcet winners and an upper bound on the probability of Pareto-dominated alternatives, respectively. We then show that the SDS that assigns probabilities proportional to Copeland scores is the only anonymous, neutral, and strategyproof SDS that can guarantee the Condorcet winner a probability of at least $2/m$. Moreover, no strategyproof SDS can exceed this bound, even when dropping anonymity and neutrality. Secondly, we prove a continuous strengthening of Gibbard’s random dictatorship theorem: the less probability we put on Pareto-dominated alternatives, the closer to a random dictatorship is the resulting SDS. Finally, we show that the only anonymous, neutral, and strategyproof SDSs that maximize the probability of Condorcet winners while minimizing the probability of Pareto-dominated alternatives are mixtures of the uniform random dictatorship and the randomized Copeland rule.

KEYWORDS

Randomized Social Choice; Social Decision Schemes; Strategyproofness; Condorcet-consistency; *ex post* efficiency

ACM Reference Format:

Felix Brandt, Patrick Lederer, and René Romen. 2022. Relaxed Notions of Condorcet-Consistency and Efficiency for Strategyproof Social Decision Schemes. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

Multi-agent systems are often faced with problems of collective decision making: how to find a group decision given the preferences of multiple individual agents. These problems, which have been traditionally studied by economists and mathematicians, are of increasing interest to computer scientists who employ the formalisms of social choice theory to analyze computational multi-agent systems [see, e.g., 8, 9, 26, 30].

A pervasive phenomenon in collective decision making is strategic manipulation: voters may be better off by lying about their preferences than reporting them truthfully. This is problematic since all desirable properties of a voting rule are in doubt when voters act dishonestly. Thus, it is important that voting rules incentivize voters to report their true preferences. Unfortunately, Gibbard [19] and Satterthwaite [28] have shown independently that dictatorships are the only non-imposing voting rules that are immune to strategic manipulations. However, these voting rules are unacceptable for most applications because they invariably return the most preferred alternative of a fixed voter. A natural question is whether more positive results can be obtained when allowing for randomization. Gibbard [20] hence introduced *social decision schemes* (SDSs), which map the preferences of the voters to a lottery over the alternatives and defined SDSs to be *strategyproof* if no voter can obtain more expected utility for any utility representation that is consistent with his ordinal preference relation. He then gave a complete characterization of strategyproof SDSs in terms of convex combinations of two types of restricted SDSs, so-called unilaterals and duples. An important consequence of this result is the *random dictatorship theorem*: random dictatorships are the only *ex post* efficient and strategyproof SDSs. Random dictatorships are convex combinations of dictatorships, i.e., each voter is selected with some fixed probability and the top choice of the chosen voter is returned. In contrast to deterministic dictatorships, the uniform random dictatorship, in which every agent is picked with the same probability, enjoys a high degree of fairness and is in fact used in many subdomains of social choice [see, e.g., 1, 12]. As a consequence of these observations, Gibbard’s theorem has been the point of departure for a lot of follow-up work. In addition to several alternative proofs of the theorem [e.g., 14, 24, 31], there have been extensions with respect to manipulations by groups [4], cardinal preferences [e.g., 16, 23, 25], weaker notions of strategyproofness [e.g., 2, 5, 7, 29], and restricted domains of preference [e.g., 11, 15].

Random dictatorships suffer from the disadvantage that they do not allow for compromise. For instance, suppose that voters strongly disagree on the best alternative, but have a common second best alternative. In such a scenario, it seems reasonable to choose the second best alternative but random dictatorships do not allow for this compromise. On a formal level, this observation is related to the fact that random dictatorships violate *Condorcet-consistency*, which demands that an alternatives that beats all other alternatives in pairwise majority comparisons should be selected. Motivated by this observation, we analyze the limitations of strategyproof SDSs by relaxing two classic conditions: Condorcet-consistency and *ex post* efficiency. To this end, we say that an SDS is *α -Condorcet-consistent* if a Condorcet winner always receives a probability of

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

at least α and β -*ex post efficient* if a Pareto-dominated alternative always receives a probability of at most β . Moreover, we say a strategyproof SDS is γ -*randomly dictatorial* if it can be represented as a convex combination of two strategyproof SDSs, one of which is a random dictatorship that will be selected with probability γ . All of these axioms are discussed in more detail in Section 2.2.

Building on an alternative characterization of strategyproof SDSs by Barberà [3], we then show the following results (m is the number of alternatives and n the number of voters):

- Let $m, n \geq 3$. There is no strategyproof SDS that satisfies α -Condorcet-consistency for $\alpha > 2/m$. Moreover, the *randomized Copeland rule*, which assigns probabilities proportional to Copeland scores, is the only strategyproof SDS that satisfies anonymity, neutrality, and $2/m$ -Condorcet-consistency.
- Let $0 \leq \epsilon \leq 1$ and $m \geq 3$. Every strategyproof SDS that is $\frac{1-\epsilon}{m}$ -*ex post efficient* is γ -randomly dictatorial for $\gamma \geq \epsilon$. If we additionally require anonymity, neutrality, and $m \geq 4$, then only mixtures of the uniform random dictatorship and the uniform lottery satisfy this bound tightly.
- Let $m \geq 4$ and $n \geq 5$. No strategyproof SDS that is α -Condorcet-consistent is β -*ex post efficient* for $\beta < \frac{m-2}{m-1}\alpha$. If we additionally require anonymity and neutrality, then only mixtures of the uniform random dictatorship and the randomized Copeland rule satisfy $\beta = \frac{m-2}{m-1}\alpha$.

The first statement characterizes the randomized Copeland rule as the “most Condorcet-consistent” SDS that satisfies strategyproofness, anonymity, and neutrality. In fact, no strategyproof SDS can guarantee more than $2/m$ probability to the Condorcet winner, even when dropping anonymity and neutrality. The second point can be interpreted as a continuous strengthening of Gibbard’s random dictatorship theorem: the less probability we put on Pareto-dominated alternatives, the more randomly dictatorial is the resulting SDS. In particular, this theorem indicates that we cannot find appealing strategyproof SDSs by allowing that Pareto-dominated alternatives gain a small probability since the resulting SDS will be very similar to random dictatorships. The last statement identifies a tradeoff between α -Condorcet-consistency and β -*ex post efficiency*: the more probability a strategyproof SDS guarantees to the Condorcet winner, the less efficient it is. Thus, we can either maximize α for α -Condorcet-consistency or minimize β for β -*ex post efficiency* of a strategyproof SDS, which again highlights the central roles of the randomized Copeland rule and random dictatorships.

2 THE MODEL

Let $N = \{1, 2, \dots, n\}$ be a finite set of voters and let $A = \{a, b, \dots\}$ be a finite set of m alternatives. Every voter i has a *preference relation* \succsim_i , which is an anti-symmetric, complete, and transitive binary relation on A . We write $x \succ_i y$ if voter i prefers x strictly to y and $x \succeq_i y$ if $x \succ_i y$ or $x = y$. The set of all preference relations is denoted by \mathcal{R} . A *preference profile* $R \in \mathcal{R}^n$ contains the preference relation of each voter $i \in N$. We define the *supporting size* for x against y in the preference profile R by $n_{x,y}(R) = |\{i \in N : x \succ_i y\}|$.

Given a preference profile, we are interested in the winning chance of each alternative. We therefore analyze social decision schemes (SDSs), which map each preference profile to a lottery over the alternatives. A *lottery* p is a probability distribution over the

set of alternatives A , i.e., it assigns each alternative x a probability $p(x) \geq 0$ such that $\sum_{x \in A} p(x) = 1$. The set of all lotteries over A is denoted by $\Delta(A)$. Formally, a *social decision scheme* (SDS) is a function $f : \mathcal{R}^n \rightarrow \Delta(A)$. We denote with $f(R, x)$ the probability assigned to alternative x by f for the preference profile R .

Since there is a huge number of SDSs, we now discuss axioms formalizing desirable properties of these functions. Two basic fairness conditions are anonymity and neutrality. Anonymity requires that voters are treated equally. Formally, an SDS f is *anonymous* if $f(R) = f(\pi(R))$ for all preference profiles R and permutations $\pi : N \rightarrow N$. Here, $R' = \pi(R)$ denotes the profile with $\succsim'_{\pi(i)} = \succsim_i$ for all voters $i \in N$. *Neutrality* guarantees that alternatives are treated equally and formally requires for an SDS f that $f(R, x) = f(\tau(R), \tau(x))$ for all preference profiles R and permutations $\tau : A \rightarrow A$. This time, $R' = \tau(R)$ is the profile derived by permuting the alternatives in R according to τ , i.e., $\tau(x) \succ'_i \tau(y)$ if and only if $x \succ_i y$ for all alternatives $x, y \in A$ and voters $i \in N$.

2.1 Stochastic Dominance and Strategyproofness

This paper is concerned with strategyproof SDSs, i.e., social decision schemes in which voters cannot benefit by lying about their preferences. In order to make this formally precise, we need to specify how voters compare lotteries. To this end, we leverage the well-known notion of stochastic dominance: a voter i (weakly) prefers a lottery p to another lottery q , written as $p \succeq_i q$, if $\sum_{y \in A: y \succ_i x} p(y) \geq \sum_{y \in A: y \succ_i x} q(y)$ for every alternative $x \in A$. Less formally, a voter prefers a lottery p weakly to a lottery q if, for every alternative $x \in A$, p returns a better alternative than x with as least as much probability as q . Stochastic dominance does not induce a complete order on the set of lotteries, i.e., there are lotteries p and q such that a voter i neither prefers p to q nor q to p .

Based on stochastic dominance, we can now formalize strategyproofness. An SDS f is *strategyproof* if $f(R) \succeq_i f(R')$ for all preference profiles R and R' and voters $i \in N$ such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$. Less formally, strategyproofness requires that every voter prefers the lottery obtained by voting truthfully to any lottery that he could obtain by voting dishonestly. Conversely, we call an SDS f *manipulable* if it is not strategyproof. While there are other ways to compare lotteries with each other, stochastic dominance is the most common one [see, e.g., 2, 3, 6, 17, 20]. This is mainly due to the fact that $p \succeq_i q$ implies that the expected utility of p is at least as high as the expected utility of q for every vNM utility function that is ordinally consistent with voter i ’s preferences. Hence, if an SDS is strategyproof, no voter can manipulate regardless of his exact utility function [see, e.g., 7, 29]. This observation immediately implies that the *convex combination* $h = \lambda f + (1 - \lambda)g$ (for some $\lambda \in [0, 1]$) of two strategyproof SDSs f and g is again strategyproof: a manipulator who obtains more expected utility from $h(R')$ than $h(R)$ prefers $f(R')$ to $f(R)$ or $g(R')$ to $g(R)$.

Gibbard [20] shows that every strategyproof SDS can be represented as convex combinations of unilaterals and duples.¹ The terms “unilaterals” and “duples” refer here to special classes of SDSs: a *unilateral* is a strategyproof SDS that only depends on the

¹In order to simplify the exposition, we slightly modified Gibbard’s terminology by requiring that duples and unilaterals have to be strategyproof.

preferences of a single voter i , i.e., $f(R) = f(R')$ for all preference profiles R and R' such that $\succ_i = \succ'_i$. A *duple*, on other hand, is a strategyproof SDS that only chooses between two alternatives x and y , i.e., $f(R, z) = 0$ for all preference profiles R and alternatives $z \in A \setminus \{x, y\}$.

Theorem 1 (Gibbard [20]). *An SDS is strategyproof if and only if it can be represented as a convex combination of unilaterals and duples.*

Since duples and unilaterals are by definition strategyproof, Theorem 1 only states that strategyproof SDSs can be decomposed into a mixture of strategyproof SDSs, each of which must be of a special type. In order to circumvent this restriction, Gibbard proves another characterization of strategyproof SDSs.

Theorem 2 (Gibbard [20]). *An SDS is strategyproof if and only if it is non-pervasive and localized.*

Non-perversity and localizedness are two axioms describing the behavior of an SDS. For defining these axioms, we denote with $R^{i:y^x}$ the profile derived from R by only reinforcing y against x in voter i 's preference relation. Note that this requires that $x \succ_i y$ and that there is no alternative $z \in A$ such that $x \succ_i z \succ_i y$. Then, an SDS f is *non-pervasive* if $f(R^{i:y^x}, y) \geq f(R, y)$ for all preference profiles R , voters $i \in N$, and alternatives $x, y \in A$. Moreover, an SDS is *localized* if $f(R^{i:y^x}, z) = f(R, z)$ for all preference profiles R , voters $i \in N$, and distinct alternatives $x, y, z \in A$. Intuitively, non-perversity—which is now often referred to as monotonicity—requires that the probability of an alternative only increases if it is reinforced, and localizedness that the probability of an alternative does not depend on the order of the other alternatives. Together, Theorem 1 and Theorem 2 show that each strategyproof SDS can be represented as a mixture of unilaterals and duples, each of which is non-pervasive and localized.

Since Gibbard's results can be quite difficult to work with, we now state another characterization of strategyproof SDSs due to Barberà [3]. Barberà has shown that every strategyproof SDS that satisfies anonymity and neutrality can be represented as a convex combination of a supporting size SDS and a point voting SDS. A *point voting SDS* is defined by a scoring vector (a_1, a_2, \dots, a_m) that satisfies $a_1 \geq a_2 \geq \dots \geq a_m \geq 0$ and $\sum_{i \in \{1, \dots, m\}} a_i = \frac{1}{n}$. The probability assigned to an alternative x by a point voting SDS f is $f(R, x) = \sum_{i \in N} a_i |\{y \in A: y \succ_i x\}|$. Furthermore, *supporting size SDSs* also rely on a scoring vector $(b_n, b_{n-1}, \dots, b_0)$ with $b_n \geq b_{n-1} \geq \dots \geq b_0 \geq 0$ and $b_i + b_{n-i} = \frac{2}{m(m-1)}$ for all $i \in \{0, \dots, n\}$ to compute the outcome. The probability assigned to an alternative x by a supporting size SDS f is then $f(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)}$. Note that point voting SDSs can be seen as a generalization of (deterministic) positional scoring rules and supporting size SDSs can be seen as a variant of Fishburn's C2 functions [18].

Theorem 3 (Barberà [3]). *An SDS is anonymous, neutral, and strategyproof if and only if it can be represented as a convex combination of a point voting SDS and a supporting size SDS.*

Many well-known SDSs can be represented as point voting SDSs or supporting size SDSs. For example, the *uniform random dictatorship* f_{RD} , which chooses one voter uniformly at random and returns his best alternative, is the point voting SDS defined by the scoring vector $(\frac{1}{n}, 0, \dots, 0)$. An instance of a supporting size SDS is the

randomized Copeland rule f_C , which assigns probabilities proportional to the Copeland scores $c(x, R) = |\{y \in A \setminus \{x\} : n_{xy}(R) > n_{yx}(R)\}| + \frac{1}{2} |\{y \in A \setminus \{x\} : n_{xy}(R) = n_{yx}(R)\}|$. This SDS is the supporting size SDS defined by the vector $b = (b_n, b_{n-1}, \dots, b_0)$, where $b_i = \frac{2}{m(m-1)}$ if $i > \frac{n}{2}$, $b_i = \frac{1}{m(m-1)}$ if $i = \frac{n}{2}$, and $b_i = 0$ otherwise. Furthermore, there are SDSs that can be represented both as point voting SDSs and supporting size SDSs. An example is the *randomized Borda rule* f_B , which randomizes proportional to the Borda scores of the alternatives. This SDS is the point voting SDS defined by the vector $(\frac{2(m-1)}{nm(m-1)}, \frac{2(m-2)}{nm(m-1)}, \dots, \frac{2}{nm(m-1)}, 0)$ and equivalently the supporting size SDS defined by the vector $(\frac{2n}{nm(m-1)}, \frac{2(n-1)}{nm(m-1)}, \dots, \frac{2}{nm(m-1)}, 0)$. Both the randomized Copeland rule and the randomized Borda rule were rediscovered several times by authors who were apparently unaware of Barberà's work [see 13, 21, 22, 27].

2.2 Relaxing Classic Axioms

The goal of this paper is to identify attractive strategyproof SDSs other than random dictatorships by relaxing classic axioms from social choice theory. In more detail, we investigate how much probability can be guaranteed to Condorcet winners and how little probability must be assigned to Pareto-dominated alternatives by strategyproof SDSs. In the following we formalize these ideas using α -Condorcet-consistency and β -ex post efficiency.

Let us first consider β -ex post efficiency, which is based on Pareto-dominance. An alternative x *Pareto-dominates* another alternative y in a preference profile R if $x \succ_i y$ for all $i \in N$. The standard notion of *ex post efficiency* then formalizes that Pareto-dominated alternatives should have no winning chance, i.e., $f(R, x) = 0$ for all preference profiles R and alternatives x that are Pareto-dominated in R . As first shown by Gibbard, random dictatorships are the only strategyproof SDSs that satisfy *ex post efficiency*. These SDSs choose each voter with a fixed probability and return his best alternative as winner. However, this result breaks down once we allow that Pareto-dominated alternatives can have a non-zero chance of winning $\beta > 0$. For illustrating this point, consider a random dictatorship d and another strategyproof SDS g . Then, the SDS $f^* = (1-\beta)d + \beta g$ is strategyproof for every $\beta \in (0, 1]$ and no random dictatorship, but assigns a probability of at most β to Pareto-dominated alternatives. We call the last property β -ex post efficiency: an SDS f is β -ex post efficient if $f(R, x) \leq \beta$ for all preference profiles R and alternatives x that are Pareto-dominated in R .

A natural generalization of the random dictatorship theorem is to ask which strategyproof SDSs satisfy β -ex post efficiency for small values of β . If β is sufficiently small, β -ex post efficiency may be quite acceptable. As we show, the random dictatorship theorem is quite robust in the sense that all SDSs that satisfy β -ex post efficiency for $\beta < \frac{1}{m}$ are similar to random dictatorships. In order to formalize this observation, we introduce γ -randomly dictatorial SDSs: a strategyproof SDS f is γ -randomly dictatorial if $\gamma \in [0, 1]$ is the maximal value such that f can be represented as $f = \gamma d + (1-\gamma)g$, where d is a random dictatorship and g is another strategyproof SDS. In particular, we require that g is strategyproof as otherwise, SDSs that seem “non-randomly dictatorial” are not 0-randomly dictatorial. For instance, the uniform lottery f_U , which

1	1	1		1	1	1
a	b	c		a	b	c
c	c	a		b	c	a
b	a	b		c	a	b
R				R'		

Figure 1: Condorcet-consistent SDSs violate strategyproofness when $m = n = 3$. Due to the symmetry of R' , we may assume without loss of generality that $f(R', a) > 0$. Since f is Condorcet-consistent, it holds that $f(R, c) = 1$. Thus, voter 1 can manipulate by swapping c and b in R .

always assigns probability $\frac{1}{m}$ to all alternatives, is not 0-randomly dictatorial if g is not required to be strategyproof because it can be represented as $f_U = \frac{1}{m}d_i + \frac{m-1}{m}g$, where d_i is the dictatorial SDS of voter i and g is the SDS that randomizes uniformly over all alternatives but voter i 's favorite one. Moreover, it should be mentioned that the maximality of γ implies that g is 0-randomly dictatorial if $\gamma < 1$. Otherwise, we could also represent g as a mixture of a random dictatorship and some other strategyproof SDS h , which means that f is γ' -randomly dictatorial for $\gamma' > \gamma$.

For a better understanding of γ -randomly dictatorial SDSs, we provide next a characterization of these SDSs. Recall for the following lemma that $R^{i:yx}$ denotes the profile derived from R by only reinforcing y against x in voter i 's preference relation.

Lemma 1. *A strategyproof SDS f is γ -randomly dictatorial if and only if there are non-negative values $\gamma_1, \dots, \gamma_n$ such that:*

- i) $\sum_{i \in N} \gamma_i = \gamma$.
- ii) $f(R^{i:yx}, y) - f(R, y) \geq \gamma_i$ for all alternatives $x, y \in A$, voters $i \in N$, and preference profiles R in which voter i prefers x the most and y the second most.
- iii) for every voter $i \in N$, there are alternatives $x, y \in A$ and a profile R such that voter i prefers x the most and y the second most in R , and $f(R^{i:yx}, y) - f(R, y) = \gamma_i$.

The proof of this lemma can be found in the extended version [10]. Lemma 1 gives an intuitive interpretation of γ -randomly dictatorial SDSs: this axiom only requires that there are voters who always increase the winning probability of an alternative by at least γ_i if they reinforce it to the first place. Hence, for small values of γ , this axiom is desirable as it only formulates a variant of strict monotonicity. However, for larger values of γ , γ -randomly dictatorial SDSs become more similar to random dictatorships. Furthermore, the proof of Lemma 1 shows that the decomposition of γ -randomly dictatorial SDSs is completely determined by the values $\gamma_1, \dots, \gamma_n$: given these values for an strategyproof SDS f , it can be represented as $f = \sum_{i \in N} \gamma_i d_i + (1 - \sum_{i \in N} \gamma_i)g$, where g is a strategyproof SDS and d_i the dictatorial SDS of voter i .

Finally, we introduce α -Condorcet-consistency. To this end, we first define the notion of a Condorcet winner. A *Condorcet winner* is an alternative x that wins every majority comparison according to preference profile R , i.e., $n_{xy}(R) > n_{yx}(R)$ for all $y \in A \setminus \{x\}$. *Condorcet-consistency* demands that $f(R, x) = 1$ for all preference profiles R and alternatives x such that x is the Condorcet winner

Table 1: Values of α , β , and γ for which specific SDSs are α -Condorcet-consistent, β -ex post efficient, and γ -randomly dictatorial. Each row shows the values of α , β , and γ for which a specific SDS satisfies the corresponding axioms. f_{RD} abbreviates the uniform random dictatorship, f_U the uniform lottery, f_B the randomized Borda rule, and f_C the randomized Copeland rule.

SDS	α -Condorcet -consistency	β -ex post efficiency	γ -random dictatorship
f_{RD}	0	0	1
f_U	$\frac{1}{m}$	$\frac{1}{m}$	0
f_B	$\frac{1}{m} + \frac{2-(n \bmod 2)}{mn}$	$\frac{2(m-2)}{m(m-1)}$	$\frac{2}{m(m-1)}$
f_C	$\frac{2}{m}$	$\frac{2(m-2)}{m(m-1)}$	0

in R . Unfortunately, Condorcet-consistency is in conflict with strategyproofness, which can easily be derived from Gibbard's random dictatorship theorem. A simple two-profile proof for this fact when $m = n = 3$ is given in Figure 1. To circumvent this impossibility, we relax Condorcet-consistency: instead of requiring that the Condorcet winner always obtains probability 1, we only require that it receives a probability of at least α . This idea leads to α -Condorcet-consistency: an SDS f satisfies this axiom if $f(R, x) \geq \alpha$ for all profiles R and alternatives $x \in A$ such that x is the Condorcet winner in R . For small values of α , this axiom is clearly compatible with strategyproofness and therefore, we are interested in the maximum value of α such that there are α -Condorcet-consistent and strategyproof SDSs.

For a better understanding of α -Condorcet-consistency, β -ex post efficiency, and γ -random dictatorships, we discuss some of the values in Table 1 as examples. The uniform random dictatorship is 1-randomly dictatorial and 0-ex post efficient by definition. Moreover, it is 0-Condorcet-consistent because a Condorcet winner may not be top-ranked by any voter. The randomized Borda rule is $\frac{2(m-2)}{m(m-1)}$ -ex post efficient because it assigns this probability to an alternative that is second-ranked by every voter. Moreover, it is $\frac{2}{m(m-1)}$ -randomly dictatorial as we can represent it as $\frac{2}{m(m-1)}f_{RD} + \left(1 - \frac{2}{m(m-1)}\right)g$, where f_{RD} is the uniform random dictatorship and g is the point voting SDS defined by the scoring vector $\left(\frac{2(m-2)}{n(m(m-1)-2)}, \frac{2(m-2)}{n(m(m-1)-2)}, \frac{2(m-3)}{n(m(m-1)-2)}, \dots, 0\right)$. Finally, the randomized Copeland rule is 0-randomly dictatorial because there is for every voter a profile in which he can swap his two best alternatives without affecting the outcome. Moreover, it is $\frac{2}{m}$ -Condorcet-consistent because a Condorcet winner x satisfies that $n_{xy}(R) > \frac{n}{2}$ for all $y \in A \setminus \{x\}$ and hence, $f_C(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)} = (m-1) \frac{2}{m(m-1)} = \frac{2}{m}$. Note that Table 1 also contains a row corresponding to the uniform lottery. We consider this SDS as a threshold with respect to α -Condorcet-consistency and β -ex post efficiency because we can compute the uniform lottery without knowledge about the voters' preferences. Hence, if an SDS performs worse than the uniform lottery with respect to α -Condorcet-consistency or β -ex post efficiency, we could also dismiss the voters' preferences.

3 RESULTS

In this section, we present our results about the α -Condorcet-consistency and the β -*ex post* efficiency of strategyproof SDSs. First, we prove that no strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$ and that the randomized Copeland rule f_C is the only anonymous, neutral, and strategyproof SDS that satisfies α -Condorcet-consistency for $\alpha = \frac{2}{m}$. Moreover, we show that every $\frac{1-\epsilon}{m}$ -*ex post* efficient and strategyproof SDS is γ -randomly dictatorial for $\gamma \geq \epsilon$. This statement can be seen as a continuous generalization of the random dictatorship theorem and implies, for instance, that every 0-randomly dictatorial and strategyproof SDS can only satisfy β -*ex post* efficiency for $\beta \geq \frac{1}{m}$, i.e., such SDSs are at least as inefficient as the uniform lottery. Even more, when additionally imposing anonymity and neutrality, we prove that only mixtures of the uniform random dictatorship and the uniform lottery satisfy this bound tightly, which shows that relaxing *ex post* efficiency does not allow for appealing SDSs. In the last theorem, we identify a tradeoff between Condorcet-consistency and *ex post* efficiency: no strategyproof SDS that satisfies α -Condorcet-consistency is β -*ex post* efficient for $\beta < \frac{m-2}{m-1}\alpha$. We derive these results through a series of lemmas. Because of space restrictions, the proofs of all lemmas and Theorem 5 are deferred to an extended version of this paper [10] and we only present short proof sketches instead.

3.1 α -Condorcet-consistency

As discussed in Section 2.2, Condorcet-consistent SDSs violate strategyproofness. Therefore, we analyze the maximal α such that α -Condorcet-consistency and strategyproofness are compatible. Our results show that strategyproofness only allows for a small degree of Condorcet-consistency: we prove that no strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$. This bound is tight as the randomized Copeland rule f_C is $\frac{2}{m}$ -Condorcet-consistent, which means that it is one of the “most Condorcet-consistent” strategyproof SDSs. Even more, we can turn this observation in a characterization of f_C by additionally requiring anonymity and neutrality: the randomized Copeland rule is the only strategyproof SDS that satisfies $\frac{2}{m}$ -Condorcet-consistency, anonymity, and neutrality.

For proving these results, we derive next a number of lemmas. As first step, we show in Lemma 2 that we can use a strategyproof and α -Condorcet-consistent SDS to construct another strategyproof SDS that satisfies anonymity, neutrality, and α -Condorcet-consistency for the same α .

Lemma 2. *If a strategyproof SDS satisfies α -Condorcet-consistency for some $\alpha \in [0, 1]$, there is also a strategyproof SDS that satisfies anonymity, neutrality, and α -Condorcet-consistency for the same α .*

The central idea in the proof of Lemma 2 is the following: if there is a strategyproof and α -Condorcet-consistent SDS f , then the SDS $f^{\pi\tau}(R, x) = f(\tau(\pi(R)), \tau(x))$ is also strategyproof and α -Condorcet-consistent for all permutations $\pi : N \rightarrow N$ and $\tau : A \rightarrow A$. Since mixtures of strategyproof and α -Condorcet-consistent SDSs are also strategyproof and α -Condorcet-consistent, we can therefore construct an SDS that satisfies all requirements of the lemma by averaging over all permutations on N and A . More formally, the SDS $f^* = \frac{1}{m!n!} \sum_{\pi \in \Pi} \sum_{\tau \in \Gamma} f^{\pi\tau}$ (where Π denotes the

set of all permutations on N and Γ the set of all permutations on A) meets all criteria of the lemma.

Due to Lemma 2, we investigate next the α -Condorcet-consistency of strategyproof SDSs that satisfy anonymity and neutrality. The reason for this is that this lemma turns an upper bound on α for these SDSs into an upper bound for all strategyproof SDSs. Since Theorem 3 shows that every strategyproof, anonymous, and neutral SDS can be decomposed in a point voting SDS and a supporting size SDS, we investigate these two classes separately in the following two lemmas. First, we bound the α -Condorcet-consistency of point voting SDSs.

Lemma 3. *No point voting SDS is α -Condorcet-consistent for $\alpha \geq \frac{2}{m}$ if $n \geq 3$ and $m \geq 3$.*

The proof of this lemma relies on the observation that there can be $\lceil \frac{m}{2} \rceil$ Condorcet winner candidates, i.e., alternatives x that can be made into the Condorcet winner by keeping x at the same position in the preferences of every voter and only reordering the other alternatives. Since reordering the other alternatives does not affect the probability of x in a point voting SDS, it follows that every Condorcet winner candidate has a probability of at least α . Hence, we derive that $\alpha \leq \frac{\lceil \frac{m}{2} \rceil}{m} \leq \frac{2}{m}$ and a slightly more involved argument shows that the inequality is strict.

The last ingredient for the proof of Theorem 4 is that no supporting size SDS can assign a probability of more than $\frac{2}{m}$ to any alternative. This immediately implies that no supporting size SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$.

Lemma 4. *No supporting size SDS can assign more than $\frac{2}{m}$ probability to an alternative.*

The proof of this lemma follows straightforwardly from the definition of supporting size SDSs. Each such SDS is defined by a scoring vector (b_n, \dots, b_0) such that $b_i + b_{n-i} = \frac{2}{m(m-1)}$ for all $i \in \{0, \dots, n\}$ and $b_n \geq b_{n-1} \geq \dots \geq b_0 \geq 0$. The probability of an alternative x in a supporting size SDS f is therefore bounded by $f(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)} \leq (m-1) \frac{2}{m(m-1)} = \frac{2}{m}$.

Finally, we have all necessary lemmas for the proof of our first theorem.

Theorem 4. *The randomized Copeland rule is the only strategyproof SDS that satisfies anonymity, neutrality, and $\frac{2}{m}$ -Condorcet-consistency if $m \geq 3$ and $n \geq 3$. Moreover, no strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$ if $n \geq 3$.*

Proof. The theorem consists of two claims: the characterization of the randomized Condorcet rule f_C and the fact that no other strategyproof SDS can attain α -Condorcet-consistency for a larger α than f_C . We prove these claims separately.

Claim 1: The randomized Copeland rule is the only strategyproof SDS that satisfies $\frac{2}{m}$ -Condorcet-consistency, anonymity, and neutrality if $m, n \geq 3$.

The randomized Copeland rule f_C is a supporting size SDS and satisfies therefore anonymity, neutrality, and strategyproofness. Furthermore, it satisfies also $\frac{2}{m}$ -Condorcet-consistency because a Condorcet winner x wins every pairwise majority comparison in R . Hence, $n_{xy}(R) > \frac{n}{2}$ for all $y \in A \setminus \{x\}$, which implies that $f_C(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)} = (m-1) \frac{2}{m(m-1)} = \frac{2}{m}$.

Next, let f be an SDS satisfying anonymity, neutrality, strategyproofness, and $\frac{2}{m}$ -Condorcet-consistency. We show that f is the randomized Copeland rule. Since f is anonymous, neutral, and strategyproof, we can apply Theorem 3 to represent f as $f = \lambda f_{\text{point}} + (1 - \lambda) f_{\text{sup}}$, where $\lambda \in [0, 1]$, f_{point} is a point voting SDS, and f_{sup} is a supporting size SDS. Lemma 3 states that there is a profile R with Condorcet winner x such that $f_{\text{point}}(R, x) < \frac{2}{m}$, and it follows from Lemma 4 that $f_{\text{sup}}(R, x) \leq \frac{2}{m}$. Hence, $f(R, x) = \lambda f_{\text{point}}(R, x) + f_{\text{sup}}(R, x) < \frac{2}{m}$ if $\lambda > 0$. Therefore, f is a supporting size SDS as it satisfies $\frac{2}{m}$ -Condorcet-consistency.

Next, we show that f has the same scoring vector as the randomized Copeland rule. Since f is a supporting size SDS, there is a scoring vector $b = (b_n, \dots, b_0)$ with $b_n \geq b_{n-1} \geq \dots \geq b_0 \geq 0$ and $b_i + b_{n-i} = \frac{2}{m(m-1)}$ for all $i \in \{1, \dots, n\}$ such that $f(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)}$. Moreover, $f(R, x) = \frac{2}{m}$ if x is the Condorcet winner in R because of $\frac{2}{m}$ -Condorcet-consistency and Lemma 4. We derive from the definition of supporting size SDSs that the Condorcet winner x can only achieve this probability if $b_{n_{xy}(R)} = \frac{2}{m(m-1)}$ for every other alternatives $y \in A \setminus \{x\}$. Moreover, observe that the Condorcet winner needs to win every majority comparison but is indifferent about the exact supporting sizes. Hence, it follows that $b_i = \frac{2}{m(m-1)}$ for all $i > \frac{n}{2}$ as otherwise, there is a profile in which the Condorcet winner does not receive a probability of $\frac{2}{m}$. We also know that $b_i + b_{n-i} = \frac{2}{m(m-1)}$, so $b_i = 0$ for all $i < \frac{n}{2}$. If n is even, then $b_{\frac{n}{2}} = \frac{1}{m(m-1)}$ is required by the definition of supporting size SDSs as $\frac{n}{2} = n - \frac{n}{2}$. Hence, the scoring vector of f is equivalent to the scoring vector of the randomized Copeland rule, which proves that f is f_C .

Claim 2: No strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$ if $n \geq 3$.

The claim is trivially true if $m \leq 2$ because α -Condorcet consistency for $\alpha > 1$ is impossible. Hence, let f denote a strategyproof SDS for $m \geq 3$ alternatives. We show in the sequel that f cannot satisfy α -Condorcet-consistency for $\alpha > \frac{2}{m}$. As a first step, we use Lemma 2 to construct a strategyproof SDS f^* that satisfies anonymity, neutrality, and α -Condorcet-consistency for the same α as f . Since f^* is anonymous, neutral, and strategyproof, it follows from Theorem 3 that f^* can be represented as a mixture of a point voting SDS f_{point} and a supporting size SDS f_{sup} , i.e., $f^* = \lambda f_{\text{point}} + (1 - \lambda) f_{\text{sup}}$ for some $\lambda \in [0, 1]$.

Next, we consider f_{point} and f_{sup} separately. Lemma 3 implies for f_{point} that there is a profile R with a Condorcet winner a such that $f_{\text{point}}(R, a) < \frac{2}{m}$. Moreover, Lemma 4 shows that $f_{\text{sup}}(R, a) \leq \frac{2}{m}$ because supporting size SDSs never return a larger probability than $\frac{2}{m}$. Thus, we derive the following inequality. w

$$\alpha \leq f^*(R, a) = \lambda f_{\text{point}}(R, a) + (1 - \lambda) f_{\text{sup}}(R, a) \leq \lambda \frac{2}{m} + (1 - \lambda) \frac{2}{m} = \frac{2}{m}$$

This proves that f^* , and therefore every strategyproof SDS, fails α -Condorcet-consistency for $\alpha \geq \frac{2}{m}$. \square

Remark 1. Lemma 2 can be applied to properties other than α -Condorcet-consistency, too. For example, given a strategyproof and β -ex post efficient SDS, we can construct another SDS that satisfies these axioms as well as anonymity and neutrality.

Remark 2. All axioms in the characterization of the randomized Copeland rule are independent of each other. The SDS that picks the Condorcet winner with probability $\frac{2}{m}$ if one exists and distributes the remaining probability uniformly between the other alternatives only violates strategyproofness. The randomized Borda rule satisfies all axioms of Theorem 4 but $\frac{2}{m}$ -Condorcet-consistency. An SDS that satisfies anonymity, strategyproofness, and $\frac{2}{m}$ -Condorcet-consistency can be defined based on an arbitrary order of alternatives x_0, \dots, x_{m-1} . Then, we pick an index $i \in \{0, \dots, m-1\}$ uniformly at random and return the winner of the majority comparison between x_i and $x_{i+1 \bmod m}$ (if there is a majority tie, a fair coin toss decides the winner). Finally, we can use the randomized Copeland rule f_C to construct an SDS that fails only anonymity for even n : we just ignore one voter when computing the outcome of f_C . Note here that for even n , an alternative x is a Condorcet winner in profile R if $n_{xy}(R) \geq \frac{n+2}{2}$ for all $y \in N \setminus \{x\}$, which means that x remains the Condorcet winner after removing a single voter.

Moreover, the impossibility in Theorem 4 does not hold when there are only $n = 2$ voters because random dictatorships are strategyproof and Condorcet-consistent in this case. The reason for this is that a Condorcet winner needs to be the most preferred alternative of both voters and is therefore chosen with probability 1.

Remark 3. The randomized Copeland rule has multiple appealing interpretations. Firstly, it can be defined as a supporting size SDS as shown in Section 2.1. Alternatively, it can be defined as the SDS that picks two alternatives uniformly at random and then picks the majority winner between them; majority ties are broken by a fair coin toss. Next, Theorem 4 shows that the randomized Copeland rule is the SDS that maximizes the value of α for α -Condorcet-consistency among all anonymous, neutral, and strategyproof SDSs. Finally, the randomized Copeland rule is the only strategyproof SDS that satisfies anonymity, neutrality, and assigns 0 probability to a Condorcet loser whenever it exists.

3.2 β -ex post Efficiency

According to Gibbard's random dictatorship theorem, random dictatorships are the only strategyproof SDSs that satisfy *ex post* efficiency. In this section, we show that this result is rather robust by identifying a tradeoff between β -ex post efficiency and γ -random dictatorships. More formally, we prove that for every $\epsilon \in [0, 1]$, all strategyproof and $\frac{1-\epsilon}{m}$ -ex post efficient SDSs are γ -randomly dictatorial for $\gamma \geq \epsilon$. If we set $\epsilon = 1$, we obtain the random dictatorship theorem. On the other hand, we derive from this theorem that every 0-randomly dictatorial and strategyproof SDS is β -ex post efficient for $\beta \geq \frac{1}{m}$, i.e., every such SDS is at least as inefficient as the uniform lottery. Moreover, we prove for every $\epsilon \in [0, 1]$ that mixtures of the uniform random dictatorship and the uniform lottery are the only ϵ -randomly dictatorial SDSs that satisfy anonymity, neutrality, strategyproofness, and $\frac{1-\epsilon}{m}$ -ex post efficiency. In summary, these results demonstrate that relaxing *ex post* efficiency does not lead to particularly appealing strategyproof SDSs. Furthermore, we also identify a tradeoff between α -Condorcet-consistency and β -ex post efficiency: every α -Condorcet-consistent and strategyproof SDS fails β -ex post efficiency for $\beta < \frac{m-1}{m-2}\alpha$. Under the additional assumption of anonymity and neutrality, we characterize the strategyproof SDSs that maximize the ratio between α and β : all these

SDSs are mixtures of the randomized Copeland rule and the uniform random dictatorship.

For proving the tradeoff between β -ex post efficiency and γ -random dictatorships, we first investigate the efficiency of 0-randomly dictatorial strategyproof SDSs. In more detail, we prove next that every such SDS fails β -ex post efficiency for $\beta < \frac{1}{m}$.

Lemma 5. *No strategyproof SDS that is 0-randomly dictatorial satisfies β -ex post efficiency for $\beta < \frac{1}{m}$ if $m \geq 3$.*

The proof of this result is quite similar to the one for the upper bound on α -Condorcet-consistency in Theorem 4. In particular, we first show that all 0-randomly mixtures of duples and all 0-randomly dictatorial mixtures of unilaterals violate β -ex post efficiency for $\beta < \frac{1}{m}$. Next, we consider an arbitrary 0-randomly dictatorial SDS f and aim to show that there are a profile R and a Pareto-dominated alternative $x \in A$ such that $f(R, x) \geq \beta$. Even though Theorem 1 allows us to represent f as the convex combination of a 0-randomly dictatorial mixture of unilaterals f_{uni} and a mixture of duples f_{duple} , our previous observations have unfortunately no direct consequences for the β -ex post efficiency of f . The reason for this is that f_{uni} and f_{duple} might violate β -ex post efficiency for different profiles or alternatives. We solve this problem by transforming f into a 0-randomly dictatorial SDS f^* that is β -ex post efficient for the same β as f and satisfies additional properties. In particular, f^* can be represented as a convex combination of a 0-randomly dictatorial mixture of unilaterals f_{uni}^* and a 0-randomly dictatorial mixture of duples f_{duple}^* such that $f_{uni}^*(R, x) \geq \frac{1}{m}$ and $f_{duple}^*(R, x) \geq \frac{1}{m}$ for some profile R in which alternative x is Pareto-dominated. Consequently, f^* fails β -ex post efficiency for $\beta < \frac{1}{m}$, which implies that also f violates this axiom.

Based on Lemma 5, we can now show the tradeoff between ex post efficiency and the similarity to a random dictatorship.

Theorem 5. *For every $\epsilon \in [0, 1]$, every strategyproof and $\frac{1-\epsilon}{m}$ -ex post efficient SDS is γ -randomly dictatorial for $\gamma \geq \epsilon$ if $m \geq 3$. Moreover, if $\gamma = \epsilon$, $m \geq 4$, and the SDS satisfies additionally anonymity and neutrality, it is a mixture of the uniform random dictatorship and the uniform lottery.*

The proof of the first claim follows easily from Lemma 5: we consider a strategyproof SDS f and use the definition of γ -randomly dictatorial SDSs to represent f as a mixture of a random dictatorship and another strategyproof SDS g . Unless f is a random dictatorship, the maximality of γ entails that g is 0-randomly dictatorial. Hence, Lemma 5 implies that g can only be β -ex post efficient for $\beta \geq \frac{1-\epsilon}{m}$. Consequently, $\gamma \geq \epsilon$ must be true if f satisfies $\frac{1-\epsilon}{m}$ -ex post efficiency. For the second claim, we observe first that every anonymous, neutral, and strategyproof SDS f can be represented as a mixture of the uniform random dictatorship and another strategyproof, anonymous, and neutral SDS g . Moreover, unless f is 1-randomly dictatorial, g is 0-randomly dictatorial. Thus, Lemma 5 and the assumption that $\gamma = \epsilon$ require that g is exactly $\frac{1}{m}$ -ex post efficient. Finally, the claim follows by proving that the uniform lottery is the only 0-randomly dictatorial and strategyproof SDS that satisfies anonymity, neutrality, and $\frac{1}{m}$ -ex post efficiency if $m \geq 4$. For $m = 3$ the randomized Copeland rule also satisfies all required axioms and the uniform rule is thus not the unique choice.

Theorem 5 represents a continuous strengthening of Gibbard's random dictatorship theorem: the more ex post efficiency is required, the closer a strategyproof SDS gets to a random dictatorship. Conversely, our result also entails that γ -randomly dictatorial SDSs can only satisfy $\frac{1-\epsilon}{m}$ -ex post efficiency for $\epsilon \leq \gamma$. Moreover, the second part of the theorem indicates that relaxing ex post efficiency does not allow for particularly appealing strategyproof SDSs.

The correlation between β -ex post efficiency and γ -randomly dictatorships also suggests a tradeoff between α -Condorcet-consistency and β -ex post efficiency because all random dictatorships are 0-Condorcet-consistent for sufficiently large m and n . Perhaps surprisingly, we show next that α -Condorcet consistency and β -ex post efficiency are in relation with each other for strategyproof SDSs. As a consequence of this insight, two strategyproof SDSs are particularly interesting: random dictatorships because they are the most ex post efficient SDSs, and the randomized Copeland rule because it is the most Condorcet-consistent SDS.

Theorem 6. *Every strategyproof SDS that satisfies anonymity, neutrality, α -Condorcet consistency, and β -ex post efficiency with $\beta = \frac{m-2}{m-1}\alpha$ is a mixture of the uniform random dictatorship and the randomized Copeland rule if $m \geq 4$, $n \geq 5$. Furthermore, there is no strategyproof SDS with $\beta < \frac{m-2}{m-1}\alpha$ if $m \geq 4$, $n \geq 5$.*

Proof. Let f be a strategyproof SDS that satisfies α -Condorcet consistency for some $\alpha \in [0, \frac{2}{m}]$ and let $\beta \in [0, 1]$ denote the minimal value such that f is β -ex post efficient. We first show that $\beta \geq \frac{m-2}{m-1}\alpha$ and hence apply Lemma 2 to construct an SDS f' that satisfies strategyproofness, anonymity, neutrality, α' -Condorcet consistency for $\alpha' \geq \alpha$, and β' -ex post efficiency for $\beta' \leq \beta$. In particular, if f' is only β' -ex post efficient for $\beta' \geq \frac{m-2}{m-1}\alpha'$, then f can only satisfy β -ex post efficiency for $\beta \geq \beta' \geq \frac{m-2}{m-1}\alpha' \geq \frac{m-2}{m-1}\alpha$.

Since f' satisfies anonymity, neutrality, and strategyproofness, we can apply Theorem 3 to represent it as a mixture of a supporting size SDS and a point voting SDS, i.e., $f' = \lambda f_{point} + (1 - \lambda) f_{sup}$ for some $\lambda \in [0, 1]$. Let (a_1, \dots, a_m) and (b_0, \dots, b_n) denote the scoring vectors describing f_{point} and f_{sup} , respectively. Next, we derive lower bound for α' and an upper bound for β' by considering specific profiles. First, consider the profile R in which every voter reports a as his best alternative and b as his second best alternative; the remaining alternatives can be ordered arbitrarily. It follows from the definition of point voting SDSs that $f_{point}(R, b) = na_2$ and from the definition of supporting size SDS that $f_{sup}(R, b) = (m - 2)b_n + b_0$. Since a Pareto-dominates b in R , it follows that $\beta' \geq f(R, b) = \lambda na_2 + (1 - \lambda)((m - 2)b_n + b_0)$.

For the upper bound on α , consider the following profile R' where alternative x is never ranked first, but it is the Condorcet winner and wins every pairwise comparison only with minimal margin. We denote for the definition of R' the alternatives as $A = \{x, x_1, \dots, x_{m-1}\}$. In R' , the voters $i \in \{1, 2, 3\}$ ranks alternatives $X_i := \{x_k \in A \setminus \{x\} : k \bmod 3 = i - 1\}$ above x and all other alternatives below. Since $m \geq 4$, none of them ranks x first. If the number of voters n is even, we duplicate voters 1, 2, and 3. As last step, we add pairs of voters with inverse preferences such that no voter prefers x the most until R' consists of n voters. Since alternative x is never top-ranked in R' , it follows that $f_{point}(R', x) \leq na_2$. Furthermore, $n_{xy}(R') = \lceil \frac{n+1}{2} \rceil$ for all $y \in A \setminus \{x\}$ and therefore $f_{sup}(R', x) = (m - 1)b_{\lceil \frac{n+1}{2} \rceil}$. Finally, we

derive that $\alpha' \leq f(R', x) \leq \lambda n a_2 + (1 - \lambda)(m - 1)b_{\lceil \frac{n+1}{2} \rceil}$ because x is by construction the Condorcet winner in R' .

Using these bounds, we show next that f' is only β' -*ex post* efficiency for $\beta' \geq \frac{m-2}{m-1}\alpha'$, which proves the second claim of the theorem. In the subsequent calculation, the first and last inequality follow from our previous analysis. The second inequality is true since $\frac{m-2}{m-1} \leq 1$ and $\frac{m-2}{m-1}(m-1) = (m-2)$. The third inequality uses the definition of supporting size SDSs.

$$\begin{aligned} \beta' &\geq \lambda n a_2 + (1 - \lambda)((m - 2)b_n + b_0) \\ &\geq \frac{m-2}{m-1}\lambda n a_2 + \frac{m-2}{m-1}(1 - \lambda)((m - 1)b_n + b_0) \\ &\geq \frac{m-2}{m-1}\lambda n a_2 + \frac{m-2}{m-1}(1 - \lambda)(m - 1)b_{\lceil \frac{n+1}{2} \rceil} \\ &\geq \frac{m-2}{m-1}\alpha' \end{aligned}$$

Finally, note that, if $\beta' = \frac{m-2}{m-1}\alpha'$, all inequalities must be tight. If the second inequality is tight $a_2 = 0$ and $b_0 = 0$, and when the third inequality is tight $b_n = b_{\lceil \frac{n+1}{2} \rceil}$. These observations fully specify the scoring vectors of f_{point} and f_{sup} . For the point voting SDS, $a_2 = 0$ implies $a_i = 0$ for all $i \geq 2$ and $a_1 = \frac{1}{n}$, i.e., f_{point} is the uniform random dictatorship. Next, $b_0 = 0$ and $b_n = b_{\lceil \frac{n+1}{2} \rceil}$ imply that $b_i = \frac{2}{m(m-1)}$ for all $i \in \{\lceil \frac{n+1}{2} \rceil, \dots, b_n\}$ and $b_i = 0$ for all $i \in \{0, \dots, \lceil \frac{n+1}{2} \rceil\}$. Moreover, if n is even, the definition of supporting size SDSs requires that $b_{\frac{n}{2}} = \frac{1}{m(m-1)}$. This shows that f_{sup} is the randomized Copeland rule. Consequently, the SDS f' is a mixture of the uniform random dictatorship and the randomized Copeland rule if $\beta' = \frac{m-2}{m-1}\alpha'$. This proves that every strategyproof SDS that satisfies anonymity, neutrality, α -Condorcet consistency, and β -*ex post* efficiency with $\beta = \frac{m-2}{m-1}\alpha$ is a mixture of the uniform random dictatorship and the randomized Copeland rule. \square

Remark 4. All axioms of the characterization in Theorem 6 are independent of each other. Every mixture of random dictatorships other than the uniform one and the randomized Copeland rule only violates anonymity. An SDS that violates only neutrality can be constructed by using a variant of the randomized Copeland rule that does not split the probability equally if there is a majority tie. Finally, the correlation between α -Condorcet-consistency and β -*ex post* efficiency is required since the uniform lottery satisfies all other axioms. Moreover, all bounds on m and n in Theorem 6 are tight. If there are only $n = 2$ voters, $m = 3$ alternatives, or $m = 4$ alternatives and $n = 4$ voters, the uniform random dictatorship is not 0-Condorcet consistent since a Condorcet winner is always ranked first by at least one voter. Hence, the bound on β does not hold in these cases. In contrast, our proof shows that Theorem 6 is also true when $n = 3$.

4 CONCLUSION

In this paper, we analyzed strategyproof SDSs by considering relaxations of Condorcet-consistency and *ex post* efficiency. Our findings, which are summarized in Figure 2, show that two strategyproof SDSs perform particularly well with respect to these axioms: the uniform random dictatorship (and random dictatorships in general), and the randomized Copeland rule. In more detail, we prove that




Figure 2: Graphical summary of our results. Points in the figures correspond to SDSs and the horizontal axis indicates in both figures the value of β for which the considered SDS is β -*ex post* efficient. In the left figure, the vertical axis states the α for which the considered SDSs are α -Condorcet-consistent, and in the right figure, it shows the γ for which SDSs are γ -randomly dictatorial. Theorems 4 and 6 show that no strategyproof SDS lies in the grey area of the left figure. Theorem 5 shows that no strategyproof SDS lies in the grey area below the diagonal in the right figure. Furthermore, no SDS lies in the grey area above the diagonal since a γ -randomly dictatorial SDS can put no more than $1 - \gamma$ probability on Pareto-dominated alternatives. Finally, the following SDS are marked in the figures: d corresponds to all random dictatorships, c to the randomized Copeland rule, b to the randomized Borda rule, and u to the uniform lottery.

the randomized Copeland rule is the only strategyproof, anonymous, and neutral SDS which guarantees a probability of $\frac{2}{m}$ to the Condorcet winner. Since no other strategyproof SDS can guarantee more probability to the Condorcet winner (even if we drop anonymity and neutrality), this characterization identifies the randomized Copeland rule as one of the most Condorcet-consistent strategyproof SDSs. On the other hand, Gibbard's random dictatorship theorem shows that random dictatorships are the only *ex post* efficient and strategyproof SDSs. We present a continuous generalization of this result: for every $\epsilon \in [0, 1]$, every $\frac{1-\epsilon}{m}$ -*ex post* efficient and strategyproof SDS is γ -randomly dictatorial for $\gamma \geq \epsilon$. This means informally that, even if we allow that Pareto-dominated alternatives can get a small amount of probability, we end up with an SDS similar to a random dictatorship. Finally, we derive a tradeoff between α -Condorcet-consistency and β -*ex post* efficiency for strategyproof SDSs: every strategyproof and α -Condorcet-consistent SDS fails β -*ex post* efficiency for $\beta < \frac{m-2}{m-1}\alpha$. This theorem entails that it is not possible to jointly optimize both notions, which again highlights the special role of the randomized Copeland rule and random dictatorships.

ACKNOWLEDGMENTS

This work was supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1. We thank Dominik Peters for stimulating discussions and the anonymous reviewers for their helpful comments.

REFERENCES

- [1] A. Abdulkadiroğlu and T. Sönmez. 1998. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica* 66, 3 (1998), 689–701.
- [2] H. Aziz, F. Brandl, F. Brandt, and M. Brill. 2018. On the Tradeoff between Efficiency and Strategyproofness. *Games and Economic Behavior* 110 (2018), 1–18.
- [3] S. Barberà. 1979. Majority and Positional Voting in a Probabilistic Framework. *Review of Economic Studies* 46, 2 (1979), 379–389.
- [4] S. Barberà. 1979. A Note on Group Strategy-Proof Decision Schemes. *Econometrica* 47, 3 (1979), 637–640.
- [5] J.-P. Benoît. 2002. Strategic Manipulation in Voting Games When Lotteries and Ties Are Permitted. *Journal of Economic Theory* 102, 2 (2002), 421–436.
- [6] A. Bogomolnaia and H. Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100, 2 (2001), 295–328.
- [7] F. Brandl, F. Brandt, M. Eberl, and C. Geist. 2018. Proving the Incompatibility of Efficiency and Strategyproofness via SMT Solving. *J. ACM* 65, 2 (2018), 1–28.
- [8] F. Brandt, V. Conitzer, and U. Endriss. 2013. Computational Social Choice. In *Multiagent Systems* (2nd ed.), G. Weiß (Ed.). MIT Press, Chapter 6, 213–283.
- [9] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [10] F. Brandt, P. Lederer, and R. Romen. Relaxed notions of Condorcet-consistency and efficiency for strategyproof social decision schemes. Technical report, <https://arxiv.org/abs/2201.10418>, 2022.
- [11] S. Chatterji, A. Sen, and H. Zeng. 2014. Random dictatorship domains. *Games and Economic Behavior* 86 (2014), 212–236.
- [12] Y.-K. Che and F. Kojima. 2010. Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms. *Econometrica* 78, 5 (2010), 1625–1672.
- [13] V. Conitzer and T. Sandholm. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*. 627–634.
- [14] J. Duggan. 1996. A geometric proof of Gibbard’s random dictatorship theorem. *Economic Theory* 7, 2 (1996), 365–369.
- [15] B. Dutta, H. Peters, and A. Sen. 2002. Strategy-Proof Probabilistic Mechanisms in Economies with Pure Public Goods. *Journal of Economic Theory* 106, 2 (2002), 392–416.
- [16] B. Dutta, H. Peters, and A. Sen. 2007. Strategy-proof cardinal decision schemes. *Social Choice and Welfare* 28, 1 (2007), 163–179.
- [17] L. Ehlers, H. Peters, and T. Storcken. 2002. Strategy-Proof Probabilistic Decision Schemes for One-Dimensional Single-Peaked Preferences. *Journal of Economic Theory* 105, 2 (2002), 408–434.
- [18] P. C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489.
- [19] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.
- [20] A. Gibbard. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45, 3 (1977), 665–681.
- [21] J. C. Heckelman. 2003. Probabilistic Borda rule voting. *Social Choice and Welfare* 21 (2003), 455–468.
- [22] J. C. Heckelman and F. H. Chen. 2013. Strategy Proof Scoring Rule Lotteries for Multiple Winners. *Journal of Public Economic Theory* 15, 1 (2013), 103–123.
- [23] A. Hylland. 1980. Strategyproofness of Voting Procedures with Lotteries as Outcomes and Infinite Sets of Strategies. (1980). Mimeo.
- [24] S. Nandeibam. 1997. An alternative proof of Gibbard’s random dictatorship result. *Social Choice and Welfare* 15, 4 (1997), 509–519.
- [25] S. Nandeibam. 2013. The structure of decision schemes with cardinal preferences. *Review of Economic Design* 17, 3 (2013), 205–238.
- [26] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- [27] A. D. Procaccia. 2010. Can approximation circumvent Gibbard-Satterthwaite?. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*. 836–841.
- [28] M. A. Satterthwaite. 1975. Strategy-Proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.
- [29] A. Sen. 2011. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIEs* 2, 4 (2011), 515–527.
- [30] Y. Shoham and K. Leyton-Brown. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- [31] Y. Tanaka. 2003. An alternative proof of Gibbard’s random dictatorship theorem. *Review of Economic Design* 8 (2003), 319–328.

SUMMARY

One of the central economic paradigms in multi-agent systems is that agents should not be better off by acting dishonestly. In the context of collective decision-making, this axiom is known as strategyproofness and turns out to be rather prohibitive, even when allowing for randomization. In particular, Gibbard's random dictatorship theorem shows that random dictatorships are the only SDSs that satisfy non-imposition and strong \succ^{SD} -strategyproofness (which is subsequently only called strategyproofness). In this paper, we interpret this result as an impossibility theorem and thus try to find more attractive strategyproof SDSs.

To this end, we consider strategyproof SDSs on restricted domains. In particular, we investigate strategyproof SDSs on the Condorcet domain which consists of all preference profiles that admit a Condorcet winner. For this domain, the Condorcet rule, which always picks the Condorcet winner with probability 1, is an appealing SDS that satisfies strategyproofness. As our first result, we demonstrate that if the number of voters n is odd, every strategyproof and non-imposing SDS on the Condorcet domain can be represented as a mixture of a random dictatorship and the Condorcet rule. Furthermore, we also show that the Condorcet domain essentially is a maximal domain that allows for attractive strategyproof social choice when n is odd as only random dictatorships are strategyproof and non-imposing on connected supersets of the Condorcet domain.

By contrast, if the number of voters n is even, the Condorcet domain is no longer a maximal domain that allows for attractive strategyproof SDSs as a single voter cannot change the Condorcet winner. We thus introduce the tie-breaking Condorcet domain, which consists of all profiles that have a Condorcet winner after adding an additional voter with a fixed preference relation. For this domain, the tie-breaking Condorcet rule, which picks the Condorcet winner in the profile with the fixed extra voter with probability 1, is strategyproof. As our second result, we then show that if n is even, an SDS on a tie-breaking-Condorcet domain is strategyproof and non-imposing if and only if it is a mixture of a random dictatorship and the corresponding tie-breaking Condorcet rule. We also prove that these domains are essentially maximal domains that allow for attractive strategyproof SDSs.

Finally, we also investigate SDSs with respect to group-strategyproofness and characterize the set of group-strategyproof and non-imposing SDSs on the Condorcet domain and most of its supersets.

REFERENCE

F. Brandt, P. Lederer, and S. Tausch. **Strategyproof social decision schemes on super Condorcet domains.** In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AA-MAS)*, pages 1734–1742, 2023.

DOI: <https://dl.acm.org/doi/10.5555/3545946.3598832>

INDIVIDUAL CONTRIBUTION

I, Patrick Lederer, am the main author of this publication. In particular, I am responsible for the development and conceptual design of the research project, the proofs of most results and the write-up of the manuscript. In more detail, only the proof of Theorem 1 was jointly derived with Sascha Tausch.

COPYRIGHT AGREEMENT

The right to present this paper in a doctoral thesis has been granted by the publisher, the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), in the copyright form presented below. There, IFAAMAS grants permission to “*personal reuse of all or portions of the above article/paper in other works of their own authorship.*” This form can also be found at <https://www.ifaamas.org/AAMAS/aamas07/IFAAMASCopyrightForm.pdf> (accessed August 24, 2023).

TERMINOLOGY

In this paper, strategyproofness refers to strong \succ^{SD} -strategyproofness and preferences are always strict.

**INTERNATIONAL FOUNDATION FOR AUTONOMOUS AGENTS AND
MULTIAGENT SYSTEMS (IFAAMAS)
COPYRIGHT FORM**

Title of Article/Paper: _____

Publication in Which Article Is to Appear: _____

Author's Name(s): _____

Please type or print your name as you wish it to appear in print
(Please read and sign Part A only, unless you are a government employee and created your article/paper as part of your employment. If your work was performed under Government contract, but you are not a Government employee, sign Part A and see item 5 under returned rights.)

PART A – Copyright Transfer Form

The undersigned, desiring to publish the above article/paper in a publication of the International Foundation for Autonomous Agents and MultiAgent Systems (IFAAMAS), hereby transfer their copyrights in the above article/paper to the International Foundation for Autonomous Agents and MultiAgent Systems (IFAAMAS), in order to deal with future requests for reprints, translations, anthologies, reproductions, excerpts, and other publications.

This grant will include, without limitation, the entire copyright in the article/paper in all countries of the world, including all renewals, extensions, and reversions thereof, whether such rights current exist or hereafter come into effect, and also the exclusive right to create electronic versions of the article/paper, to the extent that such right is not subsumed under copyright.

The undersigned warrants that they are the sole author and owner of the copyright in the above article/paper, except for those portions shown to be in quotations; that the article/paper is original throughout; and that the undersigned right to make the grants set forth above is complete and unencumbered.

If anyone brings any claim or action alleging facts that, if true, constitute a breach of any of the foregoing warranties, the undersigned will hold harmless and indemnify IFAAMAS, their grantees, their licensees, and their distributors against any liability, whether under judgment, decree, or compromise, and any legal fees and expenses arising out of that claim or actions, and the undersigned will cooperate fully in any defense IFAAMAS may make to such claim or action. Moreover, the undersigned agrees to cooperate in any claim or other action seeking to protect or enforce any right the undersigned has granted to IFAAMAS in the article/paper. If any such claim or action fails because of facts that constitute a breach of any of the foregoing warranties, the undersigned agrees to reimburse whomever brings such claim or action for expenses and attorneys' fees incurred therein.

Returned Rights

In return for these rights, IFAAMAS hereby grants to the above authors, and the employers for whom the work was performed, royalty-free permission to:

1. Retain all proprietary rights other than copyright (such as patent rights).
2. Personal reuse of all or portions of the above article/paper in other works of their own authorship.
3. Reproduce, or have reproduced, the above article/paper for the author's personal use, or for company use provided that IFAAMAS copyright and the source are indicated, and that the

copies are not used in a way that implies IFAAMAS endorsement of a product or service of an employer, and that the copies per se are not offered for sale. The foregoing right shall not permit the posting of the article/paper in electronic or digital form on any computer network, except by the author or the author's employer, and then only on the author's or the employer's own web page or ftp site. Such web page or ftp site, in addition to the aforementioned requirements of this Paragraph, must provide an electronic reference or link back to the IFAAMAS electronic server, and shall not post other IFAAMAS copyrighted materials not of the author's or the employer's creation (including tables of contents with links to other papers) without IFAAMAS's written permission.

4. Make limited distribution of all or portions of the above article/paper prior to publication.

5. In the case of work performed under U.S. Government contract, IFAAMAS grants the U.S. Government royalty-free permission to reproduce all or portions of the above article/paper, and to authorize others to do so, for U.S. Government purposes.

In the event the above article/paper is not accepted and published by IFAAMAS, or is withdrawn by the author(s) before acceptance by IFAAMAS, this agreement becomes null and void.

Author's Signature

Date

Employer for whom work was performed

Title (if not author)

PLEASE FAX THIS SIGNED COPYRIGHT FORM TO:

Jay Modi

FAX: 215-895-0545

Department of Computer Science

Drexel University

Philadelphia, PA, 19104

INTERNATIONAL FOUNDATION FOR AUTONOMOUS AGENTS AND
MULTIAGENT SYSTEMS

c/o Professor Ed Durfee

University of Michigan

EECS Department

2260 Hayward Street

Ann Arbor, MI 48109 USA

PART B – U.S. Government Employee Certification

This will certify that all authors of the above article/paper are employees of the U.S. Government and performed this work as part of their employment, and that the article/paper is therefore not subject to U.S. copyright protection. The undersigned warrants that they are the sole author/translator of the above article/paper, and that the article/paper is original throughout, except for those portions shown to be in quotations.

U.S. Government Employee Authorized Signature

Date

Name of Government Organization

Title (if not author)

Strategyproof Social Decision Schemes on Super Condorcet Domains

Felix Brandt
Technical University of Munich
Munich, Germany
brandtf@in.tum.de

Patrick Lederer
Technical University of Munich
Munich, Germany
ledererp@in.tum.de

Sascha Tausch
Technical University of Munich
Munich, Germany
sascha.tausch@tum.de

ABSTRACT

One of the central economic paradigms in multi-agent systems is that agents should not be better off by acting dishonestly. In the context of collective decision-making, this axiom is known as strategyproofness and turns out to be rather prohibitive, even when allowing for randomization. In particular, Gibbard’s random dictatorship theorem shows that only rather unattractive social decision schemes (SDSs) satisfy strategyproofness on the full domain of preferences. In this paper, we obtain more positive results by investigating strategyproof SDSs on the Condorcet domain, which consists of all preference profiles that admit a Condorcet winner. In more detail, we show that, if the number of voters n is odd, every strategyproof and non-imposing SDS on the Condorcet domain can be represented as a mixture of dictatorial SDSs and the Condorcet rule (which chooses the Condorcet winner with probability 1). Moreover, we prove that the Condorcet domain is a maximal connected domain that allows for attractive strategyproof SDSs if n is odd as only random dictatorships are strategyproof and non-imposing on any sufficiently connected superset of it. We also derive analogous results for even n by slightly extending the Condorcet domain. Finally, we also characterize the set of group-strategyproof and non-imposing SDSs on the Condorcet domain and its supersets. These characterizations strengthen Gibbard’s random dictatorship theorem and establish that the Condorcet domain is essentially a maximal domain that allows for attractive strategyproof SDSs.

KEYWORDS

Randomized Social Choice; Strategyproofness; Domain Restriction; Condorcet Domain

ACM Reference Format:

Felix Brandt, Patrick Lederer, and Sascha Tausch. 2023. Strategyproof Social Decision Schemes on Super Condorcet Domains. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Strategyproofness—no agent should be better off by acting dishonestly—is one of the central economic paradigms in multi-agent systems [6, 31, 47]. An important challenge for such systems is the identification of socially desirable outcomes by letting the agents cast votes that represent their preferences over the possible alternatives. A multitude of theorems in economic theory have shown that even rather basic properties of voting rules cannot be satisfied

simultaneously. In this context, strategyproofness is known to be a particularly restrictive axiom. This is exemplified by the Gibbard-Satterthwaite theorem which states that dictatorships are the only deterministic voting rules that satisfy strategyproofness and non-imposition (i.e., every alternative is elected for some preference profile). Since dictatorships are not acceptable for most applications, this result is commonly considered an impossibility theorem.

One of the most successful escape routes from the Gibbard-Satterthwaite impossibility is to restrict the domain of feasible preference profiles. For instance, Moulin [29] prominently showed that there are attractive strategyproof voting rules on the domain of single-peaked preference profiles, and various other restricted domains of preferences have been considered since then [e.g., 4, 21, 32, 43]. The idea behind domain restrictions is that the voters’ preferences often obey structural constraints and thus, not all preference profiles are likely or plausible. A particularly significant constraint is the existence of a Condorcet winner which is an alternative that is favored to every other alternative by a majority of the voters. Apart from its natural appeal, this concept is important because there is strong empirical evidence that real-world elections usually admit Condorcet winners [23, 27, 40]. This motivates the study of the Condorcet domain which consists precisely of the preference profiles that admit a Condorcet winner. Note that the Condorcet domain is a superset of several important domains such as those of single-peaked and single-dipped preferences when the number of voters is odd. There are several results showing the existence of attractive strategyproof voting rules on the Condorcet domain. In particular, Campbell and Kelly [10] characterize the Condorcet rule, which always picks the Condorcet winner, as the only strategyproof, non-imposing, and non-dictatorial voting rule on the Condorcet domain if the number of voters is odd.

In this paper, we focus on *randomized* voting rules, so-called social decision schemes (SDSs). Gibbard [25] has shown that randomization unfortunately does not allow for much more leeway beyond the negative consequences of the Gibbard-Satterthwaite theorem: random dictatorships, which select each voter with a fixed probability and elect the favorite alternative of the chosen voter, are the only SDSs on the full domain that satisfy strategyproofness and non-imposition (which in the randomized setting requires that every alternative is chosen with probability 1 for some preference profile). Thus, these SDSs are merely “mixtures of dictatorships”.

In order to circumvent this negative result, we are interested in large domains that allow for strategyproof and non-imposing SDSs apart from random dictatorships. A natural candidate for this is the Condorcet domain and, indeed, we show that the Condorcet domain is essentially a maximal domain that allows for strategyproof, non-imposing, and “non-randomly dictatorial” social choice. In

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

more detail, we prove that, if the number of voters n is odd, every strategyproof and non-imposing SDS on the Condorcet domain can be represented as a mixture of dictatorial SDSs and the Condorcet rule (which chooses the Condorcet winner with probability 1). This result entails that the Condorcet rule is the only strategyproof, non-imposing, and completely “non-randomly dictatorial” SDS on the Condorcet domain for odd n . Moreover, we show that, if n is odd, the Condorcet domain is a maximal domain that allows for strategyproof and non-imposing SDSs other than random dictatorships. This theorem highlights the importance of Condorcet winners for the existence of attractive strategyproof SDSs.

Unfortunately, our results for the Condorcet domain fail if the number of voters n is even because, in this case, a single voter cannot change the Condorcet winner. For extending our results to an even number of voters, we consider tie-breaking Condorcet domains, which contain all preference profiles that have a Condorcet winner after majority ties are broken according to a fixed tie-breaking order. Tie-breaking Condorcet domains are supersets of the Condorcet domain for even n , and we derive analogous results for these domains as for the Condorcet domain: if n is even, only mixtures of random dictatorships and the tie-breaking Condorcet rule (which chooses the Condorcet winner after the majority ties have been broken) are strategyproof and non-imposing on these domains, and only random dictatorships satisfy these properties on connected supersets. Finally, we also characterize the set of group-strategyproof and non-imposing SDSs on the Condorcet domain and most of its supersets independently of the parity of n : while the Condorcet rule satisfies these axioms on the Condorcet domain, only dictatorships are able to do so on most of its superdomains.

In summary, our results demonstrate two important insights: (i) the Condorcet domain is essentially a maximal domain that allows for strategyproof, non-randomly dictatorial, and non-imposing SDSs, and (ii) the (deterministic) Condorcet rule is the most appealing strategyproof voting rule on this domain, even if we allow for randomization. Our characterizations can also be seen as attractive complements to classic negative results for the full domain, whereas our results for supersets of the (tie-breaking) Condorcet domain significantly strengthen these negative results. In particular, our theorems imply statements by Barberà [2] and Campbell and Kelly [10] as well as the Gibbard-Satterthwaite theorem [24, 46] and the random dictatorship theorem [25]. A more detailed comparison between our results to these classic theorems is given in Table 1.

2 RELATED WORK

Restricting the domain of preference profiles in order to circumvent classic impossibility theorems has a long tradition and remains an active research area to date. In particular, the existence of attractive deterministic voting rules that satisfy strategyproofness has been shown for a number of domains. Classic examples include the domains of single-peaked [29], single-dipped [3], and single-crossing [43] preference profiles. More recent positive results focus on broader but more technical domains such as the domains of multi-dimensionally single-peaked or semi single-peaked preference profiles [e.g., 4, 13, 30, 39]. On the other hand, domain restrictions are also used to strengthen impossibility results by proving them for smaller domains [e.g., 1, 26, 44]. In more recent research,

the possibility and impossibility results converge by giving precise conditions under which a domain allows for strategyproof and non-dictatorial deterministic voting rules [13, 14, 17, 42].

While similar results have also been put forward for SDSs, this setting is not as well understood. For instance, Ehlers et al. [19] have shown the existence of attractive strategyproof SDSs on the domain of single-peaked preference profiles [see also 37, 38]. The existence of strategyproof and non-imposing SDSs other than random dictatorships has also been investigated for a variety of other domains [35, 36, 41]. Following a more general approach, Chatterji et al. [15] and Chatterji and Zeng [16] identify criteria for deciding whether a domain admits such SDSs.

The strong interest in restricted domains also led to the study of many computational problems for restricted domains [e.g., 5, 9, 18, 20, 22, 33, 34]. For instance, Bredereck et al. [9] give an algorithm for recognizing whether a preference profile is single-crossing. Note that for the Condorcet domain, this problem can be solved efficiently as it is easy to verify the existence of a Condorcet winner.

Finally, observe that all aforementioned results are restricted to *Cartesian* domains, i.e., domains of the form $\mathcal{D} = \mathcal{X}^N$, where \mathcal{X} is a set of preference relations. However, the Condorcet domain is not Cartesian. In this sense, the only results directly related to ours are the ones by Campbell and Kelly [10] and their follow-up work [11, 12, 28]. These papers can be seen as predecessors of our work since they investigate strategyproof deterministic voting rules on the Condorcet domain. In particular, our results extend the results by Campbell and Kelly [10] in several ways: we allow for randomization, we explore the case of even n by slightly extending the domain, we demonstrate the boundary of the possibility results, and we analyze the consequences of group-strategyproofness.

3 PRELIMINARIES

Let $N = \{1, \dots, n\}$ denote a finite set of voters and $A = \{a, b, \dots\}$ be a finite set of m alternatives. Throughout the paper, we assume that there are $n \geq 3$ voters and $m \geq 3$ alternatives. Every voter $i \in N$ is equipped with a *preference relation* \succ_i which is a complete, transitive, and anti-symmetric binary relation on A . We define \mathcal{R} as the set of all preference relations on A . A *preference profile* $R \in \mathcal{R}^N$ consists of the preference relations of all voters $i \in N$. A domain of preference profiles \mathcal{D} is a subset of the full domain \mathcal{R}^N . When writing preference profiles, we represent preference relations as comma-separated lists and indicate the set of voters who share a preference relation directly before the preference relation. Finally, we use “...” to indicate that the missing alternatives can be ordered arbitrarily. For instance, $\{1, 2\} : a, b, c, \dots$ means that voters 1 and 2 prefer a to b to c to all remaining alternatives, which can be ordered arbitrarily. We omit the brackets for singleton sets of voters.

The main object of study in this paper are *social decision schemes* (SDSs) which are voting rules that may use randomization to determine the winner of an election. More formally, an SDS maps every preference profile R of a domain \mathcal{D} to a lottery over the alternatives that determines the winning chance of every alternative. A *lottery* p is a probability distribution over the alternatives, i.e., $p(x) \geq 0$ for all $x \in A$ and $\sum_{x \in A} p(x) = 1$. We define $\Delta(A)$ as the set of all lotteries over A . Formally, an SDS on a domain \mathcal{D} is then

a function of the type $f : \mathcal{D} \rightarrow \Delta(A)$. Hence, SDSs are a generalization of deterministic voting rules which choose an alternative with probability 1 in every preference profile. The term $f(R, x)$ denotes the probability assigned to x by the lottery $f(R)$. For every set $X \subseteq A$ and lottery p , we define $p(X) = \sum_{x \in X} p(x)$; in particular $f(R, X) = \sum_{x \in X} f(R, x)$. Finally, an SDS $f : \mathcal{D} \rightarrow \Delta(A)$ is a *mixture of SDSs* g_1, \dots, g_k if there are values $\lambda_i \geq 0$ for $i \in \{1, \dots, k\}$ such that $f(R) = \sum_{i=1}^k \lambda_i g_i(R)$ for all profiles $R \in \mathcal{D}$.

A natural desideratum for an SDS $f : \mathcal{D} \rightarrow \Delta(A)$ is *non-imposition* which requires that for every alternative $x \in A$ there is a profile $R \in \mathcal{D}$ such that $f(R, x) = 1$. A prominent strengthening of this property is *ex post* efficiency. In order to define this axiom, we say an alternative $x \in A$ Pareto-dominates another alternative $y \in A \setminus \{x\}$ in a profile R if $x \succ_i y$ for all voters $i \in N$. Then, an SDS $f : \mathcal{D} \rightarrow \Delta(A)$ is *ex post efficient* if $f(R, x) = 0$ for all alternatives $x \in A$ and profiles $R \in \mathcal{D}$ such that x is Pareto-dominated in R .

3.1 Strategyproofness & Random Dictatorships

Strategic manipulation is one of the central issues in social choice theory: voters might be better off by voting dishonestly. Since satisfactory collective decisions require the voters' true preferences, SDSs should incentivize honest voting. In order to formalize this, we need to specify how voters compare lotteries over alternatives. The most prominent approach for this is based on (first order) stochastic dominance [e.g., 19, 25, 35]. Let the *upper contour set* $U(\succ_i, x) = \{y \in A : y \succ_i x \vee y = x\}$ be the set of alternatives that voter i weakly prefers to x . Then, (first order) stochastic dominance states that a voter i prefers a lottery p to another lottery q , denoted by $p \succeq_i^{SD} q$, if $p(U(\succ_i, x)) \geq q(U(\succ_i, x))$ for all $x \in A$. Note that the stochastic dominance relation is transitive but not complete. Using stochastic dominance to compare lotteries is appealing because $p \succeq_i^{SD} q$ holds if and only if p guarantees voter i at least as much expected utility than q for every utility function that is ordinally consistent with his preference relation \succ_i .

Based on stochastic dominance, we now define strategyproofness: an SDS $f : \mathcal{D} \rightarrow \Delta(A)$ is *strategyproof* if $f(R) \succeq_i^{SD} f(R')$ for all preference profiles $R, R' \in \mathcal{D}$ and voters $i \in N$ such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$. Less formally, strategyproofness requires that every voter weakly prefers the lottery obtained by acting truthfully to every lottery he could obtain by lying. Conversely, an SDS is called *manipulable* if it is not strategyproof. A convenient property of strategyproofness is that mixtures of strategyproof SDSs are again strategyproof.

Note that this strategyproof notion has attained significant attention. In particular, Gibbard [25] has shown that only random dictatorships satisfy strategyproofness and non-imposition on the full domain. For defining these functions, we say an SDS d_i is *dictatorial* or a *dictatorship* if it always assigns probability 1 to the most preferred alternative of voter i . Then, a *random dictatorship* f is a mixture of dictatorial SDSs d_i .

Strategyproofness is closely related to two properties called localizedness and non-perversity. Both of these axioms are concerned with how the outcome changes if a voter only swaps two alternatives. For making this formal, let $R^{i:yx}$ denote the profile derived from another profile R by only swapping x and y in the preference relation of voter i . Note that this definition requires that $x \succ_i y$

and that there is no alternative $z \in A \setminus \{x, y\}$ with $x \succ_i z \succ_i y$. Now, an SDS f on a domain \mathcal{D} is *localized* if $f(R, z) = f(R^{i:yx}, z)$ for all distinct alternatives $x, y, z \in A$, voters $i \in N$, and profiles $R, R^{i:yx} \in \mathcal{D}$. Moreover, f is *non-perverse* if $f(R^{i:yx}, y) \geq f(R, y)$ for all distinct alternatives $x, y \in A$, voters $i \in N$, and profiles $R, R^{i:yx} \in \mathcal{D}$. More intuitively, if voter i reinforces y against x , localizedness requires that the probability assigned to the other alternatives does not change, and non-perversity that the probability of y cannot decrease. Gibbard [25] has shown for the full domain of preferences that the conjunction of localizedness and non-perversity is equivalent to strategyproofness. Furthermore, it is easy to see that every strategyproof SDS satisfies non-perversity and localizedness on every domain. We thus mainly use the latter two axioms in our proofs as they are easier to handle.

Finally, in order to disincentivize *groups* of voters from manipulating, we need a stronger strategyproofness notion: an SDS $f : \mathcal{D} \rightarrow \Delta(A)$ is *group-strategyproof* if for all preference profiles $R, R' \in \mathcal{D}$ and all non-empty sets of voters $I \subseteq N$ with $\succ_j = \succ'_j$ for all $j \in N \setminus I$, there is a voter $i \in I$ such that $f(R) \succeq_i^{SD} f(R')$. Conversely, an SDS is *group-manipulable* if it is not group-strategyproof. Note that group-strategyproofness implies strategyproofness.

3.2 Super Condorcet Domains

Since Gibbard's random dictatorship theorem shows that there are no attractive strategyproof SDSs on the full domain, we investigate the Condorcet domain and its supersets with respect to the existence of such functions. In order to define these domains, we first need to introduce some terminology. The *majority margin* $g_R(x, y) = |\{i \in N : x \succ_i y\}| - |\{i \in N : y \succ_i x\}|$ indicates how many more voters prefer x to y in the profile R than vice versa. Based on the majority margins, we define the *Condorcet winner* of a profile R as the alternative x such that $g_R(x, y) > 0$ for all $y \in A \setminus \{x\}$. Since the existence of a Condorcet winner is not guaranteed, we focus on the *Condorcet domain* $\mathcal{D}_C = \{R \in \mathcal{R}^N : \text{there is a Condorcet winner in } R\}$ which contains all profiles with a Condorcet winner for the given electorate. As explained in the introduction, this domain is of interest because real-world elections frequently admit Condorcet winners.

A particularly natural SDS on the Condorcet domain is the *Condorcet rule* (COND) which always assigns probability 1 to the Condorcet winner. However, all SDSs defined for the full domain (e.g., random dictatorships, Borda's rule, Plurality rule) are also well-defined for the Condorcet domain and there is thus a multitude of voting rules to choose from.

Note that the Condorcet domain \mathcal{D}_C is not connected with respect to strategyproofness if n is even. To make this formal, we define \mathcal{D}_C^x as the domain of profiles in which alternative x is the Condorcet winner. Then, it is impossible for distinct alternatives $x, y \in A$ that a single voter deviates from a profile $R \in \mathcal{D}_C^x$ to a profile $R' \in \mathcal{D}_C^y$. Indeed, if $R \in \mathcal{D}_C^x$ and n is even, then $g_R(x, y) \geq 2$ and $g_{R'}(x, y) \geq 0$ for all alternatives $y \in A \setminus \{x\}$ and profiles R' that differ from R only in the preference relation of a single voter. This is problematic for our analysis because the choice for a profile $R \in \mathcal{D}_C^x$ has no influence of the choice for a profile $R' \in \mathcal{D}_C^y$. For even n , we will thus consider the *tie-breaking Condorcet domain* $\mathcal{D}_C^\triangleright = \{R \in \mathcal{R}^N : \text{there is a Condorcet winner in } (R, \triangleright)\}$, which contains all profiles that have a Condorcet winner after adding a

fixed preference relation $\triangleright \in \mathcal{R}$. Note that this extra preference relation only breaks majority ties if n is even because $|g_R(x, y)| \geq 2$ if $g_R(x, y) \neq 0$. In particular, this proves that $\mathcal{D}_C \subseteq \mathcal{D}_C^\triangleright$ for even n . An attractive SDS on $\mathcal{D}_C^\triangleright$ is the *tie-breaking Condorcet rule* ($\text{COND}^\triangleright$) which assigns probability 1 to the Condorcet winner in (R, \triangleright) for all profiles $R \in \mathcal{D}_C^\triangleright$.

To show that \mathcal{D}_C and $\mathcal{D}_C^\triangleright$ are maximal domains that allow for attractive strategyproof SDSs, we will also consider supersets of them. Formally, we analyze *super Condorcet domains* which are domains \mathcal{D} with $\mathcal{D}_C \subseteq \mathcal{D}$. Just as the Condorcet domain for even n , super Condorcet domains can be disconnected. We therefore discuss connectedness notions for domains and introduce ad-paths. An *ad-path* from a profile R to a profile R' in a domain \mathcal{D} is a sequence of profiles (R^1, \dots, R^l) such that $R^1 = R$, $R^l = R'$, $R^k \in \mathcal{D}$ for all $k \in \{1, \dots, l\}$, and the profile R^{k+1} evolves out of R^k by swapping two alternatives $x, y \in A$ in the preference relation of a voter $i \in N$, i.e., $R^{k+1} = (R^k)^{i:yx}$ for all $k \in \{1, \dots, l-1\}$. Then, we say that a domain \mathcal{D} is *weakly connected* if there is an ad-path between all profiles $R, R' \in \mathcal{D}$. Unfortunately, this condition is too weak to be useful in our analysis and we therefore slightly strengthen it: a domain \mathcal{D} is *connected* if it is weakly connected and if for all alternatives $x \in A$ and profiles $R, R' \in \mathcal{D}$ such that $U(\succ'_i, x) = U(\succ_i, x)$ for all $i \in N$, there is an ad-path (R^1, \dots, R^l) from R to R' such that $U(\succ'_i, x) = U(\succ_i, x)$ for all $i \in N$ and $k \in \{1, \dots, l-1\}$. Less formally, connectedness strengthens weak connectedness by requiring that if an alternative x is at the same position in R and R' , then we can go from R to R' without moving x .

Connectedness is a very mild property and is, e.g., weaker than Sato's non-restoration property [45]. Hence, many domains, such as the full domain and the single-peaked domain, satisfy this condition. As we show next, the same holds for the Condorcet domain if n is odd and for tie-breaking Condorcet domains if n is even.

Lemma 1. *If $n \geq 3$ is odd, the Condorcet domain \mathcal{D}_C is connected. If $n \geq 4$ is even, the tie-breaking Condorcet domain $\mathcal{D}_C^\triangleright$ is connected for every preference relation $\triangleright \in \mathcal{R}$.*

Proof sketch. The proof for \mathcal{D}_C and $\mathcal{D}_C^\triangleright$ work essentially the same, and we thus focus on \mathcal{D}_C in this proof sketch. Hence, assume that $n \geq 3$ is odd, consider two profiles $R, R' \in \mathcal{D}_C$, and let c and c' be the respective Condorcet winners. We first show that \mathcal{D}_C is weakly connected and thus need to construct an ad-path from R to R' . If $c = c'$, we start at R by reinforcing c until it unanimously top-ranked, reorder the other alternatives according to R' , and weaken c until it is in the correct position. If $c \neq c'$, we can proceed similarly: starting at R , we let all voters first push up c until it is their best alternative, and then let all voters push up c' until it is their second best alternative. We can now change the Condorcet winner without leaving the Condorcet domain by letting the voters swap c and c' one after another. After this, c' is the Condorcet winner and we can now apply the same construction as for the case $c = c'$ to go from this intermediate profile to R' . For showing that \mathcal{D}_C is connected, we also need to construct ad-paths from R to R' that do not move x for every alternative $x \in A$ with $U(\succ_i, x) = U(\succ'_i, x)$ for all $i \in N$. The construction of these ad-paths relies on a tedious case distinction with respect to whether $x = c$ and $c = c'$, so we defer it to a full version of this paper [8]. \square

4 RESULTS

We are now ready to present our characterizations of strategyproof and group-strategyproof SDSs on super Condorcet domains. In more detail, we first characterize the set of strategyproof and non-imposing SDSs on the Condorcet domain for an odd number of voters n in Section 4.1. Moreover, we also demonstrate that the Condorcet domain is a maximal connected domain that allows for strategyproof and non-imposing SDSs apart from random dictatorships. Next, we derive analogous results for the tie-breaking Condorcet domain if n is even in Section 4.2. Finally, in Section 4.3, we revisit the Condorcet domain and characterize the set of group-strategyproof and non-imposing SDSs, independently of the parity of n . Due to space restrictions, we defer the proofs of Lemma 3 and Theorems 2 and 3 to a full version of this paper [8].

4.1 Condorcet Domain

In this section, we analyze the set of strategyproof and non-imposing SDSs on the Condorcet domain and its supersets for the case that n is odd. In more detail, we will show that, if n is odd, only mixtures of random dictatorships and the Condorcet rule are strategyproof and non-imposing on the Condorcet domain. As a byproduct, we also derive a characterization of the Condorcet rule as the only strategyproof, non-imposing, and “completely non-randomly dictatorial” SDS on \mathcal{D}_C . Moreover, we will also prove that, if n is odd, only random dictatorships are strategyproof and non-imposing on every connected superset of the Condorcet domain, thus demonstrating that the Condorcet domain is an inclusion-maximal connected domain that allows for attractive strategyproof SDSs.

Before proving these claims, we discuss two auxiliary lemmas. First, we show that, if n is odd, every strategyproof and non-imposing SDS on a connected super Condorcet domain is also *ex post* efficient. Analogous claims are known for, e.g., the full domain and the domain of single-peaked preferences [19, 25].

Lemma 2. *Assume $n \geq 3$ is odd, and let $\mathcal{D} \subseteq \mathcal{R}^N$ denote a connected domain with $\mathcal{D}_C \subseteq \mathcal{D}$. Every strategyproof and non-imposing SDS on \mathcal{D} is *ex post* efficient.*

Proof. Assume $n \geq 3$ is odd and let \mathcal{D} denote a connected domain with $\mathcal{D}_C \subseteq \mathcal{D}$. Moreover, consider a strategyproof and non-imposing SDS f on \mathcal{D} and assume for contradiction that f fails *ex post* efficiency. This means that there are a profile $R^1 \in \mathcal{D}$ and two alternatives $x, y \in A$ such that $x \succ_i y$ for all voters $i \in N$ but $f(R^1, y) > 0$. Now, consider the profile R^2 derived from R^1 by making x into the best alternative of every voter $i \in N$. Clearly, $R^2 \in \mathcal{D}_C \subseteq \mathcal{D}$ because x is the Condorcet winner in R^2 . Since $U(\succ_i^1, y) = U(\succ_i^2, y)$ for all voters $i \in N$, there is by connectedness an ad-path from R^1 to R^2 along which y is never swapped. Hence, we infer that $f(R^2, y) = f(R^1, y) > 0$ and $f(R^2, x) < 1$ by repeatedly applying localizedness along this ad-path.

Next, let $R^3 \in \mathcal{D}$ denote a profile such that $f(R^3, x) = 1$; such a profile exists by non-imposition. If x is the Condorcet winner in R^3 , we can reinforce this alternative until it is top-ranked by every voter without leaving the domain \mathcal{D} . This leads to a profile R^4 in which x is unanimously top-ranked, and non-perversity shows that $f(R^4, x) \geq f(R^3, x) = 1$. Finally, we can again use the connectedness of \mathcal{D} to find an ad-path from R^4 to R^2 along which x is never

swapped. Hence, localizedness requires that $f(R^2, x) = f(R^4, x) = 1$, which contradicts our previous observation.

As second case, suppose that x is not the Condorcet winner in R^3 . Since n is odd, there is an alternative $z \in A \setminus \{x\}$ and a set of voters I with $|I| > \frac{n}{2}$ such that $z \succ_i^3 x$ for all $i \in I$. Now, consider the profile R^5 derived from R^3 by making z into the best alternative of the voters $i \in I$. Clearly, $R^5 \in \mathcal{D}_C \subseteq \mathcal{D}$ because more than half of the voters top-rank z . Moreover, it holds that $U(\succ_i^3, x) = U(\succ_i^5, x)$ for all $i \in N$, and thus connectedness and localizedness imply that $f(R^5, x) = 1$. Next, let R^6 denote the profile derived from R^5 by making x into the best alternative of the voters $i \in N \setminus I$ and into the second best one of the voters $i \in I$. We can transform R^5 into R^6 by repeatedly reinforcing x , and z stays always the Condorcet winner as it is top-ranked by the voters $i \in I$. Hence, $R^6 \in \mathcal{D}$ and non-perversity shows that $f(R^6, x) = 1$. Finally, we let the voters $i \in I$ swap x and z one after another. Since all voters top-ranks x or z in R^6 , one of these alternatives is always top-ranked by more than half of the voters during these steps. Hence, we do not leave \mathcal{D} . This process terminates in a profile R^7 in which all voters top-rank x , and non-perversity shows that $f(R^7, x) = 1$. This contradicts again that $f(R^2, x) < 1$ as there is an ad-path from R^7 to R^2 along which we do not move x . Since we have a contradiction in both cases, f must be *ex post* efficient. \square

Lemma 2 is helpful for our analysis because *ex post* efficiency—in contrast to non-imposition—is inherited to subdomains. Since an analogous claim also holds for strategyproofness, we next investigate the set of strategyproof and *ex post* efficient SDSs in the domain \mathcal{D}_C^x where alternative x is always the Condorcet winner.

Lemma 3. Fix an alternative $a \in A$ and let f denote a strategyproof and *ex post* efficient SDS on a super Condorcet domain. There is a random dictatorship d and $\gamma \in \mathbb{R}_{\geq 0}$ such that $f(R) = (1 - \gamma)\text{COND}(R) + \gamma d(R)$ for all $R \in \mathcal{D}_C^a$.

Proof sketch. Consider an arbitrary super Condorcet domain \mathcal{D} , a strategyproof and *ex post* efficient SDS f on \mathcal{D} , and fix an alternative $a \in A$. For proving this lemma, we will investigate the behavior of f on several subdomains of \mathcal{D}_C^a . In particular, we first fix a set of voters $I \subseteq N$ with $|I| = \lceil \frac{n+1}{2} \rceil$ and a profile $\hat{R} \in \mathcal{R}^I$ in which all voters $i \in I$ report a as their favorite alternative. Then, we consider the domain $\mathcal{D}_1^{I, \hat{R}}$ of profiles in which the voters $i \in I$ report \hat{r}_i and the voters $i \in N \setminus I$ report arbitrary preference relations. In particular, we show that f induces an SDS $g_{\hat{R}}$ on the domain $\mathcal{R}^{N \setminus I}$ that is non-imposing and strategyproof. The random dictatorship theorem therefore shows that $g_{\hat{R}}$ is a random dictatorship. By using the relation between $g_{\hat{R}}$ and f , we then derive that there are values $\gamma_C \geq 0$ and $\gamma_i \geq 0$ for $i \in N \setminus I$ such that $f(R) = \gamma_C^{\text{COND}}(R) + \sum_{i \in N \setminus I} \gamma_i d_i(R)$ for all profiles $R \in \mathcal{D}_1^{I, \hat{R}}$. For proving the lemma from this point on, we repeatedly enlarge the domain $\mathcal{D}_1^{I, \hat{R}}$ and show that f can always be represented as a mixture of a random dictatorship and the Condorcet rule. For instance, we consider next the domain \mathcal{D}_2^I , where the voters in I have to top-rank a . Clearly, every profile $R \in \mathcal{D}_2^I$ is in $\mathcal{D}_1^{I, R'}$ for a profile $R' \in \mathcal{R}^I$. Since f can be represented for every domain $\mathcal{D}_1^{I, \hat{R}}$ as a mixture of a random dictatorship and the Condorcet rule, we hence derive an analogous

claim for \mathcal{D}_2^I by showing that it is always the same mixture. By further generalizing the domain like this, we eventually derive the lemma. \square

Lemma 3 is itself already a rather strong statement as it characterizes the behavior of all strategyproof and *ex post* efficient SDSs f on the domains \mathcal{D}_C^x . In particular, this result does neither require that n is odd nor a connectedness condition on the domain. On the other hand, Lemma 3 does not relate the behavior of f for different subdomains \mathcal{D}_C^x , and it might be that the weight on the Condorcet rule is negative. Indeed, if n is even and $m = 3$, the SDS $f(R) = \sum_{i \in N} \frac{1}{n-1} d_i(R) - \frac{1}{n-1} \text{COND}(R)$ is well-defined, non-imposing, and strategyproof for the Condorcet domain because the Condorcet winner is top-ranked by at least one voter if $m = 3$.

Nevertheless, Lemma 3 is the central tool for proving all of our theorems and we will use it next to characterize the set of strategyproof and non-imposing SDSs on the Condorcet domain and all of its connected supersets for the case that n is odd.

Theorem 1. Assume $n \geq 3$ is odd and let $\mathcal{D} \subseteq \mathcal{R}^N$ denote a connected domain. The following claims are true.

- (1) Assume $\mathcal{D}_C = \mathcal{D}$. An SDS on \mathcal{D} is strategyproof and non-imposing if and only if it is a mixture of a random dictatorship and the Condorcet rule.
- (2) Assume $\mathcal{D}_C \subsetneq \mathcal{D}$. An SDS on \mathcal{D} is strategyproof and non-imposing if and only if it is a random dictatorship.

Proof. Assume $n \geq 3$ is odd and let \mathcal{D} denote a connected domain with $\mathcal{D}_C \subseteq \mathcal{D}$.

Proof of Claim (1): First, we assume that $\mathcal{D} = \mathcal{D}_C$ and consider an SDS f on \mathcal{D} that is a mixture of a random dictatorship and the Condorcet rule. Since mixtures of strategyproof SDSs are themselves strategyproof and the Condorcet rule as well as all random dictatorships are known to satisfy this axiom on \mathcal{D}_C , it follows immediately that f is strategyproof. Moreover, all random dictatorships and the Condorcet rule choose an alternative with probability 1 if it is unanimously top-ranked. Since all these profile are in \mathcal{D}_C , we derive that f is also non-imposing.

For the other direction, assume that f is a strategyproof and non-imposing SDS on \mathcal{D}_C . Since the Condorcet domain is connected if n is odd (see Lemma 1), we derive from Lemma 2 that f is *ex post* efficient. In turn, Lemma 3 shows that for every alternative $x \in A$, there are values γ_C^x and $\gamma_i^x \geq 0$ for all $i \in N$ such that $f(R) = \gamma_C^x \text{COND}(R) + \sum_{i \in N} \gamma_i^x d_i(R)$ for all $R \in \mathcal{D}_C^x$. Hence, the theorem follows by showing that $\gamma_C^x = \gamma_C^y$ and $\gamma_i^x = \gamma_i^y$ for all $x, y \in A$ and all $i \in N$, and that $\gamma_C^x \geq 0$. First, we show that $\gamma_i^x = \gamma_i^y$ and $\gamma_C^x = \gamma_C^y$ for all $x, y \in A$. For doing so, consider three alternatives $a, b, c \in A$, a voter $i \in N$, and the profiles R^1 and R^2 .

$$\begin{aligned} R^1: & \quad i: c, a, b, \dots & N \setminus \{i\}: a, b, c, \dots \\ R^2: & \quad i: c, a, b, \dots & N \setminus \{i\}: b, a, c, \dots \end{aligned}$$

Clearly, $R^1 \in \mathcal{D}_C^a$ and $R^2 \in \mathcal{D}_C^b$ and thus, $f(R^1, c) = \gamma_i^a$ and $f(R^2, c) = \gamma_i^b$. Furthermore, since \mathcal{D}_C is connected and $U(\succ_i^1, c) = U(\succ_i^2, c)$ for all $i \in N$, there is an ad-path from R^1 to R^2 along which c is never swapped. Localizedness implies therefore that $f(R^1, c) = f(R^2, c)$ and hence, $\gamma_i^a = \gamma_i^b$. Because voter i is chosen arbitrarily, this holds for all voters $i \in N$ and we infer that $\gamma_C^a =$

$1 - \sum_{i \in N} \gamma_i^a = 1 - \sum_{i \in N} \gamma_i^b = \gamma_C^b$. This means that $\gamma_C^x = \gamma_C^y$ and $\gamma_i^x = \gamma_i^y$ for all voters $i \in N$ and alternatives $x, y \in A$.

Next, we will show that $\gamma_C^a \geq 0$ for some $a \in A$. For this step, we partition the set of voters in three disjoint subsets I_1, I_2 , and I_3 such that $|I^1| = |I^2| = \frac{n-1}{2}$ and $|I^3| = 1$. Now, let $b, c \in A \setminus \{a\}$ denote distinct alternatives and consider the profiles R^3 and R^4 .

$$\begin{array}{lll} R^3: & I^1: a, b, \dots & I^2: b, a, \dots & I^3: c, a, b, \dots \\ R^4: & I^1: a, b, \dots & I^2: b, a, \dots & I^3: c, b, a, \dots \end{array}$$

Alternative a is the Condorcet winner in R^3 and alternative b in R^4 . Hence, $f(R^3, a) = \sum_{i \in I^1} \gamma_i^a + \gamma_C^a$ and $f(R^4, a) = \sum_{i \in I^1} \gamma_i^b$. Next, non-perversity shows that $f(R^3, a) \geq f(R^4, a)$. Since $\gamma_i^a = \gamma_i^b$ for all $i \in N$, we thus infer that $\gamma_C^a \geq 0$. Now, by defining $\gamma_C = \gamma_C^a$ and $\gamma_i = \gamma_i^a$ for all $i \in N$ and some $a \in A$, we conclude that $f(R) = \gamma_C \text{COND}(R) + \sum_{i \in N} \gamma_i d_i(R)$ for all $R \in \mathcal{D}_C$.

Proof of Claim (2): For the second claim, we assume that $\mathcal{D}_C \subsetneq \mathcal{D}$. Since it is straightforward to see that random dictatorships are strategyproof and non-imposing on \mathcal{D} , we focus on the converse. For this, let f denote a strategyproof and non-imposing SDS on \mathcal{D} . By Lemma 2, f is also *ex post* efficient. As a consequence, it is non-imposing in the Condorcet domain, and we thus infer from Claim (1) that there are $\gamma_C \geq 0$ and $\gamma_i \geq 0$ for all $i \in N$ such that $f(R) = \gamma_C \text{COND}(R) + \sum_{i \in N} \gamma_i d_i(R)$ for all $R \in \mathcal{D}_C$. Now, for proving Claim (2), consider a profile $R \in \mathcal{D} \setminus \mathcal{D}_C$ and let x denote an arbitrary alternative. Since n is odd and there is no Condorcet winner in R , there is a set of voters I and an alternative $y \in A \setminus \{x\}$ such that $|I| > \frac{n}{2}$ and $y \succ_i x$ for all $i \in I$. Next, we consider the profile R' derived from R by letting all voters $i \in I$ make y into their favorite alternative. Clearly, y is the Condorcet winner in R' and thus $f(R', x) = \sum_{i \in N} \gamma_i d_i(R', x)$. On the other hand, $U(\succ_i, x) = U(\succ'_i, x)$ for all $i \in N$ because the voters in I prefer y to x in R . Hence, we can apply connectedness and localizedness to derive that $f(R, x) = f(R', x) = \sum_{i \in N} \gamma_i d_i(R, x)$. Since x is chosen arbitrarily, this means that $f(R) = \sum_{i \in N} \gamma_i d_i(R)$. In particular, it must hold that $\sum_{i \in N} \gamma_i = 1$ and thus $\gamma_C = 0$ as otherwise $\sum_{x \in A} f(R, x) < 1$. This proves $f(R) = \sum_{i \in N} \gamma_i d_i(R)$ for all $R \in \mathcal{D}$. \square

Claim (1) of Theorem 1 immediately implies that the Condorcet rule is the only “completely non-randomly dictatorial” SDS on the Condorcet domain that satisfies strategyproofness and non-imposition. To formalize this observation, we introduce the notion of γ -randomly dictatorial SDSs first suggested by Brandt et al. [7]: a strategyproof SDS f on a domain \mathcal{D} is γ -randomly dictatorial if $\gamma \in [0, 1]$ is the maximal value such that f can be represented as $f(R) = \gamma d(R) + (1 - \gamma)g(R)$ for all profiles $R \in \mathcal{D}$, where d is a random dictatorship and g is another strategyproof SDS on \mathcal{D} . It follows immediately from Theorem 1 that, if n is odd, the Condorcet rule is the only 0-randomly dictatorial, strategyproof, and non-imposing SDS on the Condorcet domain. This corollary generalizes Theorem 1 of Campbell and Kelly [10] who have characterized the Condorcet rule with equivalent axioms in the deterministic setting. Furthermore, this insight highlights the appeal of the Condorcet rule on the Condorcet domain because every other strategyproof and non-imposing SDS is a mixture of the Condorcet rule and a random dictatorship.

On the other hand, Claim (2) of Theorem 1 generalizes the random dictatorship theorem from the full domain to all connected

supersets of \mathcal{D}_C if n is odd. Since deterministic voting rules can be seen as a special case of SDSs, our result also generalizes the Gibbard-Satterthwaite theorem to these smaller domains. In particular, Claim (2) of Theorem 1 shows that adding even a single profile to the Condorcet domain can turn the positive results of Claim (1) into a negative one. This follows, for instance, by considering the domain $\mathcal{D}_1 = \mathcal{D}_C \cup \{R^*\}$. The preference profile R^* is shown below, where $I = \{4, 6, \dots, n-1\}$, $J = \{5, 7, \dots, n\}$.

$$\begin{array}{lll} R^*: & 1: a, b, c, \dots & 2: b, c, a, \dots & 3: c, a, b, \dots \\ & I: a, b, c, \dots & J: c, b, a, \dots \end{array}$$

4.2 Tie-Breaking Condorcet Domain

A natural follow-up question to Theorem 1 is to ask for the strategyproof and non-imposing SDSs on the Condorcet domain if n is even. Unfortunately, since the Condorcet domain is not connected in this case, a concise characterization of all these SDSs seems impossible. We therefore characterize the set of strategyproof and non-imposing SDS on the tie-breaking Condorcet domain $\mathcal{D}_C^\triangleright$. Moreover, the following theorem also demonstrates that, if n is even, tie-breaking Condorcet domains are inclusion-maximal connected domains that allow for strategyproof and non-imposing SDSs other than random dictatorships.

Theorem 2. Assume $n \geq 4$ is even, let $\triangleright \in \mathcal{R}$ be a preference relation, and $\mathcal{D} \subseteq \mathcal{R}^N$ be a connected domain. The following claims hold.

- (1) Assume $\mathcal{D} = \mathcal{D}_C^\triangleright$. An SDS on \mathcal{D} is strategyproof and non-imposing if and only if it is a mixture of a random dictatorship and the tie-breaking Condorcet rule $\text{COND}^\triangleright$.
- (2) Assume $\mathcal{D}_C^\triangleright \subsetneq \mathcal{D}$. An SDS on \mathcal{D} is strategyproof and non-imposing if and only if it is a random dictatorship.

Proof sketch. Assume $n \geq 4$ is even, fix a preference relation $\triangleright \in \mathcal{R}$, and consider a connected domain \mathcal{D} with $\mathcal{D}_C^\triangleright \subseteq \mathcal{D}$. First, note that random dictatorships are strategyproof and non-imposing on \mathcal{D} , regardless of whether $\mathcal{D} = \mathcal{D}_C^\triangleright$ or $\mathcal{D}_C^\triangleright \subsetneq \mathcal{D}$. Moreover, if $\mathcal{D} = \mathcal{D}_\triangleright$, then $\text{COND}^\triangleright$ is strategyproof on \mathcal{D} because every manipulation of this rule can be turned in a manipulation of the Condorcet rule for $n+1$ voters. Since mixture of strategyproof SDSs are strategyproof, it follows that all mixtures of random dictatorships and $\text{COND}^\triangleright$ are strategyproof, and it is easy to see that these rules are also non-imposing.

Next, we focus on the direction from left to right and consider for this a strategyproof and non-imposing SDS f on \mathcal{D} . Analogous to Lemma 2, it is not difficult to derive that f is *ex post* efficient on \mathcal{D} . Hence, Lemma 3 implies that there are values $\gamma_i^x \geq 0$ for all $i \in N$ and γ_C^x such that $f(R) = \gamma_C^x \text{COND}(R) + \sum_{i \in N} \gamma_i^x d_i(R)$ for all subdomains \mathcal{D}_C^x and profiles $R \in \mathcal{D}_C^x$. Next, we show analogously to the proof of Theorem 1 that $\gamma_C^x = \gamma_C^y$ and $\gamma_i^x = \gamma_i^y$ for all $x, y \in A$ and $i \in N$ and we can thus drop the superscript. Since $\text{COND}(R) = \text{COND}^\triangleright(R)$ for all $R \in \mathcal{D}_C$ if n is even, this means that $f(R) = \gamma_C \text{COND}^\triangleright(R) + \sum_{i \in N} \gamma_i d_i(R)$ for all $R \in \mathcal{D}_C$.

Now, to prove the first claim, let us assume that $\mathcal{D} = \mathcal{D}_C^\triangleright$. In this case, we first show that $f(R)$ can also be represented as $\gamma_C \text{COND}^\triangleright(R) + \sum_{i \in N} \gamma_i d_i(R)$ if there is an alternative x in R which is top-ranked by at least half of the voters and which is the Condorcet winner in (R, \triangleright) . Next, we consider a profile $R \in \mathcal{D}_C^\triangleright$ and let

x denote the Condorcet winner in (R, \succ) . This means that for every alternative $y \in A \setminus \{x\}$, there are at least $\frac{n}{2}$ voters who prefer x to y in R . If we let these voters reinforce x until it is top-ranked, we arrive at a profile R' such that $f(R') = \gamma_C \text{COND}^\succ(R') + \sum_{i \in N} \gamma_i d_i(R')$. Moreover, connectedness and localizedness imply that the probability of y does not change when going from R to R' . Since $y \in A \setminus \{x\}$ is chosen arbitrarily, we derive from this observation that $f(R) = \gamma_C \text{COND}^\succ(R) + \sum_{i \in N} \gamma_i d_i(R)$ for every profile $R \in \mathcal{D}_C^\succ$. As last step, we show that $\gamma_C \geq 0$ by using a similar argument as in the proof of Theorem 1. This completes the proof of Claim (1).

For proving Claim (2), assume that $\mathcal{D}_C^\succ \subsetneq \mathcal{D}$. By Claim (1), we infer that f can be represented as a mixture of a random dictatorship and COND^\succ for all profiles $R \in \mathcal{D}_C^\succ$. Now, consider a profile $R \in \mathcal{D} \setminus \mathcal{D}_C^\succ$. For the proof, we identify for every alternative $x \in A$ a profile $R' \in \mathcal{D}_C^\succ$ such that $U(\succ_i, x) = U(\succ'_i, x)$ for all voters $i \in N$. Once we have these profiles, the proof proceeds exactly as the proof of Claim (2) in Theorem 1. \square

First, note that Theorem 2 implies—analogously to Theorem 1—that the tie-breaking Condorcet rule is the only strategyproof, non-imposing, and 0-randomly dictatorial SDS on the tie-breaking Condorcet domain if n is even. In particular, this proves again that choosing the Condorcet winners is desirable because COND^\succ chooses the Condorcet winners whenever there is one. Moreover, since the tie-breaking Condorcet domain is only a small extension of the Condorcet domain, this result demonstrates the important role of Condorcet winners for the existence of strategyproof and non-imposing SDSs other than random dictatorships.

Furthermore, Claim (2) in Theorem 2 shows again that adding even a single profile to \mathcal{D}_C^\succ can turn the positive result into a negative one. In particular, note that this claim also implies that the domain of all profiles with a weak Condorcet winner (which is equivalent to the domain $\mathcal{D} = \{R \in \mathcal{R}^N : \exists \succ \in \mathcal{R} : \text{there is a Condorcet winner in } (R, \succ)\}$) only allows for random dictatorships as strategyproof and non-imposing SDSs.

Remark 1. An important observation of Theorems 1 and 2 is that every strategyproof and non-imposing SDS on the respective domains can be represented as a mixture of deterministic voting rules, each of which is strategyproof and non-imposing. This is sometimes called deterministic extreme point property and remarkably, many important domains satisfy this condition [41]. On the one hand, this shows that randomization does not lead to completely new strategyproof SDSs. On the other hand, the deterministic extreme point property allows for a natural interpretation of strategyproof and non-imposing SDSs: we randomly select a deterministic voting rule.

Remark 2. The connectedness condition is required for Claim (2) in Theorems 1 and 2 because there are domains \mathcal{D} with $\mathcal{D}_C \subsetneq \mathcal{D}$ (resp. $\mathcal{D}_C^\succ \subsetneq \mathcal{D}$) that allow for non-imposing and strategyproof SDSs that are no random dictatorships. For example, consider the domain \mathcal{D}_2 which is derived by adding a single preference profile R^1 to the Condorcet domain. If R^1 differs from every profile in \mathcal{D}_C in the preference relations of at least two voters, an arbitrary outcome can be returned for R^1 without violating strategyproofness.

4.3 Group-Strategyproofness

Finally, we investigate the set of group-strategyproof and non-imposing SDSs on the Condorcet domain and its supersets. In particular, we will show that only the Condorcet rule and dictatorial SDSs satisfy group-strategyproofness on the Condorcet domain. Note that this result is independent of the parity of n and group-strategyproofness thus allows for a unified characterization. Moreover, we also prove a counterpart to Claim (2) in Theorems 1 and 2, which notably does not require connectedness.

Theorem 3. Assume $n \geq 3$ and let $\mathcal{D} \subseteq \mathcal{R}^N$ denote an arbitrary domain. The following claims are true.

- (1) Assume $\mathcal{D} = \mathcal{D}_C$. An SDS on \mathcal{D} is group-strategyproof and non-imposing if and only if it is a dictatorship or the Condorcet rule.
- (2) Assume $\mathcal{D}_C \subsetneq \mathcal{D}$ and that there is a profile $R \in \mathcal{D}$ such that for each $x \in A$, there is $y \in A$ with $g_R(y, x) > 0$. An SDS on \mathcal{D} is group-strategyproof and non-imposing if and only if it is a dictatorship.

Proof sketch. For the direction from right to left of both claims, we note first that dictatorships are clearly non-imposing and group-strategyproof on every super Condorcet domain. Furthermore, it is also apparent that the Condorcet rule is non-imposing on the Condorcet domain. We hence only need to show that COND is group-strategyproof on \mathcal{D}_C . For this, let $I \subseteq N$ denote a non-empty set of voters and consider two profiles $R, R' \in \mathcal{D}_C$ such that $\succ_i = \succ'_i$ for all $i \in N \setminus I$. Moreover, let c and c' denote the respective Condorcet winners in R and R' . If $c = c'$, then $\text{COND}(R) = \text{COND}(R')$ and the Condorcet rule is clearly group-strategyproof. On the other hand, if $c \neq c'$, there must be a voter $i \in I$ with $c \succ_i c'$ and $c' \succ'_i c$; otherwise, it is impossible that $g_R(c, c') > 0$ and $g_{R'}(c', c) > 0$. However, this voter prefers $\text{COND}(R)$ to $\text{COND}(R')$, which proves that COND is also in this case group-strategyproof.

For the other direction, we consider a group-strategyproof and non-imposing SDS f on a domain \mathcal{D} with $\mathcal{D}_C \subsetneq \mathcal{D}$. First, it is not difficult to see that f must be *ex post* efficient. Since group-strategyproofness implies strategyproofness, we can now invoke Lemma 3 to derive that for every alternative $x \in A$, there are values γ_C^x and $\gamma_i^x \geq 0$ for all $i \in N$ such that $f(R) = \gamma_C^x \text{COND}(R) + \sum_{i \in N} \gamma_i^x d_i(R)$ for all $R \in \mathcal{D}_C^\succ$. Moreover, we can essentially use the same argument as in the proof of Theorem 1 to show that $\gamma_C^x = \gamma_C^y$ and $\gamma_i^x = \gamma_i^y$ for all $i \in N$ and $x, y \in A$. We hence drop the superscript from now on and write, e.g., γ_C instead of γ_C^x .

Next, we show that $\gamma_i = 1$ if $\gamma_i > 0$. For this, we assume that there is a voter $i \in N$ with $0 < \gamma_i < 1$ and consider the profiles R^1 and R^2 shown below to derive a contradiction.

$$\begin{array}{ll} R^1: & i: c, a, b, \dots \quad N \setminus \{i\}: b, a, c, \dots \\ R^2: & i: a, b, c, \dots \quad N \setminus \{i\}: a, b, c, \dots \end{array}$$

Since b is the Condorcet winner in R^1 and $\gamma_i < 1$, we have that $f(R^1, c) = \gamma_i > 0$ and $f(R^1, b) = 1 - f(R^1, c) > 0$. On the other hand, *ex post* efficiency shows that $f(R^2, a) = 1$. However, the set of all voters can now group-manipulate by deviating from R^1 to R^2 because $f(R^1, U(\succ_{-j}^1, a)) < 1 = f(R^2, U(\succ_{-j}^1, a))$ for all $j \in N$. This contradicts that f is group-strategyproof and thus proves that $\gamma_i = 1$ if $\gamma_i > 0$. Now, since there clearly cannot be different voters i, j with $\gamma_i = 1$ and $\gamma_j = 1$, we infer that for all profiles $R \in \mathcal{D}_C$,

	Full domain \mathcal{R}^N	Domains \mathcal{D} with $\mathcal{D}_C^{(\triangleright)} \subsetneq \mathcal{D}$	(tie-breaking) Condorcet domain $\mathcal{D}_C^{(\triangleright)}$
Deterministic, strategyproof, and non-imposing voting rules	Dictatorships [24, 46]	<i>Dictatorships[◊] (Theorems 1 and 2)</i>	Dictatorships and the (tie-breaking) Condorcet rule (Theorem 2 and [10])
Strategyproof and non-imposing SDSs	Random dictatorships [25]	<i>Random dictatorships[◊] (Theorems 1 and 2)</i>	<i>Mixtures of random dictatorships and the (tie-breaking) Condorcet rule (Theorems 1 and 2)</i>
Group-strategyproof and non-imposing SDSs	Dictatorial SDSs [2]	<i>Dictatorial SDSs (Theorem 3)</i>	<i>Dictatorial SDSs and the (tie-breaking) Condorcet rule (Theorem 3)</i>

Table 1: Comparison of results for the full domain \mathcal{R}^N , strict supersets of \mathcal{D}_C (resp. $\mathcal{D}_C^{(\triangleright)}$), and the (tie-breaking) Condorcet domain \mathcal{D}_C (resp. $\mathcal{D}_C^{(\triangleright)}$). Each row characterizes a set of SDSs for the full domain \mathcal{R}^N , strict supersets of \mathcal{D}_C (resp. $\mathcal{D}_C^{(\triangleright)}$), and the (tie-breaking) Condorcet domain \mathcal{D}_C (resp. $\mathcal{D}_C^{(\triangleright)}$), respectively. For the last two columns, the results rely on a case distinction with respect to n : if n is odd, we consider the results of Theorem 1 for the Condorcet domain and its supersets; if n is even, we consider the results of Theorem 2 for the tie-breaking Condorcet domain and its supersets. The results marked with a diamond (\diamond) require that the considered domain is connected. New results are italicized.

either $f(R) = d_i(R)$ for some $i \in N$ or $f(R) = \text{COND}(R)$ if $\gamma_i = 0$ for all $i \in N$. This proves Claim (1) by choosing $\mathcal{D} = \mathcal{D}_C$.

For proving Claim (2), we assume next that there is a profile $R^* \in \mathcal{D}$ such that for every alternative $x \in A$, there is another alternative $y \in A \setminus \{x\}$ such that $g_{R^*}(y, x) > 0$. Now, consider an alternative $a \in A$ with $f(R^*, a) > 0$, let b denote an alternative with $g_{R^*}(b, a) > 0$, and define $I = \{i \in N : b \succ_i^* a\}$. We let all voters $i \in I$ make b into their best alternative to derive the profile R' . Note that $R' \in \mathcal{D}_C \subseteq \mathcal{D}$ as y is the Condorcet winner in R' . If $f(R') = \text{COND}(R')$, the voters $i \in I$ can group-manipulate by deviating from R to R' because they all prefer b to a . Hence, group-strategyproofness requires that there is a voter $i \in N$ such that $f(R) = d_i(R)$ for all $R \in \mathcal{D}_C$. From here on, it is easy to see that $f = d_i(R)$ for all $R \in \mathcal{D}$, which proves Claim (2). \square

Theorem 3 generalizes Theorem 1 to super Condorcet domains for an even number of voters by using group-strategyproofness. In particular, it entails that the Condorcet rule is the only group-strategyproof, non-imposing, and non-dictatorial SDS on the Condorcet domain, regardless of the parity of n . Moreover, Claim (2) of the theorem shows that the Condorcet domain is essentially a maximal domain that allows for a group-strategyproof and non-imposing SDS apart from dictatorships. In more detail, if n is odd, every domain \mathcal{D} with $\mathcal{D}_C \subsetneq \mathcal{D}$ satisfies the conditions of Claim (2) in Theorem 3. Hence, no superset of the Condorcet domain admits group-strategyproof and non-imposing SDSs other than dictatorships if n is odd. On the other hand, if n is even, Theorem 3 can be refined. For instance, $\text{COND}^\triangleright$ is also group-strategyproof on $\mathcal{D}_C^{(\triangleright)}$. Indeed, it is possible to prove an exact equivalent of Theorem 2 for disconnected domains based on group-strategyproofness.

Remark 3. The results of Barberà [2] imply that every group-strategyproof and non-imposing SDS on the full domain is a dictatorship. Hence, Theorem 3 and Barberà's results share a common idea: group-strategyproof and non-imposing SDSs cannot rely on randomization to determine the winner. However, whereas only undesirable SDSs are group-strategyproof and non-imposing on \mathcal{R}^N , the attractive Condorcet rule satisfies these axioms on \mathcal{D}_C .

5 CONCLUSION

We study strategyproof and non-imposing social decision schemes (SDSs) on the Condorcet domain (which consists of all preference profiles with a Condorcet winner) and its supersets. These domains are of great relevance because empirical results suggest that real-world elections usually admit a Condorcet winner. In contrast to the full domain, there are attractive strategyproof SDSs on the Condorcet domain: we show that, if the number of voters n is odd, every strategyproof and non-imposing SDS on the Condorcet domain can be represented as a mixture of a random dictatorship and the Condorcet rule. An immediate consequence of this insight is that the Condorcet rule is the only strategyproof, non-imposing, and completely non-randomly dictatorial SDS on the Condorcet domain if n is odd. Moreover, we demonstrate that, if n is odd, the Condorcet domain is a maximal connected domain that allows for strategyproof and non-imposing SDSs other than random dictatorships. We also derive analogous results for even n by slightly extending the Condorcet domain. Finally, we investigate the set of group-strategyproof and non-imposing SDSs on super Condorcet domains: we prove that the Condorcet rule is the only non-dictatorial, group-strategyproof, and non-imposing SDS on the Condorcet domain, and that no SDS satisfies these axioms on larger domains.

Our results for the Condorcet domain show an astonishing similarity to classic results for the full domain but have a more positive flavor. For instance, while the random dictatorship theorem shows that only mixtures of dictatorial SDSs are strategyproof and non-imposing on the full domain, we prove in Theorem 1 that mixtures of dictatorial SDSs and the Condorcet rule are the only strategyproof and non-imposing SDSs on the Condorcet domain (if the number of voters is odd). A more exhaustive comparison between results for the full domain and for the Condorcet domain is given in Table 1. In particular, our results highlight the important role of the Condorcet rule on the Condorcet domain: even if we allow for randomization, it is still the most appealing strategyproof voting rule. Thus, our theorems make a strong case for choosing a Condorcet winner whenever one exists.

ACKNOWLEDGEMENTS

This work was supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-2 and BR 2312/12-1. We thank the anonymous reviewers for helpful comments.

REFERENCES

- [1] N. Aswal, S. Chatterji, and A. Sen. 2003. Dictatorial domains. *Economic Theory* 22, 1 (2003), 45–62.
- [2] S. Barberà. 1979. A Note on Group Strategy-Proof Decision Schemes. *Econometrica* 47, 3 (1979), 637–640.
- [3] S. Barberà, D. Berga, and B. Moreno. 2012. Domains, ranges and strategy-proofness: the case of single-dipped preferences. *Social Choice and Welfare* 39 (2012), 335–352.
- [4] S. Barberà, F. Gul, and E. Stacchetti. 1993. Generalized Median Voter Schemes and Committees. *Journal of Economic Theory* 61 (1993), 262–289.
- [5] F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. 2015. Bypassing Combinatorial Protections: Polynomial-Time Algorithms for Single-Peaked Electorates. *Journal of Artificial Intelligence Research* 53 (2015), 439–496.
- [6] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [7] F. Brandt, P. Lederer, and R. Romen. 2022. Relaxed Notions of Condorcet-Consistency and Efficiency for Strategyproof Social Decision Schemes. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [8] F. Brandt, P. Lederer, and S. Tausch. 2023. *Strategyproof Social Decision Schemes on Super Condorcet Domains*. Technical Report. <https://arxiv.org/abs/2302.12140v1>.
- [9] R. Bredereck, J. Chen, and G. J. Woeginger. 2013. A characterization of the single-crossing domain. *Social Choice and Welfare* 41, 1 (2013), 989–998.
- [10] D. E. Campbell and J. S. Kelly. 2003. A strategy-proofness characterization of majority rule. *Economic Theory* 22, 3 (2003), 557–568.
- [11] D. E. Campbell and J. S. Kelly. 2015. Anonymous, neutral, and strategy-proof rules on the Condorcet domain. *Economics Letters* 128 (2015), 79–82.
- [12] D. E. Campbell and J. S. Kelly. 2016. Correction to “A Strategy-proofness Characterization of Majority Rule”. *Economic Theory Bulletin* 4, 1 (2016), 121–124.
- [13] S. Chatterji, R. Sanver, and A. Sen. 2013. On domains that admit well-behaved strategy-proof social choice functions. *Journal of Economic Theory* 148, 3 (2013), 1050–1073.
- [14] S. Chatterji and A. Sen. 2011. Tops-only Domains. *Economic Theory* 46 (2011), 255–282.
- [15] S. Chatterji, A. Sen, and H. Zeng. 2014. Random dictatorship domains. *Games and Economic Behavior* 86 (2014), 212–236.
- [16] S. Chatterji and H. Zeng. 2018. On random social choice functions with the tops-only property. *Games and Economic Behavior* 109 (2018), 413–435.
- [17] S. Chatterji and H. Zeng. 2021. A taxonomy of non-dictatorial domains. (2021). Working paper.
- [18] V. Conitzer. 2009. Eliciting Single-Peaked Preferences Using Comparison Queries. *Journal of Artificial Intelligence Research* 35 (2009), 161–191.
- [19] L. Ehlers, H. Peters, and T. Storcken. 2002. Strategy-Proof Probabilistic Decision Schemes for One-Dimensional Single-Peaked Preferences. *Journal of Economic Theory* 105, 2 (2002), 408–434.
- [20] E. Elkind, M. Lackner, and D. Peters. 2016. Preference restrictions in computational social choice: recent progress. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*. 4062–4065.
- [21] E. Elkind, M. Lackner, and D. Peters. 2017. Structured Preferences. In *Trends in Computational Social Choice*, U. Endriss (Ed.). Chapter 10.
- [22] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. 2011. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation* 209, 2 (2011), 89–107.
- [23] W. V. Gehrlein and D. Lepelley. 2011. *Voting Paradoxes and Group Coherence*. Springer-Verlag.
- [24] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.
- [25] A. Gibbard. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45, 3 (1977), 665–681.
- [26] A. Gopakumar and S. Roy. 2018. Dictatorship on top-circular domains. *Theory and Decision* 85, 3 (2018), 479–493.
- [27] J.-F. Laslier. 2010. In Silico Voting Experiments. In *Handbook on Approval Voting*, J.-F. Laslier and M. R. Sanver (Eds.). Springer-Verlag, Chapter 13, 311–335.
- [28] L. N. Merrill. 2011. Parity dependence of a majority rule characterization on the Condorcet domain. *Economics Letters* 112, 3 (2011), 259–261.
- [29] H. Moulin. 1980. On Strategy-Proofness and Single Peakedness. *Public Choice* 35, 4 (1980), 437–455.
- [30] K. Nehring and C. Puppe. 2007. On the Structure of Strategy-Proof Social Choice. *Journal of Economic Theory* 135 (2007), 269–305.
- [31] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- [32] W. Peremans and T. Storcken. 1999. Strategy-proofness on single-dipped preference domains. In *Proceedings of the International Conference on Logic, Game Theory and Social Choice*. 296–313.
- [33] D. Peters. 2017. Recognising Multidimensional Euclidean Preferences. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 642–648.
- [34] D. Peters and M. Lackner. 2020. Preferences Single-Peaked on a Circle. *Journal of Artificial Intelligence Research* 68 (2020), 463–502.
- [35] H. Peters, S. Roy, and S. Sadhukhan. 2021. Unanimous and strategy-proof probabilistic rules for single-peaked preference profiles on graphs. *Mathematics of Operations Research* 46, 2 (2021), 811–833.
- [36] H. Peters, S. Roy, S. Sadhukhan, and T. Storcken. 2017. An extreme point characterization of random strategy-proof and unanimous probabilistic rules over binary restricted domains. *Journal of Mathematical Economics* 69 (2017), 84–90.
- [37] H. Peters, S. Roy, A. Sen, and T. Storcken. 2014. Probabilistic strategy-proof rules over single-peaked domains. *Journal of Mathematical Economics* 52 (2014), 123–127.
- [38] M. Pycia and U. Unver. 2015. Decomposing Random Mechanisms. *Journal of Mathematical Economics* 61 (2015), 21–33.
- [39] A. Reffgen. 2015. Strategy-proof social choice on multiple and multi-dimensional single-peaked domains. *Journal of Economic Theory* 157 (2015), 349–383.
- [40] M. Regenwetter, B. Grofman, A. A. J. Marley, and I. M. Tsetlin. 2006. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press.
- [41] S. Roy and S. Sadhukhan. 2020. A unified characterization of the randomized strategy-proof rules. *Journal of Economic Theory* (2020), 105–131.
- [42] S. Roy and T. Storcken. 2019. A characterization of possibility domains in strategyproof voting. *Journal of Mathematical Economics* 84 (2019), 46–55.
- [43] A. Saporiti. 2009. Strategy-proofness and single-crossing. *Theoretical Economics* 4, 2 (2009), 127–163.
- [44] S. Sato. 2010. Circular domains. *Review of Economic Design* (2010), 331–333.
- [45] S. Sato. 2013. A sufficient condition for the equivalence of strategy-proofness and non-manipulability by preferences adjacent to the sincere one. *Journal of Economic Theory* 148 (2013), 259–278.
- [46] M. A. Satterthwaite. 1975. Strategy-Proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.
- [47] Y. Shoham and K. Leyton-Brown. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.

SUMMARY

We study strategyproofness for randomized voting rules, which are called social decision schemes (SDSs). For SDSs, strategyproofness is typically formalized with the help of utility function and the standard strategyproofness notion in the literature (strong \succsim^{SD} -strategyproofness) ensures that a voter cannot increase her expected utility regardless of their exact utility function. For this strategyproofness notion, Gibbard (1977) has characterized the set of strategyproof SDSs and, in particular, demonstrates that all these SDSs are indecisive or unfair.

In this paper, we study the tradeoff between strategyproofness and decisiveness by considering weaker strategyproofness notions. In more detail, we analyze strong \succsim^U -strategyproofness (which is called U-strategyproofness in the paper and the rest of this summary), which only requires that voters with a utility function in a given set U cannot manipulate. This strategyproofness notion is motivated by the fact that often not all utility functions are plausible or likely. In such cases, strong \succsim^{SD} -strategyproofness is unnecessarily restrictive and might force us to use an undesirable SDS even though there may be more attractive voting rules that are U-strategyproof for the considered set U . We are thus interested in whether U-strategyproofness allows for the design of appealing voting rules that are both decisive and strategyproof for a large set of utility functions U .

To answer this question, we first show that variants of the uniform random dictatorship can satisfy k-unanimity (i.e., they choose an alternative with probability 1 if it is top-ranked by $n - k$ voters) for all $k < \frac{n}{2}$ and U-strategyproofness for utility functions that value the best alternative much more than the other ones. While this result requires a quite restricted set of utility functions, it constitutes a possibility theorem with respect to the tradeoff between decisiveness and strategyproofness. We furthermore show that our considered SDSs solve this tradeoff almost optimally: no rank-based and k-unanimous SDS can be U-strategyproof for a significantly larger set of utility functions U than our variants of the uniform random dictatorship.

Finally, we also prove that U-strategyproofness is incompatible with Condorcet-consistency if the set of utility functions U satisfies minimal symmetry and richness conditions. This shows that also the concept of U-strategyproofness does not allow for too decisive SDSs.

REFERENCE

P. Lederer. **Strategyproof social choice for restricted sets of utility functions.** In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 306–312, 2021.
DOI: <https://doi.org/10.24963/ijcai.2021/43>

INDIVIDUAL CONTRIBUTION

This is a single-authored publication and I, Patrick Lederer, am responsible for all parts of this publication.

COPYRIGHT AGREEMENT

The right to present this paper in a doctoral thesis has been granted by the publisher, the International Joint Conference on Artificial Intelligence (IJCAI), in the copyright form presented below. There, IJCAI grants permission to “*personal reuse of all or portions of the paper in other works of their own authorship.*” This form can also be found at https://proceedings.ijcai.org/docs/IJCAI-Copyright_Transfer_Agreement.pdf (accessed August 24, 2023).

TERMINOLOGY

In this paper, SD -strategyproofness refers to strong \succ^{SD} -strategyproofness and U -strategyproofness to strong \succ^U -strategyproofness. Moreover, the voters’ preferences are always strict.



INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE

TRANSFER OF COPYRIGHT AGREEMENT

Title of Article/Paper: _____

Publication in Which Article Is to Appear: _____

Author's Name(s): _____

Please type or print your name as you wish it to appear in print

(Please read and sign Part A only, unless you are a government employee and created your article/paper as part of your employment. If your work was performed under Government contract, but you are not a Government employee, sign Part A and see item 6 under returned rights.)

PART A—Copyright Transfer Form

The undersigned, desiring to publish the above article/paper in a publication of the International Joint Conferences on Artificial Intelligence, Inc., hereby transfer their copyrights in the above paper to the International Joint Conferences on Artificial Intelligence, Inc. (IJCAI), in order to deal with future requests for reprints, translations, anthologies, reproductions, excerpts, and other publications.

This grant will include, without limitation, the entire copyright in the paper in all countries of the world, including all renewals, extensions, and reversions thereof, whether such rights currently exist or hereafter come into effect, and also the exclusive right to create electronic versions of the paper, to the extent that such right is not subsumed under copyright. The undersigned warrants that he/she is the sole author and owner of the copyright in the above paper, except for those portions shown to be in quotations; that the paper is original throughout; and that the undersigned's right to make the grants set forth above is complete and unencumbered.

If anyone brings any claim or action alleging facts that, if true, constitute a breach of any of the foregoing warranties, the undersigned will hold harmless and indemnify IJCAI, their grantees, their licensees, and their distributors against any liability, whether under judgment, decree, or compromise, and any legal fees and expenses arising out of that claim or actions, and the undersigned will cooperate fully in any defense IJCAI may make to such claim or action. Moreover, the undersigned agrees to cooperate in any claim or other action seeking to protect or enforce any right the undersigned has granted to IJCAI in the paper. If any such claim or action fails because of facts that constitute a breach of any of the foregoing warranties, the undersigned agrees to reimburse whomever brings such claim or action for expenses and attorney's fees incurred therein.

Returned Rights

In return for these rights, IJCAI hereby grants to the above authors, and the employers for whom the work was performed, royalty-free permission to:

1. retain all proprietary rights (such as patent rights) other than copyright and the publication rights transferred to IJCAI;
2. personally reuse all or portions of the paper in other works of their own authorship;
3. make oral presentation of the material in any forum;
4. reproduce, or have reproduced, the above paper for the author's personal use, or for company use provided that IJCAI copyright and the source are indicated, and that the copies are not used in a way that implies IJCAI endorsement of a product or service of an employer, and that the copies per se are not offered for sale. The foregoing right shall not permit the posting of the paper in electronic or digital form on any computer network, except by the author or the author's employer, and then only on the author's or the employer's own World Wide Web page or ftp site. Such Web page or ftp site, in addition to the aforementioned requirements of this Paragraph, must provide an electronic reference or link back to the IJCAI electronic server (<http://www.ijcai.org>), and shall not post other IJCAI copyrighted materials not of the author's or the employer's creation (including tables of contents with links to other papers) without IJCAI's written permission;
5. make limited distribution of all or portions of the above paper prior to publication.
6. In the case of work performed under U.S. Government contract, IJCAI grants the U.S. Government royalty-free permission to reproduce all or portions of the above paper, and to authorize others to do so, for U.S. Government purposes. In the event the above paper is not accepted and published by IJCAI, or is withdrawn by the author(s) before acceptance by IJCAI, this agreement becomes null and void.

Author's (or Employer's Representative) Signature

Date

Employer for whom work was performed

Title (if not author)

Strategyproof Randomized Social Choice for Restricted Sets of Utility Functions

Patrick Lederer

Technische Universität München

ledererp@in.tum.de

Abstract

When aggregating preferences of multiple agents, strategyproofness is a fundamental requirement. For randomized voting rules, so-called social decision schemes (SDSs), strategyproofness is usually formalized with the help of utility functions. A classic result shown by Gibbard in 1977 characterizes the set of SDSs that are strategyproof with respect to all utility functions and shows that these SDSs are either indecisive or unfair. For finding more insights into the trade-off between strategyproofness and decisiveness, we propose the notion of U -strategyproofness which requires that only voters with a utility function in the set U cannot manipulate. In particular, we show that if the utility functions in U value the best alternative much more than other alternatives, there are U -strategyproof SDSs that choose an alternative with probability 1 whenever all but k voters rank it first. We also prove for rank-based SDSs that this large gap in the utilities is required to be strategyproof and that the gap must increase in k . On the negative side, we show that U -strategyproofness is incompatible with Condorcet-consistency if U satisfies minimal symmetry conditions and there are at least four alternatives. For three alternatives, the Condorcet rule can be characterized based on U -strategyproofness for the set U containing all equidistant utility functions.

1 Introduction

When a group of agents wants to find a joint decision in a structured way, they can choose from a multitude of different voting rules. However, it is not clear which rule is the best one as each one has its benefits. This problem lies at the core of social choice theory which draws increased attention by computer scientists because it can be used to reason about computational multi-agent systems (see, e.g., [Chevalleyre *et al.*, 2007; Brandt *et al.*, 2013; Brandt *et al.*, 2016b; Endriss, 2017]). A fundamental requirement for voting rules is strategyproofness, i.e., agents should not be able to benefit by lying about their preferences. In a seminal result, Gibbard [1973] and Satterthwaite [1975] have shown that every

deterministic strategyproof voting rule is dictatorial if there are at least three different outcomes possible.

Randomization allows to escape this impossibility theorem, and we analyze therefore social decision schemes (SDSs). These functions aggregate the preferences of agents to lotteries over alternatives which determine for every alternative its winning chances. The final winner is then decided by chance according to these probabilities. While this model allows to circumvent many impossibilities, it is not straightforward how to define strategyproofness because the voters' preferences over lotteries are unclear. Maybe the most prominent approach is to assume that voters use cardinal utility functions on the alternatives to compare lotteries with respect to their expected utilities. However, voters still report ordinal preference relations to the SDS and hence, strategyproofness is defined by quantifying over utility functions: an SDS is strategyproof if voting honestly maximizes the expected utility for every voter and every utility function that is consistent with his true preferences. This strategyproofness notion, often called SD -strategyproofness, has been analyzed by Gibbard [1977] and Barberà [1979] who prove that all SD -strategyproof SDSs are indecisive because they almost always randomize over multiple alternatives. Even more, Benoît [2002] has shown that SD -strategyproofness is incompatible with the basic democratic idea that an alternative should be the winner of an election if an absolute majority of the voters report it as their best alternative.

While it is unfortunate that SD -strategyproofness does not allow for decisive SDSs, this strategyproofness notion seems also too demanding for because in many applications not all utility functions are plausible. For instance, when a representative body votes about budget proposals, it seems reasonable that similar proposals have similar utilities. Thus, we might neglect utility functions with a large gap between such options when discussing strategyproofness. This observation leads to the new notion of U -strategyproofness which requires that truth telling only maximizes the expected utility of a voter if his utility function is in the set U . Note that U -strategyproofness does not forbid utility functions $u \notin U$, but voters with such utility functions might be able to manipulate.

U -strategyproofness allows for a more detailed analysis than SD -strategyproofness because we can analyze the exact set of utility functions U for which an SDS is U -strategyproof. Conversely, we can also formulate strong im-

possibility results based on U -strategyproofness for severely restricted sets U and thus, we can pinpoint the source of manipulability far more detailed than with other strategyproofness notions. Hence, U -strategyproofness offers both the possibility of positive results by finding U -strategyproof SDSs for large sets U , and of strong impossibility results by using only a small number of utility functions. Furthermore, information about U -strategyproofness can also be valuable in practice: if the social planner can roughly guess the utility functions of the voters, he might be able to choose an SDS preventing manipulations. Even if the social planner does not have such insights, he might opt for an SDS that is U -strategyproof for a large set U as such an SDS is immune to manipulations from most voters.

Other than introducing U -strategyproofness, we use this new notion to investigate the trade-off between strategyproofness and decisiveness. On the positive side, we show that there are U -strategyproof SDSs that assign an alternative probability 1 whenever all but $k > 0$ voters agree that it is the best option if the utility functions in U value the best alternative much more than the other alternatives. Moreover, we prove for rank-based SDSs that this gap in the utility functions is required to be strategyproof and that it must increase in k . On the other hand, we show that Condorcet-consistency is incompatible with U -strategyproofness if the set U satisfies minimal symmetry conditions between preference relations and there are $m \geq 4$ alternatives. If there are only three alternatives and an odd number of voters, the Condorcet rule is characterized by U -strategyproofness for the set U of all equi-distant utility functions and Condorcet-consistency. The proofs of these theorems and of all propositions are omitted because of space limitations.

2 Related Work

To our knowledge, we are the first authors who explicitly investigate U -strategyproofness. Nevertheless, ideas similar to U -strategyproofness have been used before. For instance, Sen [2011] and Mennle and Seuken [2021] define strategyproofness by considering restricted sets of utility functions and thus, their works can be interpreted as first results on U -strategyproofness. Moreover, in set-valued social choice (where the outcome of an election is a non-empty set of alternatives instead of a lottery) preferences over sets of alternatives are often derived from utility functions. For instance, Duggan and Schwartz [2000] and Benoît [2002] employ this approach to motivate their strategyproofness notions. The relationship between these results and U -strategyproofness is discussed in more detail in Section 4.

There are also various results on other strategyproofness notions in randomized social choice (see, e.g., [Gibbard, 1977; Hoang, 2017; Aziz *et al.*, 2018; Brandl *et al.*, 2018]), many of which are surveyed by Brandt [2017]. These results either prove the incompatibility of strategyproofness with other axioms or characterize specific SDSs. Our results differ from previous ones as we investigate a different question: instead of asking whether an SDS is strategyproof according to some definition, we ask for which utility functions it is strategyproof.

Moreover, strategyproofness is often considered for restricted domains of preference profiles (see, e.g., [Ehlers *et al.*, 2002; Bogomolnaia *et al.*, 2005; Chatterji and Zeng, 2018]). For instance, Bogomolnaia *et al.* [2005] discuss an attractive SD -strategyproof SDS for dichotomous preferences. U -strategyproofness can be interpreted similarly, but we focus on utility functions instead of preference profiles: U -strategyproof SDSs are immune to manipulations if we only allow utility functions in U .

Another field related to U -strategyproofness is cardinal social choice, where the input of social decision schemes consists of the utility functions of the voters. If we allow all utility functions as input, every strategyproof cardinal SDS is, under mild additional assumptions, a variant of a random dictatorship (see, e.g., [Hylland, 1980; Dutta *et al.*, 2007; Nandebam, 2013]). As noted by Dutta *et al.* [2007], these negative results break down if the domain of cardinal SDSs is restricted, but this setting is not well understood. Our results provide insights in this problem because every U -strategyproof SDS can be interpreted as a cardinal SDS that is strategyproof on the domain U .

Finally, note that our model assumptions are quite similar to those used in the analysis of the distortion of SDSs (see, e.g., [Procaccia and Rosenschein, 2006; Gross *et al.*, 2017; Abramowitz *et al.*, 2019]). Just as these authors, we assume that voters only report ordinal preferences but use utility functions to evaluate the quality of a lottery. Whereas distortion focuses on the welfare of SDSs, we investigate their resistance to strategic behavior of voters.

3 Preliminaries

Let $N = \{1, \dots, n\}$ be a finite set of voters and let A be a set containing m alternatives. A *preference relation* is an anti-symmetric, transitive, complete, and reflexive binary relation on A and R_i denotes the preference relation of voter i . We compactly represent preference relations as comma-separated lists. Let \mathcal{R} denote the set of all preference relations on A . A *preference profile* R is an n -tuple containing the preference of every voter $i \in N$, i.e., $R \in \mathcal{R}^n$. When writing preference profiles, we indicate the corresponding voter directly before the preference relation to clarify which voter submits which preference relation. For example, $1 : a, b, c$ indicates that voter 1 reports that he prefers a to b to c .

In this paper, we discuss *social decision schemes* (SDSs), which are functions that map preference profiles to lotteries on A . A *lottery* p is a function from the set of alternatives A to the interval $[0, 1]$ such that $\sum_{x \in A} p(x) = 1$. Let $\Delta(A)$ denote the set of all lotteries on A . Formally, a social decision scheme is a function $f : \mathcal{R}^n \rightarrow \Delta(A)$ and we denote with $f(R, x)$ the probability assigned to x by the lottery $f(R)$.

The definition of SDSs allows for a huge variety of functions, some of which seem not desirable. Therefore, we introduce axioms to narrow down the set of SDSs. Two basic fairness axioms are anonymity and neutrality, which require that voter and alternatives, respectively, are treated equally. More formally, an SDS f is *anonymous* if $f(R) = f(\pi(R))$ for all profiles R and permutations $\pi : N \rightarrow N$, and *neutral* if $f(R, x) = f(\tau(R), \tau(x))$ for all alternatives $x \in A$,

profiles R , and permutations $\tau : A \rightarrow A$. Another natural axiom is *unanimity*, which requires of an SDS f that $f(R, x) = 1$ for all preference profiles R in which all voters agree that x is the best choice. While this axiom is so weak that is often considered indisputable, it is also irrelevant in practice as ballots are usually not unanimous. Therefore, we introduce the stronger notion of *k-unanimity*: an SDS f is *k-unanimous* if $f(R, x) = 1$ whenever $n - k$ or more voters report x as the best alternative. By definition, unanimity is equal to 0-unanimity and note that *k-unanimity* is only well-defined if $k < \frac{n}{2}$. A well-known strengthening of *k-unanimity* is *Condorcet-consistency*. For defining this axiom, let $n_{xy}(R) = |\{i \in N : xR_iy\}| - |\{i \in N : yR_ix\}|$ denote the *majority margin* between two alternatives $x, y \in A$ in the preference profile R . An alternative x is the *Condorcet winner* in a preference profile R if $n_{xy}(R) > 0$ for all other alternatives $y \in A \setminus \{x\}$. Less formally, an alternative x is the Condorcet winner if it is preferred to every other alternative by a majority of the voters. Finally, an SDS f is *Condorcet-consistent* if $f(R, x) = 1$ for all profiles R and alternatives $x \in A$ such that x is the Condorcet winner in R .

An important class of SDSs are rank-based SDSs. The basic idea of these schemes is that voters assign ranks to the alternatives and that an SDS should only rely on these ranks, but not on which voter assigns which rank to an alternative. For formalizing this concept, we denote with $r(R_i, x) = |\{y \in A : yR_ix\}|$ the *rank* of alternative x in voter i 's preference relation. Moreover, we define the *rank vector* $r^*(R, x)$ as the vector that contains the rank of x with respect to every voter in increasing order, i.e., $r^*(R, x)_i \leq r^*(R, x)_{i+1}$ for all $i \in \{1, \dots, n-1\}$, and the *rank matrix* $r^*(R)$ as the matrix that contains the rank vectors of all alternative as rows. Finally, we call an SDS f *rank-based* if it only depends on the rank matrix, i.e., $f(R) = f(R')$ for all preference profiles R, R' with $r^*(R) = r^*(R')$. The set of rank-based SDSs contains many prominent functions such as point scoring rules and anonymous SDSs that only depend on the first-ranked alternatives of the voters.

4 U-Strategyproofness

A central problem in social choice is that of manipulability: voters may lie about their preferences to achieve a better outcome. While the definition of a manipulation is easy if an SDS never randomizes between multiple alternatives, it is not clear how to compare non-degenerate lotteries. A classical approach for this problem is to assume that voters are endowed with *utility functions* $u_i : A \rightarrow \mathbb{R}$. We impose the constraint that no voter assigns the same utility to two alternatives, i.e., $u_i(x) \neq u_i(y)$ for all voters $i \in N$ and alternatives $x, y \in A$, to ensure that the ordinal preference relation induced by a utility function is anti-symmetric. We denote with \mathcal{U} the set of all such utility functions and say that a utility function $u \in \mathcal{U}$ is *consistent* with a preference relation R if $u(x) \geq u(y)$ iff xRy for all alternatives $x, y \in A$. Finally, each voter i uses his utility function u_i to compare lotteries by their *expected utilities* $\mathbb{E}[p]_{u_i} = \sum_{x \in A} p(x)u_i(x)$, i.e., voter i prefers lottery p weakly to lottery q if $\mathbb{E}[p]_{u_i} \geq \mathbb{E}[q]_{u_i}$.

Even though we assume the existence of utility functions,

voters only report ordinal preferences. Consequently, strategyproofness is often defined by quantifying over utility functions. In particular, Gibbard [1977] employs this approach to define *SD-strategyproofness*: an SDS f is *SD-strategyproof* if $\mathbb{E}[f(R)]_{u_i} \geq \mathbb{E}[f(R')]_{u_i}$ for all voters $i \in N$, preference profiles R, R' , and *utility functions* $u_i \in \mathcal{U}$ such that u_i is consistent with R_i and $R_j = R'_j$ for all $j \in N \setminus \{i\}$. While *SD-strategyproofness* allows for strong negative results (see, e.g. [Gibbard, 1977; Barberà, 1979]), it lacks relevance for many practical applications as not all utility functions are plausible. Also, *SD-strategyproofness* provides often only shallow theoretical insights as it is not possible to pinpoint the source of manipulability.

In order to address these problems, we introduce a new strategyproofness notion by restricting the set of feasible utility functions U beforehand: an SDS f is *U-strategyproof* if $\mathbb{E}[f(R)]_{u_i} \geq \mathbb{E}[f(R')]_{u_i}$ for all voters $i \in N$, preference profiles R, R' , and *utility functions* $u_i \in U$ such that u_i is consistent with R_i and $R_j = R'_j$ for all $j \in N \setminus \{i\}$. Less formally, *U-strategyproofness* only requires that voters with a utility function in U cannot increase their expected utility by misrepresenting their preferences. Hence, *U-strategyproofness* is equal to *SD-strategyproofness* and smaller sets of utility functions result in less demanding strategyproofness notions. Note that *U-strategyproofness* solves both problems of *SD-strategyproofness*: we can investigate whether an SDS is manipulable in practice by dismissing implausible utility functions, and we can find the core of impossibility results by determining the minimally required set of utility functions. Next, we discuss an example to illustrate the difference between *U-strategyproofness* and *SD-strategyproofness*.

Example 1. Consider the profiles R^1 and R^2 shown below and let f denote an SDS such that $f(R^1, x) = \frac{1}{3}$ for $x \in \{a, b, c\}$ and $f(R^2, b) = 1$. Moreover, consider the utility functions u_1, u_2 , and u_3 with $u_1(a) = 2, u_1(b) = 1, u_1(c) = 0, u_2(a) = 3, u_2(b) = 1, u_2(c) = 0, u_3(a) = 3, u_3(b) = 2$, and $u_3(c) = 0$. These utility functions are only consistent with voter 1's preference relation in R^1 , and thus, we can check whether this voter can benefit by deviating to R^2 . A quick calculation shows that $\mathbb{E}[f(R^1)]_{u_1} = 1 = \mathbb{E}[f(R^2)]_{u_1}, \mathbb{E}[f(R^1)]_{u_2} = \frac{4}{3} > 1 = \mathbb{E}[f(R^2)]_{u_2}$, and $\mathbb{E}[f(R^1)]_{u_3} = \frac{5}{3} < 2 = \mathbb{E}[f(R^2)]_{u_3}$. Hence, voter 1 can increase his expected utility if his utility function is u_3 and thus, f is *SD-manipulable*. In contrast, voter 1 does not benefit from deviating to R^2 if his utility function is u_1 or u_2 . Since the preferences of the other voters are not consistent with u_1, u_2 , and u_3 , it follows that f is $\{u_1, u_2\}$ -strategyproof on these two profiles.

R^1 :	1: a, b, c	2: b, c, a	3: c, a, b
R^2 :	1: b, a, c	2: b, c, a	3: c, a, b

In our results, we always consider *U-strategyproofness* for *symmetric* sets U , i.e., we assume that $u \in U$ implies that $u^\pi = u \circ \pi \in U$ for every permutation π on A . This formalizes the natural condition that all preference relations should be treated equally. Moreover, the symmetry condition is rather weak since every neutral SDS is *U'-strategyproof* for a symmetric set U' if it is *U-strategyproof* for a set $U \neq \emptyset$.

Proposition 1. *If a neutral SDS is U -strategyproof for a set $U \neq \emptyset$, it is U' -strategyproof for a symmetric set U' with $U \subseteq U'$.*

A special case of our symmetry assumption is that U consists of a single utility function u and its renamings, i.e., that $U = \{u \circ \pi : \pi \in \Pi\}$, where Π denotes the set of all permutations on A . In this case, we write u^Π -strategyproofness instead of U -strategyproofness. Note that u^Π -strategyproofness associates every preference relation with exactly one utility function, whereas $\{u\}$ -strategyproofness, i.e., strategyproofness for a single utility function u , only affects a single preference relation. Since the utility of an alternative only depends on its rank for u^Π -strategyproofness, we often write $u(k)$ to denote the utility of the k -th best alternative of a voter. As the next proposition shows, it suffices to consider u^Π -strategyproofness or even $\{u\}$ -strategyproofness because for every SDS f and every preference relation R_i , the set of utility functions u that are consistent with R_i and for which f is strategyproof is convex.

Proposition 2. *For every SDS f and preference relation R_i , the set $U_{R_i} = \{u \in \mathcal{U} : u \text{ is consistent with } R_i \text{ and } f \text{ is } \{u\}\text{-strategyproof}\}$ is convex.*

We can use this proposition to show that an SDS is U -strategyproof for a large set U by proving that it is u_i^Π -strategyproof for a few utility functions $u_i \in \{u_1, \dots, u_l\}$. Assuming that u_1, \dots, u_l are all consistent with a preference relation R_i , it follows then from Proposition 2 that the SDS is \hat{u}^Π -strategyproof for every utility function \hat{u} that can be represented as a convex mixture of u_1, \dots, u_l , which means that it is U -strategyproof for a large set U .

Next, note that U -strategyproofness inherits many attractive properties from SD -strategyproofness: for instance, the convex combination of U -strategyproof SDSs is itself U -strategyproof, i.e., the set of U -strategyproof SDSs is convex for every set U . As a consequence of this observation, it is often possible to construct an anonymous U -strategyproof SDS based on a non-anonymous U -strategyproof SDS. Another similarity between U -strategyproofness and SD -strategyproofness is that both axioms disincentivize even manipulations from groups of voters with the same preferences.

Finally, observe that U -strategyproofness can be used to transfer results from set-valued social choice to the probabilistic setting. We explain this relation using the impossibility result of Benoît [2002] as example. This theorem states that strategyproofness is incompatible with 1-unanimity for set-valued social choice functions if voters prefer every subset of their best two alternatives to every other set and other in our model negligible conditions are satisfied. For formulating this result for SDSs, we have to compare lotteries only based on their support $\text{supp}(p) = \{x \in A : p(x) > 0\}$. Hence, let $\epsilon_f = \min_{x \in A, R \in \mathcal{R}^n : f(R, x) > 0} f(R, x)$ denote the smallest non-zero probability assigned to an alternative by the SDS f and note that ϵ_f is well-defined since SDSs are defined for a fixed set of alternatives and voters. Given this probability, we derive that every voter whose utility function u satisfies $u(2) > (1 - \epsilon_f)u(1) + \epsilon_f u(3)$ prefers every lottery that randomizes only over his best two alternatives to every other lottery. After rearranging this equation, we can formu-

late Benoît's impossibility as follows.

Proposition 3. *No SDS f satisfies both u^Π -strategyproofness and 1-unanimity if $u(1) - u(2) < \frac{\epsilon_f}{1 - \epsilon_f}(u(2) - u(3))$, $m \geq 3$, and $n \geq 3$.*

Note that Proposition 3 highlights the central requirement of Benoît's impossibility theorem: voters must be close to indifferent between their best two alternatives. This refines Benoît's reasoning who justifies his strategyproofness notion with voters who "like his or her two favorite alternatives "much more" than the rest of the alternatives".¹ Based on this approach, we can also formalize other impossibility results from set-valued social choice with U -strategyproofness.

5 Results

In the sequel, we employ U -strategyproofness to analyze the trade-off between strategyproofness and decisiveness. In particular, we investigate two decisiveness axioms: k -unanimity and Condorcet-consistency. The first axiom allows for positive results if suitable utility functions are considered, whereas Condorcet-consistency is incompatible with u^Π -strategyproofness for every utility function $u \in \mathcal{U}$.

5.1 k -unanimity

A central result of Gibbard [1977], who attributes it to Hugo Sonnenschein, is that the SDS called random dictatorship (henceforth RD) is the only SD -strategyproof SDS that satisfies unanimity and anonymity. This SDS assigns an alternative x in a profile R the probability $\frac{PL(R, x)}{n}$, where $PL(R, x) = |\{i \in N : \forall y \in A : x R_i y\}|$ denotes the plurality score of alternative x . A common method for executing RD is to choose a voter uniformly at random and to return his most preferred alternative as winner. While RD is one of the most attractive SD -strategyproof SDSs, it violates k -unanimity for $k > 0$. Even more, Benoît [2002] has shown that every SD -strategyproof SDS fails k -unanimity for $k > 0$.

However, we can define a variant of RD that satisfies both k -unanimity for an arbitrary $k \in \{0, \dots, \lfloor \frac{n-1}{2} \rfloor\}$ and U -strategyproofness for a large set of utility functions U . Hence, consider the following SDS, which we call k -random dictatorship (abbreviated by RD^k): if at least $n - k$ voters agree that alternative x is the best choice, assign alternative x a probability of 1; otherwise, return the outcome of RD . As we show in Theorem 1, RD^k satisfies U -strategyproofness for $U = \{u \in \mathcal{U} : u(1) - u(2) \geq k(u(2) - u(m))\}$, i.e., if voters have a strong preference for the first alternative, RD^k is strategyproof. Unfortunately, the definition of U depends on k , i.e., for large values of k , there must be an extremely large gap between $u(1)$ and $u(2)$. Another variant of RD , which we refer to as $OMNI^*$, solves this problem. This SDS assigns probability 1 to an alternative x if more than half of the voters report x as their best alternative, and otherwise randomizes uniformly among all alternatives that are

¹Benoît [2002] also discusses a variant for SDSs in which he uses the minimal non-zero probability assigned to an alternative. However, Benoît only gives an example showing that there is a suitable utility function such that the required preferences over sets extend to preferences over lotteries.

at least once top-ranked. This SDS is U -strategyproof for $U = \{u \in \mathcal{U}: u(1) - u(2) \geq \sum_{i=3}^m u(2) - u(i)\}$. While $OMNI^*$ satisfies $\lfloor \frac{n-1}{2} \rfloor$ -unanimity for all numbers of voters and alternatives, the condition on U seems only realistic if there are few alternatives.

Theorem 1. For every $k \in \{1, \dots, \lfloor \frac{n-1}{2} \rfloor\}$, RD^k satisfies U -strategyproofness for $U = \{u \in \mathcal{U}: u(1) - u(2) \geq k(u(2) - u(m))\}$ and violates $\{u\}$ -strategyproofness for every utility function $u \notin U$. Moreover, $OMNI^*$ satisfies U -strategyproofness for $U = \{u \in \mathcal{U}: u(1) - u(2) \geq \sum_{i=3}^m u(2) - u(i)\}$ and violates $\{u\}$ -strategyproofness for every utility function $u \notin U$.

The constraint on the set U for RD^k arises naturally by considering the preference profile in which $n - k - 1$ voters top-rank the second best alternative of voter i and the remaining k voters top-rank voter i 's least preferred alternative. In this situation, voter i can ensure that his second best alternative is chosen with probability 1 by reporting it as his best one. Solving the corresponding inequality required by U -strategyproofness leads to the bound on U . A similar worst-case analysis can be applied for $OMNI^*$.

While it is positive that k -unanimity and U -strategyproofness can be simultaneously satisfied at all, the bounds on the sets U in Theorem 1 become increasingly worse with large k and m . This raises the question for less demanding bounds on the utility functions. As our next theorem shows, the approach used for defining RD^k and $OMNI^*$ has not much space for improvement as both SDSs are rank-based.

Theorem 2. There is no rank-based SDS that satisfies u^Π -strategyproofness and k -unanimity for $0 < k < \frac{n}{2}$ if $m \geq 3$, $n \geq 3$, and $u(1) - u(2) < \sum_{i=\max(3, m-k+1)}^m u(2) - u(i)$.

The proof of Theorem 2 works by contradiction: we assume that there is a k -unanimous rank-based SDS f that satisfies u^Π -strategyproofness for a utility function u with $u(1) - u(2) < \sum_{i=\max(3, m-k+1)}^m u(2) - u(i)$. Moreover, let $k^* = \min(k, m - 2)$. Our analysis then starts at a profile R where $n - k^*$ voters favor a the most, which implies that $f(R, a) = 1$ due to k -unanimity. The central argument is a rather involved construction that shows that a voter can weaken alternative a from the first rank to the second one without affecting the outcome. By repeatedly applying this construction, we eventually arrive at a profile R' where only k^* voters top-rank a and the remaining voters top-rank b , but $f(R', a) = 1$. This is in conflict with k -unanimity as $n - k^* \geq n - k$ voters report b as best choice but $f(R', b) \neq 1$.

Remark 1. A computer-aided approach has shown that there are rather technical SDSs that satisfy k -unanimity and u^Π -strategyproofness for utility functions u with $u(1) - u(2) < \sum_{i=\max(3, m-k+1)}^m u(2) - u(i)$ if we dismiss rank-basedness and $m \leq 4$. Hence, rank-basedness is required for Theorem 2. Moreover, most bounds of the theorem are tight: if $m = 2$, $OMNI^*$ and RD^k are even SD -strategyproof, and if $n = 2$, k -unanimity is not well-defined for $k > 0$. Furthermore, the condition on the utility functions is almost tight: RD^1 shows that the bound is tight for 1-unanimity, and $OMNI^*$ shows that the bound is tight if




Figure 1: Illustration of Theorem 1 and Theorem 2. We assume that there are 5 alternatives and consider a utility function u with $u(2) = 3$, $u(3) = 2$, $u(4) = 1$, and $u(5) = 0$. The figure shows for which values of $u(1)$ the SDSs RD (blue area), RD^1 (green area), RD^2 (magenta area), and $OMNI^*$ (orange area) are u^Π -strategyproof on the vertical axis. The horizontal axis illustrates the values of k for which these SDSs are k -unanimous. The red area displays the impossibility of Theorem 2 and the gray area marks the values of $u(1)$ with $u(1) < u(2)$.

$k \geq m - 2$. Finally, RD^k shows that no constraint of the type $u(1) - u(2) \leq \sum_{i=m-k+1}^m u(2) - u(i) + \epsilon$ with $\epsilon > 0$ can result in an impossibility because we can always find a utility function u such that $\sum_{i=m-k+1}^m u(2) - u(i) + \epsilon \geq u(1) - u(2) \geq k(u(2) - u(m))$ by making the difference between $u(i)$ and $u(m)$ for $i \geq 3$ sufficiently small. Nevertheless, it remains open to find rank-based SDSs that satisfy U -strategyproofness and k -unanimity for $U = \{u \in \mathcal{U}: u(1) - u(2) = \sum_{i=m-k+1}^m u(2) - u(i)\}$ and $2 \leq k \leq m - 3$.

Remark 2. Theorem 1 and Theorem 2 have an intuitive interpretation: if voters strongly prefer their best alternative, it becomes possible to achieve strategyproofness and decisiveness. This follows as strategyproofness is compatible with k -unanimity if there is a sufficiently large gap between $u(1)$ and $u(2)$. In contrast, it is impossible that an SDS satisfies both axioms if voters are close to indifferent between their best two alternatives. For the class of general SDSs, this is shown by Benoît [2002], and for the class of rank-based SDSs, Theorem 2 significantly weakens the requirements on the utility functions.

Remark 3. Figure 1 illustrates the results of this section. For this figure, we assume that there are 5 alternatives and a large number of voters $n \geq 11$, and we fix all utilities but $u(1)$. Hence, we can compute the values of $u(1)$ for all SDSs of Theorem 1 such that the considered SDS is u^Π -strategyproof. The figure shows that for RD^k , the required value of $u(1)$ increases in k and the bound of $OMNI^*$ is independent of k . Moreover, the required values of $u(1)$ are quite large compared to $u(2)$ for all SDSs but RD . However, the red area shows the values of $u(1)$ for which Theorem 2 applies and hence, these large values are indeed required. The white area shows that there is a small gap between the positive results in Theorem 1 and the impossibility in Theorem 2.

5.2 Condorcet-consistency

As there are even rank-based SDS that are k -unanimous and U -strategyproof for large sets U , the question arises whether stronger decisiveness notions can be achieved by dismissing rank-basedness. Unfortunately, we find a negative answer to this question by considering Condorcet-consistency.

Theorem 3. *There is no Condorcet-consistent SDS that satisfies u^Π -strategyproofness regardless of the utility function u if $m \geq 4$, $n \geq 5$ and $n \neq 6$, $n \neq 8$.*

The proof of this result works by contradiction and relies on a case distinction on the utility function u . If $u(1) - u(2) < u(2) - u(m)$, the utility of the second best alternative is larger than the average utility, which means that a voter can manipulate by making his second best alternative into the Condorcet winner. If $u(1) - u(m-1) > u(m-1) - u(m)$, voters value their second worst alternative less than the uniform lottery. As a consequence, there is a voter who can manipulate by weakening his second worst alternative such that it is no longer the Condorcet winner. Finally, note that these two cases are exhaustive: the strictness of the utility function u entails that $u(m-1) - u(m) < u(1) - u(m-1)$ if $u(1) - u(2) \geq u(2) - u(m)$ and $m \geq 4$.

A close inspection of the proof shows that the impossibility also holds if $m = 3$ unless U only contains equi-distant utility functions, i.e., utility functions with $u(1) - u(2) = u(2) - u(3)$. This raises the question whether there is a U -strategyproof SDS that satisfies Condorcet-consistency in this special case. Indeed, the Condorcet rule (abbreviated by *COND*), which assigns probability 1 to the Condorcet winner whenever it exists and returns the uniform lottery over all alternatives otherwise, satisfies U -strategyproofness for this set. Even more, the Condorcet rule is uniquely characterized by these axioms if n is odd.

Theorem 4. *COND is the only Condorcet-consistent SDS that satisfies U -strategyproofness for $U = \{u \in \mathcal{U} : u(1) - u(2) = u(2) - u(3)\}$ if $m = 3$ and n is odd.*

It is easy to show that the Condorcet-rule is U -strategyproof for $U = \{u \in \mathcal{U} : u(1) - u(2) = u(2) - u(3)\}$ if $m = 3$ because the uniform lottery on all three alternatives has for every voter the expected utility of $u(2)$. Hence, the proof mainly focuses on why no other Condorcet-consistent SDS f satisfies U -strategyproofness for this set U . For this, we show that there is a profile R and a voter i such that voter i 's expected utility $\mathbb{E}[f(R)]_u$ is less than $u(2)$. Moreover, this voter can either make his second best alternative into the Condorcet winner or revert to a preference profile in which each alternative is chosen with a probability of $\frac{1}{3}$. As both cases yield an expected utility of $u(2)$ for voter i , we have found a contradiction to U -strategyproofness.

Remark 4. The Condorcet rule is also U -strategyproof for the set of equi-distant utility functions if $m = 3$ and n is even. However, other SDSs satisfy Condorcet-consistency and U -strategyproofness for even n , too. For instance, the SDS that assigns the Condorcet winner probability 1 whenever it exists and uniformly randomizes among the top-ranked alternatives otherwise satisfies also all required axioms. The proof for this claim relies on the insight that every voter has a utility of at least $u(2)$ in the absence of a Condorcet winner.

Remark 5. A well-known class of SDSs are tournament solutions which only depend on the majority relation $R_M = \{(x, y) \in A^2 : n_{xy}(R) \geq n_{yx}(R)\}$ of the input profile R to compute the outcome. For these SDSs, unanimity and u^Π -strategyproofness entail Condorcet-consistency. Thus, there are no unanimous and u^Π -strategyproof tournament solutions, regardless of the utility function u , if $m \geq 4$. This is in harsh contrast to results for set-valued social choice, where attractive tournament solutions satisfy various strategyproofness notions (see, e.g., [Brandt et al., 2016a]).

Remark 6. The proof of Theorem 3 also reveals more insights about the compatibility of k -unanimity and u^Π -strategyproofness for general SDSs. In particular, the first case shows that no $\lceil \frac{n}{3} \rceil$ -unanimous SDS can be u^Π -strategyproof for a utility function u with $u(1) - u(2) < u(2) - u(m)$ if $m \geq 4$ and $n \geq 3$.

6 Conclusion and Discussion

We study a new strategyproofness notion called U -strategyproofness. Whereas the common notion of SD -strategyproofness is derived by quantifying over all utility functions, U -strategyproofness is derived by quantifying only over the utility functions in a specified set U . This new strategyproofness notion arises from practical observations as often not all utility functions are plausible, and also has theoretical advantages because it allows for a much finer analysis than SD -strategyproofness. Furthermore, we analyze the compatibility of U -strategyproofness and decisiveness axioms such as k -unanimity and Condorcet-consistency. In particular, we discuss SDSs that satisfy k -unanimity for any k with $0 < k < n/2$ and U -strategyproofness if the set U only contains utility functions u for which $u(1) - u(2)$ is sufficiently large. Moreover, we show for rank-based SDSs that the large gap between $u(1)$ and $u(2)$ is required to be strategyproof and has to increase in k . We also prove that U -strategyproofness is incompatible with Condorcet-consistency if the set U is symmetric and $m \geq 4$. This impossibility also holds if $m = 3$ unless the utility functions in U are equi-distant. In this special case and if n is odd, the Condorcet rule can be characterized by U -strategyproofness and Condorcet-consistency.

Our results have a very intuitive interpretation: strategyproofness is only compatible with decisiveness if each voter has a clear best alternative. Even more, the more decisiveness is required, the stronger voters have to favor their best alternative. This conclusion is highlighted by Theorems 1 and 2 as well as the impossibility of Benoît [2002]. Moreover, it coincides with the informal argument that it is easier to manipulate for a voter who deems many alternatives acceptable as he can just report another acceptable alternative as his best one. Hence, our results show that the main source of manipulability are voters who are close to indifferent between some alternatives.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1. I thank the anonymous reviewers and Felix Brandt for helpful comments.

References

- [Abramowitz *et al.*, 2019] Ben Abramowitz, Elliot Anshelevich, and Whennan Zhu. Awareness of voter passion greatly improves the distortion of metric social choice. In *Proceedings of the 15th International Conference on Web and Internet Economics*, pages 3–16. Springer, 2019.
- [Aziz *et al.*, 2018] Haris Aziz, Florian Brandl, Felix Brandt, and Markus Brill. On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110:1–18, 2018. Preliminary results appeared in the Proceedings of AAAI-2013 and AAMAS-2014.
- [Barberà, 1979] Salvador Barberà. Majority and positional voting in a probabilistic framework. *Review of Economic Studies*, 46(2):379–389, 1979.
- [Benoît, 2002] Jean-Pierre Benoît. Strategic manipulation in voting games when lotteries and ties are permitted. *Journal of Economic Theory*, 102(2):421–436, 2002.
- [Bogomolnaia *et al.*, 2005] Anna Bogomolnaia, Hervé Moulin, and Richard. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- [Brandl *et al.*, 2018] Florian Brandl, Felix Brandt, Manuel Eberl, and Christian Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, 65(2):1–28, 2018. Preliminary results appeared in the Proceedings of IJCAI-2016.
- [Brandt *et al.*, 2013] Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, chapter 6, pages 213–283. MIT Press, 2nd edition, 2013.
- [Brandt *et al.*, 2016a] Felix Brandt, Markus Brill, and Paul Harrenstein. Tournament solutions. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2016.
- [Brandt *et al.*, 2016b] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Brandt, 2017] Felix Brandt. Rolling the dice: Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 1, pages 3–26. AI Access, 2017.
- [Chatterji and Zeng, 2018] Shurojit Chatterji and Huaxia Zeng. On random social choice functions with the tops-only property. *Games and Economic Behavior*, 109:413–435, 2018.
- [Chevaileyre *et al.*, 2007] Yann Chevaileyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, volume 4362 of *Lecture Notes in Computer Science (LNCS)*, pages 51–69. Springer-Verlag, 2007.
- [Duggan and Schwartz, 2000] John Duggan and Thomas Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.
- [Dutta *et al.*, 2007] Bhaskar Dutta, Hans Peters, and Arunava Sen. Strategy-proof cardinal decision schemes. *Social Choice and Welfare*, 28(1):163–179, 2007.
- [Ehlers *et al.*, 2002] Lars Ehlers, Hans Peters, and Ton Storcken. Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences. *Journal of Economic Theory*, 105(2):408–434, 2002.
- [Endriss, 2017] Ulle Endriss, editor. *Trends in Computational Social Choice*. AI Access, 2017.
- [Gibbard, 1973] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- [Gibbard, 1977] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- [Gross *et al.*, 2017] Stephen Gross, Elliot Anshelevich, and Lirong Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 544–550, 2017.
- [Hoang, 2017] Lê Nguyễn Hoang. Strategy-proofness of the randomized Condorcet voting system. *Social Choice and Welfare*, 48(3):679–701, 2017.
- [Hylland, 1980] Aanund Hylland. Strategyproofness of voting procedures with lotteries as outcomes and infinite sets of strategies. Mimeo, 1980.
- [Mennle and Seuken, 2021] Timo Mennle and Sven Seuken. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory*, 191:105–144, 2021.
- [Nandeibam, 2013] Shasikanta Nandeibam. The structure of decision schemes with cardinal preferences. *Review of Economic Design*, 17(3):205–238, 2013.
- [Procaccia and Rosenschein, 2006] Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of 10th International Workshop on Cooperative Information Agents*, pages 317–331. Springer, 2006.
- [Satterthwaite, 1975] Mark Allen Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- [Sen, 2011] Arunava Sen. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIES*, 2(4):515–527, 2011.

SUMMARY

Social decision schemes (SDSs) map the preferences of individual voters over multiple alternatives to a probability distribution over the alternatives. In order to study properties such as efficiency, strategyproofness, and participation for SDSs, preferences over alternatives are typically lifted to preferences over lotteries based on some lottery extension. In particular, the predominant lottery extension in the literature is stochastic dominance (*SD*), which is typically motivated by the fact that it ensures that voters cannot benefit from lying about their preferences if they use von Neumann-Morgenstern utility functions to compare lotteries. However, requiring strategyproofness or strict participation with respect to this preference extension only leaves room for rather undesirable SDSs. For instance, Gibbard (1977) has shown that only random dictatorships satisfy non-imposition and strong \succsim^{SD} -strategyproofness (which is subsequently called *SD*-strategyproofness). Since also the presentation of voter's preferences over lotteries based on utility functions has come under scrutiny, we will analyze a different lottery extension in this paper.

In more detail, we focus on the natural but little understood pairwise comparison (*PC*) preference extension, which postulates that one lottery is preferred to another if the former is more likely to return a preferred outcome. This preference extension has previously been used to infer positive results (e.g., Brandl et al., 2019; Brandl and Brandt, 2020) and we thus analyze whether it also leads to positive results with respect to \succsim^{PC} -strategyproofness and strict \succsim^{PC} -participation (we write *PC* rather than \succsim^{PC} in this publication). In particular we settle three open questions raised by Brandt (2017) and prove the following impossibility theorems:

- there is no Condorcet-consistent SDS that satisfies *PC*-strategyproofness,
- there is no anonymous and neutral SDS that satisfies *PC*-efficiency and *PC*-strategyproofness, and
- there is no anonymous and neutral SDS that satisfies *PC*-efficiency and strict *PC*-participation

All three impossibilities require $m \geq 4$ alternatives and we construct two new SDSs to show that they turn into possibilities when $m \leq 3$. In particular, these results demonstrate that *PC* preferences do not allow for significantly more attractive and strategyproof SDSs than *SD* preferences.

Additionally, we solve in this paper an open problem of Aziz et al. (2015a) by showing that there is a profile and a *PC*-dominated lottery such that any path of *PC*-improvements starting from this lottery does not lead to a *PC*-efficient lottery.

REFERENCE

F. Brandt, W. Suksompong, and P. Lederer. **Incentives in social decision schemes with pairwise comparison preferences.** In *Games and Economic Behavior*, 142:266–291, 2023.
DOI: <https://doi.org/10.1016/j.geb.2023.08.009>

INDIVIDUAL CONTRIBUTION

I, Patrick Lederer, am the main author of this publication. In particular, I am responsible for the joint development and conceptual design of the research project, proofs and write-up of most results (all but Theorem 3), and the joint write-up of the remaining manuscript.

COPYRIGHT AGREEMENT

The right to present this paper in a doctoral thesis has been granted by the publisher, Elsevier, in the publishing agreement presented below. There, it is stated that the article is published under the *Creative Commons Attribution 4.0 International License* (see <https://creativecommons.org/licenses/by/4.0/>). In particular, this means that I, as one of the authors, still have the copyright of the paper. Therefore, I am allowed to present the paper in this thesis. This agrees with the copyright statement of the journal available at the following website <https://beta.elsevier.com/about/policies-and-standards/copyright>.

TERMINOLOGY

In this paper, *PC*-strategyproofness refers to \succsim^{PC} -strategyproofness and *SD*-strategyproofness to strong \succsim^{SD} -strategyproofness. Moreover, voters' preferences are always strict and the paper considers a variable electorate framework (i.e., SDSs are defined for all numbers of voters).

ELSEVIER

Publishing Agreement

Elsevier Inc.

Incentives in Social Decision Schemes with Pairwise Comparison Preferences

Corresponding author	Mr. Patrick Lederer
E-mail address	ledererp@in.tum.de
Journal	Games and Economic Behavior
Our reference	YGAME3596
PII	S0899-8256(23)00121-5

Your Status

- I am one author signing on behalf of all co-authors of the manuscript

License of Publishing Rights

I hereby grant to Elsevier Inc. an irrevocable non-exclusive license to publish, distribute and otherwise use all or any part of the manuscript identified above and any tables, illustrations or other material submitted for publication as part of the manuscript (the "Article") in all forms and media (whether now known or later developed) in all languages, throughout the world, for the full term of copyright, and the right to license others to do the same, effective when the Article is accepted for publication.

I acknowledge the importance of the integrity, authenticity and permanence of the scholarly record and agree that the version of the Article that appears or will appear in the journal and embodies all value-adding publisher activities (including peer review coordination, copy-editing, formatting, (if relevant) pagination, and online enrichment) shall be the definitive final record of published research ("the Published Journal Article").

I further acknowledge and agree that nothing in this Agreement shall be deemed to permit redundant/duplicate publication of the Article in violation of publishing ethics principles, as further described below.

Supplemental Materials

"Supplemental Materials" shall mean materials published as a supplemental part of the Article, including but not limited to graphical, illustrative, video and audio material.

With respect to any Supplemental Materials that I submit, Elsevier Inc. shall have a perpetual worldwide non-exclusive right and license to publish, extract, reformat, adapt, build upon, index, redistribute, link to and otherwise use all or any part of the Supplemental Materials, in all forms and media (whether now known or later developed) and permit others to do so. The publisher shall apply the same end user license to the Supplemental Materials as to the Article where it publishes the Supplemental Materials with the Article in the journal on its online platforms on an Open Access basis.

Research Data

"Research Data" shall mean the result of observations or experimentation that validate research findings and that are published separate to the Article, which can include but are not limited to raw data, processed data, software, algorithms, protocols, and methods.

With respect to any Research Data that I wish to make accessible on a site or through a service of Elsevier Inc., Elsevier Inc. shall have a perpetual worldwide, non-exclusive right and license to publish, extract, reformat, adapt, build upon, index, redistribute, link to and otherwise use all or any part of the Research Data in all forms and media (whether now known or later developed), and to permit others to do so. Where I have selected a specific end user license under which the Research Data is to be made available on a site or through a service, the publisher shall apply that end user license to the Research Data on that site or service.

Scholarly Communication Rights

I understand that I retain the copyright in the Article and that no rights in patents, trademarks or other intellectual property rights are transferred to Elsevier Inc.. As the author of the Article, I understand that I shall have the same rights to reuse the Article as those allowed to third party users (and Elsevier Inc.) of the Article under the CC BY License.

User Rights

The publisher will apply the *Creative Commons Attribution 4.0 International License* (CC BY) to the Article where it publishes the Article in the journal on its online platforms on an Open Access basis. For further information, see <http://www.elsevier.com/about/open-access/open-access-options>.

The CC-BY license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article and to make commercial use of the Article (including reuse and/or resale of the Article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. The full details of the license are available at <http://creativecommons.org/licenses/by/4.0>.

Reversion of Rights

Articles may sometimes be accepted for publication but later rejected in the publication process, even in some cases after public posting in "Articles in Press" form, in which case all rights will revert to the author. See <https://www.elsevier.com/about/our-business/policies/article-withdrawal>.

Revisions and Addenda

I understand that no revisions, additional terms or addenda to this License Agreement can be accepted without Elsevier Inc.'s express written consent. I understand that this License Agreement supersedes any previous agreements I have entered into with Elsevier Inc. in relation to the Article from the date hereof.

Copyright Notice

The publisher shall publish and distribute the Article with the appropriate copyright notice.

I agree that any copy of the Article or any part of the Article that I distribute or post will include such copyright notice and a full citation to the Published Journal Article.

Author Representations / Ethics and Disclosure / Sanctions

I affirm the Author Representations noted below, and confirm that I have reviewed and complied with the relevant Instructions to Authors, Ethics in Publishing policy, Declarations of Interest disclosure and information for authors from countries affected by sanctions. Please note that some journals may require that all co-authors sign and submit Declarations of Interest disclosure forms. I am also aware of the publisher's policies with respect to retractions and withdrawal (<https://www.elsevier.com/about/our-business/policies/article-withdrawal>).

For further information see the publishing ethics page at <https://www.elsevier.com/about/our-business/policies/publishing-ethics> and the journal home page. For further information on sanctions, see <https://www.elsevier.com/about/our-business/policies/trade-sanctions>

Author representations

- The Article I have submitted to the journal for review is original, has been written by the stated authors and has not been previously published.
- The Article was not submitted for review to another journal while under review by this journal and will not be submitted to any other journal.

- The Article and the Supplemental Materials do not infringe any copyright, violate any other intellectual property, privacy or other rights of any person or entity, or contain any libellous or other unlawful matter.
- I have obtained written permission from copyright owners for any excerpts from copyrighted works that are included and have credited the sources in the Article or the Supplemental Materials.
- Except as expressly set out in this License Agreement, the Article is not subject to any prior rights or licenses which conflict with the terms of this License Agreement.
- If I and/or any of my co-authors reside in Russia, Belarus, Iran, Cuba, or Syria, the Article has been prepared in a personal, academic or research capacity and not as an official representative or otherwise on behalf of the relevant government or institution.
- If I am using any personal details or images of patients, research subjects or other individuals, I have obtained all consents required by applicable law and complied with the publisher's policies relating to the use of such images or personal information. See <https://www.elsevier.com/about/our-business/policies/patient-consent> for further information.
- Any software contained in the Supplemental Materials is free from viruses, contaminants or worms.
- If the Article or any of the Supplemental Materials were prepared jointly with other authors, I have informed the co-author(s) of the terms of this License Agreement and that I am signing on their behalf as their agent, and I am authorized to do so.

Governing Law and Jurisdiction

This License Agreement will be governed by and construed in accordance with the laws of the country or state of Elsevier Inc. ("the Governing State"), without regard to conflict of law principles, and the parties irrevocably consent to the exclusive jurisdiction of the courts of the Governing State.

For information on the publisher's copyright and access policies, please see <http://www.elsevier.com/copyright>.

☒ I have read and agree to the terms of the License Agreement.

14 August 2023

T-copyright license-v5/2017



Incentives in social decision schemes with pairwise comparison preferences



Felix Brandt^a, Patrick Lederer^{a,*}, Warut Suksompong^b

^a Technical University of Munich, Germany

^b National University of Singapore, Singapore

ARTICLE INFO

Article history:

Received 25 November 2022

Available online 18 August 2023

JEL classification:

D7

Keywords:

Randomized social choice

Pairwise comparison preferences

Strategyproofness

Participation

ABSTRACT

Social decision schemes (SDSs) map the ordinal preferences of voters over multiple alternatives to a probability distribution over the alternatives. To study the axiomatic properties of SDSs, we lift preferences over alternatives to preferences over lotteries using the natural—but little understood—pairwise comparison (PC) preference extension. This extension postulates that one lottery is preferred to another if the former is more likely to return a preferred outcome. We settle three open questions raised by Brandt (2017) and show that (i) no Condorcet-consistent SDS satisfies PC-strategyproofness; (ii) no anonymous and neutral SDS satisfies both PC-efficiency and PC-strategyproofness; and (iii) no anonymous and neutral SDS satisfies both PC-efficiency and strict PC-participation. We furthermore settle an open problem raised by Aziz et al. (2015) by showing that no path of PC-improvements originating from an inefficient lottery may lead to a PC-efficient lottery.

© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Incentives constitute a central aspect when designing mechanisms for multiple agents: mechanisms should incentivize agents to participate and to act truthfully. However, for many applications, guaranteeing these properties—usually called *participation* and *strategyproofness*—is a notoriously difficult task. This is particularly true for collective decision making, which studies the aggregation of preferences of multiple voters into a group decision, because strong impossibility theorems show that these axioms are in variance with other elementary properties (see, e.g., Gibbard, 1973; Satterthwaite, 1975; Moulin, 1988). For instance, the Gibbard-Satterthwaite theorem shows that every strategyproof voting rule is either dictatorial or imposing, and Moulin's No-Show paradox demonstrates that all Condorcet-consistent voting rules violate participation.

A natural escape route in light of these negative results is to allow for randomization in the output of the voting rule. Rather than returning a single winner, a *social decision scheme* (SDS) selects a lottery over the alternatives and the winner is eventually drawn at random according to the given probabilities. In order to study properties such as strategyproofness and participation as well as economic efficiency for SDSs, we need to make assumptions on how voters compare lotteries. The standard approach for this problem is to lift the voters' preferences over alternatives to preferences over lotteries by using the notion of *stochastic dominance* (SD): a voter prefers a lottery to another one if the expected utility of the former exceeds that of the latter for every utility representation consistent with his preferences over the alternatives (see, e.g., Gibbard, 1977; Bogomolnaia and Moulin, 2001; Brandl et al., 2018).

* Corresponding author.

E-mail address: ledererp@in.tum.de (P. Lederer).

Unfortunately, the negative results from deterministic social choice largely prevail when analyzing SDSs based on *SD* preferences. For instance, Gibbard (1977) has shown that the only SDS that satisfies *SD*-strategyproofness, unanimity, and anonymity is the *uniform random dictatorship (RD)*, which chooses a voter uniformly at random and returns his favorite alternative (see also Sen, 2011). Similarly, Brandt et al. (2017) have proven that Moulin's No-show paradox remains intact when defining participation based on *SD* preferences. Independently of these negative results, the representation of preferences over lotteries via expected utility functions has come under scrutiny in decision theory (e.g., Allais, 1953; Kahneman and Tversky, 1979; Machina, 1989; Anand, 2009).

As an alternative to traditional expected utility representations, some authors have proposed to postulate that an agent prefers one lottery to another if it is more likely that he prefers an alternative drawn from the former to an alternative drawn from the latter than *vice versa* (Blyth, 1972; Packard, 1982; Blavatsky, 2006). The resulting preference extension is known as *pairwise comparison (PC)* and represents a special case of Fishburn's skew-symmetric bilinear utility functions (Fishburn, 1982). Brandl et al. (2019) have shown that the No-Show paradox can be circumvented using *PC* preferences. Moreover, Brandl and Brandt (2020) proved that *PC* preferences constitute the *only* domain of preferences within a rather broad class of preferences over lotteries (including all expected utility representations) that allows for preference aggregation that satisfies independence of irrelevant alternatives and efficiency, thus avoiding Arrow's impossibility. In both cases, the resulting SDS is the set of *maximal lotteries (ML)*, which was proposed by Fishburn (1984a) and has recently attracted significant attention (Brandl et al., 2016, 2022; Peyre, 2013; Hoang, 2017).¹

Since *PC* preferences are one of the few natural preference extensions that lead to positive results, we will investigate social decision schemes based on this lottery extension. More specifically, we are interested in the question of whether there are attractive SDSs that satisfy *PC*-strategyproofness or strict *PC*-participation. The latter axiom demands that a voter is *strictly* better off participating unless he is already at maximum happiness. Unfortunately, our results are mainly negative and thus show the limitations of collective choice with *PC* preferences. In particular, we prove the following theorems, all of which settle open problems raised by Brandt (2017, p. 18).²

- There is no Condorcet-consistent SDS that satisfies *PC*-strategyproofness (Theorem 1).
- There is no anonymous and neutral SDS that satisfies *PC*-efficiency and *PC*-strategyproofness (Theorem 2).
- There is no anonymous and neutral SDS that satisfies *PC*-efficiency and strict *PC*-participation (Theorem 3).

All three theorems hold for strict preferences and require $m \geq 4$ alternatives; we show that they turn into possibilities when $m \leq 3$ by constructing two new SDSs. The second theorem strengthens Theorem 5 by Aziz et al. (2018), which shows an analogous statement for weak preferences.³

In the appendix, we furthermore settle an open problem concerning *PC*-efficiency raised by Aziz et al. (2015): we construct a preference profile and a *PC*-inefficient lottery p such that no sequence of *PC*-improvements starting from p leads to a *PC*-efficient lottery (Proposition 5).

2. The model

Let $A = \{a_1, \dots, a_m\}$ be a finite set of m alternatives and $\mathbb{N} = \{1, 2, 3, \dots\}$ an infinite set of voters. We denote by $\mathcal{F}(\mathbb{N})$ the set of all finite and non-empty subsets of \mathbb{N} . Intuitively, \mathbb{N} is the set of all potential voters, whereas $N \in \mathcal{F}(\mathbb{N})$ is a concrete electorate. Given an electorate $N \in \mathcal{F}(\mathbb{N})$, every voter $i \in N$ has a *preference relation* \succ_i , which is a complete, transitive, and anti-symmetric binary relation on A . In particular, we do not allow for ties, which only makes our impossibility theorems stronger. We write preference relations as comma-separated lists and denote the set of all preference relations by \mathcal{R} . A *preference profile* R on an electorate $N \in \mathcal{F}(\mathbb{N})$ assigns a preference relation \succ_i to every voter $i \in N$, i.e., $R \in \mathcal{R}^N$. When writing preference profiles, we subsume voters who share the same preference relation. To this end, we define $[j \dots k] = \{i \in N : j \leq i \leq k\}$ and note that $[j \dots k] = \emptyset$ if $j > k$. For instance, $[1 \dots 3] : a, b, c$ means that voters 1, 2, and 3 prefer a to b to c . We omit the brackets for singleton sets. Given a preference profile $R \in \mathcal{R}^N$, the *majority margin* between two alternatives $x, y \in A$ is $g_R(x, y) = |\{i \in N : x \succ_i y\}| - |\{i \in N : y \succ_i x\}|$, i.e., the majority margin indicates how many more voters prefer x to y than *vice versa*. Furthermore, we define $n_R(x)$ as the number of voters who prefer alternative x the most in the profile R . Next, we denote by $R_{-i} = (\succ_1, \dots, \succ_{i-1}, \succ_{i+1}, \dots, \succ_n)$ the profile derived from $R \in \mathcal{R}^N$ by removing voter $i \in N$. Finally, we define $\mathcal{R}^* = \bigcup_{N \in \mathcal{F}(\mathbb{N})} \mathcal{R}^N$ as the set of all possible preference profiles.

The focus of this paper lies on *social decision schemes (SDSs)*, which are functions that map a preference profile to a lottery over the alternatives. A *lottery* p is a probability distribution over the alternatives, i.e., a function $p : A \rightarrow [0, 1]$ such that

¹ In France, maximal lotteries have been popularized under the name *scrutin de Condorcet randomisé* (randomized Condorcet voting system).

² Brandt (2017) stated these problems without anonymity and neutrality. However, anonymity is obviously required for Theorem 2 since dictatorships satisfy all of the other axioms. Whether neutrality is required is open. Note that anonymity and neutrality are much less restrictive in randomized social choice than in the classic deterministic setting, where these properties can already be prohibitive on their own (Moulin, 1983, Theorem 1). In fact, the investigation of randomized voting rules is often motivated by fairness considerations such as anonymity and neutrality (e.g., Fishburn, 1984a; Ehlers et al., 2002; Brandt, 2017).

³ When ties are allowed in the voters' preferences, much stronger results hold: Brandl et al. (2018) have shown an analogous claim based on *SD* preferences, which implies the result by Aziz et al. (2018). Recent results (e.g., Brandl et al., 2021; Brandt et al., 2022a) hint at the fact that even stronger impossibilities may hold for weak preferences.

$p(x) \geq 0$ for all $x \in A$ and $\sum_{x \in A} p(x) = 1$. The set of all lotteries on A is denoted by $\Delta(A)$. Then, an SDS f is a function of the type $f : \mathcal{R}^* \rightarrow \Delta(A)$. We define $f(R, x)$ as the probability assigned to x by $f(R)$ and extend this notion to sets $X \subseteq A$ by letting $f(R, X) = \sum_{x \in X} f(R, x)$.

In the next sections, we introduce various desirable properties of SDSs. An overview of these axioms and their relationships is given in Fig. 1.

2.1. Fairness and decisiveness

Two basic fairness notions are anonymity and neutrality, which require that voters and alternatives are treated equally, respectively. Formally, an SDS f is *anonymous* if $f(\pi(R)) = f(R)$ for all electorates $N \in \mathcal{F}(\mathbb{N})$, preference profiles $R \in \mathcal{R}^N$, and permutations $\pi : N \rightarrow N$, where $R' = \pi(R)$ is defined by $\succ'_i = \succ_{\pi(i)}$ for all $i \in N$.

Analogously, *neutrality* requires that $f(\pi(R)) = \pi(f(R))$ for all electorates $N \in \mathcal{F}(\mathbb{N})$, preference profiles $R \in \mathcal{R}^N$, and permutations $\pi : A \rightarrow A$, i.e., $f(\pi(R))$ is equal to the distribution that, for each alternative $x \in A$, assigns probability $f(R, x)$ to alternative $\pi(x)$. Here, $R' = \pi(R)$ is the profile such that for all $i \in N$ and $x, y \in A$, $\pi(x) \succ'_i \pi(y)$ if and only if $x \succ_i y$.

A technical condition that many SDSs satisfy is cancellation. An SDS f satisfies *cancellation* if $f(R) = f(R')$ for all preference profiles $R, R' \in \mathcal{R}^*$ such that R' is derived from R by adding two voters with inverse preferences.

A natural further desideratum in randomized social choice concerns the decisiveness of SDSs: randomization should only be necessary if there is no sensible deterministic winner. This idea is, for example, captured in the notion of *unanimity*, which requires that $f(R, x) = 1$ for all profiles $R \in \mathcal{R}^*$ and alternatives $x \in A$ such that all voters in R prefer x the most.

Clearly, this condition is rather weak and there are natural strengthenings, demanding that so-called absolute winners or Condorcet winners need to be returned with probability 1. An *absolute winner* is an alternative x that is top-ranked by more than half of the voters in $R \in \mathcal{R}^N$, i.e., $n_R(x) > \frac{|N|}{2}$. The *absolute winner property* requires that $f(R, x) = 1$ for all profiles $R \in \mathcal{R}^*$ with absolute winner x .

An alternative x is a *Condorcet winner* in a profile R if $g_R(x, y) > 0$ for all $y \in A \setminus \{x\}$. *Condorcet-consistency* requires that the Condorcet winner is chosen with probability 1 whenever it exists, i.e., $f(R, x) = 1$ for all preference profiles $R \in \mathcal{R}^*$ with Condorcet winner x . Since absolute winners are Condorcet winners, Condorcet-consistency implies the absolute winner property, which in turn implies unanimity.

2.2. Preferences over lotteries

We assume that the voters' preferences over alternatives are lifted to preferences over lotteries via the pairwise comparison (PC) extension (see, e.g., Aziz et al., 2015, 2018; Brandt, 2017; Brandl and Brandt, 2020).⁴ According to this notion, a voter prefers lottery p to lottery q if the probability that p returns a better outcome than q is at least as large as the probability that q returns a better outcome than p , i.e.,

$$p \succsim^{PC} q \iff \sum_{x, y \in A: x \succ y} p(x)q(y) \geq \sum_{x, y \in A: x \succ y} q(x)p(y).$$

The relation \succsim^{PC} is complete but intransitive (a phenomenon known as the Steinhaus-Trybula paradox).⁵ An appealing interpretation of PC preferences is *ex ante* regret minimization, i.e., given two lotteries, a voter prefers the one which is less likely to result in *ex post* regret.

Despite the simple and intuitive definition, PC preferences are difficult to work with and cognitively demanding on behalf of the voters because probabilities are multiplied with each other. We therefore introduce a variant of the PC extension where one of the two lotteries under consideration has to be degenerate (i.e., it puts probability 1 on a single alternative) and any pair of non-degenerate lotteries are deemed incomparable. To this end, we define PC1 preferences as follows: a voter PC1-preferences lottery p to lottery q if $p \succsim^{PC} q$ and at least one of p and q is degenerate. Assuming that $p(x) = 1$ for some $x \in A$, this is equivalent to

$$p \succsim^{PC1} q \iff \sum_{y \in A: x \succ y} q(y) \geq \sum_{y \in A: y \succ x} q(y).$$

In other words, it only needs to be checked whether q puts at least as much probability on alternatives that are worse than x than on those that are better than x . On the one hand, PC1 preferences reduce the cognitive burden on voters when comparing lotteries and can thus be seen as a particularly plausible or realistic subset of the PC extension. On the other hand, one may also view this notion from a technical perspective as it suffices for many lottery comparisons in our

⁴ PC preferences constitute a special case of skew-symmetric bilinear utility functions (Fishburn, 1982) and have previously been considered in decision theory (Blyth, 1972; Packard, 1982; Blavatsky, 2006). Packard (1982) calls them the *rule of expected dominance* and Blavatsky (2006) refers to them as a *preference for the most probable winner*.

⁵ To see that PC preferences can be intransitive, suppose that $\succ = a, b, c, d$ and consider the lotteries p, q, r defined by $p(c) = 1$, $q(b) = \frac{3}{5}$, $q(d) = \frac{2}{5}$, and $r(a) = \frac{2}{5}$, $r(d) = \frac{3}{5}$. It can be checked that $p \succ^{PC} r \succ^{PC} q \succ^{PC} p$.

proofs involving PC preferences. Thus, this lottery extension allows us to simplify the presentation of our proofs for PC preferences and even to strengthen some of our claims, while simultaneously making the results more robust by avoiding more controversial comparisons between lotteries. It follows immediately from the definition that $p \succsim^{PC1} q$ implies that $p \succsim^{PC} q$ for all lotteries $p, q \in \Delta(A)$ and all preference relations \succ . Note that in contrast to PC preferences, $PC1$ preferences are acyclic, i.e., there is no cycle in the strict part of \succsim^{PC1} .

The most common way to compare lotteries when only ordinal preferences over alternatives are known is *stochastic dominance* (SD) (e.g., Gibbard, 1977; Bogomolnaia and Moulin, 2002; Brandl et al., 2018):

$$p \succsim^{SD} q \iff \forall x \in A: \sum_{y \in A: y \succ x} p(y) \geq \sum_{y \in A: y \succ x} q(y).$$

In contrast to PC preferences, the relation \succsim^{SD} is incomplete but transitive. Furthermore, it follows from a result by Fishburn (1984b) that $p \succsim^{SD} q$ implies $p \succsim^{PC} q$ for all preference relations \succ and all lotteries p and q (see also Aziz et al., 2015). In other words, the SD relation is a subrelation of the PC relation. We will sometimes leverage this in our proofs because SD preferences are easier to handle than PC preferences.

For each $\mathcal{X} \in \{PC, PC1, SD\}$, we say a voter *strictly \mathcal{X} -prefers* p to q , denoted by $p \succ^{\mathcal{X}} q$, if $p \succsim^{\mathcal{X}} q$ and not $q \succsim^{\mathcal{X}} p$. Note that both $p \succ^{SD} q$ and $p \succ^{PC1} q$ imply $p \succ^{PC} q$.

For a better understanding of these concepts, consider a voter with the preference relation $\succ = a, b, c, d$ and three lotteries p, q , and r with

$$p(b) = 1, \quad q(b) = q(c) = 1/2, \quad r(a) = 1/3, \text{ and } r(d) = 2/3.$$

First, observe that $p \succ^{SD} q$ since p is derived from q by moving probability from c to b . In particular, this implies that $p \succ^{PC} q$ and, since p is degenerate, also $p \succ^{PC1} q$. Next, r cannot be compared with p or q via SD . Since $r(a) < r(d)$ and $a \succ b \succ d$, it follows that $p \succ^{PC1} r$, which moreover implies that $p \succ^{PC} r$. Finally, the lotteries q and r can only be compared via PC preferences and we infer that $q \succ^{PC} r$ by checking that

$$q(b)r(d) + q(c)r(d) = 2/3 > 1/3 = r(a)q(b) + r(a)q(c).$$

2.3. Efficiency

Next, we discuss efficiency, which intuitively requires that we cannot make a voter better off without making another voter worse off. Since this axiom requires voters to compare lotteries, we define efficiency depending on some underlying lottery extension $\mathcal{X} \in \{PC, PC1, SD\}$. To formalize the intuition behind this property, we say a lottery p *\mathcal{X} -dominates* another lottery q in a profile $R \in \mathcal{R}^N$ if $p \succsim_i^{\mathcal{X}} q$ for all voters $i \in N$ and $p \succ_{i^*}^{\mathcal{X}} q$ for some voter $i^* \in N$. In this case, we also say that p is an *\mathcal{X} -improvement* of q . Less formally, p is an *\mathcal{X} -improvement* of q if p makes every voter weakly better off and at least one strictly better. A lottery p is *\mathcal{X} -efficient* in R if it is not \mathcal{X} -dominated by any other lottery. Similarly, an SDS f is *\mathcal{X} -efficient* if $f(R)$ is \mathcal{X} -efficient for all preference profiles $R \in \mathcal{R}^*$.

Since both $p \succsim_i^{SD} q$ and $p \succsim_i^{PC1} q$ imply $p \succsim_i^{PC} q$ for all voters i and lotteries p and q , it follows that a lottery that is SD -dominated or $PC1$ -dominated is also PC -dominated. Hence, for every profile R , the set of PC -efficient lotteries is contained in both the sets of $PC1$ -efficient and SD -efficient lotteries. This means that PC -efficiency implies SD -efficiency and $PC1$ -efficiency. Moreover, both PC -efficiency and SD -efficiency imply *ex post* efficiency. In order to define *ex post* efficiency, we say an alternative x *Pareto-dominates* another alternative y in a profile $R \in \mathcal{R}^N$ if $x \succ_i y$ for all voters $i \in N$. Recall here that ties in \succ_i are not allowed. *Ex post efficiency* then requires that $f(R, x) = 0$ for all profiles $R \in \mathcal{R}^*$ and alternatives $x \in A$ that are Pareto-dominated in R .

To illuminate the natural relationship between *ex post* efficiency and PC -efficiency, let us take a probabilistic view on *ex post* efficiency. First, observe that there is no alternative x that is preferred by all voters to an alternative drawn from an *ex post* efficient lottery p . Hence, for any other lottery q , the probability that q returns an outcome that is unanimously preferred to an outcome returned by p is 0, i.e., $\mathbb{P}(\forall i \in N: q \succ_i p) = 0$, where we view the lotteries p and q as random variables on A . Conversely, if p is not *ex post* efficient, it follows that $\mathbb{P}(\forall i \in N: q \succ_i p) > 0$ for the lottery q derived from p by shifting the probability from the Pareto-dominated alternatives to their dominators. Hence, a lottery p is *ex post* efficient in a profile $R \in \mathcal{R}^N$ if and only if there is no other lottery q such that

$$\mathbb{P}(\forall i \in N: q \succ_i p) > \mathbb{P}(\forall i \in N: p \succ_i q). \quad (\text{ex post efficiency})$$

From this inequality, one immediately obtains PC -efficiency by moving the quantification over the voters outside of the probability: a lottery p is PC -efficient in a profile $R \in \mathcal{R}^N$ if and only if there is no other lottery q such that

$$\forall i \in N: \mathbb{P}(q \succ_i p) \geq \mathbb{P}(p \succ_i q) \quad \wedge \quad \exists i \in N: \mathbb{P}(q \succ_i p) > \mathbb{P}(p \succ_i q). \quad (PC\text{-efficiency})$$

Despite its simple and intuitive definition, PC -efficiency is surprisingly complex and little understood. Aziz et al. (2015) prove that the set of PC -efficient lotteries is non-empty and connected, but they also provide examples showing that this

set may fail to be convex and can even be “curved” (i.e., it is not the union of a finite number of polytopes). Furthermore, they construct a preference profile with a PC -dominated lottery p that is not dominated by any PC -efficient lottery. In their example, however, one can find an intermediate lottery which PC -dominates p and which is PC -dominated by a PC -efficient lottery. Aziz et al. conclude their paper by writing that “it is an interesting open problem whether there always is a path of Pareto improvements from every $[PC]$ -dominated lottery to some $[PC]$ -undominated lottery” (Aziz et al., 2015, p. 129). In Appendix A.1, we answer this problem in the negative by providing a profile with five alternatives and eight voters, where following any sequence of PC -improvements from a certain lottery p will lead back to p and it is thus not possible to reach a PC -efficient outcome by only applying PC -improvements.

2.4. Incentive-compatibility

The final axioms we consider are strategyproofness and (strict) participation. Just like efficiency, both of these axioms can be defined for all lottery extensions; we thus define each of them for $\mathcal{X} \in \{PC, PC1, SD\}$.

Strategyproofness. Intuitively, strategyproofness demands that no voter can benefit by lying about his true preferences. Since SD and $PC1$ are incomplete, there are two different ways of defining this axiom depending on how incomparable lotteries are handled. The first option, which we call \mathcal{X} -strategyproofness, requires of an SDS f that $f(R) \succsim_i^{\mathcal{X}} f(R')$ for all electorates $N \in \mathcal{F}(\mathbb{N})$, voters $i \in N$, and preference profiles $R, R' \in \mathcal{R}^N$ with $R_{-i} = R'_{-i}$. In particular, this means that we interpret a deviation from a lottery p to another lottery q as a manipulation if p is incomparable to q with respect to \mathcal{X} .

\mathcal{X} -strategyproofness is predominant in the literature on SD preferences (e.g., Gibbard, 1977; Barberà, 1979, 2010), but it becomes very prohibitive for sparse preference relations over lotteries. For instance, not even the SDS which always returns the uniform lottery over the alternatives is $PC1$ -strategyproof because $PC1$ cannot compare the uniform lottery to itself. For such preferences, the notion of weak \mathcal{X} -strategyproofness is more sensible: an SDS f is *weakly \mathcal{X} -strategyproof* if $f(R') \not\succ_i^{\mathcal{X}} f(R)$ for all electorates $N \in \mathcal{F}(\mathbb{N})$, voters $i \in N$, and preference profiles $R, R' \in \mathcal{R}^N$ with $R_{-i} = R'_{-i}$. In other words, weak strategyproofness requires that no voter can obtain a strictly \mathcal{X} -preferred outcome by lying about his true preferences.⁶ Note that, since PC preferences are complete, PC -strategyproofness coincides with weak PC -strategyproofness. By contrast, for SD and $PC1$, weak \mathcal{X} -strategyproofness is significantly less demanding than \mathcal{X} -strategyproofness. We say an SDS is \mathcal{X} -manipulable if it is not \mathcal{X} -strategyproof and *strongly \mathcal{X} -manipulable* if it is not weakly \mathcal{X} -strategyproof. Furthermore, since strategyproofness does not require a variable electorate, we always specify the electorates for which we show that an SDS is strategyproof or manipulable.

Participation. Participation axioms intuitively require that voters should not be able to benefit by abstaining from the election. Analogous to strategyproofness, one could formalize this condition in two ways depending on how incomparabilities between lotteries are interpreted.⁷ Nevertheless, we will focus in our results only on the strong notion and thus say that an SDS f satisfies \mathcal{X} -participation if $f(R) \succsim_i^{\mathcal{X}} f(R_{-i})$ for all electorates $N \in \mathcal{F}(\mathbb{N})$, voters $i \in N$, and preference profiles $R \in \mathcal{R}^N$.

In this paper, we are mainly interested in *strict \mathcal{X} -participation*, as introduced by Brandt et al. (2015), which demands of an SDS f that, for all $N \in \mathcal{F}(\mathbb{N})$, $i \in N$, and $R \in \mathcal{R}^N$, it holds that $f(R) \succsim_i^{\mathcal{X}} f(R_{-i})$ and, moreover, $f(R) \succ_i^{\mathcal{X}} f(R_{-i})$ if there is a lottery p with $p \succ_i^{\mathcal{X}} f(R_{-i})$. That is, whenever possible, a voter is *strictly* better off by voting than by abstaining from an election.

Since both $p \succsim_i^{SD} q$ and $p \succsim_i^{PC1} q$ imply $p \succsim_i^{PC} q$, the concepts above are related for PC , $PC1$, and SD : SD -strategyproofness implies PC -strategyproofness which implies weak $PC1$ -strategyproofness. Furthermore, strict SD -participation is stronger than strict PC -participation, which obviously entails PC -participation (cf. Brandt, 2017). An overview of these relationships is given in Fig. 1.

3. Random dictatorship and maximal lotteries

The following two important SDSs are useful for putting our results into perspective: the uniform random dictatorship (RD) and maximal lotteries (ML). These SDSs are well-known and most of the subsequent claims are taken from the survey by Brandt (2017). The *uniform random dictatorship* (RD) assigns probabilities proportional to $n_R(x)$, i.e., $RD(R, x) = \frac{n_R(x)}{\sum_{y \in A} n_R(y)}$ for every alternative $x \in A$ and preference profile $R \in \mathcal{R}^*$. More intuitively, RD chooses a voter uniformly at random and returns his favorite alternative as the winner. This SDS satisfies strong incentive axioms, but fails efficiency and decisiveness conditions.

⁶ In the literature, \mathcal{X} -strategyproofness is sometimes called strong \mathcal{X} -strategyproofness, and weak \mathcal{X} -strategyproofness is then called \mathcal{X} -strategyproofness. This is for instance the case in the survey by Brandt (2017).

⁷ As with strategyproofness, both versions are equivalent for PC because PC preferences are complete.




Fig. 1. Overview of results. An arrow from an axiom X to another axiom Y indicates that X implies Y . The thick lines between axioms represent impossibility theorems. Note that Theorems 2 and 3 additionally require anonymity and neutrality. Axioms labeled with *ML* are satisfied by maximal lotteries, and axioms labeled with *RD* are satisfied by the uniform random dictatorship.

Proposition 1. *RD satisfies SD-strategyproofness, strict SD-participation, and SD-efficiency, but fails PC1-efficiency and the absolute winner property.*

Clearly, since *SD-strategyproofness* and *strict SD-participation* imply the corresponding concepts for *PC*, *RD* satisfies these incentive axioms also for *PC* preferences. When additionally requiring anonymity, *RD* is the only SDS that satisfies *SD-strategyproofness* and *SD-efficiency* (Gibbard, 1977). On the other hand, a result by Benoît (2002) implies that no *SD-strategyproof* SDS can satisfy the absolute winner property. For *RD*, this claim as well as its failure of *PC1-efficiency* can be observed in the following profile.

$$R: \quad [1 \dots 3]: a, b, c \quad 4: b, a, c \quad 5: c, a, b$$

For this profile, $RD(R, a) = \frac{3}{5}$ and $RD(R, b) = RD(R, c) = \frac{1}{5}$, but a is the absolute winner and the lottery that puts probability 1 on a *PC1*-dominates $RD(R)$.

The set of maximal lotteries for profile R is defined as

$$ML(R) = \{p \in \Delta(A): \sum_{x, y \in A} p(x)q(y)g_R(x, y) \geq 0 \text{ for all } q \in \Delta(A)\}.$$

$ML(R)$ is non-empty by the minimax theorem and almost always a singleton. In particular, if the number of voters is odd, there is always a unique maximal lottery. In case of multiple maximal lotteries, the claim below that *ML* satisfies weak *PC1-strategyproofness* requires a mild tie-breaking assumption: a degenerate lottery may only be returned if it is the unique maximal lottery. For all other claims, ties can be broken arbitrarily. As the next proposition shows, *ML* satisfies strong efficiency and decisiveness notions but is rather manipulable.

Proposition 2. *ML satisfies PC-efficiency, PC-participation, Condorcet-consistency, and weak PC1-strategyproofness, but fails PC-strategyproofness and strict PC-participation.*

References for all claims except the one concerning *strict PC-participation* are given by Brandl et al. (2022). The failure of *strict PC-participation* is straightforward because *ML* is Condorcet-consistent and a voter may be unable to change the Condorcet winner by joining the electorate. Brandl et al. (2022) show that *ML* is *PC*-manipulable in most profiles that admit no weak Condorcet winner. This is, for example, the case in the profiles R and R' below, where Voter 4 can *PC*-manipulate by deviating from R to R' .

$$\begin{array}{lll} R: & \{1, 2\}: a, b, c & \{3, 4\}: b, c, a \quad 5: c, a, b \\ R': & \{1, 2\}: a, b, c & 3: b, c, a \quad \{4, 5\}: c, a, b \end{array}$$

The unique maximal lotteries in R and R' , respectively, are p and q with $p(a) = q(c) = \frac{3}{5}$ and $p(b) = p(c) = q(a) = q(b) = \frac{1}{5}$. Since $\succ_i = \succ'_i$ for all $i \in \{1, 2, 3, 5\}$ and $q \succ_4^{PC} p$, Voter 4 can *PC*-manipulate by deviating from R to R' .

In contrast to its relatively poor performance in terms of strategyproofness and participation, *ML* is very efficient. In fact, maximality of lotteries can be seen as an efficiency notion itself. To this end, note that a lottery p is in $ML(R)$ if and only if there is no lottery q such that a voter that is uniformly drawn from N is more likely to prefer an outcome drawn from q to an outcome drawn from p than vice versa (Brandl and Brandt, 2020). More formally, let I denote a uniformly distributed random variable on the voters, and interpret p and q as independent random variables on the alternatives. Then, $p \in ML(R)$ if and only if there is no lottery q such that

$$\mathbb{P}(q \succ_I p) > \mathbb{P}(p \succ_I q). \quad (\text{maximality})$$

This condition is equivalent to the definition of *ML* because

$$\mathbb{P}(q \succ_I p) = \sum_{x, y \in A} q(x)p(y) \frac{|\{i \in N: x \succ_i y\}|}{|N|}$$

and thus $\mathbb{P}(q \succ_I p) > \mathbb{P}(p \succ_I q)$ if and only if $\sum_{x,y \in A} q(x)p(y)g_R(x,y) > 0$. When comparing this to the definitions of *PC*-efficiency and *ex post* efficiency given in Section 2.3, it can be seen that the “only” difference between maximal lotteries and *PC*-efficient lotteries is that for maximal lotteries a voter is uniformly drawn at random while for *PC*-efficiency the inequality has to hold for all voters. This immediately implies that *ML* satisfies *PC*-efficiency.

4. Results

We are now ready to present our results. The results for *PC*-strategyproofness are given in Section 4.1 while those for strict *PC*-participation are given in Section 4.2. For the sake of readability, we defer all lengthy proofs to the appendix.

4.1. *PC*-strategyproofness

In this section, we show that every Condorcet-consistent and every anonymous, neutral, and *PC*-efficient SDS is *PC*-manipulable when there are $m \geq 4$ alternatives. These results show that no SDS simultaneously satisfies *PC*-strategyproofness and some of the desirable properties of maximal lotteries. Moreover, since *PC*-strategyproofness is weaker than *SD*-strategyproofness, the incompatibility of *PC*-strategyproofness and Condorcet-consistency is a strengthening of the well-known incompatibility of Condorcet-consistent and *SD*-strategyproof SDSs (see, e.g., Brandt et al., 2022b). The second result is somewhat surprising: while anonymity, neutrality, *SD*-strategyproofness, and *SD*-efficiency characterize the uniform random dictatorship, the axioms become incompatible when moving from *SD* to *PC*. Both impossibilities require $m \geq 4$ alternatives and we show that they turn into possibilities when $m \leq 3$.

Theorem 1. Every Condorcet-consistent SDS is *PC*-manipulable if $|N| \geq 5$ is odd and $m \geq 4$.

Proof. Assume for contradiction that there is a Condorcet-consistent and *PC*-strategyproof SDS f for $m \geq 4$ alternatives. Subsequently, we focus on the electorate $N = \{1, \dots, 5\}$ because we can generalize the result to any larger electorate with an odd number of voters by adding pairs of voters with inverse preferences. These voters do not affect the Condorcet winner and hence will not affect our analysis.

As the first step, consider the profiles R^1 to R^4 . The $*$ symbol is a placeholder for all missing alternatives.

R^1 :	1: $a, b, d, c, *$	2: $d, b, a, c, *$	3: $a, d, c, b, *$	4: $*, c, d, b, a$	5: $c, b, a, d, *$
R^2 :	1: $a, b, d, c, *$	2: $d, b, a, c, *$	3: $a, b, c, d, *$	4: $*, c, d, b, a$	5: $c, b, a, d, *$
R^3 :	1: $a, b, d, c, *$	2: $d, b, a, c, *$	3: $a, d, c, b, *$	4: $*, c, d, a, b$	5: $c, b, a, d, *$
R^4 :	1: $a, b, d, c, *$	2: $d, b, a, c, *$	3: $d, a, c, b, *$	4: $*, c, d, b, a$	5: $c, b, a, d, *$

Note that b is the Condorcet winner in R^2 , a in R^3 , and d in R^4 . Thus, Condorcet-consistency entails that $f(R^2, b) = f(R^3, a) = f(R^4, d) = 1$. By contrast, there is no Condorcet winner in R^1 and we use *PC*-strategyproofness to derive $f(R^1)$. In more detail, since $f(R^2)$, $f(R^3)$, and $f(R^4)$ are degenerate, it suffices to use weak *PC1*-strategyproofness. This axiom implies that $\sum_{x \in A: x \succ_{\frac{2}{3}} b} f(R^1, x) \leq \sum_{x \in A: b \succ_{\frac{2}{3}} x} f(R^1, x)$, as otherwise Voter 3 can *PC*-manipulate by deviating from R^2 to R^1 . Equivalently, this means that

$$f(R^1, a) \leq f(R^1, A \setminus \{a, b\}). \quad (1)$$

Analogously, weak *PC1*-strategyproofness between R^3 and R^1 and between R^1 and R^4 entails the following inequalities because Voter 4 in R^3 needs to *PC1*-prefer $f(R^3)$ to $f(R^1)$ and Voter 3 in R^1 needs to *PC1*-prefer $f(R^1)$ to $f(R^4)$.

$$f(R^1, A \setminus \{a, b\}) \leq f(R^1, b) \quad (2)$$

$$f(R^1, A \setminus \{a, d\}) \leq f(R^1, a) \quad (3)$$

Chaining the inequalities together, we get $f(R^1, A \setminus \{a, d\}) \leq f(R^1, a) \leq f(R^1, A \setminus \{a, b\}) \leq f(R^1, b)$, so $f(R^1, A \setminus \{a, b, d\}) = 0$. Simplifying (1), (2), and (3) then results in $f(R^1, a) \leq f(R^1, d) \leq f(R^1, b) \leq f(R^1, a)$, so $f(R^1, a) = f(R^1, b) = f(R^1, d) = \frac{1}{3}$.

Next, we analyze the profiles R^5 to R^8 .

R^5 :	1: $a, b, d, c, *$	2: $b, d, a, c, *$	3: $a, d, c, b, *$	4: $*, c, d, b, a$	5: $c, b, a, d, *$
R^6 :	1: $a, b, d, c, *$	2: $b, d, a, c, *$	3: $a, d, c, b, *$	4: $*, c, d, b, a$	5: $b, c, a, d, *$
R^7 :	1: $a, b, d, c, *$	2: $b, d, a, c, *$	3: $a, d, c, b, *$	4: $*, c, d, a, b$	5: $c, b, a, d, *$
R^8 :	1: $a, b, d, c, *$	2: $b, c, d, a, *$	3: $a, d, c, b, *$	4: $*, c, d, b, a$	5: $c, b, a, d, *$

Just as for the profiles R^1 to R^4 , there is no Condorcet winner in R^5 , whereas b is the Condorcet winner in R^6 , a in R^7 , and c in R^8 . Consequently, Condorcet-consistency requires that $f(R^6, b) = f(R^7, a) = f(R^8, c) = 1$. Next, we use again weak *PC1*-strategyproofness to derive $f(R^5)$. In particular, we infer the following inequalities as Voter 5 in R^5 needs to *PC1*-prefer $f(R^5)$ to $f(R^6)$, Voter 4 in R^7 needs to *PC1*-prefer $f(R^7)$ to $f(R^5)$, and Voter 2 in R^8 needs to *PC1*-prefer $f(R^8)$ to $f(R^5)$.

$$f(R^5, A \setminus \{b, c\}) \leq f(R^5, c) \quad (4)$$

$$f(R^5, A \setminus \{a, b\}) \leq f(R^5, b) \quad (5)$$

$$f(R^5, b) \leq f(R^5, A \setminus \{b, c\}) \quad (6)$$

Analogous computations as for R^1 now show that $f(R^5, a) = f(R^5, b) = f(R^5, c) = \frac{1}{3}$. Finally, note that R^1 and R^5 only differ in the preferences of Voter 2. This means that Voter 2 can *PC*-manipulate by deviating from R^5 to R^1 since he even *SD*-prefers $f(R^1)$ to $f(R^5)$. Hence, f fails *PC*-strategyproofness, which contradicts our assumptions. \square

Before proving the incompatibility of *PC*-efficiency and *PC*-strategyproofness, we first state an auxiliary claim which establishes that the absolute winner property, *PC*-efficiency, and *PC*-strategyproofness are incompatible. The involved proof of this lemma is deferred to Appendix A.2.

Lemma 1. *Every PC-efficient SDS that satisfies the absolute winner property is PC-manipulable if $|N| \geq 3$, $|N| \notin \{4, 6\}$, and $m \geq 4$.*

Note that Lemma 1 is a rather strong impossibility itself and, in particular, does not require anonymity or neutrality. Based on this lemma, we now show that every anonymous, neutral, and *PC*-efficient SDS is *PC*-manipulable. It is sufficient to show that the given axioms imply the absolute winner property since the result then follows from Lemma 1.

Theorem 2. *Every anonymous and neutral SDS that satisfies PC-efficiency is PC-manipulable if $|N| \geq 3$, $|N| \notin \{4, 6\}$, and $m \geq 4$.*

Proof. We prove the claim for even $|N|$; the argument for odd $|N|$ is much more involved and deferred to the appendix. Our goal is to show that every SDS that satisfies *PC*-efficiency, *PC*-strategyproofness, anonymity, and neutrality for an electorate with an even number of voters $|N| \geq 8$ also satisfies the absolute winner property. Then, Lemma 1 shows that no such SDS exists. To this end, suppose that there is an SDS f that satisfies all given axioms and consider the following profile R^1 , where the $*$ symbol indicates that all missing alternatives are added in an arbitrary fixed order.

$$R^1: \quad 1: a, b, c, * \quad 2: a, c, b, * \quad [3 \dots \frac{n}{2} + 1]: b, a, c, * \quad [\frac{n}{2} + 2 \dots n]: c, a, b, *$$

First, all alternatives except a , b , and c are Pareto-dominated and *PC*-efficiency thus requires that $f(R^1, x) = 0$ for $x \notin \{a, b, c\}$. Moreover, b and c are symmetric in R^1 and anonymity and neutrality therefore imply that $f(R^1, b) = f(R^1, c)$. Finally, note that every lottery p with $p(b) = p(c) > 0$ is *PC*-dominated by the lottery q with $q(a) = 1$. Hence, it follows from *PC*-efficiency, anonymity, and neutrality that $f(R^1, a) = 1$.

Next, consider the profile R^2 , in which the voters in $[3 \dots \frac{n}{2} + 1]$ report a, c, b instead of b, a, c , and Voter 1 reports a, c, b instead of a, b, c .

$$R^2: \quad 1: a, c, b, * \quad 2: a, c, b, * \quad [3 \dots \frac{n}{2} + 1]: a, c, b, * \quad [\frac{n}{2} + 2 \dots n]: c, a, b, *$$

A repeated application of *PC*-strategyproofness shows that a must still be chosen with probability 1 in R^2 because a is the favorite alternative of the deviator after every step. In more detail, if a is assigned probability 1 in profile R and a voter deviates to a profile R' by top-ranking a , a must still have probability 1 because otherwise, the voter can manipulate in R' by going back to R .

Finally, observe that c Pareto-dominates all alternatives but a in R^2 . Using this fact, we go to the profile R^3 by letting the voters in $[\frac{n}{2} + 2 \dots n]$ one after another change their preference relation to $c, b, *, a$.

$$R^3: \quad 1: a, c, b, * \quad 2: a, c, b, * \quad [3 \dots \frac{n}{2} + 1]: a, c, b, * \quad [\frac{n}{2} + 2 \dots n]: c, b, *, a$$

We claim that $f(R^3, a) = 1$. Indeed, *PC*-efficiency shows for R^3 and all intermediate profiles that only a and c can have positive probability as all other alternatives are Pareto-dominated. Moreover, *PC*-strategyproofness shows for every step that, if a is originally chosen with probability 1, then c must have probability 0 after the manipulation because every manipulator prefers c to a and no other alternative gets any positive probability. Hence, we infer that $f(R^3, a) = 1$.

Finally, note that the voters who top-rank a can now reorder the alternatives in $A \setminus \{a\}$ arbitrarily and the voters who bottom-rank a can even reorder all alternatives without affecting the outcome. In more detail, if a is chosen with probability 1 and a voter top-ranks a after the manipulation, *PC*-strategyproofness requires that a still is assigned probability 1 because the voter can otherwise manipulate by switching back to his original preference relation. Similarly, if a voter bottom-ranks alternative a and a is assigned probability 1, he cannot affect the outcome by deviating because any other outcome induces a *PC*-manipulation. Hence, it follows that $f(R, a) = 1$ for all profiles in which the voters in $[1 \dots \frac{n}{2} + 1]$ top-rank a . Since anonymity allows us to rename the voters and neutrality to exchange the alternatives, this means that f satisfies the absolute winner property. \square

Since both Theorems 1 and 2 require $m \geq 4$ alternatives, there is still hope for a positive result when $m \leq 3$. Indeed, for $m = 2$, *ML* satisfies Condorcet-consistency, *PC*-efficiency, *PC*-strategyproofness, anonymity, and neutrality. However, as shown in Section 3, *ML* fails *PC*-strategyproofness when $m = 3$. We therefore construct another SDS that satisfies all given

axioms. To this end, let $CW(R)$ be the set of Condorcet winners in R , and $WCW(R) = \{x \in A : g_R(x, y) \geq 0 \text{ for all } y \in A \setminus \{x\}\}$ the set of weak Condorcet winners if $CW(R) = \emptyset$, and $WCW(R) = \emptyset$ otherwise. Then, define the SDS f^1 as follows.

$$f^1(R) = \begin{cases} [x : 1] & \text{if } CW(R) = \{x\} \\ [x : \frac{1}{2}; y : \frac{1}{2}] & \text{if } WCW(R) = \{x, y\} \\ [x : \frac{3}{5}; y : \frac{1}{5}; z : \frac{1}{5}] & \text{if } WCW(R) = \{x\} \\ [x : \frac{1}{3}; y : \frac{1}{3}; z : \frac{1}{3}] & \text{otherwise} \end{cases}$$

Note that, in the absence of majority ties, f^1 boils down to the rather natural SDS that selects the Condorcet winner with probability 1 and returns the uniform lottery otherwise. This SDS was already proposed by Potthoff (1970) to achieve strategyproofness in the case of three alternatives. As we show, f^1 extends this SDS to profiles with majority ties while preserving a number of desirable properties. In particular, f^1 is the *only* SDS for $m = 3$ alternatives that satisfies cancellation and the axioms of Theorem 2. We defer the proof of this claim to Appendix A.3. Moreover, f^1 is clearly Condorcet-consistent.

Proposition 3. For $m = 3$, f^1 is the only anonymous and neutral SDS that satisfies PC-efficiency, PC-strategyproofness, and cancellation.

Remark 1. All axioms are required for Theorem 1 and all axioms with the possible exception of neutrality are required for Theorem 2. *ML* only fails PC-strategyproofness, dictatorships only fail anonymity and Condorcet-consistency, and the uniform random dictatorship only fails PC-efficiency and Condorcet-consistency. In particular, since dictatorships satisfy all axioms of Theorem 2 but anonymity, this condition is required for the impossibility. The number of alternatives required for Theorems 1 and 2 is tight as shown by f^1 . We conjecture that Theorem 2 holds even without neutrality.

Remark 2. It is open whether Theorem 1 also holds for even $|N|$. However, when additionally requiring the mild condition of homogeneity (which requires that splitting each voter into k clones with the same preferences does not affect the outcome), the statement holds also for even $|N| \geq 10$. This can be shown by duplicating all voters in the proof and adding some additional profiles in the derivation from R^1 to R^5 .

Remark 3. For the well-known class of pairwise SDSs, which only depend on the majority margins to compute the outcomes, PC-strategyproofness, unanimity, and homogeneity imply Condorcet-consistency. This follows by carefully inspecting the proof of Lemma 12 by Brandt and Lederer (2023). Hence, Theorem 1 implies that no pairwise SDS satisfies unanimity, PC-strategyproofness, and homogeneity.

Remark 4. Remarkably, the proof of Theorem 1 never uses the full power of PC-strategyproofness. Instead, every step either uses weak *PC1*-strategyproofness or weak *SD*-strategyproofness and thus, our proof shows actually a stronger but more technical result where PC-strategyproofness is replaced with weak *SD*-strategyproofness and weak *PC1*-strategyproofness. Interestingly, the Condorcet rule, which chooses the Condorcet winner if it exists and randomizes uniformly over all alternatives otherwise, is Condorcet-consistent and weakly *SD*-strategyproof, and *ML* is Condorcet-consistent and weakly *PC1*-strategyproof.

4.2. Strict PC-participation

In this section, we show that strict PC-participation is incompatible with PC-efficiency.

Theorem 3. No anonymous and neutral SDS satisfies both PC-efficiency and strict PC-participation if $m \geq 4$.

Proof. We establish a stronger statement using *PC1*-efficiency instead of PC-efficiency.

Assume for contradiction that there is a neutral and anonymous SDS f that satisfies both *PC1*-efficiency and strict PC-participation. We first prove the impossibility for the case $m = 4$ and explain how to generalize the impossibility to more alternatives at the end of this proof. First, consider the following profile with ten voters.

$$\begin{array}{lllll} R^1: & 1: a, b, c, d & 2: a, b, d, c & 3: a, c, b, d & 4: a, c, d, b & 5: a, d, b, c \\ & 6: a, d, c, b & 7: b, a, c, d & 8: b, a, d, c & 9: c, a, b, d & 10: c, a, d, b \end{array}$$

Observe that b and c are symmetric in R^1 and thus, anonymity and neutrality imply that $f(R^1, b) = f(R^1, c)$. Moreover, since a Pareto-dominates d , it can be checked that every lottery p with $p(b) = p(c) > 0$ or $p(d) > 0$ is *PC1*-dominated by the lottery q that puts probability 1 on a . Indeed, voters 1 to 6 strictly *PC1*-prefer q to p since a is their favorite alternative, and voters 7 to 10 *PC1*-prefer q to p since $p(b) = p(c)$. Hence, *PC1*-efficiency requires that $f(R^1, b) = f(R^1, c) = f(R^1, d) = 0$, which means that $f(R^1, a) = 1$.

Next, consider profile R^2 , which is obtained by adding voter 11 with the preference relation d, a, b, c to R^1 . We infer from strict PC-participation that $f(R^2, d) > f(R^2, b) + f(R^2, c)$.

Finally, consider profile R^3 , which is obtained by adding voter 12 with the preference relation d, a, c, b to R^2 . Observe that b, c , and d are symmetric in R^3 , so by neutrality and anonymity, $f(R^3, b) = f(R^3, c) = f(R^3, d)$. If $f(R^3, b) = f(R^3, c) = f(R^3, d) > 0$, then f is not $PC1$ -efficient because all voters strictly prefer the degenerate lottery that puts probability 1 on a . Hence, $f(R^3, b) = f(R^3, c) = f(R^3, d) = 0$, which means that $f(R^3, a) = 1$. Since $f(R^2, d) > f(R^2, b) + f(R^2, c)$, voter 12 has a disincentive to participate in R^3 , thereby contradicting the strict PC -participation of f .

Finally, for extending this impossibility to more than $m = 4$ alternatives, we simply add the new alternatives at the bottom of the preference rankings of all voters in a fixed order. $PC1$ -efficiency, anonymity, and neutrality still require for R^1 and R^3 that a obtains probability 1, and it thus is easy to check that the impossibility still holds. \square

By contrast, multiple SDSs are known to satisfy SD -efficiency and strict SD -participation (see Brandl et al., 2015). Theorem 3 can be seen as a complement to the work of Brandl et al. (2019): ML satisfies PC -participation and PC -efficiency, but no anonymous and neutral SDS satisfies strict PC -participation and PC -efficiency.

Since Theorem 3 requires $m \geq 4$, a natural question is whether the impossibility also holds for $m \leq 3$. As we demonstrate, this is not the case. If $m = 2$, it is easy to see that the uniform random dictatorship satisfies all axioms of Theorem 3. For $m = 3$, however, the uniform random dictatorship fails $PC1$ -efficiency (see Section 3). In light of this, we construct a new SDS that satisfies all axioms used in Theorem 3 and PC -efficiency. To this end, let B denote the set of alternatives that are never bottom-ranked. Then, the SDS f^2 is defined as follows: return the uniform random dictatorship if $|B| \in \{0, 2\}$; otherwise (i.e., $|B| = 1$), we delete the alternatives $x \in A \setminus B$ that minimize $n_R(x)$ (if there is a tie, delete both alternatives) and return the outcome of the uniform random dictatorship for the reduced profile. As the following proposition shows, f^2 indeed satisfies all axioms of Theorem 3 when $m = 3$; the proof is deferred to Appendix A.3.

Proposition 4. For $m = 3$, f^2 satisfies anonymity, neutrality, PC -efficiency, and strict PC -participation.

Remark 5. Both PC -efficiency and PC -participation are required for Theorem 3 since ML and RD satisfy all but one of the axioms. Whether anonymity and neutrality are required is open.

Remark 6. Theorem 3 still holds when replacing PC -efficiency with weak PC -efficiency and letting $m \geq 5$. (A lottery fails to be weakly PC -efficient if there is another lottery in which all voters are strictly better off.) The incompatibility with strict PC -participation can be shown by a proof that uses a preference profile with five alternatives and 18 voters that are joined by 6 further voters, but is otherwise similar to that of Theorem 3.

Remark 7. f^2 also satisfies strict SD -participation, which shows that efficiency and participation are compatible even in their strongest forms when $m = 3$. However, these axioms do not uniquely characterize f^2 .

5. Conclusion

We have studied incentive properties of social decision schemes (SDSs) based on the pairwise comparison (PC) lottery extension and answered open questions raised by Brandt (2017) under the assumptions of anonymity and neutrality. In particular, we showed that PC -strategyproofness and strict PC -participation are incompatible with PC -efficiency and Condorcet-consistency when there are at least four alternatives (see also Fig. 1). When there are fewer than four alternatives, the axioms are shown to be compatible via the introduction of two new SDSs. We also settled an open problem by Aziz et al. (2015) by showing that there exist profiles and PC -inefficient lotteries such that it is not possible to reach a PC -efficient outcome by repeatedly moving from a PC -dominated lottery to one of its dominators.

We highlight three important aspects and consequences of our results. First, when moving from the standard approach of stochastic dominance (SD) to PC , previously compatible axioms become incompatible. In particular, Theorems 2 and 3 become possibilities when using SD preferences since all given axioms are satisfied by the uniform random dictatorship. Secondly, unlike Arrow's impossibility and the No-Show paradox, our results show that PC preferences offer no attractive escape route to the Gibbard-Satterthwaite impossibility. In more detail, while Brandl et al. (2019) and Brandl and Brandt (2020) show that maximal lotteries avoid Arrow's impossibility and the No-Show paradox when assuming PC preferences, our results show that PC preferences do not allow to find a new attractive strategyproof SDS. Hence, our theorems stand in contrast to the previous positive findings on PC preferences. In light of the shown tradeoff between incentive-compatibility and efficiency, two SDSs stand out: the uniform random dictatorship because it satisfies PC -strategyproofness and strict PC -participation, and maximal lotteries because it satisfies $PC1$ -strategyproofness, Condorcet-consistency, PC -efficiency, and PC -participation.

There are only few opportunities to further strengthen these results. In some cases, it is unclear whether anonymity or neutrality—two fairness properties that are often considered imperative in social choice—are required. Two challenging questions concerning Theorem 2 are whether anonymity can be weakened to non-dictatorship and whether PC -efficiency can be replaced with weak PC -efficiency (cf. Remark 6). However, if true, any such statement would require quite different proof techniques.

Declaration of competing interest

None.

Data availability

No data was used for the research described in the article.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-2 and BR 2312/12-1, by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001, and by an NUS Start-up Grant. A preliminary version of this paper appeared in the Proceedings of the 31st International Joint Conference on Artificial Intelligence (Vienna, Austria, July 2022). We thank the anonymous referees for their helpful feedback.

Appendix A. Omitted proofs

In this section, we present the proofs omitted from the main body. In particular, we prove in Appendix A.1 that it is possible that no path of *PC*-improvements originating from a *PC*-inefficient lottery leads to a *PC*-efficient lottery. Furthermore, we discuss the proofs of Theorem 2 and Lemma 1 in Appendix A.2. Finally, Appendix A.3 contains the proofs of Propositions 3 and 4.

A.1. *PC*-efficiency cycle

Aziz et al. (2015) left as an open problem whether for every *PC*-inefficient lottery p there is a sequence of *PC*-improvements that leads to a *PC*-efficient lottery. We disprove this assertion by giving a profile R and a lottery p such that it is not possible to reach a *PC*-efficient lottery q from p by repeatedly applying *PC*-improvements.

Proposition 5. *There is a profile R and a lottery p such that no sequence of *PC*-improvements that starts at p leads to a *PC*-efficient lottery.*

Proof. Before we present the profile and the lotteries that will constitute our counterexample, we discuss two general claims on *PC* preferences. To this end, consider two lotteries p and q and two voters i and j and suppose that both voters i and j prefer q to p according to *PC*. By using the definition of *PC* preferences and partitioning alternatives also with respect to the preferences of voter j , we derive the following inequality for voter i .

$$\sum_{x,y \in A: x \succ_i y \wedge x \succ_j y} q(x)p(y) + \sum_{x,y \in A: x \succ_i y \wedge y \succ_j x} q(x)p(y) \geq \sum_{x,y \in A: x \succ_i y \wedge x \succ_j y} p(x)q(y) + \sum_{x,y \in A: x \succ_i y \wedge y \succ_j x} p(x)q(y)$$

By exchanging the roles of voter i and j , we also infer the subsequent inequality.

$$\sum_{x,y \in A: x \succ_j y \wedge x \succ_i y} q(x)p(y) + \sum_{x,y \in A: x \succ_j y \wedge y \succ_i x} q(x)p(y) \geq \sum_{x,y \in A: x \succ_j y \wedge x \succ_i y} p(x)q(y) + \sum_{x,y \in A: x \succ_j y \wedge y \succ_i x} p(x)q(y)$$

Now, by summing up these two inequalities and cancelling common terms, we infer our first key insight: two voters i and j simultaneously *PC*-prefer q to p only if

$$\sum_{x,y \in A: x \succ_i y \wedge x \succ_j y} q(x)p(y) \geq \sum_{x,y \in A: x \succ_i y \wedge x \succ_j y} p(x)q(y). \quad (7)$$

On the other hand, if $\sum_{x,y \in A: x \succ_i y \wedge x \succ_j y} q(x)p(y) = \sum_{x,y \in A: x \succ_i y \wedge x \succ_j y} p(x)q(y) = 0$, our two initial inequalities simplify to

$$\sum_{x,y \in A: x \succ_i y \wedge y \succ_j x} q(x)p(y) = \sum_{x,y \in A: x \succ_i y \wedge y \succ_j x} p(x)q(y). \quad (8)$$

Using (7) and (8), we will show that we cannot reach a *PC*-efficient lottery from p^1 where $p^1(a) = p^1(b) = \frac{1}{2}$ by only making *PC*-improvements according to the profile R shown below.

$R:$	1: b, d, c, a, e	2: a, e, c, b, d	3: d, c, a, b, e	4: e, c, a, b, d
	5: b, d, e, c, a	6: a, e, d, c, b	7: b, e, d, c, a	8: a, d, e, c, b

To prove this claim, we proceed in three steps, which essentially show that from p^1 we can only go towards the lottery p^2 with $p^2(c) = 1$, from p^2 only towards the lottery p^3 with $p^3(d) = p^3(e) = \frac{1}{2}$, and from p^3 only towards p^1 . Hence, we cycle through PC-inefficient lotteries and never reach an efficient one.

Step 1: As the first step, we show that every lottery p with $p(a) = p(b) > 0$ and $p(d) = p(e) = 0$ is only PC-dominated by lotteries q with $q(a) = q(b) < p(a) = p(b)$ and $q(d) = q(e) = 0$. For this, we observe that voters 1 and 2 only agree that a is preferred to e and that b is preferred to d . By (7), q can PC-dominate p only if

$$q(b)p(d) + q(a)p(e) \geq p(b)q(d) + p(a)q(e).$$

Since $p(d) = p(e) = 0$ and $p(a) = p(b) > 0$ by assumption, this inequality can only be true if $q(d) = q(e) = 0$. Hence, we have that $q(b)p(d) + q(a)p(e) = p(b)q(d) + p(a)q(e) = 0$. Equation (8) and the fact that $p(d) = p(e) = q(d) = q(e) = 0$ therefore show that

$$q(b)p(c) + q(b)p(a) + q(c)p(a) = p(b)q(c) + p(b)q(a) + p(c)q(a).$$

Finally, since $p(a) = p(b) > 0$, it is easy to derive that this equation holds only if $q(a) = q(b)$. Using this fact and $q(d) = q(e) = 0$, it follows that voters 3 and 4 PC-prefer q to p only if $q(c) \geq p(c)$. Since $p \neq q$ if q PC-dominates p , we thus infer that q dominates p in R only if $q(a) = q(b) < p(a) = p(b)$ and $q(d) = q(e) = p(d) = p(e) = 0$. Finally, it is easy to verify that all voters indeed PC-prefer such a lottery q to p . This argument shows that every lottery p with $p = \lambda p^1 + (1 - \lambda)p^2$ is only PC-dominated by another lottery of the same form but with smaller λ . In particular, if we apply this argument at p^1 , we eventually have to go to p^2 when we aim to reach a PC-efficient lottery by applying PC-improvements.

Step 2: Next, we prove that every lottery p with $p(a) = p(b) = 0$, $p(c) > 0$, and $p(d) = p(e)$ is only PC-dominated by lotteries q with $q(a) = q(b) = 0$ and $q(d) = q(e) > p(d) = p(e)$. To this end, consider the preferences of voters 3 and 4 and note that these voters only agree on the preferences between c and a , c and b , and a and b . Hence, (7) requires that

$$q(c)p(a) + q(c)p(b) + q(a)p(b) \geq p(c)q(a) + p(c)q(b) + p(a)q(b).$$

Since $p(a) = p(b) = 0$ and $p(c) > 0$ by assumption, this inequality can be true only if $q(a) = q(b) = 0$. We hence infer that $q(c)(p(a) + p(b)) + q(a)p(b) = p(c)(q(a) + q(b)) + p(a)q(b) = 0$. Consequently, Equation (8) and the fact that $p(a) = p(b) = q(a) = q(b) = 0$ imply that

$$q(d)p(c) + q(d)p(e) + q(c)p(e) = p(d)q(c) + p(d)q(e) + p(c)q(e).$$

Since $p(d) = p(e)$ by assumption, this equation can hold only if $q(d) = q(e)$. This insight and the fact that $q(a) = q(b) = 0$ imply that voters 5 to 8 only PC-prefer q to p if $q(d) = q(e) \geq p(d) = p(e)$. Since $p \neq q$ if q PC-dominates p , this inequality must be strict. Finally, it is not difficult to verify that every lottery q that satisfies $q(a) = q(b) = 0$ and $q(d) = q(e) > p(d) = p(e)$ indeed PC-dominates p . In particular, this means that every lottery p with $p = \lambda p^2 + (1 - \lambda)p^3$ is only PC-dominated by a lottery q with $p = \lambda' p^2 + (1 - \lambda')p^3$ and $\lambda' < \lambda$. Hence, if we are at p^2 and try to find a PC-efficient lottery by applying PC-improvements, we will inevitably arrive at p^3 .

Step 3: As the last step, we prove that every lottery p with $p(a) = p(b)$, $p(c) = 0$, and $p(d) = p(e) > 0$ is only PC-dominated by lotteries q with $q(a) = q(b) > p(a) = p(b)$, $q(c) = 0$, and $q(d) = q(e)$. For this, we consider voters 5 to 8 and note that voters 5 and 6 agree only on the preference between d and c and between e and c . Moreover, the same is true for voters 7 and 8 and we thus derive from (7) that

$$q(d)p(c) + q(e)p(c) \geq p(d)q(c) + p(e)q(c).$$

Using the fact that $p(c) = 0$ and $p(d) = p(e) > 0$, this inequality can be true only if $q(c) = 0$. Hence, we have that $q(d)p(c) + q(e)p(c) = p(d)q(c) + p(e)q(c) = 0$. Using Equation (8) and the fact that $q(c) = p(c) = 0$, we derive the following two equations. The first one corresponds to voters 5 and 6, whereas the second one corresponds to voters 7 and 8.

$$\begin{aligned} & q(b)(p(d) + p(e) + p(a)) + q(d)(p(e) + p(a)) + q(e)p(a) \\ &= p(b)(q(d) + q(e) + q(a)) + p(d)(q(e) + q(a)) + p(e)q(a) \\ & q(b)(p(e) + p(d) + p(a)) + q(e)(p(d) + p(a)) + q(d)p(a) \\ &= p(b)(q(e) + q(d) + q(a)) + p(e)(q(d) + q(a)) + p(d)q(a) \end{aligned} \tag{9}$$

Subtracting these two equations from each other yields that $q(d)p(e) - q(e)p(d) = p(d)q(e) - p(e)q(d)$. Since $p(d) = p(e) > 0$, we derive from this equation that $q(d) = q(e)$. Using this observation and the assumptions that $p(a) = p(b)$ and $p(d) = p(e) > 0$, we can infer from Equation (9) that $q(a) = q(b)$. In summary, we therefore have that $q(c) = 0$, $q(a) = q(b)$,

and $q(d) = q(e)$. Next, voters 1 and 2 only *PC*-prefer q to p if $q(a) = q(b) \geq p(a) = p(b)$, which means that $q(a) = q(b) > p(a) = p(b)$ because $p \neq q$. Finally, it is easily seen that all voters indeed *PC*-prefer q to p . This means that a lottery p with $p = \lambda p^3 + (1 - \lambda)p^1$ is only *PC*-dominated by another lottery q with the same form but smaller λ . As a consequence, when applying *PC*-improvements from p^3 to reach a *PC*-efficient outcome, one will inevitably reach p^1 . \square

A.2. Proof of Theorem 2

Next, we prove Theorem 2 and start by presenting three auxiliary lemmas. In more detail, we initiate our analysis by investigating the consequences of *PC*-strategyproofness for lotteries with small support. To this end, let the *support* of a lottery be $\text{supp}(p) = \{x \in A : p(x) > 0\}$. Our first lemma then focuses on the case where $|\text{supp}(f(R))| \leq 2$ and states that, if the support does not change after a manipulation and the manipulator does not reorder the alternatives in the support, the outcome is not allowed to change.

Lemma 2. *Let $R, R' \in \mathcal{R}^N$ for some $N \in \mathcal{F}(\mathbb{N})$, $i \in N$, and $a, b \in A$ such that $R_{-i} = R'_{-i}$ and $a \succ_i b$ iff $a \succ'_i b$. Then, every *PC*-strategyproof SDS f satisfies $f(R) = f(R')$ if $\text{supp}(f(R)) \subseteq \{a, b\}$ and $\text{supp}(f(R')) \subseteq \{a, b\}$.*

Proof. Let $R, R' \in \mathcal{R}^N$, $N \in \mathcal{F}(\mathbb{N})$, $i \in N$, and $a, b \in A$ be two distinct alternatives. Without loss of generality, we assume that $a \succ_i b$ and $a \succ'_i b$; the case where voter i prefers b to a is symmetric. Moreover, let f denote a *PC*-strategyproof SDS and assume that $\text{supp}(f(R)) \subseteq \{a, b\}$ and $\text{supp}(f(R')) \subseteq \{a, b\}$. Now, assume for contradiction that $f(R) \neq f(R')$. Since both lotteries only put positive probability on a and b , this means either that $f(R, a) < f(R', a)$ and $f(R, b) > f(R', b)$, or that $f(R, a) > f(R', a)$ and $f(R, b) < f(R', b)$. First, suppose that $f(R, a) < f(R', a)$. Then, $f(R') \succ_i^{PC} f(R)$, and voter i can thus *PC*-manipulate by deviating from R to R' . On the other hand, if $f(R, a) > f(R', a)$, voter i can *PC*-manipulate by deviating from R' to R as he *PC*-prefers $f(R)$ to $f(R')$ with respect to \succ'_i . Hence, both cases result in a *PC*-manipulation, contradicting the *PC*-strategyproofness of f . This proves that $f(R) = f(R')$. \square

Next, we analyze *PC*-strategyproofness when $|\text{supp}(f(R))| \leq 3$. In this case, only a significantly weaker implication holds: if $\text{supp}(f(R)) \subseteq \{a, b, c\}$, $\text{supp}(f(R')) \subseteq \{a, b, c\}$, $f(R, a) < f(R, c)$, and $a \succ_i b \succ_i c$ for some voter i , then this voter cannot change the fact that a gets less probability than c in the resulting lottery.

Lemma 3. *Let $R, R' \in \mathcal{R}^N$ for some $N \in \mathcal{F}(\mathbb{N})$, $i \in N$, and $a, b, c \in A$ such that $R_{-i} = R'_{-i}$ and $a \succ_i b \succ_i c$. Then, every *PC*-strategyproof SDS f satisfies $f(R', a) < f(R', c)$ if $f(R, a) < f(R, c)$, $\text{supp}(f(R)) \subseteq \{a, b, c\}$, and $\text{supp}(f(R')) \subseteq \{a, b, c\}$.*

Proof. Let $R, R' \in \mathcal{R}^N$ for some $N \in \mathcal{F}(\mathbb{N})$, $i \in N$, and $a, b, c \in A$ be three distinct alternatives such that $R_{-i} = R'_{-i}$ and $a \succ_i b \succ_i c$. Furthermore, consider a *PC*-strategyproof SDS f and suppose that $\text{supp}(f(R)) \subseteq \{a, b, c\}$ and $\text{supp}(f(R')) \subseteq \{a, b, c\}$. For simplicity, we define $p = f(R)$ and $q = f(R')$ and assume for contradiction that $p(a) < p(c)$ and $q(a) \geq q(c)$. Next, we use *PC*-strategyproofness to relate p and q . In particular, we infer the following equation from the *PC*-strategyproofness between R and R' . Note that the alternatives $x \in A \setminus \{a, b, c\}$ can be omitted as $p(x) = q(x) = 0$.

$$p(a)q(b) + p(a)q(c) + p(b)q(c) \geq q(a)p(b) + q(a)p(c) + q(b)p(c)$$

Using the fact that $1 = q(a) + q(b) + q(c)$ and $1 = p(a) + p(b) + p(c)$, we have two possibilities of rewriting this inequality.

$$\begin{aligned} p(a)(1 - q(a)) + p(b)q(c) &\geq q(a)(1 - p(a)) + q(b)p(c) \\ \iff p(a) + p(b)q(c) &\geq q(a) + q(b)p(c) \\ p(a)q(b) + (1 - p(c))q(c) &\geq q(a)p(b) + (1 - q(c))p(c) \\ \iff p(a)q(b) + q(c) &\geq q(a)p(b) + p(c) \end{aligned}$$

Summing up these two inequalities results in the following inequality.

$$\begin{aligned} p(a)(1 + q(b)) + q(c)(1 + p(b)) &\geq q(a)(1 + p(b)) + p(c)(1 + q(b)) \\ \iff (p(a) - p(c))(1 + q(b)) &\geq (1 + p(b))(q(a) - q(c)) \end{aligned}$$

Our assumption that $p(a) < p(c)$ implies that the left-hand side of the inequality is smaller than 0. On the other hand, we have $q(a) \geq q(c)$, so the right-hand side is non-negative. This is a contradiction, which proves that our assumptions on p and q are in conflict with *PC*-strategyproofness. Hence, if $f(R, a) < f(R, c)$, then $f(R', a) < f(R', c)$. \square

For our proofs, we also need insights on circumstances under which a lottery is *PC*-inefficient. To this end, we analyze in the next lemma when voters prefer certain lotteries to each other. In order to succinctly formalize these results, we define the rank of an alternative x in a preference relation \succ_i as $r(\succ_i, x) = 1 + |\{y \in A \setminus \{x\} : y \succ_i x\}|$. In particular, if $r(\succ_i, x) < r(\succ_i, y)$, then $x \succ_i y$ due to the transitivity of preference relations.

Lemma 4. Let $A = \{w, x, y, z\}$, let p denote a lottery on A with $p(x) > 0$ and $p(y) > 0$, and define q as $q(x) = p(x) - \frac{\varepsilon}{p(x)+p(z)}$, $q(y) = p(y) - \frac{\varepsilon}{p(y)+p(z)}$, $q(z) = p(z) + \frac{\varepsilon}{p(x)+p(z)} + \frac{\varepsilon}{p(y)+p(z)}$, and $q(w) = p(w)$, where $\varepsilon > 0$ is sufficiently small so that $q(x) \geq 0$, $q(y) \geq 0$. Then, the following PC-preferences hold:

1. If $z \succ_i x$ and $z \succ_i y$, then $q \succ_i^{PC} p$.
2. If $x \succ_i z$ and $y \succ_i z$, then $p \succ_i^{PC} q$.
3. Assume that $x \succ_i z \succ_i y$ or $y \succ_i z \succ_i x$.
 - a) If $p(w) = 0$ or $r(\succ_i, w) \in \{1, 4\}$, then $p \succ_i^{PC} q$ and $q \succ_i^{PC} p$.
 - b) If $r(\succ_i, w) = 2$ and $p(w) > 0$, then $p \succ_i^{PC} q$.
 - c) If $r(\succ_i, w) = 3$ and $p(w) > 0$, then $q \succ_i^{PC} p$.

Proof. Let p and q be defined as in the lemma. Claims 1 and 2 follow immediately since a voter i with $z \succ_i x$ and $z \succ_i y$ (resp. $x \succ_i z$ and $y \succ_i z$) strictly SD-prefers q to p (resp. p to q).

For Claim 3, suppose that $x \succ_i z \succ_i y$; the case that $y \succ_i z \succ_i x$ is symmetric. The key insight for this case is that

$$\begin{aligned}
 & p(x)(q(z) + q(y)) + p(z)q(y) \\
 &= p(x)\left(p(z) + p(y) + \frac{\varepsilon}{p(x) + p(z)}\right) + p(z)\left(p(y) - \frac{\varepsilon}{p(y) + p(z)}\right) \\
 &= p(x)(p(z) + p(y)) + p(z)p(y) + \frac{\varepsilon p(x)}{p(x) + p(z)} - \frac{\varepsilon p(z)}{p(y) + p(z)} \\
 &= p(x)(p(z) + p(y)) + p(z)p(y) + \left(\varepsilon - \frac{\varepsilon p(z)}{p(x) + p(z)}\right) - \left(\varepsilon - \frac{\varepsilon p(y)}{p(y) + p(z)}\right) \\
 &= p(x)(p(z) + p(y)) + p(z)p(y) - \frac{\varepsilon p(z)}{p(x) + p(z)} + \frac{\varepsilon p(y)}{p(y) + p(z)} + \frac{\varepsilon p(y) - \varepsilon p(z)}{p(x) + p(z)} \\
 &= \left(p(x) - \frac{\varepsilon}{p(x) + p(z)}\right)(p(z) + p(y)) + \left(p(z) + \frac{\varepsilon}{p(x) + p(z)} + \frac{\varepsilon}{p(y) + p(z)}\right)p(y) \\
 &= q(x)(p(z) + p(y)) + q(z)p(y).
 \end{aligned}$$

As a consequence of this equation, voter i 's preference between q and p only depends on $p(w)$ and $r(\succ_i, w)$. First, if $p(w) = q(w) = 0$, the PC-comparison between p and q with respect to \succ_i reduces exactly to the above equation. Hence, $p \succ_i^{PC} q$ and $q \succ_i^{PC} p$ if $p(w) = q(w) = 0$. We therefore suppose that $p(w) = q(w) > 0$ and proceed with a case distinction with respect to $r(\succ_i, w)$. In this analysis, we use $\Delta_{p \rightarrow q}^{xyz} = p(x)(q(z) + q(y)) + p(z)q(y)$ and $\Delta_{q \rightarrow p}^{xyz} = q(x)(p(z) + p(y)) + q(z)p(y)$ as a shorthand notation.

- First, suppose that $r(\succ_i, w) = 1$, i.e., $\succ_i = w, x, z, y$. Since $p(w) = q(w)$, $p(x) + p(y) + p(z) = 1 - p(w)$, and $q(x) + q(y) + q(z) = 1 - q(w)$, it is easy to verify that $p(w)(q(x) + q(y) + q(z)) + \Delta_{p \rightarrow q}^{xyz} = q(w)(p(x) + p(y) + p(z)) + \Delta_{q \rightarrow p}^{xyz}$. This implies that $p \succ_i^{PC} q$ and $q \succ_i^{PC} p$.
- As the second case, suppose that $r(\succ_i, w) = 4$, i.e., $\succ_i = x, z, y, w$. For the same reason as in the first case it follows that $\Delta_{p \rightarrow q}^{xyz} + (p(x) + p(y) + p(z))q(w) = \Delta_{q \rightarrow p}^{xyz} + (q(x) + q(y) + q(z))p(w)$ and thus, $p \succ_i^{PC} q$ and $q \succ_i^{PC} p$.
- Next, suppose that $r(\succ_i, w) = 2$, i.e., $\succ_i = x, w, z, y$. Then, $p \succ_i^{PC} q$ as

$$\begin{aligned}
 \sum_{u, v \in A: u \succ_i v} p(u)q(v) &= p(x)q(w) + p(w)q(y) + p(w)q(z) + \Delta_{p \rightarrow q}^{xyz} \\
 &= p(x)p(w) + p(w)\left(p(y) + p(z) + \frac{\varepsilon}{p(x) + p(z)}\right) + \Delta_{p \rightarrow q}^{xyz} \\
 &= \left(p(x) + \frac{\varepsilon}{p(x) + p(z)}\right)p(w) + p(w)(p(y) + p(z)) + \Delta_{p \rightarrow q}^{xyz} \\
 &> q(x)p(w) + q(w)(p(y) + p(z)) + \Delta_{q \rightarrow p}^{xyz} \\
 &= \sum_{u, v \in A: u \succ_i v} q(u)p(v).
 \end{aligned}$$

- As the last case, suppose that $r(\succ_i, w) = 3$, i.e., $\succ_i = x, z, w, y$. In this case, a symmetric inequality as in the previous case proves that $q \succ_i^{PC} p$. \square

Note that Lemma 4 also holds for all lotteries q and p with $q(z) > 0$ and $p(x) = q(x) + \frac{\varepsilon'}{q(x)+q(z)}$, $p(y) = q(y) + \frac{\varepsilon'}{q(y)+q(z)}$, $p(z) = q(z) - \frac{\varepsilon'}{q(x)+q(z)} - \frac{\varepsilon'}{q(y)+q(z)}$, and $p(w) = q(w)$ ($\varepsilon' > 0$ is again sufficiently small so that p is a well-defined lottery on $\{w, x, y, z\}$). The reason for this is that $q(x) = p(x) - \frac{\varepsilon'}{p(x)+p(z)}$, $q(y) = p(y) - \frac{\varepsilon'}{p(y)+p(z)}$, and $q(z) = p(z) + \frac{\varepsilon'}{p(x)+p(z)} + \frac{\varepsilon'}{p(y)+p(z)}$ for $\varepsilon = \varepsilon' - \frac{\varepsilon'^2}{(q(x)+q(z)) \cdot (q(y)+q(z))}$. For instance, this follows for $p(x)$ from the following equation.

$$\begin{aligned} p(x) - \frac{\varepsilon}{p(x) + p(z)} &= p(x) - \frac{\varepsilon' - \frac{\varepsilon'^2}{(q(x)+q(z)) \cdot (q(y)+q(z))}}{q(x) + q(z) - \frac{\varepsilon'}{q(y)+q(z)}} = p(x) - \frac{\varepsilon'(q(x) + q(z)) - \frac{\varepsilon'^2}{q(y)+q(z)}}{(q(x) + q(z)) \cdot (q(x) + q(z) - \frac{\varepsilon'}{q(y)+q(z)})} \\ &= p(x) - \frac{\varepsilon'}{q(x) + q(z)} = q(x) \end{aligned}$$

Finally, we are now ready to prove Lemma 1. We prove this statement separately for electorates with an odd number of voters and for those with an even number of voters.

Lemma 1a). Every PC-efficient SDS that satisfies the absolute winner property is PC-manipulable if $|N| \geq 3$ is odd and $m \geq 4$.

Proof. Consider an arbitrary electorate $N \in \mathcal{F}(\mathbb{N})$ with an odd number of voters $n = |N| \geq 3$ and suppose there are $m \geq 4$ alternatives. We assume for contradiction that there is an PC-efficient SDS f that satisfies the absolute winner property and PC-strategyproofness for N . In the sequel, we will focus on profiles on the alternatives $\{a, b, c, d\}$; all other alternatives are always ranked below these alternatives and therefore Pareto-dominated. Hence, PC-efficiency entails for all subsequent profiles that $f(R, x) = 0$ for all $x \in A \setminus \{a, b, c, d\}$, which means that these alternatives do not affect our further analysis. In slight abuse of notation, we therefore assume that $A = \{a, b, c, d\}$.

We derive a contradiction by focusing on the profiles R and R' shown below. Specifically, our goal is to show that $f(R, a) = f(R, b) = f(R, c) = \frac{1}{3}$ and $f(R', a) = f(R', c) = f(R', d) = \frac{1}{3}$. This implies that voter $\frac{n+1}{2}$ can PC-manipulate by switching from R' to R as he even SD-prefers $f(R)$ to $f(R')$, i.e., these claims result in a contradiction to PC-strategyproofness.

$$\begin{array}{lll} R: & [1 \dots \frac{n-1}{2}]: a, d, b, c & \frac{n+1}{2}: b, c, d, a \quad [\frac{n+3}{2} \dots n]: c, a, d, b \\ R': & [1 \dots \frac{n-1}{2}]: a, d, b, c & \frac{n+1}{2}: b, d, c, a \quad [\frac{n+3}{2} \dots n]: c, a, d, b \end{array}$$

Claim 1: $f(R, a) = f(R, b) = f(R, c) = \frac{1}{3}$

For proving this claim, our first goal is to establish that $f(R, c) > 0$. Hence, assume for contradiction that this is not the case, i.e., $f(R, c) = 0$. For deriving a contradiction to this assumption, we consider the profiles R^1 and R^2 shown below.

$$\begin{array}{lll} R^1: & [1 \dots \frac{n-1}{2}]: a, d, b, c & \frac{n+1}{2}: b, c, d, a \quad [\frac{n+3}{2} \dots n-1]: c, a, d, b \quad n: a, c, d, b \\ R^2: & [1 \dots \frac{n-1}{2}]: a, d, b, c & \frac{n+1}{2}: c, b, d, a \quad [\frac{n+3}{2} \dots n]: c, a, d, b \end{array}$$

First, note that a is top-ranked by more than half of the voters in R^1 and c by more of half of the voters in R^2 . Hence, the absolute winner property requires that $f(R^1, a) = f(R^2, c) = 1$. On the other hand, R^1 is derived from R by letting voter n swap a and c . Hence, PC-strategyproofness, or more precisely PC1-strategyproofness, from R to R^1 implies that $f(R, c) \geq f(R, b) + f(R, d)$. Because we assume that $f(R, c) = 0$, this means that $f(R, b) = f(R, d) = 0$ and $f(R, a) = 1$. On the other hand, the profile R^2 is derived from R by letting voter $\frac{n+1}{2}$ swap b and c . Hence, PC-strategyproofness requires that $f(R, b) \geq f(R, a) + f(R, d)$, which conflicts with $f(R, a) = 1$. Thus, the initial assumption that $f(R, c) = 0$ is incorrect, i.e., it holds that $f(R, c) > 0$.

Departing from this insight, PC-efficiency entails that $f(R, d) = 0$. In more detail, Lemma 4 proves that every lottery q with $q(d) > 0$ and $q(c) > 0$ is PC-inefficient for R because it is dominated by the lottery p with $p(a) = q(a) + \frac{\varepsilon}{q(a)+q(d)}$, $p(b) = q(b) + \frac{\varepsilon}{q(b)+q(d)}$, $p(c) = q(c)$, and $p(d) = q(d) - \frac{\varepsilon}{q(a)+q(d)} - \frac{\varepsilon}{q(b)+q(d)}$. Indeed, Case 3a) of this lemma shows that all voters but $\frac{n+1}{2}$ are indifferent between p and q , whereas Case 3b) implies that voter $\frac{n+1}{2}$ strictly prefers p to q (see also the text after Lemma 4). Since we already know that $f(R, c) > 0$, it follows therefore that $f(R, d) = 0$.

Next, note that the inequalities derived from PC-strategyproofness on R^1 and R^2 remain valid, even if $f(R, c) > 0$. Combined with the fact that $f(R, d) = 0$, this means that $f(R, c) \geq f(R, b) \geq f(R, a)$. Hence, we prove Claim 1 by showing that $f(R, a) \geq f(R, c)$. Consider for this the profiles \bar{R}^i for $i \in \{0, \dots, \frac{n-1}{2}\}$, which are defined as follows.

$$\bar{R}^i: \quad [1 \dots i]: a, d, b, c \quad [i+1 \dots \frac{n-1}{2}]: b, a, d, c \quad \frac{n+1}{2}: b, c, d, a \quad [\frac{n+3}{2} \dots n]: c, a, d, b$$

First, note that $\bar{R}^{\frac{n-1}{2}} = R$ and that $f(\bar{R}^0, b) = 1$ because $\frac{n+1}{2}$ voters report b as their favorite alternative in this profile. Furthermore, Lemma 4 shows that $f(\bar{R}^i, d) = 0$ for every profile \bar{R}^i with $i < \frac{n-1}{2}$ because all lotteries q with $q(d) > 0$ fail PC-efficiency for \bar{R}^i . Indeed, the lottery p with $p(a) = q(a) + \frac{\varepsilon}{q(a)+q(d)}$, $p(b) = q(b) + \frac{\varepsilon}{q(b)+q(d)}$, $p(c) = q(c)$, and $p(d) = q(d) - \frac{\varepsilon}{q(a)+q(d)} - \frac{\varepsilon}{q(b)+q(d)}$ PC-dominates q . Finally, by a repeated application of Lemma 3, we derive that $f(R, a) \geq f(R, c)$. To this end, consider a fixed index $i \in \{1, \dots, \frac{n-1}{2}\}$. If $f(\bar{R}^i, a) < f(\bar{R}^i, c)$, this lemma requires that $f(\bar{R}^{i-1}, a) < f(\bar{R}^{i-1}, c)$. Hence,

if $f(R, a) < f(R, c)$, we can repeatedly apply this argument to derive that $f(\bar{R}^0, a) < f(\bar{R}^0, c)$. However, this contradicts the absolute winner property, and thus we must have that $f(R, a) \geq f(R, c)$. This proves Claim 1.

Claim 2: $f(R', a) = f(R', c) = f(R', d) = \frac{1}{3}$

As the second claim, we prove that f assigns a probability of $\frac{1}{3}$ to a , c , and d in R' . For this, we proceed analogously to Claim 1 and first show that $f(R', c) > 0$. Assume for contradiction that $f(R', c) = 0$ and consider the profiles R^1 and R^2 shown below.

$$\begin{array}{llll} R^1: & [1 \dots \frac{n-1}{2}]: a, d, b, c & \frac{n+1}{2}: b, d, c, a & [\frac{n+3}{2} \dots n-1]: c, a, d, b \quad n: a, c, d, b \\ R^2: & [1 \dots \frac{n-1}{2}]: a, d, b, c & \frac{n+1}{2}: c, d, b, a & [\frac{n+3}{2} \dots n]: c, a, d, b \end{array}$$

First, note that $f(R^1, a) = 1$ and $f(R^2, c) = 1$ because of the absolute winner property. Next, observe that R^1 is derived from R' by letting voter n swap a and c . Hence, PC -strategyproofness requires that $f(R', c) \geq f(R', b) + f(R', d)$. Since $f(R', c) = 0$ by assumption, it follows that $f(R', b) = f(R', d) = 0$ and $f(R', a) = 1$. On the other hand, we derive R^2 from R' by letting voter $\frac{n+1}{2}$ deviate. Hence, PC -strategyproofness implies that $f(R', b) + f(R', d) \geq f(R', a)$, which contradicts $f(R', a) = 1$. This shows that the initial assumption $f(R', c) = 0$ is wrong, i.e., it must be that $f(R', c) > 0$.

As the next step, we will infer from PC -efficiency that $f(R', b) = 0$. Assume that this is not the case, i.e., there is a PC -efficient lottery p with $p(b) > 0$ and $p(c) > 0$. Now, if $p(a) > 0$, then p is PC -dominated by the lottery q with $q(a) = p(a) - \frac{\varepsilon}{p(a)+p(d)}$, $q(b) = p(b) - \frac{\varepsilon}{p(b)+p(d)}$, $q(c) = p(c)$, and $q(d) = p(d) + \frac{\varepsilon}{p(a)+p(d)} + \frac{\varepsilon}{p(b)+p(d)}$ (where $\varepsilon > 0$ is so small that q is a well-defined lottery). On the other hand, if $p(a) = 0$, then p is PC -dominated by the lottery q with $q(a) = p(a)$, $q(b) = p(b) - \frac{\varepsilon}{p(b)+p(d)}$, $q(c) = p(c) - \frac{\varepsilon}{p(c)+p(d)}$, and $q(d) = p(d) + \frac{\varepsilon}{p(b)+p(d)} + \frac{\varepsilon}{p(c)+p(d)}$. Both of these claims are straightforward to verify with Lemma 4. Since p is PC -inefficient in both cases, it follows that $f(R', b) = 0$.

Just as for R , we can use the fact that $f(R', b) = 0$ to simplify the inequalities derived from PC -strategyproofness on R^1 and R^2 . In particular, we infer that $f(R', c) \geq f(R', d) \geq f(R', a)$ from these observations. Hence, Claim 2 will follow by showing that $f(R', a) \geq f(R', c)$. Consider for this the profiles \bar{R}^i for $i \in \{0, \dots, \frac{n-1}{2}\}$, which are defined as follows.

$$\bar{R}^i: \quad [1 \dots i]: a, d, b, c \quad [i+1 \dots \frac{n-1}{2}]: d, a, b, c \quad \frac{n+1}{2}: b, d, c, a \quad [\frac{n+3}{2} \dots n]: c, a, d, b$$

First, observe that $f(\bar{R}^i, b) = 0$ for all $i \in \{0, \dots, \frac{n-1}{2}\}$ because of PC -efficiency and PC -strategyproofness. Indeed, assume for contradiction that this is not true, i.e., $f(\bar{R}^i, b) > 0$ for some $i \in \{0, \dots, \frac{n-1}{2}\}$. First, we show that this assumption implies that $f(\bar{R}^i, a) = 0$ and $f(\bar{R}^i, c) = 0$ because of PC -efficiency. For this, note that every lottery p with $p(a) > 0$ and $p(b) > 0$ is PC -dominated by the lottery q with $q(a) = p(a) - \frac{\varepsilon}{p(a)+p(d)}$, $q(b) = p(b) - \frac{\varepsilon}{p(b)+p(d)}$, $q(c) = p(c)$, and $q(d) = p(d) + \frac{\varepsilon}{p(a)+p(d)} + \frac{\varepsilon}{p(b)+p(d)}$ in \bar{R}^i . Indeed, Lemma 4 shows that the voters $j \in \{i+1, \dots, \frac{n-1}{2}\}$ strictly PC -prefer q to p and all other voters at least weakly PC -prefer q to p . Moreover, using again Lemma 4, it is easy to see that every lottery p with $p(a) = 0$, $p(b) > 0$, and $p(c) > 0$ is PC -dominated by the lottery q with $q(a) = 0$, $q(b) = p(b) - \frac{\varepsilon}{p(b)+p(d)}$, $q(c) = p(c) - \frac{\varepsilon}{p(c)+p(d)}$, and $q(d) = p(d) + \frac{\varepsilon}{p(b)+p(d)} + \frac{\varepsilon}{p(c)+p(d)}$. Hence, if $f(\bar{R}^i, b) > 0$, we derive that $f(\bar{R}^i, a) = 0$ and $f(\bar{R}^i, c) = 0$, which means that $\text{supp}(f(\bar{R}^i)) \subseteq \{b, d\}$. However, this entails that one of the voters $j \in \{\frac{n+1}{2}, \dots, n\}$ can PC -manipulate. Consider for this the subsequent preference profiles $\bar{R}^{i,j}$ for $j \in \{\frac{n+1}{2}, \dots, n\}$ and $i \in \{0, \dots, \frac{n-1}{2}\}$.

$$\bar{R}^{i,j}: \quad [1 \dots i]: a, d, b, c \quad [i+1 \dots \frac{n-1}{2}]: d, a, b, c \quad \frac{n+1}{2}: b, d, c, a \quad [\frac{n+3}{2} \dots j]: c, a, d, b \quad [j+1 \dots n]: d, a, b, c$$

Note that $\bar{R}^{i,n} = \bar{R}^i$, and that $f(\bar{R}^{i, \frac{n+1}{2}}, d) = 1$ because more than half of the voters report d as their favorite choice. On the other hand, we claim that for $j \in \{\frac{n+3}{2}, \dots, n\}$, if $f(\bar{R}^{i,j}, b) > 0$, then $f(\bar{R}^{i,j-1}, b) > 0$. Observe for this that the voter types in \bar{R}^i and $\bar{R}^{i,j}$ coincide, and thus PC -efficiency also requires that $f(\bar{R}^{i,j}, a) = f(\bar{R}^{i,j}, c) = 0$ if $f(\bar{R}^{i,j}, b) > 0$. Moreover, PC -strategyproofness requires that the deviating voter j PC -prefers $f(\bar{R}^{i,j})$ to $f(\bar{R}^{i,j-1})$. Now, if $f(\bar{R}^{i,j-1}, b) = 0$, deviating to $\bar{R}^{i,j-1}$ is a PC -manipulation for voter j because $\text{supp}(f(\bar{R}^{i,j}))$ consists of his worst two alternatives. Hence, $f(\bar{R}^{i,j}, b) > 0$ implies that $f(\bar{R}^{i,j-1}, b) > 0$ and, by repeatedly applying this argument, we infer that if $f(\bar{R}^i, b) > 0$, then $f(\bar{R}^{i, \frac{n+1}{2}}, b) > 0$. However, this contradicts $f(\bar{R}^{i, \frac{n+1}{2}}, d) = 1$, so we have that $f(\bar{R}^i, b) = 0$ for all $i \in \{0, \dots, \frac{n-1}{2}\}$.

In particular, this argument also proves for the profile \bar{R}^0 that $f(\bar{R}^0, b) = 0$. We show next that $f(\bar{R}^0, d) = 1$. Consider for this the profile \hat{R} derived from \bar{R}^0 by letting voter $\frac{n+1}{2}$ swap b and d .

$$\hat{R}: \quad [1 \dots \frac{n-1}{2}]: d, a, b, c \quad \frac{n+1}{2}: d, b, c, a \quad [\frac{n+3}{2} \dots n]: c, a, d, b$$

We have that $f(\hat{R}, d) = 1$ because of the absolute winner property. Hence, PC -strategyproofness requires that $f(\bar{R}^0, b) \geq f(\bar{R}^0, c) + f(\bar{R}^0, a)$. Since $f(\bar{R}^0, b) = 0$, this means that $f(\bar{R}^0, c) = f(\bar{R}^0, a) = 0$ and therefore $f(\bar{R}^0, d) = 1$. Based on this observation, we can now use Lemma 3 to derive that $f(R', a) \geq f(R', c)$. Consider for this an index $i \in \{1, \dots, \frac{n-1}{2}\}$ and suppose that $f(\bar{R}^{i-1}, a) \geq f(\bar{R}^{i-1}, c)$. The contraposition of Lemma 3 shows that $f(\bar{R}^i, a) \geq f(\bar{R}^i, c)$ because the deviating voter i prefers a to d to c . Finally, since $f(\bar{R}^0, a) = f(\bar{R}^0, c) = 0$, repeatedly applying the previous argument and noting that $R' = \bar{R}^{\frac{n-1}{2}}$, we obtain $f(R', a) \geq f(R', c)$. This establishes Claim 2. \square

Next, we turn to the proof for electorates with an even number of voters. Note that the proof follows a similar structure but requires more involved arguments because we cannot change the absolute winner by only modifying a single preference relation.

Lemma 1b). Every PC-efficient SDS that satisfies the absolute winner property is PC-manipulable if $|N| \geq 8$ is even and $m \geq 4$.

Proof. Consider an arbitrary electorate $N \in \mathcal{F}(\mathbb{N})$ with $n = |N| \geq 8$ even and assume for contradiction that there is an SDS f for $m \geq 4$ alternatives that satisfies PC-efficiency, PC-strategyproofness, and the absolute winner property on N . We focus on the case $m = 4$ because we can generalize the constructions to larger values of m by simply ranking the additional alternatives at the bottom. Then, PC-efficiency requires that these alternatives obtain probability 0 and they therefore do not affect our analysis. We derive a contradiction by analyzing the following two profiles.

$$\begin{aligned} R: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, c, d, a & [\frac{n}{2}+2 \dots n]: c, a, d, b \\ R': & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, d, c, a & [\frac{n}{2}+2 \dots n]: c, a, d, b \end{aligned}$$

In more detail, we show in Claims 1 and 2 that $f(R, a) = f(R, b) = f(R, c) = \frac{1}{3}$ and $f(R', a) = f(R', c) = f(R', d) = \frac{1}{3}$. These two claims are in conflict with PC-strategyproofness, as the following analysis shows. Let R'' denote the profile “between” R and R' in which voter $\frac{n}{2}$ reports b, d, c, a and voter $\frac{n}{2} + 1$ reports b, c, d, a .

$$R'': [1 \dots \frac{n}{2}-1]: a, d, b, c \quad \frac{n}{2}: b, d, c, a \quad \frac{n}{2}+1: b, c, d, a \quad [\frac{n}{2}+2 \dots n]: c, a, d, b$$

Moreover, let $p = f(R)$, $q = f(R')$, and $r = f(R'')$ denote the outcome of f in these profiles. PC-strategyproofness from R' to R'' results in the following inequality because $q(a) = q(c) = q(d) = \frac{1}{3}$.

$$\begin{aligned} & q(b)(r(d) + r(c) + r(a)) + q(d)(r(c) + r(a)) + q(c)r(a) \\ & \geq r(b)(q(d) + q(c) + q(a)) + r(d)(q(c) + q(a)) + r(c)q(a) \\ \iff & \frac{1}{3}r(c) + \frac{2}{3}r(a) \geq r(b) + \frac{2}{3}r(d) + \frac{1}{3}r(c) \iff r(a) \geq \frac{3}{2}r(b) + r(d) \end{aligned}$$

Moreover, we can also use PC-strategyproofness from R'' to R and the fact that $p(a) = p(b) = p(c) = \frac{1}{3}$ to infer the following inequality.

$$\begin{aligned} & r(b)(p(d) + p(c) + p(a)) + r(d)(p(c) + p(a)) + r(c)p(a) \\ & \geq p(b)(r(d) + r(c) + r(a)) + p(d)(r(c) + r(a)) + p(c)r(a) \\ \iff & \frac{2}{3}r(b) + \frac{2}{3}r(d) + \frac{1}{3}r(c) \geq \frac{1}{3}r(d) + \frac{1}{3}r(c) + \frac{2}{3}r(a) \iff r(b) + \frac{1}{2}r(d) \geq r(a) \end{aligned}$$

Combining these two inequalities entails that $r(b) + \frac{1}{2}r(d) \geq \frac{3}{2}r(b) + r(d)$, which is true only if $r(b) = r(d) = 0$. Moreover, the second inequality upper bounds $r(a)$ and thus $r(a) = 0$. This means that $f(R'', c) = r(c) = 1$. However, c is the worst alternative of the voters $i \in [1 \dots \frac{n}{2}-1]$ and PC-strategyproofness hence requires that these voters cannot affect the outcome by misreporting their preferences. On the other hand, if we let these voters one after another change their preference relation to b, d, a, c , we arrive at a profile in which b is top-ranked by more than half of the voters. Hence, the absolute winner property requires that b is chosen with probability 1, which is in conflict with the observation that these voters are not able to affect the outcome. This is the desired contradiction. Hence, to complete the proof of Lemma 1b), it remains to show the claims for $f(R)$ and $f(R')$.

Claim 1: $f(R, a) = f(R, b) = f(R, c) = \frac{1}{3}$

Just as in the case of odd n , our first goal is to prove that $f(R, d) = 0$. As the first step in proving this statement, we show that $f(R, a) < 1$. Hence, assume for contradiction that $f(R, a) = 1$, which means that the least preferred lottery of voters $\frac{n}{2}$ and $\frac{n}{2} + 1$ is chosen. Moreover, if both of these voters swap b and c , c is top-ranked by more than half of the voters, so the absolute winner property requires c to receive a probability of 1. This is, however, in conflict with PC-strategyproofness, which requires that these voters cannot affect the outcome. Hence, the assumption that $f(R, a) = 1$ must have been wrong, i.e., $f(R, a) < 1$.

Based on this insight, we show by contradiction that $f(R, c) > 0$, i.e., suppose that $f(R, c) = 0$. Next, consider the profiles R^1 and R^2 shown below.

$$\begin{aligned} R^1: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, c, d, a & [\frac{n}{2}+2 \dots n-1]: c, a, d, b & n: a, c, d, b \\ R^2: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, c, d, a & [\frac{n}{2}+2 \dots n-2]: c, a, d, b & \{n-1, n\}: a, c, d, b \end{aligned}$$

Alternative a is top-ranked by $\frac{n}{2} + 1$ voters in R^2 , which means that $f(R^2, a) = 1$ because of the absolute winner property. Now, using PC-strategyproofness (or more precisely PC1-strategyproofness) from R^1 to R^2 , we derive that

$f(R^1, c) \geq f(R^1, d) + f(R^1, b)$. In particular, this inequality requires that $f(R^1, a) = 1$ if $f(R^1, c) = 0$. However, in that case, voter $n-1$ can PC-manipulate by deviating from R to R^1 : since $f(R, c) = 0$ and $f(R, a) < 1$, it follows that $f(R^1) \succ_{n-1}^{SD} f(R)$ and therefore also $f(R^1) \succ_{n-1}^{PC} f(R)$. Hence, it must hold that $f(R^1, c) > 0$. Next, we use PC-strategyproofness from R to R^1 and vice versa to derive the following two inequalities, where $p = f(R)$ and $q = f(R^1)$.

$$\begin{aligned} p(c)(q(a) + q(d) + q(b)) + p(a)(q(d) + q(b)) + p(d)q(b) \\ \geq q(c)(p(a) + p(d) + p(b)) + q(a)(p(d) + p(b)) + q(d)p(b) \\ q(a)(p(c) + p(d) + p(b)) + q(c)(p(d) + p(b)) + q(d)p(b) \\ \geq p(a)(q(c) + q(d) + q(b)) + p(c)(q(d) + q(b)) + p(d)q(b) \end{aligned}$$

Adding these two inequalities and cancelling common terms yields $p(c)q(a) \geq q(c)p(a)$. Since $p(c) = 0$ by assumption and $q(c) > 0$ because of our previous analysis, this inequality can only be true if $p(a) = 0$. Using the facts that $p(c) = p(a) = 0$ and $q(c) \geq q(b) + q(d)$, we can therefore vastly simplify the first inequality.

$$p(d)q(b) \geq q(c) + q(a) + q(d)p(b) \geq q(d) + q(b) + q(a) + q(d)p(b)$$

It is easy to see that this inequality can only be true if $f(R, d) = p(d) = 1$. We now derive a contradiction to this insight. First, from R , let voters $\frac{n}{2}$ and $\frac{n}{2} + 1$ make d into their favorite alternative. This leads to the profile R^3 (see below) and PC-strategyproofness (one step at a time) requires that $f(R^3, d) = 1$. Moreover, note that d Pareto-dominates b in R^3 . Next, we let voters $n-1$ and n swap a and c to obtain the profile R^4 . PC-efficiency requires that $f(R^4, b) = 0$ as this alternative is still Pareto-dominated, and PC-strategyproofness requires in turn that $f(R^4, d) = 1$ as any other lottery with support $\{a, c, d\}$ yields a PC-manipulation for voters $n-1$ and n . However, this contradicts the absolute winner property as $\frac{n}{2} + 1$ voters report a as their favorite alternative in R^4 , so it cannot be the case that $f(R, d) = 1$. Thus, no feasible outcome for $f(R)$ remains, which demonstrates that the assumption that $f(R, c) = 0$ is wrong. That is, we must have $f(R, c) > 0$.

$$\begin{aligned} R^3: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: d, b, c, a & [\frac{n}{2}+2 \dots n]: c, a, d, b \\ R^4: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: d, b, c, a & [\frac{n}{2}+2 \dots n-2]: c, a, d, b & \{n-1, n\}: a, c, d, b \end{aligned}$$

As the next step, we can use an analogous argument as in Lemma 1a) to derive that $f(R, d) = 0$ due to PC-efficiency. Indeed, this follows immediately since the profile R here and in the proof of Lemma 1a) consists of the same voter types and $f(R, c) > 0$.

Based on this insight, we show now that $f(R, c) \geq f(R, b) \geq f(R, a) \geq f(R, c)$, which implies that $f(R, a) = f(R, b) = f(R, c) = \frac{1}{3}$. For the first inequality, consider the profiles R^5 and R^6 .

$$\begin{aligned} R^5: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, c, d, a & [\frac{n}{2}+2 \dots n-1]: c, a, d, b & n: a, c, b, d \\ R^6: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, c, d, a & [\frac{n}{2}+2 \dots n-2]: c, a, d, b & \{n-1, n\}: a, c, b, d \end{aligned}$$

Lemma 4 and PC-efficiency imply that $f(R^5, d) = 0$ as we can otherwise find a lottery that PC-dominates $f(R^5)$ by redistributing probability from d to a and b . Moreover, the absolute winner property shows that $f(R^6, a) = 1$. In particular, $f(R, d) = f(R^5, d) = f(R^6, d) = 0$ and we can thus apply Lemma 3 twice to derive that $f(R^5, c) \geq f(R^5, b)$ and $f(R, c) \geq f(R, b)$ since $f(R^6, c) = f(R^6, b) = 0$.

Next, we show that $f(R, b) \geq f(R, a)$. Consider for this the profiles R^7 and R^8 derived from R by replacing the preference relations of voters $\frac{n}{2}$ and $\frac{n}{2} + 1$ sequentially with c, a, b, d .

$$\begin{aligned} R^7: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \frac{n}{2}: c, a, b, d & \frac{n}{2}+1: b, c, d, a & [\frac{n}{2}+2 \dots n]: c, a, d, b \\ R^8: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: c, a, b, d & [\frac{n}{2}+2 \dots n]: c, a, d, b \end{aligned}$$

First, observe that $f(R^8, c) = 1$ as all voters $i \in [\frac{n}{2} \dots n]$ report c as their best choice. Furthermore, Lemma 4 and PC-efficiency show that every lottery q with $q(d) > 0$ is PC-inefficient in R^7 because we can find a PC-improvement by redistributing the probability of d to a and b . Hence, $f(R, d) = f(R^7, d) = f(R^8, d) = 0$ and applying Lemma 3 twice then shows that $f(R, b) \geq f(R, a)$ because $f(R^8, b) = f(R^8, a) = 0$.

Finally, we prove that $f(R, a) \geq f(R, c)$. Consider for this the profiles \bar{R}^k for $k \in \{0, \dots, \frac{n}{2} - 1\}$ defined as follows.

$$\bar{R}^k: [1 \dots k]: a, d, b, c \quad [k+1 \dots \frac{n}{2}-1]: b, a, d, c \quad \{\frac{n}{2}, \frac{n}{2}+1\}: b, c, d, a \quad [\frac{n}{2}+2 \dots n]: c, a, d, b$$

Note that $R = \bar{R}^{\frac{n}{2}-1}$ and that $f(\bar{R}^0, b) = 1$ because of the absolute winner property. Moreover, $f(\bar{R}^k, d) = 0$ for all $k \in \{0, \dots, \frac{n}{2} - 2\}$ because of PC-efficiency: once again, Lemma 4 shows that any lottery q with $q(d) > 0$ is PC-dominated by the lottery p with $p(a) = q(a) + \frac{\varepsilon}{q(d)+q(a)}$, $p(b) = q(b) + \frac{\varepsilon}{q(d)+q(b)}$, $p(c) = q(c)$, and $p(d) = q(d) - \frac{\varepsilon}{q(d)+q(a)} - \frac{\varepsilon}{q(d)+q(b)}$ (where $\varepsilon > 0$ is so small that $p(d) \geq 0$). Hence, d receives probability 0 for all of these profiles. We also know that $f(R, d) = 0$, so $f(\bar{R}^k, d) = 0$ for all $k \in \{0, \dots, \frac{n}{2} - 1\}$. By inductively applying Lemma 3, we derive that $f(R, a) \geq f(R, c)$ because $f(\bar{R}^0, a) = f(\bar{R}^0, c) = 0$. This completes the proof of Claim 1.

Claim 2: $f(R', a) = f(R', c) = f(R', d) = \frac{1}{3}$

For proving this claim, we show as the first step that $f(R', b) = 0$. Note for this that an analogous argument as in Claim 1 proves that $f(R', c) > 0$. Based on this insight, an analogous argument as in the proof of Lemma 1a shows that $f(R', b) = 0$ because of PC-efficiency and Lemma 4. Indeed, this is straightforward as the profile R' here and in the proof of Lemma 1a consists of the same voter types and thus, the same lotteries are PC-efficient.

Using this observation, we show next that $f(R', c) \geq f(R', d) \geq f(R', a) \geq f(R', c)$, which implies that all three alternatives receive a probability of $\frac{1}{3}$. First, we prove that $f(R', c) \geq f(R', d)$ by considering the profiles R^1 and R^2 .

$$\begin{array}{llll} R^1: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, d, c, a & [\frac{n}{2}+2 \dots n-1]: c, a, d, b \quad n: a, d, c, b \\ R^2: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: b, d, c, a & [\frac{n}{2}+2 \dots n-2]: c, a, d, b \quad \{n-1, n\}: a, d, c, b \end{array}$$

Note that $f(R^2, a) = 1$ because more than half of the voters rank a top in R^2 . Consequently, PC1-strategyproofness entails that $f(R^1, c) \geq f(R^1, d) + f(R^1, b)$. Hence, if $f(R^1, c) = 0$, then $f(R^1, a) = 1$ and another application of PC1-strategyproofness shows that $f(R', c) \geq f(R', d)$ because $f(R', b) = 0$. On the other hand, if $f(R^1, c) > 0$, PC-efficiency requires that $f(R^1, b) = 0$. In more detail, every lottery with $q(a) > 0$, $q(b) > 0$, and $q(c) > 0$ is PC-dominated in R^1 by the lottery p with $p(a) = q(a) - \frac{\varepsilon}{q(a)+q(d)}$, $p(b) = q(b) - \frac{\varepsilon}{q(b)+q(d)}$, $p(c) = q(c)$, and $p(d) = q(d) + \frac{\varepsilon}{q(a)+q(d)} + \frac{\varepsilon}{q(b)+q(d)}$, whereas every lottery q with $q(a) = 0$, $q(b) > 0$, and $q(c) > 0$ is PC-dominated in R^1 by the lottery p with $p(a) = q(a)$, $p(b) = q(b) - \frac{\varepsilon}{q(b)+q(d)}$, $p(c) = q(c) - \frac{\varepsilon}{q(c)+q(d)}$, and $p(d) = q(d) + \frac{\varepsilon}{q(b)+q(d)} + \frac{\varepsilon}{q(c)+q(d)}$ (see Lemma 4). Hence, we have $f(R', b) = f(R^1, b) = f(R^2, b) = 0$ and a repeated application of Lemma 3 shows that $f(R', c) \geq f(R', d)$ because $f(R^2, c) = f(R^2, d) = 0$.

Next, we derive that $f(R', d) \geq f(R', a)$. To this end, consider the profiles R^3 and R^4 derived from R' by replacing the preference relations of voters $\frac{n}{2}$ and $\frac{n}{2} + 1$ with c, a, d, b .

$$\begin{array}{llll} R^3: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \frac{n}{2}: c, a, d, b & \frac{n}{2}+1: b, d, c, a \quad [\frac{n}{2}+2 \dots n]: c, a, d, b \\ R^4: & [1 \dots \frac{n}{2}-1]: a, d, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: c, a, d, b & [\frac{n}{2}+2 \dots n]: c, a, d, b \end{array}$$

It follows from the absolute winner property that $f(R^4, c) = 1$ as more than half of the voters report c as their best alternative. Moreover, we can use the same construction as for R (in Claim 1) to derive that $f(R^3, c) > 0$. Indeed, voters n and $n-1$ have the same preference relations in R and R^3 and they also can make a into the absolute winner by swapping a and c in R^3 . Analogously to R' , it now follows from PC-efficiency and Lemma 4 that $f(R^3, b) = 0$. Finally, a repeated application of Lemma 3 shows that $f(R', d) \geq f(R', a)$ because $f(R', b) = f(R^3, b) = f(R^4, b) = 0$ and $f(R^4, d) = f(R^4, a) = 0$.

It remains to show that $f(R', a) \geq f(R', c)$. Consider the sequence of profiles \hat{R}^k for $k \in \{0, \dots, \frac{n}{2} - 1\}$ and note that $R' = \hat{R}^{\frac{n}{2}-1}$, which means that $f(\hat{R}^{\frac{n}{2}-1}, b) = 0$.

$$\hat{R}^k: \quad [1 \dots k]: a, d, b, c \quad [k+1 \dots \frac{n}{2}-1]: d, a, b, c \quad \{\frac{n}{2}, \frac{n}{2}+1\}: b, d, c, a \quad [\frac{n}{2}+2 \dots n]: c, a, d, b$$

First, we show that $f(\hat{R}^k, b) = 0$ for all other profiles \hat{R}^k with $k \in \{0, \dots, \frac{n}{2} - 2\}$. To this end, note that Lemma 4 shows that either $f(\hat{R}^k, b) = 0$ or $f(\hat{R}^k, a) = 0$ because otherwise, we can find a PC-improvement by redistributing probability from a and b to d . Moreover, if $f(\hat{R}^k, a) = 0$, then Lemma 4 entails that $f(\hat{R}^k, b) = 0$ or $f(\hat{R}^k, c) = 0$, because otherwise the probability from b and c can be redistributed to d . Now, assume for contradiction that $f(\hat{R}^k, b) > 0$ for a fixed k and hence $f(\hat{R}^k, a) = f(\hat{R}^k, c) = 0$. We proceed with a case distinction on k . First, suppose that $\frac{n}{2} - 2 \geq k \geq 2$, which means that at least two voters in $[1 \dots k]$ top-rank a . For this case, we consider the profiles $\hat{R}^{k,j}$ for $j \in \{\frac{n}{2} + 1, \dots, n\}$.

$$\hat{R}^{k,j}: \quad [1 \dots k]: a, d, b, c \quad [k+1 \dots \frac{n}{2}-1]: d, a, b, c \quad \{\frac{n}{2}, \frac{n}{2}+1\}: b, d, c, a \quad [\frac{n}{2}+2 \dots j]: a, d, b, c \quad [j+1 \dots n]: c, a, d, b$$

It holds by definition that $\hat{R}^k = \hat{R}^{k, \frac{n}{2}+1}$. Moreover, analogous to \hat{R}^k , PC-efficiency requires that either $f(\hat{R}^{k,j}, a) + f(\hat{R}^{k,j}, c) = 0$ or $f(\hat{R}^{k,j}, b) = 0$ for every $j \in \{\frac{n}{2} + 1, \dots, n\}$. This implies for every j that if $f(\hat{R}^{k,j}, b) > 0$, then $f(\hat{R}^{k,j+1}, b) > 0$. In more detail, $f(\hat{R}^{k,j}, b) > 0$ requires that $f(\hat{R}^{k,j}, b) + f(\hat{R}^{k,j}, d) = 1$ because of PC-efficiency. Hence, if $f(\hat{R}^{k,j+1}, b) = 0$, voter $j+1$ can PC-manipulate by deviating from $\hat{R}^{k,j}$ to $\hat{R}^{k,j+1}$ since b and d are his worst two alternatives in $\hat{R}^{k,j}$. By repeatedly applying this argument, it follows that if $f(\hat{R}^k, b) > 0$, then $f(\hat{R}^{k,n}, b) > 0$. However, this is in conflict with the absolute winner property as $\frac{n}{2} - 1 + k$ voters top-rank a in $\hat{R}^{k,n}$. This proves that $f(\hat{R}^k, b) = 0$ for all $k \in \{2, \dots, \frac{n}{2} - 2\}$.

Note that the argument above fails if $k \leq 1$ as no more than $\frac{n}{2}$ voters top-rank a in $\hat{R}^{1,n}$ and $\hat{R}^{0,n}$. Hence, we investigate the case $k \leq 1$ separately and consider for this the profiles $\tilde{R}^{k,j}$ for $j \in \{\frac{n}{2} + 1, \dots, n\}$.

$$\tilde{R}^{k,j}: \quad [1 \dots k]: a, d, b, c \quad [k+1 \dots \frac{n}{2}-1]: d, a, b, c \quad \{\frac{n}{2}, \frac{n}{2}+1\}: b, d, c, a \quad [\frac{n}{2}+2 \dots j]: d, a, b, c \quad [j+1 \dots n]: c, a, d, b$$

It holds by definition that $\hat{R}^k = \tilde{R}^{k, \frac{n}{2}+1}$. Note that the profiles $\tilde{R}^{k,j}$ consist of the same preference relations as \hat{R}^k and hence, PC-efficiency once again requires that $f(\tilde{R}^{k,j}, b) = 0$ or $f(\tilde{R}^{k,j}, a) = f(\tilde{R}^{k,j}, c) = 0$. (When $j = n$, even though the preference relation c, a, d, b is not present, a similar argument still holds.) Moreover, if $f(\tilde{R}^{k,j}, b) > 0$, then $f(\tilde{R}^{k,j+1}, b) > 0$. The reason for this is that if $f(\tilde{R}^{k,j}, b) > 0$, then $f(\tilde{R}^{k,j}, b) + f(\tilde{R}^{k,j}, d) = 1$. Hence, if $f(\tilde{R}^{k,j+1}, b) = 0$, voter $j+1$ can PC-manipulate by deviating from $\tilde{R}^{k,j}$ to $\tilde{R}^{k,j+1}$ as b and d are his least preferred alternatives in $\tilde{R}^{k,j}$. By repeatedly applying this argument, we derive that if $f(\hat{R}^k, b) > 0$, then $f(\tilde{R}^{k,n}, b) > 0$. However, this contradicts the absolute winner property

as at least $(\frac{n}{2} - 1 - k) + (\frac{n}{2} - 1) \geq n - 3 > \frac{n}{2}$ voters top-rank d in $\tilde{R}^{k,n}$. (Here we use the assumption that $n \geq 8$.) Hence, we also have that $f(\hat{R}^k, b) = 0$ if $k \leq 1$.

As the last point, we prove that $f(\hat{R}^0, d) = 1$. Then, it follows from repeated application of Lemma 3 that $f(R', a) \geq f(R', c)$ because $f(\hat{R}^0, a) = f(\hat{R}^0, c) = 0$. Thus, consider the profiles R^5 and R^6 derived from \hat{R}^0 by sequentially replacing the preferences of voter $\frac{n}{2}$ and $\frac{n}{2} + 1$ with d, a, b, c .

$$\begin{array}{llll} R^5: & [1 \dots \frac{n}{2}-1]: d, a, b, c & \frac{n}{2}: d, a, b, c & \frac{n}{2}+1: b, d, c, a & [\frac{n}{2}+2 \dots n]: c, a, d, b \\ R^6: & [1 \dots \frac{n}{2}-1]: d, a, b, c & \{\frac{n}{2}, \frac{n}{2}+1\}: d, a, b, c & [\frac{n}{2}+2 \dots n]: c, a, d, b \end{array}$$

Note that an absolute majority top-ranks d in R^6 , which means that $f(R^6, d) = 1$. Hence, $PC1$ -strategyproofness entails for R^5 that $f(R^5, b) \geq f(R^5, c) + f(R^5, a)$. If $f(R^5, b) = 0$, we derive then that $f(R^5, d) = 1$, and an application of $PC1$ -strategyproofness between \hat{R}^0 and R^5 shows that $f(\hat{R}^0, b) \geq f(\hat{R}^0, c) + f(\hat{R}^0, a)$. Since we already established that $f(\hat{R}^0, b) = 0$, this proves that $f(\hat{R}^0, d) = 1$. On the other hand, if $f(R^5, b) > 0$, PC -efficiency requires that $f(R^5, a) = f(R^5, c) = 0$ (see Lemma 4). However, then voter n can PC -manipulate in R^5 by reporting d as his favorite alternative. Thereafter, d must be chosen with probability 1 because it is top-ranked by $\frac{n}{2} + 1$ voters. However, voter n PC -prefers this lottery to $f(R^5)$ if $f(R^5, b) > 0$, which contradicts the PC -strategyproofness of f . Hence, it must indeed hold that $f(R^5, b) = 0$ and therefore also $f(R^5, d) = 1$ and $f(\hat{R}^0, d) = 1$. Finally, as mentioned earlier in this paragraph, a repeated application of Lemma 3 shows now that $f(R', a) \geq f(R', c)$. Therefore, we have that $f(R', c) \geq f(R', d) \geq f(R', a) \geq f(R', c)$ and $f(R', b) = 0$, which implies that $f(R', a) = f(R', c) = f(R', d) = \frac{1}{3}$. \square

Since the conjunction of Lemma 1a) and Lemma 1b) is equivalent to Lemma 1, this concludes the first step for the proof of Theorem 2. Next, we show that every anonymous, neutral, PC -efficient, and PC -strategyproof SDS also satisfies the absolute winner property. This insight together with Lemma 1 immediately implies Theorem 2. Note that the subsequent lemma is slightly stronger than required: we show that the implication also holds for $m \geq 3$ and all electorates.

Lemma 5. Assume that $m \geq 3$. Every SDS that satisfies PC -efficiency, PC -strategyproofness, neutrality, and anonymity also satisfies the absolute winner property.

Proof. Let f denote an SDS that satisfies anonymity, neutrality, PC -efficiency, and PC -strategyproofness for $m \geq 3$ alternatives. First, note that for electorates N with $n = |N| \leq 2$, the absolute winner property requires that an alternative is chosen with probability 1 if it is top-ranked by all voters. This is clearly implied by PC -efficiency and we thus focus on the case that $n \geq 3$. Since the construction in the main body (in the proof of Theorem 2) works for every even $n \geq 4$, we only need to show the lemma for the case that the number of voters is odd. Hence, consider an electorate N with an odd number of voters $n \geq 3$. Moreover, we focus on the case that there are $m = 3$ alternatives because we can generalize all steps by simply adding additional alternatives at the bottom of all preference rankings. PC -efficiency then requires that these alternatives get probability 0 and they thus do not affect our analysis.

We start our analysis by considering the profiles R^1 and R^2 described below.

$$\begin{array}{llll} R^1: & [1 \dots \frac{n-1}{2}]: b, a, c & \frac{n+1}{2}: a, b, c & [\frac{n+3}{2} \dots n]: c, a, b \\ R^2: & [1 \dots \frac{n-1}{2}]: b, a, c & \frac{n+1}{2}: a, c, b & [\frac{n+3}{2} \dots n]: c, a, b \end{array}$$

First, note that anonymity and neutrality imply that $f(R^1, a) = f(R^2, a)$, $f(R^1, b) = f(R^2, c)$, and $f(R^1, c) = f(R^2, b)$. Furthermore, PC -efficiency shows that $f(R^1, b) = f(R^2, c) = 0$ or $f(R^1, c) = f(R^2, b) = 0$. Subsequently, we show that $f(R^1, b) = f(R^1, c) = 0$ must be true, which means that $f(R^1, a) = 1$.

Assume for contradiction that $f(R^1, c) = f(R^2, b) > 0$. Then, our previous observation implies that $f(R^2, c) = f(R^1, b) = 0$. However, this means that voter $\frac{n+1}{2}$ can manipulate by deviating from R^1 to R^2 because he PC -prefers $f(R^2)$ to $f(R^1)$ (he even SD -prefers $f(R^2)$ to $f(R^1)$). Hence, f is PC -manipulable if $f(R^1, c) > 0$, contradicting our assumptions.

As the second case, assume that $f(R^1, b) = f(R^2, c) > 0$ (note that this is *not* symmetric to the case studied in the previous paragraph) and consider the following sequence of preference profiles \bar{R}^i for $i \in \{0, \dots, \frac{n-1}{2}\}$.

$$\bar{R}^i: [1 \dots \frac{n-1}{2}]: b, a, c \quad \frac{n+1}{2}: a, b, c \quad [\frac{n+3}{2} \dots n-i]: c, a, b \quad [n-i+1 \dots n]: a, c, b$$

First, note that $R^1 = \bar{R}^0$ and PC -efficiency shows for all profiles \bar{R}^i that $f(\bar{R}^i, b) = 0$ or $f(\bar{R}^i, c) = 0$. Moreover, PC -strategyproofness and PC -efficiency imply that if $f(\bar{R}^i, b) > 0$, then $f(\bar{R}^{i+1}) = f(\bar{R}^i)$. The reason for this is that if $f(\bar{R}^i, b) > 0$, then $f(\bar{R}^i, c) = 0$ because of PC -efficiency. This means that every lottery with $f(\bar{R}^{i+1}, b) = 0$ is a PC -manipulation for the deviating voter $n - i$ as he even SD -prefers $f(\bar{R}^{i+1})$ to $f(\bar{R}^i)$. Hence, $f(\bar{R}^{i+1}, b) > 0$, and we can now use PC -efficiency to derive that $f(\bar{R}^{i+1}, c) = f(\bar{R}^i, c) = 0$. Finally, Lemma 2 implies that $f(\bar{R}^{i+1}) = f(\bar{R}^i)$. As a consequence, this sequence ends at a profile $R^3 = \bar{R}^{\frac{n-1}{2}}$ with $f(R^3) = f(R^1)$.

Next, consider the profile R^4 which is derived from R^3 by swapping b and c in the preference relation of voter $\frac{n+1}{2}$.

$$R^4: [1 \dots \frac{n-1}{2}]: b, a, c \quad \frac{n+1}{2}: a, c, b \quad [\frac{n+3}{2} \dots n]: a, c, b$$

Since a Pareto-dominates c in R^4 , it follows that $f(R^4, c) = 0$. Hence, we can use again Lemma 2 to conclude that $f(R^4) = f(R^3) = f(R^1)$.

As the last step, consider the sequence of profiles \hat{R}^i for $i \in \{0, \dots, \frac{n-1}{2}\}$, which leads from R^4 to R^2 .

$$\hat{R}^i: \quad [1 \dots \frac{n-1}{2}]: b, a, c \quad \frac{n+1}{2}: a, c, b \quad [\frac{n+3}{2} \dots n-i]: a, c, b \quad [n-i+1 \dots n]: c, a, b$$

First, observe that $\hat{R}^0 = R^4$ and $\hat{R}^{\frac{n-1}{2}} = R^2$. Moreover, PC-efficiency requires again for every profile \hat{R}^i that either $f(\hat{R}^i, b) = 0$ or $f(\hat{R}^i, c) = 0$. Even more, since $f(\hat{R}^0, c) = 0$ and $f(\hat{R}^{\frac{n-1}{2}}, b) = 0$, there is at least one index i such that $f(\hat{R}^i, c) = f(\hat{R}^{i+1}, b) = 0$. Let $i^* \in \{0, \dots, \frac{n-3}{2}\}$ denote the smallest such index, which means that $f(\hat{R}^i, c) = 0$ for all $i \in \{0, \dots, i^*\}$. Therefore, we can again use Lemma 2 to conclude that $f(\hat{R}^{i^*}) = f(\hat{R}^0) = f(R^1)$, which means in particular that $f(\hat{R}^{i^*}, b) = f(R^1, b) > 0$. Now, if $f(\hat{R}^{i^*+1}, a) \geq f(R^1, a)$, voter $n - i^*$ can PC-manipulate by deviating from \hat{R}^{i^*} to \hat{R}^{i^*+1} . This follows as voter $n - i^*$, whose preference is a, c, b in R^{i^*} , SD-prefers (and therefore also PC-prefers) $f(\hat{R}^{i^*+1})$ to $f(\hat{R}^{i^*})$ in this case. Hence, PC-strategyproofness requires that $f(\hat{R}^{i^*+1}, a) < f(\hat{R}^{i^*}, a)$. Since $f(\hat{R}^{i^*+1}, b) = f(\hat{R}^{i^*}, c) = 0$, this implies that $f(\hat{R}^{i^*+1}, c) > f(\hat{R}^{i^*}, b) = f(R^1, b)$.

Next, we prove that $f(\hat{R}^{i+1}, c) \geq f(\hat{R}^i, c)$ for all $i > i^*$. Assume for contradiction that there is an index j where this is not the case. Then, there is also a minimal index $j^* > i^*$ such that $f(\hat{R}^{j^*+1}, c) < f(\hat{R}^{j^*}, c)$. In particular, it follows from the minimality of j^* that $f(\hat{R}^{j^*}, c) \geq f(\hat{R}^{i^*+1}, c) > 0$ and PC-efficiency then shows that $f(\hat{R}^{j^*}, b) = 0$. Now, note that voter $n - j^*$'s preference relation in \hat{R}^{j^*+1} is c, a, b . Hence, if $f(\hat{R}^{j^*+1}, c) = 0$, he clearly PC-prefers $f(\hat{R}^{j^*})$ to $f(\hat{R}^{j^*+1})$. This means that $f(\hat{R}^{j^*+1}, c) > 0$ and consequently $f(\hat{R}^{j^*+1}, b) = 0$ because of PC-efficiency. However, then voter j^* still PC-prefers $f(\hat{R}^{j^*})$ to $f(\hat{R}^{j^*+1})$ because $f(\hat{R}^{j^*+1}, c) < f(\hat{R}^{j^*}, c)$. Hence, voter j^* can either way PC-manipulate by deviating from \hat{R}^{j^*+1} to \hat{R}^{j^*} . This contradicts the PC-strategyproofness of f , and so we must have $f(\hat{R}^{i+1}, c) \geq f(\hat{R}^i, c)$ for all $i > i^*$. In particular, this implies that $f(R^2, c) \geq f(\hat{R}^{i^*+1}, c) > f(\hat{R}^{i^*}, b) = f(R^1, b)$ because $R^2 = \hat{R}^{\frac{n-1}{2}}$. However, this observation is in conflict with anonymity and neutrality between R^2 and R^1 , and thus, the assumption that $f(R^1, b) > 0$ must be wrong. It follows that $f(R^1, b) = f(R^1, c) = 0$, and so $f(R^1, a) = 1$.

Finally, departing from the insight that $f(R^1, a) = 1$, we can essentially apply the same steps as in the proof for even n (in the main body) to show that f must satisfy the absolute winner property. \square

A.3. Proofs of Propositions 3 and 4

In this subsection, we prove Propositions 3 and 4, which show that the impossibilities for $m \geq 4$ in Theorems 1 to 3 turn into possibilities when $m = 3$. We first provide additional insights into PC-efficiency that facilitate the analysis of f^1 and f^2 .

Lemma 6. Consider a profile $R \in \mathcal{R}^*$ on three alternatives $A = \{a, b, c\}$. A lottery p is PC-efficient for R if it satisfies the following conditions.

1. $p(x) = 0$ if x is Pareto-dominated in R .
2. For an alternative $x \in A$ that is never bottom-ranked and at least once top-ranked in R , there is $y \in A \setminus \{x\}$ with $p(y) = 0$.
3. For an alternative $x \in A$ that is never top-ranked and at least once bottom-ranked in R , $p(x) = 0$.

Proof. Consider an arbitrary electorate $N \in \mathcal{F}(\mathbb{N})$ and let $R \in \mathcal{R}^N$ denote a profile. Moreover, let p denote a lottery that satisfies the given conditions and suppose for contradiction that there is another lottery q that PC-dominates p on R . Hence, $p \neq q$, which means that there are alternatives x, y such that $q(x) > p(x)$ and $q(y) < p(y)$. We suppose subsequently that $q(a) > p(a)$ and $q(b) < p(b)$ because our arguments are completely symmetric. Next, we proceed with a case distinction with respect to the relation between $q(c)$ and $p(c)$.

Case 1: $q(c) = p(c)$

As the first case, we suppose that $q(c) = p(c)$. Then, it follows for all voters $i \in N$ that $a \succ_i b$ because $q \succ_i^{PC} p$. Indeed, if $b \succ_i a$ for some $i \in N$, this voter strictly SD-prefers p to q . Since q PC-dominates p by assumption, it thus follows that a Pareto-dominates b . However, condition 1 then requires that $p(b) = 0$, which contradicts that $q(b) < p(b)$. Hence, q cannot PC-dominate p in this case.

Case 2: $q(c) < p(c)$

Next, suppose that $q(c) < p(c)$. Combined with $q(a) > p(a)$, $q(b) < p(b)$, and the assumption that q PC-dominates p , this means that no voter bottom-ranks a . Indeed, it is easy to see that such a voter strictly SD-prefers p to q , contradicting the PC-dominance of q . Now, if a is top-ranked by a voter in R , then condition 2 requires that either $p(c) = 0$ or $p(b) = 0$. However, this is not possible since $q(c) < p(c)$ and $q(b) < p(b)$ by assumption. Hence, every voter ranks a at the second position in R .

Furthermore, both b and c must be top-ranked at least once; otherwise, one of these alternatives is unanimously top-ranked and therefore Pareto-dominates both other alternatives, which again conflicts with $q(c) < p(c)$ and $q(b) < p(b)$.

Hence, every voter in R has the preference relation b, a, c or c, a, b , and each of these two preferences is submitted at least once. Voters of the first type PC-prefers q to p if $q(b)p(a) + q(b)p(c) + q(a)p(c) \geq p(b)q(a) + p(b)q(c) + p(a)q(c)$ and voters of the second type if $q(c)p(a) + q(c)p(b) + q(a)p(b) \geq p(c)q(a) + p(c)q(b) + p(a)q(b)$. Clearly, both inequalities are

only true if they hold with equality. However, then no voter strictly *PC*-prefers q to p and hence q does not *PC*-dominate p in this case either.

Case 3: $q(c) > p(c)$

As the last case, we assume that $q(c) > p(c)$. Since $q(a) > p(a)$, $q(b) < p(b)$, and q *PC*-dominates p in R , no voter top-ranks b in R . Indeed, such a voter strictly *SD*-prefers p to q , which contradicts that q *PC*-dominates p . Now, if b is bottom-ranked by at least one voter in R , then condition 3 requires that $p(b) = 0$, which contradicts $q(b) < p(b)$. Hence, b is second-ranked by all voters in R . Next, both a and c are top-ranked at least once in R ; otherwise, b is Pareto-dominated which again contradicts $q(b) < p(b)$. Hence, every voter in R has the preference a, b, c or c, b, a , and both preferences are reported at least once. An analogous argument as in Case 2 implies that q cannot *PC*-dominate p in this case either. \square

Next, we use Lemma 6 to prove Proposition 3. Recall the definition of f^1 (where $CW(R)$ is the set of Condorcet winners in R and $WCW(R)$ the set of weak Condorcet winners).

$$f^1(R) = \begin{cases} [x : 1] & \text{if } CW(R) = \{x\} \\ [x : \frac{1}{2}; y : \frac{1}{2}] & \text{if } WCW(R) = \{x, y\} \\ [x : \frac{3}{5}; y : \frac{1}{5}; z : \frac{1}{5}] & \text{if } WCW(R) = \{x\} \\ [x : \frac{1}{3}; y : \frac{1}{3}; z : \frac{1}{3}] & \text{otherwise} \end{cases}$$

Proposition 3. For $m = 3$, f^1 is the only anonymous and neutral SDS that satisfies *PC*-efficiency, *PC*-strategyproofness, and cancellation.

Proof. The proposition consists of two claims: on the one hand, we need to show that f^1 satisfies all axioms of the proposition, and on the other hand, that f^1 is the only SDS satisfying these axioms. We consider both claims separately and start by showing that f^1 satisfies all axioms of the proposition.

Claim 1: f^1 satisfies anonymity, neutrality, cancellation, *PC*-efficiency and *PC*-strategyproofness.

First, note that f^1 satisfies cancellation because adding two voters with inverse preferences does not affect whether an alternative is a (weak) Condorcet winner. Furthermore, the definition of f^1 immediately shows that it is anonymous and neutral.

For proving that f^1 is *PC*-efficient, we consider an arbitrary preference profile $R \in \mathcal{R}^*$. Now, if an alternative x is Pareto-dominated in R , then it is never top-ranked. Consequently, there is either a Condorcet winner $y \neq x$ (if more than half of the voters top-rank y) or the remaining two alternatives y, z are weak Condorcet winners (if both y and z are top-ranked by exactly half of the voters). In both cases, $f^1(R, x) = 0$, which shows that $f^1(R)$ satisfies condition 1 of Lemma 6 for all profiles R . Similarly, if there is an alternative x that is never top-ranked and at least once bottom-ranked, then either there is a Condorcet winner $y \neq x$, or the remaining two alternatives y, z are weak Condorcet winners. Hence, $f^1(R, x) = 0$, which proves that $f^1(R)$ also satisfies condition 3 of Lemma 6. Finally, if there is an alternative x that is never bottom-ranked and at least once top-ranked in R , then there is an alternative $z \neq x$ with $g_R(x, z) > 0$. We claim that $f^1(R, z) = 0$, which proves condition 2 of Lemma 6. If x or the third alternative y is a Condorcet winner, this follows immediately. On the other hand, if neither x nor y are Condorcet winners, both of them are weak Condorcet winners because y must be top-ranked by at least half of the voters if x is not a Condorcet winner. Hence, we have two weak Condorcet winners and the definition of f^1 again shows that $f^1(R, z) = 0$. Since all conditions of Lemma 6 hold, it thus follows that $f^1(R)$ is *PC*-efficient for all profiles R .

Finally, we need to show that f^1 is *PC*-strategyproof. Assume for contradiction that this is not the case. Then, there are an electorate N , two preference profiles $R, R' \in \mathcal{R}^N$, and a voter $i \in N$ such that $f^1(R') \succ_i^{PC} f^1(R)$ in R , and $R_{-i} = R'_{-i}$. Subsequently, we discuss a case distinction with respect to the definition of f^1 . In more detail, we have for both R and R' five different options: there is a Condorcet winner (CW), or there is no Condorcet winner but $k \in \{0, 1, 2, 3\}$ weak Condorcet winners ($kWCW$). We label the cases with a shorthand notation: for instance, $CW \rightarrow 1WCW$ is the case where there is a Condorcet winner in R and a single weak Condorcet winner in R' .

To keep the length of the proof manageable, we subsequently focus only on the case that R and R' are defined by an odd number of voters. This assumption means that there are no weak Condorcet winners and thus significantly reduces the number of cases that need to be considered. For the case that R and R' are defined by an even number of voters, we refer to a preprint of this paper (Brandt et al., 2022c). When the number of voters n is odd, there are only four possible types of manipulations.

- $CW \rightarrow CW$: Suppose that a is the Condorcet winner in R . If a is also the Condorcet winner in R' , then $f^1(R) = f^1(R')$ and deviating from R to R' is no *PC*-manipulation. On the other hand, if another alternative b is the Condorcet winner in R' , we must have $a \succ_i b$ in R . Since $f^1(R, a) = f^1(R', b) = 1$, this is no *PC*-manipulation.
- $CW \rightarrow 0WCW$: Suppose that a is the Condorcet winner in R , and there is no Condorcet winner in R' . This means that voter i reinforces an alternative b against a , i.e., a is ranked either second or third. Since $f^1(R, a) = 1$ and $f^1(R', x) = \frac{1}{3}$ for all $x \in A$, this proves that deviating from R to R' is no *PC*-manipulation.

- $0WCW \rightarrow CW$: Suppose there is no Condorcet winner in R , but a is the Condorcet winner in R' . Hence, voter i needs to reinforce a against at least one other alternative b . This means that a is not voter i 's favorite alternative in R . Since $f^1(R, x) = \frac{1}{3}$ for all $x \in A$ and $f(R', a) = 1$, this observation proves that f^1 is PC -strategyproof in this case.
- $0WCW \rightarrow 0WCW$: We have $f^1(R) = f^1(R')$ in this case, which contradicts that voter i can PC -manipulate.

Claim 2: f^1 is the only SDS that satisfies anonymity, neutrality, cancellation, PC -efficiency, and PC -strategyproofness.

Consider an arbitrary SDS f for $m = 3$ alternatives that satisfies all given axioms. We show that $f(R) = f^1(R)$ for all profiles $R \in \mathcal{R}^*$, which proves this claim. For this, we name the six possible preference relations $\succ_1 = a, b, c$, $\succ_2 = c, b, a$, $\succ_3 = b, c, a$, $\succ_4 = a, c, b$, $\succ_5 = c, a, b$, and $\succ_6 = b, a, c$. Moreover, given a profile R , let n_i denote the number of voters who report preference relation \succ_i in R . Using this notation, we can describe the majority margins of R as follows.

$$\begin{aligned} g_R(a, b) &= (n_1 - n_2) - (n_3 - n_4) + (n_5 - n_6) \\ g_R(b, c) &= (n_1 - n_2) + (n_3 - n_4) - (n_5 - n_6) \\ g_R(c, a) &= -(n_1 - n_2) + (n_3 - n_4) + (n_5 - n_6) \end{aligned}$$

It is not difficult to derive from these equations that

$$n_1 = \frac{g_R(a, b) + g_R(b, c)}{2} + n_2, \quad n_3 = \frac{g_R(b, c) + g_R(c, a)}{2} + n_4, \quad n_5 = \frac{g_R(c, a) + g_R(a, b)}{2} + n_6.$$

Next, consider an arbitrary preference profile R . Based on cancellation, we can use the above equations to remove pairs of voters with inverse preferences from R until $n_{2k} = 0$ or $n_{2k-1} = 0$ for each $k \in \{1, 2, 3\}$. Unless all majority margins are 0, this leads to a minimal profile R' , which we consider in the subsequent case distinction. Note that the removal of voters with inverse preferences does not affect the majority margins and therefore also not the (weak) Condorcet winners. In particular, this means that $f^1(R) = f^1(R')$. Analogously, cancellation yields for f that $f(R) = f(R')$. Hence, we will consider multiple cases depending on the structure of R' and prove that $f(R) = f(R') = f^1(R') = f^1(R)$ in every case. On the other hand, if all majority margins are 0, we need a separate argument, which we discuss in our first case below. Taken together, our cases imply that $f(R) = f^1(R)$ for every profile R .

Case 2.1: $g_R(a, b) = g_R(b, c) = g_R(c, a) = 0$.

First, suppose that $g_R(a, b) = g_R(b, c) = g_R(c, a) = 0$, which means that all three alternatives are weak Condorcet winners in R . Our equations show that $n_1 = n_2$, $n_3 = n_4$, and $n_5 = n_6$. Let n^* denote the maximum among all n_i . Using cancellation, we can add pairs of voters with inverse preferences until $n_k = n^*$ for every $k \in \{1, \dots, 6\}$. Moreover, cancellation implies that $f(R) = f(R'')$ for the new profile R'' . Finally, all alternatives are symmetric to each other in R'' since all preference relations appear equally often. Hence, anonymity and neutrality require that $f(R'', x) = \frac{1}{3}$ for all $x \in A$, which means that $f(R) = f(R'') = f^1(R)$.

Case 2.2: An alternative x is top-ranked by more than half of the voters in R' .

As the second case, suppose that R' is well-defined and that an alternative x is top-ranked by more than half of the voters in this profile. Then, it holds that $f(R', x) = 1$ because Lemma 5 implies that f satisfies the absolute winner property. Since x is the Condorcet winner in R' , it holds that $f(R') = f^1(R')$.

Case 2.3: Two alternatives are top-ranked by exactly half of the voters in R' .

Next, suppose that R' is well-defined and that two alternatives, say a and b , are top-ranked by exactly half of the voters in R' . Then, a and b are weak Condorcet winners. Moreover, c is not a weak Condorcet winner in R' since not all majority margins in R can be 0, which implies that there is a voter who ranks c last. Due to symmetry, we can assume that this voter's preference relation is a, b, c . Now, if there is a voter with preference relation a, c, b in R' , then the last possible preference relation is b, a, c ; otherwise, R' is not minimal. Hence, a Pareto-dominates c in R' . Similarly, if there is no voter with the preference a, c, b , all voters prefer b to c and c is again Pareto-dominated. Therefore, it follows in both cases that $f(R', c) = 0$ because of PC -efficiency. Moreover, we can let the voters with a, c, b and b, c, a (if any) push down c . Then, c stays Pareto-dominated and therefore still receives probability 0 from f . Hence, Lemma 2 shows that the probability of a and b does not change during these steps. Finally, this process results in a profile R'' in which half of the voters report a, b, c and the other half b, a, c . Anonymity, neutrality, and PC -efficiency imply for this profile R'' that $f(R'', a) = f(R'', b) = \frac{1}{2}$. Hence, we have that $f(R') = f(R'') = f^1(R')$ because a and b are the only weak Condorcet winners in R' .

Case 2.4: Each alternative is top-ranked at least once and one alternative is top-ranked by exactly half of the voters in R' .

Next, suppose that an alternative is top-ranked by exactly half of the voters and the other two alternatives are top-ranked at least once. Without loss of generality, assume that there is a voter with preference relation a, b, c in R' . Since c is top-ranked by a voter, there is also a voter with preference relation c, a, b ; note for this that no voter can report c, b, a in R' because of the minimality of R' . By an analogous argument, we also derive that there is a voter with preference relation b, c, a . In summary, we have that $n_1 > 0$, $n_3 > 0$, $n_5 > 0$, and $n_2 = n_4 = n_6 = 0$. Moreover, one alternative is top-ranked by half of the voters; suppose without loss of generality that this alternative is a . Hence, $n_1 = n_3 + n_5$. We prove that $f(R', a) = \frac{3}{5}$ and $f(R', b) = f(R', c) = \frac{1}{5}$ by considering the following preference profiles, where $l = n_1 + n_3$.

R^{1,n_3,n_5} :	$[1 \dots n_1]: a, b, c$	$[n_1+1 \dots l]: b, c, a$	$[l+1 \dots n-1]: c, a, b$	$n: c, b, a$
R^{2,n_3,n_5} :	$[1 \dots n_1]: a, b, c$	$[n_1+1 \dots l]: b, c, a$	$[l+1 \dots n]: c, a, b$	
R^{3,n_3,n_5} :	$[1 \dots n_1]: a, b, c$	$[n_1+1 \dots l-1]: b, c, a$	$l: c, b, a$	$[l+1 \dots n]: c, a, b$

Anonymity implies that $f(R') = f(R^{2,n_3,n_5})$. Hence, our goal is to show that $f(R^{2,n_3,n_5}, a) = \frac{3}{5}$ and $f(R^{2,n_3,n_5}, b) = f(R^{2,n_3,n_5}, c) = \frac{1}{5}$ for all $n_3 > 0$ and $n_5 > 0$. Note for this that, in R^{1,n_3,n_5} and R^{3,n_3,n_5} , we can use cancellation to remove voters 1 and n or voters 1 and l , respectively. This step leads to the profile R^{2,n_3,n_5-1} or R^{2,n_3-1,n_5} , which proves that $f(R^{1,n_3,n_5}) = f(R^{2,n_3,n_5-1})$ and $f(R^{3,n_3,n_5}) = f(R^{2,n_3-1,n_5})$. Moreover, note that if $n_3 = 0$, then a and c are top-ranked by half of the voters in R^{2,n_3,n_5} . Hence, we have that $f(R^{2,0,n_5}, a) = f(R^{2,0,n_5}, c) = f(R^{3,1,n_5}, a) = f(R^{3,1,n_5}, c) = \frac{1}{2}$ by Case 2.3. An analogous argument also shows that $f(R^{2,n_3,0}, a) = f(R^{2,n_3,0}, b) = f(R^{1,n_3,1}, a) = f(R^{1,n_3,1}, b) = \frac{1}{2}$. Based on these insights, we now prove our claim on $f(R^{2,n_3,n_5})$ with an induction on $n_3 + n_5$.

First, we consider the induction basis that $n_3 = n_5 = 1$. The previous paragraph implies that $f(R^{1,n_3,n_5}, a) = f(R^{1,n_3,n_5}, b) = f(R^{3,n_3,n_5}, a) = f(R^{3,n_3,n_5}, c) = \frac{1}{2}$. Hence, PC-strategyproofness from R^{1,n_3,n_5} to R^{2,n_3,n_5} and from R^{2,n_3,n_5} to R^{3,n_3,n_5} entails the following inequalities, where $p = f(R^{2,n_3,n_5})$.

$$\frac{1}{2}p(a) \geq p(c) + \frac{1}{2}p(b) \quad p(b) + \frac{1}{2}p(c) \geq \frac{1}{2}p(a)$$

Moreover, note that voter n can ensure in R^{2,n_3,n_5} that a is chosen with probability 1 by reporting it as his favorite alternative because of the absolute winner property. Hence, we also get that $p(c) \geq p(b)$ from PC-strategyproofness. Finally, it is easy to see that these three inequalities are true at the same time only if $p(a) = 3p(b) = 3p(c)$. Using the fact that $p(a) + p(b) + p(c) = 1$, we hence derive that $p(a) = \frac{3}{5}$ and $p(b) = p(c) = \frac{1}{5}$.

Next, we prove the induction step and thus consider some fixed $n_3 > 0$ and $n_5 > 0$ such that $n_3 + n_5 > 2$. The induction hypothesis is that $f(R^{2,n'_3,n'_5}, a) = \frac{3}{5}$ and $f(R^{2,n'_3,n'_5}, b) = f(R^{2,n'_3,n'_5}, c) = \frac{1}{5}$ for all $n'_3 > 0$ and $n'_5 > 0$ with $n'_3 + n'_5 = n_3 + n_5 - 1$. Now, recall that $f(R^{1,n_3,n_5}) = f(R^{2,n_3,n_5-1})$, which means that $f(R^{1,n_3,n_5}, a) = \frac{3}{5}$ and $f(R^{1,n_3,n_5}, b) = f(R^{1,n_3,n_5}, c) = \frac{1}{5}$ if $n_5 > 1$ because of the induction hypothesis. PC-strategyproofness from R^{1,n_3,n_5} to R^{2,n_3,n_5} implies then the following inequality, where $p = f(R^{2,n_3,n_5})$.

$$\frac{2}{5}p(a) + \frac{1}{5}p(b) \geq \frac{4}{5}p(c) + \frac{3}{5}p(b) \quad \Longleftrightarrow \quad \frac{1}{2}p(a) \geq p(c) + \frac{1}{2}p(b)$$

On the other hand, if $n_5 = 1$, then $f(R^{1,n_3,n_5}, a) = f(R^{1,n_3,n_5}, b) = \frac{1}{2}$, and PC-strategyproofness results in the same inequality.

Similarly, if $n_3 > 1$, then $f(R^{3,n_3,n_5}, a) = \frac{3}{5}$ and $f(R^{3,n_3,n_5}, b) = f(R^{3,n_3,n_5}, c) = \frac{1}{5}$ because of the induction hypothesis and cancellation. Hence, we derive the following inequality from PC-strategyproofness between R^{2,n_3,n_5} and R^{3,n_3,n_5} .

$$\frac{4}{5}p(b) + \frac{3}{5}p(c) \geq \frac{1}{5}p(c) + \frac{2}{5}p(a) \quad \Longleftrightarrow \quad p(b) + \frac{1}{2}p(c) \geq \frac{1}{2}p(a)$$

On the other hand, if $n_3 = 1$, then $f(R^{3,n_3,n_5}, a) = f(R^{3,n_3,n_5}, c) = \frac{1}{2}$. Applying PC-strategyproofness in this case results in the same inequality as above.

Finally, it must hold that $p(c) \geq p(b)$. Indeed, otherwise voter n could PC-manipulate in R^{2,n_3,n_5} by reporting a as his favorite option— a would then be chosen with probability 1 because of the absolute winner property. Since $p(a) + p(b) + p(c) = 1$, it can be verified that the only possible solution to the three inequalities that we have derived is $p(a) = \frac{3}{5}$ and $p(b) = p(c) = \frac{1}{5}$. This proves the induction step and therefore that $f(R') = f(R^{2,n_3,n_5}) = f^1(R')$.

Case 2.5: Every alternative is top-ranked by less than half of the voters in R' .

As the last case, suppose that every alternative is top-ranked by less than half of the voters in R' . In particular, this means that every alternative is top-ranked at least once. We suppose again without loss of generality that a voter reports a, b, c in R' and hence, the same analysis as in the previous case shows that the only possible preference relations in R' are $\succ_1 = a, b, c$, $\succ_3 = b, c, a$, and $\succ_5 = c, a, b$. In particular, we have that $n_1 > 0$, $n_3 > 0$, $n_5 > 0$, and $n_2 = n_4 = n_6 = 0$. Moreover, since no alternative is top-ranked by at least half of the voters, we have that $n_1 < n_3 + n_5$, $n_3 < n_1 + n_5$, and $n_5 < n_1 + n_3$. This shows that there is not even a weak Condorcet winner in R' , and our goal hence is to show that $f(R', x) = \frac{1}{3}$ for all $x \in A$. Suppose that this is not the case, which means that either $f(R', a) < f(R', c)$, $f(R', b) < f(R', a)$, or $f(R', c) < f(R', b)$; otherwise, $f(R', a) \geq f(R', c) \geq f(R', b) \geq f(R', a)$, which implies that all alternatives get a probability of $\frac{1}{3}$. We assume in the sequel that $f(R', a) < f(R', c)$ as all cases are symmetric. Now, in this case, we let the voters i with preference relation a, b, c one after another swap a and b . For each step, Lemma 3 implies that the probability of a remains smaller than that of c . However, this process results in a profile R'' in which $n_1 + n_3$ voters report b as their favorite alternative. Since $n_1 + n_3 > n_5$, b is the absolute winner and Case 2.2 shows that $f(R'', b) = 1$. However, this contradicts $f(R'', a) < f(R'', c)$ and hence, the claim that $f(R', a) < f(R', c)$ must be wrong. This proves that $f(R', x) = \frac{1}{3} = f^1(R', x)$ for all $x \in A$. \square

Finally, we prove Proposition 4. Recall for this that $n_R(x)$ denotes the number of voters who top-rank alternative x in R , and let $B(R)$ be the set of alternatives that are never bottom-ranked in R . Moreover, the uniform random dictatorship RD is defined by $RD(R, x) = \frac{n_R(x)}{\sum_{y \in A} n_R(y)}$ for all $x \in A$ and $R \in \mathcal{R}^*$. As discussed in Section 3, RD is known to satisfy strict

SD-participation and therefore satisfies also strict PC-participation, but fails PC-efficiency. We consider the following variant of RD called f^2 : if $|B(R)| \in \{0, 2\}$, then $f^2(R) = RD(R)$. On the other hand, if $|B(R)| = 1$, let x denote the single alternative in $B(R)$ and let C denote the set of alternatives $y \in A \setminus \{x\}$ with minimal $n_y(R)$. Then, $f^2(R, x) = \frac{n_R(x) + \sum_{y \in C} n_R(y)}{\sum_{y \in A} n_R(y)}$, $f(R^2, y) = 0$ for $y \in C$, and $f^2(R, z) = RD(R, z)$ for $z \notin C \cup \{x\}$. Intuitively, if $|B(R)| = 1$, f^2 removes the alternatives in $A \setminus B(R)$ with minimal $n_R(x)$ and then computes RD.

Proposition 4. For $m = 3$, f^2 satisfies anonymity, neutrality, PC-efficiency, and strict PC-participation.

Proof. First note that f^2 is anonymous and neutral since its definition does not depend on the identities of voters or alternatives.

Next, we show that f^2 satisfies PC-efficiency by proving that $f^2(R)$ satisfies for all profiles R the three conditions of Lemma 6. To this end, note first that f^2 is *ex post* efficient: it only puts positive probability on an alternative that is never top-ranked if it is second-ranked by all voters and both other alternatives are top-ranked at least once. In this case, all three alternatives are Pareto-optimal, and thus f^2 is *ex post* efficient. This argument also shows that an alternative that is never top-ranked and at least once bottom-ranked is always assigned probability 0. Finally, if an alternative is never bottom-ranked and at least once top-ranked, only two alternatives can have positive probability. In more detail, either $|B(R)| = 2$, which means that one alternative is bottom-ranked by all voters and receives probability 0, or $|B(R)| = 1$ and an alternative in $A \setminus B(R)$ gets probability 0 by definition of f^2 . Hence, all conditions of Lemma 6 hold, which implies that f^2 is PC-efficient.

Lastly, we discuss why f^2 satisfies strict PC-participation—in fact, we prove the even stronger claim that it satisfies strict SD-participation. Consider an arbitrary electorate $N \in \mathcal{F}(\mathbb{N})$, a voter $i \in N$, and two preference profiles $R \in \mathcal{R}^N$ and $R' \in \mathcal{R}^{N \setminus \{i\}}$ such that $R' = R_{-i}$. We need to show that if i 's top alternative is not already chosen with probability 1 in $f^2(R')$, then $f^2(R) \succ_i^{SD} f^2(R')$. First, note that this is obvious if $f^2(R) = RD(R)$ and $f^2(R') = RD(R')$ because RD satisfies strict SD-participation. Moreover, $|B(R')| - 1 \leq |B(R)| \leq |B(R')|$ because voter i can only bottom-rank a single alternative. These two observations leave us with three interesting cases: $|B(R')| = 2$ and $|B(R)| = 1$, $|B(R')| = |B(R)| = 1$, and $|B(R')| = 1$ and $|B(R)| = 0$.

First, consider the case where $|B(R')| = 1$ and $|B(R)| = 0$. Without loss of generality, we assume that $B(R') = \{a\}$, which means that a is voter i 's least preferred alternative. Moreover, we call voter i 's best alternative $z \in \{b, c\}$. The following case distinction proves that f^2 satisfies strict SD-participation under the given assumptions.

- If $n_{R'}(b) = n_{R'}(c)$, then $f^2(R', a) = 1$ and it is obvious that $f^2(R) \succ_i^{SD} f^2(R')$ because a is voter i 's least preferred outcome and $f^2(R, z) = RD(R, z) > 0$.
- If $n_{R'}(b) > n_{R'}(c)$, we have that $f^2(R', a) = \frac{n_{R'}(a) + n_{R'}(c)}{\sum_{x \in A} n_{R'}(x)} > \frac{n_{R'}(a)}{1 + \sum_{x \in A} n_{R'}(x)} = f^2(R, a)$ and $f^2(R', z) \leq \frac{n_{R'}(z)}{\sum_{x \in A} n_{R'}(x)} < \frac{1 + n_{R'}(z)}{1 + \sum_{x \in A} n_{R'}(x)} = f^2(R, z)$. It is now easy to see that $f^2(R) \succ_i^{SD} f^2(R')$.
- The case $n_{R'}(b) < n_{R'}(c)$ is symmetric to the previous one.

Next, consider the case where $|B(R')| = 2$ and $|B(R)| = 1$. Without loss of generality, we suppose that $B(R') = \{a, b\}$ and $B(R) = \{a\}$, which means that voter i bottom-ranks b . Moreover, note that all voters in $N \setminus \{i\}$ bottom-rank c as otherwise $B(R') = \{a, b\}$ is not possible. This means that $f^2(R', c) = 0$, $n_{R'}(c) = 0$, and $n_R(c) \leq 1$. We consider again several subcases.

- If $n_R(b) > n_R(c)$, then $f^2(R, c) = 0 = f^2(R', c)$, $f^2(R, a) \geq \frac{1 + n_{R'}(a)}{1 + \sum_{x \in A} n_{R'}(x)} > \frac{n_{R'}(a)}{\sum_{x \in A} n_{R'}(x)} = f^2(R', a)$, and thus $f^2(R, b) < f^2(R', b)$. Hence, $f^2(R) \succ_i^{SD} f^2(R')$ as b is voter i 's worst alternative.
- If $n_R(c) > n_R(b)$, then $f^2(R, b) = 0 \leq f^2(R', b)$ and $f^2(R, c) > 0 = f^2(R', c)$. If i top-ranks c , we have $f^2(R) \succ_i^{SD} f^2(R')$. Else, i top-ranks a , and we have $f^2(R, a) \geq \frac{1 + n_{R'}(a)}{1 + \sum_{x \in A} n_{R'}(x)} > \frac{n_{R'}(a)}{\sum_{x \in A} n_{R'}(x)} = f^2(R', a)$, so again $f^2(R) \succ_i^{SD} f^2(R')$.
- If $n_R(c) = n_R(b) = 0$, all voters (including i) report a as their best option and thus $f^2(R', a) = f^2(R, a) = 1$, which satisfies strict SD-participation because $f^2(R)$ is voter i 's favorite lottery.
- If $n_R(c) = n_R(b) = 1$, then voter i 's preference relation is c, a, b . Moreover, $f^2(R', c) = 0 = f^2(R, c)$ and $f^2(R', b) > 0 = f^2(R, b)$. This proves again that $f^2(R) \succ_i^{SD} f^2(R')$.

As the last case, suppose that $|B(R')| = |B(R)| = 1$ and let a denote the alternative in $B(R) = B(R')$. Since $a \in B(R)$, voter i does not bottom-rank a . We consider again a case distinction.

- First, suppose that voter i top-ranks a , which means that $n_R(b) = n_{R'}(b)$ and $n_R(c) = n_{R'}(c)$.
 - If $n_{R'}(b) = n_{R'}(c)$, we have that $f^2(R, a) = f^2(R', a) = 1$ and strict PC-participation holds as this is voter i 's favorite lottery.
 - If $n_{R'}(b) > n_{R'}(c)$. Then, $f^2(R, a) = \frac{n_{R'}(a) + n_{R'}(c) + 1}{1 + \sum_{x \in A} n_{R'}(x)} > \frac{n_{R'}(a) + n_{R'}(c)}{\sum_{x \in A} n_{R'}(x)} = f^2(R', a)$, $f^2(R, c) = f^2(R', c) = 0$, and hence $f^2(R, b) < f^2(R', b)$. It is now easy to verify that $f(R) \succ_i^{SD} f(R')$.
 - The case $n_{R'}(b) < n_{R'}(c)$ is symmetric to the previous one.

- Next, suppose that voter i places a second. We assume without loss of generality that $\succ_i = b, a, c$ because the case $\succ_i = c, a, b$ is symmetric. This assumption means that $n_{R'}(b) + 1 = n_R(b)$ and $n_{R'}(x) = n_R(x)$ for $x \in \{a, c\}$.
 - If $n_{R'}(b) \geq n_{R'}(c)$, then $f^2(R', b) \leq \frac{n_{R'}(b)}{\sum_{x \in A} n_{R'}(x)} < \frac{1+n_{R'}(b)}{1+\sum_{x \in A} n_{R'}(x)} = f^2(R, b)$, $f^2(R', c) = f^2(R, c) = 0$, and hence $f^2(R', a) > f^2(R, a)$, which proves that $f(R) \succ_i^{SD} f(R')$.
 - If $n_{R'}(b) + 1 = n_{R'}(c)$, then $f^2(R', b) = 0 = f^2(R, b)$, $f^2(R', c) > 0 = f^2(R, c)$, and thus $f^2(R', a) < 1 = f^2(R, a)$. It can again be verified that $f(R) \succ_i^{SD} f(R')$.
 - If $n_{R'}(b) + 1 < n_{R'}(c)$, then $f^2(R', b) = 0 = f^2(R, b)$, $f^2(R', c) = \frac{n_{R'}(c)}{\sum_{x \in A} n_{R'}(x)} > \frac{n_{R'}(c)}{1+\sum_{x \in A} n_{R'}(x)} = f^2(R, c)$, and hence $f^2(R', a) < f^2(R, a)$. Once again, it holds that $f^2(R) \succ_i^{SD} f^2(R')$. \square

References

- Allais, M., 1953. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica* 21 (4), 503–546.
- Anand, P., 2009. Rationality and intransitive preference: foundations for the modern view. In: Anand, P., Pattanaik, P.K., Puppe, C. (Eds.), *The Handbook of Rational and Social Choice*. Oxford University Press. Chapter 6.
- Aziz, H., Brandl, F., Brandt, F., 2015. Universal Pareto dominance and welfare for plausible utility functions. *J. Math. Econ.* 60, 123–133.
- Aziz, H., Brandl, F., Brandt, F., Brill, M., 2018. On the tradeoff between efficiency and strategyproofness. *Games Econ. Behav.* 110, 1–18.
- Barberà, S., 1979. Majority and positional voting in a probabilistic framework. *Rev. Econ. Stud.* 46 (2), 379–389.
- Barberà, S., 2010. Strategy-proof social choice. In: Arrow, K.J., Sen, A., Suzumura, K. (Eds.), *Handbook of Social Choice and Welfare*, vol. 2. Elsevier, pp. 731–832. Chapter 25.
- Benoît, J.-P., 2002. Strategic manipulation in voting games when lotteries and ties are permitted. *J. Econ. Theory* 102 (2), 421–436.
- Blavatsky, P.R., 2006. Axiomatization of a preference for most probable winner. *Theory Decis.* 60 (1), 17–33.
- Blyth, C.R., 1972. Some probability paradoxes in choice from among random alternatives. *J. Am. Stat. Assoc.* 67 (338), 366–373.
- Bogomolnaia, A., Moulin, H., 2001. A new solution to the random assignment problem. *J. Econ. Theory* 100 (2), 295–328.
- Bogomolnaia, A., Moulin, H., 2002. A simple random assignment problem with a unique solution. *Econ. Theory* 19 (3), 623–635.
- Brandl, F., Brandt, F., 2020. Arrovian aggregation of convex preferences. *Econometrica* 88 (2), 799–844.
- Brandl, F., Brandt, F., Hofbauer, J., 2015. Incentives for participation and abstention in probabilistic social choice. In: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1411–1419.
- Brandl, F., Brandt, F., Seedig, H.G., 2016. Consistent probabilistic social choice. *Econometrica* 84 (5), 1839–1880.
- Brandl, F., Brandt, F., Eberl, M., Geist, C., 2018. Proving the incompatibility of efficiency and strategyproofness via SMT solving. *J. ACM* 65 (2), 1–28.
- Brandl, F., Brandt, F., Hofbauer, J., 2019. Welfare maximization entices participation. *Games Econ. Behav.* 14, 308–314.
- Brandl, F., Brandt, F., Peters, D., Stricker, C., 2021. Distribution rules under dichotomous preferences: two out of three ain't bad. In: *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM-EC)*, pp. 158–179.
- Brandl, F., Brandt, F., Stricker, C., 2022. An analytical and experimental comparison of maximal lottery schemes. *Soc. Choice Welf.* 58 (1), 5–38.
- Brandt, F., 2017. Rolling the dice: recent results in probabilistic social choice. In: Endriss, U. (Ed.), *Trends in Computational Social Choice*. AI Access, pp. 3–26. Chapter 1.
- Brandt, F., Lederer, P., 2023. Characterizing the top cycle via strategyproofness. *Theor. Econ.* 18 (2), 837–883.
- Brandt, F., Geist, C., Peters, D., 2017. Optimal bounds for the no-show paradox via SAT solving. *Math. Soc. Sci.* 90, 18–27. Special Issue in Honor of Hervé Moulin.
- Brandt, F., Bullinger, M., Lederer, P., 2022a. On the indecisiveness of Kelly-strategyproof social choice functions. *J. Artif. Intell. Res.* 73, 1093–1130.
- Brandt, F., Lederer, P., Romen, R., 2022b. Relaxed notions of Condorcet-consistency and efficiency for strategyproof social decision schemes. In: *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 181–189.
- Brandt, F., Lederer, P., Suksompong, W., 2022c. Incentives in social decision schemes with pairwise comparison preferences. Technical report. <https://arxiv.org/abs/2204.12436>.
- Ehlers, L., Peters, H., Storcken, T., 2002. Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences. *J. Econ. Theory* 105 (2), 408–434.
- Fishburn, P.C., 1982. Nontransitive measurable utility. *J. Math. Psychol.* 26 (1), 31–67.
- Fishburn, P.C., 1984a. Probabilistic social choice based on simple voting comparisons. *Rev. Econ. Stud.* 51 (4), 683–692.
- Fishburn, P.C., 1984b. Dominance in SSB utility theory. *J. Econ. Theory* 34 (1), 130–148.
- Gibbard, A., 1973. Manipulation of voting schemes: a general result. *Econometrica* 41 (4), 587–601.
- Gibbard, A., 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45 (3), 665–681.
- Hoang, L.N., 2017. Strategy-proofness of the randomized Condorcet voting system. *Soc. Choice Welf.* 48 (3), 679–701.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47 (2), 263–292.
- Machina, M.J., 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. *J. Econ. Lit.* 27 (4), 1622–1668.
- Moulin, H., 1983. *The Strategy of Social Choice*. North-Holland.
- Moulin, H., 1988. Condorcet's principle implies the no show paradox. *J. Econ. Theory* 45 (1), 53–64.
- Packard, D.J., 1982. Cyclical preference logic. *Theory Decis.* 14 (4), 415–426.
- Peyre, R., 2013. *La quête du graal électoral. Images des Mathématiques*. CNRS.
- Potthoff, R.F., 1970. The problem of the three-way election. In: Rose, R.C., Chakravarti, I.M., Mahalanobis, P.C., Rao, C.R., Smith, K.J.C. (Eds.), *Essays in Probability and Statistics*. The University of North Carolina Press.
- Satterthwaite, M.A., 1975. Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *J. Econ. Theory* 10 (2), 187–217.
- Sen, A., 2011. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIEs, J. Span. Econ. Assoc.* 2 (4), 515–527.

SUMMARY

Gibbard and Satterthwaite have shown that the only single-valued social choice functions (SCFs) that satisfy non-imposition (i.e., the function's range coincides with its codomain) and strategyproofness (i.e., voters are never better off by misrepresenting their preferences) are dictatorships. In this paper, we will demonstrate how we can circumvent this result by relaxing the convenient but rather restrictive assumption that SCFs always choose a single winner. In more detail, we will study social choice correspondences (SCCs), which return sets of possible winners rather than a single winner, based on weak \succsim^F -strategyproofness (which will be called strategyproofness in the remainder of this chapter).

As our main contribution, we characterize an attractive SCC called the top cycle based on strategyproofness and several mild auxiliary conditions. The top cycle chooses the maximal elements of the transitive closure of the weak majority relation and is well-known in the literature (e.g., Brandt et al., 2016). For our characterization of this SCC, we use three additional axioms except strategyproofness: homogeneity, which requires that splitting every voter into k clones with the same preference relation does not affect the outcome; set non-imposition, which requires that every set of alternatives can be chosen for some profile; and pairwise-ness, which requires that the winning alternatives are computed only based on the weighted majority relation. Based on these axioms, we then show that the top cycle is the *only* pairwise SCC that satisfies strategyproofness, homogeneity, and set non-imposition. This result effectively turns the Gibbard-Satterthwaite impossibility into a complete characterization of the top cycle by moving from SCFs to pairwise SCCs.

Furthermore, we leverage the main ideas of the proof of this statement to obtain a more general characterization of strategyproof SCCs. To this end, we say a set is dominant in the majority relation if every alternative in the set majority-dominates every alternative outside of the set. It can be shown that the top cycle is the SCC that always returns the smallest dominant set. Finally, we introduce a class of SCCs whose defining feature is to always return a dominant set and characterize this class of SCCs based on strategyproofness, pairwise-ness, non-imposition, and neutrality.

REFERENCE

F. Brandt and P. Lederer. **Characterizing the top cycle via strategy-proofness.** In *Theoretical Economics*, 18(2):837–883, 2023.
DOI: <http://dx.doi.org/10.3982/TE5120>

INDIVIDUAL CONTRIBUTION

I, Patrick Lederer, am the main author of this publication. In particular, I am responsible for the joint development and conceptual design of the research project, proofs of all results and the joint write-up of the manuscript (in particular, sections 3 and 4 and the appendix).

COPYRIGHT AGREEMENT

The right to present this paper in a doctoral thesis has been granted by the publisher, The Econometric Society, in the publication agreement shown below. There, it is stated that the article is published under the *Creative Commons Attribution-NonCommercial License 4.0* (see <https://creativecommons.org/licenses/by-nc/4.0/legalcode>). In particular, theoretical economics does not require a copyright transfer and the authors have thus still full control over the paper. As an author, I am therefore allowed to present the paper here. This agrees with the copyright statement of the journal available at the following website <https://econtheory.org/copyright.php>.

TERMINOLOGY

In this paper, strategyproofness refers to weak \succsim^F -strategyproofness. Moreover, voters' preferences are assumed to be strict and we use a variable electorate model (i.e., SCCs are defined for a variable number of voters).

Theoretical Economics Publication Agreement

The Econometric Society
Cowles Foundation
Yale University
30 Hillhouse Avenue
POB 208281
New Haven, CT 06511
USA

Paper title: "Characterizing the top cycle via strategyproofness"

Authors: Felix Brandt and Patrick Lederer

I agree to license my article under the [Creative Commons Attribution-NonCommercial License 4.0](#), a copy of which is available to me.

I warrant that I am a joint author of the article acting with the express consent of all joint authors. I further warrant that

1. the authors have full exclusive right, title, and interest in the article, including all associated intellectual property rights;
2. the article does not infringe upon any U.S. copyright, trade secret or other proprietary rights of any third party, and is not libelous or defamatory;
3. the article has not been published or accepted for publication elsewhere.

In addition, I acknowledge the right of the Econometric Society to correct or retract the article following publication in the event of (i) a breach of these warranties, (ii) any substantive errors, or (iii) evidence of scientific misconduct or other fraudulent actions.

Name printed:

Date (format: yyyy-m-d):

Signature:

Paper metadata are submitted to [RePEc](#). We strongly encourage you to make the published version of your paper available in other Open Access repositories (e.g. your university's institutional repository).

Characterizing the top cycle via strategyproofness

FELIX BRANDT

Department of Computer Science, Technische Universität München

PATRICK LEDERER

Department of Computer Science, Technische Universität München

Gibbard and Satterthwaite have shown that the only single-valued social choice functions (SCFs) that satisfy nonimposition (i.e., the function's range coincides with its codomain) and strategyproofness (i.e., voters are never better off by misrepresenting their preferences) are dictatorships. In this paper, we consider set-valued social choice correspondences (SCCs) that are strategyproof according to Fishburn's preference extension and, in particular, the *top cycle*, an attractive SCC that returns the maximal elements of the transitive closure of the weak majority relation. Our main theorem shows that, under mild conditions, the top cycle is the *only* non-imposing strategyproof SCC whose outcome only depends on the quantified pairwise comparisons between alternatives. This result effectively turns the Gibbard–Satterthwaite impossibility into a complete characterization of the top cycle by moving from SCFs to SCCs. We also leverage key ideas of the proof of this statement to obtain a more general characterization of strategyproof SCCs.

KEYWORDS. Top cycle, strategyproofness, Condorcet, preference extension.

JEL CLASSIFICATION. D71.

1. INTRODUCTION

One of the most influential results in microeconomic theory, the Gibbard–Satterthwaite theorem, states that dictatorships are the only single-valued social choice functions (SCFs) that are nonimposing (i.e., every alternative is returned for some preference profile) and strategyproof (i.e., voters are unable to obtain a better outcome by misrepresenting their preferences) when there are at least three alternatives. The convenient but rather restrictive assumption of single-valuedness has been criticized by various scholars. For instance, [Gärdenfors \(1976\)](#) asserts that “[resoluteness] is a rather restrictive and unnatural assumption.” In a similar vein, [Kelly \(1977\)](#) writes that “the Gibbard–Satterthwaite theorem [...] uses an assumption of singlevaluedness, which is unreasonable” and [Taylor \(2005\)](#) that “if there is a weakness to the Gibbard–Satterthwaite theorem, it is the assumption that winners are unique.” The problem with single-valuedness

Felix Brandt: brandtf@in.tum.de

Patrick Lederer: ledererp@in.tum.de

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-2 and BR 2312/12-1. Results from this paper were presented at the COMSOC video seminar (November 2021) and the 16th Meeting of the Society of Social Choice and Welfare in Mexico City (June 2022). The authors thank Florian Brandl and the anonymous referees for helpful comments.

© 2023 The Authors. Licensed under the [Creative Commons Attribution-NonCommercial License 4.0](#). Available at <https://econtheory.org>. <https://doi.org/10.3982/TE5120>

is that it is in conflict with the basic fairness notions of anonymity and neutrality, which require that all voters and all alternatives are treated equally. For example, if half of the voters favor a and the other half b , there is no fair way of selecting a single winner because both alternatives are equally acceptable. In the context of social choice, these fairness conditions are imperative because elections should be unbiased. One way to deal with this problem is to identify a set of winning candidates with the understanding that one of these candidates will eventually be selected by some tie-breaking rule independent of the voters' preferences. Ties can, for example, be broken by lottery or by letting a chairperson or a committee pick the winner.¹

As a result, a large body of research investigates so-called social choice correspondences (SCCs), which return sets of alternatives. In particular, several papers have shown statements that mimic the negative consequences of the Gibbard–Satterthwaite theorem (e.g., Duggan and Schwartz (2000), Barberà, Dutta, and Sen (2001), Ching and Zhou (2002), Benoît (2002), Sato (2014)). These results are based on relatively strong assumptions about the manipulators' preferences over sets (which in turn are based on the voters' beliefs about how ties are broken). For example, all of these results rely on the assumption that a voter who prefers a to b to c will engage in a manipulation in which the outcome changes from set $\{a, c\}$ to set $\{b\}$. However, it is quite possible that no voter entertains such preferences over sets. By contrast, the voters' preferences over sets we surmise in this paper are systematically deduced from their preferences over alternatives, which leads to a weaker notion of strategyproofness. In more detail, we consider a preference extension attributed to (Fishburn (1972)), according to which a manipulation is only successful if the manipulator can change the outcome from a set Y to another set X such that he prefers all alternatives in $X \setminus Y$ to all alternatives in Y and all alternatives in X to all alternatives in $Y \setminus X$. Two natural justifications for this extension are the existence of a chairperson who breaks ties or of a priori probabilities of the voters how ties are broken (see, e.g., Gärdenfors (1979), Ching and Zhou (2002), Erdamar and Sanver (2009), Brandt, Saile, and Stricker (2022)). The resulting notion of strategyproofness, often called Fishburn strategyproofness, allows for positive results. For example, the rather indecisive *omnination rule*, which returns all alternatives that are top-ranked by at least one voter, is strategyproof according to this notion.

A particularly promising approach to construct attractive strategyproof SCCs is to focus on the pairwise comparisons between alternatives (see, e.g., Gärdenfors (1976), MacIntyre and Pattanaik (1981), Bandyopadhyay (1983), Campbell and Kelly (2003), Brandt (2015)). For instance, Brandt (2015) shows that several attractive SCCs that only rely on the pairwise majority relation, such as the uncovered set and the bipartisan set, satisfy a strategyproofness notion which is slightly weaker than Fishburn strategyproofness. In this paper, we thus focus on the class of pairwise (aka C2) SCCs, whose outcome only depends on the weighted majority comparisons. This class was introduced by Fishburn (1977) and includes many important SCCs such as Borda's rule, Copeland's

¹These tie-breaking rules are common in real-world elections. Tied elections on the state level within the U.S. are sometimes decided by lottery. The U.S. Vice President acts as the President of the Senate and frequently breaks ties in the Senate. If no candidate in a U.S. presidential election obtains an absolute majority of the voters, then the House of Representatives elects the winner among the best three candidates.

rule, the top cycle, the essential set, the Simpson–Kramer rule, Kemeny’s rule, ranked pairs, and Schulze’s rule (see Chapters 3 and 4 in Brandt, Conitzer, Endriss, Lang, and Procaccia (2016), for an overview of these SCCs). Indeed, it is well known that almost all other SCCs (e.g., positional scoring rules or runoff rules) are manipulable according to Fishburn’s preference extension (see, e.g., Taylor (2005, pp. 44–51)).

A prominent concept that arises from pairwise comparisons between alternatives is that of a Condorcet winner, an alternative that is preferred to every other alternative by a majority of voters (Condorcet (1785)). Condorcet winners need not exist, but many scholars agree that an SCC should uniquely return the Condorcet winner whenever one exists. The nonexistence of Condorcet winners can be addressed by extending the notion of Condorcet winners to so-called dominant sets of alternatives. A set of alternatives X is dominant if every element of X is preferred to every element not in X by a majority of voters. Dominant sets are guaranteed to exist since the set of all alternatives is trivially dominant, and they can be ordered by set inclusion.² These observations have led to the definition of the *top cycle*, an SCC that returns the unique smallest dominant set for any given preference profile. This set consists precisely of the maximal elements of the transitive closure of the weak majority relation. The top cycle has been reinvented several times and is known under various names such as Good set (Good (1971)), Smith set (Smith (1973)), weak closure maximality (Sen (1977)), and GETCHA (Schwartz (1986)).

In this paper, we characterize the class of strategyproof pairwise SCCs under relatively mild and common technical assumptions, namely the conditions of nonimposition, homogeneity, and neutrality. Our first result shows that every strategyproof pairwise SCC that satisfies these conditions always returns a dominant set. An important variant of this characterization is obtained when replacing nonimposition and neutrality with set not-imposition (every set of alternatives is returned for some preference profile): the top cycle is the *only* strategyproof pairwise SCC that satisfies set nonimposition and homogeneity. This result effectively turns the Gibbard–Satterthwaite impossibility theorem into a complete characterization of the top cycle by moving from SCFs to SCCs.

On top of strategyproofness, the top cycle is very robust in terms of changes to the set of feasible alternatives and preferences of the voters: it is invariant under removing losing alternatives as well as modifications of preferences between losing alternatives, and has been characterized repeatedly by choice consistency conditions implied by the weak axiom of revealed preference. Finally, it is one of the most straightforward Condorcet extensions and can be easily computed. The main disadvantage of the top cycle is its possible inclusion of Pareto-dominated alternatives. We believe that this drawback is tolerable because empirical results suggest that the top cycle only rarely contains Pareto-dominated alternatives. This is due to the fact that Pareto dominances are increasingly unlikely for large numbers of voters and the persistent observation that an overwhelming number of real-world elections admit Condorcet winners (see, e.g., Regenwetter, Grofman, Marley, and Tsetlin (2006), Laslier (2010), Gehrlein and Lepelley (2011), Brandt and Seedig (2016)). In these cases, the top cycle consists of a single

²Assume for contradiction that two dominant sets, X , Y , are not contained in each other. Then there exists $x \in X \setminus Y$ and $y \in Y \setminus X$. The definition of dominant sets requires that x is majority-preferred to y and that y is majority-preferred to x , a contradiction.

Pareto-optimal alternative. Moreover, as we point out in Remark 8, the SCC that returns the top cycle of the subset of Pareto-optimal alternatives satisfies all our axioms except pairwiseness. In particular, this SCC satisfies strategyproofness with respect to Fishburn's extension.

2. RELATED WORK

Gärdenfors (1979) initiated the study of strategyproofness with respect to Fishburn's preference extension. He attributed this extension to Fishburn because it is the weakest extension that satisfies a set of axioms proposed by Fishburn (1972). A small number of SCCs were shown to be Fishburn strategyproof, sometimes by means of stronger strategyproofness notions: the Pareto rule—which returns all Pareto-optimal alternatives (Feldman (1979)), the omninomination rule—which returns all top-ranked alternatives (Gärdenfors (1976)), the Condorcet rule—which returns the Condorcet winner whenever one exists and all alternatives otherwise (Gärdenfors (1976)), the SCC that returns the Condorcet winner whenever one exists and all Pareto-optimal alternatives otherwise (Brandt and Brill (2011)), and the top cycle (Bandyopadhyay (1983), Brandt and Brill (2011), Sanver and Zwicker (2012)). All other commonly studied SCCs fail Fishburn strategyproofness (see, e.g., Taylor (2005), Brandt, Brill, and Harrenstein (2016)). A universal example showing the Fishburn manipulability of many SCCs is given in Figure 2 of Section 3.

More recently, the limitations of Fishburn strategyproofness were explored. Brandt and Geist (2016) studied majoritarian SCCs, that is, SCCs whose outcome only depends on the pairwise majority relation, and showed that no majoritarian SCC satisfies Fishburn strategyproofness and Pareto optimality. The condition of majoritarianism can be replaced with the much weaker condition of anonymity when allowing for ties in the preferences (Brandt, Saile, and Stricker (2022)). Both results were obtained with the help of computer-aided theorem proving techniques. Brandt and Geist (2016, Remark 3) observed that, in the absence of majority ties, the top cycle could be the finest majoritarian Condorcet extension that satisfies Fishburn strategyproofness when there are at least five alternatives. A computer verified this claim for five, six, and seven alternatives using 24 hours of runtime. The claim now follows immediately from our Theorem 1 (see also Remark 3), irrespective of majority ties.

Ching and Zhou (2002) considered a much stronger notion of strategyproofness based on Fishburn's preference extension: they require that the outcome when voting honestly is comparable and preferred to every choice set obtainable by a manipulation according to Fishburn's extension. Their main result shows that only constant and dictatorial SCCs are strategyproof according to this definition. Barberà, Dutta, and Sen (2001) derive a similar conclusion for a weaker notion of strategyproofness based on Fishburn's extension (but still stronger than the one considered in this paper). In their model, voters submit preference relations over sets of alternatives that adhere to certain structural restrictions. When these restrictions are given by Fishburn's extension, they prove that only dictatorships satisfy strategyproofness and unanimity.

Several choice-theoretic characterizations of the top cycle exist. When assuming that choices from two-element sets are made according to majority rule, the influential characterization by Bordes (1976) entails that the top cycle is the finest SCC satisfying β^+ , an expansion consistency condition implied by the weak axiom of revealed preference.³ Ehlers and Sprumont (2008) have shown that, in the absence of majority ties, the refinement condition can be replaced with two contraction consistency conditions. Brandt (2011), Houy (2011), and Brandt, Brill, Seedig, and Suksompong (2018) provide further characterization using choice consistency conditions. We are not aware of a characterization of the top cycle using strategyproofness.

3. PRELIMINARIES

Let $\mathbb{N} = \{1, 2, \dots\}$ denote an infinite set of voters and A a finite set of m alternatives. Moreover, let $\mathcal{F}(\mathbb{N})$ denote the set of all finite and nonempty subsets of \mathbb{N} . Intuitively, \mathbb{N} is the set of all possible voters, whereas an element $N \in \mathcal{F}(\mathbb{N})$ represents a concrete electorate. Given an electorate $N \in \mathcal{F}(\mathbb{N})$, each voter $i \in N$ has a *preference relation* represented by a strict total order \succ_i on A . The set of all preference relations on A is denoted by $\mathcal{R}(A)$. A *preference profile* R is a vector of preference relations, that is, $R \in \mathcal{R}(A)^N$ for some electorate $N \in \mathcal{F}(\mathbb{N})$. The set of all preference profiles on A is denoted by $\mathcal{R}^*(A) = \bigcup_{N \in \mathcal{F}(\mathbb{N})} \mathcal{R}(A)^N$.

For a preference profile $R \in \mathcal{R}^*(A)$, let

$$g_R(x, y) = |\{i \in N : x \succ_i y\}| - |\{i \in N : y \succ_i x\}| \quad (\text{Majority margin})$$

be the *majority margin* of x over y in R . It describes how many more voters prefer x to y than y to x . Whenever $g_R(x, y) \geq 0$ for some pair of alternatives, we say that x weakly (majority) dominates y , denoted by $x \succsim_R y$. Note that the relation \succsim_R , which we call *majority relation*, is complete. Its strict part will be denoted by \succ_R , that is, $x \succ_R y$ if and only if $x \succsim_R y$ and not $y \succsim_R x$, and its indifference part by \sim_R , that is, $x \sim_R y$ if and only if $x \succsim_R y$ and $y \succsim_R x$. Whenever the number of voters is odd, there can be no majority ties and \succsim_R is antisymmetric. We will extend both individual preference relations and the majority relation to sets of alternatives using the shorthand notation $X \succ Y$ whenever $x \succ y$ for all $x \in X$ and $y \in Y$.

The majority relation gives rise to a number of important concepts in social choice theory. A *Condorcet winner* is an alternative x such that $x \succ_R A \setminus \{x\}$. In a similar vein, a *Condorcet loser* is an alternative x with $A \setminus \{x\} \succ_R x$. Neither Condorcet winners nor Condorcet losers need to exist, but whenever they do, each of them is unique. A natural extension of these ideas to sets of alternatives is formalized via the notion of dominant sets. A nonempty set $X \subseteq A$ is *dominant* if $X \succ_R A \setminus X$. Whenever a Condorcet winner exists, it forms a singleton dominant set. In contrast to Condorcet winners, dominant sets are guaranteed to exist since the set of all alternatives A is trivially dominant. For every majority relation \succsim_R , the set of dominant sets is totally ordered by set inclusion, that

³This result was recently rediscovered by Evren, Nishimura, and Ok (2019).




FIGURE 1. A preference profile with $N = \{1, 2, 3, 4\}$ and $A = \{a, b, c, d, e\}$ (left-hand side) and the corresponding weighted majority graph (right-hand side). An edge from x to y with weight w denotes that $g_R(x, y) = w$. Edges with weight 0 are bidirectional since, in this case, both alternatives weakly majority dominate each other. The smallest dominant set, $\{a, b, c\}$, is highlighted in gray.

is, each majority relation induces a hierarchy of dominant sets that are strictly contained in each other.²

For an illustration of these concepts, consider the example given in Figure 1, which shows a preference profile R and its majority relation. The weights on the edges of the majority relation indicate the majority margins. The profile R neither admits a Condorcet winner nor a Condorcet loser since each alternative has at least one incoming and outgoing edge. There are two dominant sets, $\{a, b, c\}$ and $\{a, b, c, d, e\}$. Note that the notions of Condorcet winners and dominant sets are independent of the exact weights of the edges but only depend on their directions.

Social choice correspondences

This paper is concerned with *social choice correspondences* (SCCs). An SCC maps a preference profile to a nonempty subset of alternatives called the choice set, that is, it is a function of the form $f : \mathcal{R}^*(A) \rightarrow 2^A \setminus \{\emptyset\}$. Note that we employ a so-called variable population framework, so SCCs are defined for all electorates. In this paper, we focus on two important classes of SCCs: *majoritarian* SCCs and *pairwise* SCCs. An SCC f is called *majoritarian* if its outcome merely depends on the majority relation, that is, $f(R) = f(R')$ for all $R, R' \in \mathcal{R}^*(A)$ with $\succ_R = \succ'_{R'}$. Furthermore, an SCC f is *pairwise* if its outcome merely depends on the majority margins, that is, $f(R) = f(R')$ for all $R, R' \in \mathcal{R}^*(A)$ with $g_R = g_{R'}$. Majoritarian SCCs can be interpreted as functions that map an unweighted graph (A, \succ_R) to a nonempty subset of its vertices, while pairwise SCCs may additionally use the majority margins $g_R(x, y)$ as weights of the edges. The classes of majoritarian and pairwise SCCs are very rich and contain a variety of well-studied SCCs. For instance, Copeland's rule, the uncovered set, and the bipartisan set are majoritarian, and Borda's rule, the Simpson–Kramer rule, the essential set, Kemeny's rule, ranked pairs, and Schulze's rule are pairwise (the interested reader may consult Chapters 3 and 4 in Brandt et al. (2016) for definitions of these SCCs). All SCCs listed above, except Borda's rule, are *Condorcet extensions*, that is, they uniquely return the Condorcet winner whenever one exists.

We say that an SCC f is a *refinement* of an SCC g if $f(R) \subseteq g(R)$ for all preference profiles $R \in \mathcal{R}^*(A)$. In this case, g is also said to be *coarser* than f . For example, Copeland's

rule, the uncovered set, the essential set, Kemeny's rule, ranked pairs, and Schulze's rule are known to be refinements of the top cycle while the Condorcet rule is a coarsening of the top cycle.

The top cycle is a majoritarian SCC that returns the smallest dominant set for a given preference profile. Every preference profile admits a unique smallest dominant set because dominant sets are ordered by set inclusion. Alternatively, the top cycle can be defined based on paths with respect to the majority relation. A path from an alternative x to an alternative y in \succsim_R is a sequence of alternatives (x_1, \dots, x_k) such that $x_1 = x$, $x_k = y$, and $x_i \succsim_R x_{i+1}$ for all $i \in \{1, \dots, k-1\}$. The transitive closure \succsim_R^* of the majority relation contains all pairs of alternatives (x, y) such that there is a path from x to y in \succsim_R . Then the top cycle can be defined as the set of alternatives that are maximal according to \succsim_R^* :

$$TC(R) = \bigcap \{X \subseteq A : X \succ_R A \setminus X\} = \{x \in A : x \succsim_R^* A\}. \quad (\text{Top cycle})$$

In other words, the top cycle consists precisely of those alternatives that reach every other alternative on some path in the majority graph. For instance, the top cycle of the example profile in Figure 1 is $\{a, b, c\}$. Here, it is important that we interpret majority ties as bidirectional edges as otherwise there would be no path from a to b .⁴

On top of majoritarianism and pairwiseism, which restrict the informational basis of SCCs, we now introduce a number of additional properties of SCCs.

- An SCC is *nonimposing* if for every alternative $x \in A$ there is a profile $R \in \mathcal{R}^*(A)$ such that $f(R) = \{x\}$.
- An SCC is *neutral* if $f(R') = \pi(f(R))$ for all electorates $N \in \mathcal{F}(\mathbb{N})$, profiles $R, R' \in \mathcal{R}(A)^N$, and permutations $\pi : A \rightarrow A$ such that $x \succ_i y$ if and only if $\pi(x) \succ'_i \pi(y)$ for all alternatives $x, y \in A$ and voters $i \in N$.
- An SCC is *homogeneous* if for all preference profiles $R \in \mathcal{R}^*(A)$, $f(R) = f(kR)$ where the profile kR consists of k copies of R .

Nonimposition is a mild decisiveness requirement demanding that every alternative will be selected uniquely for some configuration of preferences. It is weaker than Pareto optimality and unanimity (an alternative that is top-ranked by all voters has to be elected uniquely). Neutrality requires that if alternatives are relabeled in a preference profile, the alternatives in the corresponding choice set are relabeled accordingly. Homogeneity states that cloning the entire electorate will not affect the choice set. All three of these properties are very mild and satisfied by all SCCs typically considered in the literature, including the top cycle.

⁴There is a refinement of the top cycle, sometimes called the Schwartz set or GOCHA, which is defined as the union of undominated sets, or alternatively, as the set of alternatives that reach every other alternative on a path according to \succ_R (rather than \succsim_R) (see, e.g., Schwartz (1972), Deb (1977), Schwartz (1986)). We will not consider it further because it violates rather mild consistency and strategyproofness conditions.

Fishburn's extension and strategyproofness

An important desirable property of SCCs is strategyproofness, which demands that voters should never be better off by lying about their preferences. To make this formally precise for social choice *correspondences*, we need to make assumptions about the voters' preferences over sets. In this paper, we extend the voters' preferences over alternatives to incomplete preference over sets by using Fishburn's preference extension. Given two sets of alternatives $X, Y \subseteq A$, $X \neq Y$, and a preference relation \succ_i , Fishburn's extension is defined by

$$X \succ_i^F Y \quad \text{if and only if} \quad X \setminus Y \succ_i Y \quad \text{and} \quad X \succ_i Y \setminus X. \quad (\text{Fishburn's extension})$$

Fishburn's extension is frequently considered in social choice theory and can be justified in various ways (see, e.g., Gärdenfors (1979), Ching and Zhou (2002), Erdamar and Sanver (2009), Brandt, Saile, and Stricker (2022)). For example, one motivation assumes that a single alternative will eventually be selected from each choice set according to a linear tie-breaking ordering (such as the preference relation of a chairperson) and that voters are unaware of the concrete ordering used to break ties. Then set X is preferred to set Y if and only if for all tie-breaking orderings, the voter weakly prefers the alternative selected from X to that selected from Y and there is at least one ordering for which this comparison is strict. Another motivation is based on a function that assigns an a priori weight to each alternative such that each choice set can be mapped to a lottery over the chosen alternatives such that the probabilities are proportional to the alternatives' weights. Again the voters are unaware of the concrete weight function and prefer set X to set Y if and only if for all utility functions consistent with their ordinal preferences and all a priori weight functions, the expected utility derived from X is higher than that derived from Y .

Strategyproofness based on Fishburn's extension can be defined as follows. An SCC f is (*Fishburn*) *strategyproof* if for all electorates $N \in \mathcal{F}(\mathbb{N})$ and profiles $R \in \mathcal{R}(A)^N$, there is no profile $R' \in \mathcal{R}(A)^N$ such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$ and $f(R') \succ_i^F f(R)$.

Even though Fishburn strategyproofness seems like a relatively weak strategyproofness notion, it is only satisfied by very few SCCs. In particular, the omninomination rule, the Pareto rule, the top cycle, and the Condorcet rule are Fishburn strategyproof, while virtually all other commonly studied SCCs fail Fishburn strategyproofness. As an example of this claim, we give an example demonstrating the Fishburn manipulability of plurality rule (which chooses the alternatives top-ranked by most voters) and many other SCCs in Figure 2.

4. RESULTS

Our first result is a complete characterization of pairwise strategyproof SCCs in terms of dominant set rules. An SCC is a *dominant set rule* if for every preference profile R , it returns a dominant set with respect to \succsim_R . Examples of dominant set rules are the top cycle, the Condorcet rule (which returns the Condorcet winner if it exists and all alternatives otherwise), and the Condorcet nonloser rule (which returns all alternatives

$$\begin{array}{ll}
i \in \{1, 2\}: & a \succ_i b \succ_i c \\
i \in \{3, 4\}: & c \succ_i a \succ_i b \\
5: & b \succ_5 c \succ_5 a
\end{array}
\qquad
\begin{array}{ll}
i \in \{1, 2\}: & a \succ_i b \succ_i c \\
i \in \{3, 4\}: & c \succ_i a \succ_i b \\
5: & c \succ_5 b \succ_5 a
\end{array}$$

FIGURE 2. Example showing the Fishburn manipulability of plurality rule. Plurality rule chooses $\{a, c\}$ for the left profile and $\{c\}$ for the right profile, so voter 5 can manipulate by deviating from the profile on the left to the one on the right. The same example shows that many other popular SCCs are Fishburn manipulable, for example, Borda's rule, Nanson's rule, Black's rule, the maximin rule, Bucklin's rule, Young's rule, and Kemeny's rule (we refer to Brandt et al. (2016) for definitions of these SCCs).

but a Condorcet loser). Observe that—even though dominant set rules may seem rather restricted—they allow for a fair degree of freedom in the choice of dominant sets. For instance, dominant set rules can take the majority margins into account. This is demonstrated by the SCC that returns an alternative x as unique winner if $g_R(x, y) > 2$ for all $y \in A \setminus \{x\}$ and otherwise returns all alternatives. It is also possible to define rather unnatural dominant set rules such as the SCC that returns the smallest dominant set whenever the majority graph contains a cycle with identical weights and the second smallest dominant set otherwise.

For our analysis, it suffices to consider the particularly simple subclass of robust dominant set rules: a dominant set rule f is *robust* if $f(R') \subseteq f(R)$ for all preference profiles R, R' such that $f(R)$ is dominant in R' . In other words, if the choice set $f(R)$ for some profile R is also dominant in another profile R' , then no alternative outside of $f(R)$ can be chosen for R' . It is easily seen that the top cycle, the Condorcet rule, and the Condorcet nonloser rule are robust dominant set rules, whereas the two artificial examples given above fail this condition. Moreover, robust dominant set rules are majoritarian and, therefore, also homogeneous.

Our first theorem shows that—under mild additional assumptions—robust dominant set rules are the *only* pairwise SCCs that satisfy strategyproofness.

THEOREM 1. *Let f be a pairwise SCC that satisfies nonimposition, homogeneity, and neutrality. Then f is strategyproof if and only if it is a robust dominant set rule.*

The direction from right to left is relatively straightforward: every robust dominant set rule f is strategyproof because robustness prohibits successful manipulations of dominant set rules. Consider, for example, that voter i manipulates from a profile R to another profile R' such that $f(R') \subsetneq f(R)$. According to Fishburn's extension, we have that $f(R') \succ_i f(R) \setminus f(R')$. Moreover, $f(R') \succ_{R'} A \setminus f(R')$ because $f(R')$ is dominant in $\succ_{R'}$. Since voter i can only weaken alternatives in $f(R')$ against those in $f(R) \setminus f(R')$ when moving from R to R' , $f(R')$ will only be strengthened against $f(R) \setminus f(R')$ when moving from R' to R . Hence, $f(R') \succ_R f(R) \setminus f(R')$. This implies that $f(R')$ is also dominant in \succ_R because $f(R') \subsetneq f(R)$ and $f(R)$ is dominant in R . Since $f(R')$ is dominant in both R and R' , robustness implies $f(R) \subseteq f(R')$, which is at variance with our initial assumption that $f(R') \subsetneq f(R)$. A similar argument applies to the case that $f(R') \not\subseteq f(R)$ where strategyproofness now implies that $f(R') \setminus f(R) \succ_i f(R)$.

The converse direction—every pairwise SCC that satisfies nonimposition, homogeneity, neutrality, and strategyproofness is a robust dominant set rule—is much more difficult to prove. As a first step, we investigate the consequences of strategyproofness for pairwise SCCs. It turns out that we can abstract away from the concrete preferences of the voters and derive multiple axioms that describe how the choice set is affected when modifying the pairwise comparisons between pairs of alternatives. For instance, we show that rearranging unchosen alternatives in the voters' preferences does not affect the choice set of strategyproof and pairwise SCCs. The proofs of these implications heavily use the fact that two voters with preferences inverse to each other can be added to a preference profile without affecting the outcome of pairwise SCCs. As a second step, we use these axioms to derive some insights on the structure of choice sets returned by pairwise SCCs that satisfy strategyproofness, nonimposition, homogeneity, and neutrality. In more detail, we show for such an SCC f that (i) it chooses a single winner if and only if it is the Condorcet winner and (ii) for every alternative $x \in A$, either $f(R) = \{x\}$ or there is an alternative $y \in f(R) \setminus \{x\}$ such that $y \succ_R x$. The first condition is called strong Condorcet consistency, and we show that every pairwise, homogeneous, strategyproof, and strongly Condorcet-consistent coarsening of the top cycle is a robust dominant set rule. As the last step, we show that every pairwise SCC that satisfies the given axioms is a coarsening of the top cycle, which completes the proof of the theorem.

Theorem 1 has a number of important and perhaps surprising consequences. For instance, it implies that every strategyproof and pairwise SCC that satisfies the given conditions has to be majoritarian. In other words, every such SCC completely ignores the absolute values of majority margins, even though these values allow for the definition of more sophisticated SCCs.⁵ Moreover, Theorem 1 entails that many strategyproofness notions are equivalent under the assumptions of the theorem because robust dominant set rules satisfy much stronger notions of strategyproofness than Fishburn strategyproofness (see Remark 6).

Another interesting consequence of Theorem 1 is that the top cycle is the finest pairwise SCC that satisfies strategyproofness, nonimposition, homogeneity, and neutrality since it returns the smallest dominant set for any given preference profile. As shown in the sequel, we can turn this observation into a characterization of the top cycle by replacing nonimposition and neutrality with set nonimposition. An SCC f satisfies *set nonimposition* if for every nonempty set $X \subseteq A$, there is a profile R such that $f(R) = X$. In other words, every set is chosen for some preference profile R , which is in line with the original motivation of nonimposition for SCFs: the functions's image coincides with its codomain. For neutral and pairwise SCCs, set nonimposition can be interpreted as a weak efficiency notion. To see this, assume that there is some set $X \subseteq A$ that is never returned by f . Now, consider a profile R such that $X \succ_i A \setminus X$ for all $i \in N$ and $x \sim_R y$ for all $x, y \in X$ and $x, y \in A \setminus X$. Neutrality and pairwise SCCs imply that f can only return X , A , or $A \setminus X$. Since f never returns X by assumption, we have that $A \setminus X \subseteq f(R)$. However, every voter $i \in N$ prefers X to $A \setminus X$ and the choice set of f is thus very inefficient.

⁵This insight resembles the fact that individual preference intensities can usually not be used by strategyproof voting rules (see, e.g., Nandebam (2013), Ehlers, Majumdar, Mishra, and Sen (2020)). Note, however, that majority margins represent collective preference intensities.

The following lemma shows how set nonimposition can be used to single out the top cycle among all robust dominant set rules.

LEMMA 1. *The top cycle is the only robust dominant set rule that satisfies set nonimposition.*

PROOF. It has already been stated that TC is a robust dominant set rule. Moreover, TC satisfies set nonimposition because every set X is the smallest dominant set for every profile R such that $x \sim_R y$ for all $x, y \in X$, and $X \succ_R A \setminus X$; the existence of such a profile follows from McGarvey's construction (McGarvey (1953)).

For the other direction, consider a robust dominant set rule $f \neq TC$. Since f is not the top cycle, there is a profile R with dominant sets D_1, \dots, D_k such that $D_i \subseteq D_j$ if and only if $i \leq j$ and $f(R) = D_i$ with $i \geq 2$. This means that there is no profile R' such that $f(R') = D_1$ since otherwise, robustness from R' to R implies that $f(R) \subseteq f(R') = D_1$ because $f(R')$ is dominant in \succsim_R . In other words, f violates set nonimposition and the top cycle is consequently the only robust dominant set rule that satisfies this axiom. \square

The combination of Theorem 1 and Lemma 1 already characterizes the top cycle as the only pairwise SCC that satisfies strategyproofness, set nonimposition, homogeneity, and neutrality. It turns out that neutrality is not required for this characterization as the insights of the proof of Theorem 1 can be leveraged to establish that only robust dominant set rules satisfy pairwiseness, strategyproofness, set nonimposition, and homogeneity. In particular, the axioms of Theorem 2 also imply strong Condorcet consistency, which together with our previous insights and Lemma 1 yields the following characterization.

THEOREM 2. *The top cycle is the only pairwise SCC that satisfies strategyproofness, set nonimposition, and homogeneity.*

We conclude the paper with a number of remarks.

REMARK 1 (Independence of the axioms). We can show that all of the axioms, except nonimposition, are required for the direction from left to right of Theorem 1. If we only omit pairwiseness, the omninomination rule satisfies all required axioms but is no dominant set rule. If we dismiss neutrality, the following SCC based on two special alternatives a and b satisfies all requirements, but is no dominant set rule: f^{ab} returns $\{a\}$ if $a \succ_R A \setminus \{a, b\}$ and $a \succsim_R b$; otherwise it returns the outcome of the Condorcet rule. All axioms except homogeneity are satisfied by the SCC TC^* , which returns the top cycle with respect to the relation $x \succsim_R^* y$ if and only if $g_R(x, y) \geq -1$. However, TC^* is no robust dominant set rule because it depends on the majority margins. It is open whether nonimposition is required for Theorem 1. We discuss a variant of Theorem 1, which uses strong Condorcet consistency instead of neutrality and nonimposition in the [Appendix](#). For this variant, it is easy to prove that all axioms are indeed required.

For the converse direction of Theorem 1, none of the auxiliary axioms is required. In particular, every robust dominant set rule is homogeneous and pairwise because robustness entails majoritarianism. Moreover, these SCCs satisfy strategyproofness regardless of whether they are neutral or nonimposing. For instance, the SCC that chooses the set $\{a, b, c\}$ if it is a dominant set and all alternatives otherwise is neither neutral nor nonimposing, but it is a robust dominant set rule and strategyproof.

For Theorem 2, we can show the independence of all axioms. Borda's rule only violates strategyproofness, the Condorcet rule only violates set nonimposition, the omninomination rule only violates pairwise, and TC^* only violates homogeneity.

REMARK 2 (Tournaments). A significant part of the literature focuses on the special case when there are no majority ties and the majority graph is a tournament (see, e.g., Laslier (1997), Brandt, Brill, and Harrenstein (2016)). This, for example, happens when the number of voters is odd. In the absence of majority ties and when $m \leq 4$, there is a strategyproof SCC known as the uncovered set, which satisfies all requirements of Theorem 1 but is no dominant set rule. When $m \geq 5$, the uncovered set violates strategyproofness and Theorem 1 holds even in the absence of majority ties.

REMARK 3 (Dropping homogeneity). The example given in Remark 1 for the independence of homogeneity only shows that robustness might be violated if we dismiss homogeneity, but the considered SCC is still a dominant set rule. It turns out that this observation is true in general if we mildly strengthen nonimposition to *unanimity* (a unanimously top-ranked alternative will be selected uniquely): every pairwise SCC that satisfies strategyproofness, unanimity, and neutrality is a dominant set rule if $m \neq 4$. The last condition is required because of the uncovered set discussed in Remark 2. By weakening robustness, one can thus obtain an alternative characterization of strategyproof SCCs based on *weak robustness*: an SCC f is weakly robust if $f(R') \subseteq f(R)$ for all preference profiles R, R' such that $g_R(x, y) \leq g_{R'}(x, y)$ for all $x \in f(R), y \in A \setminus f(R)$. Then, if $m \neq 4$, every pairwise SCC that satisfies unanimity and neutrality is strategyproof if and only if it is a weakly robust dominant set rule.

REMARK 4 (Weakening neutrality). Another variant of Theorem 1 can be obtained by weakening neutrality to the following condition: $x \in f(R)$ if and only if $y \in f(R)$ for every preference profile R and all pairs of alternatives $x, y \in A$ such that $g_R(x, y) = 0$ and $g_R(x, z) = g_R(y, z)$ for all $z \in A \setminus \{x, y\}$. This condition is void if there is an odd number of voters. As a result, homogeneity becomes more important for the proof as some steps only work for an even number of voters.

REMARK 5 (Weakening nonimposition). In the presence of neutrality, nonimposition can be weakened to a condition that merely requires that the SCC returns a singleton set for at least one profile. If we weaken neutrality, this is no longer possible and our proof suggests that, among the three auxiliary axioms, nonimposition plays the most important role as it is crucial for deriving strong Condorcet consistency.

REMARK 6 (Strengthening strategyproofness). Fishburn strategyproofness is a rather weak strategyproofness notion, which makes the direction from left to right in our characterizations strong. However, robust dominant set rules—especially the top cycle—are actually much more resistant against manipulation. To formalize this, we introduce a new preference extension, denoted by \succsim_i^{F+} , based on the relation \succ_i^\exists over the subsets of A . This relation is defined as $X \succsim_i^\exists Y$ if and only if $X = \emptyset$, $Y = \emptyset$, or there are alternatives $x \in X$, $y \in Y$ such that $x \succ_i y$. Then

$$X \succsim_i^{F+} Y \quad \text{if and only if} \\ X \setminus Y \succ_i Y \setminus X \quad \text{and} \quad X \setminus Y \succ_i^\exists X \cap Y \quad \text{and} \quad X \cap Y \succ_i^\exists Y \setminus X.$$

Clearly, $X \succsim_i^F Y$ implies $X \succsim_i^{F+} Y$, and consequently, \succsim_i^{F+} -strategyproofness is stronger than Fishburn strategyproofness. We define an even stronger notion of strategyproofness based on the \succsim_i^{F+} extension as follows: an SCC f is *strongly* \succsim_i^{F+} -strategyproof if $f(R) \succsim_i^{F+} f(R')$ for all voters $i \in N$ and preference profiles R, R' with $\succ_j = \succ_j'$ for all $j \in N \setminus \{i\}$. Strong \succsim_i^{F+} -strategyproofness requires that all choice sets for manipulated preference profiles are comparable to the original choice set, making it much stronger than both \succsim_i^{F+} -strategyproofness and Fishburn strategyproofness. Strong Fishburn strategyproofness can be defined analogously. The top cycle is strongly \succsim_i^{F+} -strategyproof. Interestingly, Ching and Zhou (2002) have shown that only dictatorial and constant SCCs satisfy the slightly stronger notion of strong Fishburn strategyproofness, which obviously rules out the top cycle.

REMARK 7 (Group strategyproofness). An SCC f is group strategyproof if for all preference profiles R, R' and sets of voters $G \subseteq N$ such that $\succ_j = \succ_j'$ for $j \in N \setminus G$, it holds that $f(R') \not\succsim_i^F f(R)$ for some voter $i \in G$. Since every robust dominant set rule is group strategyproof, it follows from Theorem 1 that strategyproofness is equivalent to group strategyproofness for pairwise SCCs that satisfy homogeneity, nonimposition, and neutrality.

REMARK 8 (Pareto optimality). The main disadvantage of the top cycle is that it may return Pareto-dominated alternatives. In fact, *every* strategyproof pairwise SCC that satisfies our assumptions violates Pareto optimality. However, it is possible to circumvent this impossibility by first removing all Pareto-dominated alternatives and then computing the top cycle of the remaining alternatives. This SCC, $TC(PO)$ where PO stands for the Pareto rule, was already considered by Bordes (1979) and can be shown to be strategyproof. In fact, it satisfies all conditions of Theorem 1 except pairwiseness since it is not possible to compute the set of Pareto-dominated alternatives based on the majority margins only. Interestingly, the “converse” SCC, $PO(TC)$, which first computes the top cycle and then removes all Pareto-dominated alternatives, is nested in between $TC(PO)$ and TC but violates strategyproofness.

REMARK 9 (Fishburn efficiency). As discussed in the previous remark, the top cycle fails Pareto optimality. However, the top cycle satisfies the weaker notion of *Fishburn efficiency*, which requires that for every profile R , there is no set of alternatives X such that

$X \succ_i^F f(R)$ for all $i \in N$. Fishburn efficiency can be seen as a weak form of ex ante efficiency, where outcomes are compared before ties are broken. It is easy to see that the top cycle is the only robust dominant set rule satisfying this axiom since every other such rule already violates set nonimposition. It can moreover be shown that the top cycle is the coarsest majoritarian SCC that satisfies Fishburn efficiency, that is, every majoritarian SCC f that is Fishburn efficient satisfies that $f(R) \subseteq TC(R)$ for all preference profiles R . Since the top cycle is also the finest majoritarian SCC that satisfies Fishburn strategyproofness, neutrality, and nonimposition, it can be completely characterized using strategyproofness and efficiency.

REMARK 10 (Beyond the majority relation). Dominant set rules can be defined with respect to any complete binary relation derived from the preference profile. To formalize this idea, let the *information base* $I(R)$ denote a function that maps $R \in \mathcal{R}^*(A)$ to a complete binary relation $\succsim_{I(R)}$ on A . Applying a dominant set rule to $\succsim_{I(R)}$ clearly results in an SCC. Moreover, if $I(R)$ is *local* (i.e., $a \succsim_{I(R)} b$ if and only if $a \succsim_{I(R')} b$ for all $a, b \in A$ and $R, R' \in \mathcal{R}(A)^N$ such that $a \succ_i b$ if and only if $a \succ'_i b$ for all $i \in N$) and *monotone* (i.e., $a \succsim_{I(R)} b$ implies $a \succsim_{I(R')} b$ for all $a, b \in A$, and $R, R' \in \mathcal{R}^*(A)$ such that R' is derived from R by reinforcing a against b in the preference relation of a voter i), then every robust dominant set rule on $\succsim_{I(R)}$ is strategyproof. This proves, for instance, that dominant set rules based on *supermajority relations* (i.e., $a \succsim_{I(R)} b$ if and only if $g_R(a, b) \geq -k$ for some $k \in \mathbb{N}$) or on *shifted majority relations* (i.e., $a \succ_{I(R)} b$ if $g_R(a, b) > k$, $a \sim_{I(R)} b$ if $g_R(a, b) = k$, and $b \succ_{I(R)} a$ otherwise) are strategyproof.

For some information bases $I(R)$, it is even possible to prove statements analogous to Theorem 1 when demanding exclusive dependence on $I(R)$. To this end, we say an SCC is $I(R)$ -based for some information basis $I(R)$ if $f(R) = f(R')$ for all preference profiles $R, R' \in \mathcal{R}^*(A)$ such that $\succsim_{I(R)} = \succsim_{I(R')}$. This extends the definition of majoritarianism. For instance, it is easy to derive from our proof that robust dominant set rules on a supermajority relation $I(R)$ are the only neutral, nonimposing, strategyproof, and $I(R)$ -based SCCs. An equivalent statement holds for shifted majority relations $I(R)$ when defining neutrality based on $\succsim_{I(R)}$.

REMARK 11 (Fixed electorates). A nonstandard assumption in our model is that of a variable electorate. This assumption is necessary, because when fixing the number of voters, the Pareto rule satisfies all axioms of Theorem 1 and Theorem 2 but fails to be a dominant set rule. We now sketch two approaches to adapt our results to a fixed electorate framework. First, we may replace pairwise and homogeneity with majoritarianism. Since the construction of McGarvey (1953) allows us to build every majority relation with at most m^2 voters, we need at most $m^2 + 2$ voters for our results to hold under majoritarianism. The second approach is to restrict attention to profiles whose maximal majority margin is bounded by a constant $c \geq 2$. This is possible because we never need to increase the maximal majority margin to a value larger than c in our proofs.⁶ Using again

⁶This is not in conflict with the fact that we sometimes use homogeneity to duplicate preference profiles in the proof, because it is either possible to entirely avoid these homogeneity applications, or to ensure that all majority margins are 1 before duplicating the profile.

McGarvey's construction, every profile with a majority margin of at most c can be built with cm^2 voters and we thus need at most $cm^2 + 2$ voters for our proof. Hence, we can show that every nonimposing, neutral, strategyproof, homogeneous, and pairwise SCC is a robust dominant set rule for profiles with maximal majority margin of at most c if there are $cm^2 + 2$ voters (here, homogeneity is defined for majority margins).

APPENDIX: OMITTED PROOFS

This Appendix contains the proofs of Theorems 1 and 2. Proof sketches for these results were given in Section 4 and we here focus on the details. Since the proofs are rather involved, we divide them into multiple lemmas, which are organized in subsections to highlight related ideas. In particular, we discuss additional notation in Appendix A.1, some general results on the structure of the top cycle in Appendix A.2, implications of Fishburn strategyproofness for pairwise SCCs in Appendix A.3, a variant of Theorem 1 that relies on strong Condorcet consistency in Appendix A.4, and finally the proofs of our main results in Appendix A.5.

A.1 Notation

Before discussing our proofs, we need to introduce some additional notation. First, we specify how we denote preference relations. We usually write preference relations as comma-separated lists. In these lists, we use $\mathbf{lex}(X)$ and $\mathbf{lex}(X)^{-1}$ to indicate that the alternatives in a set X are ordered lexicographically or inversely lexicographically. For instance, $a, \mathbf{lex}(\{b, c\}), d$ is equivalent to a, b, c, d and means that a is preferred to b , b to c , and c to d . Similarly, $a, \mathbf{lex}(\{b, c\})^{-1}, d$ is equivalent to a, c, b, d . Furthermore, we occasionally interpret a voter's preference relation as a set of tuples and use set operations such as set intersections and set differences to form new preference relations. In particular, we write $\succ_i|_X$ to denote the restriction of \succ_i to X , that is, $\succ_i|_X = \succ_i \cap X^2$. We use the same notation for the majority relation, that is, $\succsim_R|_X$ denotes the restriction of \succsim_R to X . For instance, $\succsim_R|_X = \succsim_{R'}|_X$ means that the majority relations of R and R' agree on the alternatives in X .

The second important concept is that of cycles in the majority relation. A *cycle* in a majority relation \succsim_R is a sequence of $q \geq 2$ alternatives (a_1, \dots, a_q) such that $a_i \succsim_R a_{i+1}$ for all $i \in \{1, \dots, q-1\}$, $a_q \succsim_R a_1$, and $a_i \neq a_j$ for all distinct $i, j \in \{1, \dots, q\}$. Informally, a cycle is a path in \succsim_R that starts and ends at the same alternative and visits every alternative on the cycle (except the first one) only once. For instance, in the majority graph in Figure 1, $C = (a, b, c)$ is a cycle. While slightly overloading notation, we denote with C both the ordered sequence of alternatives that defines a cycle and the set of alternatives contained in the cycle.

Finally, we introduce the notions of connectors and connected sets. The *connected set* A_x of an alternative $x \in A$ in a profile R contains all alternatives (except x) that drop out of the top cycle if we remove x from the preference profile, that is, $A_x = TC(R) \setminus (TC(R|_{A \setminus \{x\}}) \cup \{x\})$. The notion of connected sets helps us to distinguish

the alternatives in the top cycle further: we say an alternative $x \in A$ is a *connector* in R if $A_x \neq \emptyset$. This means intuitively that, if we remove x from the preference profile, x and additional alternatives drop out of the top cycle. In other words, x connects the alternatives in A_x to the rest of the top cycle. Note that an alternative $x \notin TC(R)$ cannot be a connector since $A_x = \emptyset$ for these alternatives and that connectors only exist if $|TC(R)| \geq 3$.

A.2 Structure of the top cycle

For the proofs of our results, it will be helpful to have a deeper understanding of the structure of the top cycle. In more detail, we first show that the top cycle is closely connected to cycles in the majority relation. [Moulin \(1986\)](#) has shown such a statement under the assumption that there are no majority ties: there is a cycle in the majority relation \succsim_R that connects all the alternatives in $TC(R)$. Since we need to allow for majority ties, we generalize this result by interpreting majority ties as bidirectional edges.

LEMMA 2. *Let R be a preference profile. It holds for a set $X \subseteq A$ with $|X| \geq 2$ that $TC(R) = X$ if and only if there is a cycle $C = (a_1, \dots, a_{|X|})$ in \succsim_R such that $C = X$ and $X \succ_R A \setminus X$. Furthermore, $TC(R) = \{x\}$ if and only if x is the Condorcet winner in R .*

PROOF. We first prove that $TC(R) = \{x\}$ if and only if x is the Condorcet winner in R . Thus, note that $\{x\} = TC(R)$ implies that $x \succ_R A \setminus \{x\}$ because the top cycle returns a dominant set. Hence, x is the Condorcet winner if it is the unique winner of the top cycle. Next, let x denote the Condorcet winner in a preference profile R . It follows that $x \succ_R A \setminus \{x\}$ and, therefore, $\{x\}$ is a dominant set. Even more, it is obviously the smallest dominant set, and thus $TC(R) = \{x\}$, which proves the first claim.

Next, we focus on sets of alternatives $X \subseteq A$ with $|X| \geq 2$ and show first that if $X = TC(R)$, there is a cycle C in \succsim_R such that $C = X$ and $X \succ_R A \setminus X$. Since the latter condition directly follows from the definition of the top cycle, we only have to show that there is a cycle in \succsim_R containing all alternatives in X . Note for this that if there is an alternative $x \in X$ with $x \succ_R X \setminus \{x\}$, this alternative is the Condorcet winner and $TC(R) = \{x\} \neq X$. Consequently, for every alternative $x \in X$, there is another alternative $y \in X \setminus \{x\}$ such that $y \succsim_R x$. This means that there is a cycle in $\succsim_R|_X$. Let $C = (a_1, \dots, a_q)$ denote an inclusion maximal cycle in $\succsim_R|_X$ and assume for contradiction that there is an alternative $y \in X \setminus C$.

As a first step, consider the case that there are two distinct alternatives $a_i, a_j \in C$ such that $a_i \succsim_R y$ and $y \succsim_R a_j$. In this case, we can extend the cycle C by adding y , which contradicts the inclusion maximality of C . Note for this that we can find two alternatives $a_k, a_{k+1} \in C$ such that a_{k+1} is the successor of a_k in C , $a_k \succsim_R y$, and $y \succsim_R a_{k+1}$. Otherwise, it holds for all $a_l \in C$ that $a_l \succsim_R y$ implies for its successor a_{l+1} in C that $a_{l+1} \succ_R y$. If we start at a_i and subsequently apply this argument along the cycle C , we derive eventually that $a_l \succ_R y$ for all $a_l \in C$, which contradicts that $y \succsim_R a_j$. Hence, there must be such alternatives a_k and a_{k+1} and we can extend the cycle C to $C' = (a_1, \dots, a_k, y, a_{k+1}, \dots, a_q)$.

As a consequence of the last case, it holds for all alternatives $x \in X \setminus C$ that either $x \succ_R C$ or $C \succ_R x$. We partition the alternatives in $X \setminus C$ with respect to these two options into the sets $X_1 = \{x \in X \setminus C : x \succ_R C\}$ and $X_2 = \{x \in X \setminus C : C \succ_R x\}$. If $X_1 = \emptyset$, then $C \succ_R A \setminus C$, which contradicts that $X = TC(R)$ because C is a smaller dominant set than X . If $X_2 = \emptyset$ or $X_1 \succ_R X_2$, then $X_1 \succ_R A \setminus X_1$, which again contradicts that $X = TC(R)$ because X_1 is now a smaller dominant set than X . Thus, both X_1 and X_2 are nonempty and there is a pair of alternatives $x_1 \in X_1$, $x_2 \in X_2$ such that $x_2 \succsim_R x_1$. However, this means that we can extend the cycle C by adding x_1 and x_2 as $a_1 \succ_R x_2$, $x_2 \succsim_R x_1$, and $x_1 \succ_R a_2$. This contradicts the inclusion maximality of C and, therefore, the initial assumption that $C \neq X$ was incorrect.

Finally, we prove that $TC(R) = X$ for a set $X \subseteq A$ with $|X| \geq 2$ if there is a cycle $C = (a_1, \dots, a_{|X|})$ in \succsim_R with $C = X$ and $X \succ_R A \setminus X$. Note for this that X is a dominant set in \succsim_R if it satisfies these conditions. Since dominant sets are totally ordered by set inclusion and the top cycle is the smallest dominant set, it follows that $TC(R) \subseteq X$. Next, assume that $X \setminus TC(R) \neq \emptyset$, which means that $TC(R) \succ_R X \setminus TC(R)$ because of the definition of the top cycle. However, then there cannot be a cycle in \succsim_R that connects all alternatives in X because there is no path from an alternative in $X \setminus TC(R)$ to an alternative in $TC(R)$. This contradicts our assumptions, and thus the assumption $X \setminus TC(R) \neq \emptyset$ was incorrect. Hence, it follows that $X = TC(R)$. \square

Lemma 2 is one of the most important insights for our subsequent proofs as the existence of the cycle provides paths between all alternatives $x, y \in TC(R)$. This insight will also be used in the next lemma, where we investigate connected sets.

LEMMA 3. *Let R be a preference profile and suppose that x is a connector in R . Moreover, let $y \in A_x$ denote an alternative in the connected set of x . It holds that $A_y \subseteq A_x$ unless $x \succ_R A \setminus \{x, y\}$.*

PROOF. Consider an arbitrary preference profile R and a connector x in R . Note that the existence of a connector implies that $k = |TC(R)| \geq 3$. Thus, let $C = (a_1, \dots, a_k)$ denote a cycle connecting the alternatives in $TC(R)$; such a cycle exists because of Lemma 2. Since connectors need to be in the top cycle, it follows that there is an index i such that $x = a_i$. In the sequel, we assume without loss of generality that $x = a_1$ since we can decide on the starting point of the cycle.

As a first step, we show that there is an index $l \in \{2, \dots, k-1\}$ such that $A_x = \{a_{l+1}, \dots, a_k\}$. Consider for this the profile $R^{-x} = R|_{A \setminus \{x\}}$ derived from R by removing x from the preference profile. We next determine the top cycle in R^{-x} because $A_x = TC(R) \setminus (TC(R^{-x}) \cup \{x\})$. First, note that $TC(R^{-x}) \subseteq TC(R) \setminus \{x\}$ because all alternatives in $TC(R) \setminus \{x\} \neq \emptyset$ still strictly dominate all alternatives outside of this set. This implies that a_2 , the successor of $x = a_1$ on C , is in $TC(R^{-x})$ because it can reach every other alternative $a_i \in TC(R) \setminus \{a_1, a_2\}$ via $\succsim_{R^{-x}}$: we can simply traverse the cycle C to go from a_2 to a_i . Now, if $a_2 \succ_R TC(R) \setminus \{a_1, a_2\}$, then a_2 is the Condorcet winner in R^{-x} , and thus $A_x = \{a_3, \dots, a_k\}$, so $l = 3$ satisfies our condition. Otherwise, let $h_1 \in \{3, \dots, k\}$ denote the largest index such that $a_{h_1} \succsim_R a_2$. It follows from the definition of the top cycle

that $a_{h_1} \in TC(R^{-x})$ because $a_2 \in TC(R^{-x})$, and thus $\{a_2, \dots, a_{h_1}\} \subseteq TC(R^{-x})$ because all these alternatives can reach a_{h_1} by traversing the cycle C . It is easy to see that we can repeat this argument: if $\{a_2, \dots, a_{h_1}\} \succ_R \{a_{h_1+1}, \dots, a_k\}$, then $TC(R^{-x}) = \{a_2, \dots, a_{h_1}\}$ and $A_x = \{a_{h_1+1}, \dots, a_k\}$. Otherwise, we can find the largest index $h_2 \in \{h_1 + 1, \dots, k\}$ such that a_{h_2} dominates an alternative in $\{a_2, \dots, a_{h_1}\}$. Then it follows that $a_{h_2} \in TC(R^{-x})$, and consequently, $\{a_2, \dots, a_{h_2}\} \subseteq TC(R^{-x})$ because we can traverse the cycle C to find a path for every such alternative a_i to a_{h_2} . By repeating this argument, we eventually arrive at an index l such that $A_x = \{a_{l+1}, \dots, a_k\}$ since $A_x \neq \emptyset$.

Next, consider an arbitrary alternative $y \in A_x$. Our goal is to prove that $A_y \subseteq A_x$, and thus we consider the profile $R^{-y} = R|_{A \setminus \{y\}}$. We will show that $TC(R^{-x}) \cup \{x\} \subseteq TC(R^{-y})$ because then $A_y = TC(R) \setminus (TC(R^{-y}) \cup \{y\}) \subseteq TC(R) \setminus (TC(R^{-x}) \cup \{x\}) = A_x$. For this, we employ a case distinction with respect to y and first suppose that y is not the direct predecessor of x on C , that is, $y = a_i$ for some $i < k$. Hence, let $y' = a_{i+1}$ denote the successor of y on C and note that our previous insights show that $y' \in A_x$, too. It holds that $y' \in TC(R^{-y})$ because y' can reach every alternative $a_j \in TC(R) \setminus \{y, y'\}$ in $\succsim_{R^{-y}}$ by traversing the cycle C . Next, note that $TC(R^{-x}) \succ_{R^{-x}} y'$ because $y' \notin TC(R^{-x})$ and, therefore, also $TC(R^{-x}) \succ_{R^{-y}} y'$. This proves that $TC(R^{-x}) \subseteq TC(R^{-y})$ because $y' \in TC(R^{-y})$. In particular, $x' = a_2$, the successor of $x = a_1$ on the cycle C , is in $TC(R^{-y})$ because $x' \in TC(R^{-x})$. Since $x \succsim_{R^{-y}} x'$, it follows also that $x \in TC(R^{-y})$, which proves that $TC(R^{-x}) \cup \{x\} \subseteq TC(R^{-y})$, and thus $A_y \subseteq A_x$.

As second case, suppose that $y = a_k$, that is, y is the direct predecessor of $x = a_1$ on C . In this case, we immediately derive that $x \in TC(R^{-y})$ because we can again traverse the cycle C to find a path from x to every other alternative $a_i \in TC(R) \setminus \{x, y\}$ in $\succsim_{R^{-y}}$. Next, it is important that there is an alternative $z \in A \setminus \{x, y\}$ such that $z \succsim_R x$. If there is no such alternative, then $x \succ_R A \setminus \{x, y\}$ and we have nothing to show as this is the exception stated in the lemma. Since $z \neq y$ and $z \succsim_R x$, it follows also that $z \in TC(R)$ and $z \in TC(R^{-y})$. Now, if $z \in A_x$, then $TC(R^{-x}) \subseteq TC(R^{-y})$ because $TC(R^{-x}) \succ_R z$. Conversely, if $z \in TC(R^{-x}) = TC(R) \setminus (A_x \cup \{x\})$, we use the fact that there is a cycle C' connecting the alternatives $TC(R^{-x})$ in $\succsim_{R^{-x}}$. This cycle exists also in \succsim_R , and since $y \notin TC(R^{-x})$, also in $\succsim_{R^{-y}}$. Hence, there is a path from every alternative $a_i \in TC(R^{-x})$ to z , which proves that $TC(R^{-x}) \cup \{x\} \subseteq TC(R^{-y})$. Thus, it follows also in this case that $A_y \subseteq A_x$, which proves the lemma. \square

A.3 Implications of strategyproofness

In the context of pairwise SCCs, it is inconvenient to work with the preference relations of individual voters since the main idea of these SCCs is to abstract away from profiles. However, strategyproofness requires information about a voter's preference relation to deduce which choice sets are possible before and after a manipulation. To mitigate this tradeoff, we analyze the implications of strategyproofness for pairwise SCCs in this section. This leads to the definition of four axioms, all of which are satisfied by every pairwise and strategyproof SCC. Also, the first three of these axioms are weakened versions of a property known as set-monotonicity (see Brandt (2015), Brandt, Brill, and Harrenstein (2016)).

In more detail, we investigate how the choice set of a strategyproof and pairwise SCC is allowed to change if a voter reinforces or weakens an alternative against some other alternatives. Formally, reinforcing an alternative a against some other alternative b in the preference relation of voter i means that voter i switches from $b \succ_i a$ to $a \succ'_i b$ and nothing else changes in voter i 's preference relation or in the preference relations of other voters. Conversely, weakening an alternative a against some other alternative b in the preference relation of voter i means that voter i reinforces b against a . Note that weakening or reinforcing an alternative a against another alternative b requires that a and b are adjacent in \succ_i , that is, there is no alternative $z \in A \setminus \{a, b\}$ such that $a \succ_i z \succ_i b$ or $b \succ_i z \succ_i a$.

Depending on whether the alternatives a and b are chosen, strategyproofness has different consequences when reinforcing a against another alternative b . The first case that we consider is to reinforce a chosen alternative a against another alternative b . A natural requirement in this situation is monotonicity, which demands that a chosen alternative is still chosen after reinforcing it (see, e.g., [Moulin \(1988\)](#)). Unfortunately, we cannot show that strategyproofness implies monotonicity for pairwise SCCs. For instance, assume that a voter submits b, a, c , and $\{a, c\}$ is chosen. Next, voter i reinforces a against b and as result $\{b, c\}$ is chosen. In this example, Fishburn's set extension does not allow to compare $\{a, c\}$ to $\{b, c\}$, and hence, this is no violation of strategyproofness. As a consequence, we consider a weakened variant of monotonicity, which we refer to as weak monotonicity (WMON). This axiom requires that, if a voter reinforces a chosen alternative a against another alternative b , then a is still in the choice set unless b is chosen after the manipulation but not before.

DEFINITION 1 (Weak monotonicity (WMON)). An SCC f satisfies weak monotonicity (WMON) if $a \in f(R)$ implies $a \in f(R')$ or $b \in f(R') \setminus f(R)$ for all alternatives $a, b \in A$, and preference profiles R, R' for which there is a voter i such that $\succ'_j = \succ_j$ for all $j \in N \setminus \{i\}$ and $\succ'_i = \succ_i \setminus \{(b, a)\} \cup \{(a, b)\}$.

WMON has multiple important consequences. First, if we reinforce a chosen alternative a against another chosen alternative b , it guarantees that a remains chosen because $b \notin f(R') \setminus f(R)$. Second, if we reinforce a chosen alternative a against an unchosen alternative b , either $a \in f(R')$ and $b \notin f(R')$, or $a \notin f(R')$ and $b \in f(R')$. If both alternatives were chosen after this step, we could reinforce b against a in voter i 's preference relation to revert back to the original preference profile R , and WMON implies that b remains chosen. However, this is in conflict with the assumption that b is not chosen for R . Conversely, it follows directly from the definition of WMON that it is not possible that $a, b \notin f(R')$ if $a \in f(R)$. Finally, it should be mentioned that monotonicity implies weak monotonicity because it requires that a chosen alternative a remains chosen after reinforcing it. Or, put differently, monotonicity excludes additionally the case that a becomes unchosen and b becomes chosen after reinforcing a against b .

Unfortunately, WMON does not guarantee that weakening an unchosen alternative means that the unchosen alternative remains unchosen. We thus introduce weak set-monotonicity (WSMON) as our second axiom, which is concerned with what happens if

we weaken an unchosen alternative not against a single alternative but against all other alternatives.

DEFINITION 2 (Weak set-monotonicity (WSMON)). An SCC f satisfies weak set-monotonicity (WSMON) if $f(R) = f(R')$ for all preference profiles R, R' for which a voter $i \in N$ and an alternative $a \notin f(R)$ exist such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$, $\succ_i|_{A \setminus \{a\}} = \succ'_i|_{A \setminus \{a\}}$, $a \succ_i A \setminus \{a\}$, and $A \setminus \{a\} \succ'_i a$.

The idea of WSMON is that moving an unchosen alternative from the first place to the last place in a voter's preference relation should not affect the outcome. This is a weaker variant of set-monotonicity, which requires that weakening an unchosen alternative against a single alternative does not affect the choice set. Unfortunately, we cannot prove this stronger variant because we cannot even prove monotonicity at this point. However, pushing the top-ranked alternative to the bottom of the preference ranking is a rather common operation in the analysis of strategyproof SCCs, which is often referred to as push-down lemma (see, e.g., [Zwicker \(2016\)](#)).

The third situation that we are concerned with is that a voter only reorders unchosen alternatives. Intuitively, such an operation should not change the choice set as no relevant comparisons change. This idea is formalized as independence of unchosen alternatives (IUA).

DEFINITION 3 (Independence of unchosen alternatives (IUA)). An SCC f satisfies independence of unchosen alternatives (IUA) if $f(R) = f(R')$ for all preference profiles R, R' for which a voter $i \in N$ and alternatives $B \subseteq A \setminus f(R)$ exist such that $\succ_j = \succ'_j$ for all voters $j \in N \setminus \{i\}$ and $\succ_i \setminus \succ_i|_B = \succ'_i \setminus \succ'_i|_B$.

Independence of unchosen alternatives, also called independence of losers, is a well-known axiom (see, e.g., [Laslier \(1997\)](#), [Brandt \(2011, 2015\)](#)), which requires that the choice set is invariant with respect to modifications of preferences between unchosen alternatives. In particular, if a voter reinforces an unchosen alternative against another unchosen alternative, the choice set is not allowed to change. Just like WMON and WSMON, IUA is implied by set-monotonicity.

Finally, we introduce an axiom with a different spirit than the previous ones: instead of asking whether an alternative a is chosen after weakening or reinforcing it, we ask whether alternatives that are not involved in the swap are chosen or not. Intuitively, it seems plausible that if an alternative is not affected by a manipulation, its membership in the choice set should not change. However, this condition, whose spirit is similar to the localizedness property used in the characterization of strategyproof randomized social choice functions by [Gibbard \(1977\)](#), is extremely restrictive. Here, we consider a weaker variant: if a voter changes his preference relation between some alternatives B and the inclusion of the alternatives in B in the choice set is unaffected by this modification, then the choice set should not change at all. This idea leads to weak localizedness (WLOC), which is formalized below.

DEFINITION 4 (Weak localizedness (WLOC)). An SCC f satisfies weak localizedness (WLOC) if $f(R) = f(R')$ for all preference profiles R, R' for which a voter $i \in N$ and alternatives $B \subseteq A$ exist such that $\succ_j = \succ'_j$ for all voters $j \in N \setminus \{i\}$, $\succ_i \setminus \succ_i|_B = \succ'_i \setminus \succ'_i|_B$, and $B \cap f(R) = B \cap f(R')$.

To the best of our knowledge, neither WLOC nor similar axioms have been studied before for social choice correspondences. Furthermore, it should be mentioned that WLOC—even though it might seem weak when considered in isolation—is quite powerful when combined with other axioms. For instance, the combination of WMON and WLOC implies that swapping two chosen alternatives can only affect the choice set if the weakened alternative becomes unchosen.

We now prove that strategyproof and pairwise SCCs satisfy all axioms discussed in this section.

LEMMA 4. *Every strategyproof and pairwise SCC satisfies WMON, WSMON, IUA, and WLOC.*

PROOF. Let f denote a strategyproof and pairwise SCC. We consider each axiom listed in the lemma separately, but each proof relies on the same idea: we assume for contradiction that f fails the considered axiom, which means that there are two profiles R and R' that differ in the preference relation of a single voter i and $f(R)$ and $f(R')$ violate the conditions of the axiom. Next, we add two new voters i^* and j^* with inverse preferences such that voter i^* can make the same modification as voter i . This leads to new preference profiles R^1 and R^2 such that $f(R^1) = f(R)$ and $f(R^2) = f(R')$ due to pairwise-ness. Finally, we can choose the preference relation of voter i^* such that deviating from R^1 to R^2 is a manipulation, and thus obtain a contradiction to the strategyproofness of f .

WMON: Following the idea explained above, we assume for contradiction that f violates WMON. This means that there are preference profiles R, R' , alternatives $a \in f(R)$, $b \in A \setminus \{a\}$, and a voter $i \in N$ such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$ and $\succ'_i = \succ_i \setminus \{(b, a)\} \cup \{(a, b)\}$, but $a \notin f(R')$ and $b \notin f(R') \setminus f(R)$. Next, we let R^1 denote the profile derived from R by adding the voters i^* and j^* . The preference relations of these voters are shown below, where $\bar{A} = A \setminus \{a, b\}$ and $\bar{f}(R) = f(R) \setminus \{a, b\}$. Moreover, the profile R^2 evolves out of R^1 by letting voter i^* swap a and b :

$$\begin{aligned}\succ_{i^*}^1 &= \mathbf{lex}(\bar{A} \setminus f(R)), \mathbf{lex}(\bar{f}(R) \cap f(R')), b, a, \mathbf{lex}(\bar{f}(R) \setminus f(R')) \\ \succ_{j^*}^1 &= \mathbf{lex}(\bar{f}(R) \setminus f(R'))^{-1}, a, b, \mathbf{lex}(\bar{f}(R) \cap f(R'))^{-1}, \mathbf{lex}(\bar{A} \setminus f(R))^{-1}\end{aligned}$$

Since the preference relations of voter i^* and j^* are inverse in R^1 , pairwise-ness implies that $f(R^1) = f(R)$. Moreover, this axiom also requires that $f(R^2) = f(R')$. It then follows that voter i^* can manipulate by deviating from R^1 to R^2 as he prefers all alternatives in $f(R') \setminus f(R)$ to those in $f(R)$ and all alternatives in $f(R')$ to those in $f(R) \setminus f(R')$. This can be seen by making a case distinction on whether $b \in f(R)$: if $b \notin f(R)$, then our contradiction assumption implies that $b \notin f(R')$, too. Hence, no alternative in $(f(R) \setminus f(R')) \cup \{b\}$ is chosen, which ensures that this is a manipulation for voter i^* .

since $a \in f(R) \setminus f(R')$. Conversely, if $b \in f(R)$, then b is either in $f(R) \cap f(R')$ or in $f(R) \setminus f(R')$. Both cases constitute again a manipulation as $b \succ_{i^*} f(R) \setminus (f(R') \cup \{b\})$ and $(f(R) \cap f(R')) \setminus \{b\} \succ_{i^*} b$. Hence, switching from R^1 to R^2 is in all cases a manipulation for voter i^* , which contradicts the strategyproofness of f . Consequently, the initial assumption that f violates WMON was incorrect.

WSMON: As second case, assume that f fails WSMON. Thus, there are preference profiles R, R' , a voter $i \in N$, and an alternative $a \notin f(R)$ such that R and R' only differ in the fact that $a \succ_i A \setminus \{a\}$ and $A \setminus \{a\} \succ'_i a$, but $f(R) \neq f(R')$. Consider the profile R^1 , which is derived from R by adding the voters i^* and j^* with the preferences shown below. Moreover, R^2 evolves out of R^1 by letting voter i^* make a into his least preferred alternative:

$$\begin{aligned} \succ_{i^*}^1 &= a, \mathbf{lex}(A \setminus (\{a\} \cup f(R))), \mathbf{lex}(f(R) \cap f(R')), \mathbf{lex}(f(R) \setminus f(R')) \\ \succ_{j^*}^1 &= \mathbf{lex}(f(R) \setminus f(R'))^{-1}, \mathbf{lex}(f(R) \cap f(R'))^{-1}, \mathbf{lex}(A \setminus (\{a\} \cup f(R)))^{-1}, a \end{aligned}$$

It is again easy to verify that $f(R^1) = f(R)$ and $f(R^2) = f(R')$ because of pairwise-ness. Thus, voter i^* can manipulate f by switching from R^1 to R^2 because he prefers all alternatives in $A \setminus f(R)$ to all alternatives in $f(R)$ and all alternatives in $f(R) \cap f(R')$ to all alternatives in $f(R) \setminus f(R')$. This contradicts the strategyproofness of f and, therefore, the initial assumption that f violates WSMON was incorrect.

IUA: Third, assume that f violates IUA, which means that there are preference profiles R, R' , a voter $i \in N$, and a set of alternatives $B \subseteq A \setminus f(R)$ such that $\succ_j = \succ'_j$ for all voters $j \in N \setminus \{i\}$, $\succ_i \setminus \succ_i|_B = \succ'_i \setminus \succ'_i|_B$, and $f(R) \neq f(R')$. Now, consider the profile R^1 derived from R by adding two voters i^* and j^* . The preference relations of these two voters are shown below, where $\bar{A} = A \setminus B$, $\succ_i|_B$ indicates that the alternatives in B are ordered as in \succ_i , and $\succ_i^{-1}|_B$ that the alternatives in B are ordered exactly inverse to \succ_i . Moreover, let R^2 denote the profile derived from R^1 by letting voter i^* order the alternatives in B as voter i does in R' :

$$\begin{aligned} \succ_{i^*}^1 &= \succ_i|_B, \mathbf{lex}(\bar{A} \setminus f(R)), \mathbf{lex}(f(R) \cap f(R')), \mathbf{lex}(f(R) \setminus f(R')) \\ \succ_{j^*}^1 &= \mathbf{lex}(f(R) \setminus f(R'))^{-1}, \mathbf{lex}(f(R) \cap f(R'))^{-1}, \mathbf{lex}(\bar{A} \setminus f(R))^{-1}, \succ_i^{-1}|_B \end{aligned}$$

Just as before, we infer from pairwise-ness that $f(R^1) = f(R)$ and $f(R^2) = f(R')$. However, this means that voter i^* can manipulate by deviating from R^1 to R^2 : by construction, he prefers all alternatives in $A \setminus f(R)$ to all alternatives in $f(R)$, and all alternatives in $f(R) \cap f(R')$ to all alternatives in $f(R) \setminus f(R')$. Since $f(R) \neq f(R')$, this is in conflict with strategyproofness.

WLOC: Finally, suppose for contradiction that f violates WLOC. Hence, there are two preference profiles R, R' , a nonempty set of alternatives $B \subseteq A$, and a voter $i \in N$ such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$, $\succ_i \setminus \succ_i|_B = \succ'_i \setminus \succ'_i|_B$, $f(R) \cap B = f(R') \cap B$, but $f(R) \neq f(R')$. Once again, we derive a new profile R^1 from R by adding two voters i^* and j^* . The preferences of these voters are shown below, where $\bar{A} = A \setminus B$ and $\bar{f}(R) = f(R) \setminus$

B . Moreover, let R^2 denote the profile derived from R^1 by letting voter i^* arrange the alternatives in B according to \succ'_i :

$$\begin{aligned}\succ_{i^*}^1 &= \mathbf{lex}(\bar{A} \setminus f(R)), \succ_i|_B, \mathbf{lex}(\bar{f}(R) \cap f(R')), \mathbf{lex}(\bar{f}(R) \setminus f(R')) \\ \succ_{j^*}^1 &= \mathbf{lex}(\bar{f}(R) \setminus f(R'))^{-1}, \mathbf{lex}(\bar{f}(R) \cap f(R'))^{-1}, \succ_i^{-1}|_B, \mathbf{lex}(\bar{A} \setminus f(R))^{-1}\end{aligned}$$

Also in this case, pairwise-ness shows that $f(R^1) = f(R)$ and $f(R^2) = f(R')$. Thus, voter i^* can manipulate by switching from R^1 to R^2 because $f(R') \succ_{i^*}^1 f(R) \setminus f(R')$ and $f(R') \setminus f(R) \succ_{i^*}^1 f(R)$. For both claims, it is important that B is disjoint to both $f(R) \setminus f(R')$ and $f(R') \setminus f(R)$ since an alternative $x \in B$ is in $f(R)$ if and only if it is in $f(R')$. Hence, the first claim follows directly as the alternatives in $f(R) \setminus f(R')$ are the least preferred ones of voter i^* , and the second claim follows since $f(R') \setminus f(R) \subseteq A \setminus (B \cup f(R))$. Thus, deviating from R^1 to R^2 is a manipulation for voter i^* , which contradicts the strategyproofness of f . \square

A.4 Consequences of strong Condorcet consistency

In this section, we prove a variant of Theorem 1 which relies on *strong Condorcet consistency* instead of neutrality and nonimposition. This axiom requires of an SCC f that $f(R) = \{x\}$ if and only if x is the Condorcet winner in R . Less formally, strongly Condorcet-consistent SCCs have to elect the Condorcet winner whenever there is one, and cannot elect a single alternative in the absence of a Condorcet winner. The main result of this section states that robust dominant set rules are the only SCCs that satisfy pairwise-ness, homogeneity, strong Condorcet consistency, and strategyproofness. As we will see in Appendix A.5, the combination of pairwise-ness, homogeneity, strategyproofness, neutrality, and nonimposition implies strong Condorcet consistency, which means that this auxiliary claim is actually more general than Theorem 1. Moreover, the axioms of Theorem 2 imply strong Condorcet consistency and we can therefore use the results of this section also to characterize the top cycle.

For proving the results of this section, we rely on the lemmas of the previous subsections. In particular, we often say that we reinforce an alternative x against another alternative y without specifying which voter reinforces x against y . This is possible because we can always add two voters with inverse preferences such that one of them can perform the required manipulation. Adding these two voters does not affect the choice set because of pairwise-ness and the consequences of the deviation will be specified by the axioms of Appendix A.3. Hence, we can abstract away from the exact preference profiles and focus on the majority margins. For the readers' convenience, we repeat the four axioms of the last section because they form the basis of the following proofs. Let f be a pairwise SCC, $a, b \in A$, and $R, R' \in \mathcal{R}^*(A)$.

- **WMON:** If R' is derived from R by reinforcing a against b and $a \in f(R)$, then $a \in f(R)$ or $b \in f(R') \setminus f(R)$.
- **WSMON:** If R' is derived from R by weakening a against all other alternatives $x \in A \setminus \{a\}$ and $a \notin f(R)$, then $f(R) = f(R')$.

- IUA: If R' is derived from R by reordering some alternatives in $A \setminus f(R)$, then $f(R) = f(R')$.
- WLOC: If R' is derived from R by reordering the alternatives in $B \subseteq A$ such that $f(R) \cap B = f(R') \cap B$, then $f(R) = f(R')$.

As shown in Appendix A.3, every strategyproof and pairwise SCC satisfies these axioms. We now use these properties to show our first key insight, namely that all such SCCs that satisfy strong Condorcet consistency also satisfy a new property called Condorcet stability (COS). This axiom requires that there should be no alternative—within or outside of the choice set—that strictly dominates all other alternatives in the choice set. Note that COS implies that an alternative can only be a single winner if it weakly dominates every other alternative.⁷

DEFINITION 5 (Condorcet stability (COS)). An SCC f satisfies Condorcet stability (COS) if for every preference profile R , there is no alternative $x \in A$ such that $x \succ_R f(R) \setminus \{x\}$ whenever $f(R) \setminus \{x\}$ is nonempty.

This condition is equivalent to requiring that every alternative is weakly dominated by another chosen alternative unless it is the unique winner, that is, for every alternative $x \in A$ with $f(R) \neq \{x\}$ there is an alternative $y \in f(R) \setminus \{x\}$ such that $y \succeq_R x$. It is thus closely connected to the notion of external stability, which requires that for every alternative $x \in A \setminus f(R)$, there is an alternative $y \in f(R)$ such that $y \succeq_R x$ (see, e.g., Miller, Grofman, and Feld (1990), Duggan (2013)). Indeed, Condorcet stability is a stronger requirement than external stability as it also includes a notion of internal stability.

As we show next, the conjunction of our axioms implies COS.

LEMMA 5. *Every pairwise SCC that is strategyproof and strongly Condorcet consistent satisfies COS.*

PROOF. Let f denote a pairwise SCC that satisfies strategyproofness and strong Condorcet consistency. First, recall that strong Condorcet consistency requires that $f(R) = \{x\}$ if and only if x is the Condorcet winner in R . Hence, strong Condorcet consistency implies COS for all profiles R with a Condorcet winner x because $f(R) = \{x\}$ entails $x \succ_R A \setminus \{x\}$. Next, we focus on profiles without a Condorcet winner and assume for contradiction that f fails COS for such a profile. More formally, this assumption means that there is a profile R without a Condorcet winner and an alternative $a \in A$ such that $a \succ_R f(R) \setminus \{a\}$. It follows from the absence of a Condorcet winner that there is at least one alternative x with $x \succeq_R a$, but no such alternative is chosen. Hence, we can repeatedly use WSMON to weaken the alternatives x with $x \succeq_R a$ against all other alternatives until we arrive at a profile R' such that $a \succ_{R'} A \setminus \{a\}$. Now, WSMON implies that the choice set does not change during these steps, and thus it holds that $f(R') = f(R)$. This

⁷This axiom is quite useful for characterizing SCCs that are strategyproof in profiles that admit a Condorcet winner: a majoritarian and nonimposing SCC is strategyproof in such profiles if and only if it satisfies COS.

means in particular that $|f(R')| \geq 2$ because strong Condorcet consistency requires that $|f(R)| \geq 2$. However, a is the Condorcet winner in R' , and consequently, strong Condorcet consistency also shows that $f(R') = \{a\}$. These two observations contradict each other, and consequently, the assumption that f violates COS was incorrect. \square

COS plays an important role in our proofs because we can use it to force an alternative into the choice set. In particular, the combination of strong Condorcet consistency and COS have rather strong consequences: the first axiom states that we choose a single winner if and only if it is the Condorcet winner and the second one requires therefore that every alternative is weakly dominated by a chosen alternative if there is no Condorcet winner. We now use this interaction to prove our first lemma for pairwise SCCs that satisfy homogeneity, strategyproofness, and strong Condorcet consistency.

LEMMA 6. *Let f denote a pairwise SCC that satisfies strong Condorcet consistency, homogeneity, and strategyproofness. If $TC(R) \subseteq f(R)$ for all profiles R , then f is a robust dominant set rule.*

PROOF. Let f denote a pairwise, homogeneous, strategyproof, and strongly Condorcet-consistent SCC. We prove this lemma in three steps: first, we show that if $TC(R) \subseteq f(R)$ for all profiles R , then f is a dominant set rule. Next, we prove that, if f is a dominant set rule, our assumptions require it to be majoritarian. As last point, we show that f is even robust if it is a majoritarian dominant set rule. Combining all three steps thus shows the lemma.

Step 1: If $TC(R) \subseteq f(R)$ for all profiles R , f is a dominant set rule.

We prove this claim by contradiction, and thus assume that f always chooses a superset of TC but is no dominant set rule. This means that there is a profile R such that $f(R)$ is no dominant set in \succsim_R , which implies that $f(R) \neq TC(R)$. Moreover, there is no Condorcet winner in R , because otherwise strong Condorcet consistency would require that f chooses this alternative as unique winner, which would contradict that $f(R)$ is no dominant set. We infer from this observation that $|TC(R)| > 1$ because TC is strongly Condorcet consistent. Next, note that there are alternatives $a \in f(R)$ and $b \in A \setminus f(R)$ such that $b \succsim_R a$ because $f(R)$ is no dominant set. Even more, $a \notin TC(R)$; otherwise, b would also be in $TC(R)$ because $b \succsim_R a$, which would imply that $TC(R) \not\subseteq f(R)$ since $b \notin f(R)$. However, this contradicts our assumptions.

Next, let x denote an alternative in $TC(R) \subseteq f(R)$, which implies that $x \succ_R a$. We repeatedly reinforce a against x until we arrive at a profile R' such that $a \succ_{R'} x$. Moreover, let \bar{R}' denote the last profile constructed before R' , that is, $x \succ_{\bar{R}'} a$ and a single voter needs to reinforce a against x to derive R' . First, we show that $f(\bar{R}') = f(R)$ and for this consider two consecutive profiles \hat{R} and \hat{R}' in the sequence that leads from R to \bar{R}' . This means that \hat{R}' is derived from \hat{R} by reinforcing a against x . Moreover, it holds that $\succsim_R = \succsim_{\hat{R}} = \succsim_{\hat{R}'}$ because the majority relation between a and x has not changed yet. This implies that $TC(R) = TC(\hat{R}) = TC(\hat{R}')$ and, therefore, $x \in TC(\hat{R}) \subseteq f(\hat{R})$ and $x \in TC(\hat{R}') \subseteq f(\hat{R}')$. Now, if $a \in f(\hat{R})$, then WMON implies that $a \in f(\hat{R}')$ since a is reinforced against x to derive \hat{R}' . Finally, WLOC implies then that $f(\hat{R}) = f(\hat{R}')$ because a and x

are both chosen before and after the manipulation. Since we start this process at the profile R with $\{a, x\} \subseteq f(R)$, it follows from a repeated application of this argument that $f(\bar{R}') = f(R)$.

Finally, we show that $\{b, x\} \not\subseteq f(R')$ but $\{b, x\} \subseteq TC(R')$. This contradicts the assumption that TC is always contained in f , and thus proves that our initial assumption that f is no dominant set rule was incorrect. First, we show that $\{b, x\} \not\subseteq f(R')$. Observe for this that $b \notin f(\bar{R}') = f(R)$ and that $\{a, x\} \subseteq f(\bar{R}') = f(R)$. Since R' is derived from \bar{R}' by reinforcing a against x , WMON implies that $a \in f(R')$. Now, if $x \in f(R')$, it follows from WLOC that the choice set is not allowed to change, which implies that $b \notin f(R')$. This means that $b \in f(R')$ is only possible if $x \notin f(R')$, so $\{b, x\} \not\subseteq f(R')$. Next, we prove the second claim that $\{b, x\} \subseteq TC(R')$. For proving this, it is important that $|TC(R)| > 1$ and that $\succsim_{R'}$ differs from \succsim_R only in the fact that $a \succsim_{R'} x$ and $x \succ_R a$. The first point means that there is a cycle C in \succsim_R connecting all alternatives in $TC(R)$ because of Lemma 2 and the second one that this cycle also exists in $\succsim_{R'}$. Hence, there is a path from x to every alternative $y \in TC(R) \setminus \{x\}$. Moreover, there is a path from x to every alternative $z \in A \setminus TC(R)$ because we can go from x to another alternative $y \in TC(R) \setminus \{x\}$ using the cycle C and from y to z because $y \succ_{R'} z$. This means that $x \in TC(R')$ as it reaches every other alternative on some path. Furthermore, $a \in TC(R')$ because $a \succsim_{R'} x$, and $b \in TC(R')$ because $b \succsim_{R'} a$. This shows that $\{b, x\} \subseteq TC(R')$, even though $\{b, x\} \not\subseteq f(R')$. Hence, $TC(R') \not\subseteq f(R')$, which contradicts the assumption that $TC(R) \subseteq f(R)$ for all profiles R . This proves that the assumption that f is no dominant set rule was incorrect.

Step 2: If f is a dominant set rule, it is majoritarian.

Our goal in this step is to show that if f is a dominant set rule, it is majoritarian. Thus, assume for contradiction that f is a dominant set rule but violates majoritarianism. The latter point means that there are two preference profiles R and R' such that $\succsim_R = \succsim_{R'}$ but $f(R) \neq f(R')$. We assume that both R and R' are defined by an even number of voters. This is without loss of generality as we can just duplicate the profiles if required. The majority relations do not change by this step since the majority margins are only doubled, and the choice sets do not change because of homogeneity. Thus, we can also work with these larger profiles instead. Next, observe that $f(R) \neq f(R')$ implies that $g_R \neq g_{R'}$ because f is pairwise. Moreover, $f(R)$ and $f(R')$ are both dominant sets in \succsim_R because f is a dominant set rule. Since dominant sets are ordered by set inclusion, it follows that $f(R) \subsetneq f(R')$ or $f(R') \subsetneq f(R)$. We assume without loss of generality that $f(R)$ is a subset of $f(R')$; otherwise, we can just exchange the role of R and R' in the subsequent arguments. Our goal is to transform R into a profile R^* such that $g_{R^*} = g_{R'}$ and $f(R^*) \subseteq f(R) \subsetneq f(R')$. This is in conflict with the pairwise nature of f and shows therefore that the assumption $f(R) \neq f(R')$ was incorrect.

We use the largest majority margin $c = \max_{x, y \in A} g_R(x, y)$ in R for the derivation of R^* . In more detail, we first construct a profile R^1 such that $g_{R^1}(x, y) = c$ for all alternatives $x, y \in A$ with $x \succ_R y$. For this, we repeatedly use the following steps: first, identify a pair of alternatives $x, y \in A$ such that $x \succ_R y$ but the majority margin between x and y is not c yet. Then reinforce x against y . By repeating these steps, we eventually arrive at a profile R^1 , which satisfies $g_{R^1}(x, y) = c$ for all $x, y \in A$ with $x \succ_R y$. We show next that $f(R^1) \subseteq f(R)$ by a case distinction with respect to x and y . For this, consider a single

step of our process, and let \bar{R} denote the profile before reinforcing x against y and \bar{R}' denote the profile after reinforcing x against y . If $x \notin f(\bar{R})$ and $y \notin f(\bar{R})$, it follows from IUA that $f(\bar{R}) = f(\bar{R}')$. If $x \in f(\bar{R})$ and $y \notin f(\bar{R})$, it follows from WMON that $x \in f(\bar{R})$ and $y \notin f(\bar{R}')$ because we have $x \succ_R y$ and, therefore, also $x \succ_{\bar{R}'} y$. Hence, if $y \in f(\bar{R}')$, then $x \in f(\bar{R}')$ as f chooses a dominant set. However, this is in conflict with WMON: if we revert the swap, this axiom implies that $y \in f(\bar{R})$, which contradicts our assumptions. Thus, $x \in f(\bar{R}')$, $y \notin f(\bar{R}')$, and WLOC implies that $f(\bar{R}) = f(\bar{R}')$. As a third point, note that $x \notin f(\bar{R})$ and $y \in f(\bar{R})$ are impossible because we assume that $x \succ_R y$. Hence, this case contradicts that $f(\bar{R})$ is a dominant set. The last case is that $x \in f(\bar{R})$ and $y \in f(\bar{R})$. In this case, it follows from WMON that $x \in f(\bar{R}')$. If now also $y \in f(\bar{R}')$, WLOC implies that $f(\bar{R}) = f(\bar{R}')$. Conversely, if $y \notin f(\bar{R}')$, then $f(\bar{R}') \subsetneq f(\bar{R})$. Otherwise, an alternative $z \in A \setminus f(\bar{R})$ is in $f(\bar{R}')$, which is in conflict with the fact that $f(\bar{R}')$ is a dominant set since $y \succ_{\bar{R}'} z$ because $y \in f(\bar{R})$ and $z \notin f(\bar{R})$. Hence, we derive in all possible cases that $f(\bar{R}') \subseteq f(\bar{R})$. By repeatedly applying this argument, it follows that $f(R^1) \subseteq f(R)$.

Next, note that $\succ_{R^1} = \succ_R$ because we only increase the majority margins between alternatives $x, y \in A$ with $x \succ_R y$. Furthermore, there are only two possible majority margins in R^1 : if $x \sim_{R^1} y$, then $g_{R^1}(x, y) = 0$ and if $x \succ_{R^1} y$, then $g_{R^1}(x, y) = c$. This means that we can use homogeneity to derive a profile R^2 with smaller majority margins: we set $g_{R^2}(x, y) = 2$ for all alternatives $x, y \in A$ with $x \succ_{R^1} y$ and $g_{R^2}(x, y) = 0$ for all alternatives $x, y \in A$, with $x \sim_{R^1} y$. Such a preference profile R^2 exists because we can use McGarvey's construction to build a preference profile for all majority margins that have the same parity (McGarvey (1953)). It follows from homogeneity and pairwise-ness that $f(R^2) = f(R^1)$ because we can just multiply R^2 such that all majority margins are equal to those in R^1 . Note here that the assumption that R is defined by an even number of voters is important because it ensures that c is a multiple of 2. As last point, observe that $\succ_{R^2} = \succ_R = \succ_{R'}$ because we did not change the sign of a majority margin. Moreover, $g_{R^2}(x, y) \leq g_{R'}(x, y)$ for all $x, y \in A$ because R' is defined by an even number of voters. Hence, if $x \succ_{R'} y$, then $g_{R'}(x, y) \geq 2 = g_{R^2}(x, y)$, and if $x \sim_{R'} y$, then $g_{R^2}(x, y) = g_{R'}(x, y) = 0$.

As a last step, we derive a preference profile R^3 with $g_{R^3} = g_{R'}$ from R^2 by applying the same process as in the construction of R^1 : we repeatedly identify a pair of alternatives $x, y \in A$ such that $x \succ_R y$ and the current majority margin between x and y is less than the one in R' , and reinforce x against y . Clearly, this process results in a profile R^3 with $g_{R^3} = g_{R'}$ and the same arguments as for R^1 show that $f(R^3) \subseteq f(R^2)$. We derive therefore from pairwise-ness that $f(R') = f(R^3) \subseteq f(R^2) \subseteq f(R^1) \subseteq f(R) \subsetneq f(R')$, which is a contradiction because the last subset relation is by assumption strict. Hence, our initial assumption was incorrect and f is indeed majoritarian.

Step 3: If f is a majoritarian dominant set rule, it is robust.

As a last step, we show that f is robust if it is a majoritarian dominant set rule. Thus, assume for contradiction that f is a majoritarian dominant set rule that fails robustness. This means that there are two preference profiles R and R' such that $f(R)$ is dominant in $\succ_{R'}$, but $f(R') \not\subseteq f(R)$. As a consequence, there is an alternative $y \in f(R') \setminus f(R)$. Moreover, since $f(R)$ is dominant in $\succ_{R'}$, it follows that $f(R) \succ_{R'} y$, and hence $f(R) \subsetneq f(R')$ as f is a dominant set rule. We derive a contradiction to this

assumption by constructing two preference profiles R^2 and R^3 such that $f(R^2) = f(R)$, $f(R^3) = f(R')$, and $\succsim_{R^2} = \succsim_{R^3}$. These observations are conflicting since $\succsim_{R^2} = \succsim_{R^3}$ requires that $f(R^2) = f(R^3)$ because of majoritarianism, but $f(R) \neq f(R')$. Note that we assume in the sequel that both R and R' are defined by an even number of voters as we want to introduce majority ties. This is without loss of generality as f is homogeneous.

First, we explain how to derive R^2 from R . As a first step, we reorder the alternatives in $A \setminus f(R)$ to derive a profile R^1 with $\succsim_{R^1}|_{A \setminus f(R)} = \succsim_{R'}|_{A \setminus f(R)}$. As a consequence of IUA, it follows that $f(R^1) = f(R)$ since this step does not affect chosen alternatives. Next, let D_{i^*} denote the dominant set in R^1 that is currently chosen, that is, $f(R^1) = D_{i^*}$. We derive the profile R^2 by repeating the following procedure with R^1 as starting profile: in the current preference profile \bar{R} , we choose a pair of alternatives $x, y \in D_{i^*}$ such that $y \succ_{\bar{R}} x$ and reinforce x against y until we arrive at a profile \bar{R}' with $x \sim_{\bar{R}'} y$. It follows from WMON and majoritarianism that $x \in f(\bar{R}')$ if $x, y \in f(\bar{R})$. Moreover, as f is a dominant set rule, $x \in f(\bar{R}')$ implies $y \in f(\bar{R}')$ because $y \succ_{\bar{R}'} x$. Hence, we infer from WLOC that $f(\bar{R}) = f(\bar{R}')$ if $x, y \in f(\bar{R})$. Since $f(R^1) = D_{i^*}$, we can thus repeat this process until we arrive at a profile R^2 with $x \sim_{R^2} y$ for all $x, y \in D_{i^*}$, and it follows from the previous argument that $f(R^1) = f(R^2)$. Moreover, the majority relation of R^2 is completely specified: we have $f(R) \succ_{R^2} A \setminus f(R)$, $x \sim_{R^2} y$ for all $x, y \in f(R)$, and $\succsim_{R^2}|_{A \setminus f(R)} = \succsim_{R'}|_{A \setminus f(R)}$.

Finally, we apply the same construction as for R^2 to derive the profile R^3 from R' . In more detail, observe that, by assumption, $D_{i^*} = f(R)$ is dominant in $\succsim_{R'}$ and $D_{i^*} \subsetneq f(R')$. Hence, we can use the same construction as for R^2 to introduce majority ties between all alternatives in D_{i^*} in $\succsim_{R'}$. The same reasoning as in the previous paragraph shows that this step does not change the choice set, and it hence holds for the resulting profile R^3 that $f(R') = f(R^3)$. In particular, R^3 has now the same majority relation as R^2 , which is in conflict with majoritarianism since $\succsim_{R^2} = \succsim_{R^3}$ but $f(R^2) = f(R) \neq f(R') = f(R^3)$. This is a contradiction to our assumptions and f is therefore robust if it is a majoritarian dominant set rule. \square

Lemma 6 presents a simple criterion for deciding when a strategyproof, homogeneous, pairwise, and strongly Condorcet-consistent SCC f is a robust dominant set rule, namely when $TC(R) \subseteq f(R)$ for all profiles R . Our next goal is to prove that every such SCC meets this condition without further assumptions. Hence, suppose for contradiction that this is not the case, that is, there are an SCC f that satisfies all our axioms and a profile R such that $TC(R) \not\subseteq f(R)$. If such a profile R exists, we may as well focus on the profile R^f that minimizes the size of the top cycle among all profiles R with $TC(R) \not\subseteq f(R)$. Furthermore, for every SCC f , we define $k_f \in \{1, \dots, m+1\}$ as the maximal value such that $TC(R) \subseteq f(R)$ for all preference profiles R with $|TC(R)| < k_f$. Note that $k_f = |TC(R^f)|$ if $TC(R)$ is not always a subset of $f(R)$, and $k_f = m+1$ otherwise.

As the next step, we show that $k_f \geq 4$ for all pairwise SCCs f that satisfy strategyproofness, homogeneity, and strong Condorcet consistency. In general, this means that such SCCs can only fail to choose a superset of the top cycle if $TC(R)$ is sufficiently large. For the special case where $m \leq 3$, Lemmas 6 and 7 already imply that f needs to be a robust dominant set rule because the size of the top cycle is bounded by the number of alternatives.

LEMMA 7. *Let f denote a pairwise SCC that satisfies strong Condorcet consistency, strategyproofness, and homogeneity. Then $k_f \geq 4$.*

PROOF. Let f denote a pairwise SCC that satisfies strong Condorcet consistency, homogeneity, and strategyproofness. Furthermore, suppose for contradiction that there is a profile R^* such that $k = |TC(R^*)| \leq 3$ but $TC(R^*) \not\subseteq f(R^*)$. We proceed with a case distinction with respect to $|TC(R^*)|$ to derive a contradiction for the three possible cases.

Case 1: $|TC(R^*)| = 1$

If $|TC(R^*)| = 1$, there has to be a Condorcet winner in R^* since TC is strongly Condorcet consistent. Consequently, the strong Condorcet consistency of f requires that $f(R^*) = TC(R^*)$, which contradicts the assumption that $TC(R^*) \not\subseteq f(R^*)$.

Case 2: $|TC(R^*)| = 2$

The top cycle only elects two alternatives $x, y \in A$ if $x \sim_{R^*} y$ and $\{x, y\} \succ_{R^*} A \setminus \{x, y\}$. Hence, strong Condorcet consistency requires that f chooses at least two alternatives. In turn, COS implies then that both x and y are chosen because x is the only alternative that dominates y and y is the only alternative that dominates x . This shows that $TC(R^*) \subseteq f(R^*)$ and we again have a contradiction.

Case 3: $|TC(R^*)| = 3$

Next, assume there are three alternatives a, b , and c such that $TC(R^*) = \{a, b, c\}$, but $\{a, b, c\} \not\subseteq f(R^*)$. First, note that strong Condorcet consistency requires that $|f(R^*)| \geq 2$ because there is no Condorcet winner in R^* ; otherwise, it would hold that $|TC(R^*)| = 1$ as TC uniquely chooses the Condorcet winner whenever it exists. Next, observe that, according to Lemma 2, there has to be a cycle C that connects a, b, c since $TC(R^*) = \{a, b, c\}$. Without loss of generality, suppose that $C = (a, b, c)$, that is, $a \succ_{R^*} b$, $b \succ_{R^*} c$, $c \succ_{R^*} a$. Moreover, we also suppose without loss of generality that $a \notin f(R^*)$ because $TC(R^*) \not\subseteq f(R^*)$. COS then implies that $b \sim_{R^*} c$ and $\{b, c\} \subseteq f(R^*)$, because otherwise no chosen alternative dominates b or c . Next, consider the profile R^1 derived from R^* by reinforcing b against c , so we have $b \succ_{R^1} c$ instead of $b \sim_{R^*} c$. First, note that there is no Condorcet winner in R^1 , and thus strong Condorcet consistency requires us to choose at least two alternatives. Hence, COS implies that $a \in f(R^1)$ because it is the only alternative that dominates b in R^1 . Moreover, WMON shows that $b \in f(R^1)$ since we swap two chosen alternatives to derive R^1 from R^* . Finally, the contraposition of WLOC implies that $c \notin f(R^1)$ since $a \in f(R^1) \setminus f(R^*)$. These observations entail that $a \sim_{R^1} b$, because otherwise no chosen alternative dominates a , which violates COS. Thus, we can repeat the previous steps by reinforcing a against b , which results in a profile R^2 such that $a \succ_{R^2} b$, $b \succ_{R^2} c$, $\{a, c\} \subseteq f(R^2)$, and $b \notin f(R^2)$. Hence, COS implies once again that $c \sim_{R^2} a$ and we can again break this majority tie to derive the profile R^3 . In more detail, all edges of C are now strict, and $a \notin f(R^3)$. This contradicts COS because $b \succ_{R^3} A \setminus \{a, b\}$. Hence, no chosen alternative dominates b and the assumption that $TC(R^*) \not\subseteq f(R^*)$ was incorrect. \square

Due to Lemma 7, it follows that every pairwise SCC that satisfies strategyproofness, homogeneity, and strong Condorcet consistency can only fail to choose a superset of the top cycle if $|TC(R)| \geq 4$. Hence, we subsequently investigate profiles with a top cycle that

contains at least 4 alternatives. For this, we first need to discuss some auxiliary lemmas and start by showing that a pairwise SCC f that satisfies homogeneity, strategyproofness, and strong Condorcet consistency must choose almost all alternatives of the top cycle for profiles R with $|TC(R)| = k_f \geq 4$.

LEMMA 8. *Let f denote a pairwise SCC that satisfies strong Condorcet consistency, homogeneity, strategyproofness, and $4 \leq k_f \leq m$. It holds that $|TC(R) \cap f(R)| \geq k_f - 1$ for all profiles R with $|TC(R)| = k_f$.*

PROOF. Let f denote a pairwise SCC that satisfies all axioms of the lemma and assume that $k_f \in \{4, \dots, m\}$. Furthermore, suppose for contradiction that there is a profile R such that $|TC(R)| = k_f$ and $|f(R) \cap TC(R)| \leq k_f - 2$. This means that at least two alternatives of the top cycle are not chosen, that is, the set $X = TC(R) \setminus f(R)$ contains at least two alternatives. Next, let $x \in X$ denote one of these alternatives. We proceed with a case distinction with respect to the connected set A_x and first suppose that $X \not\subseteq A_x \cup \{x\}$. Because of the definition of connected sets, this assumption means that $X \cap TC(R^{-x}) \neq \emptyset$, where $R^{-x} = R|_{A \setminus \{x\}}$ denotes the profile derived from R by removing x . We use this fact to derive a contradiction as follows: starting at R , we repeatedly weaken x against all alternatives until we derive a profile R' in which x is the Condorcet loser. WSMON entails for every step that the choice set does not change, which means that $f(R) = f(R')$. Moreover, $TC(R') = TC(R) \setminus (A_x \cup \{x\})$ because for the top cycle it is irrelevant whether an alternative is a Condorcet loser or not present at all. However, this means that $TC(R') \not\subseteq f(R')$ because $X \cap TC(R') \neq \emptyset$, but $X \cap f(R') = X \cap f(R) = \emptyset$. Since $|TC(R')| \leq |TC(R) \setminus \{x\}| < k_f$, this contradicts the definition of k_f , which requires that $TC(\bar{R}) \subseteq f(\bar{R})$ for all profiles \bar{R} with $|TC(\bar{R})| < k_f$.

As second case, suppose that $X \subseteq A_x \cup \{x\}$. In this case, consider a second alternative $y \in X \setminus \{x\}$, which means that $y \in A_x$. We want to use Lemma 3. Note that there is an alternative $z \in f(R)$ with $z \succ_R x$ because of COS and strong Condorcet consistency. Hence, x does not dominate all other alternatives but y and Lemma 3 consequently shows that $A_y \subseteq A_x$. In particular, this means that $x \notin A_y \cup \{y\}$ and, therefore, $X \not\subseteq A_y \cup \{y\}$. Hence, we can use the same argument as in the last case to derive a contradiction by focusing on y . Since both cases result in a contradiction, it follows that the assumption $|TC(R) \cap f(R)| \leq k_f - 2$ was incorrect, that is, $|TC(R) \cap f(R)| \geq k_f - 1$ holds for all profiles R with $|TC(R)| = k_f$. \square

Lemma 8 is important because it implies for all profiles R with $|TC(R) \cap f(R)| < |TC(R)| = k_f$ that there is a single alternative of the top cycle which is unchosen. As we demonstrate next, this insight can be used to strengthen the axioms in Appendix A.3 when we restrict attention to profiles R, R' with $|TC(R)| = |TC(R')| = k_f$. In particular, the next lemma is concerned with what happens when we weaken an alternative $y \in f(R) \cap TC(R)$ against multiple alternatives $X \subseteq f(R) \cap TC(R)$ when $|TC(R)| = k_f$.

LEMMA 9. *Let f denote a pairwise SCC that satisfies strong Condorcet consistency, homogeneity, strategyproofness, and $4 \leq k_f \leq m$, and consider two preference profiles R, R' such*

that $TC(R) = TC(R')$ and $|TC(R)| = k_f$. If there are a set of alternatives $X \subseteq f(R) \cap TC(R)$ and an alternative $y \in (f(R) \cap TC(R)) \setminus X$ such that $g_{R'}(x, y) = 2 + g_R(x, y)$ for all $x \in X$, and $g_{R'}(x', y') = g_R(x', y')$ for all other pairs of alternatives, it holds that $f(R) = f(R')$ or $f(R') \cap TC(R') = TC(R') \setminus \{y\}$.

PROOF. Consider a pairwise SCC f that satisfies homogeneity, strategyproofness, and strong Condorcet consistency and let R , R' , X , and y be defined as in the lemma. In particular, it holds that $g_{R'}(x, y) = g_R(x, y) + 2$ for all $x \in X$, and $g_{R'}(x', y') = g_R(x', y')$ for all other pairs of alternatives. This means that we can transform R into a profile R^* with the same majority margins as R' by reinforcing all alternatives in X against y . Consequently, the lemma follows if we show that $X \subseteq f(R')$: if also $y \in f(R')$, then WLOC entails that $f(R) = f(R')$, and if $y \notin f(R')$, then Lemma 8 implies that $f(R') \cap TC(R') = TC(R') \setminus \{y\}$ because $|TC(R')| = k_f \geq 4$ and $y \in TC(R') \setminus f(R')$.

Thus, suppose for contradiction that $X \not\subseteq f(R')$. Then Lemma 8 shows that there is an alternative $z \in X$ such that $f(R') \cap TC(R') = TC(R') \setminus \{z\}$ because $X \subseteq TC(R')$. Moreover, we assume that $X \cup \{y\} \subseteq f(R)$. We can turn these observations into a manipulation of f by adding two voters i^* and j^* with inverse preferences to R . In more detail, the profile R^1 consists of R and the voters i^* and j^* whose preference relations are specified subsequently. In the definitions of these preference relations, we use $\tilde{f}(R) = f(R) \setminus (X \cup \{y\})$ and $\tilde{X} = X \setminus \{z\}$:

$$\begin{aligned} \succ_{i^*}^1 &= \mathbf{lex}(A \setminus f(R)), \mathbf{lex}(\tilde{f}(R) \cap f(R')), y, \mathbf{lex}(\tilde{X}), z, \mathbf{lex}(\tilde{f}(R) \setminus f(R')) \\ \succ_{j^*}^1 &= \mathbf{lex}(\tilde{f}(R) \setminus f(R'))^{-1}, z, \mathbf{lex}(\tilde{X})^{-1}, y, \mathbf{lex}(\tilde{f}(R) \cap f(R'))^{-1}, \mathbf{lex}(A \setminus f(R))^{-1} \end{aligned}$$

Since the preferences of these voters are inverse, it follows from pairwiseness that $f(R^1) = f(R)$. Next, we derive R^2 from R^1 by letting voter i^* reinforce all alternatives in X against y . Since we derive R' from R by the same modification, pairwiseness shows that $f(R^2) = f(R')$. However, this means that voter i^* can manipulate by deviating from R^1 to R^2 . Note for this that $f(R') \setminus f(R) \succ_{i^*} f(R)$ because $A \setminus f(R) \succ_{i^*} f(R)$. Furthermore, it holds that $f(R') \succ_{i^*} f(R) \setminus f(R')$ because the alternatives in $f(R) \setminus f(R') = (\tilde{f}(R) \setminus f(R')) \cup \{z\}$ are bottom-ranked by voter i^* . Finally, since $z \in f(R) \setminus f(R')$, this is indeed a manipulation for voter i^* . Hence, the assumption that $X \not\subseteq f(R')$ was incorrect, which proves the lemma. \square

Lemma 9 significantly strengthens WMON for profiles R , R' with $TC(R) = TC(R')$ and $|TC(R)| = k_f \geq 4$ and alternatives in $TC(R)$. In particular, we can now reinforce sets of alternatives against single alternatives and there are only two possible outcomes under the given assumptions. Therefore, we ensure in the following that the premises of Lemma 9 are always true: in all subsequent profiles R , it holds that $|TC(R)| = k_f$, we only modify the preferences between alternatives in the top cycle, and the top cycle will never change. As the next step, we derive a profile R for which all majority margins are known, $|TC(R)| = k_f$, and $TC(R) \not\subseteq f(R)$.

LEMMA 10. *Let f denote a pairwise SCC that satisfies strong Condorcet consistency, homogeneity, and strategyproofness. If $4 \leq k_f \leq m$, there is a profile R such that $TC(R) \not\subseteq f(R)$, $|TC(R)| = k_f$, and there is a cycle $C = (a_1, \dots, a_{k_f})$ in $\succsim_R|_{TC(R)}$ with $g_R(a_{k_f}, a_1) = 2$ and $g_R(a_i, a_j) = 2$ for all other indices $i, j \in \{1, \dots, k_f\}$ with $i < j$.*

PROOF. Let f denote a pairwise SCC that satisfies strong Condorcet consistency, homogeneity, and strategyproofness, and suppose that $k_f \in \{4, \dots, m\}$. Moreover, consider a profile R^* such that $TC(R^*) \not\subseteq f(R^*)$ and $k_f = |TC(R^*)|$; such a profile exists by the definition of k_f . Additionally, we assume in the sequel that R^* is defined by an even number of voters. This is possible as we can simply duplicate the profile R^* if it is defined by an odd number of voters. This step does neither affect the top cycle nor f since both SCCs are homogeneous, and we can thus work with this larger profile if R^* was defined by an odd number.

We prove this lemma in two steps: first, we construct a profile \hat{R}^1 such that $TC(R^*) = TC(\hat{R}^1) \not\subseteq f(\hat{R}^1)$ and there is a pair of alternatives $a, b \in TC(\hat{R}^1)$ with $b \succ_{\hat{R}^1} A \setminus \{a, b\}$. This profile is essential since COS shows now that a must be chosen, even after various manipulations. Based on this insight, we construct as the second step a profile \hat{R}^2 that satisfies all requirements of our lemma.

Step 1: Constructing the profile \hat{R}^1 .

Our first goal is to construct a profile \hat{R}^1 such that $TC(R^*) = TC(\hat{R}^1) \not\subseteq f(\hat{R}^1)$ and there is a pair of alternatives $a, b \in TC(\hat{R}^1)$ with $b \succ_{\hat{R}^1} A \setminus \{a, b\}$. For this, consider a cycle $C = (a_1, \dots, a_{k_f})$ in \succsim_{R^*} that contains all alternatives in $TC(R^*)$; such a cycle exists because of Lemma 2. Furthermore, let $b = a_{i+1}$ denote an arbitrary alternative in $TC(R^*) \cap f(R^*)$ and let $a = a_i$ denote its predecessor on the cycle C . Our goal is to reinforce b against all alternatives $A \setminus \{a, b\}$ such that $b \succ_R A \setminus \{a, b\}$.

The first key insight for this is that strong Condorcet consistency and COS entail that there is always a chosen alternative c that dominates b if there is no Condorcet winner. Based on this observation, we repeat the following steps starting at profile R^* : in the current profile R' , we identify an alternative $c \in f(R') \setminus \{a, b\}$ with $c \succ_{R'} b$ and reinforce b against c . First, note that during all of these steps, b remains chosen because of WMON and the fact that we only swap chosen alternatives. Next, observe that these steps do not affect the cycle C because c is not the predecessor of b . Hence, Lemma 2 implies that the top cycle does not change and that $c \in TC(R^*) = TC(R')$ because $c \succ_{R'} b$. The latter observation and Lemma 9 also entail that not all alternatives in the top cycle are chosen after reinforcing b against c because either c is now unchosen or the choice set is not allowed to change at all. Thus, this process terminates at a profile R^1 such that $b \succ_{R^1} f(R^1) \setminus \{a, b\}$, $TC(R^1) = TC(R^*) \not\subseteq f(R^1)$, and $\{a, b\} \subseteq f(R^1)$. The last point is true since WMON shows that $b \in f(R^1)$ and COS requires that $a \in f(R^1)$ because $a \succ_{R^1} b$ and $b \succ_{R^1} f(R^1) \setminus \{a, b\}$. We are done after this step if $b \succ_{R^1} A \setminus \{a, b\}$ but this is not guaranteed.

Hence, assume that there are alternatives $x \in A \setminus f(R^1)$ with $x \succ_{R^1} b$. Note that this assumption implies that $x \in TC(R^1)$ because $b \in TC(R^*) = TC(R^1)$. We want to repeatedly identify such an alternative $x \in TC(R^1) \setminus f(R^1)$ with $x \succ_{R^1} b$ and reinforce b against x . WMON and WLOC imply for each of these steps that either the choice set does not change, or b becomes unchosen and x chosen. In particular, this means that

after such a step, not all alternatives of the top cycle are chosen because the top cycle is not affected by these changes. However, we cannot guarantee that b remains chosen during these steps and, therefore, we need to treat the case that we arrive at a profile R^2 with $b \notin f(R^2)$ separately. Given such a profile R^2 , we show how we can find another profile R^3 such that $b \in f(R^3)$, $TC(R^*) = TC(R^3) \not\subseteq f(R^3)$, and $g_{R^3}(b, x) = g_{R^2}(b, x)$ for all $x \in A$. For this, note that the cycle $C = (a_1, \dots, a_{k_f})$ exists also in R^2 , and thus $TC(R^2) = TC(R^*)$. Since $|TC(R^*)| = k_f$, Lemma 8 shows that $TC(R^2) \setminus \{b\} \subseteq f(R^2)$ if $b \notin f(R^2)$. In particular, this means that $c = a_{i+2}$ (i.e., the successor of b on C) is in $f(R^2)$. We apply next a similar idea as in the construction of R^1 : at every preference profile R' , we reinforce c against a chosen alternative $x \in f(R') \setminus \{b, c\}$ with $x \succ_{R'} c$. Just as in the first step, Lemma 9 implies that c remains chosen during these steps and that we never choose all alternatives of $TC(R')$. Also, we do not flip any edge in the cycle C during this process because we never reinforce c against its predecessor b . Hence, neither the top cycle nor a majority margin involving b change. Finally, this process terminates at a profile R^3 such that $c \succ_{R^3} f(R^3) \setminus \{b, c\}$ and $c \in f(R^3)$. Moreover, COS now requires that $b \in f(R^3)$ because $c \succ_{R^3} A \setminus \{b, c\}$, that is, if $b \notin f(R^3)$, no chosen alternative dominates c . Hence, profile R^3 indeed satisfies all our requirements.

Thus, if b drops out of the choice set after reinforcing it against an unchosen alternative, we can apply this construction to derive a profile R^3 with $b \in f(R^3)$ and $TC(R^*) = TC(R^3) \not\subseteq f(R^3)$. At this point, we can simply repeat the same constructions used in the derivation of R^1 and R^2 , and eventually, we will arrive at a profile \hat{R}^1 such that $b \succ_{\hat{R}^1} A \setminus \{a, b\}$ because the majority margins of b are non-decreasing during all steps and strictly decreasing during the constructions of R^1 and R^2 . Also, none of the constructions requires us to invert edges of the cycle C , and thus $TC(\hat{R}^1) = TC(R^*)$, whereas Lemma 9 shows that $TC(\hat{R}^1) \not\subseteq f(\hat{R}^1)$.

Step 2: Constructing the profile \hat{R}^2 .

As a second step, we construct the profile \hat{R}^2 that satisfies all requirements of the lemma. In more detail, \hat{R}^2 has to satisfy that $TC(R^*) = TC(\hat{R}^2) \not\subseteq f(\hat{R}^2)$ and that there is a cycle $C = (a_1, \dots, a_{k_f})$ in $\succ_{\hat{R}^2}$ that connects all alternatives in $TC(\hat{R}^2)$ such that $g_{\hat{R}^2}(a_{k_f}, a_i) = 2$ and $g_{\hat{R}^2}(a_i, a_j) = 2$ for all other indices $i, j \in \{1, \dots, k_f\}$ with $i < j$. For the construction of this profile, let R^1 denote the profile constructed in the last step, and let $C = (a_1, \dots, a_{k_f})$ denote a cycle that connects all alternatives $x \in TC(R^*) = TC(R^1)$ in \succ_{R^1} . By construction, there are alternatives $a, b \in TC(R^1)$ such that $b \succ_{R^1} A \setminus \{a, b\}$. In the sequel, we assume without loss of generality that $b = a_1$ since we can pick the starting point of the cycle. This means that $a = a_{k_f}$, that is, a is the predecessor of b on the cycle C , because a is the only alternative that dominates b . Finally, recall that R^* and, therefore, also R^1 are defined by an even number of voters, which implies that the majority margins are even.

The central observation for the construction of \hat{R}^2 is that COS and strong Condorcet consistency guarantee that $a \in f(R^1)$ because a is the only alternative that dominates b . Even more, this is true as long as $b \succ_R A \setminus \{a, b\}$ and $a \succ_R b$. We use this observation to reinforce $a = a_{k_f}$ against the alternatives $x \in TC(R^1) \setminus \{a_1, a_{k_f-1}, a_{k_f}\}$: in each step, we identify an alternative $x \in TC(R^1) \setminus \{a_1, a_{k_f-1}, a_{k_f}\}$ with $x \succ_{R'} a$ in the current profile R' and reinforce a against x . As mentioned before, COS implies that a has to be chosen

during all steps. Moreover, if $x \notin f(R')$, it follows from WMON that x remains unchosen after this step; otherwise, we could revert the modification and WMON entails that $x \in f(R')$, contradicting our previous assumption. Hence, WLOC implies that the choice set cannot change in this case. Conversely, if $x \in f(R')$, it follows either that x is no longer chosen after this step, or that the choice set is not allowed to change because of WLOC. In particular, this shows that not all alternatives in $TC(R^*) = \{a_1, \dots, a_{k_f}\}$ are chosen after this step. Hence, we can repeat these steps until we arrive at a profile R^2 such that $a \succ_{R^2} A \setminus \{a, a_1, a_{k_f-1}\}$. Also note that $g_{R^2}(a, x) = 2$ for all $x \in TC(R^1) \setminus \{a, a_1, a_{k_f-1}\}$ with $x \succ_{R^1} a$ because we only reinforce a against such alternatives x until a strictly dominates them. Finally, none of these steps involves an edge of the cycle C , which implies that $TC(R^2) = TC(R^1) = TC(R^*)$. Hence, $TC(R^2) = TC(R^*) \not\subseteq f(R^2)$.

As the next step, we reinforce a against b if $a \sim_{R^2} b$, and against its predecessor a_{k_f-1} on the cycle C until $g_{R^3}(a, a_{k_f-1}) = -2$ if $g_{R^2}(a, a_{k_f-1}) \leq -4$. This results in a new profile R^3 and, by the same arguments as before, it follows that not all alternatives in $TC(R^*)$ are chosen. Also, it is easy to see that the top cycle did not change since $a \succ_{R^3} b$ and $a_{k_f-1} \succ_{R^3} a$. As a last point, observe that all new outgoing edges $a \succ_{R^2} x$ have weight 2 and that the incoming edge from a_{k_f-1} has a weight of at most 2.

Finally, note that a dominates each alternative $x \in A \setminus \{a, a_{k_f-1}\}$ in R^3 , and thus COS and strong Condorcet consistency imply now that a_{k_f-1} must be chosen. Hence, we can repeat the previous steps for a_{k_f-1} , or more generally, we can traverse along the cycle C using these steps. Thus, we repeat this process until we applied our constructions to a_1 . It follows from the construction that each edge in the final profile \hat{R}^2 has weight 2. Furthermore, the cycle C also exists in the final profile \hat{R}^2 and thus, $TC(\hat{R}^2) = TC(R^*)$. Moreover, it is a consequence of COS, WMON, and WLOC that $TC(\hat{R}^2) \not\subseteq f(\hat{R}^2)$. Finally, it follows for the profile \hat{R}^2 that a_1 dominates each alternative but a_{k_f} , and each alternative a_i with $1 < i < k_f$ dominates all alternatives a_j with $j > i$. This claim follows by inspecting our construction in more detail: if $j = i + 1$, that is, if a_j is the successor of a_i in C , this follows immediately as we do not break the cycle. If $j > i + 1$, we first apply our construction to a_j ensuring that a_j dominates a_i . Later, we apply our construction to a_i , which reverts this edge and ensures that it has a weight of 2. Since this majority margin will not be modified anymore, this proves that the profile \hat{R}^2 indeed satisfies all criteria of the lemma. \square

If we consider a pairwise SCC f that satisfies all required axioms but is no robust dominant set rule, Lemma 10 states the exact majority margins of a profile R^* such that $TC(R^*) \not\subseteq f(R^*)$ and $|TC(R^*)| = k_f \leq m$. As a last step, we derive a contradiction to this by showing that $TC(R^*) \subseteq f(R^*)$ is required. By considering the contraposition of Lemma 10, we infer from this that $k_f \notin \{4, \dots, m\}$. Together with Lemma 7, this means that $k_f = m + 1$, which shows that every pairwise SCC that satisfies strong Condorcet consistency, homogeneity, and strategyproofness is a robust dominant set rule.

LEMMA 11. *Every pairwise SCC that satisfies strong Condorcet consistency, homogeneity, and strategyproofness is a robust dominant set rule.*

PROOF. Assume for contradiction that there is a pairwise SCC f that satisfies strategyproofness, homogeneity, and strong Condorcet consistency, but is no robust dominant set rule. The contraposition of Lemma 6 shows that there is a profile R such that $TC(R) \not\subseteq f(R)$. On the other hand, Lemma 7 shows that $TC(R) \subseteq f(R)$ for all profiles R with $|TC(R)| \leq 3$. These claims contradict each other if $m \leq 3$, and thus we focus on the case that $m \geq 4$. Hence, let $k_f \in \{4, \dots, m\}$ denote the maximal value such that $TC(R) \subseteq f(R)$ for all profiles R with $|TC(R)| < k_f$. Moreover, let \bar{R} denote a profile such that $TC(\bar{R}) \not\subseteq f(\bar{R})$ and $|TC(\bar{R})| = k_f$; such a profile exists because of the definition of k_f . Next, we apply Lemma 10 to derive a profile R^* such that $TC(R^*) = TC(\bar{R})$, $TC(R^*) \not\subseteq f(R^*)$, and the alternatives $TC(R^*) = \{a_1, \dots, a_{k_f}\}$ can be ordered such that $g_{R^*}(a_{k_f}, a_1) = 2$ and $g_{R^*}(a_i, a_j) = 2$ for all other indices $i, j \in \{1, \dots, k_f\}$ with $i < j$. Furthermore, $TC(R^*) \setminus f(R^*)$ contains a single alternative a_j because of Lemma 8.

For deriving a contradiction to this assumption, we will consider a number of profiles related to R^* . In particular, our subsequent construction will mimic neutrality since f needs not be neutral. Thus, we define profile R^π given some permutation $\pi : TC(R^*) \rightarrow TC(R^*)$ as follows: $g_{R^\pi}(x, y) = g_{R^*}(x, y)$ if $x \in A \setminus TC(R^*)$ or $y \in A \setminus TC(R^*)$ and $g_{R^\pi}(\pi(a_i), \pi(a_j)) = g_{R^*}(a_i, a_j)$ for all $a_i, a_j \in TC(R^*)$. Less formally, R^π is constructed as follows: we derive the majority margins of R^π by reordering the edges between alternatives $x, y \in TC(R^*)$ according to π but we do not reorder the edges to alternatives outside of the top cycle. For a better readability, we refer to $\pi(a_i)$ as a_i^π from now on. In particular, the construction of R^π implies that $g_{R^\pi}(a_{k_f}^\pi, a_1^\pi) = 2$ and $g_{R^\pi}(a_i^\pi, a_j^\pi) = 2$ for all $i, j \in \{1, \dots, k_f\}$ with $i < j$. Furthermore, we define Π_π^l as the set of permutations π' with $\pi(a_i) = \pi'(a_i)$ for $i \in \{1, \dots, k_f\}$, that is, all permutations $\pi' \in \Pi_\pi^l$ agree with π on the alternatives $\{a_1, \dots, a_{k_f}\}$.

Based on the profiles R^π , we next prove the lemma. For this, let j^* denote the smallest index such that $a_{j^*}^\pi \notin f(R^\pi)$ for some permutation π on $TC(R^*)$. Moreover, let π^* denote the corresponding permutation, that is, $a_{j^*}^{\pi^*} \notin f(R^{\pi^*})$. Given the value j^* and the profile R^{π^*} , we prove the lemma in three steps. First, we show that $\{a_1^\pi, a_2^\pi, a_{k_f}^\pi\} \subseteq f(R^\pi)$ for all permutations π . This means in particular that $j^* \in \{3, \dots, k_f - 1\}$. Next, we show that $f(R^\pi) = f(R^{\pi^*})$ for all permutations $\pi \in \Pi_{\pi^*}^{j^*}$. This observation implies that $a_{j^*}^\pi = a_{j^*}^{\pi^*} \notin f(R^\pi)$ for all these permutations. Finally, we use this insight to derive a contradiction. All profiles used for Steps 2 and 3 are depicted exemplarily in Figure 3 for the case that $k_f = 6$.

Step 1: $\{a_1^\pi, a_2^\pi, a_{k_f}^\pi\} \subseteq f(R^\pi)$ for all permutations $\pi : TC(R^*) \rightarrow TC(R^*)$

Consider an arbitrary preference profile R^π . First, note that there is no Condorcet winner in this profile since there is no Condorcet winner in R^* , and thus strong Condorcet consistency requires that $|f(R^\pi)| \geq 2$. As a consequence, COS requires that $a_1^\pi \in f(R^\pi)$ and $a_{k_f}^\pi \in f(R^\pi)$ because a_1^π is the only alternative that dominates a_2^π and $a_{k_f}^\pi$ is the only alternative that dominates a_1^π . As a last point, suppose for contradiction that $a_2^\pi \notin f(R^\pi)$. Since $|TC(R^\pi)| = k_f$, it follows from Lemma 8 that $TC(R^\pi) \setminus \{a_2^\pi\} \subseteq f(R^\pi)$,

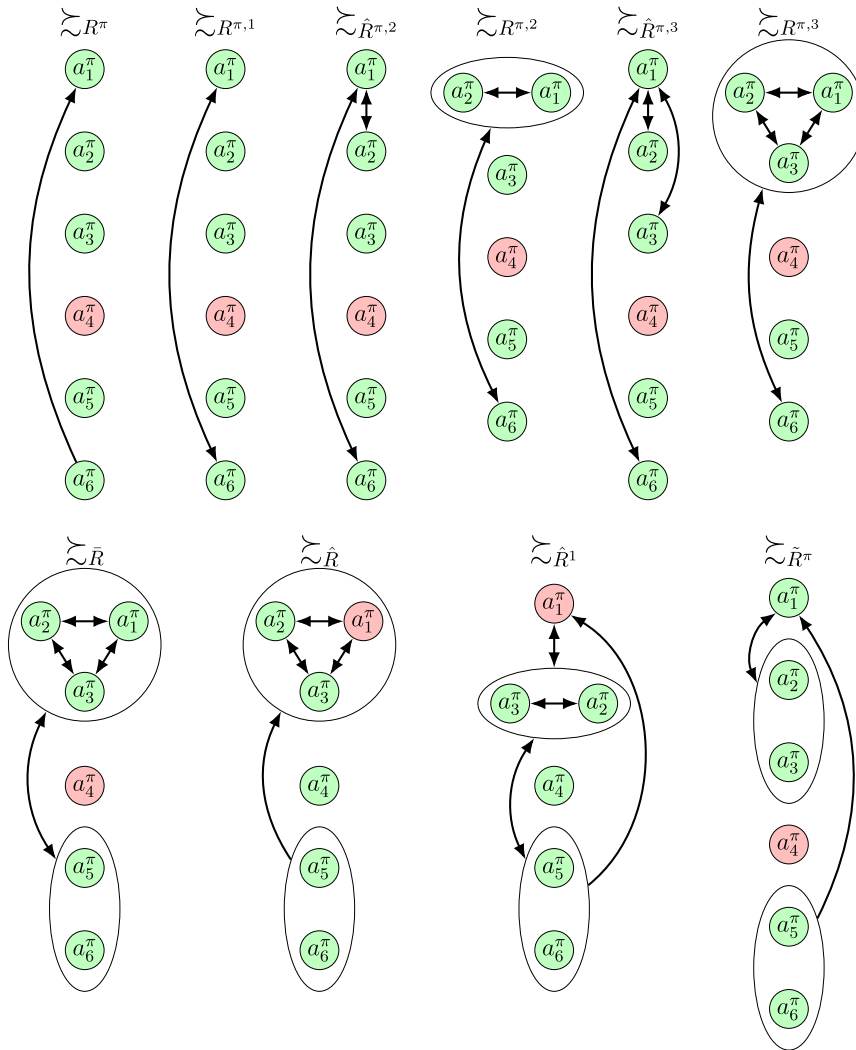


FIGURE 3. The (weighted) majority relations used in the proof of Lemma 11 for $k_f = 6$. Alternatives outside of the top cycle are not depicted, and we assume that $j^* = 4$ and $a_1^\pi \notin f(\hat{R})$, where π denotes a permutation in $\Pi_{\pi^*}^{j^*}$. Alternatives placed in an ellipse have identical relationships to all alternatives outside of the ellipse, and all missing edges point downwards. All directed edges indicate a majority margin of 2 and all bidirectional edges indicate a majority margin of 0. Green (light gray) alternatives are chosen and red (dark gray) ones are unchosen by f .

in particular that $a_3^\pi \in f(R^\pi)$. As the next step, we reinforce a_3^π twice against a_1^π to derive a profile R' with $g_{R'}(a_3^\pi, a_1^\pi) = 2$. Note that a_1^π needs to stay chosen during these steps because it is still the only alternative dominating a_2^π and WMON implies that a_3^π remains also chosen. Hence, it follows from WLOC that $f(R') = f(R^\pi)$, which means that $a_2^\pi \notin f(R')$. However, a_2^π is the only alternative that dominates a_3^π in R' , and thus

COS is violated. This is a contradiction, and thus the assumption that $a_2^\pi \notin f(R^\pi)$ was incorrect.

Step 2: $f(R^\pi) = f(R^{\pi^*})$ for all permutations $\pi \in \Pi_{\pi^*}^{j^*}$

For proving this step, we consider the profiles $R^{\pi,l}$ for every $l \in \{1, \dots, j^* - 1\}$, which differs from R^π in the fact that $g_{R^{\pi,l}}(a_i^\pi, a_j^\pi) = 0$ for all $i, j \in \{1, \dots, l, k_f\}$. Intuitively, $R^{\pi,l}$ is derived R^π by introducing a large set of tied alternatives $\{a_1^\pi, \dots, a_l^\pi, a_{k_f}^\pi\}$ in the majority relation. Our goal is to show that $f(R^\pi) = f(R^{\pi,j^*-1})$ for all permutations $\pi \in \Pi_{\pi^*}^{j^*}$. This implies that $f(R^\pi) = f(R^{\pi'})$ for all such permutation π, π' because $g_{R^{\pi,l}} = g_{R^{\pi',l}}$ for all $\pi, \pi' \in \Pi_{\pi^*}^{l+1}$ and $l \in \{1, \dots, j^* - 1\}$. For deriving this statement, we show inductively that $f(R^\pi) = f(R^{\pi,l})$ for all $\pi \in \Pi_{\pi^*}^{j^*}$ and all $l \in \{1, \dots, j^* - 1\}$.

First, we focus on the induction basis $l = 1$ and consider therefore an arbitrary permutation $\pi \in \Pi_{\pi^*}^{j^*}$. Note that $R^{\pi,1}$ only differs from R^π by the fact that $g_{R^{\pi,1}}(a_{k_f}^\pi, a_1^\pi) = 0$ instead of 2. Hence, we can derive $R^{\pi,1}$ from R^π by reinforcing a_1^π against $a_{k_f}^\pi$. Furthermore, we have shown in the last step that $\{a_1^\pi, a_{k_f}^\pi\} \subseteq f(R^\pi)$, and COS requires that both alternatives are chosen in $f(R^{\pi,1})$ because $a_{k_f}^\pi$ is still the only alternative that dominates a_1^π and a_1^π is the only alternative that dominates a_2^π . Consequently, WLOC shows that $f(R^\pi) = f(R^{\pi,1})$.

For the induction step, assume that there is a value $l \in \{1, \dots, j^* - 2\}$ such that $f(R^\pi) = f(R^{\pi,l})$ for all $\pi \in \Pi_{\pi^*}^{j^*}$. We need to prove that this claim is also true for $l + 1$. Hence, note that for every permutation $\pi \in \Pi_{\pi^*}^{j^*}$, it holds that $\{a_1^\pi, \dots, a_{j^*-1}^\pi, a_{k_f}^\pi\} \subseteq f(R^\pi)$ because of the definition of j^* and Step 1. Next, we explain how to derive $R^{\pi,l+1}$ from R^π for an arbitrary permutation $\pi \in \Pi_{\pi^*}^{j^*}$: first, we reinforce a_1^π against $a_{k_f}^\pi$ to derive the profile \bar{R}^π . It follows from the same argument as in the induction basis that $f(\bar{R}^\pi) = f(R^\pi)$. Next, we reinforce all alternatives a_i^π with $i \in \{2, \dots, l + 1\}$ against a_1^π . COS requires for the resulting profile \hat{R}^π that a_1^π is chosen because it is still the only alternative dominating a_2^π and Lemma 9 shows therefore that $f(\hat{R}^\pi) = f(R^\pi)$. Furthermore, observe that $g_{\hat{R}^\pi}(a_1^\pi, a_i^\pi) = 0$ for all $i \in \{2, \dots, l + 1, k_f\}$. Finally, we can derive $R^{\pi,l+1}$ from \hat{R}^π by letting a voter i with preference relation $\succ_i = a_2^\pi, \dots, a_{l+1}^\pi, a_{k_f}^\pi, \text{lex}(\mathcal{A} \setminus \{a_2^\pi, \dots, a_{l+1}^\pi, a_{k_f}^\pi\})$ change his preference relation to $\succ'_i = a_{k_f}^\pi, a_{l+1}^\pi, \dots, a_2^\pi, \text{lex}(\mathcal{A} \setminus \{a_2^\pi, \dots, a_{l+1}^\pi, a_{k_f}^\pi\})$. We can assume that such a voter exists since pairwiseiness allows us to add voters with inverse preferences without affecting the choice set. This step ensures that $g_{R^{\pi,l+1}}(a_i^\pi, a_j^\pi) = 0$ for all $i, j \in \{2, \dots, l + 1, k_f\}$ and it therefore transforms \hat{R}^π into $R^{\pi,l+1}$.

Now, since $\{a_2^\pi, \dots, a_{l+1}^\pi, a_{k_f}^\pi\} \subseteq f(R^\pi) = f(\hat{R}^\pi)$, WLOC implies that $f(R^\pi) = f(R^{\pi,l+1})$ if $\{a_2^\pi, \dots, a_{l+1}^\pi, a_{k_f}^\pi\} \subseteq f(R^{\pi,l+1})$. Hence, our next goal is to prove this set inclusion and we assume for contradiction that there is an alternative a_j^π with $j \in \{2, \dots, l + 1, k_f\}$ such that $a_j^\pi \notin f(R^{\pi,l+1})$. First, suppose that $a_j^\pi \in \{a_2^\pi, \dots, a_{l+1}^\pi\}$, which means that $TC(R^\pi) \setminus \{a_j^\pi\} \subseteq f(R^\pi)$ because of Lemma 8. In this case, we derive a contradiction by considering the permutation π' with $a_1^{\pi'} = a_j^\pi, a_j^{\pi'} = a_1^\pi$,

and $a_i^{\pi'} = a_i^{\pi}$ for all other $i \in \{1, \dots, k_f\}$. In more detail, we can use the same construction as for π to transform $R^{\pi'}$ into $R^{\pi', l+1} = R^{\pi, l+1}$. In particular, the analysis of the previous paragraph shows that $\{a_1^{\pi'}, \dots, a_{l+1}^{\pi'}, a_{k_f}^{\pi'}\} \subseteq f(R^{\pi'}) = f(\hat{R}^{\pi'})$ and we derive $R^{\pi', l+1} = R^{\pi, l+1}$ from the profile $\hat{R}^{\pi'}$ by a manipulation that only involves the alternatives $\{a_2^{\pi'}, \dots, a_{l+1}^{\pi'}, a_{k_f}^{\pi'}\}$. Hence, the assumption that $TC(R^{\pi}) \setminus \{a_j^{\pi}\} = TC(R^{\pi'}) \setminus \{a_1^{\pi'}\} \subseteq f(R^{\pi, l+1})$ implies that $f(\hat{R}^{\pi'}) = f(R^{\pi, l+1})$ because of WLOC. However, this contradicts that $a_1^{\pi'} = a_j^{\pi} \notin f(R^{\pi, l+1})$, which proves that $\{a_2^{\pi}, \dots, a_{l+1}^{\pi}\} \subseteq f(R^{\pi, l+1})$.

As a second case, suppose that $a_{k_f}^{\pi} \notin f(R^{\pi, l+1})$. In this case, we derive a contradiction to the induction hypothesis by deriving $R^{\pi, l+1}$ from $R^{\pi, l}$. Thus, note that these two preference profiles only differ in majority margins involving a_{l+1}^{π} : we have $g_{R^{\pi, l}}(a_{l+1}^{\pi}, a_{k_f}^{\pi}) = 2$ but $g_{R^{\pi, l+1}}(a_{l+1}^{\pi}, a_{k_f}^{\pi}) = 0$, and for all $i \in \{1, \dots, l\}$, $g_{R^{\pi, l}}(a_i^{\pi}, a_{l+1}^{\pi}) = 2$ but $g_{R^{\pi, l+1}}(a_i^{\pi}, a_{l+1}^{\pi}) = 0$. Also, observe that the induction hypothesis implies that $f(R^{\pi}) = f(R^{\pi, l})$, which means that $\{a_1^{\pi}, \dots, a_{l+1}^{\pi}, a_{k_f}^{\pi}\} \subseteq f(R^{\pi, l})$. Hence, we can transform $R^{\pi, l}$ into $R^{\pi, l+1}$ as follows: first, we reinforce a_{l+1}^{π} one after another against all alternatives a_i^{π} with $i \in \{1, \dots, l\}$. For each swap, Lemma 9 implies that the choice set does either not change at all, or all alternatives in $TC(R^{\pi}) \setminus \{a_i^{\pi}\}$ are chosen (where a_i^{π} denotes the weakened alternative). In particular, this shows that a_{l+1}^{π} and $a_{k_f}^{\pi}$ stay chosen during this process. Finally, we reinforce $a_{k_f}^{\pi}$ against a_{l+1}^{π} to derive $R^{\pi, l+1}$. Then WMON implies that $a_{k_f}^{\pi} \in f(R^{\pi, l+1})$, contradicting our assumption. Hence, it follows that $\{a_2^{\pi}, \dots, a_{l+1}^{\pi}, a_{k_f}^{\pi}\} \subseteq f(R^{\pi, l+1})$, which proves the induction step. As a consequence, we infer that $f(R^{\pi}) = f(R^{\pi, j^*-1}) = f(R^{\pi', j^*-1}) = f(R^{\pi'})$ for all permutations $\pi, \pi' \in \Pi_{\pi^*}^{j^*}$.

Step 3: Deriving the contradiction

As a last step, we derive a contradiction by showing that $a_{j^*}^{\pi^*} \in f(R^{\pi^*})$. This claim is in conflict with the definitions of j^* and R^{π^*} , which require that $a_{j^*}^{\pi^*} \notin f(R^{\pi^*})$. For proving this claim, we consider first the profile \bar{R} , which differs in the following majority margins from R^{π^*} : $g_{\bar{R}}(a_i^{\pi^*}, a_{j^*}^{\pi^*}) = 0$ for all $i, j \in \{1, \dots, j^* - 1\}$ and $g_{\bar{R}}(a_i^{\pi^*}, a_{j^*}^{\pi^*}) = 0$ for all $i \in \{1, \dots, j^* - 1\}$, $j \in \{j^* + 1, \dots, k_f\}$. A similar analysis as in Step 2 shows that $f(R^{\pi^*}) = f(\bar{R})$. In more detail, we can use an induction on the profiles $\bar{R}^{\pi, l}$ for all $l \in \{1, \dots, j^* - 1\}$ and $\pi \in \Pi_{\pi^*}^{j^*}$, which are defined by the following majority margins: $g_{\bar{R}^{\pi, l}}(a_i^{\pi}, a_j^{\pi}) = 0$ for all $i, j \in \{1, \dots, l\}$, $g_{\bar{R}^{\pi, l}}(a_i^{\pi}, a_j^{\pi}) = 0$ for all $i \in \{1, \dots, l\}$, $j \in \{j^* + 1, \dots, k_f\}$, and $g_{\bar{R}^{\pi, l}}(x, y) = g_{R^{\pi}}(x, y)$ for all remaining majority margins. More intuitively, the profiles $\bar{R}^{\pi, l}$ differ from the profiles $R^{\pi, l}$ only in the fact that all alternatives $\{a_1^{\pi}, \dots, a_l^{\pi}\}$ are in a majority tie with all alternatives in $\{a_{j^*+1}^{\pi}, \dots, a_{k_f}^{\pi}\}$ instead of just $a_{k_f}^{\pi}$. Since the last step shows that $TC(R^{\pi}) \setminus \{a_{j^*}^{\pi}\} \subseteq f(R^{\pi})$ for all permutations $\pi \in \Pi_{\pi^*}^{j^*}$, an almost identical induction as in Step 2 shows that $f(R^{\pi}) = f(\bar{R}^{\pi, l})$ for all permutations $\pi \in \Pi_{\pi^*}^{j^*}$ and $l \in \{1, \dots, j^* - 1\}$. This means in particular that $f(R^{\pi^*}) = f(\bar{R}^{\pi^*, j^*-1}) = f(\bar{R})$.

Departing from this observation, we now consider profile \hat{R} which is derived from \bar{R} by reinforcing all alternatives in $X_2 = \{a_{j^*+1}^{\pi^*}, \dots, a_{k_f}^{\pi^*}\}$ against all alternatives in $X_1 =$

$\{a_1^\pi, \dots, a_{j^*-1}^\pi\}$. This means for the majority margins that $g_{\hat{R}}(x, y) = 2$ for all $x \in X_2$, $y \in X_1$. As a consequence of this observation, $a_{j^*}^\pi$ is now the only alternative that dominates $a_{j^*+1}^\pi$, and thus COS requires that $a_{j^*}^\pi \in f(\hat{R})$. Next, note that a repeated application of Lemma 9 shows that $X_2 \subseteq f(\hat{R})$ because we can transform \bar{R} into \hat{R} by reinforcing the alternatives X_2 against each alternative $x \in X_1$ individually. For each of these steps, Lemma 9 shows that either the choice set does not change or all alternatives in $TC(R^\pi) \setminus \{x\}$ are chosen. In particular, this means that $X_2 \subseteq f(\hat{R})$ and that $TC(\hat{R}) \not\subseteq f(\hat{R})$. Hence, there is an alternative $a_j^\pi \in X_1$ such that $a_j^\pi \notin f(\hat{R})$.

As the next step, consider the profile \hat{R}^j derived from \bar{R} by reinforcing the alternatives in X_2 only against a_j^π . We show that $a_j^\pi \notin f(\hat{R}^j)$ and assume for the sake of contradiction that this is not the case. Hence, Lemma 9 implies that $f(\hat{R}^j) = f(\bar{R})$. Moreover, we can now transform \hat{R}^j into \hat{R} by reinforcing the alternatives in X_2 once against each alternative $x \in X_1 \setminus \{a_j^\pi\}$. For every step, Lemma 9 shows that a_j^π needs to stay chosen, and thus we have a contradiction to the assumption that $a_j^\pi \notin f(\hat{R})$. Hence, it must hold that $a_j^\pi \notin f(\hat{R}^j)$, which implies that $TC(\hat{R}^j) \setminus \{a_j^\pi\} \subseteq f(\hat{R}^j)$ because of Lemma 8.

Finally, we derive a contradiction to this observation. Consider for this a permutation $\pi \in \Pi_{\pi^*}^{j^*}$ such that $a_1^\pi = a_j^\pi$. We show that $a_j^\pi \in f(\hat{R}^j)$ by transforming R^π into \hat{R}^j and observe for this $X_1 \cup X_2 \subseteq f(R^\pi)$ because of Step 2. As a first step, we reinforce all alternatives in $X_2 \setminus \{a_{k_f}^\pi\}$ against a_1^π twice, and the alternatives in $X_1 \setminus \{a_1^\pi\}$ once against a_1^π . During all these steps, COS requires that a_1^π remains chosen because it is the only alternative that dominates a_2^π . In turn, Lemma 9 implies therefore that the choice set cannot change, that is, this process results in a profile \tilde{R}^π with $f(\tilde{R}^\pi) = f(R^\pi)$. Moreover, note that $g_{\tilde{R}^\pi}(a_1^\pi, x) = g_{\hat{R}^j}(a_1^\pi, x)$ for all $x \in A \setminus \{a_1^\pi\}$. For the next step, consider a voter i with the preference relation $\succ_i = a_2^\pi, \dots, a_{j^*-1}^\pi, a_{j^*+1}^\pi, \dots, a_{k_f}^\pi, \mathbf{lex}(X)$, where X contains all missing alternatives. We can assume that such a voter exists as pairwiseness allows us to add pairs of voters with inverse preferences without affecting the choice set. Next, we let voter i deviate to the preference relation $a_{j^*+1}^\pi, \dots, a_{k_f}^\pi, a_{j^*-1}^\pi, \dots, a_2^\pi, \mathbf{lex}(X)$, which transforms the profile \tilde{R}^π into \hat{R}^j . In particular, we know that all alternatives in $\{a_2^\pi, \dots, a_{j^*-1}^\pi, a_{j^*+1}^\pi, \dots, a_{k_f}^\pi\}$ are chosen both in $f(\tilde{R}^\pi)$ (because $f(R^\pi) = f(\tilde{R}^\pi)$) and in $f(\hat{R}^j)$ (because $TC(R^\pi) \setminus \{a_j^\pi\} = TC(R^\pi) \setminus \{a_1^\pi\} \subseteq f(\hat{R}^j)$). Thus, WLOC implies that $f(\tilde{R}^\pi) = f(\hat{R}^j)$, which conflicts with $a_j^\pi = a_1^\pi \notin f(\hat{R}^j)$. This contradiction proves the lemma. \square

A.5 Proofs of the main results

Finally, we are ready to prove our main results. First, we discuss the proof of Theorem 1: a pairwise, nonimposing, neutral, and homogeneous SCC is strategyproof if and only if it is a robust dominant set rule. To be able to use the results of the previous section, we

show that every SCC which satisfies these requirements is strongly Condorcet consistent.

LEMMA 12. *Every pairwise SCC that satisfies strategyproofness, nonimposition, homogeneity, and neutrality is strongly Condorcet consistent.*

PROOF. Consider a pairwise SCC f that satisfies nonimposition, homogeneity, neutrality, and strategyproofness. We need to show two claims: if there is a Condorcet winner, it is chosen uniquely by f , and if an alternative is the unique winner of f , it is the Condorcet winner. We prove these claims separately.

Claim 1: If x is the Condorcet winner in R , then $f(R) = \{x\}$.

Consider an arbitrary profile R with Condorcet winner x . The claim follows by showing that $f(R) = \{x\}$, and thus let R' denote a profile such that $f(R') = \{x\}$. Such a profile exists because f is nonimposing. Next, we repeatedly use WSMON to push down the best alternative of every voter until we arrive at a profile R^1 such that every voter top-ranks x . Since R^1 is constructed by repeated application of WSMON, it follows that $f(R^1) = f(R') = \{x\}$. As the next step, we let all voters order the alternatives in $A \setminus \{x\}$ lexicographically. This leads to the profile R^2 and IUA implies that the choice set does not change, so $f(R^2) = \{x\}$. Moreover, all voters have the same preference relation in R^2 . Thus, it follows from homogeneity that $f(R^3) = \{x\}$, where R^3 consists of a single voter who has the same preference relation as the voters in R^2 . Next, let $c = \min_{y \in A \setminus \{x\}} g_R(x, y)$ denote the smallest majority margin of x in R and note that $c \geq 1$ because x is the Condorcet winner in R . We use again homogeneity to construct a profile R^4 that consists of c copies of R^3 , which means that $f(R^4) = \{x\}$. Furthermore, observe that the parity of the number of voters used in R^4 is equal to the parity of the number of voters used in R . The reason for this is that c is odd if and only if R is defined by an odd number of voters.

As the last step, we need to set the majority margins to their values in R . For this, we repeat the following procedure on each pair of alternatives y, z with $g_{R^4}(y, z) < g_R(y, z)$ until we arrive at a profile R^5 with $g_{R^5} = g_R$. First, we add two voters i and j to the preference profile such that voter i prefers x the least and ranks z directly over y , and voter j 's preference relation is inverse to voter i 's. Observe that we can assign such a preference relation to voter i because $g_{R^4}(x, z') = c \leq g_R(x, z')$ for all $z' \in A \setminus \{x\}$ implies that $x \neq z$. Since the preference relations of these two voters are inverse, the majority margins do not change and pairwise consistency requires that x is still the unique winner. Next, we let voter i swap y and z , which increases the majority margin between y and z by 2. Moreover, the choice set cannot change during this step because x is voter i 's least preferred alternative. Hence, if another set would be chosen, this step is a manipulation for voter i , which contradicts strategyproofness. Therefore, we can repeat this process for every pair of alternatives until we derive a profile R^5 with $g_{R^5} = g_R$ and our arguments show that $f(R^5) = \{x\}$. This proves that $f(R) = \{x\}$ because of pairwise consistency, which shows that f chooses the Condorcet winner uniquely whenever it exists.

Claim 2: If $f(R) = \{x\}$, then x is the Condorcet winner in R .

Next, we focus on the opposite direction and show that if an alternative is chosen as unique winner by f , then it is the Condorcet winner. Assume for contradiction that this is not the case, which means that there is a preference profile R and an alternative x such $f(R) = \{x\}$ even though x is not the Condorcet winner in R . Then there is an alternative $y \in A \setminus \{x\}$ such that $g_R(y, x) \geq 0$. We continue with a case distinction with respect to whether $g_R(y, x) = 0$ or $g_R(y, x) > 0$. First, assume that $g_R(y, x) > 0$. In this case, we can repeatedly reinforce y against all other alternatives $z \in A \setminus \{x\}$. This process eventually results in a profile R' in which y is the Condorcet winner. Since Claim 1 proves that f elects the Condorcet winner whenever it exists, we derive that $f(R') = \{y\}$. On the other side, it follows from IUA that these steps do not change the choice set because we only swap unchosen alternatives, so $f(R') = \{x\}$. These two observations contradict each other, and thus the assumption that $g_R(y, x) > 0$ was incorrect.

Next, assume that $g_R(y, x) = 0$. In this case, we partition the voters N according to their preferences between x and y : we denote with $N_{x>y} = \{i \in N: x \succ_i y\}$ the set of voters who prefer x to y , and with $N_{y>x} = \{i \in N: y \succ_i x\}$ the set of voters who prefer y to x . We let all voters in $N_{x>y}$ change their preferences such that y is directly below x , and all voters in $N_{y>x}$ change their preferences such that y is directly above x . For these steps, IUA implies that x remains the unique winner as we only reorder unchosen alternatives. Hence, it follows for the resulting profile R' that $f(R') = \{x\}$. However, it holds that $g_{R'}(x, z) = g_R(y, z)$ for all $z \in A \setminus \{x, y\}$ and $g_{R'}(x, y) = 0$. Neutrality and pairwise consistency thus require that either $\{x, y\} \subseteq f(R')$ or $\{x, y\} \cap f(R') \neq \emptyset$ because renaming x and y does not change the majority margins. This is in conflict with the previous claim, and hence the assumption that $f(R) = \{x\}$ and $g_R(x, y) = 0$ was incorrect. We have derived a contradiction in both cases, which proves that $f(R) = \{x\}$ can only be true if x is the Condorcet winner in R . \square

Since we established that every SCC that satisfies the requirements of Theorem 1 is strongly Condorcet consistent, this result follows now easily from Lemma 11.

THEOREM 1. *Let f be a pairwise SCC that satisfies nonimposition, homogeneity, and neutrality. Then f is strategyproof if and only if it is a robust dominant set rule.*

PROOF. Let f denote a pairwise SCC that satisfies homogeneity, neutrality, and nonimposition. The direction from left to right follows from Lemma 11 and Lemma 12: if f is additionally strategyproof, Lemma 12 shows that f is strongly Condorcet consistent and, in turn, Lemma 11 implies that f is a robust dominant set rule.

Next, we discuss the direction from right to left and assume thus that f is a robust dominant set rule. Furthermore, suppose for contradiction that f is not strategyproof. Hence, there are two preference profiles R and R' and a voter i such that $\succ_j = \succ'_j$ for all $j \in N \setminus \{i\}$ and $f(R') \succ_i^F f(R)$. First, assume that $f(R') \setminus f(R) \neq \emptyset$ and observe that $f(R) \succ_R f(R') \setminus f(R)$ since $f(R)$ is a dominant set. Deviating from R to

R' is only a manipulation for voter i if $f(R') \setminus f(R) \succ_i f(R)$. However, this means that $f(R) \succ_{R'} f(R') \setminus f(R)$ as voter i can only weaken the alternatives in $f(R') \setminus f(R)$ against those in $f(R)$. Since f is a dominant set rule and $f(R') \setminus f(R) \neq \emptyset$, this implies that $f(R) \subseteq f(R')$. Hence, $f(R) \succ_{R'} A \setminus f(R')$ because $f(R')$ is a dominant set, which proves that $f(R)$ is also a dominant set in $\succ_{R'}$. As a consequence, robustness requires that $f(R') \subseteq f(R)$, which contradicts the assumption that $f(R') \setminus f(R) \neq \emptyset$. Hence, no manipulation is possible in this case.

As a second case, suppose that $f(R') \subsetneq f(R)$. First, note that $f(R') \succ_{R'} A \setminus f(R')$ because $f(R')$ is a dominant set. Moreover, since deviating from R to R' is a manipulation for voter i , it holds that $f(R') \succ_i f(R) \setminus f(R')$. As a consequence of these two observations, it follows that $f(R') \succ_R f(R) \setminus f(R')$ because voter i can only weaken the alternatives in $f(R')$ against those in $f(R) \setminus f(R')$. Finally, since $f(R') \subseteq f(R)$, it follows that $f(R') \succ_R A \setminus f(R)$, and thus $f(R')$ is a dominant set in \succ_R . Hence, robustness from R' to R implies that $f(R) \subseteq f(R')$, which contradicts our assumption that $f(R') \subsetneq f(R)$. Thus, f is also in this case not manipulable, which shows that it is strategyproof. \square

Next, we focus on Theorem 2 and prove this result using Lemma 11. As a first step, we show that pairwise consistency, strategyproofness, homogeneity, and set nonimposition imply strong Condorcet consistency.

LEMMA 13. *Every pairwise SCC that satisfies set nonimposition, homogeneity, and strategyproofness is strongly Condorcet consistent.*

PROOF. Let f denote an SCC as specified by the lemma. First, note the proof of Claim 1 in Lemma 12 does not require neutrality and it thus shows that f is Condorcet consistent. Hence, we focus on the converse direction and show that $f(R) = \{x\}$ can only be true if x is the Condorcet winner in R . For this, suppose for contradiction that there is a profile R and an alternative x such that $f(R) = \{x\}$, but x is not the Condorcet winner in R . This means that there is another alternative $y \in A \setminus \{x\}$ such that $g_R(y, x) \geq 0$. If $g_R(y, x) > 0$, we can use the same construction as in the proof of Lemma 12 to derive that f violates Condorcet consistency since this construction works again without neutrality. Hence, suppose that $g_R(x, y) = 0$. In this case, we first weaken y in the preference relation of every voter $i \in N$ with $y \succ_i x$ such that it is directly over x and reinforce y in the preference relation of every voter $i \in N$ with $x \succ_i y$ such that it is placed directly below x . We infer from IUA that x is still the unique winner. Next, we iterate over the voters $i \in N$ and use WSMON to repeatedly push down voter i 's best alternative until he top-ranks x or y . It follows for the resulting profile R^1 that $f(R^1) = \{x\}$ because of WSMON and that all voters report x and y as their best two alternatives. Thereafter, we let the voters reorder the alternatives in $A \setminus \{x, y\}$ lexicographically. IUA implies that this step does not affect the choice set, and thus it holds for the new profile R^2 that $f(R^2) = f(R^1) = \{x\}$.

Next, we show that $f(R^2) = \{x\}$ is in conflict with set nonimposition. For this, consider a preference profile R^3 with $f(R^3) = \{x, y\}$; such a profile exists since f is

set nonimposing. Our goal is to transform R^3 into R^2 while showing that both x and y must be chosen. As a first step, we repeatedly add voters with inverse preferences and use WSMON to weaken the alternatives $z \in A \setminus \{x, y\}$ until we derive a profile R^4 with $\{x, y\} \succ_{R^4} A \setminus \{x, y\}$. It follows from pairwise and WSMON that $f(R^4) = f(R^3) = \{x, y\}$. This implies that $x \sim_{R^4} y$ because otherwise there is a Condorcet winner, which must be chosen uniquely. Even more, note that, as long as $x \sim_R y$ and $\{x, y\} \succ_R A \setminus \{x, y\}$, it holds that either $\{x, y\} \subseteq f(R)$ or $f(R) \subseteq \{x, y\}$. Otherwise, there is a profile R' and alternatives z_1, z_2 such that $z_1 \in f(R') \setminus \{x, y\}$ and $z_2 \in \{x, y\} \setminus f(R')$. Hence, if a voter reinforces $z_3 \in \{x, y\} \setminus \{z_2\}$ against z_2 , z_3 is the Condorcet winner and it must thus be chosen uniquely. However, this is in conflict with strategyproofness because WLOC (if $z_3 \in f(R')$) or IUA (if $z_3 \notin f(R')$) is violated.

We use the last observation to repeatedly reinforce the alternatives $\{x, y\}$ against the alternatives $A \setminus \{x, y\}$ in R^4 until all voters report x and y as their best two alternatives. For each swap, WMON implies that either $z_1 \in \{x, y\}$ remains chosen and $z_2 \in A \setminus \{x, y\}$ remains unchosen, or z_1 becomes unchosen and z_2 chosen. However, the latter is impossible because of our previous observation, and thus we derive from WMON and WLOC that the choice set is not allowed to change. Hence, it holds for the resulting profile R^5 that $f(R^5) = \{x, y\}$ and that all voters report x and y as their best two alternatives. Thereafter, we derive the profile R^6 from R^5 by arranging the alternatives in $A \setminus \{x, y\}$ in lexicographic order, which does not affect the choice set because of IUA. Finally, note that in R^6 , half of the voters report $\succ_1 = x, y, \text{lex}(A \setminus \{x, y\})$ and the other half report $\succ_2 = y, x, \text{lex}(A \setminus \{x, y\})$. Using homogeneity, it follows therefore that $f(R^7) = f(R^6)$, where R^7 consists of two voters who report \succ_1 and \succ_2 , respectively. Finally, the profile R^2 consists of multiple copies of R^7 , and hence it again follows from homogeneity that $f(R^2) = f(R^7) = \{x, y\}$. However, this contradicts the previous observation that $f(R^2) = \{x\}$, and hence f can only choose a single winner if it is the Condorcet winner. \square

Finally, we prove Theorem 2 based on Lemmas 1, 11, and 13.

THEOREM 2. *The top cycle is the only pairwise SCC that satisfies strategyproofness, set nonimposition, and homogeneity.*

PROOF. We have already shown in Lemma 1 that the top cycle satisfies set nonimposition. Moreover, by definition, TC is majoritarian and, therefore, also pairwise and homogeneous. Finally, the top cycle is a robust dominant set rule and hence strategyproof by Theorem 1. For the other direction, consider an arbitrary pairwise SCC f that satisfies strategyproofness, set nonimposition, and homogeneity. Since all criteria of Lemma 13 are satisfied, it follows that f is strongly Condorcet consistent. Next, we use Lemma 11 to derive that f is a robust dominant set rule. As the last step, Lemma 1 shows that f is the top cycle since this is the only robust dominant set rule that satisfies set nonimposition. \square

REFERENCES

- Bandyopadhyay, Taradas (1983), “Multi-valued decision rules and coalitional non-manipulability.” *Economics Letters*, 13, 37–44. [838, 840]
- Barberà, Salvador, Bhaskar Dutta, and Arunava Sen (2001), “Strategy-proof social choice correspondences.” *Journal of Economic Theory*, 101, 374–394. [838, 840]
- Benoît, Jean-Pierre (2002), “Strategic manipulation in voting games when lotteries and ties are permitted.” *Journal of Economic Theory*, 102, 421–436. [838]
- Bordes, Georges (1976), “Consistency, rationality and collective choice.” *Review of Economic Studies*, 43, 451–457. [841]
- Bordes, Georges (1979), “Some more results on consistency, rationality and collective choice.” In *Aggregation and Revelation of Preferences* (J. J. Laffont, ed.), 175–197, North-Holland. [849]
- Brandt, Felix (2011), “Minimal stable sets in tournaments.” *Journal of Economic Theory*, 146, 1481–1499. [841, 856]
- Brandt, Felix (2015), “Set-monotonicity implies Kelly-strategyproofness.” *Social Choice and Welfare*, 45, 793–804. [838, 854, 856]
- Brandt, Felix and Markus Brill (2011), “Necessary and sufficient conditions for the strategyproofness of irresolute social choice functions.” In *Proceedings of The 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 136–142. [840]
- Brandt, Felix, Markus Brill, and Paul Harrenstein (2016), “Tournament solutions.” In *Handbook of Computational Social Choice* (F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, eds.), Cambridge University Press. [840, 848, 854]
- Brandt, Felix, Markus Brill, Hans Georg Seedig, and Warut Suksompong (2018), “On the structure of stable tournament solutions.” *Economic Theory*, 65, 483–507. [841]
- Brandt, Felix, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, eds (2016), *Handbook of Computational Social Choice*, Cambridge University Press. [839, 842, 845]
- Brandt, Felix and Christian Geist (2016), “Finding strategyproof social choice functions via SAT solving.” *Journal of Artificial Intelligence Research*, 55, 565–602. [840]
- Brandt, Felix, Christian Saile, and Christian Stricker (2022), “Strategyproof social choice when preferences and outcomes may contain ties.” *Journal of Economic Theory*, 202, 105447. [838, 840, 844]
- Brandt, Felix and Hans Georg Seedig (2016), “On the discriminative power of tournament solutions.” In *Selected Papers of the International Conference on Operations Research, OR2014, Operations Research Proceedings*, 53–58, Springer-Verlag. [839]
- Campbell, Donald E. and Jerry S. Kelly (2003), “A strategy-proofness characterization of majority rule.” *Economic Theory*, 22, 557–568. [838]

- Ching, Stephen and Lin Zhou (2002), “Multi-valued strategy-proof social choice rules.” *Social Choice and Welfare*, 19, 569–580. [838, 840, 844, 849]
- Condorcet, Marquis (1785), *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale. Facsimile published in 1972 by Chelsea Publishing Company, New York. [839]
- Deb, Rajat (1977), “On Schwartz’s rule.” *Journal of Economic Theory*, 16, 103–110. [843]
- Duggan, John (2013), “Uncovered sets.” *Social Choice and Welfare*, 41, 489–535. [860]
- Duggan, John and Thomas Schwartz (2000), “Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized.” *Social Choice and Welfare*, 17, 85–93. [838]
- Ehlers, Lars, Dipjyoti Majumdar, Debasis Mishra, and Arunava Sen (2020), “Continuity and incentive compatibility in cardinal mechanisms.” *Journal of Mathematical Economics*, 88, 31–41. [846]
- Ehlers, Lars and Yves Sprumont (2008), “Weakened WARP and top-cycle choice rules.” *Journal of Mathematical Economics*, 44, 87–94. [841]
- Erdamar, Bora and M. Remzi Sanver (2009), “Choosers as extension axioms.” *Theory and Decision*, 67, 375–384. [838, 844]
- Evren, Ozgur, Hiroki Nishimura, and Efe A. Ok (2019), “Top-cycles and revealed preference structures.” Report. [841]
- Feldman, Allan (1979), “Manipulation and the Pareto rule.” *Journal of Economic Theory*, 21, 473–482. [840]
- Fishburn, Peter C. (1972), “Even-chance lotteries in social choice theory.” *Theory and Decision*, 3, 18–40. [838, 840]
- Fishburn, Peter C. (1977), “Condorcet social choice functions.” *SIAM Journal on Applied Mathematics*, 33, 469–489. [838]
- Gärdenfors, Peter (1976), “Manipulation of social choice functions.” *Journal of Economic Theory*, 13, 217–228. [837, 838, 840]
- Gärdenfors, Peter (1979), “On definitions of manipulation of social choice functions.” In *Aggregation and Revelation of Preferences* (J. J. Laffont, ed.). North-Holland. [838, 840, 844]
- Gehrlein, William V. and Dominique Lepelley (2011), *Voting Paradoxes and Group Coherence. Studies in Choice and Welfare*. Springer-Verlag. [839]
- Gibbard, Allan (1977), “Manipulation of schemes that mix voting with chance.” *Econometrica*, 45, 665–681. [856]
- Good, I. J. John (1971), “A note on Condorcet sets.” *Public Choice*, 10, 97–101. [839]
- Houy, Nicolas (2011), “Common characterizations of the untrapped set and the top cycle.” *Theory and Decision*, 70, 501–509. [841]

- Kelly, Jerry S. (1977), "Strategy-proofness and social choice functions without single-valuedness." *Econometrica*, 45, 439–446. [837]
- Laslier, Jean-François (1997), *Tournament Solutions and Majority Voting*. Springer-Verlag. [848, 856]
- Laslier, Jean-François (2010), "In silico voting experiments." In *Handbook on Approval Voting* (J.-F. Laslier and M. R. Sanver, eds.), 311–335, Springer-Verlag. [839]
- MacIntyre, I. and Prasanta K. Pattanaik (1981), "Strategic voting under minimally binary group decision functions." *Journal of Economic Theory*, 25, 338–352. [838]
- McGarvey, David C. (1953), "A theorem on the construction of voting paradoxes." *Econometrica*, 21, 608–610. [847, 850, 851, 863]
- Miller, Nicholas R., Bernard Grofman, and Scott L. Feld (1990), "The structure of the Banks set." *Public Choice*, 66, 243–251. [860]
- Moulin, Hervé (1986), "Choosing from a tournament." *Social Choice and Welfare*, 3, 271–291. [852]
- Moulin, Hervé (1988), *Axioms of Cooperative Decision Making*. Cambridge University Press. [855]
- Nandeibam, Shasikanta (2013), "The structure of decision schemes with cardinal preferences." *Review of Economic Design*, 17, 205–238. [846]
- Regenwetter, Michel, Bernhard Grofman, A. A. J. Marley, and Ilia M. Tsetlin (2006), *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press. [839]
- Sanver, M. Remzi and William S. Zwicker (2012), "Monotonicity properties and their adaption to irresolute social choice rules." *Social Choice and Welfare*, 39, 371–398. [840]
- Sato, Shin (2014), "A fundamental structure of strategy-proof social choice correspondences with restricted preferences over alternatives." *Social Choice and Welfare*, 42, 831–851. [838]
- Schwartz, Thomas (1972), "Rationality and the myth of the maximum." *Noûs*, 6, 97–117. [843]
- Schwartz, Thomas (1986), *The Logic of Collective Choice*. Columbia University Press. [839, 843]
- Sen, Amartya (1977), "Social choice theory: A re-examination." *Econometrica*, 45, 53–89. [839]
- Smith, John H. (1973), "Aggregation of preferences with variable electorate." *Econometrica*, 41, 1027–1041. [839]
- Taylor, Alan D. (2005), *Social Choice and the Mathematics of Manipulation*. Cambridge University Press. [837, 839, 840]

Zwicker, William S. (2016), “Introduction to the theory of voting.” In *Handbook of Computational Social Choice* (F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, eds.). Cambridge University Press. [856]

Co-editor Federico Echenique handled this manuscript.

Manuscript received 24 November, 2021; final version accepted 25 May, 2022; available online 11 July, 2022.

SUMMARY

The Gibbard-Satterthwaite theorem has established that only undesirable social choice functions are strategyproof when there are more than two alternatives: these rules are either dictatorial or some alternatives can never be chosen. In this paper, we thus investigate voting rules that choose a set of alternatives instead of a single winner. These voting rules are called social choice correspondences in most parts of this thesis, but we refer to them as (set-valued) social choice functions (SCFs) in the following publication. For set-valued SCFs, the consequences of strategyproofness are less clear: while most results are negative (e.g., Duggan and Schwartz, 2000; Benoît, 2002; Brandt et al., 2022c), there are also some possibility theorems (e.g., Nehring, 2000; Brandt, 2015). In particular, the simple and intuitive notion of weak \succsim^K -strategyproofness (henceforth Kelly-strategyproofness or simply strategyproofness) has turned out to be compelling as it allows for positive results. For example, the Pareto rule is Kelly-strategyproof even when the voters' preferences are weak, and several attractive SCFs (such as the top cycle, the uncovered set, and the essential set) are strategyproof for strict preferences.

In this paper, we aim to better understand Kelly-strategyproofness for weak preferences. To this end, we will show three far-reaching impossibility theorems which show that only indecisive voting rules are Kelly-strategyproof. In more detail, we prove that:

- every strategyproof rank-based SCF violates Pareto-optimality,
- every strategyproof and support-based SCF (which generalize Fishburn's C2 SCFs) fails Pareto-optimality or returns at least one most preferred alternative of every voter, and
- every strategyproof and non-imposing SCF returns the Condorcet loser in at least one profile.

From these results, we also derive several corollaries. For instance, our theorem on support-based SCCs can be used to show that no majority-based and non-imposing SCF satisfies strategyproofness. Moreover, we also discuss the consequences of our results for randomized social choice.

REFERENCE

F. Brandt, M. Bullinger, and P. Lederer. **On the indecisiveness of Kelly-strategyproof social choice functions.** In *Journal of Artificial Intelligence Research*, 73:1093–1130, 2022.
DOI: <https://doi.org/10.1613/jair.1.13449>

INDIVIDUAL CONTRIBUTION

I, Patrick Lederer, am the main author of this publication. In particular, I am responsible for the joint development and conceptual design of the research project, proofs of all results and the joint write-up of the manuscript (in particular, sections 2 to 4 and the appendix).

COPYRIGHT AGREEMENT

The right to present this paper in a doctoral thesis has been granted by the publisher, AI Access Foundation, in the copyright form presented below. There, it is stated that “*the author retains the right to use this material in future works of his or her own authorship.*” This form can also be found at <http://jair.org/public/resources/copyright.pdf> (accessed August 24, 2023).

TERMINOLOGY

In this paper, strategyproofness refers to weak \succsim^K -strategyproofness and voters report weak preferences. Moreover, social choice correspondences are called (set-valued) social choice functions in this publication.

AI Access Foundation

Transfer of Copyright

Title of Contribution:

Author:

The undersigned (Author), desiring to publish the article named above, in the Journal of Artificial Intelligence Research, hereby grants and assigns exclusively to the AI Access Foundation (Publisher) all rights of copyright in this article, and the exclusive right to copy and distribute the article throughout the world, and the authority to exercise or to dispose of all subsidiary rights in all countries and in all languages.

The author retains the right to use this material in future works of his or her own authorship. The author retains the right to post an electronic copy of the article on his or her web site, on web sites operated by the author's institution, and on scholarly repositories, provided there is no commercial purpose involved, and with appropriate bibliographic citation to the JAIR article.

The Author agrees that the contribution does not contain any material created by others, or from other copyrighted works, unless the written consent of the owner of such material is attached.

The undersigned warrants that s/he has full power to make this agreement, and has not previously granted, assigned, or encumbered any of the rights granted and assigned herein; that the material submitted for publication is the work of the author, and is original, unpublished, and not previously the subject of any application for copyright registration (except material for which the Publisher has written grants of permission to include, as described above).

Author

Date

On the Indecisiveness of Kelly-Strategyproof Social Choice Functions

Felix Brandt

Martin Bullinger

Patrick Lederer

Institut für Informatik

Technische Universität München

Boltzmannstr. 3, 85748 Garching, Germany

BRANDTF@IN.TUM.DE

BULLINGE@IN.TUM.DE

LEDERERP@IN.TUM.DE

Abstract

Social choice functions (SCFs) map the preferences of a group of agents over some set of alternatives to a non-empty subset of alternatives. The Gibbard-Satterthwaite theorem has shown that only extremely restrictive SCFs are strategyproof when there are more than two alternatives. For set-valued SCFs, or so-called social choice correspondences, the situation is less clear. There are miscellaneous—mostly negative—results using a variety of strategyproofness notions and additional requirements. The simple and intuitive notion of Kelly-strategyproofness has turned out to be particularly compelling because it is weak enough to still allow for positive results. For example, the Pareto rule is strategyproof even when preferences are weak, and a number of attractive SCFs (such as the top cycle, the uncovered set, and the essential set) are strategyproof for strict preferences. In this paper, we show that, for weak preferences, only indecisive SCFs can satisfy strategyproofness. In particular, *(i)* every strategyproof rank-based SCF violates Pareto-optimality, *(ii)* every strategyproof support-based SCF (which generalize Fishburn’s C2 SCFs) that satisfies Pareto-optimality returns at least one most preferred alternative of every voter, and *(iii)* every strategyproof non-imposing SCF returns the Condorcet loser in at least one profile. We also discuss the consequences of these results for randomized social choice.

1. Introduction

Whenever a group of agents aims at reaching a joint decision in a fair and principled way, they need to aggregate their individual preferences using a social choice function (SCF). SCFs are traditionally studied by economists and mathematicians, but have also come under increasing scrutiny from computer scientists who are interested in their computational properties or want to utilize them in computational multiagent systems (see, e.g., Brandt et al., 2016b; Endriss, 2017).

An important phenomenon in social choice is that agents misrepresent their preferences in order to obtain a more preferred outcome. An SCF that is immune to strategic misrepresentation of preferences is called strategyproof. Gibbard (1973) and Satterthwaite (1975) have shown that only extremely restrictive single-valued SCFs are strategyproof: either the range of the SCF is restricted to only two outcomes or the SCF always returns the most preferred alternative of the same voter. Perhaps the most controversial assumption of the Gibbard-Satterthwaite theorem is that the SCF must always return a single alternative (see, e.g., Gärdenfors, 1976; Kelly, 1977; Barberà, 1977b; Duggan & Schwartz, 2000;

Nehring, 2000; Barberà et al., 2001; Ching & Zhou, 2002; Taylor, 2005). This assumption is at variance with elementary fairness conditions such as anonymity and neutrality. For instance, consider an election with two alternatives and two voters such that each alternative is favored by a different voter. Clearly, both alternatives are equally acceptable, but single-valuedness forces us to pick a single alternative based on the preferences only.

We therefore study the manipulability of *set-valued* SCFs (or so-called *social choice correspondences*). When SCFs return sets of alternatives, there are various notions of strategyproofness, depending on the circumstances under which one set is considered to be preferred to another. When the underlying notion of strategyproofness is sufficiently strong, the negative consequences of the Gibbard-Satterthwaite theorem remain largely intact (see, e.g., Duggan & Schwartz, 2000; Barberà et al., 2001; Ching & Zhou, 2002; Benoît, 2002; Sato, 2014).¹ In this paper, we are concerned with a rather weak—but natural and intuitive—notation of strategyproofness attributed to Kelly (1977). Several attractive SCFs have been shown to be strategyproof for this notion when preferences are strict (Brandt, 2015; Brandt et al., 2016a). These include the top cycle, the uncovered set, the minimal covering set, and the essential set. However, when preferences are weak, these results break down and strategyproofness is not well understood in general.

Feldman (1979) has shown that the Pareto rule is strategyproof according to Kelly’s definition, even when preferences are weak. Moreover, the omninomination rule and the intersection of the Pareto rule and the omninomination rule are strategyproof as well (Brandt et al., 2022, Remark 1). These results are encouraging because they rule out impossibilities using Pareto-optimality and other weak properties.² In the context of strategic abstention (i.e., manipulation by deliberately abstaining from an election), even more positive results can be obtained. Brandl et al. (2019) have shown that all of the above mentioned SCFs that are strategyproof for strict preferences are immune to strategic abstention even when preferences are weak.

A number of negative results were shown for severely restricted classes of SCFs. Kelly (1977) and Barberà (1977b) have shown independently that there is no strategyproof SCF that satisfies quasi-transitive rationalizability. However, this result suffers from the fact that quasi-transitive rationalizability is almost prohibitive on its own (see, e.g., Mas-Colell & Sonnenschein, 1972).³ In subsequent work by MacIntyre and Pattanaik (1981) and Bandyopadhyay (1983), quasi-transitive rationalizability has been replaced with weaker conditions such as minimal binariness or quasi-binariness, which are still very demanding and violated by most SCFs. Barberà (1977a) has shown that positively responsive SCFs fail to be strategyproof under mild assumptions. However, positively responsive SCFs are almost always single-valued and of all commonly considered SCFs only Borda’s rule and Black’s rule satisfy this criterion. More recently, Taylor (2005, Theorem 8.1.2) has proven that every SCF that returns the set of all weak Condorcet winners whenever this set is non-empty fails to be strategyproof. This result was strengthened by Brandt (2015), who showed that ev-

1. We refer to Barberà (2010) and Brandt et al. (2022) for a more detailed overview over this extensive stream of research.

2. For example, Brandt et al. (2022) have shown that Pareto-optimality is incompatible with anonymity and a notion of strategyproofness that is slightly stronger than Kelly’s.

3. This is acknowledged by Kelly (1977) who writes that “one plausible interpretation of such a theorem is that, rather than demonstrating the impossibility of reasonable strategy-proof social choice functions, it is part of a critique of the regularity [rationalizability] conditions.”

ery SCF that uniquely returns the (strict) Condorcet winner whenever one exists fails to be strategyproof. Brandt et al. (2022) have shown with the help of computers that every Pareto-optimal SCF whose outcome only depends on the pairwise majority margins can be manipulated.

Note that—in contrast to most other work—the strong impossibility theorems by Brandt (2015) and Brandt et al. (2022) require weak preference relations, i.e., these authors assume that preference relations are transitive and complete, but not necessarily anti-symmetric. We follow this approach because ties arise quite naturally in many applications. In fact, we see little justification to assume that all agents entertain strict preferences. For example, a voter who strongly cares about the environment may deem all parties that deny climate change equally unacceptable. Moreover, preferential voting rules are often criticized for being impractical because they put an unduly heavy burden on voters by asking them to submit a complete and strict ranking of, say, 20 alternatives. This burden can be reduced by allowing voters to express indifferences between similar alternatives. The case of indifferences is even more striking when the set of alternatives consists of partitions of agents or assignments of objects to agents. In these settings, agents are likely to be indifferent between coalitions or assignments in which they are grouped with the same agents or in which they receive the same objects.

For these reasons, we study strategyproofness based on the assumption that voters can express indifferences between alternatives. In particular, we investigate three broad classes of SCFs: *rank-based* SCFs (which include all scoring rules), *support-based* SCFs (which generalize Fishburn’s C2 SCFs), and *non-imposing* SCFs (which return every alternative as the unique winner for some preference profile). An overview of the three classes and typical examples of SCFs belonging to these classes are given in Figure 1. The classes are unrelated in a set-theoretic sense: for any subset of classes, there exist SCFs which lie precisely in these classes. For instance, Borda’s rule is contained in all three classes. Taken together, they cover virtually all SCFs commonly considered in the literature.

For rank-based and support-based SCFs, we show that Pareto-optimality and strategyproofness imply that every voter is a nominator, i.e., the resulting choice sets contain at least one most preferred alternative of every voter. In the case of rank-based SCFs, this entails an impossibility (Theorem 1) whereas for support-based SCFs it demonstrates a high degree of indecisiveness in the sense that the SCF tends to return large choice sets (Theorem 2). For non-imposing SCFs, we show that strategyproofness implies that the Condorcet loser has to be returned in at least one preference profile (Theorem 3). The latter result remarkably holds without imposing fairness conditions such as anonymity or neutrality and can again be phrased in terms of indecisiveness: every strategyproof SCF that satisfies the Condorcet loser property will never return certain alternatives alone. Hence, our main theorems can be summarized by the observation that strategyproofness requires an SCF to return unreasonably large choice sets for some preference profiles.

All these results rely on two auxiliary statements (Lemmas 1 and 2) which discuss the relationship between decisive, nominating, and vetoing groups of voters. Roughly, these concepts address how much influence a group of voters has when acting unanimously, e.g., a group of voters is decisive if it can ensure a subset of its best alternatives to be chosen and vetoing if it can ensure its unique least preferred alternative not to be chosen. Due to their universality, these lemmas may be of independent interest. Our results can also

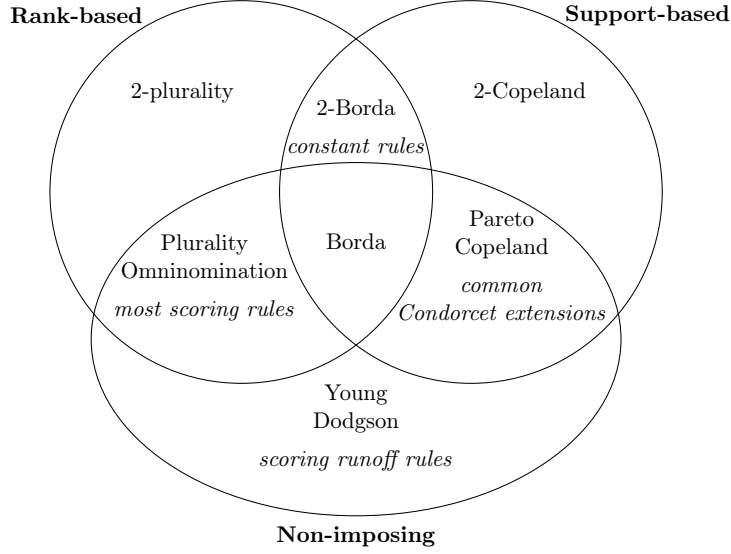


Figure 1: The classes of rank-based, support-based, and non-imposing SCFs and typical examples. 2-plurality, 2-Copeland, and 2-Borda return all alternatives whose respective score is at least as large as the second-highest score. All scoring rules except Borda’s rule are rank-based, non-imposing, but not support-based. Common Condorcet extensions include the top cycle, the uncovered set, the minimal covering set, the essential set, the Simpson-Kramer rule, Nanson’s rule, Schulze’s rule, and Kemeny’s rule. We refer to Brandt et al. (2016b, Chapters 2-5) for definitions of these SCFs.

be interpreted in the context of randomized social choice (where the outcome is a lottery over the alternatives instead of a set of alternatives). In more detail, all our axioms can be transferred to the randomized setting, and thus we also derive strong impossibilities for randomized social choice.

Even though our main results are rather negative, they are important to improve our understanding of strategyproof SCFs. Much more positive results are obtained by making minuscule adjustments to the assumptions such as restricting the domain of preferences to strict preferences, weakening the underlying notion of strategyproofness, or replacing strategic manipulation with strategic abstention (see, e.g., Nehring, 2000; Brandt, 2015; Brandl et al., 2019). In all of these cases, a small number of support-based Condorcet extensions such as the top cycle, the uncovered set, the minimal covering set, and the essential set constitute appealing positive examples.

2. The Model

Let $N = \{1, \dots, n\}$ denote a finite set of voters and let $A = \{a, b, \dots\}$ denote a finite set of m alternatives. Moreover, let $[x \dots y] = \{i \in N : x \leq i \leq y\}$ denote the subset of voters from x to y and note that $[x \dots y]$ is empty if $x > y$. Every voter $i \in N$ is equipped

with a *weak preference relation* \succsim_i , i.e., a complete and transitive binary relation on A . We denote the strict part of \succsim_i by \succ_i , i.e., $x \succ_i y$ if and only if $x \succsim_i y$ and $y \not\succsim_i x$, and the indifference part by \sim_i , i.e., $x \sim_i y$ if and only if $x \succsim_i y$ and $y \succsim_i x$. We compactly represent preference relations as comma-separated lists, where sets of alternatives express indifferences. For example, $x \succ y \sim z$ is represented by $x, \{y, z\}$. Furthermore, we call a preference relation \succsim *strict* if its irreflexive part is equal to its strict part \succ . The set of all weak preference relations on A is called \mathcal{R} . A *preference profile* $R \in \mathcal{R}^n$ is an n -tuple containing the preference relation of every voter $i \in N$. When defining preference profiles, we specify a set of voters who share the same preference relation by writing the set directly before the preference relation. For instance, $[x \dots y]: a, b, c$ means that all voters $i \in [x \dots y]$ prefer a to b to c . We omit brackets for singleton sets.

Our central object of study are *social choice functions (SCFs)*, or so-called social choice correspondences, which map preference profiles to non-empty sets of alternatives, i.e., functions of the form $f : \mathcal{R}^n \rightarrow 2^A \setminus \{\emptyset\}$.

2.1 Axioms for Social Choice Functions

We now introduce axioms that formalize desirable properties for SCFs, all of which are well-known in the literature. A basic fairness condition is anonymity, which requires that all voters are treated equally: an SCF f is *anonymous* if $f(R) = f(R')$ for all preference profiles R, R' for which there is a permutation $\pi : N \rightarrow N$ such that $\succsim'_i = \succsim_{\pi(i)}$ for all $i \in N$.

Perhaps one of the most prominent axioms in economic theory is Pareto-optimality, which is based on the notion of Pareto-dominance: an alternative x *Pareto-dominates* another alternative y if $x \succsim_i y$ for all $i \in N$ and there is a voter $j \in N$ with $x \succ_j y$. An alternative is *Pareto-optimal* if it is not Pareto-dominated by any other alternative. This idea leads to the *Pareto rule* which returns all Pareto-optimal alternatives. Finally, an SCF f is *Pareto-optimal* if it never returns Pareto-dominated alternatives.

An axiom that is closely related to Pareto-optimality is *near unanimity*, as introduced by Benoît (2002). Near unanimity requires that $f(R) = \{x\}$ for all alternatives $x \in A$ and preference profiles R in which at least $n - 1$ voters uniquely top-rank x . The more voters there are, the more compelling near unanimity is.

A natural weakening of these axioms is *non-imposition* which requires that for every alternative $x \in A$, there is a profile R such that $f(R) = \{x\}$. For single-valued SCFs, non-imposition is almost imperative because it merely requires that the SCF is surjective. For set-valued SCFs, as considered in this paper, this is not necessarily the case. For example, every SCF that always returns at least two alternatives fails non-imposition (see, for example, 2-plurality, 2-Borda, and 2-Copeland in Figure 1).

An influential concept in social choice theory is that of a Condorcet winner, which is an alternative that wins all pairwise majority comparisons. For formally defining this term, let the pairwise *support* of x over y denote the number of voters who strictly prefer x to y , i.e., $s_{xy}(R) = |\{i \in N : x \succ_i y\}|$. Then, an alternative $a \in A$ is a *Condorcet winner* if $s_{ax}(R) > s_{xa}(R)$ for all $x \in A \setminus \{a\}$. An SCF is *Condorcet-consistent* or a so-called *Condorcet extension* if it uniquely returns the Condorcet winner whenever one exists.

Analogously, an alternative a is called *Condorcet loser* if $s_{xa}(R) > s_{ax}(R)$ for all $x \in A \setminus \{a\}$. An SCF f satisfies the *Condorcet loser property* if $x \notin f(R)$ whenever x is the

Condorcet loser in R . While there are Condorcet extensions that violate the Condorcet loser property (e.g., the Simpson-Kramer rule) and SCFs that satisfy the Condorcet loser property but fail Condorcet-consistency (e.g., Borda's rule), the Condorcet loser property “feels” weaker. This follows the intuition that both properties affect exactly the same number of preference profiles, but the Condorcet loser property only excludes a single alternative (and otherwise leaves a lot of freedom) whereas Condorcet-consistency completely determines the (singleton) choice set.

2.2 Strategyproofness

One of the central problems in social choice theory is manipulation, i.e., voters may lie about their true preferences to obtain a more preferred outcome. For single-valued SCFs, it is clear what constitutes a more preferred outcome. In the case of set-valued SCFs, there are various ways to define a manipulation depending on the assumptions about the voters' preferences over sets of alternatives. Here, we make a simple and natural assumption first considered by Kelly (1977): a voter i weakly prefers a set X to another set Y , denoted by $X \succsim_i Y$, if and only if $x \succsim_i y$ for all $x \in X, y \in Y$. Thus, the strict part of this preference extension is

$$X \succ_i Y \text{ if and only if for all } x \in X, y \in Y, x \succsim_i y \text{ and} \\ \text{there are } x' \in X, y' \in Y \text{ with } x' \succ_i y'.$$

An SCF is manipulable if a voter can improve his outcome by lying about his preferences. Formally, an SCF f is *manipulable* if there are a voter $i \in N$ and preference profiles R, R' such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$ and $f(R') \succ_i f(R)$. Moreover, f is *strategyproof* if it is not manipulable.

These assumptions can, for example, be justified by considering a randomized tie-breaking procedure (a so-called lottery) that is used to select a single alternative from every set of alternatives returned by the SCF. We then have that $X \succ_i Y$ if and only if all lotteries with support X yield strictly more expected utility than all lotteries with support Y for all utility functions that are ordinally consistent with \succsim_i (see, e.g., Gärdenfors, 1979; Brandt et al., 2022).

Note that, in the proofs of this paper, we often do not need the full power of strategyproofness. Instead, we mainly consider two types of manipulations: either the original choice set only consists of the manipulator's least preferred alternatives, or the new choice set only consists of the manipulator's most preferred alternatives. In order to formalize these situations, we define $T_i(R)$ as the set of voter i 's top-ranked alternatives in R and $B_i(R)$ as the set of voter i 's bottom-ranked alternatives in R . We then derive the following two consequences of strategyproofness, where R and R' are two preference profiles that only differ in the preference relation of voter i :

(SP1) If $f(R) \subseteq B_i(R)$, then $f(R') \subseteq B_i(R)$.

(SP2) If $f(R) \subseteq T_i(R')$, then $f(R') \subseteq T_i(R')$.

SP1 states that, if a subset of voter i 's least preferred alternatives is the choice set for R , then the choice set after a manipulation is also a subset of voter i 's least preferred

alternatives; any other outcome constitutes a manipulation for voter i . On the other hand, *SP2* states that turning the current choice set $f(R)$ into a subset of voter i 's most preferred alternatives in R' results into a choice set $f(R')$ that is a subset of $T_i(R')$. If this was not true, voter i could manipulate f by deviating from R' to R .

For a better understanding of strategyproofness, *SP1*, and *SP2*, consider the following profiles R and R' and assume that f is a strategyproof SCF.

R :	1: $\{a, b\}, c, \{d, e\}$	2: $\{a, c\}, \{d, e\}, b$	3: d, c, a, b, e
R' :	1: $\{b, c\}, a, \{d, e\}$	2: $\{a, c\}, \{d, e\}, b$	3: d, c, a, b, e

If $f(R) = \{d\}$, then *SP1* requires that $f(R') \subseteq \{d, e\}$ since voter 1 prefers every set X with $X \not\subseteq \{d, e\}$ to $f(R)$. Furthermore, if $f(R) = \{c\}$, then *SP2* implies that $f(R') \subseteq \{b, c\}$; otherwise voter 1 can manipulate by reverting from R' to R . Finally, note that strategyproofness is stronger than the conjunction of *SP1* and *SP2*, e.g., if $f(R) = \{c, d\}$, then $f(R') \neq \{b, c\}$ as voter 1 could manipulate otherwise.

2.3 Decisive, Nominating, and Vetoing Groups of Voters

A common concern in social choice theory is that single voters or groups of voters might be more influential than others (see, e.g., Le Breton & Weymark, 2011). Perhaps the most prominent example of such a notion are dictators: a voter $i \in N$ is a *dictator* for an SCF f if $f(R)$ always chooses a subset of voter i 's most preferred alternatives, i.e., if $f(R) \subseteq T_i(R)$ for all profiles R . The existence of dictators is usually undesirable because it means that a single voter can determine the outcome of the election alone.

A related but far less restrictive concept concerns the notion of nominators: a voter $i \in N$ is a *nominator* for an SCF f if $f(R)$ always contains at least one of his most preferred alternatives. More formally, a voter $i \in N$ is a nominator for an SCF f if $f(R) \cap T_i(R) \neq \emptyset$ for all preference profiles R . Nominators are weak dictators in the sense that they can always force an alternative into the choice set by reporting it as their unique top choice.

Finally, we formalize the converse idea that a voter might be able to prohibit an alternative from being chosen. This leads to the notion of a *vetoer*, which is a voter $i \in N$ such that $f(R)$ does never contain the uniquely least preferred alternative of voter i . If a vetoer does not report a uniquely least preferred alternative, the corresponding SCF is not restricted. Formally, a voter $i \in N$ is a vetoer for an SCF f if $f(R) \cap B_i(R) = \emptyset$ for all preference profiles R with $|B_i(R)| = 1$.

In the context of social choice, the existence of dictators, nominators, and vetoers is often undesirable as these notions formalize that some voters have an undesirably large impact on the outcome. For instance, if a voter is a nominator, he can force an alternative x into the choice set even if all other voters agree that x is the worst option. To avoid these problems, it is natural to consider generalizations of dictators, nominators, and vetoers to groups of voters. We opt for a rather weak generalization of these axioms and require that all voters in the group need to report the same preference relation to influence the SCF.

We say that a non-empty set of voters $I \subseteq N$ is *decisive* if $f(R)$ is a subset of the best alternatives of the voters $i \in I$ for all profiles R such that $\succsim_i = \succsim_j$ for all $i, j \in I$. The concept of decisive groups of voters is best known from Arrow's proof of his impossibility theorem (Arrow, 1951). Similarly, a non-empty set of voters $I \subseteq N$ is *nominating* if $f(R)$ contains at least one of the most preferred alternatives of the voters $i \in I$ for all profiles R .

such that $\succsim_i = \succsim_j$ for all $i, j \in I$, and *vetoing* if $f(R)$ excludes the uniquely least preferred alternative of the voters $i \in I$ for all such profiles.

The notions of decisive, nominating, and vetoing groups of voters are far less restrictive than the corresponding single voter notions. More precisely, if a single voter fulfills any of these properties, then also every group containing this agent does. In fact, many desirable axioms even imply that sufficiently large groups of voters need to be decisive or vetoing. For instance, Pareto-optimality implies that the set of all voters is decisive, and near unanimity for strategyproof SCFs is equivalent to the requirement that every group of $n - 1$ voters is decisive.⁴ Furthermore, the Condorcet loser property implies that every group I with $|I| > \frac{n}{2}$ is vetoing. As we will show, there are strong relationships between the notions of decisive, nominating, and vetoing groups for strategyproof SCFs.

2.4 Rank-Basedness and Support-Basedness

In this section, we introduce two classes of anonymous SCFs that capture many of the SCFs commonly studied in the literature: rank-based and support-based SCFs. The basic idea of rank-basedness is that voters assign ranks to the alternatives and that an SCF should only depend on the ranks of the alternatives, but not on which voter assigns which rank to an alternative. In order to formalize this idea, we first need to define the rank of an alternative. In the case of strict preferences, this is straightforward: the rank of alternative x according to \succsim_i is $\bar{r}(\succsim_i, x) = |\{y \in A : y \succsim_i x\}|$ (Laslier, 1996). By contrast, there are multiple possibilities how to define the rank in the presence of ties. We define a new and very weak notion of rank-basedness for weak preferences, making our results only stronger. To this end, define the *rank tuple* of x with respect to \succsim_i as

$$\begin{aligned} r(\succsim_i, x) &= (\bar{r}(\succ_i, x), \bar{r}(\sim_i, x)) \\ &= (|\{y \in A : y \succ_i x\}|, |\{y \in A : y \sim_i x\}|). \end{aligned}$$

The rank tuple contains more information than many other generalizations of the rank and therefore, it leads to a more general definition of rank-basedness. Next, we define the *rank vector* of an alternative x which contains the rank tuple of x with respect to every voter in increasing lexicographic order, i.e., $r^*(R, x) = (r(\succsim_{i_1}, x), r(\succsim_{i_2}, x), \dots, r(\succsim_{i_n}, x))$ where $\bar{r}(\succ_{i_j}, x) \leq \bar{r}(\succ_{i_{j+1}}, x)$ and if $\bar{r}(\succ_{i_j}, x) = \bar{r}(\succ_{i_{j+1}}, x)$, then $\bar{r}(\sim_{i_j}, x) \leq \bar{r}(\sim_{i_{j+1}}, x)$. Finally, the *rank matrix* $r^*(R)$ of the preference profile R contains the rank vectors as rows. An SCF f is called *rank-based* if $f(R) = f(R')$ for all preference profiles $R, R' \in \mathcal{R}^n$ with $r^*(R) = r^*(R')$. The class of rank-based SCFs contains many popular SCFs such as all scoring rules or the omninomination rule, which returns all top-ranked alternatives.⁵

A similar line of thought leads to support-basedness, which is based on the pairwise support of an alternative x against another one y . Recall that the pairwise support refers to the number of voters who strictly prefer x to y , i.e., $s_{xy}(R) = |\{i \in N : x \succ_i y\}|$. We define the *support matrix* $s^*(R) = (s_{xy}(R))_{x, y \in A}$ which contains the supports for all pairs of alternatives. Finally, an SCF f is *support-based* if $f(R) = f(R')$ for all preference profiles

4. The claim on near unanimity may fail for manipulable SCFs as near unanimity only affects profiles where $n - 1$ voters *uniquely* top-rank the same alternative, whereas a decisive group I affects the outcome for all profiles R with $\succsim_i = \succsim_j$ for all $i, j \in I$.

5. We refer to Brandt et al. (2016b, Chapters 2-5) for the definitions of these and all following SCFs.

$R, R' \in \mathcal{R}^n$ with $s^*(R) = s^*(R')$. Note that support-basedness is a new generalization of Fishburn's C2 to weak preferences (Fishburn, 1977). Hence, many well-known SCFs such as Borda's rule, Kemeny's rule, the Simpson-Kramer rule, Nanson's rule, Schulze's rule, the Pareto rule, and the top cycle are support-based.

Support-basedness is less restrictive than *pairwiseness*, which requires that $f(R) = f(R')$ for all preference profiles $R, R' \in \mathcal{R}^n$ with $s_{ab}(R) - s_{ba}(R) = s_{ab}(R') - s_{ba}(R')$ for all $a, b \in A$ (see, e.g., Brandt et al., 2022). For example, the Pareto rule is support-based, but fails to be pairwise. Another important subclass of support-based SCFs are majoritarian ones, which are merely based on the majority relation. To this end, we define the *majority relation* \succsim_R of a profile R as $\succsim_R = \{(a, b) \in A^2 : s_{ab}(R) \geq s_{ba}(R)\}$. Then, an SCF f is *majoritarian* if $f(R) = f(R')$ for all preference profiles R and R' with $\succsim_R = \succsim_{R'}$ (see, e.g., Brandt et al., 2019). For instance, the top cycle is majoritarian, whereas all other previous examples rely on the exact supports for computing the outcomes and thus fail this axiom.

In order to illustrate the definitions of rank-based, support-based, and majoritarian SCFs, we discuss two classical examples. First, consider the plurality rule, which returns all alternatives x that maximize $|\{i \in N : \bar{r}(\succ_i, x) = 0\}|$. By definition, this SCF is rank-based, but it is not support-based. The latter claim follows by considering the following preference profiles R and R' because $s^*(R) = s^*(R')$ but the plurality rule chooses $\{a\}$ for R and $\{a, b\}$ for R' .

R :	1: $\{a, b\}, c$	2: a, b, c	3: c, b, a
R' :	1: $\{a, b\}, c$	2: a, c, b	3: b, c, a

As second example, consider the Pareto rule, which chooses all Pareto-optimal alternatives. This SCF is support-based because an alternative x Pareto-dominates another alternative y if and only if $s_{xy}(R) > 0$ and $s_{yx}(R) = 0$. On the other hand, it violates rank-basedness because there are profiles with the same rank matrix but different sets of Pareto-optimal alternatives (see Claim 1 in the proof of Theorem 1 for details). Finally, the Pareto rule is not majoritarian since it chooses $\{a\}$ for \bar{R} and $\{a, b\}$ for \bar{R}' , but $\succsim_{\bar{R}} = \succsim_{\bar{R}'}$.

\bar{R} :	1: $\{a, b\}, c$	2: $\{a, b\}, c$	3: a, b, c
\bar{R}' :	1: b, a, c	2: a, b, c	3: a, b, c

3. Results

The unifying theme of our results is that strategyproofness requires a high degree of indecisiveness. In more detail, we show that every voter is a nominator for all rank-based and support-based SCFs that satisfy Pareto-optimality and strategyproofness. Consequently, such SCFs have to choose a large number of alternatives for most preference profiles as one of the best alternatives of every voter needs to be in the choice set. For the very broad class of non-imposing SCFs, we show that every strategyproof SCF violates the Condorcet loser property. Put differently, for every strategyproof SCF that satisfies the Condorcet loser property, there is an alternative x that is not returned as unique winner even if it is unanimously top-ranked.

In order to prove the claim for rank-based and support-based SCFs, we focus on its contrapositive, i.e., we assume that there is a rank-based or support-based SCF f that satisfies Pareto-optimality and strategyproofness and for which a voter $i \in N$ is not a

nominator. We first show that a group of voters is not nominating for a Pareto-optimal and strategyproof SCF if and only if its complement is decisive.

Lemma 1. *Let f be a Pareto-optimal and strategyproof SCF that is defined for $m \geq 3$ alternatives and $n \geq 2$ voters. A group of voters I with $\emptyset \subsetneq I \subsetneq N$ is not nominating for f if and only if $N \setminus I$ is decisive for f .*

Proof. Let f denote a Pareto-optimal and strategyproof SCF and consider an arbitrary set of voters I with $\emptyset \subsetneq I \subsetneq N$. First, we show that I is not nominating for f if $N \setminus I$ is decisive. This follows immediately by considering a preference profile R in which all voters in $N \setminus I$ report an alternative a as best choice and are indifferent between the alternatives in $A \setminus \{a\}$, and the voters in I report another alternative b as their unique top choice and are indifferent between all alternatives in $A \setminus \{b\}$. Then, $f(R) = \{a\}$ because $N \setminus I$ is decisive for f , which proves that I is not nominating as this condition requires that $b \in f(R)$.

For the other direction, suppose that the group I is not nominating for f . Our goal is to show that the group $N \setminus I$ is decisive for f and we observe for this that there is a profile R^0 such that $f(R^0) \cap T_i(R^0) = \emptyset$ and $\succsim_i^0 = \succsim_j^0$ for all voters $i, j \in I$ because I is not nominating for f . Subsequently, we will apply multiple transformations to R^0 : first, we deduce a profile R' such that $f(R') \cap T_i(R') = \emptyset$ for all $i \in I$ and $f(R') = \{x\}$ for some alternative $x \in f(R^0)$. Secondly, we infer from this profile that $f(R) = \{x\}$ for all preferences profiles R in which the voters $j \in N \setminus I$ prefer x uniquely the most. As third step, we generalize this observation from a single alternative to all alternatives. This is reminiscent of the so-called *field expansion lemma* in proofs of Arrow's theorem (see, e.g., Sen, 1986). Finally, we extend our analysis to the case where the voters in $N \setminus I$ may top-rank multiple alternatives. For an easier notation of the subsequent arguments, we assume that $I = \{1, \dots, k\}$ for some $k \in \{1, \dots, n-1\}$; this is without loss of generality because all our arguments are independent of the naming of the voters.

Step 1: As first step, we let the voters $j \in N \setminus I = \{k+1, \dots, n\}$ replace their preference relations in R^0 sequentially such that they prefer the alternatives in $f(R^0)$ the most. More formally, this means that we consider a sequence of preference profiles $R^{0,0}, \dots, R^{0,n-k}$ such that $R^{0,0} = R^0$ and $R^{0,i}$ evolves out of $R^{0,i-1}$ by assigning voter $k+i$ a preference relation such that $T_{k+i}(R^{0,i}) = f(R^0)$. For each $i \in \{1, \dots, n-k\}$, *SP2* implies that $f(R^{0,i}) \subseteq f(R^0)$ if $f(R^{0,i-1}) \subseteq f(R^0)$ because $f(R^{0,i-1}) \subseteq T_{k+i}(R^{0,i-1})$. Since we start this process at the profile R^0 , we derive a preference profile $R^1 = R^{0,n-k}$ with $f(R^1) \subseteq f(R^0)$. Next, let $a \in f(R^0)$ denote an alternative such that $a \succsim_i b$ for all $b \in f(R^0)$ and $i \in I$, i.e., a is the most preferred alternative of the voters $i \in I$ in $f(R^0)$. Such an alternative exists since $\succsim_i^1 = \succsim_j^1$ for all $i, j \in I$. Moreover, let B denote the set of alternatives such that $b \sim_i^1 a$ for all $b \in B$ and $i \in I$. As next step, we sequentially replace the preference relations of the voters $i \in I = \{1, \dots, k\}$ in R^1 with a preference relation where all alternatives in $T_i(R^1)$ are preferred to a , which, in turn, is preferred to all alternatives in $A \setminus (T_i(R^1) \cup \{a\})$. More formally, we consider again a sequence of preference profiles $R^{1,0}, \dots, R^{1,k}$ such that $R^{1,0} = R^1$ and $R^{1,i}$ is derived from $R^{1,i-1}$ by modifying the preference relation of voter i as described in the last sentence. Next, we will show for all $i \in \{1, \dots, k\}$ that if $f(R^{1,i-1}) \subseteq B$, then $f(R^{1,i}) \subseteq B$. Observe for this that, for all profiles $R^{1,i}$, alternative a Pareto-dominates all alternatives $x \in A$ with $a \succ_j^1 x$ for all $j \in I$. Hence, Pareto-optimality ensures that

$x \notin f(R^{1,i})$ for all these alternatives. This means that if $f(R^{1,i-1}) \subseteq B$, then $f(R^{1,i}) \subseteq B$ because voter i strictly prefers every Pareto-optimal alternative $x \in A \setminus B$ to all alternatives in B , i.e., every set of Pareto-optimal alternatives $X \not\subseteq B$ constitutes a manipulation for voter i . Finally, observe that $f(R^1) \subseteq B$ because $f(R^1) \subseteq f(R^0)$ and all alternatives $x \in f(R^0)$ with $a \succ_i^1 x$ are Pareto-dominated in R^1 . We can therefore repeatedly apply the previous argument to derive that $f(R^2) \subseteq B$ for the profile $R^2 = R^{1,k}$. Moreover, observe that in R^2 , all voters $i \in I$ prefer a to all alternatives $y \in A \setminus (T_i(R^0) \cup \{a\})$ and the voters $j \in N \setminus I$ top-rank a . Hence, a Pareto-dominates all other alternatives in B , which implies that $f(R^2) = \{a\}$.

Step 2: Given the preference profile R^2 from the last step, we show that $f(R) = \{a\}$ for all preferences profiles R in which the voters in $N \setminus I$ prefer a uniquely the most. We deduce this result by modifying and analyzing the profile R^2 . First, we sequentially change the preference relation of all voters $j \in N \setminus I$ such that they prefer a uniquely the most and an alternative $b \in T_i(R^2)$ (for $i \in I$) uniquely the second most. Formally, this yields another sequence of preference profiles $R^{2,0}, \dots, R^{2,n-k}$ such that $R^{2,0} = R^2$ and $R^{2,i}$ is derived from $R^{2,i-1}$ by making a into the uniquely best alternative and b into the uniquely second best alternative of voter $k+i$. Just as for the sequence $R^{0,i}$, *SP2* implies that if $f(R^{2,i-1}) = \{a\}$, then $f(R^{2,i}) = \{a\}$ because $f(R^{2,i-1}) = \{a\} = T_{k+i}(R^{2,i})$. Since $f(R^2) = \{a\}$, we infer for the profile $R^3 = R^{2,n-k}$ that $f(R^3) = \{a\}$ by repeatedly applying this argument. Furthermore, every alternative in $A \setminus \{a, b\}$ is Pareto-dominated by b in R^3 . We use this observation to sequentially replace the current preference relations of the voters $i \in I$ with a preference relation in which b is the uniquely most preferred alternative and a is his uniquely least preferred alternative. Formally, this leads to another sequence of profiles $R^{3,0}, \dots, R^{3,k}$ that starts at R^3 and one by one changes the preference relation of the voters $i \in I$ as described. For every profile $R^{3,i}$, it holds that only a and b can be chosen because of Pareto-optimality. Moreover, if $f(R^{3,i-1}) = \{a\}$, then $f(R^{3,i}) = \{a\}$ as any other subset of $\{a, b\}$ constitutes a manipulation for voter i . Hence, this process results in a profile $R^4 = R^{3,k}$ such that $f(R^4) = \{a\}$, all voters $i \in I$ prefer a uniquely the least, and all voters $j \in N \setminus I$ prefer a uniquely the most. It follows now from *SP1* and *SP2* that $f(R) = \{a\}$ for all preference profiles R in which the voters $j \in N \setminus I$ prefer a uniquely the most: *SP1* allows the voters $i \in I$ to deviate to any other preference relation without changing the choice set because $B_i(R^4) = f(R^4) = \{a\}$ and *SP2* allows the voters $i \in N \setminus I$ to reorder the alternatives in $A \setminus \{a\}$ arbitrarily because $f(R^4) = \{a\}$ is after the deviation still their set of top-ranked alternatives.

Step 3: As next step, we show that the voters in $N \setminus I$ can make every alternative win uniquely if they report it as their common top choice. Thus, consider the preference profile R^5 in which all voters in $N \setminus I$ prefer a uniquely the most, and the voters in I prefer c uniquely the most, b uniquely second most, and a uniquely the least. It follows from the last step that $f(R^5) = \{a\}$. Next, let the voters $j \in N \setminus I$ change their preferences sequentially such that they prefer a and b the most. Formally, this leads to another sequence of preference profiles $R^{5,0}, \dots, R^{5,n-k}$ and *SP2* implies that if $f(R^{5,i-1}) \subseteq \{a, b\}$, then $f(R^{5,i}) \subseteq \{a, b\}$ because $T_i(R^{5,i}) = \{a, b\}$. Hence, it holds for the profile $R^6 = R^{5,n-k}$ that $f(R^6) = \{b\}$: our previous argument implies that $f(R^6) \subseteq \{a, b\}$ and b Pareto-dominates a in this profile. Thereafter, we replace the preference relation of every voter $i \in N \setminus I$ with a new preference

in which he prefers b uniquely the most. *SP2* shows for the corresponding sequence of profiles $R^{6,0}, \dots, R^{6,n-k}$ that $f(R^{6,i}) = \{b\}$ if $f(R^{6,i-1}) = \{b\}$. Therefore, this sequence results in a new preference profile $R^7 = R^{6,n-k}$ with $f(R^7) = \{b\}$. Since the voters $i \in I$ do not top-rank b , we can now apply the constructions in Step 2 to deduce that b is uniquely chosen if all voters in $N \setminus I$ voters prefer it uniquely the most.

Step 4: Finally, it remains to prove that $f(R) \subseteq T_i(R)$ for all voters $i \in N \setminus I$ and preference profiles R such that $\succsim_i = \succsim_j$ for all $i, j \in N \setminus I$. If the voters in $N \setminus I$ only report a single alternative x as their top choice, this claim follows from Step 3, which requires that x is the unique winner. Hence, consider a profile R^8 such that $\succsim_i^8 = \succsim_j^8$ for $i, j \in N \setminus I$, $|T_i(R)| \geq 2$ for all $i \in N \setminus I$, and let a denote one of the top-ranked alternatives of these voters. Moreover, define R^9 as the profile derived from R^8 by making a into the unique best alternative of every voter $i \in N \setminus I$. Step 3 implies for R^9 that $f(R^9) = \{a\}$. Moreover, we can go from R^9 to R^8 by letting the voters $i \in N \setminus I$ one after another revert back to the preference relation \succsim_i^8 . Since all these voters have the same preference relation in R^8 and $a \in T_i(R^8)$ for all $i \in N \setminus I$, it follows from a repeated application *SP2* that $f(R^8) \subseteq T_i(R^8)$, which proves the lemma. \square

Lemma 1 has a number of interesting consequences. First of all, it shows that, for every non-empty set of voters $I \subseteq N$, either I is nominating or $N \setminus I$ is decisive for a strategyproof and Pareto-optimal SCF. Since anonymity implies that no set $I \subseteq N$ with $|I| \leq \frac{n}{2}$ can be decisive, it follows that every set with more than half of the voters is nominating for a Pareto-optimal, strategyproof, and anonymous SCF. Furthermore, this lemma shows that, under Pareto-optimality, strategyproofness, and anonymity, indecisiveness for a single preference profile of a particularly simple type entails a large degree of indecisiveness for the entire domain of preference profiles: if an alternative is not chosen uniquely even if $n - l$ voters prefer it uniquely the most, then all groups of size l are nominating. This already indicates that strategyproof SCFs are rather indecisive under mild additional assumptions. For our subsequent proofs, the inverse direction of Lemma 1 is more interesting: if a voter i is not a nominator, then the set $N \setminus \{i\}$ is decisive. This means that the absence of nominators implies near unanimity for strategyproof and Pareto-optimal SCFs. Furthermore, if we additionally assume anonymity, we have near unanimity even if a single voter is not a nominator.

Remark 1. Remarkably, many impossibility results rule out that every voter is a nominator. For instance, Duggan and Schwartz (2000), Benoît (2002), and Sato (2008) invoke axioms prohibiting that every voter is a nominator. Moreover, a crucial step in the computer-generated proofs by Brandl et al. (2018, Theorem 3.1) and Brandt et al. (2022, Theorem 1) is to show that there is a voter who is not a nominator. Lemma 1 shows that these assumptions and observations imply the existence of a decisive group of size $n - 1$, which is in conflict with strategyproofness as defined by the above authors. Intuitively, a decisive group of size $n - 1$ is already too small to allow for their notions of strategyproofness.

3.1 Rank-Based SCFs

In this section, we prove that there is no rank-based SCF that satisfies Pareto-optimality and strategyproofness. This result follows from the observation that Pareto-optimality,

strategyproofness, and rank-basedness require that every voter is a nominator, but Pareto-optimality and rank-basedness do not allow for such SCFs.

It is possible to show Theorem 1—as well as Theorem 2—by induction proofs where completely indifferent voters and universally bottom-ranked alternatives are used to generalize the statement to arbitrarily many voters and alternatives (see, e.g., Brandl et al., 2018, 2019; Brandt et al., 2022). Instead, we prefer to give universal proofs for any number of voters and alternatives to stress the robustness of the respective constructions. As a consequence, our proofs usually hold when restricting the domain of admissible profiles by prohibiting artificial constructs such as completely indifferent voters. Note that we often assume that all voters are indifferent between all but a few alternatives $A \setminus X$. This assumption is not required and is only used for the sake of simplicity. In fact, the preferences between alternatives in X can be arbitrary and may differ from voter to voter and often even between profiles. The only restriction is that the preferences involving alternatives in $A \setminus X$ are not modified.

Theorem 1. *There is no rank-based SCF that satisfies Pareto-optimality and strategyproofness if $m \geq 4$ and $n \geq 3$, or if $m \geq 5$ and $n \geq 2$.*

Proof. Consider fixed numbers of voters n and alternatives m such that $m \geq 4$ and $n \geq 3$, or $m \geq 5$ and $n \geq 2$. Furthermore, suppose for contradiction that there is a rank-based SCF f that satisfies strategyproofness and Pareto-optimality for the given values of n and m . We derive a contradiction to this assumption by proving two claims: on the one hand, there is a voter who is not a nominator for f . On the other hand, the assumptions on the SCF require that every voter is a nominator. These two claims contradict each other and therefore f cannot exist.

Claim 1: Not every voter is a nominator for f .

First, we prove that not every voter is a nominator for f . For this, we use a case distinction and first suppose that $m \geq 4$ and $n \geq 3$. In this case, consider the following three profiles, where $X = A \setminus \{a, b, c, d\}$.

$$\begin{array}{lll} R^1: & 1: \{a, b\}, X, \{c, d\} & 2: \{c, d\}, X, \{a, b\} \quad [3 \dots n]: a, \{b, c, d\}, X \\ R^2: & 1: \{a, c\}, X, \{b, d\} & 2: \{b, d\}, X, \{a, c\} \quad [3 \dots n]: a, \{b, c, d\}, X \\ R^3: & 1: \{a, d\}, X, \{b, c\} & 2: \{b, c\}, X, \{a, d\} \quad [3 \dots n]: a, \{b, c, d\}, X \end{array}$$

It can be easily verified that $r^*(R^1) = r^*(R^2) = r^*(R^3)$ and that a Pareto-dominates b in R^1 , c in R^2 , and d in R^3 . This means that $f(R^1) = f(R^2) = f(R^3) \subseteq \{a\} \cup X$ because of rank-basedness and Pareto-optimality. Consequently, voter 2 is not a nominator for f .

Next, we focus on the case that $m \geq 5$ and $n = 2$ and consider the profiles R^4 , R^5 , and R^6 , where $X = A \setminus \{a, b, c, d, e\}$.

$$\begin{array}{ll} R^4: & 1: \{a, b\}, X, e, \{c, d\} \quad 2: \{c, d\}, X, a, \{b, e\} \\ R^5: & 1: \{a, c\}, X, e, \{b, d\} \quad 2: \{b, d\}, X, a, \{c, e\} \\ R^6: & 1: \{a, d\}, X, e, \{b, c\} \quad 2: \{b, c\}, X, a, \{d, e\} \end{array}$$

Analogous to the last case, it can be verified that $r^*(R^4) = r^*(R^5) = r^*(R^6)$, and that a Pareto-dominates b in R^4 , c in R^5 , and d in R^6 . Consequently, rank-basedness and Pareto-

optimality imply that $f(R^4) = f(R^5) = f(R^6) \subseteq \{a, e\} \cup X$, which proves that voter 2 is not a nominator for f .

Claim 2: Every voter is a nominator for f .

Assume for contradiction that a voter is not a nominator for f and let $X = A \setminus \{a, b, c, d\}$. Consequently, it follows from rank-basedness that no voter is a nominator and therefore, Lemma 1 shows that all sets of $n - 1$ voters are decisive. Furthermore, we want to point out that the subsequent construction works for all $n \geq 2$ and $m \geq 4$, which means that no case distinction is required. Our proof focuses on the profiles $R^{1,k}$ and $R^{2,k}$ for $k \in \{1, \dots, n\}$ shown below.

$$\begin{array}{lll} R^{1,k}: & 1: \{c, d\}, X, b, a & [2 \dots k]: \{a, b\}, X, c, d \quad [k+1 \dots n]: a, X, b, c, d \\ R^{2,k}: & 1: \{b, d\}, X, c, a & [2 \dots k]: \{a, b\}, X, c, d \quad [k+1 \dots n]: a, X, b, c, d \end{array}$$

We prove by induction on $k \in \{1, \dots, n\}$ that $f(R^{1,k}) = f(R^{2,k}) = \{a\}$. The case $k = n$ yields a contradiction to Pareto-optimality as a is Pareto-dominated by b in $R^{1,n}$.

The base case $k = 1$ follows because $n - 1$ voters prefer a uniquely the most in both $R^{1,1}$ and $R^{2,1}$. Therefore, our previous observation that every set of $n - 1$ voters is decisive shows that $f(R^{1,1}) = f(R^{2,1}) = \{a\}$.

Assume now that the induction hypothesis is true for some fixed $k \in \{1, \dots, n - 1\}$, i.e., $f(R^{1,k}) = f(R^{2,k}) = \{a\}$. By induction and *SP2*, $f(R^{1,k+1}) \subseteq \{a, b\}$ since otherwise voter $k + 1$ can manipulate by switching back to $R^{1,k}$. Next, we derive the profile $R^{3,k}$ shown below from $R^{2,k}$ by assigning voter $k + 1$ the preference relation $\{a, c\}, X, b, d$.

$$\begin{array}{lll} R^{3,k}: & 1: \{b, d\}, X, c, a & [2 \dots k]: \{a, b\}, X, c, d \\ & k+1: \{a, c\}, X, b, d & [k+2 \dots n]: a, X, b, c, d \end{array}$$

The induction hypothesis entails that $f(R^{2,k}) = \{a\}$ and therefore, *SP2* implies that $f(R^{3,k}) \subseteq \{a, c\}$ because $f(R^{2,k}) \subseteq T_{k+1}(R^{3,k})$. Next, we apply rank-basedness to conclude that $f(R^{1,k+1}) = \{a\}$ as $r^*(R^{1,k+1}) = r^*(R^{3,k})$. Finally, $R^{2,k+1}$ evolves from $R^{1,k+1}$ by having voter 1 change his preferences. Since $B_1(R^{1,k+1}) = \{a\} = f(R^{1,k+1})$, *SP1* implies that $f(R^{2,k+1}) = \{a\}$ as any other outcome benefits voter 1. This proves the induction step and therefore also the theorem. \square

Remark 2. The axioms used in Theorem 1 are independent: the Pareto rule satisfies all axioms except rank-basedness, the trivial SCF which always returns all alternatives only violates Pareto-optimality, and Borda's rule only violates strategyproofness.⁶ Furthermore, the Pareto rule is rank-based if $m \leq 3$, and if $m = 4$ and $n \leq 2$ (cf. Proposition 2 in the appendix), which entails that the bounds on m and n are tight.

Remark 3. Theorem 1 is only an impossibility because of the lack of compatibility of rank-basedness and Pareto-optimality in Claim 1, independently of strategyproofness. By contrast, the main consequence of strategyproofness is indecisiveness as captured in Claim 2. Indeed, Theorem 1 breaks down once we weaken Pareto-optimality to weak Pareto-optimality (which only excludes alternatives for which another alternative is strictly preferred by every voter) as then the omninomination rule satisfies all required axioms (Brandt et al., 2022,

6. We define Borda's rule as the SCF that chooses all alternatives x that minimize $\sum_{i \in N} \bar{r}(\succ_i, x) + \frac{1}{2} \bar{r}(\sim_i, x)$. This definition of Borda's rule for weak preferences is equivalent to the one suggested by Young (1974).

Remark 6). By contrast, Claim 2 is rather robust since a number of variations are true: for instance, it is easy to adapt the proof of this claim to show that no neutral, strategyproof, and rank-based SCF satisfies near unanimity if $m \geq 4$ and $n \geq 3$, or $m \geq 5$ and $n = 2$.⁷ Furthermore, the proof also reveals that a rank-based SCF that satisfies neutrality and strategyproofness can only choose a unique winner if this alternative is never uniquely bottom-ranked by a voter.

Remark 4. Theorem 1 also holds under weaker versions of rank-basedness. First, our proof uses rank-basedness only in very specific situations, namely when two voters rename exactly two alternatives. Moreover, the only real restriction on the rank function r is independence of the naming of other alternatives, i.e., $r(\succsim_i, a) = r(\succsim'_i, a)$ for all preference relations \succsim_i, \succsim'_i that only differ in the naming of alternatives in $A \setminus \{a\}$. Hence, we may also define rank-basedness based on a rank function other than the rank tuple and the result still holds.

Remark 5. Theorem 1 does not hold when preferences are strict. For instance, the omninomination rule satisfies all required axioms for arbitrary numbers of voters and alternatives for strict preferences. It can even be shown that Claim 2 of the proof no longer holds for strict preferences as the following SCF is rank-based, Pareto-optimal, and strategyproof: if an alternative is top-ranked by every voter, this alternative is the unique winner; otherwise, return the alternatives which are top-ranked by the most and second most voters (in case of a tie return all alternatives with the second highest plurality score). However, no voter is a nominator for this rule. A proof of these claims and a formal definition of this SCF can be found in Proposition 1 in the appendix.

3.2 Support-Based SCFs

It is not possible to replace rank-basedness with support-basedness in Theorem 1 since the Pareto rule is strategyproof, Pareto-optimal, and support-based. Note that the Pareto rule always chooses one of the most preferred alternatives of every voter. Consequently, Claim 1 in the proof of Theorem 1 cannot be true in general for support-based SCFs. Nevertheless, we show next that an analogue statement to Claim 2 remains true for such SCFs, i.e., every voter is a nominator for every support-based SCF that satisfies Pareto-optimality and strategyproofness.

Theorem 2. *In every support-based SCF that satisfies Pareto-optimality and strategyproofness, every voter is a nominator if $m \geq 3$.*

Proof. Let f be a support-based SCF satisfying Pareto-optimality and strategyproofness for fixed numbers of voters $n \geq 1$ and alternatives $m \geq 3$. For $n = 1$, the theorem follows immediately from Pareto-optimality as only the most preferred alternatives of the single voter are Pareto-optimal. Moreover, Lemma 1 proves the theorem for $n = 2$. Indeed, if a voter is not a nominator, support-basedness shows that no voter is a nominator. Hence, Lemma 1 shows that every voter is a dictator, which means that $f(R) = \{a\}$ and $f(R) = \{b\}$ are simultaneously true if voter 1 prefers a uniquely the most and voter 2 prefers b uniquely the most. This is a contradiction and proves the theorem if $n = 2$.

7. An SCF is neutral if $f(\pi(R)) = \pi(f(R))$ for all permutations $\pi: A \rightarrow A$ and preference profiles R, R' .

Therefore, we focus on the case that $n \geq 3$ and assume for contradiction that a voter is not a nominator for f . We derive from this assumption by an induction on $k \in \{1, \dots, n-1\}$ that every set of $n - k$ voters is decisive. This results in a contradiction when $k \geq n/2$ because then, two alternatives can be simultaneously top-ranked by $n - k \leq n/2$ voters, and both of them must be the unique winner.

The induction basis $k = 1$ follows directly from Lemma 1: support-basedness implies that if a single voter is not a nominator for f , no voter is a nominator for f as we can just rename the voters. Hence, every set of size $n - 1$ is decisive. Next, we assume that our claim holds for a fixed $k \in \{1, \dots, n - 2\}$ and prove that also every set of $n - (k + 1)$ voters is decisive. For this, we focus only on three alternatives a, b, c and on a certain partition of the voters. This is possible as the induction hypothesis allows us to exchange the roles of the alternatives without affecting the proof and support-basedness allows us to reorder the voters. Thus, consider the profile $R^{k,1}$, in which $X = A \setminus \{a, b, c\}$, and note that $f(R^{k,1}) = \{a\}$ because of Lemma 1.

$$R^{k,1}: \quad [1 \dots k]: a, X, c, b \quad k+1: c, X, b, a \quad [k+2 \dots n]: a, b, X, c$$

Next, we aim to reverse the preferences of the voters $i \in [k+2 \dots n]$ over a and b . This is achieved by the repeated application of the following steps explained for voter $k + 2$. First, voter $k + 2$ changes his preference to $\{a, b\}, c, X$ to derive the profile $R^{k,2}$. Since a subset of $\{a, b\}$ was chosen before this step, *SP2* implies that $f(R^{k,2}) \subseteq \{a, b\}$. Next, we use support-basedness to exchange the preferences of voter $k + 1$ and $k + 2$ over a and b . This leads to the profile $R^{k,3}$ and support-basedness implies that $f(R^{k,3}) = f(R^{k,2}) \subseteq \{a, b\}$. Since $\{a, b\} = B_{k+1}(R^{k,3})$, *SP1* implies that this voter cannot make another alternative win by manipulating. Thus, he can switch back to his original preference to derive $R^{k,4}$ and the fact that $f(R^{k,4}) \subseteq \{a, b\}$.

$$R^{k,2}: \quad [1 \dots k]: a, X, c, b \quad k+1: c, X, b, a \quad k+2: \{a, b\}, X, c \quad [k+3 \dots n]: a, b, X, c$$

$$R^{k,3}: \quad [1 \dots k]: a, X, c, b \quad k+1: c, X, \{a, b\} \quad k+2: b, a, X, c \quad [k+3 \dots n]: a, b, X, c$$

$$R^{k,4}: \quad [1 \dots k]: a, X, c, b \quad k+1: c, X, b, a \quad k+2: b, a, X, c \quad [k+3 \dots n]: a, b, X, c$$

It is easy to see that we can repeat these steps for every voter $i \in [k+3 \dots n]$. This process results in the profile $R^{k,5}$ and shows that $f(R^{k,5}) \subseteq \{a, b\}$. Moreover, consider the profile $R^{k,6}$ derived from $R^{k,5}$ by letting voter $k + 1$ make b his best alternative. Because $n - k$ voters prefer b uniquely the most in $R^{k,6}$, the induction hypothesis entails that $f(R^{k,6}) = \{b\}$. This means that voter $k + 1$ can manipulate by switching from $R^{k,5}$ to $R^{k,6}$ if $f(R^{k,5}) = \{a\}$ or $f(R^{k,5}) = \{a, b\}$. Consequently, $f(R^{k,5}) = \{b\}$ is the only valid choice set for $R^{k,5}$.

$$R^{k,5}: \quad [1 \dots k]: a, X, c, b \quad k+1: c, X, b, a \quad [k+2 \dots n]: b, a, X, c$$

$$R^{k,6}: \quad [1 \dots k]: a, X, c, b \quad k+1: b, X, a, c \quad [k+2 \dots n]: b, a, X, c$$

So far, we have found a profile in which b is uniquely chosen when the voters $i \in [k+2 \dots n]$ prefer it uniquely the most. Next, we show that this set of voters is therefore decisive. Hence, consider the profile $R^{k,7}$ which is derived from $R^{k,5}$ by letting the voters $i \in [1 \dots k]$ subsequently change their preference to c, X, a, b . Since $f(R^{k,5}) = \{b\}$ and b is the worst alternative for these voters, *SP1* implies that $f(R^{k,7}) = \{b\}$. As last step, we change the preferences of voter $k + 1$ such that b is his least preferred alternative. For

this, we first let all voters $i \in [k+2 \dots n]$ subsequently change their preference to b, X, c, a . This modification results in the profile $R^{k,8}$ and *SP2* implies that $f(R^{k,8}) = \{b\}$. Moreover, observe that alternative a is Pareto-dominated by c in $R^{k,8}$. Therefore, voter $k+1$ can now swap a and b to derive the profile $R^{k,9}$ and Pareto-optimality implies that $a \notin f(R^{k,9})$. Then, strategyproofness implies that $f(R^{k,9}) = \{b\}$ as any other subset of $A \setminus \{a\}$ is a manipulation for voter $k+1$.

$R^{k,7}$:	$[1 \dots k]: c, X, a, b$	$k+1: c, X, b, a$	$[k+2 \dots n]: b, a, X, c$
$R^{k,8}$:	$[1 \dots k]: c, X, a, b$	$k+1: c, X, b, a$	$[k+2 \dots n]: b, X, c, a$
$R^{k,9}$:	$[1 \dots k]: c, X, a, b$	$k+1: c, X, a, b$	$[k+2 \dots n]: b, X, c, a$

Finally, observe that the voters $i \in [1 \dots k+1]$ can change their preferences in $R^{k,9}$ arbitrarily without affecting the choice set because of *SP1*, and the voters $i \in [k+2 \dots n]$ can reorder all alternatives in $A \setminus \{b\}$ without affecting the choice set because of *SP2*. Thus, b is always the unique winner if all voters $i \in [k+2 \dots n]$ prefer b uniquely the most. Moreover, interchanging the roles of alternatives in this proof shows that every alternative is chosen uniquely if it is uniquely top-ranked by all voters in $[k+2 \dots n]$.

Next, we show that this set of voters is decisive and consider an arbitrary profile R such that $\succsim_i = \succsim_j$ for all $i, j \in [k+2 \dots n]$. If these voters only report a single alternative x as top choice, $f(R) = \{x\}$ follows from our previous analysis. Otherwise, we start at a profile R' in which the voters $i \in [k+2 \dots n]$ uniquely top-rank an alternative $x \in T_i(R)$. Once again, our previous analysis shows that $f(R') = \{x\}$ and if we let the voters $i \in [k+2 \dots n]$ sequentially deviate to their preference relation in R , *SP2* shows that $f(R) \subseteq T_i(R)$ for $i \in [k+2 \dots n]$. Hence, this set is indeed decisive. Since support-basedness allows us to reorder the voters to derive that every set of $n - (k+1)$ voters is decisive, the induction step is proven. As a consequence, every voter is a nominator for a support-based SCF that satisfies strategyproofness and Pareto-optimality. \square

Theorem 2 shows that every support-based SCF that satisfies Pareto-optimality and strategyproofness chooses one of the most preferred alternatives of every voter. Since the Pareto rule indeed satisfies all these criteria, this is no impossibility but demonstrates a high degree of indecisiveness. However, we can turn this result into an impossibility by strengthening support-basedness. For instance, if we require pairwiseness instead of support-basedness, Theorem 2 turns into an impossibility since pairwiseness and Pareto-optimality rule out that every voter is a nominator. For seeing this, consider the following two profiles R^1 and R^2 .

R^1 :	1: $\{a, b\}, X$	2: $\{a, b\}, X$	$[3 \dots n]: a, b, X$
R^2 :	1: b, a, X	2: a, b, X	$[3 \dots n]: a, b, X$

It can be verified that $f(R^2) = f(R^1) = \{a\}$ for every pairwise and Pareto-optimal SCF f , which shows that voter 1 is not a nominator. Hence, Theorem 2 implies that there is no pairwise, strategyproof, and Pareto-optimal SCF.⁸ As a consequence of this result, it follows also that no majoritarian SCF can satisfy Pareto-optimality and strategyproofness. As the next corollary shows, this impossibility is even true if we weaken Pareto-optimality to non-imposition.

8. This impossibility was first observed by Brandt et al. (2022, Theorem 2).

Corollary 1. *There is no majoritarian SCF that satisfies non-imposition and strategyproofness if $m \geq 3$ and $n \geq 3$.*

Proof. Assume for contradiction that there is a majoritarian SCF f that satisfies non-imposition and strategyproofness for $m \geq 3$ and $n \geq 3$. As a first step, we show that f satisfies Condorcet-consistency. Hence, choose an arbitrary alternative a and consider a profile R such that $f(R) = \{a\}$; such a profile exists by non-imposition. Next, we let the voters $i \in N$ one after another report a as their favorite alternative. For each step, *SP2* shows that the choice set does not change, and thus, this process results in a profile R' with $f(R') = \{a\}$ and $T_i(R') = a$ for all voters $i \in N$. Furthermore, we infer from *SP2* also that all voters can reorder all alternatives in $A \setminus \{a\}$ in R' without affecting the choice set. Since f is majoritarian, this means that a is the unique winner for all profiles in which a is the Condorcet winner, i.e., f is Condorcet-consistent.

As a consequence of this observation, every group I of at least $\lceil \frac{n+1}{2} \rceil$ voters can enforce that f chooses an alternative x uniquely if all voters in I report x as their unique top choice. Moreover, an analogous argument as in Step 4 of Lemma 1 shows that every such group I is decisive for f . Note that this statement is equivalent to the induction hypothesis in the proof of Theorem 2 when $k = \lfloor \frac{n-1}{2} \rfloor$. Since the first steps of this proof do not require Pareto-optimality, we derive that $f(R^1) = \{b\}$ because the profile R^1 corresponds to the profile $R^{k,5}$ in the proof of Theorem 2, where $k = \lfloor \frac{n-1}{2} \rfloor$.

$$R^1: \quad [1 \dots k]: a, X, c, b \quad k+1: c, X, b, a \quad [k+2 \dots n]: b, a, X, c$$

As next step, we let the voters $i \in [1 \dots k]$ deviate one after another by reporting c, b, X, a . Since b is the least preferred alternative of these voters in R^1 and $f(R^1) = \{b\}$, *SP1* requires that the choice set does not change during these steps. Hence, it holds for the resulting profile R^2 that $f(R^2) = \{b\}$. Furthermore, one after another, we let the voters $i \in [k+2 \dots n]$ make c into their second best alternative. This results in the profile R^3 and *SP2* implies that $f(R^3) = \{b\}$.

$$\begin{aligned} R^2: & \quad [1 \dots k]: c, b, X, a & k+1: c, X, b, a & [k+2 \dots n]: b, a, X, c \\ R^3: & \quad [1 \dots k]: c, b, X, a & k+1: c, X, b, a & [k+2 \dots n]: b, c, a, X \end{aligned}$$

Finally, note that $f(R^3) = \{b\}$ is a contradiction. If n is odd, then c is the Condorcet winner in $f(R^3)$ and thus, Condorcet-consistency requires that $f(R^3) = \{c\}$. On the other hand, if n is even, we can exchange the roles of b and c in the derivation of R^3 to derive that $f(R^3) = \{c\}$ must also be true. This is possible as $c \sim_{R^3} b$ and $x \succ_{R^3} y$ for all $x \in \{b, c\}$, $y \in A \setminus \{b, c\}$. Hence, if we exchange the role of b and c in the derivation of the profile R^3 , we end up with another profile $R^{3'}$ with the same majority relation, but our proof shows that $f(R^{3'}) = \{c\}$. This is in conflict with f being majoritarian and thus, no majoritarian SCF satisfies both strategyproofness and non-imposition if $m \geq 3$ and $n \geq 3$. \square

Remark 6. All axioms used in Theorem 2 are required as the following SCFs show. Every constant SCF satisfies support-basedness and strategyproofness, and violates Pareto-optimality and that every voter is a nominator. The SCF that always chooses a unique Pareto-optimal alternative according to a fixed tie-breaking order satisfies Pareto-optimality and support-basedness but violates strategyproofness and that every voter is a nominator. An SCF that satisfies Pareto-optimality and strategyproofness but violates support-basedness and that every voter is a nominator can be found as follows. We define a transitive

dominance relation by slightly strengthening Pareto-dominance. Therefore, we additionally allow that an alternative a that is among the most preferred alternatives of $n - 1$ voters can dominate another alternative b , even if a single voter strictly prefers b to a . More formally, we say that an alternative a *dominates* alternative b if a Pareto-dominates b or $n - 1$ voters prefer a the most while $s_{ab}(R) \geq 2$ and $s_{ba}(R) \leq 1$. It should be stressed that it is not required that a is uniquely top-ranked by $n - 1$ voters, but only that it is among their best alternatives. The SCF f^* that chooses all maximal elements with respect to this dominance relation satisfies all required properties (see Proposition 3 in the appendix for more details). Also the bound on m is tight as the majority rule satisfies all axioms if $m = 2$ but no voter is a nominator for this SCF.

Remark 7. Theorem 2 implies an impossibility for $m \geq 4$ and $n \geq 4$ if we strengthen Pareto-optimality to *SD*-efficiency (also known as ordinal efficiency, see Bogomolnaia & Moulin, 2001). This result follows by considering the preference profiles R^1 and R^2 shown below. In this profile, $X = A \setminus \{a, b, c, d\}$ and all voters $i \in [5 \dots n]$ are assumed to be indifferent between all alternatives.

R^1 :	1: a, c, b, d, X	2: a, d, b, c, X	3: b, c, a, d, X	4: b, d, a, c, X
R^2 :	1: a, b, c, d, X	2: a, b, d, c, X	3: c, b, a, d, X	4: d, b, a, c, X

Next, consider a support-based and *SD*-efficient SCF f . *SD*-efficiency implies that $f(R^1) \cap X = \emptyset$ and that either $c \notin f(R^1)$ or $d \notin f(R^1)$. Moreover, support-basedness implies that $f(R^1) = f(R^2)$, which means that voter 3 or 4 is not a nominator for f . Thus, Theorem 2 implies that f is not strategyproof, which shows the incompatibility of strategyproofness, support-basedness, and *SD*-efficiency. Note that there is also no rank-based, strategyproof, and *SD*-efficient SCF if $m \geq 4$ and $n \geq 3$ because of Theorem 1. This leads to the important and challenging open question whether there is an anonymous, *SD*-efficient, and strategyproof SCF.

Remark 8. Just as in the proof of Theorem 1, we make only very restricted use of support-basedness in the proof of Theorem 2. It suffices if two voters are allowed to exchange their preferences over two alternatives. This technical restriction is significantly weaker than support-basedness, which allows any number of voters to change their preferences.

Remark 9. If preferences are required to be strict, Theorem 2 does not hold. Several SCFs including the uncovered set, the minimal covering set, and the essential set are strategyproof, Pareto-optimal and support-based, but no voter is a nominator for them (for more details, see, e.g., Brandt et al., 2016b, Chapter 3). Remarkably, all these SCFs are majoritarian and thus affected by the stronger impossibility in Corollary 1 if we allow for ties in voters' preferences.

Remark 10. Theorems 1 and 2 raise the question whether all voters must be nominators for every anonymous, Pareto-optimal, and strategyproof social choice function. This is not the case because the SCF f^* , as defined in Remark 6, satisfies near unanimity and therefore represents a counterexample. This leads to the intriguing question on the minimal value l such that all groups of l voters are nominating for every anonymous SCF that satisfies Pareto-optimality and strategyproofness. An upper bound for this problem is provided by Lemma 1, which shows that $l \leq \lceil \frac{n+1}{2} \rceil$.

3.3 Non-Imposing SCFs

Finally, we consider the class of non-imposing SCFs. Recall that an SCF is non-imposing if every alternative is returned as the unique winner in some profile. Among the SCFs typically studied in social choice theory, there are only very few that fail to be non-imposing, e.g., SCFs that never return certain alternatives (such as constant SCFs) or SCFs that never return singletons. We will show a rather strong consequence of strategyproofness for non-imposing SCFs: every such function has to return the Condorcet loser in at least one preference profile and thus violates the Condorcet loser property. In the presence of neutrality, non-imposition can be seen as a decisiveness requirement. More precisely, if a neutral SCF fails non-imposition, it can never return a singleton choice set, which means that the SCF has to choose unreasonably large choice sets for many preference profiles. For instance, even if all voters agree on a best alternative, it cannot be chosen uniquely. Accordingly, the theorem identifies a tradeoff between decisiveness and the undesirable property of selecting Condorcet losers.

Similar to the proofs of Theorem 1 and Theorem 2, we start with a general lemma on strategyproof SCFs. This time, we investigate the relationship between vetoing and decisive groups of voters and show that these notions coincide for strategyproof and non-imposing SCFs.

Lemma 2. *Let f denote a strategyproof and non-imposing SCF for $m \geq 2$ alternatives. A group of voters I with $\emptyset \subsetneq I \subseteq N$ is vetoing for f if and only if it is decisive for f .*

Proof. Let f denote a strategyproof and non-imposing SCF and consider a non-empty set of voters $I \subseteq N$. First, observe that if I is decisive for f , then it is also vetoing: if all voters in I agree on the same preference relation in a profile R and report an alternative x as their unique last choice, then decisiveness requires that $f(R) \subseteq T_i(R)$ for $i \in I$. This means that $x \notin f(R)$ because $x \notin T_i(R)$, which shows that the group I is also vetoing for f .

Next, suppose that I is vetoing for f . If $m = 2$, every vetoing group of voters is decisive because such a group can determine the choice set by vetoing out an alternative. Thus, we focus in the sequel on the case $m \geq 3$ and show that the set I is decisive for f . We derive this claim in three steps: first, we show that $f(R) = \{x\}$ if all voters report x as their unique choice. Next, we prove that $f(R) = \{x\}$ is also true if only the voters $i \in I$ report x as their best choice. Finally, we infer from this insight that the set I is decisive for f .

Step 1: As a first step, we show that $f(R) = \{x\}$ for all alternatives $x \in A$ and preference profiles R such that all voters report x as unique top choice. For proving this claim, consider an arbitrary alternative a and a preference profile R^1 such that all voters uniquely top-rank a in R^1 . Since a and R^1 are chosen arbitrarily, the claim follows by showing that $f(R^1) = \{a\}$. Note for this that there is a profile R^2 such that $f(R^2) = \{a\}$ because f is non-imposing. As next step, we derive R^1 from R^2 by sequentially replacing the preference relation of all voters $i \in N$ with \succsim_i^1 . In more detail, consider the sequence $R^{2,0}, \dots, R^{2,n}$ such that $R^{2,0} = R^2$, $R^{2,n} = R^1$, and $R^{2,i}$ evolves out of $R^{2,i-1}$ by replacing $\succsim_i^{2,i-1}$ with \succsim_i^1 . Now, if $f(R^{2,i-1}) = \{a\}$ for some $i \in \{1, \dots, n\}$, *SP2* implies that $f(R^{2,i}) = \{a\}$ because $T_i(R^{2,i}) = \{a\}$. Since $f(R^2) = \{a\}$, we can repeatedly use this argument to derive that $f(R^1) = \{a\}$, which proves this step.

Step 2: Our next goal is to show that $f(R) = \{x\}$ for all alternatives $x \in A$ and profiles R such that all voters in I report x as their unique top choice. Hence, we suppose that $I \subsetneq N$; otherwise, we can directly proceed with Step 3. Just as in the last step, consider an arbitrary alternative a . Subsequently, we will show that $f(R^3) = \{a\}$ for the following profile R^3 .

$$R^3: \quad I: a, A \setminus \{a\} \quad N \setminus I: A \setminus \{a\}, a$$

This insight suffices to prove Step 2 since *SP2* allows the voters in I to reorder all alternatives in $A \setminus \{a\}$ arbitrarily without affecting the choice set, and *SP1* allows the voters in $N \setminus I$ to reorder all alternatives without affecting the choice set.

For proving the claim on R^3 , let $A = \{a, a_1, \dots, a_{m-1}\}$ denote the set of alternatives and, given $l \in \{1, \dots, m-1\}$ and a set of alternatives $X \subseteq A \setminus \{a, a_l\}$, define the profiles $R^{l,X}$ as shown below.

$$R^{l,X}: \quad I: a, A \setminus \{a\} \quad N \setminus I: a_l, \{a\} \cup X, A \setminus (X \cup \{a, a_l\})$$

In particular, the preference profiles $R^{l,\emptyset}$ and $R^{l,A \setminus \{a, a_l\}}$ are defined as follows.

$$\begin{aligned} R^{l,\emptyset}: & \quad I: a, A \setminus \{a\} & N \setminus I: a_l, a, A \setminus \{a, a_l\} \\ R^{l,A \setminus \{a, a_l\}}: & \quad I: a, A \setminus \{a\} & N \setminus I: a_l, A \setminus \{a_l\} \end{aligned}$$

Note that strategyproofness implies that, if $f(R^{l,A \setminus \{a, a_l\}}) = \{a\}$ for all $l \in \{1, \dots, m-1\}$, then $f(R^3) = \{a\}$. The reason for this is the following: starting at an arbitrary profile $R^{l,A \setminus \{a, a_l\}}$, we can derive the profile R^3 by letting the voters $i \in N \setminus I$ change their preference relation one after another to \succsim_i^3 . For each step, *SP1* implies that a subset of $A \setminus \{a_l\}$ needs to be chosen as otherwise, the deviation is a manipulation. Hence, we derive for every $l \in \{1, \dots, m-1\}$ that $a_l \notin f(R^3)$, which means that $f(R^3) = \{a\}$.

Subsequently, we will prove by induction on $z = |X|$ that $f(R^{l,X}) = \{a\}$ for all $l \in \{1, \dots, m-1\}$ and $X \subseteq A \setminus \{a, a_l\}$. Because of our previous insights, this completes the proof of Step 2. Two observations are central for the subsequent argument: firstly, our argument is closed under renaming alternatives in $A \setminus \{a\}$. This means that if we can show that $f(R^{l,X}) = \{a\}$ for some $l \in \{1, \dots, m-1\}$ and $X \subseteq A \setminus \{a, a_l\}$ with $|X| = z$, this result holds for all $l' \in \{1, \dots, m-1\}$ and subsets of $A \setminus \{a, a_{l'}\}$ with size z . Secondly, we can ensure that any alternative $a_l \in \{a_1, \dots, a_{m-1}\}$ is unchosen by letting the voters $i \in I$ report it as their unique bottom choice. Furthermore, this step does not change the choice set if a is the unique winner because of *SP2*.

First, we prove the base case $z = 0$, i.e., we show that $f(R^{l,\emptyset}) = \{a\}$ for all $l \in \{1, \dots, m-1\}$. Consider for this the following profiles and note that we display again $R^{l,\emptyset}$ so that all relevant profiles are shown.

$$\begin{aligned} \hat{R}^{l,\emptyset}: & \quad I: a, A \setminus \{a, a_l\}, a_l & N \setminus I: a, a_l, A \setminus \{a, a_l\} \\ \tilde{R}^{l,\emptyset}: & \quad I: a, A \setminus \{a, a_l\}, a_l & N \setminus I: a_l, a, A \setminus \{a, a_l\} \\ R^{l,\emptyset}: & \quad I: a, A \setminus \{a\} & N \setminus I: a_l, a, A \setminus \{a, a_l\} \end{aligned}$$

Recall that, by Step 1, $f(R) = \{a\}$ if all voters $i \in N$ uniquely top-rank a . Consequently, it holds that $f(\hat{R}^{l,\emptyset}) = \{a\}$. Next, observe that a_l is the uniquely least preferred alternative of the voters in I . Now, let the voters in $N \setminus I$ swap a and a_l one after another to derive $\tilde{R}^{l,\emptyset}$. For each step, a_l cannot be chosen because the voters in I veto it. Hence, if a is chosen before the deviation of a voter $i \in N \setminus I$, it follows that a needs to be chosen afterwards;

otherwise, a set X with $\{a\} \subsetneq X \subseteq A \setminus \{a_l\}$ is chosen and thus, voter i can manipulate by reverting this modification. Since $f(\hat{R}^{l,\emptyset}) = \{a\}$, we infer therefore that $f(\tilde{R}^{l,\emptyset}) = \{a\}$. Finally, note that $\tilde{R}^{l,\emptyset}$ only differs from $R^{l,\emptyset}$ in the preferences of the voters in I on the alternatives $A \setminus \{a\}$. Since *SP2* allows us to reorder the preferences of these voters on $A \setminus \{a\}$ arbitrarily without affecting the choice set, we derive that $f(R^{l,\emptyset}) = \{a\}$ for all $l \in \{1, \dots, m-1\}$.

Next, we focus on the induction step, i.e., we assume that $f(R^{l,X}) = \{a\}$ for all $l \in \{1, \dots, m-1\}$ and all $X \subseteq A \setminus \{a, a_j\}$ with $|X| = z-1$ and show that the same is true for all sets X' of size z . Recall for this that the derivation of $f(R^{l,X'}) = \{a\}$ is independent of the naming of the alternatives in $A \setminus \{a\}$, and thus, it suffices to show that $f(R^{z+1,\{a_1, \dots, a_z\}}) = \{a\}$. For this, let $Z = \{a_1, \dots, a_z\}$, $Z_{+a} = Z \cup \{a\}$, and $Z_{-l} = Z \setminus \{a_l\}$ for every $l \in \{1, \dots, z\}$, and consider the following profiles, where $l \in \{1, \dots, z\}$.

$$\begin{array}{lll} R^{l,Z_{-l}}: & I: a, A \setminus \{a\} & N \setminus I: a_l, Z_{-l} \cup \{a\}, A \setminus Z_{+a} \\ \hat{R}^{l,Z_{-l}}: & I: a, A \setminus \{a, a_{z+1}\}, a_{z+1} & N \setminus I: a_l, Z_{-l} \cup \{a\}, A \setminus Z_{+a} \\ \tilde{R}^{z+1,Z}: & I: a, A \setminus \{a, a_{z+1}\}, a_{z+1} & N \setminus I: a_{z+1}, Z_{+a}, A \setminus (Z \cup \{a, a_{z+1}\}) \\ R^{z+1,Z}: & I: a, A \setminus \{a\} & N \setminus I: a_{z+1}, Z_{+a}, A \setminus (Z \cup \{a, a_{z+1}\}) \end{array}$$

Now, consider an arbitrary $l \in \{1, \dots, z\}$ and note that our induction hypothesis implies that $f(R^{l,Z_{-l}}) = \{a\}$ since $|Z_{-l}| = z-1$. We derive $\hat{R}^{l,Z_{-l}}$ from $R^{l,Z_{-l}}$ by letting the voters $i \in I$ sequentially change their preference relations such that a_{z+1} is their uniquely least preferred alternative. For each step, *SP2* shows that the choice set is not allowed to change and thus, we infer that $f(\hat{R}^{l,Z_{-l}}) = \{a\}$. In particular, since all voters in I report a_{z+1} as their unique least preferred alternative, it follows that $a_{z+1} \notin f(\hat{R}^{l,Z_{-l}})$ regardless of the preference relations of the voters in $N \setminus I$.

Hence, we derive the profile $\tilde{R}^{z+1,Z}$ from $\hat{R}^{l,Z_{-l}}$ by replacing the preference relations of the voters in $N \setminus I$ one after another with $a_{z+1}, Z_{+a}, A \setminus (Z_{+a} \cup \{a_{z+1}\})$. Next, we will investigate one of these steps in detail and thus, let R^l and \bar{R}^l denote two consecutive profiles in the derivation of $\tilde{R}^{z+1,Z}$ from $\hat{R}^{l,Z_{-l}}$. Moreover, let i denote the voter whose preference relation is different in R^l and \bar{R}^l . First, note that $a_{z+1} \notin f(R^l)$ and $a_{z+1} \notin f(\bar{R}^l)$ because this alternative is vetoed out. Thus, if $f(R^l) \subseteq Z_{+a} \setminus \{a_l\}$, then $f(\bar{R}^l) \subseteq Z_{+a}$ because voter i can manipulate f by switching from \bar{R}^l to R^l otherwise. Additionally, a_l cannot be chosen in \bar{R}^l ; otherwise, voter i can manipulate f by deviating from R^l to \bar{R}^l . This means that if $f(R^l) \subseteq Z_{+a} \setminus \{a_l\}$, then $f(\bar{R}^l) \subseteq Z_{+a} \setminus \{a_l\}$. Since $f(\hat{R}^{l,Z_{-l}}) = \{a\}$, it follows from a repeated application of this argument that $f(\tilde{R}^{z+1,Z}) \subseteq Z_{+a} \setminus \{a_l\}$. Finally, note that we can apply this argument for every $l \in \{1, \dots, z\}$. This entails that $f(\tilde{R}^{z+1,Z}) \subseteq Z_{+a} \setminus Z$, i.e., $f(\tilde{R}^{z+1,Z}) = \{a\}$.

As last step, we derive $R^{z+1,Z}$ from $\tilde{R}^{z+1,Z}$ by reordering the preference relations of the voters $i \in I$. Because a stays their best alternative, it follows from a repeated application of *SP2* that $f(R^{z+1,Z}) = \{a\}$. Since the argument is closed under renaming alternatives in $A \setminus \{a\}$, this proves the induction step.

Step 3: It remains to show that the set of voters I is indeed decisive. Hence, consider an arbitrary profile R in which all voters $i \in I$ report the same preference relation. We need to show that $f(R) \subseteq T_i(R)$ for all $i \in I$. For this, let a denote an alternative in $T_i(R)$ for some voter $i \in I$, and let R' denote a profile such that all voters in I report a as their

uniquely best alternative, and all voters in $N \setminus I$ report the same preference relation as in R . By Step 2, it follows that $f(R') = \{a\}$. Next, we let the voters $i \in I$ revert one after another back to \succsim_i . Since $a \in T_i(R)$ and $\succsim_i = \succsim_j$ for all $i, j \in I$, it follows from a repeated application of *SP2* that $f(R) \subseteq T_i(R)$. This proves that every vetoing group is also decisive for f . \square

Lemma 2 has several interesting consequences. First of all, it shows that the notions of vetoing and decisive groups are equivalent for strategyproof SCFs that satisfy non-imposition. Consequently, no strategyproof, non-imposing, and non-dictatorial SCF can have a vetoer. Furthermore, our lemma entails that there cannot be two disjoint vetoing groups of voters for such SCFs. The reason for this is that both sets need to be decisive for such an SCF, but there cannot be two disjoint decisive groups. In particular, this means that no group of voters I with $|I| \leq \frac{n}{2}$ can be vetoing for an anonymous, strategyproof, and non-imposing SCF.

Furthermore, Lemma 2 shows that every group of voters I with $|I| > \frac{n}{2}$ is decisive for a strategyproof SCF that satisfies non-imposition and the Condorcet loser property. The reason for this is that the Condorcet loser property entails that such groups are vetoing. Next, we use this insight to show that there is no strategyproof SCF that satisfies the Condorcet loser property and non-imposition. Note that we present here a simplified proof with completely indifferent voters. In the appendix, we give a more involved proof which avoids such artificial voters.

Theorem 3. *There is no strategyproof SCF that satisfies the Condorcet loser property and non-imposition if $m \geq 3$ and $n \geq 4$.*

Proof. We prove the statement by induction over $n \geq 4$.

Induction basis: Assume for contradiction that there is a strategyproof SCF f for $n = 4$ voters and $m \geq 3$ alternatives that satisfies the Condorcet loser property and non-imposition. First, consider the profiles R^1 to R^5 shown below.

R^1 :	1: a, c, X, b	2: a, b, X, c	3: a, b, X, c	4: b, X, c, a
R^2 :	1: $\{a, c\}, X+b$	2: a, b, X, c	3: a, b, X, c	4: b, X, c, a
R^3 :	1: $\{a, c\}, X+b$	2: a, c, X, b	3: a, b, X, c	4: $c, X, \{a, b\}$
R^4 :	1: $\{a, c\}, X+b$	2: a, c, X, b	3: b, a, X, c	4: $c, X, \{a, b\}$
R^5 :	1: $\{a, c\}, X+b$	2: a, c, X, b	3: b, a, X, c	4: c, b, X, a

Lemma 2 shows that $f(R^1) = \{a\}$ since every group of 3 voters needs to be decisive for f . Moreover, c is the Condorcet loser in R^1 , even if voter 1 is indifferent between a and c . Thus, we replace next the preference relation of voter 1 with $\{a, c\}, X+b$, where $X+b = X \cup \{b\}$, to derive the profile R^2 . *SP2* implies that $f(R^2) \subseteq \{a, c\}$ because otherwise voter 1 can manipulate by reverting back to R^1 . Moreover, $c \notin f(R^2)$ due to the Condorcet loser property and we hence deduce that $f(R^2) = \{a\}$. As second step, we let voter 2 change his preference relation to a, c, X, b and voter 4 change his preference to $c, X, \{a, b\}$ in order to make b into the Condorcet loser. It follows from *SP1* and *SP2* that the choice set does not change during these steps since the current winner a is the best alternative of voter 2 after the manipulation and the worst alternative of voter 4 before the manipulation. Hence,

these steps result in the profile R^3 with $f(R^3) = \{a\}$. Furthermore, observe that b is the Condorcet loser in R^3 , even if voter 3 swaps a and b . Thus, we derive the profile R^4 from R^3 by applying this modification. The Condorcet loser property requires that $b \notin f(R^4)$, which entails, in turn, that $f(R^4) = \{a\}$; otherwise, voter 3 can manipulate f by deviating from R^4 to R^3 . As last step, we let voter 4 change his preference relation to c, b, X, a to derive the profile R^5 . Since $f(R^4) = \{a\} \subseteq B_4(R^4)$, it follows from *SP1* that $f(R^5) \subseteq \{a, b\}$.

Next, observe that we can apply analogous steps for profiles that are symmetric with respect to the voters or alternatives. Thus, we infer for the choice sets of the profiles R^6 , R^7 , and R^8 that $f(R^6) \subseteq \{a, c\}$, $f(R^7) \subseteq \{a, b\}$, and $f(R^8) \subseteq \{b, c\}$.

R^6 :	1: $\{b, c\}, X+a$	2: a, c, X, b	3: b, a, X, c	4: c, b, X, a
R^7 :	1: a, b, X, c	2: c, a, X, b	3: $\{b, c\}, X+a$	4: b, c, X, a
R^8 :	1: a, b, X, c	2: c, a, X, b	3: $\{a, c\}, X+b$	4: b, c, X, a

Note that, if $b \in f(R^5)$, then voter 1 can manipulate by switching to R^6 as $f(R^6) \subseteq \{a, c\}$. Hence, we derive that $f(R^5) = \{a\}$. By a symmetric argument for R^7 and R^8 , it follows that $f(R^7) = \{b\}$.

Finally, consider the profile R^9 shown below.

R^9 :	1: a, b, X, c	2: a, b, X, c ,	3: b, a, X, c	4: b, a, X, c
---------	-----------------	-------------------	-----------------	-----------------

We can derive the profile R^9 from R^5 and R^7 . In more detail, we obtain R^9 from R^5 by replacing the preference relations of voters 1 and 2 with a, b, X, c and the preference relation of voter 4 with b, a, X, c . If we apply these steps one after another, *SP1* and *SP2* imply that $f(R^9) = \{a\}$. On the other hand, we obtain R^9 from R^7 by replacing the preference relation of voters 3 and 4 with b, a, X, c and the preference relation of voter 2 with a, b, X, c and obtain $f(R^9) = \{b\}$ by an analogous argument. This is a contradiction since $f(R^9) = \{a\}$ and $f(R^9) = \{b\}$ cannot be simultaneously true, which shows that there is no strategyproof SCF that satisfies non-imposition and the Condorcet loser property if $n = 4$ and $m \geq 3$.

Induction step: Assume for contradiction that there is a strategyproof SCF f for $n > 4$ voters and $m \geq 3$ alternatives that satisfies non-imposition and the Condorcet loser property. Consider the following SCF g for $n - 1$ voters and m alternatives: given a profile R on $n - 1$ voters, g adds a new voter who is indifferent between all alternatives to derive a profile R' on n voters and returns $g(R) = f(R')$. Clearly, g is strategyproof and inherits the Condorcet loser property from f . Furthermore, Lemma 2 shows that g is non-imposing because every set of $n - 1 > \frac{n}{2}$ voters is decisive for f . Hence, we can construct a strategyproof SCF for $n - 1$ voters that satisfies the Condorcet loser property and non-imposition if there is such an SCF for n voters. Since our induction hypothesis states that no such SCF exists, we derive from the contraposition of this implication that there is no SCF satisfying all required axioms for $n > 4$ voters. \square

Remark 11. The axioms used in Theorem 3 are independent of each other. An SCF that only violates the Condorcet loser property is the Pareto rule. The SCF that returns all alternatives except the Condorcet loser only violates non-imposition. The SCF that returns all Pareto-optimal alternatives except the Condorcet loser only violates strategyproofness. The bounds on n and m are also tight. The majority rule satisfies all axioms if $m = 2$, the Pareto rule satisfies all axioms if $n \leq 2$, and a computer-aided approach proves that there

is an SCF that satisfies all required axioms if $n = 3$ and $m = 3$. It is also possible to extend the SCF by hand to $m > 3$ alternatives, but the resulting SCFs are purely technical and we therefore do not define them explicitly.

Remark 12. Brandt (2015, Theorem 2) has shown that no Condorcet extension can be strategyproof if $m \geq 3$ and $n \geq 3m$. By replacing the Condorcet loser property and non-imposition with Condorcet-consistency, careful inspection of the proof of Theorem 3 reveals that Condorcet-consistency and strategyproofness are already incompatible if $m \geq 3$ and $n \geq 4$. In particular, observe for this that a is the Condorcet winner in the profile R^4 and thus, every Condorcet-consistent SCF satisfies $f(R^4) = \{a\}$. Departing from this insight, we can apply the same steps as in the proof of Theorem 3 since the Condorcet loser property is not used anymore.

Remark 13. Lemma 2 and Theorem 3 do not require the full power of non-imposition. For Lemma 2, the following weakening holds: if an alternative x can be uniquely chosen by a strategyproof SCF, then every vetoing group of voters I can ensure that x is the unique winner if they unanimously report it as their best alternative. For Theorem 3, we can weaken non-imposition to the requirement that at least three alternatives can be returned as unique winner.

Remark 14. A desirable strengthening of Theorem 3 would be to weaken the Condorcet loser property by only demanding that an alternative that is uniquely bottom-ranked by a majority of voters should not be chosen. A computer analysis has shown that this property is compatible with non-imposition and strategyproofness when $m \leq 3$ and $n \leq 6$, even if we additionally impose anonymity. We nevertheless believe that there may be an impossibility for larger values of m and n .

4. Consequences for Randomized Social Choice

So far, we have discussed our theorems in the context of set-valued social choice, but they also have consequences for randomized social choice, which is concerned with the study of *social decision schemes* (SDSs), i.e., functions that map preference profiles to *lotteries* (i.e., probability distributions) over the alternatives. Since the notions of rank-basedness and support-basedness are independent of the type of the output of the function and merely define an equivalence relation over preference profiles, they can be straightforwardly extended to SDSs. For our other axioms, we consider variants in randomized social choice based on the support of lotteries, i.e., the set of alternatives with positive probability. For example, Pareto-optimality and the Condorcet loser property require that Pareto-dominated alternatives and Condorcet losers are always assigned probability 0. In other words, Pareto-optimality demands that Pareto-dominated alternatives are not in the support of any chosen lottery, a condition that is usually referred to as *ex post efficiency*. Similarly, an SDS satisfies the Condorcet loser property if the Condorcet loser is never in the support of any chosen lottery. Next, an SDS satisfies non-imposition if every alternative is chosen with probability 1 for some profile. Finally, Kelly-strategyproofness translates to the notion of *DD-strategyproofness* (Brandt, 2017). To this end, we say that, given a preference relation \succsim , a lottery p *deterministically dominates* a lottery q if and only if $\text{supp}(p) \succsim \text{supp}(q)$.

Then, an SDS f is called *DD-strategyproof* if $f(R')$ does not strictly deterministically dominate $f(R)$ for all voters $i \in N$ and all profiles R, R' such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$. Note that *DD-strategyproofness* is weaker than most strategyproofness notions considered in the literature (see, e.g., Gibbard, 1977; Brandt, 2017; Aziz et al., 2018; Brandl et al., 2018). In particular, it is weaker than weak *SD-strategyproofness* as used by Brandl et al. (2018) to prove a rather sweeping impossibility: no anonymous and neutral SDS is weakly *SD-strategyproof* and *SD-efficient* if $n \geq 4$ and $m \geq 4$.

Based on these axioms for SDSs, we can translate our results to the randomized context. Note for this that we can turn every SDS f into an SCF g by returning the support of $f(R)$ instead of the lottery itself. Moreover, it is easy to verify that all our axioms carry over from the SDS f to the SCF g because they are only defined based on the support. Therefore, we derive the following corollaries.

Corollary 2. *There is no rank-based SDS that satisfies ex post efficiency and DD-strategyproofness if $m \geq 4$ and $n \geq 3$, or if $m \geq 5$ and $n \geq 2$.*

Corollary 3. *Every support-based SDS that satisfies ex post efficiency and DD-strategyproofness assigns positive probability to at least one most preferred alternative of every voter if $m \geq 3$.*

Corollary 4. *There is no SDS that satisfies the Condorcet loser property, non-imposition, and DD-strategyproofness if $m \geq 3$ and $n \geq 4$.*

Corollary 2 can be seen as a strengthening of the impossibility of Brandl et al. (2018) for the class of rank-based SDSs as we require both a weaker strategyproofness notion and a weaker efficiency notion. However, our result only holds for rank-based SDSs rather than the more general class of anonymous SDSs. Corollary 3 implies that at least one of the most preferred alternatives of every voter receives positive probability, a property that is known as *positive share* in the context of dichotomous preferences (Bogomolnaia et al., 2005; Brandl et al., 2021). When strengthening Pareto-optimality to *SD-efficiency*, we derive an impossibility (see Remark 7) and this impossibility can be interpreted as a strengthening of the result by Brandl et al. (2018) for the subclass of support-based SDSs. Finally, Corollary 4 is unrelated to the aforementioned results as the Condorcet loser property is independent of the other axioms. This result can be interpreted as a new far-reaching impossibility for SDSs.

5. Conclusion

We have studied which SCFs satisfy strategyproofness according to Kelly’s preference extension and obtained results for three broad classes of SCFs. A common theme of our results is that strategyproofness entails that potentially “bad” alternatives need to be chosen. In particular, we have shown that (i) every strategyproof rank-based SCF returns a Pareto-dominated alternative in at least one profile, (ii) every strategyproof support-based SCF that satisfies Pareto-optimality returns at least one most preferred alternative of every voter, and (iii) every strategyproof non-imposing SCF returns the Condorcet loser in at least one profile. All of these impossibilities rely on general insights about decisive, nominating, and vetoing groups of voters for strategyproof SCFs. Taken together, our results

show that there is only room for rather indecisive strategyproof SCFs such as the Pareto rule, the omninomination rule, the SCF that returns all top-ranked alternatives that are Pareto-optimal, or the SCF that returns all alternatives except Condorcet losers. Furthermore, since we require sufficiently weak axioms, our results directly extend to randomized social choice and we therefore derive three impossibilities as corollaries for this setting.

In comparison to other results on the strategyproofness of set-valued SCFs, we employ a very weak notion of strategyproofness. In particular, our notion of strategyproofness is weaker than those used by Duggan and Schwartz (2000), Barberà et al. (2001), Ching and Zhou (2002), Rodríguez-Álvarez (2007), and Sato (2008). This is possible because we consider the more general domain of weak preferences, which explicitly allows for ties. Interestingly, all proofs except that of Claim 1 in Theorem 1 can be transferred to the domain of strict preferences by carefully breaking ties and replacing Kelly-strategyproofness with the significantly stronger strategyproofness notion introduced by Duggan and Schwartz (2000). While the resulting theorems are covered by the Duggan-Schwartz impossibility, this raises intriguing questions concerning the relationship between strategyproofness results for weak and strict preferences.

In contrast to previous impossibilities for Kelly’s preference extension (Brandl et al., 2019; Brandt et al., 2022), our proofs do not rely on the availability of artificial voters who are completely indifferent between all alternatives. Moreover, the results are tight in the sense that they cease to hold if we remove an axiom, reduce the number of alternatives or voters, weaken the notion of strategyproofness, or require strict preferences. For example, the essential set (Dutta & Laslier, 1999; Laslier, 2000) and a handful of other support-based Condorcet extensions satisfy strategyproofness if preferences are strict and participation for unrestricted preferences (Brandt, 2015; Brandl et al., 2019). Our results thus provide important insights on when and why strategyproofness can be attained.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1. We thank the anonymous reviewers for helpful comments. A preliminary version of this article appeared in the Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (May 2021). Results from this article were presented at the 8th International Workshop on Computational Social Choice (June 2021).

Appendix A. Alternative Proof of Theorem 3

Subsequently, we discuss an alternative proof for Theorem 3 which does not rely on completely indifferent voters.

Theorem 3. *There is no strategyproof SCF that satisfies the Condorcet loser property and non-imposition if $m \geq 3$ and $n \geq 4$.*

Proof. Assume for contradiction that there is a non-imposing SCF f that satisfies the Condorcet loser property and strategyproofness for $n \geq 4$ voters and $m \geq 3$ alternatives. Since the Condorcet loser property strongly depends on the parity of the number of voters,

we proceed with a case distinction on n . For both cases, it is important that Lemma 2 and *SP2* show that an alternative is the unique winner if it is uniquely top-ranked by at least $l = \lceil \frac{n+1}{2} \rceil$ voters. This follows from the observation that every set $I \subseteq N$ with $|I| \geq l$ is vetoing for f . Thus, Lemma 2 shows that such sets are also decisive, i.e., f needs to choose an alternative x as a unique winner if all voters in I report the same preference relation with x as unique top choice. Finally, *SP2* allows to reorder the alternatives $y \in A \setminus \{a\}$ without affecting the choice set, which proves this auxiliary claim.

Case 1: n is odd

First, assume that f is defined for an odd number of voters $n \geq 4$, and consider the following profiles, where $X = A \setminus \{a, b, c\}$.

$$\begin{aligned}
 R^1: \quad & 1: a, b, X, c & [2 \dots l]: a, c, X, b & [l+1 \dots n]: b, X, \{a, c\} \\
 R^2: \quad & 1: a, b, X, c & [2 \dots l-1]: a, c, X, b & l: c, a, X, b & [l+1 \dots n]: b, X, \{a, c\} \\
 R^3: \quad & 1: a, b, X, c & [2 \dots l-1]: a, c, X, b & l: c, a, X, b & [l+1 \dots n]: b, c, X, a \\
 R^4: \quad & [1 \dots l-1]: a, c, X, b & l: c, a, X, b & [l+1 \dots n]: c, b, X, a \\
 R^5: \quad & [1 \dots l-1]: a, c, X, b & l: c, a, X, b & [l+1 \dots n]: b, c, X, a
 \end{aligned}$$

First, note that Lemma 2 and *SP2* show that $f(R^1) = \{a\}$. Moreover, c is the Condorcet loser in R^1 because every voter prefers a weakly to c and all voters in $[l+1 \dots n]$ and voter 1 prefer all alternatives in $A \setminus \{a, c\}$ strictly to c . Alternative c even remains the Condorcet loser if voter l swaps a and c . Hence, let R^2 denote the resulting profile and observe that $c \notin f(R^2)$ because of the Condorcet loser property. In turn, strategyproofness implies that $f(R^2) = \{a\}$ if $c \notin f(R^1)$; otherwise, voter l can manipulate by reverting back to R^1 as he prefers $\{a\}$ to every other subset of $A \setminus \{c\}$.

As the next step, we subsequently replace the preference relations of the voters $i \in [l+1 \dots n]$ with b, c, X, a . *SP1* implies for each of these steps that a subset of $\{a, c\}$ is chosen if it has been chosen before the step. Since $f(R^2) = \{a\}$, we deduce that this process results in a profile R^3 with $f(R^3) \subseteq \{a, c\}$. Moreover, $f(R^3) \neq \{c\}$ as otherwise voter 1 can manipulate by swapping a and b : after this step, b is uniquely top-ranked by more than half of the voters and therefore Lemma 2 and *SP2* imply that it is the unique winner. Since voter 1 prefers $\{b\}$ to $\{c\}$, strategyproofness requires that $f(R^3) \in \{\{a\}, \{a, c\}\}$.

Next, we discuss another derivation for $f(R^3)$ which proves that $f(R^3) \notin \{\{a\}, \{a, c\}\}$. For this, consider the profile R^4 and note that $f(R^4) = \{c\}$ because more than half of the voters report c as their favorite choice. Moreover, b is the Condorcet loser in R^4 as it is uniquely bottom-ranked by the voters $i \in [1 \dots l]$. This even holds if the voters in $i \in [l+1 \dots n]$ change their preference. Thus, we let these voters swap b and c , and the Condorcet loser property always implies for the resulting profile that b is not chosen. Just as for R^3 , strategyproofness implies then that c remains the unique winner after every step because otherwise, a voter can manipulate by reverting this modification. Thus, this process results in the profile R^5 with $f(R^5) = \{c\}$.

Finally, we derive the profile R^3 from R^5 by replacing the preference relation of voter 1 with a, b, X, c . Strategyproofness from R^5 to R^3 implies that $f(R^3) \neq \{a\}$ and $f(R^3) \neq \{a, c\}$ as otherwise, voter 1 can manipulate by deviating from R^5 to R^3 . This is in conflict

with our previous observation and hence, there is no strategyproof SCF for odd $n \geq 5$ that satisfies non-imposition and the Condorcet loser property.

Case 2: n is even

As second case, we assume that f is defined for an even number of voters $n \geq 4$. First, note that the induction basis in the proof of Theorem 3 shows that no strategyproof and non-imposing SCF for $m \geq 3$ alternatives and $n = 4$ voters satisfies the Condorcet loser property, even if we forbid completely indifferent voters. The reason for this is that no such voters are required for the proof. Subsequently, we demonstrate how we can reduce the case with $n > 4$ voters to the case with $n = 4$ voters. Hence, assume that there is a strategyproof SCF f for $n > 4$ voters, n even, and $m \geq 3$ alternatives that satisfies the Condorcet loser property and non-imposition. We use this SCF f to define another SCF g for $n = 4$ voters as follows (where $X = A \setminus \{a, b, c\}$): given a profile R on 4 voters, g adds $(n - 4)/2$ voters whose preference relation is c, X, b, a and $(n - 4)/2$ whose preference relation is a, b, X, c . Then, g returns the choice set of f on the resulting profile R' , i.e., $g(R) = f(R')$. Subsequently, we prove that g satisfies all criteria required for deriving the impossibility in the 4 alternative case. As a consequence, g cannot exist, which implies that f also violates one of the required axioms.

First, note that g inherits the strategyproofness of f because any manipulation of g is by definition also a manipulation of f . Moreover, g cannot return the Condorcet loser because the Condorcet loser in a profile R on 4 voters is also the Condorcet loser in the profile R' that is obtained after g adds the $n - 4$ extra voters. The reason for this is that the preference of the first half of these $n - 4$ voters is inverse to the other half. In more detail, adding these $n - 4$ voters increases every support $s_{xy}(R)$ by $(n - 4)/2$ if $x \in \{a, b, c\}$ or $y \in \{a, b, c\}$ and the supports $s_{xy}(R)$ with $x, y \in X$ do not change at all. Consequently, the Condorcet loser does not change and g inherits the Condorcet loser property from f .

The last axiom required for the proof of Theorem 3 is non-imposition. However, a close inspection of the proof shows that we actually do not need full non-imposition, but it is sufficient if there are three alternatives that can be chosen uniquely. Hence, we only show that g can return a , b , and c as unique winner. For a and c , this follows from Lemma 2 because g adds $(n - 4)/2$ voters with preference a, b, X, c and $(n - 4)/2$ voters with preference c, X, b, a to derive the input profile R' for f . Hence, if all of the four original voters report a, b, X, c , then $f(R') = \{a\}$, and if all four voters report c, X, b, a , then $f(R') = \{c\}$. The reason for this is that in both cases, the corresponding alternative is uniquely top-ranked by more than half of the voters in R' and Lemma 2 and *SP2* thus show that this alternative needs to be chosen uniquely.

A slightly more complicated argument is required for showing that g returns can return b as unique winner. Thus, consider the profiles R and R' shown below.

$$R: \quad [1 \dots 4]: b, X, c, a \quad [5 \dots 2+n/2]: b, a, X, c \quad [3+n/2 \dots n]: c, X, b, a$$

$$R': \quad [1 \dots 4]: b, X, c, a \quad [5 \dots 2+n/2]: a, b, X, c \quad [3+n/2 \dots n]: c, X, b, a$$

First, note that Lemma 2 and *SP2* imply that $f(R) = \{b\}$. Moreover, alternative a is uniquely bottom-ranked by all voters in $[1 \dots 4] \cup [3 + n/2 \dots n]$ and it thus is the Condorcet loser. This is also true if the voters $i \in [5 \dots 2+n/2]$ swap a and b one after another. Hence, the Condorcet loser property implies that a is not chosen after these swaps and strategyproofness entails then that b is still the unique winner since all voters in $[5 \dots 2+n/2]$

prefer $\{b\}$ to every other subset of $A \setminus \{a\}$. This means that $f(R') = \{b\}$. Finally, note that $g(R'') = f(R') = \{b\}$ for the profile $R'' = (\succsim'_1, \succsim'_2, \succsim'_3, \succsim'_4)$ because the preferences of the last $n - 4$ voters are equal to those used by g to extend profiles consisting of 4 voters to profiles for n voters. This proves that g can also return b as unique winner, and thus, the proof in the main body shows that g cannot exist. On the other hand, we have shown that, if there is a strategyproof SCF for an even number of voters $n > 4$ that satisfies the Condorcet loser property and non-imposition, g exists. By the contraposition of this implication, the impossibility generalizes to all even numbers of voters $n > 4$. \square

Appendix B. Examples for the Tightness of our Results

In this appendix, we discuss the SCFs that have been used to show that our results are tight. First, we deal with rank-basedness under strict preferences. Therefore, we consider the variant of the 2-plurality rule mentioned in Remark 5, which we call 2*-plurality. For introducing this rule, we define the plurality score $PL(a, R)$ of an alternative a in a profile R as the number of voters that top-rank alternative a in the profile R . Given a profile R , let a_R denote the alternative with the second highest plurality score. Then, the 2*-plurality rule, abbreviated by $2^*-PL(R)$, chooses precisely the alternatives x with $PL(x, R) \geq PL(a_R, R)$ and $PL(x, R) > 0$, i.e., $2^*-PL(R) = \{x \in A : PL(x, R) \geq PL(a_R, R) \wedge PL(x, R) > 0\}$.

Proposition 1. *For strict preferences, the 2*-plurality rule is rank-based, Pareto-optimal, and strategyproof, but no voter is a nominator if $m \geq 3$ and $n \geq 5$.*

Proof. First, note that 2*-plurality is by definition rank-based and it satisfies Pareto-optimality as it only returns alternatives that are top-ranked by some voters. This criterion entails Pareto-optimality as we assume strict preferences. Moreover, no voter is a nominator for 2*-plurality if there are at least 5 voters and at least 3 alternatives because the top-ranked alternative c of a voter can have plurality score 1 and two other alternatives may have plurality score 2 or more. Hence, it only remains to show that 2*-PL is strategyproof. We assume for contradiction that this is not the case, i.e., that there are preference profiles R and R' and a voter i such that $\succsim_j = \succsim_j'$ for all $j \in N \setminus \{i\}$ and $2^*-PL(R') \succ_i 2^*-PL(R)$. We proceed with a case distinction on whether voter i 's most preferred alternative in R , denoted by a , is chosen.

First, assume that $a \in 2^*-PL(R)$. This means that voter i can only manipulate if $2^*-PL(R') = \{a\}$ as otherwise, there is an alternative $x \in 2^*-PL(R')$ with $a \succ_i x$. Moreover, if $2^*-PL(R) = \{a\}$, voter i can also not manipulate as his best alternative is the unique winner. Hence, another alternative b is chosen by 2*-plurality, which implies that another voter reports b as his most preferred alternative in R . As a consequence, $PL(b, R') > 0$ and therefore, $2^*-PL(R') \neq \{a\}$ as 2*-plurality only returns a single winner if all voters report it as their best choice. Hence, no manipulation is possible in this case.

Next, assume that $a \notin 2^*-PL(R)$ and let b denote voter i 's best alternative in R' . Note that $a \neq b$ because the plurality scores, and therefore also the choice set of 2*-plurality, do not change otherwise. We use another case distinction with respect to the plurality score of b in R . First, assume that $PL(b, R) \geq PL(a_R, R) > 0$, which means that $b \in 2^*-PL(R)$. In particular, the claim that $PL(a_R, R) > 0$ is true as $a \notin 2^*-PL(R)$ but $PL(a, R) > 0$. This means also that 2*-plurality elects at least two alternatives in R ,

and we choose $c \in 2^*-PL(R) \setminus \{b\}$ as the alternative with the highest plurality score in $A \setminus \{b\}$. Now, if $PL(c, R) > PL(b, R)$, alternative b has the second highest plurality score in R , i.e., $PL(b, R) \geq PL(x, R)$ for all $x \in A \setminus \{b, c\}$. Since $PL(b, R') = PL(b, R) + 1$, $PL(a, R') = PL(a, R) - 1$, and $PL(x, R') = PL(x, R)$ for all $x \in A \setminus \{a, b\}$, it follows therefore that $PL(c, R') \geq PL(b, R')$ and $PL(b, R') > PL(x, R')$ for all $x \in A \setminus \{b, c\}$. Hence, $2^*-PL(R') = \{b, c\}$. On the other hand, if $PL(b, R) \geq PL(c, R)$, b is the alternative with the highest plurality score in R , and c the one with the second highest plurality score. Since $PL(b, R') = PL(b, R) + 1$, $PL(a, R') = PL(a, R) - 1$, and $PL(x, R') = PL(x, R)$ for all $x \in A \setminus \{a, b\}$, it follows that $PL(b, R') > PL(c, R') \geq PL(x, R')$ for all $x \in A \setminus \{b, c\}$, which shows that c is also in R' the alternative with the second highest plurality score. Hence, we derive also in this case that $\{b, c\} \subseteq 2^*-PL(R')$. Hence, we have in both cases that $\{b, c\} \subseteq 2^*-PL(R) \cap 2^*-PL(R')$, which contradicts that voter i benefits by deviating from R to R' because he is not indifferent between b and c , i.e., there are alternatives $x \in 2^*-PL(R)$, $y \in 2^*-PL(R')$ such that $x \succ_i y$.

Finally, assume that $PL(b, R) < PL(a_R, R)$ and note that this assumption entails that there are at least two alternatives with a higher plurality score than a and b , i.e., $PL(a_R, R) > PL(a, R)$ and $PL(a_R, R) > PL(b, R)$. Hence, $PL(b, R') = PL(b, R) + 1 \leq PL(a_R, R)$. This means that $PL(a_R, R) = PL(a_{R'}, R')$ as b has a plurality score of at most $PL(a_R, R)$ in R' . Since the plurality scores of all alternatives x with $PL(x, R) \geq PL(a_R, R)$ have not been affected by the manipulation, it follows that every alternative chosen in $2^*-PL(R)$ is also chosen after the manipulation, i.e., $2^*-PL(R) \subseteq 2^*-PL(R')$. Since $|2^*-PL(R)| \geq 2$, deviating from R to R' is no manipulation because we can find alternatives $x \in 2^*-PL(R)$, $y \in 2^*-PL(R) \subseteq 2^*-PL(R')$ such that $x \succ_i y$. Hence, no case allows for a manipulation, which means that 2^* -plurality is strategyproof for strict preferences. \square

Next, we consider Remark 2 in which we claim that the bounds on n and m in Theorem 1 are tight as the Pareto rule is rank-based for small values of n and m . We prove this statement subsequently.

Proposition 2. *The Pareto rule is rank-based, Pareto-optimal, and strategyproof if $m \leq 3$, or if $m \leq 4$ and $n \leq 2$.*

Proof. The Pareto rule is known to satisfy Pareto-optimality and strategyproofness, regardless of the number of alternatives or voters (see, e.g., Brandt et al., 2022). Hence, it only remains to show that it also satisfies rank-basedness under the restrictions on n and m , for which we use a case distinction.

Case 1: $m \leq 2$

For $m = 1$, rank-basedness is obviously no restriction, and if $m = 2$, the rank vector of the single alternative determines all preference relations (except that we do not know which voter submits which preference relation). If an alternative a is uniquely top-ranked by a voter, its rank vector contains a $(0, 1)$ entry, and thus, the rank vector of the other alternative b must contain a $(1, 1)$ entry. Similarly, if a has a $(0, 2)$ entry, a voter is indifferent between both alternatives and thus, b has also a $(0, 2)$. Finally, we can apply a symmetric argument to the first case if the rank vector a contains a $(1, 1)$ entry, and thus, we can reconstruct a unique profile (up to renaming the voters) given a rank matrix. Hence, the Pareto rule is rank-based if $m = 2$.

Case 2: $m = 3$

Next, we focus on the case that $m = 3$ and consider an arbitrary rank matrix Q . First note that Q can only have the following entries: $(0, 3)$, $(0, 2)$, $(1, 2)$, $(0, 1)$, $(1, 1)$, and $(2, 1)$. Many of these entries specify the preferences of the voters. For instance, the $(0, 3)$ entry entails that a voter is completely indifferent between all alternatives. Consequently, we can add a completely indifferent voter for every $(0, 3)$ entry in the rank vector of an alternative a , and remove all these entries from Q afterwards. Also, the $(0, 2)$ entries in the rank vector of a specify a lot of information: there must be a voter who top-ranks a and another alternative x and bottom-ranks the last alternative y uniquely. We use this observation to formulate a system of linear equations. Let therefore n_a , n_b , and n_c denote the number of $(0, 2)$ entries in the rank vector of the respective alternatives. Moreover, let x_{ab} , x_{ac} , and x_{bc} denote the number of voters who top-rank both alternatives in the index. The following equations must hold for every profile R with $r^*(R) = Q$.

$$\begin{aligned} n_a &= x_{ab} + x_{ac} \\ n_b &= x_{ab} + x_{bc} \\ n_c &= x_{bc} + x_{ac} \end{aligned}$$

It can easily be checked that the unique solution of this system of equations is $x_{ab} = \frac{n_a + n_b - n_c}{2}$, $x_{bc} = \frac{n_b + n_c - n_a}{2}$, and $x_{ac} = \frac{n_a + n_c - n_b}{2}$. Since this solution is unique, these entries determine the preference relations of several voters: for instance, there must be x_{ab} voters who are indifferent between a and b , and prefer both alternatives to c . A symmetric argument applies also for all $(1, 2)$ entries. Hence, we can now remove these entries from Q , as well as the corresponding $(2, 1)$ and $(0, 1)$ entries, to derive a reduced rank matrix.

After the last step, Q only consists of $(0, 1)$, $(1, 1)$, and $(2, 1)$ entries, which means that all remaining preferences are strict. Unfortunately, these entries do not necessarily entail a unique profile, but we can use all our observations so far to check for an arbitrary pair of alternatives a and b whether a Pareto-dominates b . For this, we first construct the preferences involving ties as explained before and check whether one of the voters prefers b strictly to a . If this is the case, a cannot Pareto-dominate b and we are done. Otherwise, we consider the remaining entries in Q . First, if Q is empty (i.e., there are no strict preference relations), we can check the Pareto-dominance by considering the preference profile constructed so far. Else, a Pareto-dominates b if and only if a has no $(2, 1)$ entry and b has no $(0, 1)$ entry. If a has an $(2, 1)$ entry or b has an $(0, 1)$ entry, then a is uniquely last-ranked or b is uniquely top-ranked by some voter, which prohibits that a Pareto-dominates b . Conversely, if none of these entries exist, then b has to be last-ranked whenever a is second-ranked, and thus a Pareto-dominates b . Since a and b were chosen arbitrary, we can check Pareto-dominance between alternatives only based on the rank matrix if $m = 3$, which shows that the Pareto rule is rank-based in this case.

Case 3: $m = 4$ and $n \leq 2$

Finally, we show that the Pareto rule is also rank-based if $m = 4$ and $n \leq 2$. First, if $n = 1$, it is trivial to compute the Pareto rule because only the top-ranked alternatives of the single voter are Pareto-optimal, and this information is contained in the rank matrix. Hence, we focus on the case that $n = 2$ and show that $PO(R) = PO(R')$ for all profiles

R, R' with $r^*(R) = r^*(R')$. Given a rank matrix Q , we can therefore compute the Pareto rule on an arbitrary profile R with $r^*(R) = Q$ as the outcome is independent of the choice of R . Hence, consider a profile R and assume that alternative b Pareto-dominates a in R . The result follows by proving that a is Pareto-dominated in all preference profiles R' with $r^*(R) = r^*(R')$.

For this, let $(s_{xi}, t_{xi}) = r(\succsim_i, x)$ denote the rank tuple of alternative x in the preference relation of voter i and note that $s_{xi} \leq s_{yi}$ if and only if $x \succsim_i y$ for all alternatives $x, y \in A$ and voters $i \in N$. We suppose subsequently that $s_{b1} \leq s_{b2}$; this is without loss of generality as we can just reorder the voters in our analysis. Next, note that the assumption that b Pareto-dominates a implies that $b \succsim_i a$ for all $i \in \{1, 2\}$ and that this preference is strict for at least one voter. This means equivalently that $s_{bi} \leq s_{ai}$ for all $i \in I$ and that this inequality is strict for at least one voter. Now, if $s_{b1} \leq s_{b2} \leq \min_{i \in \{1, 2\}} s_{ai}$, b Pareto-dominates a in all profiles R' with $r^*(R) = r^*(R')$ because $s_{bi} \leq s_{ai}$ for all $i \in N$ in all such profiles R' and one of these inequalities must be strict.

Hence, assume that $\min_{i \in \{1, 2\}} s_{ai} < s_{b2}$, which means that $s_{b1} \leq s_{a1} < s_{b2} \leq s_{a2}$ because b Pareto-dominates a in R . Next, consider a profile R' with $r^*(R') = r^*(R)$ such that $a \succ'_i b$ for some voter $i \in \{1, 2\}$. If no such profile R' exists, it is obvious that b Pareto-dominates a in every profile R' with $r^*(R) = r^*(R')$. First, note that a cannot be the uniquely most preferred alternative of voter i in R' because otherwise, $r^*(R) = r^*(R')$ cannot be true. Hence, there is an alternative $c \in A \setminus \{a, b\}$ such that $c \succ'_i a$. Analogously, voter i cannot uniquely bottom-rank b , which means that his preference relation in R' is $c \succ'_i a \succ'_i b \succ'_i d$. Furthermore, we have that $s_{b2} \leq s_{a2}$, which means that a is among the least preferred alternatives of the second voter j in R' because there are two alternatives that strictly dominate b in \succ'_i and $r(\succ'_i, b) = r(\succ_2, b)$. This means that c Pareto-dominates a in R' because either $s_{b1} < s_{a1}$ or $s_{b2} < s_{a2}$, which means that either voter i strictly prefers c to a , or voter j uniquely bottom-ranks a . Hence, a is Pareto-dominated in R' , which proves our claim. \square

As last result, we discuss the SCF f^* that satisfies Pareto-optimality and strategyproofness but violates support-basedness and that every voter is a nominator. As described in Remark 6, this SCF chooses the maximal alternatives of a transitive dominance relation which slightly strengthens Pareto-dominance. In more detail, we say that an alternative a *dominates* alternative b in a profile R if a Pareto-dominates b or $n - 1$ voters prefer a the most while $s_{ab}(R) \geq 2$ and $s_{ba}(R) = 1$. It should be stressed that it is not required that a is uniquely top-ranked by $n - 1$ voters, but only that it is among their best alternatives. Subsequently, we show that f^* satisfies all axioms that we claim.

Proposition 3. *The SCF f^* satisfies Pareto-optimality and strategyproofness but violates support-basedness and that every voter is a nominator if $n \geq 3$.*

Proof. Before discussing the axioms, we first show that f^* is a well-defined SCFs by proving that it chooses the maximal elements of a transitive dominance relation. Hence, consider an arbitrary profile R and three alternatives a, b, c and assume that a dominates b and b dominates c . As there are two possibilities on how an alternative dominates another one (i.e., a either Pareto-dominates b , or $s_{ab}(R) \geq 2$, $s_{ba}(R) = 1$, and $n - 1$ voter top-rank a), we proceed with a case distinction with respect to the dominance relations between a and

b and between b and c . First, consider the case that a Pareto-dominates b and b Pareto-dominates c . Then, a Pareto-dominates c as the Pareto-dominance relation is transitive.

Next, consider the case that a Pareto-dominates b and b dominates c because $s_{bc}(R) \geq 2$, $s_{cb}(R) = 1$, and $n - 1$ voter top-rank b . Since every voter prefers a (weakly) to b , it follows that $s_{ac}(R) \geq s_{bc}(R) \geq 2$, $s_{ca}(R) \leq s_{cb}(R) = 1$ and that $n - 1$ voters top-rank a . Hence, a either Pareto-dominates c if $s_{ca}(R) = 0$ or satisfies the second dominance criterion if $s_{ca}(R) = 1$. This means that the dominance relation is also in this case transitive.

As third case, assume that b Pareto-dominates c , and that $s_{ab}(R) \geq 2$, $s_{ba}(R) = 1$, and $n - 1$ voters top-rank a . Since b Pareto-dominates c , it follows that $s_{ac}(R) \geq s_{ab}(R) \geq 2$ and $s_{ca}(R) \leq s_{ba}(R) = 1$. Hence, transitivity is also in this case satisfied.

Finally, assume that neither a Pareto-dominates b nor b Pareto-dominates c , but a dominates b and b dominates c . Consequently, we derive that both a and b are top-ranked by $n - 1$ voters. However, this means that at most a single voter prefers a strictly to b and thus, $s_{ab}(R) \leq 1$. This contradicts that a dominates b and thus, this case cannot occur. Hence, the resulting dominance relation is transitive and f^* is a well-defined SCF.

Next, note that f^* satisfies Pareto-optimality as it is defined by a dominance relation that refines Pareto-dominance. Moreover, no voter is a nominator for f^* because $f^*(R) = \{a\}$ for all profiles R in which $n - 1$ voters report a as their uniquely best alternative. The SCF f^* is also not support-based. To this end, consider the profiles R^1 and R^2 , where $X = A \setminus \{a, b, c\}$, and note that $f^*(R^1) = \{a\} \neq \{a, b, c\} = f^*(R^2)$ even though $s^*(R^1) = s^*(R^2)$.

R^1 :	1: c, b, a, X	2: a, b, c, X	$[3 \dots n]$: a, b, c, X
R^2 :	1: c, a, b, X	2: b, a, c, X	$[3 \dots n]$: a, b, c, X

Finally, it remains to show that f^* is strategyproof. Assume for contradiction that this is not the case, i.e., there are preference profiles R and R' and a voter i such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$ and $f^*(R') \succ_i f^*(R)$. Moreover, recall that $T_i(R)$ denotes voter i 's favorite alternatives in R . We proceed with a case distinction with respect to whether $T_i(R) \cap f^*(R)$ is empty or not. First, assume that $T_i(R) \cap f^*(R)$ is non-empty. This means that voter i can only manipulate by deviating to R' if $f^*(R') \subseteq T_i(R)$ and $f^*(R) \not\subseteq T_i(R)$. Since the dominance relation defining f^* is transitive, it follows that there are alternatives $x \in T_i(R)$, $y \in f^*(R) \setminus T_i(R)$ such that x dominates y in R' but not in R . However, this is not possible. If x does not Pareto-dominate y in R , there is a voter $j \neq i$ with $y \succ_j x$ and thus, x cannot Pareto-dominate y in R' . Furthermore, since $x \succ_i y$, it follows that $s_{xy}(R) \geq s_{xy}(R')$ and $s_{yx}(R) \leq s_{yx}(R')$, and since $x \in T_i(R)$, voter i can also not increase the number of voters who top-rank x . Consequently, since x does not dominate y in R , it does not dominate y in R' . Hence, it follows from the transitivity of the dominance relation defining f^* that $f^*(R')$ cannot be a subset of $T_i(R)$ if $f^*(R) \not\subseteq T_i(R)$, which means that no manipulation is possible in this case.

Next, assume that $T_i(R) \cap f^*(R) = \emptyset$, i.e., none of voter i 's best alternatives are chosen. Because at least one of voter i 's best alternatives is Pareto-optimal, it follows that there is a non-empty set of alternatives B such that all voters $j \in N \setminus \{i\}$ top-rank all alternatives in B . Moreover, let a denote one of voter i 's most preferred alternatives in $f^*(R)$ and let b denote one of voter i 's most preferred alternatives in B . Observe that all alternatives x with $b \succ_i x$ are Pareto-dominated by b because all voters but i top-rank b and thus, these

alternatives are not in $f^*(R)$. Moreover, it holds that $b \in f^*(R)$. Indeed, it could only be Pareto-dominated by alternatives in B , but it is voter i 's best alternative among these. Moreover, $s_{yb}(R) \leq 1$ for all $y \in A$ because $n - 1$ voters top-rank b and hence, it is not dominated.

As next step, we show that for all alternatives $y \in A$ with $y \succ_i a$ that $y \notin f^*(R')$. We prove this claim by showing that there is for every such alternative y an alternative $z \in B$ such that $s_{zy}(R) \geq 2$ and $s_{yz}(R) \leq 1$. This implies that $s_{zy}(R') \geq 2$ and $s_{yz}(R') \leq 1$ for all these alternatives because $y \succ_i a \succsim_i b \succsim_i z$, which means that $y \notin f^*(R')$. Hence, assume for contradiction that there is an alternative $c \in A \setminus f^*(R)$ such that $c \succ_i a$ and $s_{xc}(R) \leq 1$ for all $x \in B$ ($s_{cx}(R) \leq 1$ must be true for all $x \in B$ since $n - 1$ voters top-rank these alternatives). Since $c \notin f^*(R)$ and $s_{xc}(R) \leq 1$ for all $x \in B$, it is Pareto-dominated by an alternative d ; otherwise c must be chosen. As a consequence of Pareto-dominance, we derive that $s_{xd}(R) \leq s_{xc}(R) \leq 1$ for all $x \in B$ and that $d \succsim_i c \succ_i a$. In particular, the last point implies that $d \notin f^*(R)$ because of the definition of a and hence, we can apply the same argument as for c . In more detail, by repeating this argument, we will eventually find a Pareto-optimal alternative e with $s_{xe}(R) \leq 1$ for all $x \in B$ and $e \succ_i a$ because the Pareto-dominance relation is transitive. The definition of f^* shows then that $e \in f^*(R)$, contradicting that $a \succsim_i x$ for all $x \in f^*(R)$. This is the desired contradiction and hence, there is for all alternatives $y \in A$ with $y \succ_i a$ an alternative $z \in B$ such that $s_{zy}(R) \geq 2$ and $s_{yz}(R) \leq 1$. This shows that no alternative with $y \succ_i a$ is in $f^*(R')$.

As a consequence of the last observation, voter i can only manipulate by deviating from R to R' if $x \sim_i a$ for all $x \in f^*(R')$ and there is an alternative $y \in f^*(R)$ with $a \succ_i y$. The latter observation implies that $a \succ_i b$ because all alternatives x with $b \succ_i x$ are Pareto-dominated. By the definition of b , we can therefore derive a contradiction by proving that $B \cap f^*(R') \neq \emptyset$. Note for this that all alternatives in B are also in R' top-ranked by $n - 1$ voters and thus $s_{xy}(R') \leq 1$ for all $x \in A$, $y \in B$. This means that an alternative $x \in B$ is only not chosen in $f^*(R')$ if it is Pareto-dominated. However, an alternative $x \in B$ can only be Pareto-dominated by another alternative in B because for every alternative $y \in A \setminus B$, there is a voter $j \in N \setminus \{i\}$ such that $x \succ_j y$. Finally, as the Pareto-dominance relation is transitive, it follows that there is a Pareto-optimal alternative in B , and thus, $B \cap f^*(R') \neq \emptyset$. Altogether, this proves that f^* is strategyproof. \square

References

- Arrow, K. J. (1951). *Social Choice and Individual Values* (1st edition). New Haven: Cowles Foundation. 2nd edition 1963.
- Aziz, H., Brandl, F., Brandt, F., & Brill, M. (2018). On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110, 1–18.
- Bandyopadhyay, T. (1983). Coalitional manipulation and the Pareto rule. *Journal of Economic Theory*, 29(2), 359–363.
- Barberà, S. (1977a). The manipulation of social choice mechanisms that do not leave “too much” to chance. *Econometrica*, 45(7), 1573–1588.

- Barberà, S. (1977b). Manipulation of social decision functions. *Journal of Economic Theory*, 15(2), 266–278.
- Barberà, S. (2010). Strategy-proof social choice. In Arrow, K. J., Sen, A. K., & Suzumura, K. (Eds.), *Handbook of Social Choice and Welfare*, Vol. 2, chap. 25, pp. 731–832. Elsevier.
- Barberà, S., Dutta, B., & Sen, A. (2001). Strategy-proof social choice correspondences. *Journal of Economic Theory*, 101(2), 374–394.
- Benoît, J.-P. (2002). Strategic manipulation in voting games when lotteries and ties are permitted. *Journal of Economic Theory*, 102(2), 421–436.
- Bogomolnaia, A., & Moulin, H. (2001). A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2), 295–328.
- Bogomolnaia, A., Moulin, H., & Stong, R. (2005). Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2), 165–184.
- Brandl, F., Brandt, F., Eberl, M., & Geist, C. (2018). Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, 65(2), 1–28.
- Brandl, F., Brandt, F., Geist, C., & Hofbauer, J. (2019). Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. *Journal of Artificial Intelligence Research*, 66, 1031–1056.
- Brandl, F., Brandt, F., Peters, D., & Stricker, C. (2021). Distribution rules under dichotomous preferences: Two out of three ain’t bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM-EC)*, pp. 158–179.
- Brandt, F. (2015). Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4), 793–804.
- Brandt, F. (2017). Rolling the dice: Recent results in probabilistic social choice. In Endriss, U. (Ed.), *Trends in Computational Social Choice*, chap. 1, pp. 3–26. AI Access.
- Brandt, F., Brill, M., & Harrenstein, P. (2016a). Tournament solutions. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. D. (Eds.), *Handbook of Computational Social Choice*, chap. 3. Cambridge University Press.
- Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. (Eds.). (2016b). *Handbook of Computational Social Choice*. Cambridge University Press.
- Brandt, F., Saile, C., & Stricker, C. (2022). Strategyproof social choice when preferences and outcomes may contain ties. *Journal of Economic Theory*. Forthcoming.
- Ching, S., & Zhou, L. (2002). Multi-valued strategy-proof social choice rules. *Social Choice and Welfare*, 19(3), 569–580.
- Duggan, J., & Schwartz, T. (2000). Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1), 85–93.
- Dutta, B., & Laslier, J.-F. (1999). Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4), 513–532.
- Endriss, U. (Ed.). (2017). *Trends in Computational Social Choice*. AI Access.

- Feldman, A. (1979). Manipulation and the Pareto rule. *Journal of Economic Theory*, 21, 473–482.
- Fishburn, P. C. (1977). Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3), 469–489.
- Gärdenfors, P. (1976). Manipulation of social choice functions. *Journal of Economic Theory*, 13(2), 217–228.
- Gärdenfors, P. (1979). On definitions of manipulation of social choice functions. In Laffont, J. J. (Ed.), *Aggregation and Revelation of Preferences*. North-Holland.
- Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica*, 41(4), 587–601.
- Gibbard, A. (1977). Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3), 665–681.
- Kelly, J. S. (1977). Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2), 439–446.
- Laslier, J.-F. (1996). Rank-based choice correspondences. *Economics Letters*, 52(3), 279–286.
- Laslier, J.-F. (2000). Interpretation of electoral mixed strategies. *Social Choice and Welfare*, 17, 283–292.
- Le Breton, M., & Weymark, J. A. (2011). Arrovian social choice theory on economic domains. In Arrow, K. J., Sen, A. K., & Suzumura, K. (Eds.), *Handbook of Social Choice and Welfare*, Vol. 2, chap. 17. North-Holland.
- MacIntyre, I., & Pattanaik, P. K. (1981). Strategic voting under minimally binary group decision functions. *Journal of Economic Theory*, 25(3), 338–352.
- Mas-Colell, A., & Sonnenschein, H. (1972). General possibility theorems for group decisions. *Review of Economic Studies*, 39(2), 185–192.
- Nehring, K. (2000). Monotonicity implies generalized strategy-proofness for correspondences. *Social Choice and Welfare*, 17(2), 367–375.
- Rodríguez-Álvarez, C. (2007). On the manipulation of social choice correspondences. *Social Choice and Welfare*, 29(2), 175–199.
- Sato, S. (2008). On strategy-proof social choice correspondences. *Social Choice and Welfare*, 31, 331–343.
- Sato, S. (2014). A fundamental structure of strategy-proof social choice correspondences with restricted preferences over alternatives. *Social Choice and Welfare*, 42(4), 831–851.
- Satterthwaite, M. A. (1975). Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2), 187–217.
- Sen, A. K. (1986). Social choice theory. In Arrow, K. J., & Intriligator, M. D. (Eds.), *Handbook of Mathematical Economics*, Vol. 3, chap. 22, pp. 1073–1181. Elsevier.

- Taylor, A. D. (2005). *Social Choice and the Mathematics of Manipulation*. Cambridge University Press.
- Young, H. P. (1974). An axiomatization of Borda's rule. *Journal of Economic Theory*, 9(1), 43–52.

SUMMARY

In party-approval committee elections, the goal is to allocate the seats of a fixed-size committee to parties based on the approval ballots of the voters over the parties. In particular, each voter can approve multiple parties and each party can be assigned multiple seats. Two central requirements in this setting are proportional representation and strategyproofness. Intuitively, proportional representation requires that every sufficiently large group of voters with similar preferences is represented in the committee. On the other hand, strategyproofness demands that no voter can benefit by misreporting her true preferences. In this paper, we show that these two concepts are incompatible for anonymous party-approval committee voting rules.

In more detail, we formalize the idea of proportional presentation with weak representation and weak proportional representation, which require that if a group of at least $\frac{n}{k}$ (resp. $\ell \frac{n}{k}$) voters uniquely approve a party, then this party gets at least 1 (resp. ℓ) seat(s) in the chosen committee. Based on these axioms, we then show the following two impossibility theorems for party-approval committee voting rules (which are subsequently only called party-approval rules):

- No anonymous party-approval rule satisfies both weak representation and strategyproofness if $k \geq 3$, $m \geq k + 1$, and $n = 2\ell k$ for some $\ell \in \mathbb{N}$.
- No anonymous party-approval rule satisfies both weak proportional representation and strategyproofness if $k \geq 3$, $m \geq 4$, and $n = 2\ell k$ for some $\ell \in \mathbb{N}$.

We note that both of these impossibilities are obtained by a computer-aided approach called SAT solving.

Additionally, we consider an escape route to these negative results by studying a weakening of strategyproofness which requires that only voters who do not approve any of the elected parties cannot manipulate. On the one hand, we show that numerous party-approval rules fail even this very weak notion of strategyproofness. On the other hand, we prove that Chamberlin-Courant approval voting satisfies this axiom. Based on this insight, we characterize this rule within the class of Thiele rules based on strategyproofness for unrepresented voters and weak representation, which demonstrates that at least weak variants of our strategyproofness and proportional representation are compatible.

REFERENCE

T. Delemazure, T. Demeulemeester, M. Eberl, J. Israel, and P. Lederer. **Strategyproofness and proportionality in party-approval multiwinner elections.** In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5591–5599, 2023.
DOI: <https://doi.org/10.1609/aaai.v37i5.25694>

INDIVIDUAL CONTRIBUTION

My coauthor Jonas Israel and I, Patrick Lederer, are the joint main authors of this paper. In particular, we are jointly responsible for the conceptual design of the paper and the write-up. Moreover, while my coauthors have derived the results in Section 4, I am responsible for the results in Section 3.

COPYRIGHT AGREEMENT

The right to present this paper in a doctoral thesis has been granted by the publisher, AI Access Foundation, in the copyright form presented below. There, it is stated that “*personal reuse of all or portions of the above article/paper in other works of their own authorship.*” This form can also be found at <https://proceedings.aaai.org/docs/AAAI-23-copyrightform.pdf> (accessed August 24, 2023).

TERMINOLOGY

In this paper, we refer to party-approval committee voting rules simply as party-approval rules. Moreover, the committee size is part of the input of these rules, while we defined PAC voting rules for a fixed committee size. This does not affect our results as none of our axioms relates different committee sizes with each other.



Association for the Advancement of Artificial Intelligence

1900 Embarcadero Road, Suite 101, Palo Alto, California 94303

Palo Alto, California 94303 USA

AAAI COPYRIGHT FORM

Title of Article/Paper: _____

Publication in Which Article/Paper Is to Appear: _____

Author's Name(s): _____

Please type or print your name(s) as you wish it (them) to appear in print

PART A – COPYRIGHT TRANSFER FORM

The undersigned, desiring to publish the above article/paper in a publication of the Association for the Advancement of Artificial Intelligence, (AAAI), hereby transfer their copyrights in the above article/paper to the Association for the Advancement of Artificial Intelligence (AAAI), in order to deal with future requests for reprints, translations, anthologies, reproductions, excerpts, and other publications.

This grant will include, without limitation, the entire copyright in the article/paper in all countries of the world, including all renewals, extensions, and reversions thereof, whether such rights current exist or hereafter come into effect, and also the exclusive right to create electronic versions of the article/paper, to the extent that such right is not subsumed under copyright.

The undersigned warrants that they are the sole author and owner of the copyright in the above article/paper, except for those portions shown to be in quotations; that the article/paper is original throughout; and that the undersigned right to make the grants set forth above is complete and unencumbered.

If anyone brings any claim or action alleging facts that, if true, constitute a breach of any of the foregoing warranties, the undersigned will hold harmless and indemnify AAAI, their grantees, their licensees, and their distributors against any liability, whether under judgment, decree, or compromise, and any legal fees and expenses arising out of that claim or actions, and the undersigned will cooperate fully in any defense AAAI may make to such claim or action. Moreover, the undersigned agrees to cooperate in any claim or other action seeking to protect or enforce any right the undersigned has granted to AAAI in the article/paper. If any such claim or action fails because of facts that constitute a breach of any of the foregoing warranties, the undersigned agrees to reimburse whomever brings such claim or action for expenses and attorneys' fees incurred therein.

Returned Rights

In return for these rights, AAAI hereby grants to the above author(s), and the employer(s) for whom the work was performed, royalty-free permission to:

1. Retain all proprietary rights other than copyright (such as patent rights).
2. Personal reuse of all or portions of the above article/paper in other works of their own authorship. This does not include granting third-party requests for reprinting, republishing, or other types of reuse. AAAI must handle all such third-party requests.
3. Reproduce, or have reproduced, the above article/paper for the author's personal use, or for company use provided that AAAI copyright and the source are indicated, and that the copies are not used in a way that implies AAAI endorsement of a product or service of an employer, and that the copies per se are not offered for sale. The foregoing right shall not permit the posting of the article/paper in electronic or digital form on any computer network, except by the author or the author's employer, and then only on the author's or the employer's own web page or ftp site. Such web page or ftp site, in addition to the aforementioned requirements of this Paragraph, shall not post other AAAI copyrighted materials not of the author's or the employer's creation (including tables of contents with links to other papers) without AAAI's written permission.
4. Make limited distribution of all or portions of the above article/paper prior to publication.
5. In the case of work performed under a U.S. Government contract or grant, AAAI recognized that the U.S. Government has royalty-free permission to reproduce all or portions of the above Work, and to authorize others to do so, for official U.S. Government purposes only, if the contract or grant so requires.

In the event the above article/paper is not accepted and published by AAAI, or is withdrawn by the author(s) before acceptance by AAAI, this agreement becomes null and void.

(1)

Author/Authorized Agent for Joint Author's Signature

Date (MM/DD/YYYY)

Employer for whom work was performed

Title (if not author)

(For jointly authored Works, all joint authors should sign unless one of the authors has been duly authorized to act as agent for the others.)

Strategyproofness and Proportionality in Party-Approval Multiwinner Elections

Théo Delemazure¹, Tom Demeulemeester², Manuel Eberl³, Jonas Israel⁴, Patrick Lederer⁵

¹ Paris Dauphine University, PSL, CNRS, France,

² Research Center for Operations Research & Statistics, KU Leuven, Belgium,

³ Computational Logic Group, University of Innsbruck, Austria,

⁴ Research Group Efficient Algorithms, Technische Universität Berlin, Germany,

⁵ Technical University of Munich, Germany,

theo.delemazure@dauphine.eu, tom.demeulemeester@kuleuven.be, manuel.eberl@uibk.ac.at, j.israel@tu-berlin.de, lederer@in.tum.de

Abstract

In party-approval multiwinner elections the goal is to allocate the seats of a fixed-size committee to parties based on the approval ballots of the voters over the parties. In particular, each voter can approve multiple parties and each party can be assigned multiple seats. Two central requirements in this setting are proportional representation and strategyproofness. Intuitively, proportional representation requires that every sufficiently large group of voters with similar preferences is represented in the committee. Strategyproofness demands that no voter can benefit by misreporting her true preferences. We show that these two axioms are incompatible for anonymous party-approval multiwinner voting rules, thus proving a far-reaching impossibility theorem. The proof of this result is obtained by formulating the problem in propositional logic and then letting a SAT solver show that the formula is unsatisfiable. Additionally, we demonstrate how to circumvent this impossibility by considering a weakening of strategyproofness which requires that only voters who do not approve any elected party cannot manipulate. While most common voting rules fail even this weak notion of strategyproofness, we characterize Chamberlin–Courant approval voting within the class of Thiele rules based on this strategyproofness notion.

1 Introduction

A central problem in multi-agent systems is collective decision making: given the preferences of multiple agents over a set of alternatives, a joint decision has to be made. While classic literature for this problem focuses on the case of choosing a single alternative as the winner, there is also a wide range of scenarios where a set of winners needs to be elected. For instance, this is the case in parliamentary elections, where the seats of a parliament are assigned to parties based on the voters’ preferences. In the literature, parliamentary elections are studied under the term *apportionment* and a crucial assumption for their analysis is that voters are only allowed to vote for a single party (Balinski and Young 2001; Pukelsheim 2014). However, this assumption has recently been criticized because of its lack of flexibility and expressiveness (Brill, Laslier, and Skowron 2018; Brill et al. 2020). Following Brill et al. (2020), we thus study *party-approval elections*. In this setting, the parliament, or more generally a multiset of fixed

size, is elected based on the approval ballots of the voters, i.e., each voter reports a set of approved parties instead of only her most preferred one.

Two central desiderata for party-approval elections are *proportional representation* and *strategyproofness*. The former requires that the chosen committee should proportionally reflect the voters’ preferences. The latter postulates that no voter can benefit by misreporting her preferences. While Brill et al. (2020) have shown that even core-stable committees, which satisfy one of the highest degrees of proportionality, always exist in party-approval elections, strategyproofness is not yet well-understood for this setting. We thus analyze the trade-off between strategyproofness and proportional representation for party-approval elections in this paper.

Our research question also draws motivation from related models (see Figure 1 for details). Firstly, party-approval elections can be seen as a special case of *approval-based committee (ABC) elections*, where voters approve individual candidates instead of parties and the outcome is a subset of the candidates instead of a multiset. For ABC elections, proportionality and strategyproofness have received significant attention (see, e.g., the survey by Lackner and Skowron (2022)). Unfortunately, these axioms are jointly incompatible for ABC voting rules (Peters 2018) and our study can thus be seen as an attempt to circumvent this impossibility. Even more, there are hints that these axioms could be compatible for party-approval elections: this setting lies logically between ABC elections on one side, and either *apportionment* (where voters can only approve a single party instead of multiple ones (Balinski and Young 2001; Pukelsheim 2014)) or *fair mixing* (where the outcome is a probability distribution over the parties instead of a multiset (Bogomolnaia, Moulin, and Stong 2005; Aziz, Bogomolnaia, and Moulin 2019)) on the other side. Since strategyproofness and proportionality are compatible in the latter two models, it seems reasonable to conjecture positive results for party-approval elections.

Our contribution. Unfortunately, it turns out that strategyproofness conflicts even with minimal notions of proportional representation in party-approval elections. To prove this, we introduce the notions of weak representation and weak proportional representation, which require that a party is assigned at least 1 (resp. ℓ) out of k available seats if it is uniquely approved by at least an $\frac{1}{k}$ (resp. $\frac{\ell}{k}$) fraction of

the voters. Then, we show in Section 3 the following impossibility theorems (k , m , and n denote the numbers of seats, parties, and voters, respectively):

- No anonymous party-approval rule satisfies weak representation and strategyproofness if $k \geq 3$, $m \geq k + 1$, and $2k$ divides n .
- No anonymous party-approval rule satisfies weak proportional representation and strategyproofness if $k \geq 3$, $m \geq 4$, and $2k$ divides n .

The first result shows that the incompatibility of strategyproofness and proportional representation first observed for ABC elections also prevails for party-approval elections. Even more, our result implies such an impossibility for ABC elections as our setting is more general. The main drawback of the first result is that it requires more parties than seats in the committee. While this assumption is true for many applications inspired from ABC voting, this is not the case in our initial example of parliamentary elections. However, our second impossibility shows that strategyproofness is still in conflict with proportional representation if $k > m$.


We prove both of these results with a computer-aided approach based on SAT solving, which has recently led to a number of sweeping impossibility results (e.g., Brandl et al. 2021; Brandt, Saile, and Stricker 2022). In particular, our computer proof relies on 635 profiles, which makes it the largest computer proof in social choice theory (the previous record is due to Brandl et al. (2021) and uses 386 profiles).

Finally, in Section 4 we investigate a weakening of strategyproofness that only requires that voters who do not approve any party in the elected committee cannot manipulate. Perhaps surprisingly, many commonly studied voting rules fail this condition. We can thus characterize Chamberlin–Court approval voting as the only Thiele rule satisfying this strategyproofness notion and weak representation, proving an attractive escape route to our impossibility results.

Omitted proofs and further details can be found in the full version of the paper (Delemazure et al. 2022b).

Related work. Party-approval elections have been introduced by Brill et al. (2020) who showed that strong proportionality axioms can be satisfied in this setting, but we are not aware of any follow-up paper. We thus draw much inspiration from ABC elections for which there is a large amount of work on proportional representation (e.g., Aziz et al. 2017; Sánchez-Fernández et al. 2017; Peters and Skowron 2020; Brill et al. 2022) and strategyproofness (e.g., Aziz et al. 2015; Peters 2018; Lackner and Skowron 2018). For instance, Aziz et al. (2017) analyze ABC voting rules with respect to more restrictive variants of weak representation. The main message from work on proportional representation is that there are few ABC voting rules that guarantee strong representation axioms. The results on strategyproofness are mostly negative: after early results (Aziz et al. 2015; Lackner and Skowron 2018) proving that no known rule satisfies both strategyproofness and proportional representation, Peters (2018) showed that these axioms are inherently incompatible for ABC voting rules (see also Duddy 2014; Kluiving et al. 2020). Our impossibility theorems are closely connected to this result but logically independent: we need stronger strategyproofness

Models ordered by domain restrictions:



Models ordered by output type:



Figure 1: Relation of party-approval elections to other voting settings. An arrow from X to Y means that model X is more general than model Y . In the settings in the top row, elections return sets of alternatives but the models impose different restrictions on the input profiles: for ABC voting every input profile is allowed, for party-approval profiles each voter can for each party (viewed as a set of alternatives) either approve all of its members or none, and for apportionment each voter must approve all members of exactly one party. In the bottom row, the models are ordered with respect to their output type: all of fair mixing, party-approval elections, and ABC elections can take arbitrary approval profiles as input, but fair mixing rules return a probability distribution over the alternatives, party-approval rules choose a multiset of the alternatives, and ABC rules choose a subset of the alternatives. This shows that party-approval elections can be seen both as generalization and special case of ABC elections.

and representation axioms and additionally anonymity, but use a more flexible setting and no efficiency condition.

2 Preliminaries

Let $N = \{1, \dots, n\}$ denote a set of n voters and $\mathcal{P} = \{a, b, c, \dots\}$ a set of m parties. Each voter $i \in N$ is assumed to have a *dichotomous preference relation* over the parties, i.e., she partitions the parties into approved and disapproved ones. The *approval ballot* $A_i \subseteq \mathcal{P}$ of a voter i is the non-empty set of her approved parties. With slight abuse of notation we omit commas and brackets when writing approval ballots. Let \mathcal{A} denote the set of all possible approval ballots. An *approval profile* $A \in \mathcal{A}^n$ is the collection of the approval ballots of all voters. Given an approval profile A , the goal in *party-approval elections* is to assign a fixed number of seats to the parties. We call such an outcome a *committee*, which is formally a multiset of parties $W : \mathcal{P} \rightarrow \mathbb{N}$, and $W(x)$ denotes the number of seats assigned to party x . We extend this notation also to sets of parties $X \subseteq \mathcal{P}$ by defining $W(X) = \sum_{x \in X} W(x)$. Furthermore, we indicate specific committees by square brackets, e.g., $[a, a, b]$ is the committee containing party a twice and party b once. Let \mathcal{W}_k denote the set of all committees of size k .

A *party-approval rule* is a function f which takes an approval profile $A \in \mathcal{A}^n$ and a target committee size k as input and returns a winning committee $W \in \mathcal{W}_k$. In particular, party-approval rules are resolute, i.e., there is always a single winning committee. We define $f(A, k, x)$ as the number of seats assigned to party x by f for the profile A when choosing a committee of size k . We extend this notion also to sets by defining $f(A, k, X) = \sum_{x \in X} f(A, k, x)$.

Two well-known properties of voting rules are anonymity and Pareto optimality. Intuitively, anonymity requires that

all voters are treated equally, i.e., a party-approval rule f is *anonymous* if $f(A, k) = f(A', k)$ for all committee sizes $k \in \mathbb{N}$ and all approval profiles $A, A' \in \mathcal{A}^n$ such that there is a permutation $\pi : N \rightarrow N$ with $A'_i = A_{\pi(i)}$.

Next, we say that a party x *Pareto dominates* another party y in an approval profile A if $y \in A_i$ implies $x \in A_i$ for all $i \in N$ and there is a voter $i \in N$ with $x \in A_i$ and $y \notin A_i$. Then, a party-approval rule f is *Pareto optimal* if $f(A, k, y) = 0$ for all approval profiles A , committee sizes k , and parties y that are Pareto dominated in A .

2.1 Proportional Representation

One of the central desiderata in committee elections is to choose a committee that proportionally represents the voters' preferences. The notion of *justified representation*, introduced by Aziz et al. (2017), formalizes this idea by requiring that in a committee of size k , any group of voters $G \subseteq N$ with $|G| \geq \frac{n}{k}$ that agrees on a party should be represented. In this paper, we will consider a weakening of this property which we call *weak representation*. Intuitively, weak representation weakens justified representation by only considering cases where all voters in G uniquely approve a single party x .

Definition 1 (Weak Representation). A party-approval rule f satisfies *weak representation* if $f(A, k, x) \geq 1$ for every profile A , committee size k , party x , and group of voters G such that $|G| \geq \frac{n}{k}$ and $A_i = \{x\}$ for all $i \in G$.

Weak representation can easily be satisfied if we have more seats in the committee than parties by simply assigning at least one seat to every party. This, however, contradicts the idea of proportional representation since a large part of the chosen committee is independent of the voters' preferences. To address this issue, we consider weak proportional representation, which is a weakening of *proportional justified representation* (Sánchez-Fernández et al. 2017). Clearly, weak proportional representation implies weak representation.

Definition 2 (Weak Proportional Representation). A party-approval rule f satisfies *weak proportional representation* if $f(A, k, x) \geq \ell$ for every $\ell \in \mathbb{N}$, profile A , committee size k , party x , and group of voters G such that $|G| \geq \ell \frac{n}{k}$ and $A_i = \{x\}$ for all $i \in G$.

2.2 Strategyproofness

Intuitively, strategyproofness requires that a voter cannot benefit by lying about her true preferences. Consequently, if a party-approval rule fails strategyproofness, we cannot expect the voters to submit their true preferences, which may lead to socially undesirable outcomes.

Definition 3 (Strategyproofness). A party-approval rule f is *strategyproof* if $f(A, k, A_i) \geq f(A', k, A_i)$ for all approval profiles A, A' , committee sizes k , and voters $i \in N$ such that $A_j = A'_j$ for all $j \in N \setminus \{i\}$.

The motivation for this strategyproofness notion stems from the assumption that voters are indifferent between their approved parties. Then, only the number of seats assigned to these parties matters to the voters. This strategyproofness notion is commonly used in ABC voting under the name cardinality strategyproofness (e.g., Lackner and Skowron

2018; Botan 2021), and equivalent notions are used for fair mixing (e.g., Bogomolnaia, Moulin, and Stong 2005; Aziz, Bogomolnaia, and Moulin 2019).

Since we will show that strategyproofness is in conflict with minimal representation axioms, we also consider the following weakening which requires that only voters without representation in the committee cannot manipulate.

Definition 4 (Strategyproofness for Unrepresented Voters). A party-approval rule f is *strategyproof for unrepresented voters* if $f(A, k, A_i) \geq f(A', k, A_i)$ for all approval profiles A, A' , committee sizes k , and voters $i \in N$ such that $A_j = A'_j$ for all $j \in N \setminus \{i\}$ and $f(A, k, A_i) = 0$.

We believe this to be a sensible relaxation of strategyproofness because voters without any representation in the committee are more prone to manipulate. Firstly, voters who do have some representation may be more cautious to manipulate because they fear losing their representation when misstating their preferences. Secondly, the benefit of having additional representation in the committee is less straightforward than that of being represented at all.

2.3 Party-Approval Rules

Finally, we introduce three classes of party-approval rules. Note that even though we define party-approval rules for a fixed numbers of voters n and parties m , all subsequent rules are independent of such details.

Thiele rules. Thiele rules are arguably the most well-studied class of rules in the ABC setting. Introduced by Thiele (1895), a *w-Thiele rule* f is defined by a non-increasing and non-negative vector $w = (w_1, w_2, \dots)$ and chooses for each committee size k the committee $W \in \mathcal{W}_k$ that maximizes the score $s_w(W, A) = \sum_{i \in N} \sum_{j=1}^{W(A_i)} w_j$. Throughout the paper, we assume without loss of generality that $w_1 = 1$. There are many well-known Thiele rules, such as:

- *approval voting* (AV): $w = (1, 1, 1, \dots)$,
- *proportional approval voting* (PAV): $w = (1, \frac{1}{2}, \frac{1}{3}, \dots)$,
- *Chamberlin–Courant approval voting* (CCAV): $w = (1, 0, 0, \dots)$.

Sequential Thiele rules. Sequential Thiele rules are closely related to Thiele rules: instead of optimizing the score of the committee, these rules proceed in rounds and greedily choose in each iteration the party that increases the score of the committee the most. An important example of sequential Thiele rules is *sequential proportional approval voting* (seqPAV) defined by $w = (1, \frac{1}{2}, \frac{1}{3}, \dots)$.

Divisor methods based on majoritarian portioning. Brill et al. (2020) introduced the concept of composite party-approval rules, which combine a *portioning method* with an *apportionment method*. In this paper, we focus on an important subclass of such composite rules, namely *divisor methods based on majoritarian portioning*, because many of these rules satisfy strong representation axioms (Brill et al. 2020). These methods first apply *majoritarian portioning* to compute a weight w_x for each party x . Majoritarian portioning works in rounds and in each round, we determine the party x that is approved by the most voters. Then, we set its

weight w_x to the number of voters who approve x and remove all corresponding voters from the profile. This process is repeated until no voters are left. Finally, for all parties x that have no weight after all voters were removed, we set $w_x = 0$. After the portioning, we use a *divisor method* to allocate the seats to the parties based on the weights w_x . Divisor methods are defined by a monotone function $g : \mathbb{N}_0 \rightarrow \mathbb{R}_{>0}$ and proceed in rounds: in the i -th round, the next seat is assigned to the party x that maximizes $\frac{w_x}{g(t_{i-1}^x)}$, where t_{i-1}^x is the number of seats allocated to x in the previous $i - 1$ rounds. An example of a divisor method is Jefferson's method, where $g(x) = x + 1$.

Note that all rules defined above are irresolute, i.e., they may declare multiple committees as tied winners of an election. Since we investigate resolute voting rules in this paper, we assume that ties are broken lexicographically: for every $k \in \mathbb{N}$, there is a linear tie-breaking order \succ_k on the committees $W \in \mathcal{W}_k$ and, if a party-approval rule f declares multiple committees as tied winners, we choose the best one according to \succ_k . Similarly, if any rule is tied between multiple parties in a step, the tie is broken according to \succ_1 . The assumption of lexicographic tie-breaking is standard in the literature on strategyproofness (e.g., Faliszewski, Hemaspaandra, and Hemaspaandra 2010; Aziz et al. 2015).

3 Impossibility Results

In this section, we discuss the incompatibility of strategyproofness and proportional representation for party-approval rules by proving two sweeping impossibility theorems.

Theorem 1. *No party-approval rule simultaneously satisfies anonymity, weak representation, and strategyproofness if $k \geq 3$, $m \geq k + 1$, and $2k$ divides n .*

Note that Theorem 1 does not hold for all combinations of k , m , and n : we require that $2k$ divides n and that $m \geq k + 1$. The first assumption is mainly a technical one as we—just like other authors (Peters 2018; Kluiving et al. 2020)—could not find an argument to generalize the impossibility to arbitrary values of n . However, many party-approval rules (e.g., all Thiele rules and sequential Thiele rules) do not change their outcome when adding voters who approve all parties. For such rules, we can extend Theorem 1 to all $n \geq 2k$ by simply adding voters who approve all parties.

On the other side, the assumption that $m \geq k + 1$ is crucial for Theorem 1: if $m \leq k$, every rule that constantly returns a fixed committee W with $W(x) \geq 1$ for all $x \in \mathcal{P}$ satisfies the considered axioms. Nevertheless, we can restore the impossibility by strengthening weak representation to weak proportional representation.

Theorem 2. *No party-approval rule simultaneously satisfies anonymity, weak proportional representation, and strategyproofness if $k \geq 3$, $m \geq 4$, and $2k$ divides n .*

We believe that also the proofs of our results are of interest: for showing Theorems 1 and 2, we rely on a computer-aided approach called SAT solving. In the realm of social choice, this technique was pioneered by Tang and Lin (2009) and has by now been used to prove a wide variety of results (e.g.,

Peters 2018; Endriss 2020; Brandl et al. 2021). We refer to Geist and Peters (2017) for an overview of this technique.

To apply SAT solving to our problems, we proceed in three steps: first, we encode the problem of finding an anonymous party-approval rule that satisfies strategyproofness and weak representation for committees of size $k = 3$, $m = 4$ parties, and $n = 6$ voters as logical formula. By letting a computer program, a so-called SAT solver, show the formula unsatisfiable, we prove the base case of Theorems 1 and 2 for the given parameters. Next, we generalize the impossibility to larger values of k , m , and n based on inductive arguments. Finally, we verify the computer proof. The following subsections discuss each of these steps in detail.

Remark 1. AV satisfies all axioms of Theorem 1 except weak representation, and CCAV satisfies all axioms except strategyproofness. These examples show that these axioms are required for the impossibility. On the other hand, we could not show that anonymity is necessary for the impossibility and we conjecture that this axiom can be omitted.

Remark 2. For electorates where the committee size k is a multiple of the number of voters n , there are voting rules that satisfy weak proportional representation, anonymity, and strategyproofness. We can simply let every voter choose $\frac{k}{n}$ parties of the committee independently of the ballots of other voters. This is an important difference to the impossibility by Peters (2018), which also holds in the case that $n = k$.

Remark 3. If $k = 2$, a variant of AV satisfies all axioms of Theorems 1 and 2. For introducing this rule f , let \succ denote a linear order over the parties and $mAV(A)$ the maximal approval score of a party in the profile A . As first step, f removes clones according to \succ , i.e., for all parties x, y such that $x \in A_i$ if and only if $y \in A_i$ for all $i \in N$ and $x \succ y$, we remove y from A . This results in a reduced profile A' . Now, if $mAV(A') \neq \frac{n}{2}$ or there is only a single party with approval score of $\frac{n}{2}$, f assigns both seats to the approval winner. Else, f assigns the seats to the best and second best party with respect to \succ that have an approval score of $\frac{n}{2}$.

3.1 Computer-Aided Theorem Proving

The core observation for computer-aided theorem proving is that for a fixed committee size k and fixed numbers of parties m and voters n , there is a very large but finite number of party-approval rules. Hence, we could, at least theoretically, enumerate all rules and check whether they satisfy our requirements. However, the search space grows extremely fast (for $k = 3$, $m = 4$, and $n = 6$, there are roughly 6.2×10^{14819544} party-approval rules) and we thus use a different idea: we construct a logical formula which is satisfiable if and only if there is an anonymous party-approval rule that satisfies weak representation and strategyproofness for the given parameters of k , m , and n . By showing that the formula is unsatisfiable, we prove Theorems 1 and 2 for fixed parameters. Moreover, we can use computer programs, so-called SAT solvers, to show this.

Subsequently, we specify the variables and explain how we construct the formula. The idea is to introduce a variable $x_{A,W}$ for each profile $A \in \mathcal{A}^n$ and committee $W \in \mathcal{W}_k$, with the interpretation that $x_{A,W}$ is true if and only if

$f(A, k) = W$. However, for this formulation the mere number of profiles becomes prohibitive when $k = 3$, $m = 4$, and $n = 6$ and we thus apply several optimizations. First, we use anonymity to drastically reduce the number of considered profiles. This axiom states that the order of the voters does not matter for the outcomes and we thus view approval profiles from now on as multisets of approval ballots instead of ordered tuples. Next, we exclude certain approval profiles from the domain by imposing three conditions: (i) no voter is allowed to approve all parties, (ii) no party can be approved by more than four voters, and (iii) the total number of approvals given by all voters does not exceed eleven. We call the domain of all anonymous profiles that satisfy these conditions \mathcal{A}_{SAT}^n . Clearly, if there is no anonymous party-approval rule satisfying strategyproofness and weak representation on \mathcal{A}_{SAT}^n , there is also no such function on the full domain \mathcal{A}^n . For our final optimization, we note that weak representation requires that a committee W cannot be returned for a profile A if there is a party x with $W(x) = 0$ that is uniquely approved by $\frac{n}{k}$ or more voters. Hence, all corresponding variables $x_{A,W}$ must be set to false and we can equivalently omit them. To formalize this, we define $WR(A, k)$ as the set of committees of size k that satisfy weak representation for the profile A . Then, we add for every profile $A \in \mathcal{A}_{SAT}^n$ and every committee $W \in WR(A, k)$ a variable $x_{A,W}$.

Next, we turn to the constraints of our formula. First, we specify that the formula encodes a function f on \mathcal{A}_{SAT}^n , i.e., for every profile $A \in \mathcal{A}_{SAT}^n$, there is exactly one committee $W \in WR(A, k)$ such that $x_{A,W} = 1$. For this, we add two types of clauses for every profile A : the first one specifies that at least one committee is chosen for A and the second one that no more than one committee can be chosen.

$$\begin{aligned} \bigvee_{W \in WR(A, k)} x_{A,W} & \quad \forall A \in \mathcal{A}_{SAT}^n \\ \bigwedge_{V, W \in WR(A, k): V \neq W} \neg x_{A,V} \vee \neg x_{A,W} & \quad \forall A \in \mathcal{A}_{SAT}^n \end{aligned}$$

Since weak representation and anonymity are encoded in the choice of variables, we only need to add the subsequent constraints for strategyproofness. Here, $A^{A_i \rightarrow A_j}$ is the profile derived from A by changing a ballot A_i to A_j .

$$\begin{aligned} \neg x_{A,V} \vee \neg x_{A',W} & \quad \forall A, A' \in \mathcal{A}_{SAT}^n, V \in WR(A, k), \\ & \quad W \in WR(A', k) : \exists A_i, A_j \in \mathcal{A} : \\ & \quad A' = A^{A_i \rightarrow A_j} \wedge W(A_i) > V(A_i) \end{aligned}$$

For committees of size $k = 3$, $m = 4$ parties, and $n = 6$ voters, this construction results in a formula containing 21,418,593 constraints and a state-of-the-art SAT solver, such as Glucose (Audemard and Simon 2018), needs less than a minute to prove its unsatisfiability. Our code also provides options which further reduce the size of the formula to speed up the SAT solving (see the full version for details). Consequently, we derive the following result.

Proposition 1. *There is no party-approval rule that satisfies anonymity, weak representation, and strategyproofness if $k = 3$, $m = 4$, and $n = 6$.*

3.2 Inductive Arguments

Since weak proportional representation implies weak representation, Proposition 1 proves Theorems 1 and 2 for fixed parameters k , m , and n . To complete the proofs of these theorems, we use inductive arguments to generalize the impossibilities to larger parameters and subsequently present them for Theorem 1. For Theorem 2, only the third claim needs to be adapted (see the full version).

Lemma 1. *Assume there is no anonymous party-approval rule f that satisfies weak representation and strategyproofness for committees of size k , m parties, and n voters. The following claims hold:*

- (1) *For every $\ell \in \mathbb{N}$, there is no such rule for committees of size k , m parties, and $\ell \cdot n$ voters.*
- (2) *There is no such rule for committees of size k , $m + 1$ parties, and n voters.*
- (3) *If k divides n , there is no such rule for committees of size $k + 1$, $m + 1$ parties, and $\frac{n(k+1)}{k}$ voters.*

Proof sketch. For all three claims, we prove the contrapositive: if there is an anonymous party-approval rule f that satisfies strategyproofness and weak representation for the increased parameters, there is also such a rule g for committees of size k , m parties, and n voters. Subsequently, we discuss how to define the rule g for the three different cases:

- (1) Assume an $\ell \in \mathbb{N}$ such that f is defined for committees of size k , m parties, and $\ell \cdot n$ voters. Given a profile A for m parties and n voters, g copies every voter ℓ times to derive the profile A' . Then, $g(A, k) = f(A', k)$.
- (2) Assume f is defined for committees of size k , $m + 1$ parties, and n voters. Given a profile A for m parties and n voters, g first constructs the profile A^{xy} by cloning a party $x \in \mathcal{P}$ into a new party $y \notin \mathcal{P}$. More formally, $A^{xy} = A_i$ if $x \notin A_i$ and $A_i^{xy} = A_i \cup \{y\}$ otherwise. Finally, $g(A, k, z) = f(A^{xy}, k, z)$ for all $z \neq x$ and $g(A, k, x) = f(A^{xy}, k, xy)$.
- (3) Assume k divides n and f is defined for committees of size $k + 1$, $m + 1$ parties, and $\frac{n(k+1)}{k}$ voters. In this case, g maps a profile A for m parties and n voters to the profile \bar{A}^{xy} defined as follows: first g derives A^{xy} as explained in the previous case and then it adds $\frac{n}{k}$ voters with ballot xy . Finally, $g(A, k, z) = f(\bar{A}^{xy}, k + 1, z)$ for all $z \neq x$ and $g(A, k, x) = f(\bar{A}^{xy}, k + 1, xy) - 1$.

For all cases, it remains to show that g is a well-defined party-approval rule that satisfies anonymity, weak representation, and strategyproofness. Due to space restrictions, we explain this only for case (1) and defer the remaining cases to the full version. In this case, we first observe that g clearly inherits anonymity from f . Also, g satisfies weak representation: if $\frac{n}{k}$ or more voters uniquely approve a party x in a profile A , at least $\frac{\ell \cdot n}{k}$ voters uniquely approve x in A' . Thus, $g(A, x) = f(A', x) \geq 1$ because f satisfies weak representation. Finally, we prove that g is strategyproof. Note for this that $f(\bar{A}, k, \bar{A}_i) \geq f(\bar{A}', k, \bar{A}_i)$ for all profiles \bar{A}, \bar{A}' that only differ in the ballots of voters who report \bar{A}_i in \bar{A} . This is true because we can transform \bar{A} into \bar{A}' by letting voters with ballot \bar{A}_i manipulate one after another, and

strategyproofness shows for every step that the number of seats assigned to parties in \bar{A}_i cannot increase. Hence, g is strategyproof because if A and A' only differ in a single ballot A_i , the enlarged profiles \bar{A} and \bar{A}' differ in ℓ voters with ballot A_i . Thus, g meets all requirements in case (1). \square

3.3 Verification

Since Proposition 1 is proved by automated SAT solving, there is no complete human-readable proof for verifying Theorems 1 and 2. The standard approach for addressing this issue is to analyze minimal unsatisfiable subsets (MUSes) of the original formula, i.e., subsets of the formula which are unsatisfiable but removing a single constraint makes them satisfiable. Such MUSes are typically much smaller than the original formula, which makes it possible to translate them into a human-readable proof. Unfortunately, this technique does not work for Proposition 1 because all MUSes that we found (by using the programs haifamuc and muser2 (Belov and Marques-Silva 2012; Nadel, Ryvchin, and Strichman 2014)) are huge: even after applying several optimizations, the smallest MUS still contained over 20,000 constraints and 635 profiles. Because of the size of the MUSes, any human-readable proof would be unreasonably long and we thus verify our results by other means.

Firstly, we have published the code used for proving Proposition 1 (Delemazure et al. 2022c), thus enabling other researchers to reproduce the impossibility.

Secondly, we provide a human-readable proof for a weakening of Proposition 1 that additionally uses Pareto optimality. This proof is derived by applying the computer-aided approach explained in Section 3.1 and by analyzing MUSes of the corresponding formula. Hence, it showcases the correctness of our code. Unfortunately, the proof of this weaker claim still takes 11 pages (even though the used MUSes only consist of roughly 500 constraints), and we thus have to defer it to the full version.

Thirdly, we have—analogueous to Brandl et al. (2018) and Brandt, Saile, and Stricker (2022)—verified the correctness of our results with the interactive theorem prover Isabelle/HOL (Nipkow, Paulson, and Wenzel 2002). Such interactive theorem provers support much more expressive logics and we can hence formalize the entire theorems with all the mathematical notions expressed in a similar way as in Section 2. For instance, Figure 2 displays our Isabelle formalization of weak representation. Our Isabelle/HOL implementation thus directly derives Proposition 1 as well as Theorems 1 and 2 from the definitions of the axioms. This releases us from the need to check any intermediate steps encoded in Isabelle because Isabelle checks the correctness of these steps for us. Moreover, Isabelle/HOL is highly trustworthy as all proofs have to pass through an inference kernel, which only supports the most basic logical inference steps. Thus, to trust the correctness of our result, one only needs to trust the faithfulness of our Isabelle implementation to the definitions in Section 2. Such formal proofs are widely considered to be the “gold standard” of increasing the trustworthiness of a mathematical result (e.g., Hales et al. 2017). Our formal proof development is available in the *Archive of Formal Proofs* (Delemazure et al. 2022a).

```
weak_rep_for_anon_papp_rules n P k f =
  (anon_PAPP_rule n P k f ∧
   (∀A x. anon_papp_profile n P A ∧
    k * count A {x} ≥ n → count f(A) x ≥ 1))
```

Figure 2: The Isabelle/HOL code for weak representation. Given the number of voters n , the set of parties \mathcal{P} , a target committee size k , and a function f , the code first verifies that f is an anonymous party-approval rule for the given parameters and then requires for every profile A (that is valid for n and \mathcal{P}) and every party x that x has at least one seat in $f(A)$ if at least $\frac{n}{k}$ voters uniquely approve x .

4 Strategyproofness for Unrepresented Voters

Since strategyproofness does not allow for attractive party-approval rules, we consider strategyproofness for unrepresented voters (Definition 4) in this section. Instead of prohibiting all voters from manipulating, this property requires that only voters who do not approve any party in the elected committee cannot manipulate.

As a first result, we prove that CCAV satisfies this axiom and can even be characterized based on strategyproofness for unrepresented voters and weak representation within the class of Thiele rules. Hence, CCAV offers an attractive escape route to Theorem 1.

Theorem 3. *CCAV is the only Thiele rule that satisfies weak representation and strategyproofness for unrepresented voters for all committee sizes k , numbers of parties m , and numbers of voters n .*

Proof. For proving this theorem, we show that CCAV satisfies the given axioms for all k , m , and n (Claim 1), and that no other Thiele rule does so (Claim 2).

Claim 1: We start by proving that CCAV satisfies weak representation and note for this that Aziz et al. (2017) have shown that CCAV satisfies justified representation in the ABC setting. It thus satisfies weak representation for party-approval elections as this axiom is weaker than justified representation and party-approval elections can be seen as special case of ABC elections.

Next, we prove by contradiction that CCAV satisfies strategyproofness for unrepresented voters. Hence, suppose that there are a voter $i \in N$, profiles A^1 and A^2 , and a committee size k such that $CCAV(A^2, k, A_i^1) > CCAV(A^1, k, A_i^1) = 0$ and $A_j^1 = A_j^2$ for all $j \in N \setminus \{i\}$. To simplify notation, let $W^1 = CCAV(A^1, k)$ and $W^2 = CCAV(A^2, k)$, and define $s(W, A) = |\{i \in N : W(A_i) > 0\}|$ as the CCAV-score of a committee W in a profile A . Now, the definition of CCAV requires that $s(W^1, A^1) \geq s(W^2, A^1)$ and $s(W^2, A^2) \geq s(W^1, A^2)$. Moreover, since $W^1(A_i^1) = 0$ and $A_j^1 = A_j^2$ for all voters $j \in N \setminus \{i\}$, it follows that $s(W^1, A^2) \geq s(W^1, A^1)$. Finally, we assumed that $W^2(A_i^1) > 0$, which implies that $s(W^2, A^1) \geq s(W^2, A^2)$ since $A_j^1 = A_j^2$ for all $j \in N \setminus \{i\}$. By combining these inequalities, we obtain $s(W^2, A^2) \geq s(W^1, A^2) \geq s(W^1, A^1) \geq s(W^2, A^1) \geq s(W^2, A^2)$, which implies that all scores are equal. However, lexicographic tie-breaking

implies then that we choose either W^1 or W^2 for both A^1 and A^2 , which contradicts that $W^1 = \text{CCAV}(A^1, k)$ and $W^2 = \text{CCAV}(A^2, k)$.

Claim 2: Next, we show that no other Thiele rule but CCAV satisfies weak representation and strategyproofness for unrepresented voters for all k, m , and n . First, observe that AV clearly fails weak representation. Thus, let f be a w -Thiele rule other than AV and CCAV. We will show that f fails strategyproofness for unrepresented voters. Note for this that there is an index j with $w_1 > w_j$ since f is not AV. We denote with j_0 the smallest such index, which means that $\forall j < j_0, w_j = w_1 = 1$. If $w_{j_0} = 0$, then $j_0 \geq 3$ because f is not CCAV. Let $\mathcal{P} = \{a_1, \dots, a_{j_0}, b_1, \dots, b_{j_0}\}$ be a set of $m = 2j_0$ parties. We construct the profile A with $n = 2 \cdot \binom{2j_0}{j_0} - 2$ voters and set the target committee size to $k = j_0$. The approval ballots of the voters are defined as follows: voter 1 reports $\{a_1, \dots, a_{j_0}\}$, voter 2 reports $\{b_1\}$ and for every set $X \subseteq \mathcal{P}$ with $|X| = j_0$, $X \neq \{a_1, \dots, a_{j_0}\}$, and $X \neq \{b_1, \dots, b_{j_0}\}$, there are two voters who report X as their ballot.

First, note that every party appears in exactly $n_c = 2 \binom{2j_0-1}{j_0-1} - 2$ ballots of the voters $N_c = N \setminus \{1, 2\}$. Consequently, every committee W of size j_0 gets a total of $\sum_{x \in \mathcal{P}} W(x) |\{i \in N_c : x \in A_i\}| = j_0 n_c$ approvals from these voters. We use this fact to compute the scores of a committee W derived from these voters. Observe that the committees $W_A = [a_1, \dots, a_{j_0}]$ and $W_B = [b_1, \dots, b_{j_0}]$ receive a score of $j_0 n_c$ from the voters in N_c because none of them approves all parties in the committee and $w_1 = \dots = w_{j_0-1} = 1$. On the other hand, for every other committee W , there are at least two voters who approve all parties in W . Hence, these voters assign a score of $j_0 - 1 + w_{j_0}$ to the committee. Since the total sum of approvals is constant we derive that the remaining voters in N_c assign at most a score of $j_0(n_c - 2)$ to W . Hence, the score of W among voters in N_c is upper bounded by $j_0 n_c - 2(1 - w_{j_0})$. Finally, if we add the first two voters, W_A obtains a score of $j_0 n_c + j_0 - 1 + w_{j_0}$, W_B of $j_0 n_c + 1 < j_0 n_c + j_0 - 1 + w_{j_0}$ (because either $j_0 \geq 3$ or $j_0 = 2$ and $w_{j_0} > 0$), and the scores of other committees is at most $j_0 n_c - 2(1 - w_{j_0}) + j_0 < j_0 n_c + j_0 - 1 + w_{j_0}$ (since $w_{j_0} < 1$). Hence, $f(A, j_0) = W_A$.

Now, consider the profile A' derived from A by changing the approval ballot of voter 2 to $\{b_1, \dots, b_{j_0}\}$. Then, the score of the committee W_A does not change and the score of W_B is now equal to the score of W_A . Moreover, the same argument as before shows that the score of all other committees is still strictly lower. Hence, committees W_A and W_B are now tied for the win. If the tie-breaking favors W_B over W_A , we thus have $f(A', j_0) = W_B$ and voter 2 can manipulate even though $f(A, j_0, A_2) = 0$. Otherwise, we can exchange the roles of $\{a_1, \dots, a_{j_0}\}$ and $\{b_1, \dots, b_{j_0}\}$. Hence, f fails strategyproofness for unrepresented voters. \square

A natural follow-up question to Theorem 3 is whether party-approval rules other than Thiele rules satisfy strategyproofness for unrepresented voters. We partially answer this question by showing that all sequential Thiele rules (except AV) and all divisor methods based on majoritarian portioning

(except AV) fail this axiom. Hence, even this weak notion of strategyproofness is a challenging axiom for party-approval elections. We defer the proof of this theorem completely to the full version; it works by constructing counterexamples similar to Claim 2 in Theorem 3.

Theorem 4. *All sequential Thiele rules except AV and all divisor methods based on majoritarian portioning except AV fail strategyproofness for unrepresented voters for some committee size k , number of parties m , and number of voters n .*

Remark 4. CCAV becomes highly indecisive if $k \geq m$ since every voter will approve at least one party in the chosen committee. Thus, many seats of the committee will be assigned by the tie-breaking. Hence, CCAV is no attractive rule if $k > m$. Similar arguments show that all w -Thiele rules that have an index j with $w_j = 0$ are strategyproof for unrepresented voters if $k \geq (j - 1)m$: in this case, these rules always choose a committee which guarantees every voter $j - 1$ representatives and strategyproofness for unrepresented voters is trivially satisfied. Consequently, Theorem 3 needs to quantify over the committee size, number of parties, and number of voters.

Remark 5. All results of this section carry over into the ABC setting. For the negative results this follows from the fact that party-approval elections can be seen as a special case of ABC elections (see Figure 1). The first claim of Theorem 3 holds since our proof directly translates into the ABC setting.

5 Conclusion

We study the compatibility of strategyproofness and proportional representation for party-approval multiwinner elections, where a multiset of the parties is chosen based on the voters' approval ballots. First, we prove based on a computer-aided approach that strategyproofness and minimal notions of proportional representation are incompatible for anonymous party-approval rules. Thus, the incompatibility of strategyproofness and proportional representation first observed by Peters (2018) for approval-based committee voting rules (which return sets instead of multisets) also prevails in our more flexible setting. As a second contribution, we investigate a weakening of strategyproofness which requires that only voters who do not approve any member of the committee cannot manipulate. Perhaps surprisingly, almost all commonly studied party-approval rules fail even this very weak strategyproofness notion. Conversely, we can characterize Chamberlin–Courant approval voting as the unique Thiele rule that satisfies strategyproofness for unrepresented voters and a weak representation axiom, thus offering an attractive escape route to our previous impossibility theorem.

Our work offers several directions for future extensions. In particular, we feel that strategyproofness for unrepresented voters deserves more attention; for example, we have to leave it open whether weak proportional representation is compatible with this axiom. Furthermore, one can see strategyproofness and strategyproofness for unrepresented voters as two extreme cases of a parameterization of strategyproofness and it thus might be interesting to consider quantified strategyproofness notions for party-approval elections.

Acknowledgments

This work emerged from a joint retreat of the DSS group of TU Munich, the ALGO group of TU Berlin, and invited visiting scholars. We would like to thank Felix Brandt, Markus Brill, Dominik Peters, and René Romen for their helpful comments. Théo Delemazure was supported by the PRAIRIE 3IA Institute under grant ANR-19-P3IA-0001 (e). Tom Demeulemeester was supported by Research Foundation - Flanders under grant 11J8721N. Jonas Israel was supported by the Deutsche Forschungsgemeinschaft under grant BR 4744/2-1. Patrick Lederer was supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1.

References

- Audemard, G.; and Simon, L. 2018. On the glucose SAT solver. *International Journal on Artificial Intelligence Tools*, 27(01): 1–25.
- Aziz, H.; Bogomolnaia, A.; and Moulin, H. 2019. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*, 753–781.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare*, 48(2): 461–485.
- Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 107–115. IFAAMAS.
- Balinski, M.; and Young, H. P. 2001. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Brookings Institution Press, 2nd edition.
- Belov, A.; and Marques-Silva, J. 2012. MUSer2: An efficient MUS extractor. *Journal on Satisfiability, Boolean Modeling and Computation*, 8: 123–128.
- Bogomolnaia, A.; Moulin, H.; and Stong, R. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2): 165–184.
- Botan, S. 2021. Manipulability of Thiele Methods on Party-List Profiles. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 223–231.
- Brandl, F.; Brandt, F.; Eberl, M.; and Geist, C. 2018. Proving the Incompatibility of Efficiency and Strategyproofness via SMT Solving. *Journal of the ACM*, 65(2).
- Brandl, F.; Brandt, F.; Peters, D.; and Stricker, C. 2021. Distribution Rules Under Dichotomous Preferences: Two Out of Three Ain’t Bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM-EC)*.
- Brandt, F.; Saile, C.; and Stricker, C. 2022. Strategyproof Social Choice When Preferences and Outcomes May Contain Ties. *Journal of Economic Theory*, 202.
- Brill, M.; Gözl, P.; Peters, D.; Schmidt-Kraepelin, U.; and Wilker, K. 2020. Approval-Based Apportionment. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 1854–1861. AAAI Press.
- Brill, M.; Israel, J.; Micha, E.; and Peters, J. 2022. Individual Representation in Approval-Based Committee Voting. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 4892–4899. AAAI press.
- Brill, M.; Laslier, J.-F.; and Skowron, P. 2018. Multiwinner Approval Rules as Apportionment Methods. *Journal of Theoretical Politics*, 30(3): 358–382.
- Delemazure, T.; Demeulemeester, T.; Eberl, M.; Israel, J.; and Lederer, P. 2022a. The Incompatibility of Strategy-Proofness and Representation in Party-Approval Multi-Winner Elections. *Archive of Formal Proofs*. https://isa-afp.org/entries/PAPP_Impossibility.html, Formal proof development.
- Delemazure, T.; Demeulemeester, T.; Eberl, M.; Israel, J.; and Lederer, P. 2022b. Strategyproofness and Proportionality in Party-Approval Multiwinner Elections. *arXiv preprint arXiv:2211.13567*.
- Delemazure, T.; Demeulemeester, T.; Eberl, M.; Israel, J.; and Lederer, P. 2022c. Supplementary material for the paper ”Strategyproofness and Proportionality in Party-Approval Multiwinner Voting”. <https://doi.org/10.5281/zenodo.7356204>. Accessed: 2023-03-07.
- Duddy, C. 2014. Electing a representative committee by approval ballot: An impossibility result. *Economic Letters*, 124(1): 14–16.
- Endriss, U. 2020. Analysis of One-to-One Matching Mechanisms via SAT Solving: Impossibilities for Universal Axioms. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2010. Using Complexity to Protect Elections. *Communications of the ACM*, 53(11): 74–82.
- Geist, C.; and Peters, D. 2017. Computer-aided Methods for Social Choice Theory. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 13, 249–267. AI Access.
- Hales, T.; Adams, M.; Bauer, G.; Dang, T. D.; Harrison, J.; Le Truong, H.; Kaliszyk, C.; Magron, V.; McLaughlin, S.; Nguyen, T. T.; et al. 2017. A formal proof of the Kepler conjecture. In *Forum of mathematics, Pi*, volume 5. Cambridge University Press.
- Kluiving, B.; de Vries, A.; Vrijbergen, P.; Boixel, A.; and Endriss, U. 2020. Analysing Irresolute Multiwinner Voting Rules with Approval Ballots via SAT Solving. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*, 131–138. IOS Press.
- Lackner, M.; and Skowron, P. 2018. Approval-Based Multi-Winner Rules and Strategic Voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 340–346. IJCAI.
- Lackner, M.; and Skowron, P. 2022. Multi-Winner Voting with Approval Preferences. Technical report, arXiv:2007.01795 [cs.GT].
- Nadel, A.; Ryvchin, V.; and Strichman, O. 2014. Accelerated Deletion-based Extraction of Minimal Unsatisfiable Cores. *Journal of Satisfiability, Boolean Modeling, and Computation*, 9: 27–51.

- Nipkow, T.; Paulson, L. C.; and Wenzel, M. 2002. *Isabelle/HOL – A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag.
- Peters, D. 2018. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 1549–1557.
- Peters, D.; and Skowron, P. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*, 793–794. ACM Press.
- Pukelsheim, F. 2014. *Proportional Representation: Apportionment Methods and Their Applications*. Springer.
- Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J. A.; Basanta Val, P.; and Skowron, P. 2017. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 670–676. AAAI Press.
- Tang, P.; and Lin, F. 2009. Computer-aided proofs of Arrow’s and other impossibility theorems. *Artificial Intelligence*, 173(11): 1041–1053.
- Thiele, T. N. 1895. Om Flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, 415–441.

BIBLIOGRAPHY

- F. Aleskerov and E. Kurbanov. Degree of manipulability of social choice procedures. In *Current trends in economics*, pages 13–27. Springer, 1999. [p. 6]
- F. Aleskerov, D. Karabekyan, M. R. Sanver, and V. Yakuba. On the manipulability of voting rules: The case of 4 and 5 alternatives. *Mathematical Social Sciences*, 64:67–73, 2012. [p. 6]
- M. Allais. Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’école américaine. *Econometrica*, 21(4):503–546, 1953. [p. 39]
- P. Anand. Rationality and intransitive preference: Foundations for the modern view. In P. Anand, P. K. Pattanaik, and C. Puppe, editors, *The Handbook of Rational and Social Choice*, chapter 6. Oxford University Press, 2009. [p. 39]
- K. J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition, 1951. 2nd edition 1963. [p. 4]
- N. Aswal, S. Chatterji, and A. Sen. Dictatorial domains. *Economic Theory*, 22(1):45–62, 2003. [pp. 5, 34]
- G. Audemard and L. Simon. On the Glucose SAT solver. *International Journal of Artificial Intelligence Tools*, 27(1):1–25, 2019. [p. 53]
- H. Aziz, F. Brandl, and F. Brandt. Universal Pareto dominance and welfare for plausible utility functions. *Journal of Mathematical Economics*, 60:123–133, 2015a. [pp. 17, 101]
- H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 107–115, 2015b. [pp. 25, 45]
- H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017. [p. 24]
- H. Aziz, F. Brandl, F. Brandt, and M. Brill. On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110:1–18, 2018. [pp. 5, 6, 16, 28, 37, 39, 40, 43, and 48]
- H. Aziz, A. Bogomolnaia, and H. Moulin. Fair mixing: the case of dichotomous preferences. *ACM Transactions on Economics and Computation*, 8(4):18:1–18:27, 2020. [pp. 23, 45]
- M. Balinski and H. P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Brookings Institution Press, 2nd edition, 2001. [pp. 23, 45]
- T. Bandyopadhyay. Multi-valued decision rules and coalitional non-manipulability. *Economics Letters*, 13(1):37–44, 1983. [p. 40]
- S. Barberà. Manipulation of social decision functions. *Journal of Economic Theory*, 15(2): 266–278, 1977a. [p. 40]
- S. Barberà. The manipulation of social choice mechanisms that do not leave “too much” to chance. *Econometrica*, 45(7):1573–1588, 1977b. [pp. 3, 5, 8, 40, and 47]
- S. Barberà. A note on group strategy-proof decision schemes. *Econometrica*, 47(3):637–640, 1979a. [pp. 5, 28]
- S. Barberà. Majority and positional voting in a probabilistic framework. *Review of Economic Studies*, 46(2):379–389, 1979b. [pp. 5, 16, 21, 27, 32, and 47]

- S. Barberà. Strategy-proof social choice. In K. J. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, chapter 25, pages 731–832. Elsevier, 2010. [p. 4]
- S. Barberà, F. Gul, and E. Stacchetti. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61:262–289, 1993. [pp. 5, 34]
- S. Barberà, B. Dutta, and A. Sen. Strategy-proof social choice correspondences. *Journal of Economic Theory*, 101(2):374–394, 2001. [pp. 7, 41]
- S. Barberà, D. Berga, and B. Moreno. Domains, ranges and strategy-proofness: the case of single-dipped preferences. *Social Choice and Welfare*, 39:335–352, 2012. [p. 34]
- S. Barberà, D. Berga, and B. Moreno. Arrow on domain conditions: A fruitful road to travel. *Social Choice and Welfare*, 54:237–258, 2020. [p. 5]
- A. Belov and J. Marques-Silva. MUSer2: An efficient MUS extractor. *Journal on Satisfiability, Boolean Modeling and Computation*, 8:123–128, 2012. [p. 54]
- J.-P. Benoît. Strategic manipulation in voting games when lotteries and ties are permitted. *Journal of Economic Theory*, 102(2):421–436, 2002. [pp. 7, 8, 28, 37, 38, 40, 47, 48, and 183]
- A. Biere. PicoSAT essentials. *Journal on Satisfiability, Boolean Modeling and Computation*, 4: 75–79, 2008. [p. 53]
- D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1): 23–34, 1948. [p. 4]
- D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958. [pp. 4, 40]
- P. R. Blavatsky. Axiomatization of a preference for most probable winner. *Theory and Decision*, 60(1):17–33, 2006. [p. 18]
- C. R. Blyth. Some probability paradoxes in choice from among random alternatives. *Journal of the American Statistical Association*, 67(338):366–373, 1972. [p. 18]
- N. Boehmer and N. Schaar. Collecting, classifying, analyzing, and using real-world ranking data. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1706–1715, 2023. [p. 35]
- A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001. [p. 16]
- A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005. [p. 45]
- A. Boixel and U. Endriss. Automated justification of collective decision making via constraint solving. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 168–176, 2020. [p. 51]
- C. Borda. *Memoire sur les Elections au Scrutin*. Histoire de l’Academie Royale des Sciences, 1784. [pp. 4, 11]
- S. Botan. Manipulability of Thiele methods on party-list profiles. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 223–231, 2021. [p. 45]
- S. J. Brams and P. C. Fishburn. Approval voting. *The American Political Science Review*, 72 (3):831–847, 1978. [pp. 4, 5]
- F. Brandl and F. Brandt. Arrowian aggregation of convex preferences. *Econometrica*, 88(2): 799–844, 2020. [pp. 18, 39, and 101]
- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016. [p. 12]
- F. Brandl, F. Brandt, M. Eberl, and C. Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, 65(2):1–28, 2018. [pp. 5, 16, 17, 28, 37, 39, 40, 43, 44, 48, 49, 51, 54, 57, and 59]

- F. Brandl, F. Brandt, and J. Hofbauer. Welfare maximization entices participation. *Games and Economic Behavior*, 14:308–314, 2019. [pp. 18, 39, and 101]
- F. Brandl, F. Brandt, D. Peters, and C. Stricker. Distribution rules under dichotomous preferences: Two out of three ain’t bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM-EC)*, pages 158–179, 2021. [pp. 37, 40, 43, 45, 51, and 54]
- F. Brandl, F. Brandt, and C. Stricker. An analytical and experimental comparison of maximal lottery schemes. *Social Choice and Welfare*, 58(1):5–38, 2022. [p. 28]
- F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4):793–804, 2015. [pp. 5, 7, 41, 43, 48, 57, and 183]
- F. Brandt. Rolling the dice: Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 1, pages 3–26. AI Access, 2017. [pp. 5, 6, 16, 18, 22, 28, 37, 39, 40, and 101]
- F. Brandt and C. Geist. Finding strategyproof social choice functions via SAT solving. *Journal of Artificial Intelligence Research*, 55:565–602, 2016. [pp. 41, 51]
- F. Brandt and P. Lederer. Characterizing the top cycle via strategyproofness. *Theoretical Economics*, 18(2):837–883, 2023. [p. 42]
- F. Brandt, M. Brill, and P. Harrenstein. Tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2016. [pp. 22, 133]
- F. Brandt, C. Geist, and D. Peters. Optimal bounds for the no-show paradox via SAT solving. *Mathematical Social Sciences*, 90:18–27, 2017. Special Issue in Honor of Hervé Moulin. [p. 51]
- F. Brandt, M. Bullinger, and P. Lederer. On the indecisiveness of Kelly-strategyproof social choice functions. *Journal of Artificial Intelligence Research*, 73:1093–1130, 2022a. [pp. 43, 44]
- F. Brandt, P. Lederer, and R. Romen. Relaxed notions of Condorcet-consistency and efficiency for strategyproof social decision schemes. Technical report, <https://arxiv.org/abs/2201.10418>, 2022b. [pp. 33, 34]
- F. Brandt, C. Saile, and C. Stricker. Strategyproof social choice when preferences and outcomes may contain ties. *Journal of Economic Theory*, 202:105447, 2022c. [pp. 3, 5, 6, 19, 41, 43, 44, 48, 49, 51, and 183]
- F. Brandt, M. Greger, and R. Romen. Towards a characterization of random serial dictatorship. 2023a. Working paper. [p. 51]
- F. Brandt, P. Lederer, and W. Suksompong. Incentives in social decision schemes with pairwise comparison preferences. *Games and Economic Behavior*, 142:266–291, 2023b. [pp. 39, 40, and 50]
- F. Brandt, P. Lederer, and S. Tausch. Strategyproof social decision schemes on super Condorcet domains. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1734–1742, 2023c. [pp. 35, 36]
- M. Brill, P. Gözl, D. Peters, U. Schmidt-Kraepelin, and K. Wilker. Approval-based apportionment. *Mathematical Programming*, 2022. [pp. 8, 23, 25, 44, and 57]
- O. Cailloux and U. Endriss. Arguing about voting rules. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 287–295, 2016. [p. 51]
- D. E. Campbell and J. S. Kelly. A strategy-proofness characterization of majority rule. *Economic Theory*, 22(3):557–568, 2003. [p. 35]
- D. E. Campbell and J. S. Kelly. Gains from manipulating social choice rules. *Economic Theory*, 40(3):349–371, 2009. [p. 6]
- D. E. Campbell and J. S. Kelly. Anonymous, neutral, and strategy-proof rules on the condorcet domain. *Economics Letters*, 128:79–82, 2015. [p. 35]

- D. E. Campbell and J. S. Kelly. Correction to “a strategy-proofness characterization of majority rule”. *Economic Theory Bulletin*, 4(1):121–124, 2016. [p. 35]
- J. R. Chamberlin. An investigation of the relative manipulability of four voting systems. *Behavioral Science*, 30:195–203, 1985. [p. 6]
- S. Chatterji and A. Sen. Tops-only domains. *Economic Theory*, 46:255–282, 2011. [p. 34]
- S. Chatterji and H. Zeng. On random social choice functions with the tops-only property. *Games and Economic Behavior*, 109:413–435, 2018. [p. 34]
- S. Chatterji and H. Zeng. A taxonomy of non-dictatorial domains. *Games and Economic Behavior*, 137:228–269, 2023. [pp. 5, 34]
- S. Chatterji, R. Sanver, and A. Sen. On domains that admit well-behaved strategy-proof social choice functions. *Journal of Economic Theory*, 148(3):1050–1073, 2013. [p. 34]
- S. Chatterji, A. Sen, and H. Zeng. Random dictatorship domains. *Games and Economic Behavior*, 86:212–236, 2014. [pp. 5, 34]
- S. Ching and L. Zhou. Multi-valued strategy-proof social choice rules. *Social Choice and Welfare*, 19(3):569–580, 2002. [p. 40]
- W. J. Cho. Incentive properties for ordinal mechanisms. *Games and Economic Behavior*, 95:168–177, 2016. [pp. 6, 16]
- V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110, 2002. [p. 51]
- V. Conitzer and T. Sandholm. Applications of automated mechanism design. In *Proceedings of the UAI workshop on Bayesian Modeling Applications*, 2003. [p. 51]
- A. H. Copeland. A ‘reasonable’ social welfare function. Mimeo, University of Michigan Seminar on Applications of Mathematics to the Social Sciences, 1951. [p. 11]
- T. Delemazure, T. Demeulemeester, M. Eberl, J. Israel, and P. Lederer. The incompatibility of strategy-proofness and representation in party-approval multi-winner elections. *Archive of Formal Proofs*, 2022. URL https://isa-afp.org/entries/PAPP_Impossibility.html. [p. 54]
- T. Delemazure, T. Demeulemeester, M. Eberl, J. Israel, and P. Lederer. Strategyproofness and proportionality in party-approval multiwinner elections. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5591–5599, 2023. [pp. 45, 46, 53, and 55]
- A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, 2004. [p. 6]
- J. Duggan. A geometric proof of Gibbard’s random dictatorship theorem. *Economic Theory*, 7(2):365–369, 1996. [p. 28]
- J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000. [pp. 3, 5, 7, 8, 40, 47, and 183]
- M. Dummett and R. Farquharson. Stability in voting. *Econometrica*, 29(1):33–43, 1961. [p. 4]
- B. Dutta, H. Peters, and A. Sen. Strategy-proof probabilistic mechanisms in economies with pure public goods. *Journal of Economic Theory*, 106(2):392–416, 2002. [p. 57]
- B. Dutta, H. Peters, and A. Sen. Strategy-proof cardinal decision schemes. *Social Choice and Welfare*, 28(1):163–179, 2007. [p. 28]
- S. Ebadian, A. Kahng, D. Peters, and N. Shah. Optimized distortion and proportional fairness in voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (ACM-EC)*, pages 563–600, 2022. [p. 32]

- L. Ehlers, H. Peters, and T. Storcken. Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences. *Journal of Economic Theory*, 105(2):408–434, 2002. [pp. 5, 16, 21, 34, and 57]
- E. Elkind and M. Lackner. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2019–2025, 2015. [p. 59]
- E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48:599–632, 2017. [p. 22]
- U. Endriss. Analysis of one-to-one matching mechanisms via sat solving: Impossibilities for universal axioms. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020. [p. 51]
- B. Erdamar and M. R. Sanver. Choosers as extension axioms. *Theory and Decision*, 67(4): 375–384, 2009. [p. 19]
- P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. 2017. [pp. 22, 57]
- R. Farquharson. Straightforwardness in voting procedures. *Oxford Economic Papers*, 8(1): 80–89, 1956a. [p. 4]
- R. Farquharson. Strategic information in games and voting. *Information Theory*. Stoneham: Butterworths, 1956b. [p. 4]
- R. Farquharson. *Theory of Voting*. Yale University Press, 1969. [p. 4]
- P. Favardin and D. Lepelley. Some further results on the manipulability of social choice rules. *Social Choice and Welfare*, 26(3):485–509, 2006. [p. 6]
- A. M. Feldman. Manipulation and the Pareto rule. *Journal of Economic Theory*, 21:473–482, 1979a. [p. 41]
- A. M. Feldman. Nonmanipulable multi-valued social choice decision functions. *Public Choice*, 34:177–188, 1979b. [pp. 5, 7]
- A. Filos-Ratsikas and P. B. Miltersen. Truthful approximations to range voting. In *Proceedings of the 10th International Conference on Web and Internet Economics (WINE)*, Lecture Notes in Computer Science (LNCS), pages 175–188. Springer-Verlag, 2014. [p. 32]
- P. C. Fishburn. The irrationality of transitivity in social choice. *Behavioral Science*, 15:119–123, 1970. [p. 4]
- P. C. Fishburn. Even-chance lotteries in social choice theory. *Theory and Decision*, 3(1): 18–40, 1972. [p. 19]
- P. C. Fishburn. *The Theory of Social Choice*. Princeton University Press, 1973. [p. 4]
- P. C. Fishburn. Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26(1): 31–67, 1982. [p. 18]
- P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984. [p. 12]
- C. Gangl, M. Lackner, J. Maly, and S. Woltran. Aggregating expert opinions in support of medical diagnostic decision-making. In *Knowledge Representation for Health Care/Pro-Health (KR4HC)*, pages 56–62, 2019. [p. 22]
- P. Gärdenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13(2): 217–228, 1976. [pp. 3, 41, and 47]
- P. Gärdenfors. On definitions of manipulation of social choice functions. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979. [pp. 7, 19]
- G. Gawron and P. Faliszewski. Using multiwinner voting to search for movies. In *Proceedings of the 19th European Conference on Multi-Agent Systems (EUMAS)*, Lecture Notes in Computer Science (LNCS), pages 134–151. Springer-Verlag, 2022. [p. 22]

- W. V. Gehrlein and D. Lepelley. *Voting Paradoxes and Group Coherence*. Studies in Choice and Welfare. Springer-Verlag, 2011. [p. 35]
- C. Geist and U. Endriss. Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial Intelligence Research*, 40:143–174, 2011. [p. 51]
- C. Geist and D. Peters. Computer-aided methods for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 13, pages 249–267. AI Access, 2017. [p. 51]
- A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973. [pp. v, vii, 3, 4, 27, 42, and 57]
- A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3): 665–681, 1977. [pp. 4, 5, 7, 16, 21, 27, 28, 29, 32, 33, 37, 47, 57, 63, 91, and 101]
- I. J. Good. A note on Condorcet sets. *Public Choice*, 10(1):97–101, 1971. [p. 13]
- A. Gopakumar and S. Roy. Dictatorship on top-circular domains. *Theory and Decision*, 85(3):479–493, 2018. [p. 34]
- U. Grandi and U. Endriss. First-order logic formalisation of impossibility theorems in preference aggregation. *Journal of Philosophical Logic*, 42(4):595–618, 2013. [p. 51]
- B. Grofman. Some notes on voting schemes and the will of majority. *Public Choice*, 7:65–80, 1969. [p. 4]
- T. Hales, M. Adams, G. Bauer, T. D. Dang, J. Harrison, L. T. Hoang, C. Kaliszyk, V. Magron, S. McLaughlin, T. T. Nguyen, Q. T. Nguyen, T. Nipkow, S. Obua, J. Pleso, J. Rute, A. Solovyew, T. H. A. Ta, N. T. Tran, T. D. Trieu, J. Urban, K. Vu, and R. Zumkeller. A formal proof of the Kepler conjecture. In *Forum of Mathematics, Pi*, volume 5, 2017. [p. 54]
- A. Hylland. Strategyproofness of voting procedures with lotteries as outcomes and infinite sets of strategies. Mimeo, 1980. [pp. 4, 5, and 28]
- J. L. Jimeno, J. Pérez, and E. García. An extension of the Moulin No Show Paradox for voting correspondences. *Social Choice and Welfare*, 33(3):343–459, 2009. [p. 19]
- D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979. [p. 39]
- J. Kavner and L. Xia. Strategic behavior is bliss: Iterative voting improves social welfare. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021. [p. 6]
- J. S. Kelly. Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2):439–446, 1977. [pp. 3, 5, 18, 19, and 40]
- B. Kluiving, A. de Vries, P. Vrijbergen, A. Boixel, and U. Endriss. Analysing irresolute multiwinner voting rules with approval ballots via SAT solving. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*, 2020. [pp. 45, 51]
- U. Kumar, S. Roy, A. Sen, S. Yadav, and H. Zeng. Local global equivalence in voting models: A characterization and applications. *Theoretical Economics*, 16(1195–1220), 2021. [p. 6]
- M. Lackner and P. Skowron. Approval-based multi-winner rules and strategic voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 340–346, 2018. [pp. 25, 45]
- M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences*. Springer-Verlag, 2023. [pp. 22, 23, and 57]
- G. Laffond, J.-F. Laslier, and M. Le Breton. A theorem on symmetric two-player zero-sum games. *Journal of Economic Theory*, 72(2):426–431, 1997. [p. 12]
- J.-F. Laslier. In silico voting experiments. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, chapter 13, pages 311–335. Springer-Verlag, 2010. [p. 35]

- J.-F. Laslier and M. R. Sanver, editors. *Handbook on Approval Voting*. Studies in Choice and Welfare. Springer-Verlag, 2010. [p. 22]
- M. Le Breton. On the uniqueness of equilibrium in symmetric two-player zero-sum games with integer payoffs. *Économie publique*, 17(2):187–195, 2005. [p. 12]
- P. Lederer. Strategyproof randomized social choice for restricted sets of utility functions. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 306–312, 2021. [pp. 37, 38, and 49]
- M. J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, 1989. [p. 39]
- I. MacIntyre and P. K. Pattanaik. Strategic voting under minimally binary group decision functions. *Journal of Economic Theory*, 25(3):338–352, 1981. [p. 40]
- T. Majumdar. Choice and revealed preference. *Econometrica*, 24(1):71–73, 1956. [p. 4]
- S. Maus, H. Peters, and T. Storcken. Anonymous voting and minimal manipulability. *Journal of Economic Theory*, 135(1):533–544, 2007. [p. 6]
- K. May. A set of independent, necessary and sufficient conditions for simple majority decisions. *Econometrica*, 20(4):680–684, 1952. [p. 5]
- R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 4. 2017. [p. 6]
- R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 823–828, 2010. [p. 6]
- J. Mellon. Tactical voting and electoral pacts in the 2019 UK general election. *Political Review Studies*, 20(3):504–516, 2022. [p. 3]
- T. Mennle and S. Seuken. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory*, 191:105–144, 2021. [p. 17]
- L. N. Merrill. Parity dependence of a majority rule characterization on the condorcet domain. *Economics Letters*, 112(3):259–261, 2011. [p. 35]
- M. Mittelman, B. Maubert, A. Murano, and L. Perrussel. Automated synthesis of mechanisms. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 426–432, 2022. [p. 51]
- H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980. [pp. 4, 5, and 34]
- H. Moulin. *The Strategy of Social Choice*. North-Holland, 1983. [p. 3]
- E. Muller and M. A. Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14(2):412–418, 1977. [p. 6]
- Y. Murakami. *Logic and Social Choice*. Monographs in Modern Logic. Routledge and Kegan Paul, 1968. [p. 4]
- A. Nadel, V. Ryzhkin, and O. Strichman. Accelerated deletion-based extraction of minimal unsatisfiable cores. *Journal of Satisfiability, Boolean Modeling, and Computation*, 9:27–51, 2014. [p. 54]
- S. Nandeibam. An alternative proof of Gibbard’s random dictatorship result. *Social Choice and Welfare*, 15(4):509–519, 1997. [p. 28]
- S. Nandeibam. The structure of decision schemes with cardinal preferences. *Review of Economic Design*, 17(3):205–238, 2013. [pp. 5, 28]
- H. Narasimhan, S. Agarwal, and D. C. Parkes. Automated mechanism design without money via machine learning. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 433–439, 2016. [p. 51]
- J. F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences (PNAS)*, 36:48–49, 1950. [p. 4]

- K. Nehring. Monotonicity implies generalized strategy-proofness for correspondences. *Social Choice and Welfare*, 17(2):367–375, 2000. [pp. 3, 5, 36, 41, 57, and 183]
- K. Nehring and C. Puppe. Efficient and strategy-proof voting rules: A characterization. *Games and Economic Behavior*, 59(1):132–153, 2007. [pp. 5, 34]
- T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL – A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, 2002. [p. 54]
- S. Özyurt and M. R. Sanver. A general impossibility result on strategy-proof social choice hyperfunctions. *Games and Economic Behavior*, 66(2):880–892, 2009. Special Section in Honor of David Gale. [p. 41]
- D. J. Packard. Cyclical preference logic. *Theory and Decision*, 14(4):415–426, 1982. [p. 18]
- P. K. Pattanaik. On the stability of sincere voting situations. *Journal of Economic Theory*, 6: 558–574, 1973. [p. 4]
- E. A. Pazner and E. Wesley. Cheatproofness properties of the plurality rule in large elections. *The Review of Economic Studies*, 45(1):85–91, 1978. [p. 5]
- B. Peleg. A note on manipulability of large voting schemes. *Theory and Decision*, 11:401–412, 1979. [p. 5]
- D. Peters. Proportionality and strategyproofness in multiwinner elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 1549–1557, 2018. [pp. 25, 45, and 59]
- D. Peters, A. D. Procaccia, A. Psomas, and Z. Zhou. Explainable voting. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, pages 1525–1534, 2020. [p. 51]
- H. Peters, S. Roy, A. Sen, and T. Storcken. Probabilistic strategy-proof rules over single-peaked domains. *Journal of Mathematical Economics*, 52:123–127, 2014. [p. 34]
- H. Peters, S. Roy, S. Sadhukhan, and T. Storcken. An extreme point characterization of random strategy-proof and unanimous probabilistic rules over binary restricted domains. *Journal of Mathematical Economics*, 69:84–90, 2017. [p. 34]
- H. Peters, S. Roy, and S. Sadhukhan. Unanimous and strategy-proof probabilistic rules for single-peaked preference profiles on graphs. *Mathematics of Operations Research*, 46(2): 811–833, 2021. [p. 34]
- J. Picot and A. Sen. An extreme point characterization of random strategy-proof social choice functions: The two alternative case. *Economics Letters*, 115(1):49–52, 2012. [p. 5]
- G. Pritchard and M. C. Wilson. Exact results on the manipulability of positional scoring rules. *Social Choice and Welfare*, 29:487–513, 2007. [p. 6]
- A. D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 836–841, 2010. [p. 32]
- F. Pukelsheim. *Proportional Representation: Apportionment Methods and Their Applications*. Springer, 2014. [p. 45]
- M. Pycia and M. U. Ünver. Decomposing random mechanisms. *Journal of Mathematical Economics*, 61:21–33, 2015. [pp. 34, 35]
- S. Rabinovich, S. Obraztsova, O. Lev, E. Markakis, and J. S. Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1007–1013, 2015. [p. 6]
- M. Regenwetter, B. Grofman, A. A. J. Marley, and I. M. Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006. [p. 35]
- A. Reinjoud and U. Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 635–644, 2012. [p. 6]

- S. Roy and S. Sadhukhan. A unified characterization of the randomized strategy-proof rules. *Journal of Economic Theory*, pages 105–131, 2020. [pp. 5, 34, and 35]
- S. Roy and T. Storcken. A characterization of possibility domains in strategyproof voting. *Journal of Mathematical Economics*, 84:46–55, 2019. [pp. 5, 34]
- S. Roy, S. Sadhukhan, and A. Sen. Recent results on strategy-proofness of random social choice functions. In S. Borkotokey, R. Kumar, D. Mukherjee, K. S. M. Rao, and S. Sarangi, editors, *Game Theory and Networks*, chapter 3, pages 63–87. Springer, 2022. [pp. 5, 21, 28, and 34]
- L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017. [p. 25]
- M. R. Sanver and W. S. Zwicker. One-way monotonicity as a form of strategy-proofness. *International Journal of Game Theory*, 38(4):553–574, 2009. [p. 6]
- A. Saporiti. Strategy-proofness and single-crossing. *Theoretical Economics*, 4(2):127–163, 2009. [p. 34]
- S. Sato. On strategy-proof social choice correspondences. *Social Choice and Welfare*, 31:331–343, 2008. [p. 40]
- S. Sato. Circular domains. *Review of Economic Design*, pages 331–33, 2010. [p. 34]
- S. Sato. A sufficient condition for the equivalence of strategy-proofness and non-manipulability by preferences adjacent to the sincere one. *Journal of Economic Theory*, 148:259–278, 2013. [pp. 6, 36]
- S. Sato. A fundamental structure of strategy-proof social choice correspondences with restricted preferences over alternatives. *Social Choice and Welfare*, 42(4):831–851, 2014. [p. 40]
- M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975. [pp. v, vii, 3, 4, 27, 42, and 57]
- M. C. Schmidtlein and U. Endriss. Voting by axioms. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 2067–2075, 2023. [p. 51]
- T. Schwartz. *The Logic of Collective Choice*. Columbia University Press, 1986. [p. 13]
- A. Sen. A possibility theorem on majority decisions. *Econometrica*, 34(2):491–499, 1966. [p. 4]
- A. Sen. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIEs*, 2(4):515–527, 2011. [pp. 17, 28]
- A. K. Sen. *Collective Choice and Social Welfare*. North-Holland, 1970. [p. 4]
- A. K. Sen. Social choice theory: A re-examination. *Econometrica*, 45(1):53–89, 1977. [p. 13]
- D. A. Smith. Manipulability measures of common social choice functions. *Social Choice and Welfare*, 16(4):639–661, 1999. [p. 6]
- J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973. [p. 13]
- E. S. Staveley. *Greek and Roman Voting and Elections*. Cornell Univserity Press, 1972. [p. 3]
- L. B. Stephenson, J. H. Aldrich, and A. Blais. *The Many Faces of Strategic Voting: Tactical Behavior in Electoral Systems Around the World*. University of Michigan Press, 2018. [p. 3]
- Y. Tanaka. An alternative proof of Gibbard’s random dictatorship theorem. *Review of Economic Design*, 8:319–328, 2003. [p. 28]
- P. Tang and F. Lin. Computer-aided proofs of Arrow’s and other impossibility theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009. [p. 51]
- A. D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005. [pp. 3, 4, and 19]

- T. N. Thiele. Om flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, pages 415–441, 1895. [p. 23]
- W. Vickrey. Utility, strategy, and social decision rules. *Quarterly Journal of Economics*, 74(4): 507–535, 1960. [p. 4]
- R. Wilson. An axiomatic model of logrolling. *The American Economic Review*, 59(3):331–341, 1969. [p. 4]
- L. Xia. The impact of a coalition: Assessing the likelihood of voter influence in large elections. In *Proceedings of the 24th ACM Conference on Economics and Computation (ACM-EC)*, page 1156, 2023. [p. 6]
- R. Zeckhauser. Voting systems, honest preferences, and Pareto optimality. *American Political Science Review*, 67(3):934–946, 1973. [p. 4]