On the Indecisiveness of Kelly-Strategyproof Social Choice Functions

Felix Brandt
Technische Universität München
München, Germany
brandtf@in.tum.de

Martin Bullinger
Technische Universität München
München, Germany
bullinge@in.tum.de

Patrick Lederer
Technische Universität München
München, Germany
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 unattractive single-valued 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 a Condorcet loser in at least one profile.

KEYWORDS

Strategyproofness; Social Choice Correspondences; Kelly's Preference Extension

ACM Reference Format:
Felix Brandt, Martin Bullinger, and Patrick Lederer. 2021. On the Indecisiveness of Kelly-Strategyproof Social Choice Functions . In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

Whenever a group of multiple 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., 14, 19].

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 [24] and Satterthwaite [34] have shown that only extremely restricted 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., 3, 5, 16, 17, 22, 25, 30, 36]. 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., 5, 6, 16, 17, 33].[1] In this paper, we are concerned with a rather weak—but natural and intuitive—notion of strategyproofness attributed to Kelly [25]. Several attractive SCFs have been shown to be strategyproof for this notion when preferences are strict [10, 12]. 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 [20] 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 [15, 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. [9] 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 [25] and Barberà [3] 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

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

[1] We refer to Barberà [4] and Brandt et al. [15] for a more detailed overview over this extensive stream of research.
[2] For example, Brandt et al. [15] have shown that Pareto-optimality is incompatible with anonymity and a notion of strategyproofness that is slightly stronger than Kelly's.

own [see, e.g., 29].[3] In subsequent work by MacIntyre and Pattanaik [28] and Bandyopadhyay [1], 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 [2] 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. Taylor [36, Th. 8.1.2] has shown 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 [10], who showed that every SCF that returns a (strict) Condorcet winner whenever one exists fails to be strategyproof. More recently, Brandt et al. [15] have shown with the help of computers that every Pareto-optimal SCF whose outcome only depends on the pairwise majority margins can be manipulated.

In this paper, we study strategyproofness in three broad classes of SCFs. These classes are *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 pair of classes, their intersection is non-empty, and 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 ranked-based SCFs, this entails an impossibility whereas for support-based SCFs it demonstrates a high degree of indecisiveness. For non-imposing SCFs, we show that strategyproofness implies that a Condorcet loser has to be returned in at least one preference profile. The latter result remarkably holds without imposing fairness conditions such as anonymity or neutrality. Even though these 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., 9, 10, 30]. 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.

Our results can also be interpreted in the context of randomized social choice. When transferred to this setting, Kelly-strategyproofness is weaker than weak *SD*-strategyproofness and we thus obtain three strong impossibilities.

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, transitive, and reflexive 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., if $x \sim_i y$ if and only if $x \succsim_i y$ and $y \succsim_i x$. We compactly represent a preference relation as a comma-separated list, 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, $[x \dots y]\colon a, b, c$ means that all voters $i \in [x \dots y]$ prefer a to b and b to c. We omit the brackets for singleton sets. For two alternatives $x, y \in A$, the pairwise *support* of x over y is defined as the number of voters who strictly prefer x to y, i.e., $s_{xy}(R) = |\{i \in N : x \succ_i y\}|$.

Our central objects of study are *social choice functions (SCFs)*, or so-called social choice correspondences, which map a preference profile to a non-empty set of alternatives, i.e., functions of the form $f : \mathcal{R}^n \mapsto 2^A \setminus \emptyset$. The mere mathematical description of SCFs is so general that it allows for rather undesirable functions. We now introduce a number of axioms in order to narrow down the set of SCFs. The most basic fairness condition is anonymity, which

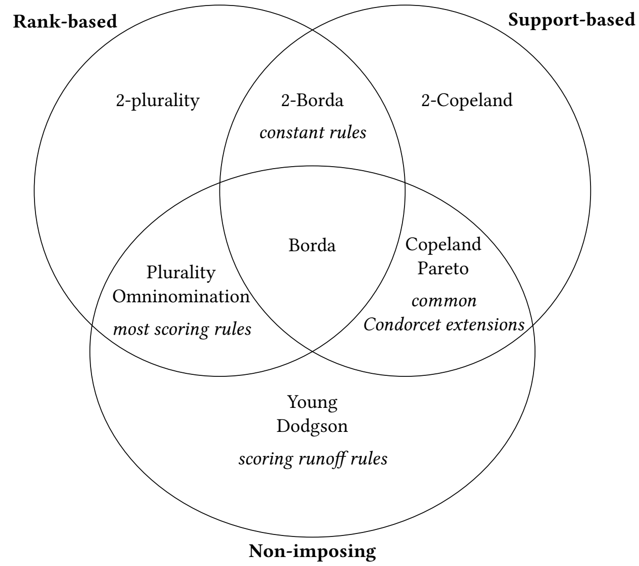

Figure 1: The classes of rank-based, support-based, and non-imposing SCFs and typical examples. 2-plurality, 2-Copeland, and 2-Borda return all alternatives whose respective score is at least as large as the second-highest score. All scoring rules except Borda's rule are rank-based, non-imposing, but 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.

[3]This is acknowledged by Kelly [25] 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."

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 $R_i = R'_{\pi(i)}$ 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. 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 [6]. 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 is near-unanimity. 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 onto. 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 $a \in A$ that wins all pairwise majority comparisons, i.e., $s_{ax}(R) > s_{xa}(R)$ for all $x \in A \setminus \{a\}$. An SCF is *Condorcet-consistent* or a so-called *Condorcet extension* if it uniquely returns a Condorcet winner whenever one exists. Analogously, one can define a Condorcet loser by requiring that $s_{xa}(R) > s_{ax}(R)$ for all $x \in A \setminus \{a\}$. An SCF f satisfies the *Condorcet loser property* if $x \notin f(R)$ whenever x is a 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 could be justified by arguing that both properties affect exactly the same number of preference profiles, but the Condorcet loser property only excludes a single alternative (and leaves otherwise a lot of freedom) whereas Condorcet-consistency completely determines the (singleton) choice set.

While the axioms so far make reference to the entire preference profile, there are also concepts that only refer to the preferences of a single voter. One such concept that is particularly important in our context is that of a nominator. A voter is a *nominator* if $f(R)$ always contains at least one of his most preferred alternatives. A nominator is a weak dictator in the sense that he can always force an alternative into the choice set by declaring it his uniquely most preferred one.

2.1 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 $\bar{r}(\succsim_i, x) = |\{y \in A \colon y \succsim_i x\}|$ [26]. In contrast, there are multiple possibilities how to define the rank in the presence of ties. We define a very weak notion of ranked-basedness for weak preferences, making our results only stronger. To this end, define the *rank tuple* of x with respect to $\succsim_i$ as

$$\begin{aligned} r(\succsim_i, x) &= (\bar{r}(\succ_i, x), \bar{r}(\sim_i, x)) \\ &= (|\{y \in A \colon y \succ_i x\}|, |\{y \in A \colon y \sim_i x\}|). \end{aligned}$$

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 a which contains the rank tuple of a 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.

A similar line of thought leads to support-basedness, which is based on the pairwise support of an alternative x against another one 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 it yields $f(R) = f(R')$ for all preference profiles $R, R' \in \mathcal{R}^n$ with $s^*(R) = s^*(R')$. Note that support-basedness generalizes Fishburn's C2 to weak preferences [21]. Hence, many well-known SCFs such as Borda's rule, Kemeny's rule, the Simpson-Kramer rule, Nanson's rule, Schulze's rule, and the essential set 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., 15]. For example, the Pareto rule is support-based, but fails to be pairwise.

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 manipulation depending on what is assumed about the voters' preferences over sets of alternatives. Here, we make a simple and natural assumption: voter i weakly prefers set X to set Y, denoted by $X \succsim_i Y$, if and only if $x \succsim y$ for all $x \in X, y \in Y$. Thus, the strict part of this preference extension is

$$X \succ_i Y \text{ if and only if for all } x \in X, y \in Y, x \succsim_i y \text{ and there are } x' \in X, y' \in Y \text{ with } x' \succ_i y'.$$

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., 15, 23].

3 RESULTS

The unifying theme of our results is that strategyproofness requires a large degree of indecisiveness. In more detail, we show that every voter is a nominator for all ranked-based and support-based SCFs that satisfy Pareto-optimality and strategyproofness. For the very broad class of non-imposing SCFs, we show that every strategyproof SCF violates the Condorcet loser property. Due to space restrictions, we defer the proofs of all auxiliary lemmas and non-trivial claims in the remarks to the extended version of this paper [13].

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 such a function f and a voter $i \in N$ who is no nominator for f. Our first lemma shows that these assumptions imply that f satisfies near unanimity.

Lemma 1. *Let f be an anonymous, Pareto-optimal, and strategyproof SCF that is defined for $m \geq 3$ alternatives and $n \geq 2$ voters. If some voter is no nominator for f, then f satisfies near unanimity.*

For the proof of this lemma, we consider an arbitrary SCF f that satisfies all required axioms, and a voter i who is no nominator for f. This means that there is a profile R such that $f(R)$ does not contain any of voter i's most preferred alternatives. Next, we construct a profile R' in which $f(R') = \{a\}$ but a is not among the most preferred alternatives of voter i. Based on the profile R', we derive then that $n-1$ voters can ensure that a is the unique winner by submitting it as a uniquely most preferred alternative. Finally, we show that f satisfies near unanimity by generalizing this observation from a single alternative to all alternatives.

Lemma 1 can be interpreted in various appealing ways. For instance, one can see it as a push-down lemma that allows a single voter to weaken the unique winner in his preference relation. Moreover, this lemma shows that, under the given assumptions, indecisiveness for a single preference profile of a particularly simple type entails a large degree of indecisiveness for the entire domain of preference profiles. More precisely, if an alternative is not chosen uniquely even if $n-1$ voters uniquely prefer it the most, then all voters are nominators.

Remark 1. There is also a variant of Lemma 1 without anonymity. Then, an alternative is the unique winner if all voters but the non-nominator prefer it uniquely the most. Thus, requiring the absence of nominators for a strategyproof and Pareto-optimal SCF implies near unanimity.

Remark 2. Remarkably, many impossibility results rule out that every voter is a nominator. For instance, Duggan and Schwartz [17], Benoît [6], and Sato [32] invoke axioms prohibiting that every voter is a nominator. Moreover, a crucial step in the computer-generated proofs of Theorem 3.1 by Brandl et al. [8] and Theorem 1 by Brandt et al. [15] is to show that there is some voter who is no nominator. Lemma 1 gives intuition about why these assumptions and observations are important.

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 the theorem—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., 8, 9, 15]. 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 often hold when restricting the domain of admissible profiles by prohibiting artificial constructs such as completely indifferent voters. Note that, in our proofs, 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 $n \geq 3$.*

Proof. Assume for contradiction that there is a rank-based SCF f that satisfies strategyproofness and Pareto-optimality and that is defined for fixed numbers of voters $n \geq 3$ and alternatives $m \geq 4$. We derive a contradiction to this assumption by proving two claims: on the one hand, there is a voter who is no 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. Consider therefore the following three profiles in which $X = A \setminus \{a, b, c, d\}$.

R^1: 1: $\{a, b\}, X, \{c, d\}$ 2: $\{c, d\}, X, \{a, b\}$
$[3 \dots n]$: $a, \{b, c, d\}, X$

R^2: 1: $\{a, c\}, X, \{b, d\}$ 2: $\{b, d\}, X, \{a, c\}$
$[3 \dots n]$: $a, \{b, c, d\}, X$

R^3: 1: $\{a, d\}, X, \{b, c\}$ 2: $\{b, c\}, X, \{a, d\}$
$[3 \dots n]$: $a, \{b, c, d\}, X$

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 no nominator for f.

Claim 2: Every voter is a nominator for f

Assume for contradiction that a voter is no nominator for f and consider the profiles $R^{k,1}$ and $R^{k,2}$ for $k \in \{1, \dots, n\}$.

$R^{k,1}$: 1: $\{c, d\}, X, b, a$ $[2 \dots k]$: $\{a, b\}, X, c, d$
$[k{+}1 \dots n]$: a, X, b, c, d

$R^{k,2}$: 1: $\{b,d\}, X, c, a$ $[2 \ldots k]$: $\{a,b\}, X, c, d$
$[k{+}1 \ldots n]$: a, X, b, c, d

We prove by induction on $k \in \{1,\ldots,n\}$ that $f(R^{k,1}) = f(R^{k,2}) = \{a\}$. The case $k = n$ yields a contradiction to Pareto-optimality as a is Pareto-dominated by b in $R^{n,1}$.

The base case $k = 1$ follows because $n-1$ voters prefer a uniquely the most in both $R^{1,1}$ and $R^{1,2}$. Therefore, Lemma 1 implies that $f(R^{1,1}) = f(R^{1,2}) = \{a\}$. Assume now that the claim is true for some fixed $k \in \{1,\ldots,n-1\}$.

By induction and strategyproofness, $f(R^{k+1,1}) \subseteq \{a,b\}$ since otherwise voter $k+1$ can manipulate by switching back to $R^{k,1}$. Next, we derive the profile $R^{k,3}$ from $R^{k,2}$ by assigning voter $k+1$ the preference $\{a,c\}, X, b, d$. Formally, $R^{k,3}$ is defined as follows.

$R^{k,3}$: 1: $\{b,d\}, X, c, a$ $[2 \ldots k]$: $\{a,b\}, X, c, d$
$k{+}1$: $\{a,c\}, X, b, d$ $[k{+}2 \ldots n]$: a, X, b, c, d

The induction hypothesis entails that $f(R^{k,2}) = \{a\}$ and therefore, strategyproofness implies that $f(R^{k,3}) \subseteq \{a,c\}$; otherwise, voter $k+1$ could manipulate by switching back to $R^{k,2}$. Next, we apply rank-basedness to conclude that $f(R^{k+1,1}) = \{a\}$ as $r^*(R^{k+1,1}) = r^*(R^{k,3})$. Finally, $R^{k+1,2}$ evolves from $R^{k+1,1}$ by having voter 1 change his preferences. As a is the uniquely least preferred alternative of this voter, strategyproofness implies that $f(R^{k+1,2}) = \{a\}$ as any other outcome benefits voter 1. □

Remark 3. 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.[4] Furthermore, the Pareto rule is rank-based if $m \leq 3$, and if $m = 4$ and $n \leq 2$, which entails that the bounds on m and n are tight if considered simultaneously. By contrast, the theorem is also true if $m \geq 5$ and $n = 2$. More details can be found in the extended version.

Remark 4. Theorem 1 is only an impossibility because of the bad compatibility of rank-basedness and Pareto-optimality in Claim 1, independently of strategyproofness. In contrast, the main consequence of strategyproofness is indecisiveness as captured in Claim 2. This follows as 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 the omninomination rule satisfies then all required axioms [15, Remark 6]. In contrast, the proof of Claim 2 shows that more decisive rank-based SCFs violate strategyproofness if near unanimity is already sufficient for a contradiction.

Remark 5. Theorem 1 holds also under weaker versions of rank-basedness. First, it 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 preferences $\succsim_i, \succsim_i'$ 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.

[4]We define Borda's rule as the SCF that chooses all alternatives that maximize $m \cdot n - \sum_{i \in N} \bar{r}(\succ_i, a)$. This definition agrees with the standard notation used in literature on the strict domain and generalizes it to the weak domain.

Remark 6. 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 a variant of the 2-plurality rule, which chooses the two alternatives that are top-ranked by the most voters, is rank-based, Pareto-optimal, and strategyproof. However, no voter is a nominator for this rule. A formal definition and proofs for the properties can be found in the extended version.

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 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 no nominator, a single voter can determine the choice set. However, this 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.

Therefore, we focus on the case $n \geq 3$ and assume for contradiction that a voter is no nominator for f. We derive from this assumption by an induction on $k \in \{1,\ldots,n-1\}$ that $n-k$ voters can determine a unique winner by uniquely top-ranking it. 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 as this lemma states that f satisfies near unanimity. Next, we assume that our claim holds for a fixed $k \in \{1,\ldots n-2\}$ and prove that also $n-(k+1)$ voters can determine the winner uniquely. 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 near unanimity.

$R^{k,1}$: $[1 \ldots k]$: a, X, c, b $k{+}1$: c, X, b, a
$[k{+}2 \ldots n]$: a, b, X, c

Next, we aim to reverse the preferences of the voters $i \in [k{+}2 \ldots n]$ 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, strategyproofness implies that $f(R^{k,2}) \subseteq \{a, b\}$ as otherwise, voter $k+2$ can manipulate by reverting this modification. 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 $f(R^{k,3}) = f(R^{k,2}) \subseteq \{a, b\}$. As a subset of the least preferred alternatives of voter $k+1$ is chosen for $R^{k,3}$, strategyproofness 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 $f(R^{k,4}) \subseteq \{a, b\}$.

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

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

$R^{k,4}$: $[1 \dots k]$: a, X, c, b $\quad$ k+1: c, X, b, a
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. As $n-k$ 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 $R^{k,5}$.

$R^{k,5}$: $[1 \dots k]$: a, X, c, b $\quad$ k+1: c, X, b, a
$[k+2 \dots n]$: b, a, X, c

$R^{k,6}$: $[1 \dots k]$: a, X, c, b $\quad$ k+1: b, X, a, c
$[k+2 \dots n]$: b, a, X, c

So far, we have found a profile in which b is uniquely chosen when only $n-(k+1)$ voters prefer it uniquely the most. Next, we show that b is always the unique winner if the voters $i \in [k+2 \dots n]$ prefer it uniquely the most. Therefore, 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. As $f(R^{k,5}) = \{b\}$ and b is the worst alternative for these voters, strategyproofness implies that $f(R^{k,7}) = \{b\}$.

$R^{k,7}$: $[1 \dots k]$: c, X, a, b $\quad$ k+1: c, X, b, a
$[k+2 \dots n]$: b, a, X, c

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 strategyproofness 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 $k+1$.

$R^{k,8}$: $[1 \dots k]$: c, X, a, b $\quad$ k+1: c, X, b, a
$[k+2 \dots n]$: b, X, c, a

$R^{k,9}$: $[1 \dots k]$: c, X, a, b $\quad$ 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, and the voters $i \in [k+2 \dots n]$ can reorder all alternatives in $A \setminus \{b\}$ without affecting the choice set because of strategyproofness. 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 and reordering the voters shows that every alternative is chosen if it is uniquely top-ranked by $n-(k+1)$ voters. This completes the induction step and consequently, we derive that every voter is a nominator for a support-based SCF that satisfies strategyproofness and Pareto-optimality. ◻

Remark 7. 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 chooses the lexicographic smallest Pareto-optimal alternative satisfies Pareto-optimality and support-basedness but violates strategyproofness and that every voter is a nominator. For defining an SCF that satisfies Pareto-optimality and strategyproofness but violates support-basedness and that every voter is a nominator, we define a transitive dominance relation by slightly strengthening Pareto-dominance by allowing additionally 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 prefers b strictly to a. Therefore, 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 the extended version 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 8. Brandt et al. [15, Th. 5.4] have shown that there is no pairwise, Pareto-optimal and strategyproof SCF if $m \geq 3$ and $n \geq 3$. This result immediately follows from Theorem 2 by observing that strategyproofness, pairwiseness, and Pareto-optimality rule out that every voter is a nominator. For this, it suffices to find a preference profile in which a Pareto-dominates b and another profile with the same majority margins where b is uniquely top-ranked by a voter.

Remark 9. 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 10. An important subclass of support-based SCFs are majority-based SCFs, which Fishburn [21] calls C1 functions. They only rely on the majority relation $R_M = \{(a, b) \in A^2 \colon s_{ab}(R) \geq s_{ba}(R)\}$ of the input profile R to compute the choice set. For majority-based SCFs, an even more severe impossibility holds: there is no majority-based SCF that satisfies non-imposition and strategyproofness if $m \geq 3$ and $n \geq 3$. Even though this statement does not require Pareto-optimality and therefore cannot use Lemma 1, the result follows from a proof similar to the one of Theorem 2. See the extended version for more details.

Remark 11. 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 (see, e.g, [12] for more details).

Remark 12. 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 7, represents a counterexample.

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 a Condorcet loser in at least one preference profile and thus violate the Condorcet loser property. In the presence of neutrality (symmetry among alternatives), non-imposition can be seen as a decisiveness requirement. Accordingly, the theorem identifies a tradeoff between decisiveness and the undesirable property of selecting Condorcet losers. Similarly to Theorem 1 and Theorem 2, we start with a lemma that allows a voter to push down the unique winner. Since we do not require Pareto-optimality in this section, we cannot use Lemma 1 and therefore propose a new lemma that uses non-imposition instead.

Lemma 2. *Let f denote a strategyproof SCF for $n \geq 3$ voters that satisfies non-imposition and the Condorcet loser property. Then, $f(R) = \{a\}$ for every preference profile R and every alternative $a \in A$ such that more than half of the voters in R prefer the alternative a uniquely the most.*

As first step of the proof of this lemma, we show that an alternative is the unique winner for such an SCF if it is uniquely top-ranked by every voter. Next, we prove that a single voter can make his uniquely best alternative a into his uniquely worst one if more than half of the voters still prefer a uniquely the most after this modification. Repeatedly applying this argument leads to a profile in which $\lceil \frac{n+1}{2} \rceil$ voters prefer a uniquely the most, whereas the remaining voters prefer a uniquely the least and a is the unique winner. Finally, we can apply strategyproofness to turn this profile into any other profile in which a is uniquely top-ranked by a majority of the voters without changing the choice set.

Note that the Condorcet loser property allows for a significantly stronger push-down lemma than Pareto-optimality, even though it only requires that a single alternative may not be chosen. The reason is that an absolute majority of voters can exclude every alternative from the choice set. We use Lemma 2 to show that there is no strategyproof SCF that satisfies the Condorcet loser property and non-imposition. In the interest of space and simplicity, we prove the theorem with an induction using completely indifferent voters. Nevertheless, the theorem—just as our other results—is rather robust to domain restrictions and holds, for instance, also without 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. We prove the statement by induction over $n \geq 4$.

Induction basis: Assume for contradiction that f is a strategyproof SCF for $n = 4$ voters and $m \geq 3$ alternatives that satisfies the Condorcet loser property and non-imposition. Consider the profile R^1 shown in the sequel. By Lemma 2, $f(R^1) = \{a\}$.

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

Moreover, c is the Condorcet loser in R^1, even if voter 1 is indifferent between a and c. Thus, we replace the preference of voter 1 with $\{a, c\}, X + b$, where $X + b = X \cup \{b\}$, to derive the profile R^2. As consequence, $c \notin f(R^2)$ due to the Condorcet loser property, and strategyproofness implies that $f(R) \subseteq \{a, c\}$. Otherwise, an alternative in $X + b$ is chosen and voter 1 can manipulate by reverting back to R^1. Hence, we deduce that $f(R^2) = \{a\}$.

R^2: 1: $\{a, c\}, X{+}b$ 2: a, b, X, c 3: a, b, X, c 4: b, X, c, a

As next step, we let voter 2 change his preference to a, c, X, b and voter 4 change his preference to $c, X, \{a, b\}$ in order to make b the Condorcet loser. By applying these modifications subsequently, it follows from strategyproofness that the choice set does not change: otherwise, voter 2 can manipulate by undoing this step since a is his best alternative after the modification, or voter 4 can manipulate by applying the modification since a is his least preferred alternative in R^2. Hence, these steps result in the profile R^3 with $f(R^3) = \{a\}$.

R^3: 1: $\{a, c\}, X{+}b$ 2: a, c, X, b 3: a, b, X, c 4: $c, X, \{a, b\}$

Note that b is the Condorcet loser, even if voter 3 swaps a and b. Hence, we derive the profile R^4 with $b \notin f(R^4)$ and by strategyproofness, $f(R^4) = \{a\}$.

R^4: 1: $\{a, c\}, X{+}b$ 2: a, c, X, b 3: b, a, X, c 4: $c, X, \{a, b\}$

Now, we let voter 4 change his preference to c, b, X, a to derive the profile R^5. As $f(R^4) = \{a\}$ and a is among the least preferred alternatives of voter 4, it follows that $f(R^5) \subseteq \{a, b\}$. Otherwise, voter 4 can manipulate by applying this modification.

R^5: 1: $\{a, c\}, X{+}b$ 2: a, c, X, b 3: b, a, X, c 4: c, b, X, a

We can apply the same steps for profiles symmetric with respect to the voters or alternatives. Thus, we can infer the choice sets for the profiles R^6, R^7, and R^8 as $f(R^6) \subseteq \{a, c\}$, $f(R^7) \subseteq \{a, b\}$, 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 in the sequel.

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 by replacing the preferences of some voters. In more detail, we obtain R^9 from R^5 by replacing the preference of voters 1 and 2 with a, b, X, c and the preference of voter 4 with b, a, X, c. If we apply these steps one after another, strategyproofness entails that the choice set is not allowed to change. Hence, $f(R^9) = \{a\}$. Moreover, we obtain R^9 from R^7 by replacing the preferences of voters 3 and 4 with b, a, X, c and the preference of voter 2 with a, b, X, c and obtain $f(R^9) = \{b\}$, a contradiction. Therefore, f cannot exist and there is no strategyproof SCF that satisfies non-imposition and the Condorcet loser property if $n \geq 4$ 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. By Lemma 2, g is non-imposing because f returns an alternative as unique winner if all voters prefer it uniquely the most and one voter is completely indifferent. 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. □

Remark 13. 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 rather technical SCF based on a case distinction on the maximal plurality score of an alternative satisfies all axioms if $n = 3$, $m \geq 3$.

Remark 14. Brandt [10, Th. 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 $n \geq 4$.

4 CONCLUSION AND DISCUSSION

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 a Condorcet loser in at least one profile. These results only leave 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.

Our results also have consequences for so-called *social decision schemes (SDSs)*, which map a preference profile to a *lottery* over alternatives. Since the notions of ranked-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. When extended to the support of lotteries, Kelly-strategyproofness is weaker than the well-studied notion of (weak) *SD*-strategyproofness [11]. Hence, Theorem 1 implies that there is no rank-based SDS that satisfies Pareto-optimality and *SD*-strategyproofness. Furthermore, Theorem 2 implies that every support-based SDS that satisfies Pareto-optimality and *SD*-strategyproofness puts positive probability on at least one most preferred alternative of every voter, a property that is known as *positive share* in the context of dichotomous preferences [7]. Finally, we can also define the Condorcet loser property for SDSs by requiring that Condorcet losers should always receive probability 0, and non-imposition by demanding that for every alternative, there is a profile such that this alternative receives probability 1. Then, Theorem 3 implies that there is no SDS that satisfies the Condorcet loser property, non-imposition, and *SD*-strategyproofness.

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 [17], Barberà et al. [5], Ching and Zhou [16], Rodríguez-Álvarez [31], and Sato [32]. 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 [17]. 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 [9, 15], 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 [18, 27] and a handful of other support-based Condorcet extensions satisfy strategyproofness if preferences are strict and participation for unrestricted preferences [9, 10]. 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.

REFERENCES

[1] T. Bandyopadhyay. 1983. Coalitional manipulation and the Pareto rule. *Journal of Economic Theory* 29, 2 (1983), 359–363.

[2] S. Barberà. 1977. The Manipulation of Social Choice Mechanisms That Do Not Leave "Too Much" to Chance. *Econometrica* 45, 7 (1977), 1573–1588.

[3] S. Barberà. 1977. Manipulation of Social Decision Functions. *Journal of Economic Theory* 15, 2 (1977), 266–278.

[4] S. Barberà. 2010. Strategy-proof social choice. In *Handbook of Social Choice and Welfare*, K. J. Arrow, A. K. Sen, and K. Suzumura (Eds.). Vol. 2. Elsevier, Chapter 25, 731–832.

[5] S. Barberà, B. Dutta, and A. Sen. 2001. Strategy-proof social choice correspondences. *Journal of Economic Theory* 101, 2 (2001), 374–394.

[6] J.-P. Benoît. 2002. Strategic Manipulation in Voting Games When Lotteries and Ties Are Permitted. *Journal of Economic Theory* 102, 2 (2002), 421–436.

[7] A. Bogomolnaia, H. Moulin, and R. Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.

[8] 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. Preliminary results appeared in the Proceedings of IJCAI-2016.

[9] F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. 2019. Strategic Abstention based on Preference Extensions: Positive Results and Computer-Generated Impossibilities. *Journal of Artificial Intelligence Research* 66 (2019), 1031–1056. Preliminary results appeared in the Proceedings of IJCAI-2015.

[10] F. Brandt. 2015. Set-Monotonicity Implies Kelly-Strategyproofness. *Social Choice and Welfare* 45, 4 (2015), 793–804. Preliminary results appeared in the Proceedings of IJCAI-2011.

[11] F. Brandt. 2017. Rolling the Dice: Recent Results in Probabilistic Social Choice. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access, Chapter 1, 3–26.

[12] F. Brandt, M. Brill, and P. 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, Chapter 3.

[13] F. Brandt, M. Bullinger, and P. Lederer. 2021. *On the Indecisiveness of Kelly-Strategyproof Social Choice Functions*. Technical Report. https://arxiv.org/abs/2102.00499.

[14] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

[15] F. Brandt, C. Saile, and C. Stricker. 2021. Strategyproof Social Choice When Preferences and Outcomes May Contain Ties. Working Paper. Preliminary results appeared in the Proceedings of AAMAS-2018.

[16] S. Ching and L. Zhou. 2002. Multi-valued strategy-proof social choice rules. *Social Choice and Welfare* 19, 3 (2002), 569–580.

[17] J. Duggan and T. Schwartz. 2000. Strategic Manipulability without Resoluteness or Shared Beliefs: Gibbard-Satterthwaite Generalized. *Social Choice and Welfare* 17, 1 (2000), 85–93.

[18] B. Dutta and J.-F. Laslier. 1999. Comparison Functions and Choice Correspondences. *Social Choice and Welfare* 16, 4 (1999), 513–532.

[19] U. Endriss (Ed.). 2017. *Trends in Computational Social Choice*. AI Access.

[20] A. Feldman. 1979. Manipulation and the Pareto Rule. *Journal of Economic Theory* 21 (1979), 473–482.

[21] P. C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489.

[22] P. Gärdenfors. 1976. Manipulation of Social Choice Functions. *Journal of Economic Theory* 13, 2 (1976), 217–228.

[23] P. Gärdenfors. 1979. On definitions of manipulation of social choice functions. In *Aggregation and Revelation of Preferences*, J. J. Laffont (Ed.). North-Holland.

[24] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.

[25] J. S. Kelly. 1977. Strategy-Proofness and Social Choice Functions Without Single-Valuedness. *Econometrica* 45, 2 (1977), 439–446.

[26] J.-F. Laslier. 1996. Rank-Based Choice Correspondences. *Economics Letters* 52, 3 (1996), 279–286.

[27] J.-F. Laslier. 2000. Interpretation of electoral mixed strategies. *Social Choice and Welfare* 17 (2000), 283–292.

[28] I. MacIntyre and P. K. Pattanaik. 1981. Strategic voting under minimally binary group decision functions. *Journal of Economic Theory* 25, 3 (1981), 338–352.

[29] A. Mas-Colell and H. Sonnenschein. 1972. General Possibility Theorems for Group Decisions. *Review of Economic Studies* 39, 2 (1972), 185–192.

[30] K. Nehring. 2000. Monotonicity implies generalized strategy-proofness for correspondences. *Social Choice and Welfare* 17, 2 (2000), 367–375.

[31] C. Rodríguez-Álvarez. 2007. On the manipulation of social choice correspondences. *Social Choice and Welfare* 29, 2 (2007), 175–199.

[32] S. Sato. 2008. On strategy-proof social choice correspondences. *Social Choice and Welfare* 31 (2008), 331–343.

[33] S. Sato. 2014. A fundamental structure of strategy-proof social choice correspondences with restricted preferences over alternatives. *Social Choice and Welfare* 42, 4 (2014), 831–851.

[34] 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.

[35] A. K. Sen. 1986. Social Choice Theory. In *Handbook of Mathematical Economics*, K. J. Arrow and M. D. Intriligator (Eds.). Vol. 3. Elsevier, Chapter 22, 1073–1181.

[36] A. D. Taylor. 2005. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press.