

Weak Strategyproofness in Randomized Social Choice

Felix Brandt¹, Patrick Lederer²

¹Technical University of Munich

²UNSW Sydney

brandtf@in.tum.de, p.lederer@unsw.edu.au

Abstract

An important—but very demanding—property in collective decision-making is strategyproofness, which requires that voters cannot benefit from submitting insincere preferences. Gibbard (1977) has shown that only rather unattractive rules are strategyproof, even when allowing for randomization. However, Gibbard’s theorem is based on a rather strong interpretation of strategyproofness, which deems a manipulation successful if it increases the voter’s expected utility for *at least one* utility function consistent with his ordinal preferences. In this paper, we study weak strategyproofness, which deems a manipulation successful if it increases the voter’s expected utility for *all* utility functions consistent with his ordinal preferences. We show how to systematically design attractive, weakly strategyproof social decision schemes (SDSs) and explore their limitations for both strict and weak preferences. In particular, for strict preferences, we show that there are weakly strategyproof SDSs that are either *ex post* efficient or Condorcet-consistent, while neither even-chance SDSs nor pairwise SDSs satisfy both properties and weak strategyproofness at the same time. By contrast, for the case of weak preferences, we discuss two sweeping impossibility results that preclude the existence of appealing weakly strategyproof SDSs.

1 Introduction

Any mechanism that relies on the private information of agents should incentivize agents to report their private information truthfully. However, designing mechanisms that satisfy this property—known as *strategyproofness*—is a challenging task in many domains of economic interest. This is particularly true for the field of social choice, which studies voting rules that aggregate the voters’ preferences into a collective decision: a seminal result by Gibbard (1973) and Satterthwaite (1975) shows that voters can benefit by lying about their true preferences in *any* reasonable deterministic voting rule. Early hopes that more positive results can be achieved for randomized voting rules were shattered by Gibbard (1977). Gibbard considered social decision schemes (SDSs), which return a probability distribution for each profile of individual preferences, and the final winner will be chosen by chance according to this distribution.

In particular, Gibbard (1977) has shown that the only strategyproof and *ex post* efficient SDSs are *random dictatorships*,

which choose each voter with a fixed probability and implement this voter’s favorite alternative as the winner of the election. While this result allows for more rules than the Gibbard-Satterthwaite theorem, random dictatorships suffer from a large variance and fail to identify good compromise alternatives. The latter observation is related to the fact that random dictatorships violate *Condorcet-consistency*, i.e., they may fail to select an alternative that beats all other alternatives in a pairwise majority comparison.

Like all results on strategyproof SDSs, Gibbard’s random dictatorship theorem crucially hinges on the exact definition of strategyproofness, which, in turn, depends on the assumptions of the voters’ preferences over lotteries. Gibbard (1977) postulates that a voter prefers one lottery to another if the former yields at least as much expected utility as the latter for every utility function that is consistent with his true preferences. Then, an SDS is called strategyproof if voters always prefer the outcome when voting truthfully to every outcome they could obtain by lying about their true preferences. This strategyproofness notion, which we will call *strong strategyproofness*, is predominant in the literature (e.g., Barberà 1979; Procaccia 2010; Brandt, Lederer, and Romen 2024) because it guarantees that voters cannot manipulate regardless of their exact utility functions. However, as demonstrated by the random dictatorship theorem, strong strategyproofness mainly leads to negative results.

In this paper, we will thus study a weaker notion of strategyproofness, first considered by Postlewaite and Schmeidler (1986) and later popularized by Bogomolnaia and Moulin (2001) in the context of random assignment. To this end, we observe that the preferences over lotteries defined by Gibbard (1977) are incomplete, i.e., there are lotteries such that a voter prefers neither of them to the other. Hence, there are two ways to define strategyproofness, depending on how we interpret incomparable lotteries. Strong strategyproofness, as defined by Gibbard (1977), views a deviation to an incomparable lottery as a successful manipulation. By contrast, we only consider deviations to comparable lotteries as successful. The resulting strategyproofness notion is called *weak strategyproofness* and requires that voters cannot obtain a strictly preferred lottery by lying about their true preferences.

Our contribution. We improve our understanding of weak strategyproofness by contributing various positive and nega-

tive results. A summary of these results is given in Table 1.

In our first theorem, we introduce the large class of *score-based* SDSs and show that all SDSs within this class are weakly strategyproof. This result allows, e.g., to construct appealing weakly strategyproof SDSs that satisfy Condorcet-consistency or *ex post* efficiency, or that approximate deterministic voting rules arbitrarily close. Both of these objectives are impossible to achieve for strong strategyproofness (Procaccia 2010; Brandt, Lederer, and Romen 2024). Secondly, we present characterizations of weakly strategyproof tops-only SDSs, which only have access to the voters' favorite alternatives. In this context, we also cover the design of weakly strategyproof *even-chance* SDSs, which always return a uniform lottery over some subset of alternatives. It has often been argued that such SDSs are more acceptable because uniform lotteries are easier to grasp cognitively and to implement in practice (e.g., Fishburn 1972; Gärdenfors 1979; Brandt, Saile, and Stricker 2022). However, no attractive even-chance SDS satisfies strong strategyproofness.

We also analyze the limitations of weakly strategyproof SDSs. In particular, for strict preferences, we show that no weakly strategyproof, even-chance SDS simultaneously satisfies *ex post* efficiency and Condorcet-consistency, and that no weakly strategyproof, pairwise, and neutral SDS (which can only access the pairwise majority comparisons between alternatives) satisfies *ex post* efficiency. These results indicate that Condorcet-consistency and *ex post* efficiency may be incompatible for *all* weakly strategyproof SDSs. However, such a result seems very difficult to obtain as, for every $\varepsilon > 0$, there are weakly strategyproof and Condorcet-consistent SDSs that assign at most ε probability to Pareto-dominated alternatives.

Finally, we also consider weakly strategyproof SDSs for weak preferences. In this setting, we provide the first easily verifiable proof of an important impossibility theorem by Brandl et al. (2018), showing that no weakly strategyproof SDS simultaneously satisfies anonymity, neutrality, and *ex ante* efficiency. When restricting attention to even-chance SDSs, we strengthen this result by proving that all weakly strategyproof and *ex post* efficient SDSs always randomize over the favorite alternatives of at most two fixed voters. These results show that, if we allow weak preferences, even mild forms of strategyproofness preclude attractive SDSs.

Related work. Studying weaker forms of strategyproofness is an active area in social choice theory (e.g., Bogomolnaia and Moulin 2001; Balbuzanov 2016; Aziz et al. 2018; Brandl et al. 2018; Lederer 2021; Mennle and Seuken 2021). Unfortunately, in the realm of voting, this approach mainly led to strengthened impossibility results: for instance, Lederer (2021) studies a strategyproofness notion that lies logically between weak and strong strategyproofness and shows that it is still incompatible with Condorcet-consistency. Despite these negative results, weak strategyproofness has received little attention in social choice theory. In particular, only few SDSs, such as the Condorcet rule (Postlewaite and Schmeidler 1986) or the egalitarian simultaneous reservation rule by Aziz and Stursberg (2014), are known to be weakly strategyproof. In more recent works, weak strategyproofness was used to prove impossibility theorems for the case of weak

preferences (Aziz et al. 2018; Brandl, Brandt, and Suksompong 2016; Brandl et al. 2018). This approach culminated in a sweeping impossibility theorem for weak preferences: no weakly strategyproof SDS satisfies anonymity, neutrality, and *ex ante* efficiency (Brandl et al. 2018).

Some of our results can also be compared to results for set-valued voting rules (which return sets of winning alternatives instead of lotteries). In particular, even-chance SDSs can be interpreted as set-valued voting rules and weak strategyproofness then translates to a strategyproofness notion called even-chance strategyproofness (e.g., Gärdenfors 1979; Brandt, Saile, and Stricker 2022). This strategyproofness notion is slightly stronger than commonly considered set-valued strategyproofness notions such as Fishburn-strategyproofness (Fishburn 1972) or Kelly-strategyproofness (Kelly 1977), and our paper is thus related to recent work on set-valued voting rules (e.g., Botan and Endriss 2021; Brandt, Saile, and Stricker 2022; Brandt and Lederer 2023).

2 Preliminaries

Let $N = \{1, \dots, n\}$ denote a set of n voters and let $A = \{a_1, \dots, a_m\}$ denote a set of $m \geq 2$ alternatives. Every voter $i \in N$ reports a (weak) *preference relation* \succsim_i , which is a complete and transitive binary relation on A . The strict part of \succsim_i is denoted by \succ_i (i.e., $x \succ_i y$ iff $x \succsim_i y$ and not $y \succsim_i x$) and the indifference part by \sim_i (i.e., $x \sim_i y$ iff $x \succsim_i y$ and $y \succsim_i x$). A preference relation \succsim_i is called *strict* if its irreflexive part coincides with \succ_i . We represent preference relations by comma-separated lists, where brackets indicate that a voter is indifferent between some alternatives. For instance, $a, \{b, c\}$ denotes that the considered voter prefers a to both b and c and is indifferent between the latter two alternatives. We denote the set of all strict preference relations by \mathcal{L} and the set of all weak preference relations by \mathcal{R} .

A (weak) *preference profile* $R = (\succsim_1, \dots, \succsim_n)$ is a vector that specifies the preference relations \succsim_i of all voters $i \in N$. A preference profile is *strict* if the preference relations of all voters are strict. The set of all strict preference profiles is given by \mathcal{L}^N , and the set of all weak preference profiles is \mathcal{R}^N . We represent preference profiles as collections of preference relations, where the set of voters that report a preference relation is stated directly before the preference relation. For instance, $\{1, 2, 3\} : a, b, c$ means that voters 1, 2, and 3 prefer a to b to c . To improve readability, we omit curly brackets for singleton sets.

The main objects of study in this paper are social decision schemes, which intuitively are voting rules that may use chance to determine the winner of an election. To formalize this, we define *lotteries* as probability distributions over the alternatives, i.e., a lottery p is a function of the type $A \rightarrow [0, 1]$ such that $\sum_{x \in A} p(x) = 1$. Moreover, by $\Delta(A)$ we denote the set of all lotteries over A . Then, a *social decision scheme* (SDS) on a domain $\mathcal{D} \in \{\mathcal{L}^N, \mathcal{R}^N\}$ is a function that maps every preference profile $R \in \mathcal{D}$ to a lottery $p \in \Delta(A)$. We denote by $f(R, x)$ the probability that the SDS f assigns to alternative x in the profile R and extend this notion to sets of alternatives X by defining $f(R, X) = \sum_{x \in X} f(R, x)$.

We will sometimes restrict our attention to even-chance SDSs. These SDSs pick a set of alternatives and random-

ize uniformly over these alternatives. Formally, an SDS f is *even-chance* if it chooses for every profile R a non-empty set of alternatives X such that $f(R, x) = \frac{1}{|X|}$ if $x \in X$ and $f(R, x) = 0$ otherwise. Even-chance lotteries are appealing because of their simplicity and because non-uniform randomization may be difficult to implement in the real world. On top of that, even-chance SDSs can naturally be interpreted as set-valued voting rules, which have been studied before.

2.1 Strategyproofness

The central axiom for our analysis is strategyproofness which demands that voters cannot benefit by lying about their true preferences. To define this axiom for SDSs, we need to specify how voters compare lotteries over the alternatives. We assume for this that voters have utility functions $u_i : A \rightarrow \mathbb{R}$ and aim to maximize their expected utility. However, the exact utility functions are not known as voters only reveal their ordinal preferences over alternatives. We will thus quantify over all utility functions u_i that are consistent with the voter's preference relation \succsim_i , i.e., that satisfy that $u_i(x) \geq u_i(y)$ if and only if $x \succsim_i y$ for all $x, y \in A$. A voter then prefers lottery p to lottery q , denoted by $p \succsim_i q$, if p guarantees him at least as much expected utility as q for *every* utility function u_i that is consistent with his preference relation, i.e., if $\mathbb{E}[u_i(p)] \geq \mathbb{E}[u_i(q)]$ for all consistent utility functions u_i . Alternatively, this lottery extension can also be defined based on stochastic dominance. To state this definition, we let the *upper contour set* $U(\succsim_i, x) = \{y \in A : y \succsim_i x\}$ of x denote the set of alternatives that voter i weakly prefers to x . It then holds that $p \succsim_i q$ if and only if $p(U(\succsim_i, x)) \geq q(U(\succsim_i, x))$ for all $x \in A$ (see, e.g., Sen 2011; Brandl et al. 2018).

Importantly, the voters' preferences over lotteries, as defined above, are incomplete, i.e., there are lotteries p and q and a preference relation \succsim_i such that neither $p \succsim_i q$ nor $q \succsim_i p$. For example, for the preference relation a, b, c , the lotteries p and q defined by $p(a) = p(b) = p(c) = 1/3$ and $q(b) = 1$ are incomparable as neither of them stochastically dominates the other. Consequently, there are two ways to define strategyproofness depending on how we handle incomparable lotteries. The approach suggested by Gibbard (1977) counts a deviation to an incomparable lottery as a successful manipulation and strategyproofness hence prohibits such deviations. This results in *strong strategyproofness*, which requires of an SDS f that $f(R) \succsim_i f(R')$ for all profiles R, R' and voters $i \in N$ such that $\succsim_j = \succsim_j'$ for all $j \in N \setminus \{i\}$. By contrast, we will not count the deviation to an incomparable lottery as a successful manipulation. This leads to a weaker form of strategyproofness: an SDS f is *weakly strategyproof* if $f(R') \not\succeq_i f(R)$ for all profiles R, R' and voters $i \in N$ such that $\succsim_j = \succsim_j'$ for all $j \in N \setminus \{i\}$.

2.2 Further Axioms

We conclude this section by stating four standard axioms.

Anonymity. Anonymity is a basic fairness property that states that the identities of the voters should not matter. Formally, an SDS f is *anonymous* if $f(R) = f(\pi(R))$ for all permutations $\pi : N \rightarrow N$ and profiles R , where the profile $R' = \pi(R)$ is given by $\succsim'_{\pi(i)} = \succsim_i$ for all voters $i \in N$.

Neutrality. Similar to anonymity, neutrality is a fairness property that requires that alternatives are treated equally. In more detail, an SDS f is *neutral* if $f(\tau(R), \tau(x)) = f(R, x)$ for all permutations $\tau : A \rightarrow A$ and profiles R . This time, the profile $R' = \tau(R)$ is defined by $\tau(x) \succsim'_i \tau(y)$ if and only if $x \succsim_i y$ for all $x, y \in A$ and $i \in N$.

Ex post efficiency. *Ex post* efficiency postulates that alternatives should have no chance of being selected if there is another alternative that makes at least one voter better off without making any other voter worse off. To this end, we say 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 $i \in N$. Conversely, an alternative is *Pareto-optimal* in a profile R if it is not Pareto-dominated by any other alternative. Finally, an SDS f is *ex post efficient* if $f(R, x) = 0$ for all profiles R and Pareto-dominated alternative x .

Condorcet-consistency. Condorcet-consistency demands that a Condorcet winner, an alternative that beats every other alternative in a pairwise majority comparison, should be selected with probability 1 whenever it exists. To formalize this, we define the *majority relation* \succsim_M of a profile R by $x \succsim_M y$ if and only if $|\{i \in N : x \succ_i y\}| \geq |\{i \in N : y \succ_i x\}|$ for all $x, y \in A$. Moreover, \succ_M denotes the strict part of \succsim_M and \sim_M its indifferent part. Then, an alternative x is a *Condorcet winner* in a profile R if $x \succ_M y$ for all $y \in A \setminus \{x\}$, and an SDS f is *Condorcet-consistent* if $f(R, x) = 1$ whenever x is the Condorcet winner in R .

3 Results

We are now ready to state our results. We first present theorems that allow the design attractive weakly strategyproof SDSs (Section 3.1), and then discuss the limitations of weakly strategyproof SDSs (Sections 3.2 and 3.3). Due to space constraints, we defer most of the proofs to an extended version of this paper (Brandt and Lederer 2024).

3.1 Possibility Theorems for Strict Preferences

In this section, we will show how to design weakly strategyproof SDSs when voters have strict preferences. In more detail, we will first present a large class of weakly strategyproof SDSs (cf. Theorem 1) and then give two characterizations of weakly strategyproof SDSs that only depend on the voters' favorite alternatives (cf. Theorem 2).

We first introduce the class of score-based SDSs and show that all these SDSs are weakly strategyproof. For this, let $R^{i:yx}$ be the profile derived from another profile R by only changing voter i 's preferences from $x \succ_i y$ to $y \succ_i^{i:yx} x$. In particular, this requires that there is no $z \in A$ with $x \succ_i z \succ_i y$. Next, a function $s : \mathcal{L}^N \times A \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is a *score function* if it satisfies for all profiles $R \in \mathcal{L}^N$, distinct alternatives $x, y, z \in A$, and voters $i \in N$ that

- $s(R, x) = \infty$ implies $s(R, y) \neq \infty$ (at most one infinity),
- $s(R, z) = s(R^{i:yx}, z)$ (localizedness),
- $s(R, y) \leq s(R^{i:yx}, y)$ (monotonicity), and
- $s(R, y) = s(R^{i:yx}, y)$ implies $s(R, x) = s(R^{i:yx}, x)$ unless $s(R, y) = \infty$, or $s(R, x) = \infty$ and $s(R, y) > 0$ (balancedness).

We note that a score function can assign a score of infinity to at most one alternative. We thus assume the usual arithmetic rules for infinity: for all $x \in \mathbb{R}$, it holds that $\infty > x$, $\infty + x = \infty$, $\frac{x}{\infty} = 0$, and $\frac{\infty}{\infty} = 1$. Finally, an SDS f on \mathcal{L}^N is *score-based* if there is a score function s such that $\sum_{y \in A} s(R, y) > 0$ and $f(R, x) = \frac{s(R, x)}{\sum_{y \in A} s(R, y)}$ for all alternatives $x \in A$ and profiles $R \in \mathcal{L}^N$.

We will next discuss several examples of score-based SDSs to illustrate this class and its versatility. To this end, we first consider two classical score functions, namely the Copeland score function $s_C(R, x) = |\{y \in A \setminus \{x\} : x \succ_M y\}| + \frac{1}{2}|\{y \in A \setminus \{x\} : x \sim_M y\}|$ and the plurality score function $s_P(R, x) = |\{i \in N : \forall y \in A \setminus \{x\} : x \succ_i y\}|$. Both of these functions are indeed score functions according to our definition and the corresponding SDSs are thus score-based. Moreover, for every strictly monotonically increasing function $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, it holds that $s_C^g(R, x) = g(s_C(R, x))$ and $s_P^g(R, x) = g(s_P(R, x))$ are also score functions. For example, this means that the SDSs defined by the functions s_C^k and s_P^k , which take the k -th power of $s_C(R, x)$ and $s_P(R, x)$, are score-based. Even SDSs that seem rather unrelated to score functions belong to our class. For instance, the Condorcet rule (which chooses the Condorcet winner with probability 1 whenever it exists and randomizes uniformly over all alternatives otherwise) is the score-based SDS defined by the score function s with $s(R, x) = \infty$ if x is the Condorcet winner in R and $s(R, x) = 1$ otherwise. Similarly, the function $s_C^{k, CW}$, which assigns a score of infinity to the Condorcet winner and otherwise coincides with s_C^k , satisfies all our conditions and is thus a score function.

We will now prove that all score-based SDSs are weakly strategyproof.

Theorem 1. *Every score-based SDS on \mathcal{L}^N satisfies weak strategyproofness.*

Proof sketch. Let f be a score-based SDS and let s be its score function. Moreover, we consider two profiles $R, R' \in \mathcal{L}^N$ and a voter $i \in N$ such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$ and $f(R) \neq f(R')$. To simplify this proof sketch, we additionally assume that $0 < s(R, x) < \infty$ and $s(R', x) < \infty$ for all $x \in A$, and that $\succsim_i = x_1, x_2, \dots, x_m$. Next, we define $s_{total}(R) = \sum_{x \in A} s(R, x)$ and consider three cases. First, if $s_{total}(R) < s_{total}(R')$, we use the monotonicity and localizedness of s to show that $s(R, x_1) \geq s(R', x_1)$ by transforming R to R' with a swap sequence that never reinforces x_1 . Since $s_{total}(R) < s_{total}(R')$, it follows that $f(R, x_1) = \frac{s(R, x_1)}{s_{total}(R)} > \frac{s(R', x_1)}{s_{total}(R')} = f(R', x_1)$ and thus $f(R') \not\succeq_i f(R)$. If $s_{total}(R) > s_{total}(R')$, we can use a similar argument by showing that the score of voter i 's least preferred alternative x_m weakly increases when going from R to R' . Finally, if $s_{total}(R) = s_{total}(R')$, we let x_h denote the alternative such that $s(R, x_h) = s(R', x_h)$ for all $\ell < h$ and $s(R, x_h) \neq s(R', x_h)$. Then, we prove that $s(R, x_h) > s(R', x_h)$, which shows that $f(R, U(\succsim_i, x_h)) > f(R', U(\succsim_i, x_h))$ and thus $f(R') \not\succeq_i f(R)$. Finally, slightly more involved arguments extend this analysis to the case that $s_{total}(R) = \infty$ or $s_{total}(R') = \infty$. \square

Theorem 1 has a number of important consequences. Firstly, this result implies that the score-based SDSs defined by s_P^k and s_C^k are weakly strategyproof. Since all these rules fail strong strategyproofness when $k \neq 1$, this demonstrates that the space of weakly strategyproof SDSs is significantly richer than the one of strongly strategyproof SDSs. Secondly, we note that the score-based SDSs defined by s_P^k and s_C^k approximate the Plurality rule and Copeland rule (which simply choose the alternatives with maximal Plurality and Copeland score, respectively) arbitrarily closely by increasing the exponent k . This stands in sharp contrast to a result by Procaccia (2010) who has shown that strongly strategyproof SDSs are poor approximations of common deterministic voting rules. Thirdly, Theorem 1 implies that there are interesting weakly strategyproof SDSs that are *ex post* efficient or Condorcet-consistent. For instance, all score-based SDSs defined by a score-function $s_P^g(R, x) = g(s_P(R, x))$ are *ex post* efficient when g satisfies that $g(0) = 0$ and $g(x) > g(y)$ for all $x, y \in \mathbb{N}_0$ with $x > y$, and the Condorcet rule as well as the score-based SDS defined by $s_C^{k, CW}$ are Condorcet-consistent. This stands again in contrast to results for strong strategyproofness because Brandt, Lederer, and Romen (2024) have shown that all strongly strategyproof SDSs can put at most probability $2/m$ on Condorcet winners, and that all strongly strategyproof SDSs that assign at most a probability of less than $1/m$ to Pareto-dominated alternatives have a random dictatorship component.

A natural follow-up question for Theorem 1 is whether the class of score-based SDSs is equivalent to the set of weakly strategyproof SDSs. This is not the case since the omninomination rule f^O , which randomizes uniformly over the set of top-ranked alternatives $OMNI(R) = \{x \in A : s_P(R, x) > 0\}$, is weakly strategyproof but not score-based. In particular, every score function that induces this SDS fails balancedness or localizedness, so it cannot be score-based.

To give some characterizations for weakly strategyproof SDSs, we will next focus on the class of tops-only SDSs, which only depend on the voters' favorite alternatives. More formally, let $T_i(R) = \{x \in A : x \succsim_i y \text{ for all } y \in A\}$ denote the set of voter i 's favorite alternatives and note that $|T_i(R)| = 1$ if R is strict. Then, an SDS f is *tops-only* if $f(R) = f(R')$ for all preference profiles R and R' such that $T_i(R) = T_i(R')$ for all $i \in N$. We will now provide two characterizations of weak strategyproofness for tops-only SDSs on \mathcal{L}^N : firstly, we will show that, for tops-only SDSs, weak strategyproofness is equivalent to a monotonicity property. Furthermore, we will characterize the class of weakly strategyproof SDSs that are tops-only, even-chance, anonymous, and neutral as parameterized omninomination rules. These SDSs are defined by two parameters $\theta_1 > \frac{n}{2}$ and θ_2 and coincide with f^O except for two special cases: (i) if a single alternative is top-ranked by at least θ_1 voters, then this alternative is assigned probability 1, and (ii) if no such alternative exists and more than θ_2 alternatives are top-ranked in total, then we randomize uniformly over all alternatives. More formally, an SDS f is a *parameterized omninomination rule* if there are two parameters $\theta_1 \in \{\lceil \frac{n+1}{2} \rceil, \dots, n+1\}$ and $\theta_2 \in \{0, \dots, m-1\}$ such that

- $f(R, x) = 1$ for all profiles R and alternatives $x \in A$ such that $s_P(R, x) \geq \theta_1$,
- $f(R) = f^O(R)$ for all profiles R such that $\max_{x \in A} s_P(R, x) < \theta_1$ and $|OMNI(R)| \leq \theta_2$,
- $f(R, x) = \frac{1}{m}$ for all profiles R and alternatives $x \in A$ such that $\max_{x \in A} s_P(R, x) < \theta_1$ and $|OMNI(R)| > \theta_2$.

Theorem 2. Let f denote a tops-only SDS on \mathcal{L}^N .

- 1) f is weakly strategyproof if and only if $f(R) = f(R')$ or $f(R, T_i(R)) > f(R', T_i(R))$ for all profiles $R, R' \in \mathcal{L}^N$ and voters $i \in N$ such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$.
- 2) f is weakly strategyproof, even-chance, anonymous, and neutral if and only if it is a parameterized omninomination rule.

Proof. We will only prove the first claim here and defer the proof of the second part of the theorem to the full version. Thus, let f denote a tops-only SDS. We first show the direction from right to left and hence suppose that f satisfies the given condition. Now, let R and R' denote two preference profiles and i a voter such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$. Moreover, let x denote voter i 's favorite alternative in R and y his favorite alternative in R' . If $f(R) = f(R')$, voter i clearly cannot manipulate by deviating from R to R' . Hence, we suppose that $f(R) \neq f(R')$, which requires that $x \neq y$ due to tops-onlyness. In turn, the condition of our theorem implies that $f(R', x) < f(R, x)$ if $f(R) \neq f(R')$ and $x \neq y$. This implies that $f(R') \neq f(R)$, so f is weakly strategyproof.

Next, we will show that f fails weak strategyproofness if it fails the condition in the theorem. To this end, assume there are profiles R and R' , a voter i , and an alternative x such that $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$, $T_i(R) = \{x\}$, $f(R) \neq f(R')$, and $f(R, x) \leq f(R', x)$. We define $Z^+ = \{z \in A \setminus \{x\} : f(R', z) \geq f(R, z)\}$ and $Z^- = \{z \in A \setminus \{x\} : f(R', z) < f(R, z)\}$ and observe that $Z^- \neq \emptyset$ since $f(R) \neq f(R')$. Moreover, we consider the profile R^* derived from R by assigning voter i a strict preference relation \succsim_i^* with $x \succ_i^* z^+ \succ_i^* z^-$ for all $z^+ \in Z^+$ and $z^- \in Z^-$. By tops-onlyness, $f(R) = f(R^*)$. On the other hand, it holds by construction that $f(R') \neq f(R^*)$ and $f(R', U(\succsim_i^*, y)) \geq f(R^*, U(\succsim_i^*, y))$ for all $y \in A$, so $f(R') \succ_i^* f(R^*)$ and f fails weak strategyproofness. \square

Remark 1. The second claim of Theorem 2 can be used to characterize the SDS that assigns probability 1 to an alternative if it is top-ranked by a strict majority of voters and randomizes uniformly over $OMNI(R)$ if no such alternative exists: among all SDSs that are tops-only, even-chance, weakly strategyproof, anonymous, and neutral, it is the one that uses the least amount of randomization.

Remark 2. Every strongly strategyproof SDS is score-based. This follows from a result by Gibbard (1977), which states that an SDS is strongly strategyproof if and only if the SDS itself is localized and monotonic. In our terminology, this means that an SDS is strongly strategyproof if and only if it is defined by score function $s : \mathcal{L}^N \times A \rightarrow \mathbb{R}_{\geq 0}$ that satisfies monotonicity, localizedness, and that $s(R, x) - s(R^{i:yx}, x) = s(R^{i:yx}, y) - s(R, y)$ for all

$x, y \in A$ and $R \in \mathcal{L}^N$. By replacing the last constraint by balancedness (and even allowing at most one alternative with $s(R, x) = \infty$), we derive significantly more flexibility in the design of weakly strategyproof SDSs.

3.2 Impossibility Theorems for Strict Preferences

We will now turn to the limitations of weakly strategyproof SDSs for the case that voters report strict preferences. To this end, we observe that, while our results in Section 3.1 allow to construct weakly strategyproof SDSs that are arbitrarily close to simultaneously satisfying *ex post* efficiency and Condorcet-consistency (e.g., the score-based rule defined by s_C^k for very large k), we were not able to construct a weakly strategyproof SDS that satisfies both axioms at the same time. As it turns out, constructing such an SDS may be impossible. In more detail, we subsequently prove that no even-chance SDS (cf. Theorem 3) and no neutral and pairwise SDSs (cf. Theorem 4) simultaneously satisfies weak strategyproofness, Condorcet-consistency, and *ex post* efficiency. This shows that the most common approaches for designing Condorcet-consistent SDSs do not allow to simultaneously satisfy weak strategyproofness, *ex post* efficiency, and Condorcet-consistency.

Let us first consider even-chance SDSs.

Theorem 3. Assume that $m \geq 5$ and $n \geq 5$ is odd. No even-chance SDS on \mathcal{L}^N satisfies weak strategyproofness, Condorcet-consistency, and *ex post* efficiency.

Proof Sketch. For the proof of this result, we first show two auxiliary claims: assuming that the number of voters n is odd, we prove (i) that every weakly strategyproof and Condorcet-consistent even-chance SDS assigns probability 1 to an alternative x if and only if x is the Condorcet winner, and (ii) that such SDSs can never randomize over exactly two alternatives. We then consider the following two preference profiles; additional alternatives are bottom-ranked by all voters.

$$\begin{aligned}
 R^1: & \quad 1: b, e, d, c, a & 2: a, b, c, e, d & 3: e, d, c, a, b \\
 & \quad \{4, \dots, \frac{n+3}{2}\}: b, c, a, e, d & \{\frac{n+5}{2}, \dots, n\}: e, d, a, b, c \\
 \hat{R}^1: & \quad 1: b, e, d, c, a & 2: a, b, c, e, d & 3: d, a, e, b, c \\
 & \quad \{4, \dots, \frac{n+3}{2}\}: b, c, a, e, d & \{\frac{n+5}{2}, \dots, n\}: e, d, a, b, c
 \end{aligned}$$

For these profiles, we show based on our auxiliary claims that every even-chance SDS that is weakly strategyproof, Condorcet-consistent, and *ex post* efficient uniformly randomizes over $\{a, b, c, e\}$ for R^1 and over $\{a, b, d, e\}$ for \hat{R}^1 . This means that voter 3 can manipulate by deviating from R^1 to \hat{R}^1 , contradicting weak strategyproofness. \square

Let us now turn to pairwise and neutral SDSs. An SDS is *pairwise* if $f(R) = f(R')$ for all profiles R, R' such that $|\{i \in N : x \succ_i y\}| - |\{i \in N : y \succ_i x\}| = |\{i \in N : x \succ'_i y\}| - |\{i \in N : y \succ'_i x\}|$ for all $x, y \in A$. In other words, an SDS is pairwise if it only depends on the weighted majority relation. There are many important pairwise SDSs (see, e.g., Brandt et al. 2016, Chapter 3 and 4) and most Condorcet-consistent voting rules are pairwise. Hence, showing that all pairwise, neutral, and weakly strategyproof SDSs fail *ex post* efficiency can be seen as evidence

that no SDS simultaneously satisfies weak strategyproofness, Condorcet-consistency, and *ex post* efficiency. We need to slightly extend the domain of SDSs for the following result: $\mathcal{L}^* = \bigcup_{N \subseteq \mathbb{N} \text{ is finite and non-empty}} \mathcal{L}^N$ contains all strict preference profiles for every finite and non-empty electorate. This is required because pairwise relationships establishes relationships between profiles with different numbers of voters.

Theorem 4. *Assume that $m \geq 5$. No pairwise, neutral, and weakly strategyproof SDS on \mathcal{L}^* satisfies *ex post* efficiency.*

Proof Sketch. For this proof sketch, we focus on the case that there are $m = 5$ alternatives. Moreover, we will only consider profiles with $n = 3$ voters; this suffices as we work in \mathcal{L}^* . Nevertheless, the impossibility theorem can be extended to all odd values $n > 3$ by adding voters with inverse preferences as pairwise relationships requires that the outcome does not change when adding such voters. Now, assume for contradiction that there is an SDS f that satisfies all axioms of our theorem. First, we show that f is invariant under weakening alternatives that obtain probability 0, i.e., if R' arises from a profile R by only weakening an alternative x with $f(R, x) = 0$, then $f(R) = f(R')$. Next, we focus on the profiles R and R' .

$$\begin{array}{lll} R: & 1: a, b, e, c, d & 2: b, c, e, a, d & 3: e, c, a, b, d \\ R': & 1: a, b, e, c, d & 2: b, c, a, e, d & 3: e, c, a, b, d \end{array}$$

In particular, we show that $f(R, a) = f(R, b) = f(R, e) = \frac{1}{3}$ and $f(R', a) = f(R', b) = f(R', c) = \frac{1}{3}$ by analyzing several auxiliary profiles. However, this means that voter 2 can manipulate by deviating from R to R' as $f(R') \succ_2 f(R)$, so f fails weak strategyproofness. \square

Remark 3. All axioms except of the even-chance condition are required for Theorem 3: the omninomination rule f^O satisfies *ex post* efficiency and weak strategyproofness, the Condorcet rule satisfies Condorcet-consistency and weak strategyproofness, and uniformly randomizing over known majoritarian choice sets such as the uncovered set satisfies *ex post* efficiency and Condorcet-consistency. Moreover, the bounds on n and m are tight: if $n \leq 4$, the SDS that chooses the Condorcet winner with probability 1 if there is one and otherwise randomizes uniformly over the top-ranked alternatives satisfies all given axioms; if $m \leq 4$, the SDS that chooses the Condorcet winner with probability 1 if there is one, randomizes uniformly over the top-ranked alternative if there are two that are first-ranked by exactly half of the voters, and otherwise randomizes uniformly over the set of *ex post* efficient alternatives meets all conditions of Theorem 3. Finally, we conjecture that the even-chance assumption is not required for the impossibility.

Remark 4. Just like numerous other results on Condorcet-consistency (e.g., Botan and Endriss 2021; Brandt, Lederer, and Suksompong 2023), we cannot prove Theorem 3 for even n . However, we show in the full version that there is no even-chance SDS that satisfies weak strategyproofness, *ex post* efficiency, and strong Condorcet-consistency. The last condition requires that an alternative is chosen with probability 1 if and only if it is the Condorcet winner, and it is typically satisfied by all strategyproof, Condorcet-consistent SDSs.

Remark 5. The main result of Brandt and Lederer (2023) implies that, under mild additional assumptions, no pairwise set-valued voting rule satisfies Condorcet-consistency, *ex post* efficiency, and a strategyproofness notion called Fishburn-strategyproofness. When interpreting even-chance SDSs as set-valued voting rules, Theorem 3 extends this observation to *all* set-valued voting rules at the expense of using a slightly stronger notion of strategyproofness.

3.3 Impossibility Theorems for Weak Preferences

In this section, we prove two impossibility theorems for weakly strategyproof SDSs on \mathcal{R}^N : firstly, we present a simplified proof of the main result of Brandl et al. (2018) (cf. Theorem 5) in the full version; secondly, we prove an even more severe impossibility for even-chance SDSs (cf. Theorem 6). These results show that there are no attractive weakly strategyproof SDSs when voters have weak preferences.

We start by revisiting the impossibility theorem by Brandl et al. (2018). It shows that no anonymous, neutral, and weakly strategyproof SDS is *ex ante* efficient. The last condition, also known as SD-efficiency, is a strengthening of *ex post* efficiency focusing on lotteries rather than individual alternatives. In more detail, a lottery p *ex ante* dominates a lottery q in a profile R if $p \succsim_i q$ for all $i \in N$ and $p \succ_i q$ for some $i \in N$. Conversely, a lottery is *ex ante* efficient if it is not *ex ante* dominated by any lottery, and an SDS f is *ex ante* efficient if $f(R)$ is *ex ante* efficient for every profile R . *Ex ante* efficiency ensures that there is no lottery that weakly increases the expected utility of all voters and strictly for at least one voter. Hence, the impossibility theorem by Brandl et al. (2018) shows that no weakly strategyproof SDS on \mathcal{R}^N satisfies mild efficiency constraints.

Theorem 5 (Brandl et al. (2018)). *Assume $n \geq 4$ and $m \geq 4$. No anonymous and neutral SDS on \mathcal{R}^N satisfies *ex ante* efficiency and weak strategyproofness.*

Brandl et al. (2018) have shown this result by a computer-generated proof, which reasons over 47 (canonical) preference profiles. As it is very difficult for humans to verify the correctness of this 14-page proof, Brandl et al. had the proof checked by the interactive theorem prover Isabelle/HOL. By contrast, we give a rather simple proof of this result in the full version (Brandt and Lederer 2024), which argues over only 13 profiles (10 canonical profiles) and takes less than two pages. Its correctness can easily be verified by the avid reader.

Theorem 5 crucially relies on *ex ante* efficiency. Indeed, if *ex ante* efficiency is replaced with *ex post* efficiency, random serial dictatorship satisfies all the axioms (Aziz et al. 2018). This SDS randomly chooses an order over the voters and each voter in the sequence then acts as dictator, breaking the ties left by the previous dictators. This leads to the question whether there are also reasonable *even-chance* SDSs on \mathcal{R}^N that satisfy weak strategyproofness and *ex post* efficiency. Unfortunately, it turns out that this is not the case: every such SDS can only randomize over the top-ranked alternatives of at most two voters. To make this formal, we say an SDS f is *dictatorial* if there is a voter $i \in N$ such that $f(R, T_i(R)) = 1$ for all profiles R , and *bidictatorial* if there are two distinct voters $i, j \in N$ such that $f(R, T_i(R) \cup T_j(R)) = 1$ for all

	Strict preferences	Weak preferences
<i>ex post</i> efficiency	\oplus Various tops-only SDSs (Thm 1,2) \ominus No pairwise SDS* (Thm 4)	\oplus Random serial dictatorship (Aziz et al. 2018) \ominus No <i>ex ante</i> efficient SDSs* (Brandl et al. 2018) \ominus Only (bi)dictatorial even-chance SDSs (Thm 6)
Condorcet-consistency	\oplus Variants of Copeland’s rule (Thm 1)	\ominus No SDS (Brandt 2015)
<i>ex post</i> efficiency and Condorcet-consistency	\oplus Approximately <i>ex post</i> efficient and Condorcet-consistent SDSs (Thm 1) \ominus No even-chance SDS (Thm 3)	\ominus No SDS (Brandt 2015)

Table 1: Summary of our results. Each cell states which weakly strategyproof SDSs satisfy which axioms for strict preferences (left column) and weak preferences (right column). Results marked by a \oplus symbol are possibility theorems, whereas the \ominus symbol indicates impossibility theorems. Results with an asterisk (*) additionally need anonymity and/or neutrality.

profiles R . Clearly, dictatorial and bidictatorial SDSs are undesirable as at most two voters can influence which alternatives are returned with positive probability.

Theorem 6. *Assume $m \geq 3$ and $n \geq 3$. Every *ex post* efficient and weakly strategyproof even-chance SDS on \mathcal{R}^N is dictatorial or bidictatorial.*

Proof Sketch. Let f denote an even-chance SDS that is *ex post* efficient and weakly strategyproof. The proof of this theorem is focusing on the decisive groups and weak dictators of f . To this end, we say that a voter i is a weak dictator for f if $f(R, T_i(R)) > 0$ for all profiles R and a group of voters $G \subseteq N$ is decisive for f if $f(R, T_i(R)) = 1$ for all voters $i \in G$ and profiles R such that all voters in G report the same preference relation in R . First, a result of Brandt, Bullinger, and Lederer (2022) implies that if a voter i is not a weak dictator for f , then $N \setminus \{i\}$ is decisive for f . We next show that there are at least one and at most two weak dictators i, j for f , so we get for every voter $h \notin \{i, j\}$ that $N \setminus \{h\}$ is decisive. The last insight for our proof is a contraction lemma stating that if there are two decisive groups G and G' for f such that $|G| = |G'|$ and $|G \cap G'| = |G| - 1$, then $G \cap G'$ is also decisive for f . By applying this to our decisive groups, we infer that the set of weak dictators is decisive. Based on this observation, we finally show that f is (bi)dictatorial. \square

Remark 6. All axioms are required for Theorem 6. Every constant even-chance SDS only fails *ex post* efficiency, the SDS that randomizes uniformly over the Pareto-optimal alternatives is *ex post* efficient and not (bi)dictatorial. Finally, random serial dictatorship satisfies all axioms but it is not even-chance nor (bi)dictatorial.

Remark 7. Theorem 6 has interesting connections to known results. Firstly, based on much stronger strategyproofness notions, Feldman (1980) and Barberà, Dutta, and Sen (2001) show that all strategyproof and *ex post* efficient even-chance SDSs are dictatorial or bidictatorial when the voters’ preferences are strict. For instance, Feldman (1980) uses strong strategyproofness and his result is thus a corollary of the theorem by Gibbard (1977). Theorem 6 demonstrates that a much weaker strategyproofness notion still allows to

deduce the same result when allowing for weak preferences. Secondly, Theorem 6 is related to a result by Brandt, Saile, and Stricker (2022) who show that no set-valued voting rule on \mathcal{R}^N satisfies anonymity, *ex post* efficiency, and a strategyproofness notion due to Fishburn (1972). When interpreting even-chance SDSs as set-valued voting rules, weak strategyproofness is only slightly stronger than Fishburn-strategyproofness, but it yields the much more restrictive conclusion of (bi)dictatorial rules. Moreover, Theorem 6 is stronger than Corollary 2 of Brandt, Saile, and Stricker.

4 Conclusion

In this paper, we study randomized voting rules, so-called social decision schemes (SDSs), with respect to weak strategyproofness. This strategyproofness notion only deems a manipulation successful if it increases the voter’s expected utility for all utility functions consistent with his ordinal preferences. We show that weak strategyproofness allows for some positive results. For example, in contrast to results on strong strategyproofness (Brandt, Lederer, and Romen 2024), there are Condorcet-consistent weakly strategyproof SDSs that are approximately *ex post* efficient. We also explore the limitations of weak strategyproofness and show, e.g., that no even-chance SDS simultaneously satisfies weak strategyproofness, Condorcet-consistency, and *ex post* efficiency when preferences are strict. Moreover, we prove much more severe impossibility theorems for weak preferences, highlighting a sharp contrast between strict and weak preferences. We refer to Table 1 for a complete overview of our results.

Our work points to a number of interesting and challenging open questions: firstly, based on our results in Section 3.2, we conjecture that no weakly strategyproof SDS satisfies both Condorcet-consistency and *ex post* efficiency. If this conjecture was true, this result would effectively unify several results analyzing the existence of strategyproof and Condorcet-consistent SDSs (e.g., Lederer 2021; Brandt, Lederer, and Suksompong 2023; Brandt and Lederer 2023). Secondly, our negative results for the case of weak preferences lead to the question of whether all weakly strategyproof and *ex post* efficient SDSs on the domain of weak preferences only randomize over the top-ranked alternatives of the voters.

Acknowledgements

We thank Alexander Thole and René Romen for helpful discussions, and four anonymous reviewers for their feedback. This work was supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-2 and BR 2312/12-1, and by the NSF-CSIRO grant on "Fair Sequential Collective Decision-Making" (RG230833).

References

- Aziz, H.; Brandl, F.; Brandt, F.; and Brill, M. 2018. On the Tradeoff between Efficiency and Strategyproofness. *Games and Economic Behavior*, 110: 1–18. A preliminary version of this paper appeared in the Proceedings of AAMAS-2013.
- Aziz, H.; and Stursberg, P. 2014. A Generalization of Probabilistic Serial to Randomized Social Choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 559–565.
- Balbuzanov, I. 2016. Convex strategyproofness with an application to the probabilistic serial mechanism. *Social Choice and Welfare*, 46(3): 511–520.
- Barberà, S. 1979. A Note on Group Strategy-Proof Decision Schemes. *Econometrica*, 47(3): 637–640.
- Barberà, S.; Dutta, B.; and Sen, A. 2001. Strategy-proof social choice correspondences. *Journal of Economic Theory*, 101(2): 374–394.
- Bogomolnaia, A.; and Moulin, H. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100(2): 295–328.
- Botan, S.; and Endriss, U. 2021. Preserving Condorcet Winners under Strategic Manipulation. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5202–5210.
- 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): 1–28. A preliminary version of this paper appeared in the Proceedings of IJCAI-2016.
- Brandl, F.; Brandt, F.; and Suksompong, W. 2016. The Impossibility of Extending Random Dictatorship to Weak Preferences. *Economics Letters*, 141: 44–47.
- Brandt, F. 2015. Set-Monotonicity Implies Kelly-Strategyproofness. *Social Choice and Welfare*, 45(4): 793–804.
- Brandt, F.; Bullinger, M.; and Lederer, P. 2022. On the Indecisiveness of Kelly-Strategyproof Social Choice Functions. *Journal of Artificial Intelligence Research*, 73: 1093–1130. A preliminary version of this paper appeared in the Proceedings of AAMAS-2021.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Brandt, F.; and Lederer, P. 2023. Characterizing the Top Cycle via Strategyproofness. *Theoretical Economics*, 18(2): 837–883.
- Brandt, F.; and Lederer, P. 2024. Weak Strategyproofness in Randomized Social Choice. <https://arxiv.org/abs/2412.11977>
- Brandt, F.; Lederer, P.; and Romen, R. 2024. Relaxed Notions of Condorcet-Consistency and Efficiency for Strategyproof Social Decision Schemes. *Social Choice and Welfare*. A preliminary version of this paper appeared in the Proceedings of AAMAS-2022.
- Brandt, F.; Lederer, P.; and Suksompong, W. 2023. Incentives in Social Decision Schemes with Pairwise Comparison Preferences. *Games and Economic Behavior*, 142: 266–291. A preliminary version of this paper appeared in the Proceedings of IJCAI-2022.
- Brandt, F.; Saile, C.; and Stricker, C. 2022. Strategyproof Social Choice When Preferences and Outcomes May Contain Ties. *Journal of Economic Theory*, 202: 105447. A preliminary version of this paper appeared in the Proceedings of AAMAS-2018.
- Feldman, A. M. 1980. Strongly nonmanipulable multi-valued collective choice rules. *Public Choice*, 35: 503–509.
- Fishburn, P. C. 1972. Even-chance lotteries in social choice theory. *Theory and Decision*, 3(1): 18–40.
- 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.
- Lederer, P. 2021. Strategyproof Randomized Social Choice for Restricted Sets of Utility Functions. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 306–312. A preliminary version of this paper appeared in the Proceedings of IJCAI-2021.
- Mennle, T.; and Seuken, S. 2021. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory*, 191: 105–144.
- Postlewaite, A.; and Schmeidler, D. 1986. Strategic behaviour and a notion of Ex Ante efficiency in a voting model. *Social Choice and Welfare*, 3(1): 37–49.
- Procaccia, A. D. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 836–841.
- 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. 2011. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIEs*, 2(4): 515–527.