Strategyproof Randomized Social Choice for Restricted Sets of Utility Functions

Patrick Lederer

Technische Universtiät München

ledererp@in.tum.de

Abstract

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

1 Introduction

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

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

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

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

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

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

2   Related Work

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

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

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

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

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

3   Preliminaries

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

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

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

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

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

4 U-Strategyproofness

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

Even though we assume the existence of utility functions,

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

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

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

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

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

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

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

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

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

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

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

late Benoît's impossibility as follows.

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

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

5 Results

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

5.1 k-unanimity

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

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

---

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

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

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

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

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

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

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

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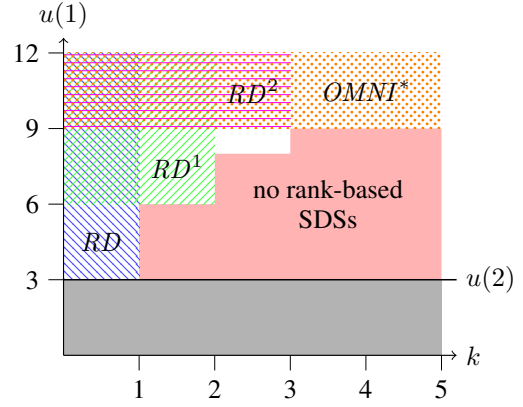

Figure 1: Illustration of Theorem 1 and Theorem 2. We assume that there are 5 alternatives and consider a utility function u with $u(2) = 3$, $u(3) = 2$, $u(4) = 1$, and $u(5) = 0$. The figure shows for which values of $u(1)$ the SDSs RD (blue area), RD^1 (green area), RD^2 (magenta area), and $OMNI^*$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

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

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

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

5.2 Condorcet-consistency

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

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

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

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

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

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

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

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

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

6 Conclusion and Discussion

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

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

Acknowledgments

References

[Abramowitz *et al.*, 2019] Ben Abramowitz, Elliot Anshele-vich, and Whennan Zhu. Awareness of voter passion greatly improves the distortion of metric social choice. In *Proceedings of the 15th International Conference on Web and Internet Economics*, pages 3–16. Springer, 2019.

[Aziz *et al.*, 2018] Haris Aziz, Florian Brandl, Felix Brandt, and Markus Brill. On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110:1–18, 2018. Preliminary results appeared in the Proceedings of AAAI-2013 and AAMAS-2014.

[Barberà, 1979] Salvador Barberà. Majority and positional voting in a probabilistic framework. *Review of Economic Studies*, 46(2):379–389, 1979.

[Benoît, 2002] Jean-Pierre Benoît. Strategic manipulation in voting games when lotteries and ties are permitted. *Journal of Economic Theory*, 102(2):421–436, 2002.

[Bogomolnaia *et al.*, 2005] Anna Bogomolnaia, Hervé Moulin, and Richard. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.

[Brandl *et al.*, 2018] Florian Brandl, Felix Brandt, Manuel Eberl, and Christian Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, 65(2):1–28, 2018. Preliminary results appeared in the Proceedings of IJCAI-2016.

[Brandt *et al.*, 2013] Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, chapter 6, pages 213–283. MIT Press, 2nd edition, 2013.

[Brandt *et al.*, 2016a] Felix Brandt, Markus Brill, and Paul Harrenstein. Tournament solutions. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2016.

[Brandt *et al.*, 2016b] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[Brandt, 2017] Felix Brandt. Rolling the dice: Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 1, pages 3–26. AI Access, 2017.

[Chatterji and Zeng, 2018] Shurojit Chatterji and Huaxia Zeng. On random social choice functions with the tops-only property. *Games and Economic Behavior*, 109:413–435, 2018.

[Chevaleyre *et al.*, 2007] Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, volume 4362

of *Lecture Notes in Computer Science (LNCS)*, pages 51–69. Springer-Verlag, 2007.

[Duggan and Schwartz, 2000] John Duggan and Thomas Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.

[Dutta *et al.*, 2007] Bhaskar Dutta, Hans Peters, and Arunava Sen. Strategy-proof cardinal decision schemes. *Social Choice and Welfare*, 28(1):163–179, 2007.

[Ehlers *et al.*, 2002] Lars Ehlers, Hans Peters, and Ton Storcken. Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences. *Journal of Economic Theory*, 105(2):408–434, 2002.

[Endriss, 2017] Ulle Endriss, editor. *Trends in Computational Social Choice*. AI Access, 2017.

[Gibbard, 1973] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

[Gibbard, 1977] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.

[Gross *et al.*, 2017] Stephen Gross, Elliot Anshelevich, and Lirong Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 544–550, 2017.

[Hoang, 2017] Lê Nguyên Hoang. Strategy-proofness of the randomized Condorcet voting system. *Social Choice and Welfare*, 48(3):679–701, 2017.

[Hylland, 1980] Aanund Hylland. Strategyproofness of voting procedures with lotteries as outcomes and infinite sets of strategies. Mimeo, 1980.

[Mennle and Seuken, 2021] Timo Mennle and Sven Seuken. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory*, 191:105–144, 2021.

[Nandeibam, 2013] Shasikanta Nandeibam. The structure of decision schemes with cardinal preferences. *Review of Economic Design*, 17(3):205–238, 2013.

[Procaccia and Rosenschein, 2006] Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of 10th International Workshop on Cooperative Information Agents*, pages 317—331. Springer, 2006.

[Satterthwaite, 1975] Mark Allen Satterthwaite. Strategyproofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[Sen, 2011] Arunava Sen. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIEs*, 2(4):515–527, 2011.