

Welfare Maximization Entices Participation

Florian Brandl Felix Brandt Johannes Hofbauer

Technische Universität München

{brandlf1,brandtf,hofbauej}@in.tum.de

We consider randomized public good mechanisms with optional participation. Preferences over lotteries are modeled using skew-symmetric bilinear (SSB) utility functions, a generalization of classic von Neumann-Morgenstern utility functions. We show that every welfare-maximizing mechanism entices participation and that the converse holds under additional assumptions. As a corollary, we obtain a characterization of an attractive randomized voting rule that satisfies Condorcet-consistency and entices participation. This stands in contrast to Moulin’s well-known *no-show paradox* (J. of Econ. Theory, 45, 53–64, 1988), which shows that no deterministic voting rule can satisfy both properties simultaneously.

Keywords: SSB utility, no show paradox, maximal lotteries, stochastic dominance
JEL Classifications Codes: D7, D6

1 Introduction

The question we pursue in this paper is whether public good mechanisms that maximize utilitarian social welfare entice participation in the sense that no group of agents is ever better off by not participating in the mechanism. For standard utility representations, where each agent assigns a numerical value to each alternative, this is obviously true. In the context of voting, this entails the well-known fact that scoring rules (such as plurality rule or Borda’s rule) entice participation.

Our focus in this paper lies on mechanisms that return *lotteries* over alternatives. Clearly, if preferences over lotteries are given by von Neumann-Morgenstern utility functions, every randomized mechanism that maximizes some affine combination of these utility functions will also entice participation. We consider welfare maximization under much loosened assumptions about preferences over lotteries. In particular, we assume that preferences over lotteries are given by *skew-symmetric bilinear (SSB) utility functions*, which assign a numerical value to each *pair* of lotteries. One lottery is preferred

to another lottery if the SSB utility for this pair is positive. SSB utility theory is a generalization of von Neumann-Morgenstern utility theory, which neither requires the controversial independence axiom nor transitivity (see, e.g., Fishburn, 1982, 1984b, 1988). Independence requires that if lottery p is preferred to lottery q , then a coin toss between p and a third lottery r is preferred to a coin toss between q and r (with the same coin used in both cases). There is experimental evidence that both independence and transitivity are violated systematically by human decision makers. The Allais Paradox (Allais, 1953) is perhaps the most famous example pointing out violations of independence. Detailed reviews of such violations, including those reported by Kahneman and Tversky (1979), have been provided by Machina (1983, 1989) and McClennen (1988).¹ Mas-Colell et al. (1995, p. 181) conclude that “because of the phenomena illustrated [...] the search for a useful theory of choice under uncertainty that does not rely on the independence axiom has been an active area of research”. Even the widely accepted transitivity axiom has come under increasing scrutiny (see, e.g., May, 1954; Fishburn, 1970; Bar-Hillel and Margalit, 1988; Fishburn, 1991; Anand, 1993, 2009). Anand (2009, p. 156) concludes that “once considered a cornerstone of rational choice theory, the status of transitivity has been dramatically reevaluated by economists and philosophers in recent years”.

SSB utility theory dispenses with independence and transitivity and therefore can accommodate both effects, the Allais Paradox and the preference reversal phenomenon. Despite a lack of transitivity, the Minimax Theorem (von Neumann, 1928) implies that, for every SSB utility function and every compact and convex set of lotteries, there is a maximal lottery, i.e., a lottery that is weakly preferred to any other lottery within the set. In other words, the main appeal of transitivity—the existence of maximal elements—remains intact. For a more thorough discussion of SSB utility theory, the reader is referred to Fishburn (1984b, 1988).

Our main theorems show that every SSB welfare-maximizing mechanism entices participation and that the converse holds under additional assumptions. This has interesting consequences for the special case of randomized voting rules, i.e., mechanisms that map an ordinal preference profile to a lottery. As first observed by Fishburn and Brams (1983), there are voting rules where agents are better off by abstaining from the election. This phenomenon is called the *no-show paradox* and Moulin (1988) has famously shown that it pertains to all *deterministic* Condorcet-consistent voting rules.² By contrast, our first theorem implies that *maximal lotteries*, a randomized Condorcet-consistent voting rule due to Fishburn (1984a), entices participation in a very natural sense. The underlying notion of participation we consider is stronger than group-participation with respect to stochastic dominance. Moreover, the second theorem implies that, under a mild technical assumption, maximal lotteries can even be characterized using participation and Condorcet-consistency.

¹Fishburn and Wakker (1995) give an interesting historical perspective on the independence axiom.

²This result has triggered a number of extensions by strengthening the theorem (Holzman, 1988; Sanver and Zwicker, 2009; Peters, 2017; Brandt et al., 2017a), allowing for weak preferences (Duddy, 2014), and considering set-valued or randomized voting rules (Pérez, 2001; Jimeno et al., 2009; Brandl et al., 2015a,b; Brandt et al., 2017a) and random assignment rules (Brandl et al., 2017).

2 Preliminaries

Let $\mathbb{N} = \{1, 2, \dots\}$ be an infinite set of agents and $\mathcal{F}(\mathbb{N})$ the set of all finite and non-empty subsets of \mathbb{N} . Moreover, A is a finite set of alternatives and $\Delta(A)$ the set of all *lotteries* (or *probability distributions*) over A . A lottery is *degenerate* if it puts all probability on a single alternative. We assume that preferences over lotteries are given by *skew-symmetric bilinear (SSB) utility* functions as introduced by Fishburn (1982) (see, also, Fishburn, 1984b, 1988). An SSB function ϕ is a function from $\Delta(A) \times \Delta(A) \rightarrow \mathbb{R}$ that is skew-symmetric and bilinear, i.e.,

$$\begin{aligned}\phi(p, q) &= -\phi(q, p) \quad \text{and} \\ \phi(\lambda p + (1 - \lambda)q, r) &= \lambda\phi(p, r) + (1 - \lambda)\phi(q, r)\end{aligned}$$

for all $p, q, r \in \Delta(A)$ and $\lambda \in [0, 1]$. Lottery p is weakly preferred to lottery q if and only if $\phi(p, q) \geq 0$. Note that, by skew-symmetry, linearity in the first argument implies linearity in the second argument and that, due to bilinearity, ϕ is completely determined by its function values for degenerate lotteries. Thus, every SSB function can be represented by a skew-symmetric matrix in $\mathbb{R}^{A \times A}$. As mentioned in Section 1, SSB utility theory is more general than linear expected utility theory due to von Neumann and Morgenstern (1947), henceforth vNM. In particular, every vNM function u is equivalent to an SSB function ϕ^u , where $\phi^u(p, q) = u(p) - u(q)$, in the sense that both functions induce the same preferences over lotteries.

We denote by $\Phi \subseteq \mathbb{R}^{A \times A}$ the set of possible SSB functions called the *domain*. For every $N \in \mathcal{F}(\mathbb{N})$, let $\phi_N = (\phi_i)_{i \in N} \in \Phi^N$ be a vector of SSB functions. If $N = \{i\}$, we write ϕ_i with slight abuse of notation. For every $N \in \mathcal{F}(\mathbb{N})$ and $\phi_N \in \Phi^N$, we define $\phi_N^\Sigma = \sum_{i \in N} \phi_i$. A lottery p is *welfare-maximizing* for ϕ_N if

$$\phi_N^\Sigma(p, q) \geq 0 \text{ for all } q \in \Delta(A). \quad (\text{welfare maximization})$$

Thus, a lottery is welfare-maximizing if it is weakly preferred to any other lottery with respect to the accumulated SSB functions. Such lotteries can be identified with mixed maximin strategies of the symmetric zero-sum game given by ϕ_N^Σ .³ If agents are endowed with vNM functions, welfare maximization is equivalent to maximizing the sum of expected utilities. Hence, welfare maximization for SSB functions generalizes welfare maximization for vNM functions. While there always exists a *degenerate* welfare-maximizing lottery for vNM functions, this does not hold anymore in the more general model of SSB utilities. However, the existence of a (not necessarily degenerate) welfare-maximizing lottery is guaranteed by the Minimax Theorem (von Neumann, 1928).

³ ϕ_N^Σ can be identified with a symmetric zero-sum game as follows: the set of mixed strategies of both players is given by $\Delta(A)$. The payoff of the row player when he plays strategy p and the column player plays strategy q is given by $\phi_N^\Sigma(p, q) = p^T \phi_N^\Sigma q$.

The following example illustrates the definitions. Let $N = \{1, 2\}$, $A = \{a, b, c\}$, and

$$\phi_1 = \begin{pmatrix} a & b & c \\ 0 & 2 & 4 \\ -2 & 0 & 2 \\ -4 & -2 & 0 \end{pmatrix} \begin{matrix} a \\ b \\ c \end{matrix} \quad \text{and} \quad \phi_2 = \begin{pmatrix} a & b & c \\ 0 & -3 & -3 \\ 3 & 0 & -3 \\ 3 & 3 & 0 \end{pmatrix} \begin{matrix} a \\ b \\ c \end{matrix} .$$

The SSB function ϕ_1 is equivalent to a vNM utility function that assign utilities 4, 2, and 0 to a , b , and c , respectively. The unique most preferred lottery of agent 1 is the lottery with probability 1 on alternative a . The SSB function ϕ_2 is an example where the intensities of all pairwise comparisons are identical in the sense that they are assigned the same numerical value. This special case will be discussed later in the paper. Agent 2 prefers the lottery with probability 1 on c the most. The sum of both SSB functions is

$$\phi_N^\Sigma = \phi_1 + \phi_2 = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

and the unique welfare maximizing lottery according to ϕ_N^Σ is $p = 1/3 a + 1/3 b + 1/3 c$ since $\phi_N^\Sigma(p, x) = 0$ for all $x \in A$ and for every $p' \in \Delta(A) \setminus \{p\}$, there is some $x \in A$ such that $\phi_N^\Sigma(p', x) < 0$. In particular, there is no degenerate lottery that is welfare maximizing for ϕ_N^Σ .

Our central objects of study are *mechanisms* that map a vector of SSB functions to a lottery. A mechanism f is *welfare-maximizing* if it always returns welfare-maximizing lotteries for vectors of SSB functions from Φ , i.e., for all $N \in \mathcal{F}(\mathbb{N})$ and $\phi_N \in \Phi^N$, $f(\phi_N)$ is welfare-maximizing.

We will consider welfare maximization in settings with optional participation. A mechanism satisfies *welfare participation* if participating in the mechanism is always at least as good as not participating in terms of accumulated SSB welfare. Formally, for every $N \in \mathcal{F}(\mathbb{N})$, $S \subsetneq N$, and $\phi_N \in \Phi^N$,

$$\phi_S^\Sigma(f(\phi_N), f(\phi_{N \setminus S})) \geq 0. \quad (\text{welfare participation})$$

As we will see in Section 4, this strong notion of participation has important consequences even in settings in which the interpersonal comparison of utility is problematic, such as in voting.

3 Welfare Maximization and Participation

We are now ready to prove two theorems that highlight the relationship between welfare maximization and welfare participation. The first result shows that welfare maximization implies welfare participation.

Theorem 1. *Every welfare-maximizing mechanism satisfies welfare participation.*

Proof. Let $N \in \mathcal{F}(\mathbb{N})$, $S \subsetneq N$, $\phi_N \in \Phi^N$, and f be a welfare-maximizing mechanism. For $p = f(\phi_N)$ and $p' = f(\phi_{N \setminus S})$, we then have that

$$\begin{aligned}\phi_N^\Sigma(p, q) &\geq 0 \text{ for all } q \in \Delta(A), \text{ and} \\ \phi_{N \setminus S}^\Sigma(p', q) &\geq 0 \text{ for all } q \in \Delta(A),\end{aligned}$$

since, by assumption, f is welfare-maximizing for ϕ_N and $\phi_{N \setminus S}$. Thus, it follows that

$$\phi_S^\Sigma(p, p') = \phi_N^\Sigma(p, p') - \phi_{N \setminus S}^\Sigma(p, p') = \underbrace{\phi_N^\Sigma(p, p')}_{\geq 0} + \underbrace{\phi_{N \setminus S}^\Sigma(p', p)}_{\geq 0} \geq 0.$$

The second equality follows from skew-symmetry of $\phi_{N \setminus S}^\Sigma$. The inequality follows from the fact that f is welfare-maximizing for ϕ_N and $\phi_{N \setminus S}$. Hence, f satisfies welfare participation. \square

Clearly, Theorem 1 also holds for Cartesian domains that are not symmetric among agents. The converse of Theorem 1 does not hold in full generality as every constant function satisfies welfare participation but fails to be welfare-maximizing. However, for sufficiently rich domains, the converse holds for mechanisms that satisfy two additional properties.

These properties are homogeneity and weak welfare maximization. Homogeneity is well-known in social choice theory (see, e.g., Smith, 1973; Young, 1977).⁴ In order to define it, we introduce notation for making copies of utility vectors. For all $N \in \mathcal{F}(\mathbb{N})$ and $k \in \mathbb{N}$, let $kN = \{i + l \max(N) : i \in N \text{ and } l \in \{0, \dots, k-1\}\}$ and $\phi_i = \phi_j$ if $i \equiv j \pmod{\max(N)}$. A mechanism f is *homogeneous* if replicating the entire set of agents does not affect the outcome, i.e., $f(\phi_N) = f(\phi_{kN})$ for all $k \in \mathbb{N}$ and $\phi_N \in \Phi^N$. Homogeneity thus prescribes that the outcome of a mechanism should not depend on the absolute number of agents as long as the relative composition of preferences is completely identical.

In certain cases, there is no need for randomization in order to achieve welfare maximization, because there is a degenerate lottery that is weakly preferred to every other lottery in terms of accumulated welfare. A mechanism f is *weakly welfare-maximizing* if $f(\phi_N) = p$ whenever p is degenerate and the unique welfare-maximizing lottery for ϕ_N . Weak welfare maximization can thus be seen as a weak version of *ex post* welfare maximization (recall that degenerate welfare-maximizing lotteries need not exist for SSB utility functions). It is weaker than *ex post* welfare maximization because it is only applicable when there is a *unique* welfare-maximizing alternative.

The following lemma shows that a degenerate lottery is the unique welfare-maximizing lottery if and only if it is strictly preferred to every other degenerate lottery in terms of welfare. We slightly abuse notation by identifying degenerate lotteries with alternatives.

Lemma 1. *Let $\phi_N \in \Phi^N$ and $x \in A$. Then, x is the unique welfare-maximizing lottery for ϕ_N if and only if $\phi_N^\Sigma(x, y) > 0$ for all $y \in A \setminus \{x\}$.*

⁴In the context of fair division, this property is often called replication invariance.

Proof. For the direction from left to right, let $x \in A$ be the unique welfare-maximizing lottery for ϕ_N . Assume for contradiction that $B = \{y \in A \setminus \{x\} : \phi_N^\Sigma(x, y) \leq 0\} \neq \emptyset$. Let $p' \in \Delta(B)$ denote some lottery on B that is welfare-maximizing for ϕ_N^Σ restricted to B , i.e., $\phi_N^\Sigma(p', q) \geq 0$ for all $q \in \Delta(B)$ and $p \in \Delta(A)$ be the lottery that is equal to p' on B and 0 otherwise, i.e., $p(y) = p'(y)$ for all $y \in B$ and $p(y) = 0$ for all $y \in A \setminus B$. By the choice of B , we have that $\phi_N^\Sigma(p, y) \geq 0$ for all $y \in B \cup \{x\}$. Moreover, for $\varepsilon > 0$ small enough, we have $\phi_N^\Sigma(\varepsilon p + (1 - \varepsilon)x, y) \geq 0$ for all $y \in A \setminus B$, since $\phi_N^\Sigma(x, y) > 0$. Hence, $\varepsilon p + (1 - \varepsilon)x$ is also welfare-maximizing for ϕ_N^Σ and x cannot be the unique welfare-maximizing lottery.

The direction from right to left follows from linearity of ϕ_N^Σ . \square

Next, we aim at showing that every homogeneous and weakly welfare-maximizing mechanism that satisfies welfare participation is welfare-maximizing. This result does not hold without making some assumptions about the richness of the domain of preferences. We therefore impose two domain conditions. A domain Φ is *symmetric* if, for all $\phi \in \Phi$, $-\phi \in \Phi$. Symmetry prescribes that it is always possible to completely disagree with any of the other agents. It can thus be seen as an axiom that ensures a minimal degree of freedom with respect to the structure of preferences. A domain Φ is *non-imposing* if, for all $x \in A$, there are $N \in \mathcal{F}(\mathbb{N})$ and $\phi_N \in \Phi^N$ such that $\phi_N^\Sigma(x, y) > 0$ for all $y \in A \setminus \{x\}$. Non-imposition demands that for every *degenerate* lottery x there is a preference profile such that x is the unique welfare-maximizing lottery. Hence, non-imposition, prevents an extreme bias against certain alternatives in the domain of admissible preferences.

Theorem 2. *Let Φ be a symmetric and non-imposing domain. Every homogeneous, weakly welfare-maximizing mechanism on Φ that satisfies welfare participation is welfare-maximizing.*

Proof. Let f be a homogeneous, weakly welfare-maximizing mechanism that satisfies welfare participation. Assume for contradiction that f is not welfare-maximizing for some $N \in \mathcal{F}(\mathbb{N})$ and $\phi_N \in \Phi^N$, i.e., there is a lottery q such that $\phi_N^\Sigma(p, q) < 0$, where $p = f(\phi_N)$. By linearity of ϕ_N^Σ , there is an alternative x such that $\phi_N^\Sigma(p, x) = c < 0$. Now, let \bar{N} be a set of agents disjoint from N and $\phi_{\bar{N}} \in \Phi^{\bar{N}}$ such that $\phi_{\bar{N}} = -\phi_N$. Furthermore, let J be a set of agents that has empty intersection with N and \bar{N} and $\phi_J \in \Phi^J$ such that $\phi_J^\Sigma(x, y) > 0$ for all $y \in A \setminus \{x\}$. The sets \bar{N} and J exist, since \mathbb{N} is infinite and Φ is symmetric and non-imposing by assumption. Moreover, let $d = \phi_J^\Sigma(x, p)$ and k be an integer such that $kc + d < 0$. It follows from homogeneity of f that $f(\phi_{kN}) = f(\phi_N) = p$. By definition of $\phi_{\bar{N}}$ it follows that $\phi_{kN}^\Sigma + \phi_{k\bar{N}}^\Sigma + \phi_J^\Sigma = \phi_J^\Sigma$. Hence, x is the unique welfare-maximizing lottery for $\phi_{kN \cup k\bar{N} \cup J}$. This implies that $x = f(\phi_{kN \cup k\bar{N} \cup J})$ by Lemma 1 and weak welfare maximization of f . Furthermore, we have

$$\phi_{k\bar{N}}^\Sigma(p, x) + \phi_J^\Sigma(p, x) = -(kc + d) > 0.$$

Hence, the set of agents $k\bar{N} \cup J$ prefers abstaining (which yields p) to participating (which yields x). This contradicts welfare participation. \square

4 Randomized Voting Rules

We now turn to the important special case in which only ordinal preferences between alternatives are known and consider *randomized voting rules*, i.e., functions that map an ordinal preference profile to a lottery. Ordinal preferences are given in the form of complete binary relations, which can be conveniently represented by SSB functions whose entries are restricted to $\{-1, 0, +1\}$, where $\phi_i(x, y) = +1$ if agent i prefers x to y , $\phi_i(x, y) = -1$ if he prefers y to x , and $\phi_i(x, y) = 0$ if he is indifferent. We refer to this representation as the *canonical* utility representation of ordinal preferences and define randomized voting rules on the domain

$$\Phi_{PC} = \{\phi \in \{-1, 0, +1\}^{A \times A} : \phi(x, y) = -\phi(y, x) \text{ for all } x, y \in A\}.$$

Every such representation entails a complete preference relation over lotteries of alternatives (called the *pairwise comparison (PC) preference extension*). The natural interpretation of this relation is that lottery p is preferred to lottery q if it is more likely that p yields a better alternative than q than *vice versa*. This preference relation cannot be represented by a vNM utility function and may be intransitive, even when preferences over alternatives are transitive. For more details, please see Blavatskyy (2006), Aziz et al. (2015, 2018), and Brandl and Brandt (2020). Aziz et al. (2015) have shown that the preference relation induced by the canonical utility representation is a refinement of stochastic dominance, i.e., if p stochastically dominates q , then p is also preferred to q with respect to the canonical utility representation. A lottery *stochastically dominates* another if the former yields more expected utility than the latter for every vNM function that is consistent with the ordinal preferences. A randomized voting rule satisfies *group-participation with respect to stochastic dominance (SD-group-participation)* if no group of agents can abstain from f such that *each* of the agents is individually better off with respect to stochastic dominance. When only requiring this for singleton groups, the corresponding property is called *SD-participation*.

Proposition 1. *Every randomized voting rule that satisfies welfare participation satisfies SD-group-participation.*

Proof. Let $N \in \mathcal{F}(\mathbb{N})$, $S \subsetneq N$, $\phi_N \in \Phi_{PC}^N$, and f a randomized voting rule that satisfies welfare participation. Welfare participation of f implies $\sum_{i \in N} \phi_i(f(\phi_N), f(\phi_{N \setminus S})) \geq 0$. In particular, there is $i \in S$ such that $\phi_i(f(\phi_N), f(\phi_{N \setminus S})) \geq 0$. This implies that $f(\phi_{N \setminus S})$ does not stochastically dominate $f(\phi_N)$. \square

SD-participation is not easily satisfied. For example, Brandl et al. (2015b) have shown that no majoritarian randomized voting rule can satisfy *SD-participation* and *ex post efficiency*.⁵ By leveraging the results obtained in Section 3, we can derive a number of statements concerning randomized voting rules that return so-called maximal lotteries. A lottery is *maximal* for a given ordinal preference profile if it is welfare-maximizing

⁵A voting rule is majoritarian if its output only depends on the pairwise majority relation. A mechanism is *ex post efficient* if it always assigns probability 0 to Pareto dominated alternatives.

for the canonical utility representation. In the context of voting, maximal lotteries are almost always unique and randomized voting rules that return maximal lotteries form an attractive class of randomized voting rules (Fishburn, 1984a; Brandl et al., 2016; Aziz et al., 2018; Brandl and Brandt, 2020). Maximal lotteries have also been considered in the context of private good settings, such as randomized matching markets and house allocation (e.g., Kavitha et al., 2011; Aziz et al., 2013; Brandt et al., 2017b).

First, Theorem 1 and Proposition 1 imply that returning maximal lotteries satisfies *SD*-group-participation, independently of how ties between maximal lotteries are broken.

Corollary 1. *Every randomized voting rule that returns maximal lotteries satisfies *SD*-group-participation.*

To illustrate maximal lotteries as well as the notions of welfare participation and *SD*-participation consider the following example. Let $N = \{1, 2, 3\}$, $A = \{a, b, c\}$, and agents' preferences be defined such that first agent prefers a to b to c , the second b to c to a , and the third c to a to b . Then,

$$\phi_N^\Sigma = \sum_{i \in N} \phi_i = \begin{pmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{pmatrix} + \begin{pmatrix} 0 & -1 & -1 \\ 1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & -1 \\ 1 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}.$$

Just like in the example given in Section 2, the unique maximal lottery according to ϕ_N^Σ is $p = 1/3 a + 1/3 b + 1/3 c$. Now, assume that the first agent leaves the electorate. In the resulting two-agent profile, every lottery that randomizes between b and c happens to be maximal. Observe that p is preferred to any such lottery $q = \lambda b + (1 - \lambda)c$, $\lambda \in [0, 1]$, with respect to the canonical utility representation of the first agent because

$$\phi_1(p, q) = \begin{pmatrix} 1/3 & 1/3 & 1/3 \end{pmatrix} \begin{pmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ \lambda \\ 1 - \lambda \end{pmatrix} = 2/3(1 - \lambda) \geq 0.$$

Furthermore, this preference is strict unless $\lambda = 1$. As a result, p is preferred to q with respect to stochastic dominance whenever $\lambda \leq 2/3$; otherwise, the lotteries are incomparable. Hence, *SD*-participation is also not violated.

Theorem 2 entails an axiomatic characterization of randomized voting rules that return maximal lotteries. It is easily seen that Φ_{PC} satisfies symmetry and non-imposition. Alternative x is called a *Condorcet winner* for a given preference profile $\phi_N \in \Phi_{PC}^N$ if a majority of agents prefers it to any other alternative, i.e., $\phi_N^\Sigma(x, y) > 0$ for all $y \in A \setminus \{x\}$. A randomized voting rule is *Condorcet-consistent* if it always puts probability 1 on a Condorcet winner. It follows from Lemma 1 that Condorcet-consistency is equivalent to weak welfare maximization for the domain Φ_{PC} . We thus obtain the following characterization as a corollary of Theorem 2.

Corollary 2. *Every homogeneous, Condorcet-consistent, randomized voting rule that satisfies welfare participation returns maximal lotteries.*

As discussed in Section 1, Corollary 2 can be contrasted with a classic result by Moulin (1988) who showed that no Condorcet-consistent deterministic voting rule satisfies participation.

5 Concluding Comments

Remark 1. The axioms used in Corollary 2 (and therefore also the ones used in Theorem 2) are independent of each other. The rule that always returns the uniform lottery over all alternatives satisfies homogeneity and welfare participation, but violates Condorcet-consistency. The rule that puts probability 1 on the Condorcet winner whenever it exists and otherwise returns the uniform lottery satisfies homogeneity and Condorcet-consistency, but violates welfare participation. A rule that satisfies Condorcet-consistency and welfare participation, but violates homogeneity, can be constructed as follows. Let $|A| = 3$, add some small $\varepsilon > 0$ to every positive entry and $-\varepsilon$ to every negative entry of ϕ_N^Σ and return a welfare-maximizing lottery of the resulting matrix ϕ' . If there are several such lotteries, there has to be a degenerate welfare-maximizing lottery (since $|A| = 3$) and the uniform lottery over those degenerate welfare-maximizing lotteries for which the corresponding rows in ϕ' contain the largest number of positive entries should be returned.

Remark 2. Note that Corollary 2 also holds for any symmetric and non-imposing subdomain of Φ_{PC} , and hence, in particular, for the domain consisting of all transitive and complete preference relations and for the domain consisting of all transitive, complete, and anti-symmetric preference relations.

Remark 3. Corollary 2 does not hold if welfare participation is weakened to *SD*-group-participation. For example, the voting rule that returns the Condorcet winner if one exists and the uniform lottery over all alternatives otherwise is homogeneous, Condorcet-consistent, and satisfies *SD*-group-participation.⁶

Remark 4. Using similar arguments as in the proof of Theorem 2, maximal lotteries can also be characterized by replacing Condorcet-consistency with non-imposition and cancellation in Corollary 2. Non-imposition requires every *degenerate* lottery to be chosen for at least one preference profile; cancellation prescribes that agents with completely opposed preferences cancel each other out (see Young, 1974).

Remark 5. It is also possible to define a stronger notion of *SD*-participation where the outcome when participating has to stochastically dominate the outcome when abstaining. For this notion, Moulin's (1988) result remains intact (Brandt et al., 2017a, Theorem 9).

Remark 6. The proof of Theorem 1 can be adapted to show that welfare-maximizing mechanisms satisfy *one-way monotonicity* (Sanver and Zwicker, 2009). As a consequence, randomized voting rules that return maximal lotteries satisfy *SD*-one-way-monotonicity. This stands in contrast to Sanver and Zwicker (2009) and Peters (2017) who have shown that no deterministic voting rule satisfies Condorcet-consistency and half-way-monotonicity, a weakening of both one-way-monotonicity and participation.

Remark 7. Fishburn (1988, p. 85) has shown that a preference relation on $\Delta(A)$ is continuous and convex if and only if it can be represented by a utility function that maps

⁶However, this rule violates *ex post* efficiency let alone the stronger notion of *SD*-efficiency, which is satisfied by maximal lotteries.

from $\Delta(A) \times \Delta(A)$ to \mathbb{R} and is sign skew-symmetric and linear in the first argument. Theorem 1 and Remark 6 also hold for this more general class of utility functions. The existence of maximal elements in this framework is guaranteed by a result due to Sonnenschein (1971).

Acknowledgments

This material is based upon work supported by Deutsche Forschungsgemeinschaft under grants BR 2312/10-1 and BR 2312/11-1. We are grateful to an anonymous editor and two anonymous reviewers for helpful comments.

References

- M. Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21(4):503–546, 1953.
- P. Anand. The philosophy of intransitive preference. *The Economic Journal*, 103(417):337–346, 1993.
- P. Anand. Rationality and intransitive preference: Foundations for the modern view. In P. Anand, P. K. Pattanaik, and C. Puppe, editors, *The Handbook of Rational and Social Choice*, chapter 6. Oxford University Press, 2009.
- H. Aziz, F. Brandt, and P. Stursberg. On popular random assignments. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, volume 8146 of *Lecture Notes in Computer Science (LNCS)*, pages 183–194. Springer-Verlag, 2013.
- H. Aziz, F. Brandl, and F. Brandt. Universal Pareto dominance and welfare for plausible utility functions. *Journal of Mathematical Economics*, 60:123–133, 2015.
- H. Aziz, F. Brandl, F. Brandt, and M. Brill. On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110:1–18, 2018.
- M. Bar-Hillel and A. Margalit. How vicious are cycles of intransitive choice? *Theory and Decision*, 24(2):119–145, 1988.
- P. R. Blavatsky. Axiomatization of a preference for most probable winner. *Theory and Decision*, 60(1):17–33, 2006.
- F. Brandl and F. Brandt. Arrovian aggregation of convex preferences. *Econometrica*, 88(2):799–844, 2020.
- F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 18–24, 2015a.

- F. Brandl, F. Brandt, and J. Hofbauer. Incentives for participation and abstention in probabilistic social choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1411–1419, 2015b.
- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.
- F. Brandl, F. Brandt, and J. Hofbauer. Random assignment with optional participation. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 326–334, 2017.
- F. Brandt, C. Geist, and D. Peters. Optimal bounds for the no-show paradox via SAT solving. *Mathematical Social Sciences*, 90:18–27, 2017a. Special Issue in Honor of Hervé Moulin.
- F. Brandt, J. Hofbauer, and M. Suderland. Majority graphs of assignment problems and properties of popular random assignments. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 335–343, 2017b.
- C. Duddy. Condorcet’s principle and the strong no-show paradoxes. *Theory and Decision*, 77(2):275–285, 2014.
- P. C. Fishburn. The irrationality of transitivity in social choice. *Behavioral Science*, 15:119–123, 1970.
- P. C. Fishburn. Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26(1):31–67, 1982.
- P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984a.
- P. C. Fishburn. SSB utility theory: An economic perspective. *Mathematical Social Sciences*, 8(1):63–94, 1984b.
- P. C. Fishburn. *Nonlinear preference and utility theory*. The Johns Hopkins University Press, 1988.
- P. C. Fishburn. Nontransitive preferences in decision theory. *Journal of Risk and Uncertainty*, 4(2):113–134, 1991.
- P. C. Fishburn and S. J. Brams. Paradoxes of preferential voting. *Mathematics Magazine*, 56(4):207–214, 1983.
- P. C. Fishburn and P. Wakker. The invention of the independence condition for preferences. *Management Science*, 41(7):1130–1144, 1995.
- R. Holzman. To vote or not to vote: What is the quota? *Discrete Applied Mathematics*, 22(2):133–141, 1988.

- J. L. Jimeno, J. Pérez, and E. García. An extension of the Moulin No Show Paradox for voting correspondences. *Social Choice and Welfare*, 33(3):343–459, 2009.
- D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
- T. Kavitha, J. Mestre, and M. Nasre. Popular mixed matchings. *Theoretical Computer Science*, 412(24):2679–2690, 2011.
- M. J. Machina. Generalized expected utility analysis and the nature of observed violations of the independence axiom. In B. Stigum and F. Wenstop, editors, *Foundations of Utility and Risk Theory with Applications*, chapter 5. Springer, 1983.
- M. J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, 1989.
- A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- K. May. Intransitivity, utility, and the aggregation of preference patters. *Econometrica*, 22(1):1–13, 1954.
- E. F. McClennen. Sure-thing doubts. In P. Gärdenfors and N.-E. Sahlin, editors, *Decision, Probability and Utility*, chapter 10. Cambridge University Press, 1988.
- H. Moulin. Condorcet’s principle implies the no show paradox. *Journal of Economic Theory*, 45(1):53–64, 1988.
- J. Pérez. The Strong No Show Paradoxes are a common flaw in Condorcet voting correspondences. *Social Choice and Welfare*, 18(3):601–616, 2001.
- D. Peters. Condorcet’s principle and the preference reversal paradox. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 455–469, 2017.
- M. R. Sanver and W. S. Zwicker. One-way monotonicity as a form of strategy-proofness. *International Journal of Game Theory*, 38(4):553–574, 2009.
- J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.
- H. Sonnenschein. Demand theory without transitive preference with applications to the theory of competitive equilibrium. In J. Chipman, L. Hurwicz, M. Richter, and H. Sonnenschein, editors, *Preferences, Utility and Demand*. Houghton Mifflin Harcourt, 1971.
- J. von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100(1):295–320, 1928.
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947.

- H. P. Young. An axiomatization of Borda's rule. *Journal of Economic Theory*, 9(1): 43–52, 1974.
- H. P. Young. Extending Condorcet's rule. *Journal of Economic Theory*, 16:335–353, 1977.