

On the Tradeoff between Efficiency and Strategyproofness

Haris Aziz Florian Brandl Felix Brandt Markus Brill
Data61 and UNSW TU München TU München TU Berlin

We study social decision schemes (SDSs), i.e., functions that map a collection of individual preferences over alternatives to a lottery over the alternatives. Depending on how preferences over alternatives are extended to preferences over lotteries, there are varying degrees of efficiency and strategyproofness. In this paper, we consider four such preference extensions: stochastic dominance (*SD*), a strengthening of *SD* based on pairwise comparisons (*PC*), a weakening of *SD* called bilinear dominance (*BD*), and an even weaker extension based on Savage’s sure-thing principle (*ST*). While *random serial dictatorships* are *PC*-strategyproof, they only satisfy *ex post* efficiency. On the other hand, we show that *strict maximal lotteries* are *PC*-efficient and *ST*-strategyproof. We also prove the incompatibility of (i) *PC*-efficiency and *PC*-strategyproofness for anonymous and neutral SDSs, (ii) *ex post* efficiency and *BD*-strategyproofness for pairwise SDSs, and (iii) *ex post* efficiency and *BD*-group-strategyproofness for anonymous and neutral SDSs.

Keywords: strategyproofness, efficiency, stochastic dominance, random serial dictatorship, maximal lotteries

JEL Classifications Codes: D71, D81

1. Introduction

Two fundamental notions in microeconomic theory are *efficiency*—no agent can be made better off without making another one worse off—and *strategyproofness*—no agent can obtain a more preferred outcome by misrepresenting his preferences. The conflict between these two notions is already apparent in Gibbard and Satterthwaite’s seminal theorem, which states that the only single-valued social choice functions (SCFs) that satisfy non-imposition—a weakening of efficiency—and strategyproofness are dictatorships (Gibbard, 1973; Satterthwaite, 1975). In this paper, we study efficiency and strategyproofness in the

context of *social decision schemes (SDSs)*, i.e., functions that map a preference profile to a probability distribution (or lottery) over a fixed set of alternatives (e.g., Gibbard, 1977; Barberà, 1979a). Randomized voting methods have a surprisingly long tradition going back to ancient Greece and have recently gained increased attention in social choice (see, e.g., Bogomolnaia et al., 2005; Chatterji et al., 2014; Brandl et al., 2016a) and political science (see, e.g., Goodwin, 2005; Stone, 2011; Guerrero, 2014). Randomization is particularly natural in subdomains of social choice that are concerned with the assignment of objects to agents such as house allocation (see, e.g., Abdulkadiroğlu and Sönmez, 1998; Bogomolnaia and Moulin, 2004; Che and Kojima, 2010; Budish et al., 2013). Positive results, such as our results on maximal lotteries, are inherited from the general social choice domain to these subdomains.

In order to identify efficient lotteries and argue about incentives with respect to lotteries, one needs to know the agents’ preferences over lotteries. There are various problems associated with asking the agents to submit complete preference relations over all lotteries. For example, the preferences may not allow for a concise representation and agents may not even be aware of these preferences in the first place.¹ We therefore follow the common approach which only assumes that ordinal preferences over alternatives are available. These preferences are then systematically extended to possibly incomplete preferences over lotteries. We will refer to such extensions as *lottery extensions* (see also Cho, 2016). One of the most studied lottery extensions is *stochastic dominance (SD)*, which states that one lottery is preferred to another iff the former first-order stochastically dominates the latter. This extension is of particular importance because it coincides with the extension in which one lottery is preferred to another iff, for any von Neumann-Morgenstern (vNM) utility function consistent with the ordinal preferences, the former yields at least as much expected utility as the latter. Settings in which the existence of an underlying vNM utility function cannot be assumed may call for other lottery extensions. A natural candidate is the *pair-wise comparison (PC)* extension, which arises as a special case of skew-symmetric bilinear (SSB) utility functions, a generalization of vNM utility functions proposed by Fishburn (1982). According to this extension, one lottery is preferred to another iff it is more likely that the former yields a better alternative than the latter. Clearly, each of these lottery extensions gives to rise to different variants or degrees of efficiency and strategyproofness.

Since many lottery extensions are incomplete, i.e., some pairs of lotteries are incomparable, there are two fundamentally different ways how to define strategyproofness. The

¹We surmise that, even if humans *think* they can competently assign von Neumann-Morgenstern utilities to alternatives, these assignments are prone to be based on arbitrary choices because of missing information and the inability to fully grasp the consequences of these choices. A similar sentiment is expressed by Abramowitz and Anshelevich (2018): “Human beings are terrible at expressing their feelings quantitatively. For example, when forming collaborations people may be able to order their peers from ‘best to collaborate with’ to worst, but would have a difficult time assigning exact numeric values to the acuteness of these preferences. In other words, even when numerical (possibly latent) utilities exist, in many settings it is much more reasonable to assume that we only know ordinal preferences: every agent specifies the order of their preferences over the alternatives, instead of a numerical value for each alternative.”

strong notion, first advocated by Gibbard (1977), requires that every misreported preference relation of an agent will result in a lottery that is comparable and weakly less preferred by that agent to the original lottery. According to the weaker notion, first used by Postlewaite and Schmeidler (1986) and then popularized by Bogomolnaia and Moulin (2001), no agent can misreport his preferences to obtain another lottery that is strictly preferred to the original one. In other words, the strong version always interprets incomparabilities in the worst possible manner (such that they violate strategyproofness) while the weak version interprets them as actual incomparabilities that cannot be resolved.² Usually, the strong notion is much more demanding than the weak one. Whenever a lottery extension is *complete*, however, both notions coincide.

Perhaps the most well-known SDS, which is only defined for linear preferences, is *random dictatorship* (*RD*). In *RD*, one of the agents is chosen uniformly at random and this agent's most preferred alternative is selected. Gibbard (1977) has shown that *RD* is the only strongly *SD*-strategyproof SDS that never puts positive probability on Pareto-dominated alternatives. The latter property is known as *ex post efficiency*.³ It is easily verified that *RD* even satisfies the stronger condition of *SD*-efficiency.

Gibbard's proof requires the universal domain of linear preferences and cannot be extended to arbitrary subdomains (see, e.g., Chatterji et al., 2014). In many important subdomains of social choice such as house allocation, matching, and coalition formation, ties are unavoidable because agents are indifferent among all outcomes in which their allocation, match, or coalition is the same (see Section 2). In the presence of ties, *RD* is typically extended to *random serial dictatorship* (*RSD*), where dictators are invoked sequentially and ties between most-preferred alternatives are broken by subsequent dictators. While *RSD* still satisfies strong *SD*-strategyproofness, it violates *SD*-efficiency. This was first observed by Bogomolnaia and Moulin (2001) in the restricted domain of house allocation. The example by Bogomolnaia and Moulin (2001) can be translated to a preference profile with 24 alternatives in the general social choice domain. We give a minimal example for four agents and four alternatives and completely characterize in which settings *RSD* satisfies *SD*-efficiency. Recently, Brandl et al. (2018a) have shown a sweeping impossibility: no anonymous and neutral SDS simultaneously satisfies *SD*-efficiency and *SD*-strategyproofness whenever there are at least four alternatives and four agents. This result was obtained with the help of computers and is tedious to verify. We give manual and simpler proofs for two related statements: there is no anonymous, neutral, *PC*-efficient, and *PC*-strategyproof SDS and there is no *ex post* efficient and *BD*-strategyproof pairwise SDS, where *BD*-strategyproofness is a weakening of *SD*-strategyproofness. While the first result uses stronger conditions than the one by Brandl et al. (2018a), it requires less

²The weak notion of strategyproofness has often been considered in the context of set-valued social choice where preferences over alternatives are extended to incomplete preference relations over sets of alternatives (see, e.g., Gärdenfors, 1976; Barberà, 1977a,b; Kelly, 1977; Feldman, 1979a,b). For relatively weak preference extensions, it allows for rather positive results (Nehring, 2000; Brandt and Brill, 2011; Brandt, 2015).

³Alternative proofs of this theorem were given by Duggan (1996), Nandeibam (1997), and Tanaka (2003).

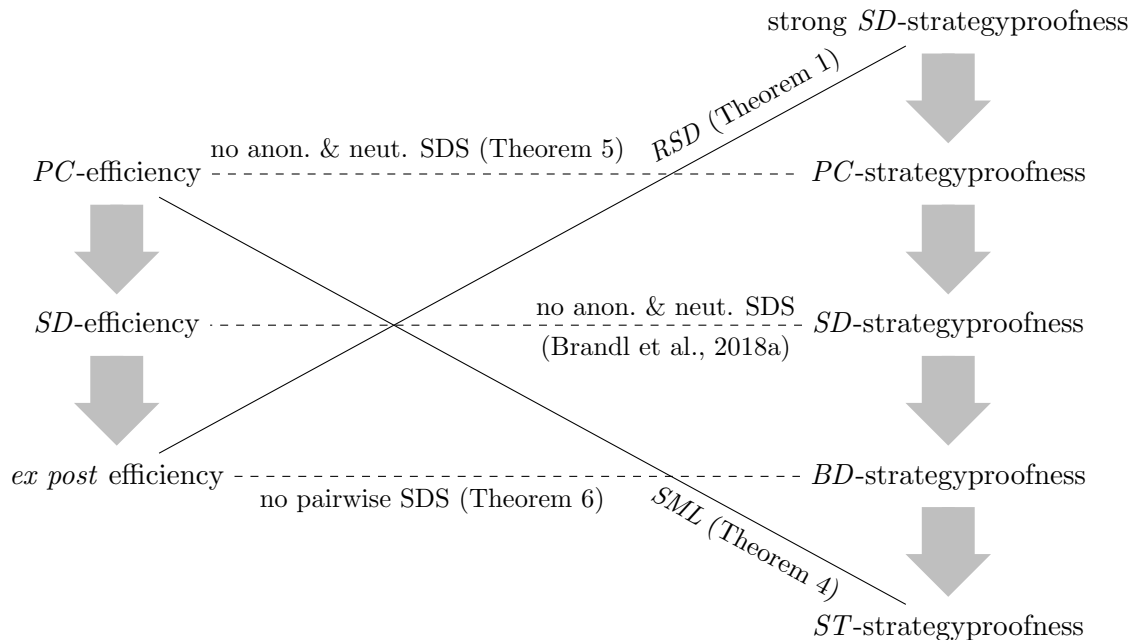


Figure 1: Relationships between varying degrees of efficiency and strategyproofness. An arrow from one notion of efficiency or strategyproofness to another denotes that the former implies the latter. Solid lines represent positive results whereas dashed lines represent impossibility results. The result by Brandl et al. (2018a) was shown using computer-aided solving techniques.

agents and alternatives. The second result uses weaker conditions, but only holds for the restricted class of pairwise SDSs.

In order to obtain positive results we then introduce a new lottery extension that is weaker than stochastic dominance and is based on Savage’s sure-thing principle (ST). All three lottery extensions are then used to demonstrate an interesting tradeoff (see Figure 1): Random serial dictatorship is strongly SD -strategyproof, but only satisfies $ex post$ efficiency. On the other hand, *strict maximal lotteries* (SML) as defined by Kreweras (1965) and Fishburn (1984a), satisfy PC -efficiency and ST -strategyproofness. Strict maximal lotteries correspond to the quasi-strict mixed equilibria of the symmetric zero-sum game induced by the pairwise majority margins. While ST -strategyproofness is quite weak, it is important to note that most common $ex post$ efficient SDSs (except RSD) violate much weaker notions of strategyproofness. Moreover, SML satisfies a number of other desirable properties violated by RSD such as Condorcet-consistency and composition-consistency (Laslier, 2000; Brandl et al., 2016a). Figure 1 summarizes our findings.

We also consider manipulation by groups of agents. We prove that both RSD and SML




Figure 2: Relationships between efficiency and group-strategyproofness concepts. Solid lines represent positive results whereas dashed lines represent impossibility results.

satisfy ST -group-strategyproofness and that no anonymous, neutral, and $ex post$ efficient SDS satisfies the slightly stronger notion of BD -group-strategyproofness. These results are visualized in Figure 2.

All of our impossibility results assume anonymity. *Serial dictatorship*, an SDS that clearly violates anonymity, is defined for a fixed sequence of the agents and lets each agent narrow down the set of alternatives by picking his most preferred of the alternatives selected by the previous agents. Serial dictatorship trivially satisfies all reasonable notions of efficiency and strategyproofness. Since lotteries can guarantee $ex ante$ fairness via randomization, anonymity and neutrality are typically two minimal conditions that fair SDSs are expected to satisfy.

2. Related Work

Starting with the Gibbard-Satterthwaite impossibility (Gibbard, 1973; Satterthwaite, 1975), there is remarkable number of results that reveal a tradeoff between efficiency and strategyproofness.

As mentioned above, RD , as proposed by Gibbard (1977), satisfies SD -efficiency and (strong) SD -strategyproofness when preferences are linear. Brandl et al. (2016b) have shown that RD cannot be extended to the full domain of weak preferences without violating at least one of these properties or anonymity or neutrality.⁴ This theorem has been strengthened by Brandl et al. (2018a), who showed that no anonymous and neutral SDS satisfies SD -efficiency and SD -strategyproofness by leveraging computer-aided solving techniques.⁵

⁴A natural candidate is RSD , which—as we discuss in Section 4—violates SD -efficiency. Aziz (2013) proposes another SDS that satisfies a stronger notion of efficiency and a weaker notion of strategyproofness than RSD . However, it also violates SD -efficiency.

⁵The theorem by Brandl et al. (2018a) also implies analogous impossibilities for the upward lexicographic UL and downward lexicographic DL extensions introduced by Cho (2016). The impossibility for UL even holds for linear preferences while this is not the case for DL since RD satisfies both DL -efficiency and DL -strategyproofness (Brandl, 2013).

When preferences are *dichotomous*, efficiency and strategyproofness are compatible. The utilitarian mechanism described by Bogomolnaia et al. (2005) satisfies the strongest degrees of efficiency and strategyproofness considered in this paper (*PC*-efficiency and strong *SD*-strategyproofness). Interestingly, this mechanism always returns a maximal lottery (which happens to be of a particularly simple form for dichotomous preferences). When replacing strategyproofness with group-strategyproofness and weakening *SD*-efficiency to *ex post* efficiency, this possibility turns into an impossibility (Bogomolnaia et al., 2005). We strengthen this result in Section 7.

A subdomain of social choice that has been thoroughly studied in the literature is the *assignment* (aka house allocation or two-sided matching with one-sided preferences) domain. An assignment profile can be associated with a social choice profile by letting the set of alternatives be the set of deterministic allocations and postulating that agents are indifferent among all allocations in which they receive the same object (see, e.g., Aziz et al., 2013).⁶ Thus, impossibility results for the assignment setting imply impossibility results for the social choice setting.

Bogomolnaia and Moulin (2001) have shown that no random assignment rule satisfies *SD*-efficiency, strong *SD*-strategyproofness, and equal treatment of equals. The result by Bogomolnaia and Moulin even holds when preferences over objects are linear and identical up to a single object (Chang and Chun, 2017). In a related paper, Katta and Sethuraman (2006) proved that no assignment rule satisfies *SD*-efficiency, *SD*-strategyproofness, and strong *SD*-envy-freeness for the full domain of weak preferences over objects.⁷ Nesterov (2017) showed similar impossibilities: *ex post* efficiency, strong *SD*-strategyproofness, and strong *SD*-envy-freeness as well as *SD*-efficiency, strong *SD*-strategyproofness, and weak *SD*-envy-freeness are incompatible with each other. Bogomolnaia and Moulin (2001) have introduced the probabilistic serial (PS) assignment rule which satisfies *SD*-efficiency and *SD*-strategyproofness when preferences over objects are linear. Natural extensions of the probabilistic serial rule to the full assignment domain and to the general social choice domain fail to satisfy *SD*-strategyproofness (Katta and Sethuraman, 2006; Aziz and Stursberg, 2014). Thus, it is an interesting open question whether there is any *SD*-efficient and *SD*-strategyproof assignment rule on the full assignment domain that satisfies equal treatment of equals. We conjecture that no such rule exists.

There is an extensive literature showing impossibility results for *set-valued* SCFs. Just like in probabilistic social choice, strategyproofness is defined by lifting the preference relation, in this case to sets of alternatives. It turns out that some of the resulting notions of strategyproofness are logically related to notions considered in this paper. For example, every strongly *SD*-strategyproof SDS induces a set-valued SCF (by just taking the support of the resulting lottery) that is strategyproof with respect to the optimist and pessimist extensions as used by Duggan and Schwartz (2000), Rodríguez-Álvarez (2007), Rodríguez-

⁶Note that this transformation turns assignment profiles with linear preferences over k objects into social choice profiles with non-linear preferences over $k!$ allocations.

⁷Strong envy-freeness is a fairness property that is stronger than equal treatment of equals as used by Bogomolnaia and Moulin (2001). Weak envy-freeness and equal treatment of equals are incomparable.

Álvarez (2009), Sato (2008), and others. Similarly, every *ST*-strategyproof SDS induces a set-valued SCF that is strategyproof with respect to the simple extension where one set is preferred to another iff all alternatives in the former are strictly preferred to all alternatives in the latter (see, e.g., Nehring, 2000; Brandt, 2015, Remark 6). The relationship between *BD*-strategyproofness and strategyproofness with respect to Fishburn’s set extension as used by Feldman (1979a), Ching and Zhou (2002), Sanver and Zwicker (2012), Brandt and Geist (2016), and others works the other way round: Every Fishburn-strategyproof SCF induces a *BD*-strategyproof SDS by taking the uniform lottery over the resulting choice set. As a consequence, Theorem 6 implies Theorem 3 by Brandt and Geist (2016). However, the proof of Theorem 6 uses weak preferences while the result by Brandt and Geist (2016) even holds for linear preferences.

Mennle and Seuken (2015) proposed a different approach to trade off efficiency and strategyproofness by quantifying manipulation losses and considering convex combinations of random assignment rules. There also has been some recent work on the tradeoff between participation (resistance against strategic abstention) and efficiency (Brandl et al., 2015a,b).

Preliminary proofs of results presented in this paper have appeared in conference proceedings. Please see the Acknowledgments for details.

3. Preliminaries

Let $N = \{1, \dots, n\}$ be a set of agents who entertain ordinal preferences over a finite set A of m alternatives. Every agent $i \in N$ is equipped with a complete and transitive *preference relation* $R_i \subseteq A \times A$. The set of all preference relations will be denoted by \mathcal{R} . In accordance with conventional notation, we write P_i for the strict part of R_i , i.e., $a P_i b$ if $a R_i b$ but not $b R_i a$ and I_i for the indifference part of R_i , i.e., $a I_i b$ if $a R_i b$ and $b R_i a$. Preference relations are straightforwardly extended to sets of alternatives X, Y where $X R_i Y$ denotes that $x R_i y$ holds for all $x \in X$ and $y \in Y$. Similarly, $X P_i Y$ iff $x P_i y$ for all $x \in X$ and $y \in Y$. We will compactly represent a preference relation as a comma-separated list where all alternatives among which an agent is indifferent are represented by a set. For example $a P_i b I_i c$ is written as $R_i: a, \{b, c\}$. A preference relation R_i is *linear* if $x P_i y$ or $y P_i x$ for all distinct alternatives $x, y \in A$. A preference relation R_i is *dichotomous* if $x R_i y R_i z$ implies $x I_i y$ or $y I_i z$. A *preference profile* $R = (R_1, \dots, R_n)$ is an n -tuple containing a preference relation R_i for every agent $i \in N$. The set of all preference profiles is thus given by \mathcal{R}^n . By R_{-i} we denote the preference profile obtained from R by removing the preference relation of agent i , i.e., $R_{-i} = R \setminus \{(i, R_i)\}$.

The set of all *lotteries* (or *probability distributions*) over A is denoted by $\Delta(A)$. We will write lotteries as convex combinations of alternatives, e.g., $1/2 a + 1/2 b$ denotes the uniform distribution over $\{a, b\}$. For a given lottery p and alternative x , $p(x)$ denotes the probability that p assigns to x . The support of a lottery $p \in \Delta(A)$, denoted by \hat{p} , is the set of all alternatives to which p assigns positive probability, i.e., $\hat{p} = \{x \in A: p(x) > 0\}$.

A lottery p is *degenerate* if $|\widehat{p}| = 1$.

Our central object of study are social decision schemes, i.e., functions that map a preference profile to a lottery. Thus, a *social decision scheme (SDS)* is a function

$$f: \mathcal{R}^n \rightarrow \Delta(A).$$

A minimal fairness condition for SDSs is *anonymity*, which requires that $f(R) = f(R \circ \pi)$ for all $R \in \mathcal{R}^n$ and all permutations $\pi: N \rightarrow N$. Another fairness requirement is *neutrality*. For a permutation π of A and a preference relation R_i , we define R_i^π as the preference relation where alternatives are renamed according to π , i.e., $\pi(x) R_i^\pi \pi(y)$ iff $x R_i y$. An SDS f is *neutral* if, for all $R \in \mathcal{R}^n$, $x \in A$, and all permutations $\pi: A \rightarrow A$, $f(R)(x) = f(R^\pi)(\pi(x))$.

3.1. Lottery Extensions

In order to reason about the outcomes of SDSs, we need to make assumptions on how agents compare lotteries given their preferences over alternatives. A *lottery extension* maps a preference relation to a (possibly incomplete) preference relation over lotteries. We will now define the lottery extensions considered in this paper. Throughout this section, let $R_i \in \mathcal{R}$ and $p, q \in \Delta(A)$. For the examples we assume that the underlying preference relation is $R_i: a, b, c$.

A very simple and crude lottery extension prescribes that p is preferred to q iff every alternative in the support of p is preferred to every alternative in the support of q , i.e., $\widehat{p} P_i \widehat{q}$. This extension only allows the comparison of lotteries with disjoint supports. We slightly generalize this definition by requiring that p and q assign the same probability to all alternatives that are contained in both supports and $(\widehat{p} \setminus \widehat{q}) P_i (\widehat{p} \cap \widehat{q}) P_i (\widehat{q} \setminus \widehat{p})$. Following Savage's sure-thing principle, the resulting lottery extension will be referred to as the *sure-thing (ST)* extension. Formally,

$$p R_i^{ST} q \quad \text{iff} \quad (\widehat{p} \setminus \widehat{q}) P_i (\widehat{p} \cap \widehat{q}) P_i (\widehat{q} \setminus \widehat{p}) \wedge \forall x \in \widehat{p} \cap \widehat{q}: p(x) = q(x). \quad (ST)$$

The idea underlying *ST* is that the comparison of two lotteries should be independent of the part in which they coincide. This is related to von Neumann and Morgenstern's independence axiom (von Neumann and Morgenstern, 1947) and has also been used for defining preference extensions from alternatives to sets of alternatives (Fishburn, 1972; Gärdenfors, 1979). For example, $1/2 a + 1/2 b P_i^{ST} 1/2 b + 1/2 c$.⁸

The second extension we consider, called *bilinear dominance (BD)*, requires that for every pair of alternatives the probability that p yields the more preferred alternative and q the less preferred alternative is at least as large as the other way round. Formally,

$$p R_i^{BD} q \quad \text{iff} \quad \forall x, y \in A: (x P_i y \Rightarrow p(x)q(y) \geq p(y)q(x)). \quad (BD)$$

⁸A refinement of the *ST* extension can be defined by demanding that $(\widehat{p} \setminus \widehat{q}) R_i (\widehat{p} \cap \widehat{q}) R_i (\widehat{q} \setminus \widehat{p})$ instead of $(\widehat{p} \setminus \widehat{q}) P_i (\widehat{p} \cap \widehat{q}) P_i (\widehat{q} \setminus \widehat{p})$. However, Theorems 4 and 7 do not hold under this extension.

Apart from its intuitive appeal, the main motivation for BD is that p bilinearly dominates q iff p is preferable to q for every SSB utility function consistent with R_i (cf. Fishburn, 1984b; Aziz et al., 2015). For example, $1/2 a + 1/2 b P_i^{BD} 1/3 a + 1/3 b + 1/3 c$.

Perhaps the best-known lottery extension is *stochastic dominance* (SD), which prescribes that for each alternative $x \in A$, the probability that p selects an alternative that is at least as good as x is greater or equal to the probability that q selects such an alternative (Hadar and Russell, 1969). Formally,

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

It is well-known that $p R_i^{SD} q$ iff the expected utility for p is at least as large as that for q for every von-Neumann-Morgenstern utility function compatible with R_i . For example, $1/2 a + 1/2 c P_i^{SD} 1/2 b + 1/2 c$.

A novel strengthening of SD is the *pairwise comparison* (PC) extension (Aziz et al., 2015). The reasoning behind PC is that p should be preferred to q if the probability that p yields an alternative preferred to the alternative returned by q is at least as large as the other way round. In other words, p is preferred to q if choosing p results in lower *ex ante* regret. Formally,

$$p R_i^{PC} q \quad \text{iff} \quad \sum_{x R_i y} p(x)q(y) \geq \sum_{x R_i y} q(x)p(y). \quad (PC)$$

For example, $2/3 a + 1/3 c P_i^{PC} b$.

The PC extension can alternatively be defined using skew-symmetric bilinear (SSB) utility functions as defined by Fishburn (1982). Blavatsky (2006) gave a characterization of the PC extension which relies on the axioms that characterize SSB utility functions (cf. Fishburn, 1982, 1988) plus an additional axiom that singles out PC . Moreover, he cites empirical evidence showing that decision makers' preferences adhere to the PC extension (see also Butler et al., 2016). In contrast to the previous three extensions, PC yields *complete* preference relations over lotteries.

The four lottery extensions introduced here form a hierarchy.⁹ For all $R_i \in \mathcal{R}$,

$$R_i^{ST} \subseteq R_i^{BD} \subseteq R_i^{SD} \subseteq R_i^{PC}.$$

3.2. Efficiency and Strategyproofness

Arguably one of the most fundamental axioms in microeconomic theory is Pareto-efficiency. An alternative *Pareto-dominates* another alternative if every agent weakly prefers the former to the latter and at least one agent strictly prefers the former to the latter. *Pareto-efficiency* prescribes that Pareto-dominated alternatives are not chosen. There are various

⁹ To see that $R_i^{ST} \subseteq R_i^{BD}$, let $p, q \in \Delta(A)$ such that $p R_i^{ST} q$ and $x, y \in A$ with $x P_i y$. We need to show that $p(x)q(y) \geq p(y)q(x)$. If $p(y) = 0$ or $q(x) = 0$ this is trivial, since the right hand side is zero. So we consider the case where $y \in \hat{p}$ and $x \in \hat{q}$. Because $p R_i^{ST} q$ and $x P_i y$ we have that $x \in \hat{p}$ and $y \in \hat{q}$. This implies that $p(x) = q(x)$ and $p(y) = q(y)$ and in particular $p(x)q(y) = p(y)q(x)$. We refer to Aziz et al. (2015) for the remaining inclusions.

reasonable ways to define Pareto-efficiency in probabilistic social choice. In particular, every lottery extension defines a corresponding notion of Pareto-efficiency.

Definition 1. Let $\mathcal{E} \in \{ST, BD, SD, PC\}$, $R \in \mathcal{R}^n$, and $p, q \in \Delta(A)$. Then, p \mathcal{E} -dominates q if $p R_i^{\mathcal{E}} q$ for all $i \in N$ and $p P_i^{\mathcal{E}} q$ for some $i \in N$. An SDS f is \mathcal{E} -efficient if, for every $R \in \mathcal{R}^n$, there does not exist a lottery that \mathcal{E} -dominates $f(R)$.

Since the lottery extensions we consider form a hierarchy, PC -efficiency implies SD -efficiency which in turn implies BD -efficiency which in turn implies ST -efficiency.

A standard efficiency notion that cannot be formalized using lottery extensions is *ex post* efficiency. An SDS is *ex post efficient* if, for every preference profile, it assigns probability zero to all Pareto-dominated alternatives. It can be shown that SD -efficiency implies *ex post* efficiency and *ex post* efficiency implies BD -efficiency (cf. Aziz et al., 2015).

Efficiency essentially requires that outcomes are *socially* optimal. This can be contrasted with strategyproofness, which is concerned with the *individual* behavior of agents. Strategyproofness prescribes that no agent can obtain a more preferred outcome by misrepresenting his preferences. Again, we obtain varying degrees of this property depending on the underlying lottery extension.

Definition 2. Let $\mathcal{E} \in \{ST, BD, SD, PC\}$. An SDS f is \mathcal{E} -manipulable if there are $R, R' \in \mathcal{R}^n$ and $i \in N$ with $R_j = R'_j$ for all $j \neq i$ such that $f(R') P_i^{\mathcal{E}} f(R)$. An SDS is \mathcal{E} -strategyproof if it is not \mathcal{E} -manipulable.

PC -strategyproofness implies SD -strategyproofness which in turn implies BD -strategyproofness which in turn implies ST -strategyproofness (see Figure 1). Note that our definition of strategyproofness does *not* require that $f(R) R_i^{\mathcal{E}} f(R')$ for all R' with $R'_j = R_j$ for all $j \neq i$. We refer to this stronger notion as *strong strategyproofness*, but only use it in the context of the SD extension. For coarser extensions, in which most lotteries are incomparable, it seems unduly restrictive. The weaker notion employed here is for example also used by Postlewaite and Schmeidler (1986) and Bogomolnaia and Moulin (2001) for the SD extension.

An SDS is strongly SD -strategyproof if no agent is better off by manipulating his preferences for *some* expected utility representation of his ordinal preferences. This condition is quite demanding because an SDS may be deemed manipulable just because it can be manipulated for a contrived and highly unlikely utility representation. An SDS is weakly SD -strategyproof if no agent is better off by manipulating his preferences for *all* expected utility representations of his preferences.

Note that due to the completeness of the PC extension, strong PC -strategyproofness and PC -strategyproofness coincide. Moreover, strong SD -strategyproofness is stronger than PC -strategyproofness while (weak) SD -strategyproofness is weaker.

4. Random Serial Dictatorship

In this section, we examine random serial dictatorship (*RSD*)—an extension of random dictatorship to the case where agents may express indifference among alternatives. *RSD* is commonly used in house allocation, matching, and coalition formation domains, where ties between outcomes naturally arise because agents are assumed to be indifferent between outcomes in which their individual allocation is the same. (see, e.g., Abdulkadiroğlu and Sönmez, 1998; Bogomolnaia and Moulin, 2004; Che and Kojima, 2010; Budish et al., 2013). In these contexts, *RSD* is sometimes also referred to as the *random priority* mechanism. *RSD* is defined by picking a sequence of the agents uniformly at random and then invoking serial dictatorship (i.e., each agent narrows down the set of alternatives by picking his most preferred of the alternatives selected by the previous agents).

For a formal definition of *RSD*, let $\max_{R_i}(A)$ denote the set of maximal alternatives according to R_i , i.e., $\max_{R_i}(A) = \{x \in A : x R_i y \text{ for all } y \in A\}$. For given $R \in \mathcal{R}$, $A' \subseteq A$, and $N' \subseteq N$, *RSD* is then recursively defined as follows.¹⁰

$$RSD(R, A', N') = \begin{cases} \Delta(A') & \text{if } N' = \emptyset, \\ \frac{1}{|N'|} \sum_{i \in N'} RSD(R, \max_{R_i}(A'), N' \setminus \{i\}) & \text{otherwise.} \end{cases}$$

Then, $RSD(R) = RSD(R, A, N)$.

Clearly, the set $RSD(R)$ can only contain more than one lottery if there are two alternatives among which all agents are indifferent. Otherwise, $RSD(R)$ consists of a single lottery. An SDS is called an *RSD* scheme if it always selects a lottery from the set $RSD(R)$ that furthermore only depends on $RSD(R)$.

Definition 3. An SDS f is an *RSD* scheme if for every $R \in \mathcal{R}^n$, $f(R) \in RSD(R)$, and for all $R, R' \in \mathcal{R}^n$, $RSD(R) = RSD(R')$ implies $f(R) = f(R')$.

For a given notion of efficiency or strategyproofness, we write that *RSD* satisfies the notion if every *RSD* scheme satisfies it. Similarly, we say that *RSD* violates the notion if every *RSD* scheme violates it.

It is well-known that, if the outcome is determined by *RSD*, truth telling is a weakly dominant strategy for every agent when lotteries are compared according to the *SD* extension (see, e.g., Bogomolnaia et al., 2005). Moreover, *RSD* is *ex post* efficient.

Theorem 1. *RSD* is *ex post* efficient and strongly *SD*-strategyproof.

The proofs of all theorems are deferred to the appendix.

It has been shown by Bogomolnaia and Moulin (2001) that *RSD* violates *SD*-efficiency within the domain of house allocation. The example by Bogomolnaia and Moulin can be translated to a preference profile on 24 alternatives in the general social choice setting. We

¹⁰Here, the sum of two sets A and B is defined as the Minkowski sum, i.e., $A + B = \{x + y : x \in A, y \in B\}$.

give an independent example with four alternatives and four agents and completely characterize in which settings RSD satisfies SD -efficiency.¹¹ Consider the following preference profile.

$$R_1 : \{a, c\}, b, d$$

$$R_2 : \{a, d\}, b, c$$

$$R_3 : \{b, c\}, a, d$$

$$R_4 : \{b, d\}, a, c$$

The unique RSD lottery is $p = 1/3 a + 1/3 b + 1/6 c + 1/6 d$, which is SD -dominated by $1/2 a + 1/2 b$. In fact, it is even the case that *all* agents *strictly* prefer the latter lottery according to SD . In other words, there exists a lottery which gives *strictly more* expected utility than p to *each* agent and for *every* utility representation consistent with the ordinal preferences of the agents. Therefore RSD is not SD -efficient for $n = 4$ and $m = 4$.

It turns out that RSD is only SD -efficient in settings where SD -efficiency is equivalent to the weaker notion of *ex post* efficiency.

Theorem 2. *RSD is SD -efficient iff $n \leq 2$, or $m \leq 3$, or $n = 3$ and $m \leq 5$.*

The stronger notion of PC -efficiency is violated by RSD even when preferences are linear. In other words, RD already fails to satisfy PC -efficiency. To see this, consider the following preference profile with three agents and three alternatives.

$$R_1 : a, b, c$$

$$R_2 : b, c, a$$

$$R_3 : c, b, a$$

Clearly, RD (and hence RSD) returns $1/3 a + 1/3 b + 1/3 c$, which is PC -dominated by the degenerate lottery b .¹² RD does satisfy SD -efficiency, however.

While implementing RSD by uniformly selecting a sequence of agents and then running serial dictatorship is straightforward, it was recently shown that computing the *ex ante* RSD probabilities is $\#P$ -complete and therefore computationally intractable (Aziz et al., 2013). Even the question of whether the RSD probability of a given alternative exceeds some fixed value $\lambda \in (0, 1)$ is NP -complete and hence cannot be answered in polynomial time unless P equals NP .

5. Maximal Lotteries

Maximal lotteries were first considered by Kreweras (1965) and independently proposed and studied in more detail by Fishburn (1984a). Interestingly, maximal lotteries or variants

¹¹Bogomolnaia et al. (2005) provide an example with six agents and five alternatives for the special case of dichotomous preferences.

¹²A similar example in which all agents are *strictly* better off can easily be constructed with four alternatives and four agents.

thereof have been rediscovered again by economists (Laffond et al., 1993), mathematicians (Fisher and Ryan, 1995), political scientists (Felsenthal and Machover, 1992), and computer scientists (Rivest and Shen, 2010). An axiomatic characterization of maximal lotteries was recently given by Brandl et al. (2016a).

In order to define maximal lotteries, we need some notation. For a preference profile $R \in \mathcal{R}^n$ and two alternatives $x, y \in A$, the *majority margin* $g_R(x, y)$ is defined as the difference between the number of agents who prefer x to y and the number of agents who prefer y to x , i.e.,

$$g_R(x, y) = |\{i \in N : x R_i y\}| - |\{i \in N : y R_i x\}|.$$

Thus, $g_R(y, x) = -g_R(x, y)$ for all $x, y \in A$. A *maximal alternative*, aka (*weak*) *Condorcet winner*, is an alternative $x \in A$ with $g_R(x, y) \geq 0$ for all alternatives $y \in A$. It is well known that maximal alternatives may fail to exist. This drawback can however be remedied by considering lotteries over alternatives. The function g_R can be naturally extended to pairs of lotteries by considering its bilinear form, which corresponds to *expected* majority margins. For $p, q \in \Delta(A)$, let

$$g_R(p, q) = \sum_{x, y \in A} p(x)q(y)g_R(x, y).$$

The set of maximal lotteries is then defined as

$$ML(R) = \{p \in \Delta(A) : g_R(p, q) \geq 0 \text{ for all } q \in \Delta(A)\}.$$
¹³

The Minimax Theorem (von Neumann, 1928) implies that $ML(R)$ is non-empty for all $R \in \mathcal{R}^n$. In fact, g_R can be interpreted as the payoff matrix of a symmetric zero-sum game and maximal lotteries as the mixed maximin strategies (or Nash equilibria) of this game. As a consequence, elements of $ML(R)$ can be found in polynomial time using linear programming. Interestingly, $ML(R)$ is a singleton in most cases (see Brandl et al., 2016a). This holds, in particular, if all agents have linear preferences and the number of agents is odd (Laffond et al., 1997; Le Breton, 2005).

We first show that no maximal lottery is *PC*-dominated.


Theorem 3. *Every SDS that returns maximal lotteries is PC-efficient.*

This result contrasts with our earlier observation that *RSD* fails to be *SD*-efficient and hence *PC*-efficient.

While *ML* satisfies a very high degree of efficiency, it does not do as well in terms of strategyproofness. In fact, *ML* already violates *BD*-strategyproofness. To see this, let $A = \{a, b, c\}$ and consider the preference profiles given below. The set of maximal lotteries only depends on g_R , which can be nicely represented as a weighted majority graph: For

¹³Since *ML* only depends on g_R , it does neither require transitivity nor completeness of individual preferences.

every pair of alternatives x and y with $g_R(x, y) > 0$, there is an edge from x to y labeled with $g_R(x, y)$.



It can be verified that $ML(R) = \{1/3 a + 1/3 b + 1/3 c\}$. However, if agent 1 misrepresents his preferences between b and c by reporting $R'_1: a, b, c$, the outcome for the new preference profile R' is $ML(R') = \{3/5 a + 1/5 b + 1/5 c\}$. Thus, $f(R') P_1^{BD} f(R)$ for any SDS f that returns maximal lotteries.

More severe violations of strategyproofness can be constructed by using preference profiles that admit more than one maximal lottery and by breaking ties in an unfavorable way. In order to illustrate this, consider the following preference profiles.



Here, $ML(R) = \Delta(\{a, c\})$ and $ML(R') = \Delta(\{a, b, c\})$. Hence, when letting $f(R) = \{c\}$ and $f(R') = \{a\}$, f is manipulable for *any* reasonable lottery extension (including *ST*). In order to avoid this, we define a useful and natural subclass of $ML(R)$ called *strict* maximal lotteries $SML(R)$. $SML(R)$ corresponds to the set of *quasi-strict* Nash equilibria of g_R , i.e., all equilibria p in which every action in the support of p yields strictly more payoff than every action outside of the support of p .¹⁴ In zero-sum games, quasi-strict equilibria constitute a subset of equilibria with maximal support (see, e.g., Dutta and Laslier, 1999; Brandt and Fischer, 2008).

Definition 4. Let $R \in \mathcal{R}^n$ and $p \in \Delta(A)$. Then $p \in SML(R)$ if, for all $x \in A$,

$$\begin{aligned} p(x) > 0 & \text{ iff } g_R(x, p) = 0, \text{ and} \\ p(x) = 0 & \text{ iff } g_R(x, p) < 0. \end{aligned}$$

An SDS is called an *SML scheme* if it always selects a lottery from the set $SML(R)$ and furthermore only depends on $SML(R)$.¹⁵

Definition 5. An SDS f is an *SML scheme* if for every $R \in \mathcal{R}^n$, $f(R) \in SML(R)$, and for all $R, R' \in \mathcal{R}^n$, $SML(R) = SML(R')$ implies $f(R) = f(R')$.

¹⁴Geometrically, $SML(R)$ is the relative interior of $ML(R)$.

¹⁵The second assumption is not critical because, as mentioned above, $SML(R)$ is almost always a singleton.

For a given notion of efficiency or strategyproofness, we write that *SML* satisfies the notion if every *SML* scheme satisfies it. Similarly, we say that *SML* violates the notion if every *SML* scheme violates it. Note that every *SML* scheme only depends on g_R , since $SML(R)$ only depends on g_R .

Theorem 3 obviously implies that *SML* satisfies *PC*-efficiency. It turns out that *SML* is also *ST*-strategyproof.

Theorem 4. *SML is PC-efficient and ST-strategyproof.*

The proof of Theorem 4 provides interesting insights into *SML*. It seems as if this is the highest degree of strategyproofness one can hope for when also insisting on *PC*-efficiency.¹⁶ While *ST*-strategyproofness does allow manipulators to skew the resulting distribution, crude manipulative attacks such as distorting the outcome from one degenerate lottery to another—an attack that many common SDSs suffer from (see Section 6)—or from one support to another disjoint one are futile. Also note that Theorem 4 holds for *every SML* scheme, i.e., *ST*-strategyproofness holds irrespectively of how ties between strict maximal lotteries are broken.

Brandt et al. (2018b) have shown that *ML* is *PC*-strategyproof for all preference profiles that admit a Condorcet winner.¹⁷

6. Negative Results

As shown in Section 4, *RSD* violates *PC*-efficiency, even when preferences are assumed to be linear. In fact, no anonymous, *PC*-efficient, and *PC*-strategyproof SDS is known under this assumption. Randomizing over the winning sets of various commonly used SCFs such as *Borda's rule*, *Copeland's rule*, or *Hare's rule* (aka instant runoff) fails to be *PC*-strategyproof because all these rules can be manipulated with respect to *any* lottery extension (Taylor, 2005, pp. 44–51). Known *PC*-strategyproof SDSs that are assigning probabilities to alternatives in proportion to their Borda or Copeland scores (see, e.g., Barberà, 1979b), on the other hand, trivially fail to satisfy *ex post* efficiency (and therefore also *PC*-efficiency). Still, *PC*-efficiency is not unduly restrictive as it is satisfied by *ML*.

We now show that *PC*-efficiency and *PC*-strategyproofness are indeed incompatible with each other.

Theorem 5. *There is no anonymous, neutral, PC-efficient, and PC-strategyproof SDS for $n \geq 3$ and $m \geq 3$.*

¹⁶For example, *SML* does not satisfy strategyproofness with respect to the strengthening of *ST* given in Footnote 8. When preferences are strict, *SML* satisfies strategyproofness with respect to the Kelly extension (p is preferred to q iff $\hat{p} R_i \hat{q}$) (Brandt, 2015). This extension is incomparable to the *ST* extension, but only allows the comparison of lotteries whose supports overlap in at most one alternative.

¹⁷This is based on an earlier observation by Peyre (2013) and Hoang (2017), who showed a similar statement for a variant of *ML* that only uses the sign of majority margins.

We do not know whether neutrality is required for the impossibility. Apart from neutrality, all axioms are necessary: serial dictatorship satisfies all axioms except anonymity, *RSD* satisfies all axioms except *PC*-efficiency, and *SML* satisfies all axioms except *PC*-strategyproofness. We conjecture that Theorem 5 even holds for linear preferences and even when giving up neutrality.

Brandl et al. (2018a) have leveraged computer-aided solving techniques to prove a similar statement for *SD*-efficiency and *SD*-strategyproofness.¹⁸ While their theorem uses considerably weaker notions of efficiency and strategyproofness, their computer-generated proof is extremely tedious to check. We give a more accessible manual proof of Theorem 5 in the Appendix.

An important subclass of SDSs consists of pairwise SDSs. An SDS f is *pairwise* (or a neutral C2 function) if it is neutral and $f(R) = f(R')$ for all $R, R' \in \mathcal{R}^n$ such that for all $x, y \in A$,

$$|\{i \in N : x R_i y\}| - |\{i \in N : y R_i x\}| = |\{i \in N : x R'_i y\}| - |\{i \in N : y R'_i x\}|.$$

In other words, the outcome of a pairwise SDS only depends on the anonymized comparisons between pairs of alternatives (see, e.g., Young, 1974; Fishburn, 1977; Zwicker, 1991). Hence, pairwise-ness is stronger than both anonymity and neutrality. Many SCFs such Borda’s rule, Copeland’s rule, and Kemeny’s rule are pairwise. Moreover, *ML* is pairwise.

The following theorem shows that the conditions in Theorem 5 (and also those in the theorem by Brandl et al., 2018a) can be significantly weakened when restricting attention to pairwise SDSs.

Theorem 6. *There is no pairwise, ex post efficient, and BD-strategyproof SDS for $n \geq 4$ and $m \geq 4$.*

The three axioms are logically independent from each other: *RSD* satisfies all axioms except pairwise-ness, always returning the uniform lottery over all alternatives satisfies all axioms except *ex post* efficiency, and *SML* satisfies all axioms except *BD*-strategyproofness.

Our results, both positive and negative, concerning the tradeoff between efficiency and strategyproofness are summarized in Figure 1.

7. Group-strategyproofness

A strengthening of strategyproofness that is often considered is *group-strategyproofness*. It requires that no group of agents should be able to jointly benefit from misrepresenting their preferences.

Definition 6. Let $\mathcal{E} \in \{ST, BD, SD, PC\}$. An SDS f is \mathcal{E} -*group-manipulable* if there are $R, R' \in \mathcal{R}^n$ and $S \subseteq N$ with $R_j = R'_j$ for all $j \notin S$ and $f(R') P_i^{\mathcal{E}} f(R)$ for all $i \in S$. An SDS is \mathcal{E} -*group-strategyproof* if it is not \mathcal{E} -group-manipulable.

¹⁸Strictly speaking, Theorem 5 is not implied by the theorem of Brandl et al. (2018a) because their theorem requires at least four agents and four alternatives.

Clearly, *SML* violates *BD*-group-strategyproofness because it already violates *BD*-strategyproofness. It may be more surprising that *RSD* also violates *BD*-group-strategyproofness. This can be seen by letting $A = \{a, b, c\}$ and considering the following two preference profiles.

$$\begin{array}{ll} R_1: \{a, b\}, c & R'_1: a, \{b, c\} \\ R_2: \{a, c\}, b & R'_2: a, \{b, c\} \\ R_3: \{b, c\}, a & R'_3: \{b, c\}, a \end{array}$$

$RSD(R) = 1/3 a + 1/3 b + 1/3 c$ and $RSD(R') = 2/3 a + 1/6 b + 1/6 c$. Furthermore, $RSD(R') P_1^{BD} RSD(R)$ and $RSD(R') P_2^{BD} RSD(R)$. Hence, *RSD* is *BD*-group-manipulable by agents 1 and 2.

It can easily be seen that *RSD* is *ST*-group-strategyproof because every *RSD* lottery assigns positive probability to at least one most preferred alternative of every agent. There can be no lottery that an agent prefers to this lottery according to the *ST* extension.

For the case of *SML*, it can be verified that the proof of Theorem 4 straightforwardly carries over to group-strategyproofness. As a consequence, *RSD* and *SML* satisfy the same degree of group-strategyproofness (with respect to the preference extensions considered in this paper): they both satisfy *ST*-group-strategyproofness and violate *BD*-group-strategyproofness.

Theorem 7. *RSD and SML satisfy ST-group-strategyproofness*

For the next result we consider the same conditions as in Theorem 6, but replace *BD*-strategyproofness with *BD*-group-strategyproofness. It turns out that pairwiseness is no longer required for an impossibility.

Theorem 8. *For $n \geq 3$ and $m \geq 3$, there is no anonymous, neutral, ex post efficient, and BD-group-strategyproof SDS, even when preferences are dichotomous.*

We do not know whether neutrality is required for this impossibility and conjecture that it is not. The other three axioms are necessary: serial dictatorship satisfies all axioms except anonymity, always returning the uniform lottery over all alternatives satisfies all axioms except *ex post* efficiency, and *SML* (as well as *RSD*) satisfies all axioms except *BD*-group-strategyproofness.

Theorem 8 is a strengthening of a theorem by Bogomolnaia et al. (2005), who showed that same statement for *SD*-group-strategyproofness and at least four agents and six alternatives.

Let *BD*-strong-group-strategyproofness be the strengthening of *BD*-group-strategyproofness in which only one of the deviating agents has to be strictly better off. For this notion of group-strategyproofness, the statement of Theorem 8 holds even without requiring anonymity and neutrality.

Theorem 9. *For $n \geq 3$ and $m \geq 3$, there is no ex post efficient and BD-strong-group-strategyproof SDS, even when preferences are dichotomous.*

The proof of Theorem 9 is based on a construction by Bogomolnaia et al. (2005), but uses a weaker notion of strategyproofness. Independence of the axioms can be shown using the examples given after Theorem 8.

Our results on group-strategyproofness are summarized in Figure 2.

Acknowledgements

This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/7-1, BR 2312/7-2, and BR 2312/10-1.

The authors thank Anna Bogomolnaia, Paul Harrenstein, Hervé Moulin, Klaus Nehring, Shasi Nandeibam, Arunava Sen, and William Thomson for helpful discussions and the anonymous referees for valuable feedback.

Results of this paper were presented at the Optimization and Incentives Seminar (University of Cambridge, UK, November 2012), the Algorithmic Game Theory and Computational Social Choice Workshop (National University of Singapore, Singapore, January 2013), the Dagstuhl Seminar on the Interface of Computation, Game Theory, and Economics (Dagstuhl, Germany, April 2013), the 12th International Conference on Autonomous Agents and Multi-Agent Systems (Saint Paul, USA, May 2013), the Algorithms Seminar (University of Sydney, Australia, May 2013), the Dagstuhl Seminar on Electronic Markets and Auctions (Dagstuhl, Germany, November 2013), the 15th ACM Conference on Economics and Computation (Palo Alto, USA, June 2014), the 28th AAAI Conference on Artificial Intelligence (Québec City, Canada, July 2014), the 5th World Congress of the Game Theory Society (Maastricht, The Netherlands, July 2016), the 4th Day on Computational Game Theory (ETH Zürich, Switzerland, February 2017), and the Dagstuhl Seminar on Voting (Dagstuhl, Germany, June 2017).

A preliminary version of the proof of Theorem 4 appeared in the Proceedings of AAMAS 2013 and preliminary versions of the proofs of Theorems 5, 6, and 8 appeared in the Proceedings of AAAI 2014. In order to enable a self-contained presentation of the material, proofs of these theorems are contained in the Appendix.

References

- A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- B. Abramowitz and E. Anshelevich. Utilitarians without utilities: Maximizing social welfare for graph problems using only ordinal preferences. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018. Forthcoming.
- H. Aziz. Maximal Recursive Rule: A New Social Decision Scheme. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 34–40. AAAI Press, 2013.

- H. Aziz and P. Stursberg. A generalization of probabilistic serial to randomized social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 559–565. AAAI Press, 2014.
- H. Aziz, F. Brandt, and M. Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 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.
- S. Barberà. Manipulation of social decision functions. *Journal of Economic Theory*, 15(2):266–278, 1977a.
- S. Barberà. The manipulation of social choice mechanisms that do not leave “too much” to chance. *Econometrica*, 45(7):1573–1588, 1977b.
- S. Barberà. A note on group strategy-proof decision schemes. *Econometrica*, 47(3):637–640, 1979a.
- S. Barberà. Majority and positional voting in a probabilistic framework. *Review of Economic Studies*, 46(2):379–389, 1979b.
- P. R. Blavatsky. Axiomatization of a preference for most probable winner. *Theory and Decision*, 60(1):17–33, 2006.
- A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
- A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.
- A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- F. Brandl. Efficiency and incentives in randomized social choice. Master’s thesis, Technische Universität München, 2013.
- 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. AAAI Press, 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. IFAAMAS, 2015b.

- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016a.
- F. Brandl, F. Brandt, and W. Suksompong. The impossibility of extending random dictatorship to weak preferences. *Economics Letters*, 141:44–47, 2016b.
- F. Brandl, F. Brandt, M. Eberl, and C. Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, 65(2), 2018a.
- F. Brandl, F. Brandt, and C. Stricker. An analytical and experimental comparison of maximal lottery schemes. 2018b. Working paper.
- F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4):793–804, 2015.
- F. Brandt and M. Brill. Necessary and sufficient conditions for the strategyproofness of irresolute social choice functions. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 136–142. ACM Press, 2011.
- F. Brandt and F. Fischer. On the hardness and existence of quasi-strict equilibria. In *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT)*, volume 4997 of *Lecture Notes in Computer Science (LNCS)*, pages 291–302. Springer-Verlag, 2008.
- F. Brandt and C. Geist. Finding strategyproof social choice functions via SAT solving. *Journal of Artificial Intelligence Research*, 55:565–602, 2016.
- E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- D. Butler, G. Pogrebna, and P. R. Blavatskyy. Predictably intransitive preferences. 2016. Working paper.
- H.-I. Chang and Y. Chun. Probabilistic assignment of indivisible objects when agents have the same preferences except the ordinal ranking of one object. *Mathematical Social Sciences*, 90:80–82, 2017.
- S. Chatterji, A. Sen, and H. Zeng. Random dictatorship domains. *Games and Economic Behavior*, 86:212–236, 2014.
- Y.-K. Che and F. Kojima. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica*, 78(5):1625–1672, 2010.
- S. Ching and L. Zhou. Multi-valued strategy-proof social choice rules. *Social Choice and Welfare*, 19(3):569–580, 2002.

- W. J. Cho. Incentive properties for ordinal mechanisms. *Games and Economic Behavior*, 95:168–177, 2016.
- J. Duggan. A geometric proof of Gibbard’s random dictatorship theorem. *Economic Theory*, 7(2):365–369, 1996.
- J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.
- B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
- A. Feldman. Manipulation and the Pareto rule. *Journal of Economic Theory*, 21:473–482, 1979a.
- A. Feldman. Nonmanipulable multi-valued social choice decision functions. *Public Choice*, 34:177–188, 1979b.
- D. S. Felsenthal and M. Machover. After two centuries should Condorcet’s voting procedure be implemented? *Behavioral Science*, 37(4):250–274, 1992.
- P. C. Fishburn. Even-chance lotteries in social choice theory. *Theory and Decision*, 3(1):18–40, 1972.
- P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- 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. Dominance in SSB utility theory. *Journal of Economic Theory*, 34(1):130–148, 1984b.
- P. C. Fishburn. *Nonlinear preference and utility theory*. The Johns Hopkins University Press, 1988.
- D. C. Fisher and J. Ryan. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2):217–236, 1995.
- P. Gärdenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13(2):217–228, 1976.
- P. Gärdenfors. On definitions of manipulation of social choice functions. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.

- A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4): 587–601, 1973.
- A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3): 665–681, 1977.
- B. Goodwin. *Justice by Lottery*. University of Chicago Press, 2005. 1st edition 1992.
- A. Guerrero. Against elections: The lottocratic alternative. *Philosophy and Public Affairs*, 42(2):135–178, 2014.
- J. Hadar and W. R. Russell. Rules for ordering uncertain prospects. *American Economic Review*, 59(1):2–34, 1969.
- L. N. Hoang. Strategy-proofness of the randomized Condorcet voting system. *Social Choice and Welfare*, 48(3):679–701, 2017.
- A.-K. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131(1):231–250, 2006.
- J. S. Kelly. Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2):439–446, 1977.
- G. Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3–27 July 1962)*, pages 73–79, 1965.
- G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993.
- G. Laffond, J.-F. Laslier, and M. Le Breton. A theorem on symmetric two-player zero-sum games. *Journal of Economic Theory*, 72(2):426–431, 1997.
- J.-F. Laslier. Aggregation of preferences with a variable set of alternatives. *Social Choice and Welfare*, 17(2):269–282, 2000.
- M. Le Breton. On the uniqueness of equilibrium in symmetric two-player zero-sum games with integer payoffs. *Économie publique*, 17(2):187–195, 2005.
- T. Mennle and S. Seuken. Hybrid mechanisms: Trading off strategyproofness and efficiency of random assignment mechanisms. 2015. Working paper.
- S. Nandeibam. An alternative proof of Gibbard’s random dictatorship result. *Social Choice and Welfare*, 15(4):509–519, 1997.
- K. Nehring. Monotonicity implies generalized strategy-proofness for correspondences. *Social Choice and Welfare*, 17(2):367–375, 2000.

- A. Nesterov. Fairness and efficiency in strategy-proof object allocation mechanisms. *Journal of Economic Theory*, 2017. Forthcoming.
- R. Peyre. La quête du graal électoral. Images des Mathématiques, CNRS, 2013.
- A. Postlewaite and D. Schmeidler. Strategic behaviour and a notion of ex ante efficiency in a voting model. *Social Choice and Welfare*, 3(1):37–49, 1986.
- R. L. Rivest and E. Shen. An optimal single-winner preferential voting system based on game theory. In *Proceedings of the 3rd International Workshop on Computational Social Choice (COMSOC)*, pages 399–410, 2010.
- C. Rodríguez-Álvarez. On the manipulation of social choice correspondences. *Social Choice and Welfare*, 29(2):175–199, 2007.
- C. Rodríguez-Álvarez. On strategy-proof social choice correspondences: a comment. *Social Choice and Welfare*, 32(1):29–35, 2009.
- M. R. Sanver and W. S. Zwicker. Monotonicity properties and their adaption to irresolute social choice rules. *Social Choice and Welfare*, 39(2–3):371–398, 2012.
- S. Sato. On strategy-proof social choice correspondences. *Social Choice and Welfare*, 31: 331–343, 2008.
- M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- P. Stone. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press, 2011.
- Y. Tanaka. An alternative proof of Gibbard’s random dictatorship theorem. *Review of Economic Design*, 8:319–328, 2003.
- A. D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.
- 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.
- W. S. Zwicker. The voter’s paradox, spin, and the Borda count. *Mathematical Social Sciences*, 22(3):187–227, 1991.

APPENDIX

A. Proofs

A.1. Random Serial Dictatorship

Theorem 1. *RSD is ex post efficient and strongly SD-strategyproof.*

Proof. RSD can be seen as a convex combination of serial dictatorships for all permutations of agents. Serial dictatorships are Pareto-efficient because whenever x Pareto-dominates y , y will never be removed before x and x can only be removed together with y . The convex combination of Pareto-efficient SCFs constitutes an *ex post* efficient SDS.

Serial dictatorships are strategyproof because at any stage of the mechanism, an agent can only get a less preferred alternatives selected by lying about his preferences. The convex combination of strategyproof SCFs constitutes a strongly SD-strategyproof SDSs by the linearity of expectation. \square

Theorem 2. *RSD is SD-efficient iff $n \leq 2$, or $m \leq 3$, or $n = 3$ and $m \leq 5$.*

Proof. For the cases where RSD is SD-efficient, we in fact show that SD-efficiency is equivalent to *ex post* efficiency. The statement then follows from the fact that RSD is *ex post* efficient. We assume throughout the proof that there are no two alternatives among which every agent is indifferent. This assumption is without loss of generality, any two such alternatives may be regarded as the same alternative for all conclusions drawn throughout this proof. Let $R \in \mathcal{R}^n$ be a preference profile and let $p \in \Delta(A)$ be *ex post* efficient in R . Assume for contradiction that p is not SD-efficient, i.e., there is $q \in \Delta(A)$ that SD-dominates p , i.e., $q R_i^{SD} p$ for all $i \in N$ and $q P_i^{SD} p$ for some $i \in N$. Let $S^- = \{x \in A: q(x) < p(x)\}$ and $S^+ = \{x \in A: q(x) > p(x)\}$. Since $q \neq p$, both S^- and S^+ are non-empty. All elements of S^- are Pareto-efficient because $S^- \subseteq \hat{p}$. Moreover, observe that, for every agent $i \in N$, there are $x \in S^+$ and $y \in S^-$ such that x is ranked highest and y is ranked lowest among the elements of $S^+ \cup S^-$, since $q R_i^{SD} p$. We now consider the three cases from the statement of the theorem.

- $n \leq 2$: The argument for $n = 1$ is trivial. Consider $n = 2$. Since all $x, x' \in S^-$ are Pareto-efficient and $n = 2$, we have that $x R_1 x'$ implies $x' P_2 x$. Let $x \in S^-$ be the highest ranked alternative in R_1 among the alternatives in S^- and let $y \in S^+$ such that $y R_1 x$. Such a y exists because $q R_1^{SD} p$. Since x is Pareto-efficient, we have that $x P_2 y$. But this implies that $x' P_2 y$ for all $x' \in S^-$, which contradicts $q R_2^{SD} p$.
- $m \leq 3$: Clearly, $S^- \cap S^+ = \emptyset$. Hence either $|S^-| = 1$ or $|S^+| = 1$. But then x Pareto-dominates y for all $x \in S^+$ and $y \in S^-$, which contradicts *ex post* efficiency of p .

- $n = 3$ and $m \leq 5$: Clearly, $S^- \cap S^+ = \emptyset$. If $|S^-| = 1$ or $|S^+| = 1$, then x Pareto-dominates y for all $x \in S^+$ and $y \in S^-$, which contradicts *ex post* efficiency of p . Consider the case $|S^-| = 2$ and $|S^+| = 2$. For every $x \in S^-$ and $y \in S^+$, there is $i \in N$ such that $x P_i y$, since x is Pareto optimal. From $n = 3$, it follows that there are $i \in N$ such that either $x P_i y$ for some $x \in S^-$ and all $y \in S^+$ or $x P_i y$ for all $x \in S^-$ and some $y \in S^+$. Either case contradicts $q R_i^{SD} p$. Now consider the case $|S^-| = 2$ and $|S^+| = 3$ and let $S^- = \{a, b\}$. Without loss of generality, $a P_1 b$ and $b R_i a$ for $i \in \{2, 3\}$. As stated above, $q R_i^{SD} p$ and $b R_i a$ imply that a is ranked last in R_i for $i \in \{2, 3\}$. Also, for every $y \in S^+$, there is $i \in N$ such that $a P_i y$. It follows that $a P_1 y$ for all $y \in S^+$. This contradicts the fact that $q R_1^{SD} p$. The case $|S^-| = 3$ and $|S^+| = 2$ is analogous.

To show that *RSD* does not satisfy *SD*-efficiency for $n = 3$ and $m = 6$, consider the following preference profile R .

$$\begin{aligned} R_1 &: \{a, d\}, c, e, b, f \\ R_2 &: \{b, e\}, a, f, c, d \\ R_3 &: \{c, f\}, b, d, a, e \end{aligned}$$

$RSD(R) = 1/6 a + 1/6 b + 1/6 c + 1/6 d + 1/6 e + 1/6 f = p$. However, for the lottery $q = 1/3 a + 1/3 b + 1/3 c$ we have $q P_i^{SD} p$ for all $i \in \{1, 2, 3\}$, i.e., p is not *SD*-efficient.

An example showing that *RSD* does not satisfy *SD*-efficiency for $n = 4$ and $m = 4$ has been given in Section 4.

If *RSD* is *SD*-inefficient for some number of agents and alternatives, then it is also inefficient for any larger numbers. To see this for the number of agents, observe that adding agents who are indifferent between all alternatives does not change the set of *SD*-efficient lotteries. And to see this for the number of alternatives, observe that adding an alternative at the bottom of each agent's preference relation does not change the set of *SD*-efficient lotteries. Hence, the statements follow from induction on the number of agents and the number of alternatives with the above examples for inefficiency of *RSD* as base cases. \square

A.2. Maximal Lotteries

Theorem 3. *Every SDS that returns maximal lotteries is PC-efficient.*

Proof. With each preference relation $R_i \in \mathcal{R}$ we can associate a function $\phi_i : A \times A \rightarrow \{-1, 0, 1\}$ such that for all $x, y \in A$,

$$\phi_i(x, y) = \begin{cases} 1 & \text{if } x P_i y, \\ -1 & \text{if } y P_i x, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

With slight abuse of notation, we extend ϕ_i to a function from $\Delta(A) \times \Delta(A)$ to $[0, 1]$. For $p, q \in \Delta(A)$, let

$$\phi_i(p, q) = \sum_{x, y \in A} p(x)q(y)\phi_i(x, y).$$

Observe that $p R_i^{PC} q$ iff $\phi_i(p, q) \geq 0$. Now, let $R \in \mathcal{R}^n$ and $q \in ML(R)$. If q is not PC -efficient, there are $p \in \Delta(A)$ and $j \in N$ such that $p P_j^{PC} q$ and $p R_i^{PC} q$ for all $i \in N$. Thus, $\phi_j(p, q) > 0$ and $\phi_i(p, q) \geq 0$ for all $i \in N$. This implies that

$$g_R(p, q) = \sum_{i \in N} \phi_i(p, q) > 0.$$

However, this contradicts the assumption that $q \in ML(R)$. Thus, q is PC -efficient. \square

For the proof of Theorem 4 we need the following lemma, which states that weakening alternatives outside $\widehat{SML}(R)$ does not affect the set of strict maximal lotteries.¹⁹

Lemma 1. *Let $R \in \mathcal{R}^n$ and $a \in A$ with $a \notin \widehat{SML}(R)$. Let furthermore $R' \in \mathcal{R}^n$ be such that $g_{R'}(b, a) > g_R(b, a)$ for some $b \in A \setminus \{a\}$ and $g_{R'}(x, y) = g_R(x, y)$ for all $x, y \in A$ with $\{x, y\} \neq \{a, b\}$. Then, $SML(R') = SML(R)$.*

Proof. Let R, R' , and a be as stated. First, assume for contradiction that $a \in \widehat{SML}(R')$. It follows from Dutta and Laslier (1999), Theorem 4.4, that $a \in \widehat{SML}(R)$ which is a contradiction. Now, by definition, $p \in SML(R)$ if $g_R(x, p) = 0$ for all $x \in \widehat{SML}(R)$ and $g_R(x, p) < 0$ for all $x \notin \widehat{SML}(R)$. Since for all $x \in A \setminus \{a\}$,

$$g_R(x, p) = \sum_{y \in A} p(y)g_R(x, y) = \sum_{y \in \widehat{p}} p(y)g_R(x, y) = \sum_{y \in \widehat{p}} p(y)g_{R'}(x, y) = g_{R'}(x, p),$$

it follows that $SML(R) = SML(R')$. \square

Theorem 4. *SML is PC-efficient and ST-strategyproof.*

Proof. The first statement directly follows from Theorem 3 and the fact that $SML(R) \subseteq ML(R)$ for all $R \in \mathcal{R}^n$.

For a proof of the second statement, assume for contradiction that there is an SML scheme f that is not ST -strategyproof. Then, there are two preference profiles R and R' and an agent i such that $R_j = R'_j$ for all $j \neq i$ and $q = f(R') P_i^{ST} f(R) = p$.

For two alternatives $x, y \in A$ we say that x is *strengthened against* y if either (1) $y R_i x$ and $x P'_i y$, or (2) $y P_i x$ and $x R'_i y$. Define $\Delta(R, R') =$

¹⁹This proof recycles an argument that was used by Brandt (2015) to characterize set-valued social choice functions that are strategyproof with respect to Kelly's set extension.

$\{(x, y) : x \text{ is strengthened against } y\}$. This set can be partitioned into the following four subsets.

$$\begin{aligned}\Delta_1 &= \{(x, y) \in \Delta(R, R') : y \notin \widehat{p}\} \\ \Delta_2 &= \{(x, y) \in \Delta(R, R') : x \notin \widehat{q}\} \setminus \Delta_1 \\ \Delta_3 &= \{(x, y) \in \Delta(R, R') : x \in \widehat{q}, y \in \widehat{p}, \text{ and } \{x, y\} \not\subseteq \widehat{p} \cap \widehat{q}\} \\ \Delta_4 &= \{(x, y) \in \Delta(R, R') : x \in \widehat{q}, y \in \widehat{p}, \text{ and } \{x, y\} \subseteq \widehat{p} \cap \widehat{q}\}\end{aligned}$$

We now construct two new preference profiles \tilde{R} and \tilde{R}' based on R and R' . The idea behind this construction is to make R and R' agree on as many pairs as possible, while maintaining the invariant that the outcomes are p and q , respectively.

\tilde{R} is identical to R except that for all pairs $(x, y) \in \Delta_1$, we strengthen x against y in the preferences of agent i such that \tilde{R}_i agrees with R'_i on all such pairs.²⁰ Lemma 1 implies that $f(\tilde{R}) = f(R) = p$. Analogously, \tilde{R}' is identical to R' except that for all pairs $(x, y) \in \Delta_2$, we strengthen y against x in the preferences of agent i such that \tilde{R}'_i agrees with R_i on all such pairs. Lemma 1 implies that $f(\tilde{R}') = f(R') = q$.

By definition, \tilde{R} and \tilde{R}' differ only on pairs that are contained in Δ_3 or Δ_4 . Observe, however, that $\Delta_3 = \emptyset$. To see this, assume for contradiction that there is a pair $(x, y) \in \Delta(R, R')$ with $x \in \widehat{q}$, $y \in \widehat{p}$, and $\{x, y\} \not\subseteq \widehat{p} \cap \widehat{q}$. There are three cases:

- $x \in \widehat{q} \setminus \widehat{p}$ and $y \in \widehat{p} \setminus \widehat{q}$,
- $x \in \widehat{q} \setminus \widehat{p}$ and $y \in \widehat{p} \cap \widehat{q}$, and
- $x \in \widehat{p} \cap \widehat{q}$ and $y \in \widehat{p} \setminus \widehat{q}$.

In each case, $q P_i^{ST} p$ implies $x P_i y$. Since $(x, y) \in \Delta(R, R')$ implies $y R_i x$, we have a contradiction.

We thus have that $\Delta_3 = \emptyset$, and, consequently, that \tilde{R} and \tilde{R}' only differ on pairs of alternatives that are contained in Δ_4 . In particular, $g_{\tilde{R}}$ and $g_{\tilde{R}'}$ agree on all pairs of alternatives that do not lie in $\widehat{p} \cap \widehat{q}$, i.e.,

$$g_{\tilde{R}}(a, b) = g_{\tilde{R}'}(a, b) \text{ for all } a, b \text{ with } \{a, b\} \not\subseteq \widehat{p} \cap \widehat{q}.$$

For such pairs, we omit the subscript and write $g(a, b)$ instead of $g_{\tilde{R}}(a, b)$. Likewise, we write $g(a, p)$ for $g_{\tilde{R}}(a, p)$ whenever $a \notin \widehat{p} \cap \widehat{q}$ and $p \in \Delta(A)$.

Let $x \in A$ and define a function $s : A \rightarrow [0, 1]$ via

$$s(x) = \begin{cases} p(x) & \text{if } x \in \widehat{p} \\ q(x) & \text{if } x \in \widehat{q} \\ 0 & \text{otherwise.} \end{cases}$$

²⁰Note that \tilde{R}_i might not be transitive. Therefore, we do not assume transitivity of preferences in this proof. In fact, the statement of Theorem 4 becomes stronger but is easier to prove for general—possibly intransitive—preferences.

Note that s is well-defined because $p(z) = q(z)$ for all $z \in \hat{p} \cap \hat{q}$, and that s may *not* correspond to a lottery because the individual probabilities may not add up to one.

For $a \notin \hat{p} \cap \hat{q}$ and a subset $B \subseteq A$ of alternatives, let furthermore $s(a, B) = \sum_{b \in B} s(b)g(a, b)$. If $B = \hat{q}$, we have $s(a, \hat{q}) = \sum_{b \in \hat{q}} s(b)g(a, b) = \sum_{b \in \hat{q}} q(b)g(a, b) = \sum_{b \in A} q(b)g(a, b) = g(a, q)$. Analogously, $s(a, \hat{p})$ equals $g(a, p)$.

By definition, $g(x, p) = 0$ and $g(x, q) < 0$ for all $x \in \hat{p} \setminus \hat{q}$, as well as $g(y, p) < 0$ and $g(y, q) = 0$ for all $y \in \hat{q} \setminus \hat{p}$. Therefore, we get

$$\begin{aligned} s(x, \hat{q}) &= g(x, q) < 0 = g(x, p) = s(x, \hat{p}) \text{ for all } x \in \hat{p} \setminus \hat{q}, \text{ and} \\ s(y, \hat{p}) &= g(y, p) < 0 = g(y, q) = s(y, \hat{q}) \text{ for all } y \in \hat{q} \setminus \hat{p}. \end{aligned}$$

The inequality $s(x, \hat{q}) < s(x, \hat{p})$ remains valid if $s(x, \hat{p} \cap \hat{q})$ is subtracted from both sides. Since $s(x, \hat{q}) - s(x, \hat{p} \cap \hat{q}) = s(\hat{q} \setminus \hat{p})$ and $s(x, \hat{p}) - s(x, \hat{p} \cap \hat{q}) = s(\hat{p} \setminus \hat{q})$, we obtain

$$\begin{aligned} s(x, \hat{q} \setminus \hat{p}) &< s(x, \hat{p} \setminus \hat{q}) \text{ for all } x \in \hat{p} \setminus \hat{q}, \text{ and} \\ s(y, \hat{p} \setminus \hat{q}) &< s(y, \hat{q} \setminus \hat{p}) \text{ for all } y \in \hat{q} \setminus \hat{p}. \end{aligned}$$

Multiplying both sides of these inequalities with a positive number and writing $s'(a, B)$ for $s(a) \cdot s(a, B)$ results in

$$\begin{aligned} s'(x, \hat{q} \setminus \hat{p}) &< s'(x, \hat{p} \setminus \hat{q}) \text{ for all } x \in \hat{p} \setminus \hat{q}, \text{ and} \\ s'(y, \hat{p} \setminus \hat{q}) &< s'(y, \hat{q} \setminus \hat{p}) \text{ for all } y \in \hat{q} \setminus \hat{p}. \end{aligned}$$

We finally summarize over $\hat{q} \setminus \hat{p}$ and $\hat{p} \setminus \hat{q}$, respectively, and get

$$\sum_{x \in \hat{p} \setminus \hat{q}} s'(x, \hat{q} \setminus \hat{p}) < \sum_{x \in \hat{p} \setminus \hat{q}} s'(x, \hat{p} \setminus \hat{q}), \text{ and} \quad (1)$$

$$\sum_{y \in \hat{q} \setminus \hat{p}} s'(y, \hat{p} \setminus \hat{q}) < \sum_{y \in \hat{q} \setminus \hat{p}} s'(y, \hat{q} \setminus \hat{p}). \quad (2)$$

In order to arrive at a contradiction, we state two straightforward identities that are based on the skew-symmetry of g .

$$\sum_{a \in B} s'(a, B) = 0 \text{ for all } B \subseteq A \setminus (\hat{p} \cap \hat{q}), \text{ and} \quad (3)$$

$$\sum_{b \in B} s'(b, C) + \sum_{c \in C} s'(c, B) = 0 \text{ for all } B, C \subseteq A \setminus (\hat{p} \cap \hat{q}). \quad (4)$$

Now (3) implies that the right hand side of both (1) and (2) is zero, and therefore

$$\begin{aligned} \sum_{x \in \hat{p} \setminus \hat{q}} s'(x, \hat{q} \setminus \hat{p}) &< 0, \text{ and} \\ \sum_{y \in \hat{q} \setminus \hat{p}} s'(y, \hat{p} \setminus \hat{q}) &< 0. \end{aligned}$$

However, (4) implies that

$$\sum_{x \in \widehat{p} \setminus \widehat{q}} s'(x, \widehat{q} \setminus \widehat{p}) + \sum_{y \in \widehat{q} \setminus \widehat{p}} s'(y, \widehat{p} \setminus \widehat{q}) = 0,$$

a contradiction. □

A.3. Negative Results

Theorem 5. *There is no anonymous, neutral, PC-efficient, and PC-strategyproof SDS for $n \geq 3$ and $m \geq 3$.*

Proof. This result is established by reasoning about a set of preference profiles for a fixed number of agents and alternatives and deriving a contradiction. We prove the statement for $n = 3$ and $m = 3$. It can be generalized to any larger number of agents and alternatives as described in the proof of Theorem 2. To show a statement for more agents, we add agents that are indifferent between all alternatives. To show a statement for more alternatives, we add alternatives at the bottom of each agent's preference ranking. Both constructions do not affect the incentives of agents and the set of efficient lotteries. Hence, the proof carries through with the same arguments.

Let f be an SDS that satisfies anonymity, neutrality, PC-efficiency, and PC-strategyproofness. First, we consider the following preference profile.

$$R_1^1: a, \{b, c\} \qquad R_2^1: b, a, c \qquad R_3^1: c, a, b$$

Anonymity and neutrality imply that $f(R^1)(b) = f(R^1)(c)$. The only PC-efficient lottery which puts equal weight on b and c is the degenerate lottery a , since every other lottery of this form is dominated by a (agent 2 and 3 are indifferent while agent 1 is strictly better off). Hence, $f(R^1) = a$. Now, consider the following profile.

$$R_1^2: a, \{b, c\} \qquad R_2^2: b, a, c \qquad R_3^2: \{a, c\}, b$$

In this profile, a Pareto-dominates c , hence $f(R^2)(c) = 0$. If agent 3 reports R_3^1 instead of R_3^2 , he receives one of his most preferred alternatives, namely a , with probability 1. Therefore, PC-strategyproofness implies that $f(R^2) = a$. Next, consider the following preference profile.

$$R_1^3: a, \{b, c\} \qquad R_2^3: b, \{a, c\} \qquad R_3^3: \{a, c\}, b$$

Again, PC-efficiency implies that $f(R^3)(a) = 0$, since a Pareto-dominates c . If $f(R^3)(b) > 0$, agent 2 has an incentive to report R_2^3 instead of R_2^2 in R^2 . Thus, $f(R^3) = a$.

Since we will need it later, we state an observation for the following preference profile.

$$R_1^4: c, a, b \qquad R_2^4: a, b, c \qquad R_3^4: b, c, a$$

Anonymity and neutrality imply that $f(R^4) = 1/3 a + 1/3 b + 1/3 c$. Also notice that agent 1 prefers any lottery with higher probability on c than on b to the uniform lottery according to the PC -extension if his preferences are R_1^4 . Now, consider another preference profile.

$$R_1^5: \{a, c\}, b \quad R_2^5: a, b, c \quad R_3^5: b, c, a$$

Here we distinguish two cases. First, assume $f(R^5) = a$ and consider a deviation by agent 3.

$$R_1^6: \{a, c\}, b \quad R_2^6: a, b, c \quad R_3^6: c, b, a$$

Anonymity and neutrality imply that $f(R^6)(a) = f(R^6)(c)$. Any lottery of this form other than $1/2 a + 1/2 c$ is PC -dominated by the latter. Thus, PC -efficiency implies that $f(R^6) = 1/2 a + 1/2 c$. But agent 3 prefers $1/2 a + 1/2 c$ to a if his preferences are R_3^6 . This is a contradiction to PC -strategyproofness. The second case is $f(R^5) \neq a$. If $f(R^5)(c) > f(R^5)(b)$, then by the above observation, agent 1 prefers $f(R^5)$ to $f(R^4)$ if his preferences are R_1^4 . This is a contradiction to PC -strategyproofness. Hence, $f(R^5)(c) \leq f(R^5)(b)$ and, since $f(R^5) \neq a$, $f(R^5)(b) > 0$.

$$R_1^7: \{a, c\}, b \quad R_2^7: a, b, c \quad R_3^7: b, \{a, c\}$$

It follows from $f(R^5)(b) > 0$ that $f(R^7)(b) > 0$, since otherwise agent 3 can benefit from reporting R_3^7 instead of R_3^5 . In particular, we get $f(R^7) \neq a$. Finally, consider the following preference profile.

$$R_1^8: \{a, c\}, b \quad R_2^8: a, \{b, c\} \quad R_3^8: b, \{a, c\}$$

It follows from anonymity that $f(R^8) = f(R^3) = a$. But this implies that agent 2 can successfully deviate from R_2^7 to R_2^8 , since he prefers a to any other lottery if his preferences are R_2^7 . Hence, we obtain the desired contradiction. \square

Theorem 6. *There is no pairwise, ex post efficient, and BD -strategyproof SDS for $n \geq 4$ and $m \geq 4$.*

Proof. We prove the statement for $n = 4$ and $m = 4$. It can be generalized to any larger number of agents and alternatives as described in the proof of Theorem 2. Let f be a pairwise, *ex post* efficient, and BD -strategyproof SDS. We first consider the preference profile R^1 and its weighted majority graph depicted in Figure 3 (i).

$$R_1^1: a, c, \{b, d\} \quad R_2^1: b, d, \{a, c\}$$

Both, c and d are Pareto-dominated in R^1 and, thus, *ex post* efficiency implies $f(R^1)(c) = f(R^1)(d) = 0$. Since f is pairwise and, in particular, anonymous and neutral, it follows that $f(R^1) = 1/2 a + 1/2 b = p$. Now we consider the preference profile R^2 and its weighted majority graph as in Figure 3 (ii).

$$R_1^2: a, c, \{b, d\} \quad R_2^2: \{b, d\}, \{a, c\}$$

Both agents are indifferent between b and d and again c is Pareto-dominated. Thus, pairwiseness and *ex post* efficiency imply that $f(R^2)(b) = f(R^2)(d)$ and $f(R^2)(c) = 0$. Hence, $f(R^2) = (1 - 2\lambda)a + \lambda b + \lambda d = q$ for some $\lambda \in [0, 1/2]$.

First, assume for contradiction that $\lambda > 1/3$. We consider the following preference profile and its weighted majority graph depicted in Figure 3 (*iii*).

$$R_1^3: a, \{b, c, d\} \qquad R_2^3: \{b, d\}, \{a, c\}$$

Pairwiseness implies that $f(R^3) = 1/3 a + 1/3 b + 1/3 d = r$. But $r (P_1^2)^{BD} q$ if $\lambda > 1/3$, which contradicts *BD*-strategyproofness of f since agent 1 can manipulate in R^2 by reporting R_1^3 instead of R_1^2 .

Now assume for a contradiction that $\lambda = 1/3$.

$$\begin{array}{ll} R_1^4: a, c, b, d & R_2^4: \{b, d\}, \{a, c\} \\ R_3^4: a, c, \{b, d\} & R_4^4: \{b, d\}, c, a \end{array}$$

The weighted majority graph of R^4 is equal to that of R^1 and, thus, $f(R^4) = f(R^1) = p$.

$$\begin{array}{ll} R_1^5: a, c, b, d & R_2^5: \{b, d\}, \{a, c\} \\ R_3^5: a, c, \{b, d\} & R_4^5: d, b, c, a \end{array}$$

The majority graph of R^5 is equal to that of R^2 and, hence, $f(R^5) = q$. But then, agent 4 in R^4 can manipulate by reporting R_4^5 instead of R_4^4 since $q (P_4^4)^{BD} p$. This again contradicts *BD*-strategyproofness of f .

Finally, we assume $\lambda < 1/3$ and consider the preference profile R^6 .

$$\begin{array}{ll} R_1^6: a, c, \{b, d\} & R_2^6: \{b, d\}, \{a, c\} \\ R_3^6: \{b, c, d\}, a & R_4^6: a, \{b, c, d\} \end{array}$$

The weighted majority graph of R^6 is equal to that of R^2 and, therefore, $f(R^6) = f(R^2) = q$. We consider one last preference profile.

$$\begin{array}{ll} R_1^7: a, c, \{b, d\} & R_2^7: \{b, d\}, \{a, c\} \\ R_3^7: \{b, d\}, c, a & R_4^7: a, \{b, c, d\} \end{array}$$

The majority graph of R^7 is equal to that of R^1 , which implies that $f(R^7) = f(R^3) = r$. But $r (P_3^6)^{BD} q$ if $\lambda < 1/3$. Thus, agent 3 in R^6 can benefit from reporting R_3^7 instead of R_3^6 . In any case, we found a successful manipulation, contradicting *BD*-strategyproofness of f . \square

A.4. Group-strategyproofness

Theorem 8. *For $n \geq 3$ and $m \geq 3$, there is no anonymous, neutral, ex post efficient, and *BD*-group-strategyproof SDS, even when preferences are dichotomous.*




Figure 3: Graphs depicting pairwise comparisons. An edge from x to y is labeled with $g_R(x, y)$, the number of agents preferring x to y minus the number of agents preferring y to x in preference profile R . All missing edges denote majority ties.

Proof. We prove the statement for $n = 3$ and $m = 3$. It can be generalized to any larger number of agents and alternatives as described in the proof of Theorem 2. Assume for contradiction there is an SDS f with the properties as stated and consider the following preference profile.

$$R_1^1: \{a, b\}, c \qquad R_2^1: \{a, c\}, b \qquad R_3^1: \{b, c\}, a$$

By neutrality and anonymity, $f(R^1) = 1/3 a + 1/3 b + 1/3 c$. Now let agents 1 and 2 change their preferences and consider the profile R^2 .

$$R_1^2: a, \{b, c\} \qquad R_2^2: a, \{b, c\} \qquad R_3^2: \{b, c\}, a$$

Again by neutrality and anonymity, $f(R^2) = (1 - 2\lambda)a + \lambda b + \lambda c$. If $\lambda > 1/3$, then agents 1 and 2 would rather report R_1^1 and R_2^1 respectively if their true preferences were R_1^2 and R_2^2 . On the other hand, if $\lambda < 1/3$ and their true preferences were R_1^1 and R_2^1 , they would rather report R_1^2 and R_2^2 . Hence, $\lambda = 1/3$ and $f(R^2) = 1/3 a + 1/3 b + 1/3 c$.

$$R_1^3: a, \{b, c\} \qquad R_2^3: \{a, b\}, c \qquad R_3^3: b, \{a, c\}$$

In R^3 , c is Pareto-dominated, thus by neutrality and anonymity, $f(R^3) = 1/2 a + 1/2 b$. To this end, we consider the following profile.

$$R_1^4: a, \{b, c\} \qquad R_2^4: \{a, b\}, c \qquad R_3^4: \{b, c\}, a$$

If agent 3 changes his preferences from R_3^3 to R_3^4 , c is still Pareto-dominated and his preferences over a and b remain unchanged. Hence, by *BD*-strategyproofness, $f(R^4) = f(R^3)$. But then agent 2 in R^2 would have an incentive to report R_2^4 instead of R_2^2 , a contradiction. \square

Theorem 9. *There is no ex post efficient and *BD*-strong-group-strategyproof SDS for $n \geq 3$ and $m \geq 3$, even when preferences are dichotomous.*

Proof. We prove the statement for $n = 3$ and $m = 3$. It can be generalized to any larger number of agents and alternatives as described in the proof of Theorem 2. Assume for contradiction that f is an SDS with properties as stated and consider the following preference profile.

$$R_1^1: a, \{b, c\} \qquad R_2^1: b, \{a, c\} \qquad R_3^1: c, \{a, b\}$$

Let $f(R^1) = p$. We assume without loss of generality that $p(a) > 0$. Now consider a variation of the previous preference profile in which agent 1 is completely indifferent and let $f(R^2) = q$.

$$R_1^2: \{a, b, c\} \qquad R_2^2: b, \{a, c\} \qquad R_3^2: c, \{a, b\}$$

Clearly, a is Pareto-dominated by both b and c and therefore $q(a) = 0$. If agent 1 claims his preferences are as in R_1^3 , alternative a remains Pareto-dominated by b .

$$R_1^3: \{a, b\}, \{c\} \qquad R_2^3: b, \{a, c\} \qquad R_3^3: c, \{a, b\}$$

Let $f(R^3) = r$ where $r(a) = 0$. If we assume that R_1^2 is the true preference relation of agent 1, group-strategyproofness requires that agent 2 should not prefer r to q because otherwise this may be seen as a beneficial group deviation of agents 1 and 2. As a consequence, $r(b) \leq q(b)$. Similarly, a group deviation by agents 1 and 3 implies that $r(c) \leq q(c)$ and consequently that $r = q$.

$$R_1^4: \{a, c\}, \{b\} \qquad R_2^4: b, \{a, c\} \qquad R_3^4: c, \{a, b\}$$

If we consider the profile R^4 , a is Pareto-dominated by c and analogous arguments imply that $f(R^4) = q$.

Finally, consider the preference profile R^3 again. Strategyproofness implies that agent 1 should not benefit from deviating to R_1^1 . It can be shown that

$$\neg(p (R_1^3)^{BD} q) \text{ iff } q(c) < \frac{p(c)}{p(b) + p(c)}.$$

Similarly, agent 1 should not benefit from deviating to R_1^1 in profile R^4 and

$$\neg(p (R_1^4)^{BD} q) \text{ iff } q(b) < \frac{p(b)}{p(b) + p(c)}.$$

Adding both inequalities yields that

$$q(b) + q(c) = 1 < \frac{p(b)}{p(b) + p(c)} + \frac{p(c)}{p(b) + p(c)} = 1,$$

a contradiction. □