Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence

On the Incompatibility of Efficiency and Strategyproofness in Randomized Social Choice

Haris Aziz
NICTA and UNSW
Sydney 2033, Australia

Florian Brandl
Technische Universität München
85748 Garching, Germany

Felix Brandt
Technische Universität München
85748 Garching, Germany

Abstract

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—are two cornerstones of economics and ubiquitous in important areas such as voting, auctions, or matching markets. Within the context of random assignment, Bogomolnaia and Moulin have shown that two particular notions of efficiency and strategyproofness based on stochastic dominance are incompatible. However, there are various other possibilities of lifting preferences over alternatives to preferences over lotteries apart from stochastic dominance. In this paper, we give an overview of common preference extensions, propose two new ones, and show that the abovementioned incompatibility can be extended to various other notions of strategyproofness and efficiency in randomized social choice.

1 Introduction

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—are two cornerstones of economics and ubiquitous in important areas such as voting, auctions, or matching markets. The conflict between both notions is already apparent in Gibbard and Satterthwaite's seminal theorem, which states that the only single-valued social choice functions 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 (Gibbard 1977; Barberà 1979). Randomized voting methods have a surprisingly long tradition going back to ancient Greece and have recently gained increased attention in political science (Stone 2011). Within computer science, randomization is a very successful technique in algorithm design and is being considered more and more often in the context of voting (Conitzer and Sandholm 2006; Procaccia 2010; Walsh and Xia 2012; Service and Adams 2012; Birrell and Pass 2011; Aziz, Brandt, and Brill 2013b; Aziz 2013).

There are various ways of extending preferences over alternatives to preferences over lotteries. We refer to these extensions as *lottery extensions*. Perhaps the most wide-spread lottery extension is *stochastic dominance* (SD). This extension is of particular importance because one lottery stochastically dominates another one iff the former yields at least as much expected utility as the latter for any von-Neumann-Morgenstern (vNM) utility representation consistent with the ordinal preferences. However, settings in which the existence of an underlying vNM utility function cannot be assumed may call for other lottery extensions. For instance, in this paper, we put forward a particularly natural new extension called *pairwise comparison* (PC), which arises as the special case of skew-symmetric bilinear (SSB) utility functions as proposed by Fishburn (1982).[1] According to this extension lottery p is preferred to lottery q iff it is more likely that p yields a better alternative than q. The PC extension is more powerful than the SD extension in the sense that, for the same preference relation over alternatives, the SD preference relation is contained in the PC relation. Apart from PC, we consider two other completions of SD due to Cho (2012), namely the *upward lexicographic* (UL) and the *downward lexicographic* (DL) extension. We furthermore consider a weakening of SD that we call *bilinear dominance* (BD) and which is again based on Fishburn's SSB utility (Fishburn 1984). 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 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 orig-

[1]SSB utility functions are a generalization of vNM utility functions.

inal 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. Usually, the strong notion is much more demanding than the weak one. Whenever a lottery extension is *complete*, however, both notions coincide.

One of the best known results about SDSs is a consequence of a characterization by Gibbard (1977), who attributes it to Hugo Sonnenschein: when individual preferences are linear, every Pareto optimal and strongly SD-strategyproof SDS is a random dictatorship, i.e., one of the agents is chosen at random and then picks his most preferred alternative.[2] Gibbard's proof requires the universal domain of linear preferences and cannot be extended to arbitrary subdomains (Chatterji, Sen, and Zeng 2014). Moreover, in many important subdomains of social choice such as house allocation, matching, and coalition formation, ties are unavoidable since agents are indifferent among all outcomes in which their allocation, match, or coalition is the same (Sönmez and Ünver 2011; Aziz, Brandt, and Seedig 2013; Bouveret and Lang 2008; Elkind and Wooldridge 2009; Aziz, Brandt, and Seedig 2013). Within the special domain of random assignment, Bogomolnaia and Moulin (2001) have been able to show that there is no anonymous, SD-efficient, and strongly SD-strategyproof SDS. As a consequence, all generalizations of random dictatorship to weak preferences, violate SD-efficiency or strong SD-strategyproofness.[3] Aziz, Brandt, and Brill (2013b) recently conjectured that the impossibility by Bogomolnaia and Moulin even holds when only requiring *weak* SD-strategyproofness and proved this for the rather limited class of *majoritarian* SDSs.[4] We complement and strengthen these results by proving the following theorems (always assuming anonymity):

1. PC-strategyproofness is incompatible with PC-efficiency in the context of neutral SDSs.
2. UL-strategyproofness is incompatible with UL-efficiency.
3. BD-strategyproofness is incompatible with Pareto optimality in the context of pairwise SDSs.
4. BD-group-strategyproofness is incompatible with Pareto-optimality in the context of neutral SDSs.

The first result is a proof of a particularly natural weakening of the above mentioned conjecture by Aziz, Brandt, and Brill (2013b). The second result might be surprising because the corresponding statement for the DL-extension does *not* hold (random dictatorship satisfies both DL-efficiency and DL-strategyproofness). The third and the fourth result significantly strengthen theorems by Aziz, Brandt, and Brill (2013b) (Theorem 1) and Bogomolnaia, Moulin, and Stong (2005) (Proposition 3), respectively.

The assumption of anonymity is crucial as all our impossibilities fail to hold when omitting anonymity. *Serial dictatorship*, an extreme example of a non-anonymous SDS, 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

Apart from some early precursors (Zeckhauser 1969; Fishburn 1972), the first formal study of strategyproof randomized social choice was conducted by Gibbard (1977). A recent survey of randomized social choice is contained in a book chapter by Barberà (2010). Using stochastic dominance for strategyproofness, efficiency, and fairness conditions was popularized by Bogomolnaia and Moulin (2001). They focussed on a subdomain of randomized social choice called *random assignment*, in which each outcome is a one-to-one assignment of objects to agents. Recently, Cho (2012) extended the approach of Bogomolnaia and Moulin (2001) by introducing new lottery extensions such as ones based on lexicographic preferences. Aziz, Brandt, and Brill (2013b) examined the tradeoff between efficiency and strategyproofness for social decision schemes and initiated the analysis of *strict maximal lotteries*, a little known SDS due to Kreweras and Fishburn. Recently, Aziz (2013) proposed a new SDS that compromises between RSD and efficient but manipulable SDSs.

A line of inquiry that has been especially popular in AI and multi-agent systems is to check how well strategyproof SDSs approximate common deterministic voting rules such as Borda's rule (Conitzer and Sandholm 2006; Procaccia 2010; Birrell and Pass 2011; Service and Adams 2012).

3 Preliminaries

Let $N = \{1, \ldots, n\}$ be a set of agents with preferences over a finite set A with $|A| = m$. The preferences of agent $i \in N$ are represented by 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 $b \; R_i \; a$. A *preference profile* $R = (R_1, \ldots, R_n)$ is an n-tuple containing a preference relation R_i for each agent $i \in N$. The set of all preference profiles is thus given by $\mathcal{R}^n$. We will compactly represent a preference relation as a comma-separated list with all alternatives among which an agent is indifferent placed in a set.

[2]While random dictatorship is strongly SD-strategyproof, it only satisfies *weak* SD-*group*-strategyproofness.

[3]Random *serial* dictatorship, for instance, somewhat surprisingly violates SD-efficiency (Bogomolnaia and Moulin 2001; Bogomolnaia, Moulin, and Stong 2005; Aziz, Brandt, and Brill 2013b).

[4]Within the domain of random assignment with unit-demand and the domain of dichotomous preferences, respectively, the conditions of the conjecture are compatible (Bogomolnaia and Moulin 2001; Bogomolnaia, Moulin, and Stong 2005).

For example $a\ P_i\ b\ I_i\ c$ is represented by $R_i\colon a, \{b, c\}$. 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 $y\ I_i\ z$.

Our central object of study are social decision schemes, i.e., functions that map the individual preferences of the agents to a lottery (or *probability distribution*) over alternatives. A *social decision scheme (SDS)* is a function $f\colon \mathcal{R}^n \to \Delta(A)$.

A minimal fairness condition for SDSs is *anonymity*, which requires that $f(R) = f(R')$ for all $R, R' \in \mathcal{R}^n$ and permutations $\pi\colon N \to N$ such that $R'_i = R_{\pi(i)}$ for all $i \in N$. Another fairness requirement is *neutrality*. For a permutation π of A and a preference relation R_i, $\pi(x)\ R_i^\pi\ \pi(y)$ if and only if $x\ R_i\ y$. Then, an SDS f is *neutral* if for all $R \in \mathcal{R}^n$, $f(R)(x) = f(R^\pi)(\pi(x))$ for all $x \in A$.

An SDS f is *pairwise* (or a neutral C2 function) if it is neutral and for all preference profiles R and R', $f(R) = f(R')$ whenever for all alternatives x, y,

$$\begin{aligned}&|\{i \in N \mid x\ R_i\ y\}| - |\{i \in N \mid y\ R_i\ x\}|\\ = {}&|\{i \in N \mid x\ R'_i\ y\}| - |\{i \in N \mid y\ R'_i\ x\}|.\end{aligned}$$

In other words, the outcome of a pairwise SDS only depends on the anonymized comparisons between pairs of alternatives (Young 1974; Zwicker 1991).

An SDS f is *majoritarian* (or a neutral C1 function) if it is neutral and for all preference profiles R and R', $f(R) = f(R')$ whenever for all alternatives x, y,

$$\begin{aligned}&|\{i \in N \mid x\ R_i\ y\}| \geq |\{i \in N \mid y\ R_i\ x\}|\\ \text{iff } &|\{i \in N \mid x\ R'_i\ y\}| \geq |\{i \in N \mid y\ R'_i\ x\}|.\end{aligned}$$

It is easy to see that the three classes of SDSs form a hierarchy: every majoritarian SDS is pairwise and every pairwise SDS is anonymous.

Two anonymous SDSs that have been recently analyzed in a framework similar to this paper are *random serial dictatorship* (*RSD*) and *strict maximal lotteries* (*SML*) (Aziz, Brandt, and Brill 2013b). *RSD* is the canonical generalization of random dictatorship to weak preferences. It 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). *SML* is a little known class of pairwise SDSs due to Kreweras and Fishburn that return a mixed quasistrict Nash equilibrium of the plurality game. Computing *RSD* was recently shown to be #P-complete (Aziz, Brandt, and Brill 2013a) while *SML* can be computed efficiently using linear programming.

4 Lottery Extensions

In order to reason about the outcomes of SDSs, we need to make assumptions on how agents compare lotteries. A *lottery extension* maps preferences over alternatives to (possibly incomplete) preferences over lotteries. We will now define the lottery extensions considered in this paper. For a more detailed account of the lottery extensions SD, DL, and UL, we refer to Cho (2012).

$$\begin{tikzcd}
R_i^{DL} & R_i^{PC} & R_i^{UL} \\
 & R_i^{SD} \arrow[ul] \arrow[u] \arrow[ur] & \\
 & R_i^{BD} \arrow[u] &
\end{tikzcd}$$

Figure 1: Inclusion relationships between lottery extensions. An arrow denotes set inclusion between two relations, e.g., $R_i^{BD} \subset R_i^{SD}$. DL, PC, and UL are extensions that yield complete preference relations over sets.

Throughout this section, let $R_i \in \mathcal{R}$ be a preference relation and $p, q \in \Delta(A)$.

The first extension we propose is called *bilinear dominance (BD)* and 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, $p\ R_i^{BD}\ q$ iff

$$\forall x, y, x\ P_i\ y\colon p(x)q(y) \geq p(y)q(x). \qquad (BD)$$

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 (Fishburn 1984).

Stochastic dominance (SD) 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. Formally, $p\ R_i^{SD}\ q$ iff

$$\forall x\colon \sum_{y\colon yR_ix} p(y) \geq \sum_{y\colon yR_ix} q(y). \qquad (SD)$$

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.

A novel strengthening of SD, considered in this paper for the first time, is the *pairwise comparison (PC)* extension. The reasoning behind PC is to prefer p to q if the probability that p yields an alternative preferred to the alternative returned by q is at least as large than the other way round.[5] Formally, $p\ R_i^{PC}\ q$ iff

$$\sum_{xR_iy} p(x)q(y) \geq \sum_{xR_iy} q(x)p(y). \qquad (PC)$$

Finally, we define the *downward lexicographic (DL)* extension and the *upward lexicographic (UL)* extension introduced by Cho (2012). According to DL the lottery with higher probability on the top ranked alternative is preferred, in case of equality, the one with higher probability on the second ranked alternative, and so on. Formally, $p\ R_i^{DL}\ q$ if $p = q$ or

$$\exists x\colon (p(x) > q(x) \text{ and } \forall y, y\ R_i\ x\colon p(y) = q(y)). \qquad (DL)$$

[5]Interestingly, this extension may lead to intransitive preferences over lotteries, even when the preferences over alternatives are transitive (Blyth 1972; Fishburn 1988).

Upward lexicographic ordering is dual to the former, i.e., $p\ R_i^{UL}\ q$ if $p = q$ or

$$\exists x\colon (p(x) < q(x) \text{ and } \forall y, x\ R_i\ y\colon p(y) = q(y)). \quad (UL)$$

The following example illustrates the definitions of the extensions. Consider the preference relation $R_i\colon a, b, c$ and lotteries

$$p = 2/3a + 0b + 1/3c \quad \text{and}$$
$$q = 0a + 1b + 0c.$$

Then, $p\ P_i^{PC}\ q$; $p\ P_i^{DL}\ q$; $q\ P_i^{UL}\ p$; $\neg[p\ R_i^{SD}\ q]$; $\neg[q\ R_i^{SD}\ p]$; $\neg[p\ R_i^{BD}\ q]$; and $\neg[q\ R_i^{BD}\ p]$.

The inclusion relationships between the lottery extensions are depicted in Figure 1.

5 Efficiency and Strategyproofness

In this section, we present general definitions of efficiency and strategyproofness which give rise to varying levels of efficiency and strategyproofness depending on which lottery extension is used to define them. The relationships between these concepts are depicted in Figure 2.

Efficiency prescribes that there is no lottery that all agents prefer to the one returned by the SDS. Each lottery extension yields a corresponding notion of efficiency. Let $\mathcal{E}$ be a lottery extension. Given a preference profile R, a lottery p $\mathcal{E}$-*dominates* another lottery 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 $f(R)$. An alternative is *Pareto dominated* if there exists another alternative such that all agents strictly prefer the latter to the former. An SDS is *Pareto optimal* (or *ex post* efficient) if it puts probability zero on all Pareto dominated alternatives. It is well-known that SD-efficiency implies Pareto optimality. Our first theorem, the proof of which is omitted due to space restrictions, shows that Pareto optimality is stronger than BD-efficiency.

Theorem 1. *Pareto optimality implies BD-efficiency but the converse does not hold.*

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. Let $\mathcal{E}$ be a lottery extension. An SDS f is $\mathcal{E}$-*manipulable* if there exist preference profiles R and R' 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. An SDS is *strongly $\mathcal{E}$-strategyproof* if for all R and R' with $R_j = R'_j$ for all $j \neq i$ such that $f(R)\ R_i^{\mathcal{E}}\ f(R')$.[6] For complete lottery extensions (DL, PC, and UL), the weak and the strong notions of strategyproofness coincide.

An SDS f is $\mathcal{E}$-*group-manipulable* if there exists an $S \subseteq N$ and preference profiles R and R' with $R_j = R'_j$ for all $j \in N \setminus S$ such that $f(R')\ P_i^{\mathcal{E}}\ f(R)$ for all $i \in S$. An SDS is $\mathcal{E}$-*group-strategyproof* if it is not $\mathcal{E}$-group-manipulable.

[6] Please note that in some papers (e.g., Bogomolnaia and Moulin 2001) the term "strategyproofness" refers to the strong notion of strategyproofness.

6 Results and Discussion

Recent research has shown that there exists an interesting tradeoff between efficiency and strategyproofness in randomized social choice (Aziz, Brandt, and Brill 2013b). For example, RSD satisfies strong SD-strategyproofness and Pareto optimality, but violates SD-efficiency. SML, on the other hand, satisfies PC-efficiency and ST-strategyproofness, where ST is a weakening of BD, but violates SD-strategyproofness. This section contains four impossibility results that improve our understanding of the interplay between efficiency and strategyproofness and have nontrivial consequences on concrete SDSs such as RSD and SML.

We prove each of these results by reasoning about a set of preference profiles and deriving a contradiction. In particular, the proofs assume a specific number of agents and alternatives, but can be generalized to any (larger) number of agents and alternatives as follows. For more alternatives, we add the additional alternatives as tied for last rank in each agent's preference relation and every preference profile. Notice that we do not leave the domain in case of dichotomous preferences. To show a statement for more agents, we add agents that are indifferent between all alternatives. Both constructions do not affect the set of efficient lotteries and incentives of agents. Hence, the proofs with the same arguments carry through.

Our first result states that efficiency and strategyproofness are incompatible when preferences over lotteries are given by the natural PC extension.

Theorem 2. *There is no anonymous, neutral, PC-efficient, and PC-strategyproof SDS for $n \geq 3$ and $m \geq 3$.*

Proof. Assume for a contradiction that f is an SDS with properties as stated above and consider the following preference profile.

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

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.

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

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 by PC-strategyproofness, $f(R^2) = a$. Next, we consider the profile R^3.

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

PC-efficiency implies that f puts probability 0 on c when applied to R^3, since a Pareto dominates c. If $f(R^3) \neq f(R^2)$, agent 2 would have an incentive to deviate in one direction or the other. Thus, $f(R^3) = a$.

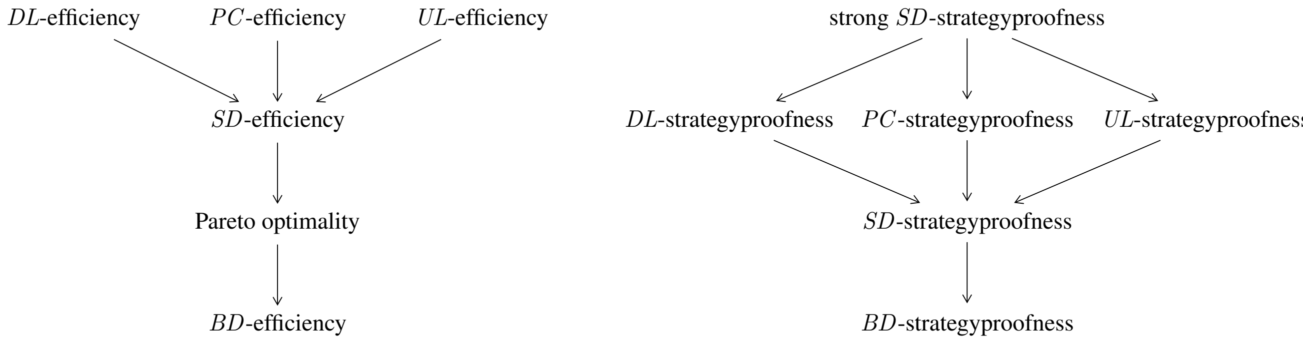

Figure 2: Relationships between efficiency and strategyproofness concepts.

Since we will need it later, we state an observation at this point. By anonymity and neutrality, f has to choose the uniform lottery $1/3a + 1/3b + 1/3c$ in the following profile.

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

Also notice that in this profile agent 1 prefers any lottery with higher probability on c than on b to the uniform lottery according to $(R_1^4)^{PC}$.

Now we consider another preference profile.

$$R_1^5\colon \{a, c\}, b \qquad R_2^5\colon a, b, c \qquad R_3^5\colon b, c, a$$

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

$$R_1^6\colon \{a, c\}, b \qquad R_2^6\colon a, b, c \qquad R_3^6\colon c, b, a$$

Anonymity and neutrality imply that $f(R^6)(a) = f(R^6)(c)$. Any lottery of this form other than $1/2a + 1/2c$ is PC-dominated by the latter. Thus, $f(R^6) = 1/2a + 1/2c$ by PC-efficiency. But agent 3 prefers $1/2a + 1/2c$ to a if his preferences are R_3^5. Hence, 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. A contradiction to PC-strategyproofness. Hence, $f(R^5)(c) \leq f(R^5)(b)$ and thus, by the assumption in the second case, $f(R^5)(b) > 0$.

$$R_1^7\colon \{a, c\}, b \qquad R_2^7\colon a, b, c \qquad R_3^7\colon b, \{a, c\}$$

It follows from $f(R^5)(b) > 0$ that $f(R^7)(b) > 0$, otherwise agent 3 would deviate from R_3^7 to R_3^5. In particular $f(R^7) \neq a$. Finally, consider the following profile.

$$R_1^8\colon \{a, c\}, b \qquad R_2^8\colon a, \{b, c\} \qquad R_3^8\colon b, \{a, c\}$$

By anonymity, $f(R^8) = f(R^3) = a$. But this implies that agent 2 can successfully deviate from R_2^7 to R_2^8 and receive a instead. Hence, the desired contradiction. $\square$

It can be shown that random dictatorship violates PC-efficiency, even when preferences are linear. This emphasizes the efficiency problems of random dictatorship. Previously, it was only known that RSD violates SD-efficiency for weak preferences. Still, PC-efficiency is not unduly restrictive as SML is known to satisfy it.

Next, we prove a similarly negative result for the UL-extension, which only requires two agents.

Theorem 3. *There is no anonymous, UL-efficient, and UL-strategyproof SDS for $n \geq 2$ and $m \geq 3$.*

The proof of this theorem is omitted to meet space constraints.

Interestingly, DL-efficiency—the dual notion of UL-efficiency—*is* compatible with SD-strategyproofness (and hence DL-strategyproofness) because random dictatorship satisfies both conditions when preferences are linear.

The next result is a strengthening of Theorem 1 by Aziz, Brandt, and Brill (2013b) in two respects: it uses a weaker notion of strategyproofness and it holds for the set of all pairwise SDSs rather than only majoritarian SDSs.[7]

Theorem 4. *There is no pairwise, Pareto optimal, and BD-strategyproof SDS for $n \geq 4$ and $m \geq 4$.*

Proof. Let f be a pairwise, Pareto optimal, and BD-strategyproof SDS. We first consider the preference profile R and its (weighted) majority graph depicted in Figure 3 (i).

$$R_1^1\colon a, c, \{b, d\} \qquad R_2^1\colon b, d, \{a, c\}$$

Since f is Pareto optimal and pairwise, $f(R^1) = p = 1/2a + 1/2b$. Now we consider the profile R^2 with majority graph as in Figure 3 (ii).

$$R_1^2\colon a, c, \{b, d\} \qquad R_2^2\colon \{b, d\}, \{a, c\}$$

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

First assume for a contradiction $\lambda > 1/3$. We consider profile R^3.

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

The majority graph of R^3 is as in Figure 3 (iii). Hence, by anonymity and neutrality, $f(R^3) = r = 1/3a + 1/3b + 1/3d$. But $r \; (P_1^2)^{BD} \; q$ if $\lambda > 1/3$, which contradicts BD-strategyproofness of f since voter 1 in R^2 can manipulate by reporting R_1^3 instead of R_1^2.

[7]Note, however, that the proof of Theorem 1 by Aziz, Brandt, and Brill (2013b) only requires linear preferences.

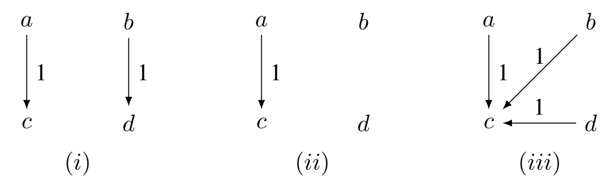

Figure 3: Graphs depicting pairwise comparisons. An edge from x to y is labeled with the number of agents preferring x to y minus the number of agents preferring y to x. All missing edges denote majority ties.

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

$$R_1^4\colon a, c, b, d \qquad R_2^4\colon \{b, d\}, \{a, c\}$$
$$R_3^4\colon a, c, \{b, d\} \qquad R_4^4\colon \{b, d\}, c, a$$

The profile R^4 has a majority graph as in Figure 3 (i), thus $f(R^4) = p$.

$$R_1^5\colon a, c, b, d \qquad R_2^5\colon \{b, d\}, \{a, c\}$$
$$R_3^5\colon a, c, \{b, d\} \qquad R_4^5\colon d, b, c, a$$

The majority graph of R^5 is as in Figure 3 (ii) and hence $f(R^5) = q$. But then voter 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 profile R^6.

$$R_1^6\colon a, c, \{b, d\} \qquad R_2^6\colon \{b, d\}, \{a, c\}$$
$$R_3^6\colon \{b, c, d\}, a \qquad R_4^6\colon a, \{b, c, d\}$$

Its majority graph is as in Figure 3 (ii) and therefore $f(R^6) = q$. We consider one last profile.

$$R_1^7\colon a, c, \{b, d\} \qquad R_2^7\colon \{b, d\}, \{a, c\}$$
$$R_3^7\colon \{b, d\}, c, a \qquad R_4^7\colon a, \{b, c, d\}$$

The majority graph of this profile is as in Figure 3 (i) which implies $f(R^7) = r$. But $r\ (P_3^6)^{BD}\ q$ if $\lambda < 1/3$. Thus, agent 3 in R^6 benefits from reporting R_3^7 instead of R_3^6. In any case, we found a successful manipulation, contradicting BD-strategyproofness of f. ☐

As a consequence of the previous theorem, the pairwise SDS SML does not satisfy BD-strategyproofness and, as a matter of fact, *no* reasonable pairwise SDS can.

For our final result we again consider the rather weak notion of BD-strategyproofness, but also consider deviations by *groups* of agents. It turns out that BD-group-strategyproofness is already incompatible with Pareto optimality when also requiring anonymity and neutrality.[8]

[8]Bogomolnaia, Moulin, and Stong (2005) (Proposition 3) proved that for $n \geq 4$ and $m \geq 6$, there exists no anonymous, neutral, Pareto optimal, and SD-group-strategyproof SDS for dichotomous preferences. We strengthen their result by weakening SD-group-strategyproofness to the significantly weaker notion of BD-group-strategyproofness and by using less alternatives and agents.

Theorem 5. *There is no anonymous, neutral, Pareto-optimal, and BD-group-strategyproof SDS for $n \geq 3$ and $m \geq 3$, even when preferences are dichotomous.*

Proof. Assume for contradiction that f is an SDS with the properties as stated. Consider a setting with three agents and three alternatives and the following preference profile.

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

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

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

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 $f(R^2) = 1/3a + 1/3b + 1/3c$.

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

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

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

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. ☐

For the stronger (but less reasonable) notion of group-strategyproofness in which only one of the deviating agents has to be strictly better off, we were able to show the previous impossibility even without requiring anonymity and neutrality. The proof is omitted due to limited space.

Theorem 5 implies that RSD violates BD-group-strategyproofness. As a matter of fact, both RSD and SML only satisfy the rather weak notion of ST-group-strategyproofness where ST is a weakening of BD introduced by Aziz, Brandt, and Brill (2013b). Put in a nutshell, RSD does better in terms of individual strategyproofness (strong SD-strategyproofness vs. ST-strategyproofness) while SML is more efficient (PC-efficiency vs. Pareto optimality).

Acknowledgments

This material is based upon work supported by Deutsche Forschungsgemeinschaft under grants BR 2312/7-2 and BR 2312/10-1 and by NICTA, which is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program. The authors thank Markus Brill, Hervé Moulin, Shasi Nandeibam, and Arunava Sen for helpful discussions.

References

Aziz, H.; Brandt, F.; and Brill, M. 2013a. The computational complexity of random serial dictatorship. *Economics Letters* 121(3):341–345.

Aziz, H.; Brandt, F.; and Brill, M. 2013b. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 455–462. IFAAMAS.

Aziz, H.; Brandt, F.; and Seedig, H. G. 2013. Computing desirable partitions in additively separable hedonic games. *Artificial Intelligence* 195:316–334.

Aziz, H. 2013. Maximal Recursive Rule: A New Social Decision Scheme. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 34–40.

Barberà, S. 1979. A note on group strategy-proof decision schemes. *Econometrica* 47(3):637–640.

Barberà, S. 2010. Strategy-proof social choice. In Arrow, K. J.; Sen, A. K.; and Suzumura, K., eds., *Handbook of Social Choice and Welfare*, volume 2. Elsevier. chapter 25, 731–832.

Birrell, E., and Pass, R. 2011. Approximately strategy-proof voting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 67–72.

Blyth, C. R. 1972. Some probability paradoxes in choice from among random alternatives. *Journal of the American Statistical Association* 67(338):366–373.

Bogomolnaia, A., and Moulin, H. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100(2):295–328.

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

Bouveret, S., and Lang, J. 2008. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. *Journal of Artificial Intelligence Research* 32(1):525–564.

Chatterji, S.; Sen, A.; and Zeng, H. 2014. Random dictatorship domains. *Games and Economic Behavior*. Forthcoming.

Cho, W. J. 2012. Probabilistic assignment: A two-fold axiomatic approach. Unpublished manuscript.

Conitzer, V., and Sandholm, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, 627–634. AAAI Press.

Elkind, E., and Wooldridge, M. 2009. Hedonic coalition nets. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 417–424.

Fishburn, P. C. 1972. Lotteries and social choice. *Journal of Economic Theory* 5:189–207.

Fishburn, P. C. 1982. Nontransitive measurable utility. *Journal of Mathematical Psychology* 26(1):31–67.

Fishburn, P. C. 1984. Dominance in SSB utility theory. *Journal of Economic Theory* 34(1):130–148.

Fishburn, P. C. 1988. *Nonlinear preference and utility theory*. The Johns Hopkins University Press.

Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41:587–602.

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

Postlewaite, A., and Schmeidler, D. 1986. Strategic behaviour and a notion of ex ante efficiency in a voting model. *Social Choice and Welfare* 3:37–49.

Procaccia, A. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 836–841. AAAI Press.

Satterthwaite, M. A. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187–217.

Service, T. C., and Adams, J. A. 2012. Strategyproof approximations of distance rationalizable voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 569–576.

Sönmez, T., and Ünver, M. U. 2011. Matching, allocation, and exchange of discrete resources. In Benhabib, J.; Jackson, M. O.; and Bisin, A., eds., *Handbook of Social Economics*, volume 1. Elsevier. chapter 17, 781–852.

Stone, P. 2011. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press.

Walsh, T., and Xia, L. 2012. Lot-based voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 603–610.

Young, H. P. 1974. An axiomatization of Borda's rule. *Journal of Economic Theory* 9:43–52.

Zeckhauser, R. 1969. Majority rule with lotteries on alternatives. *Quarterly Journal of Economics* 83(4):696–703.

Zwicker, W. S. 1991. The voter's paradox, spin, and the Borda count. *Mathematical Social Sciences* 22:187–227.