

Relaxed Notions of Condorcet-Consistency and Efficiency for Strategyproof Social Decision Schemes

Felix Brandt
 Technical University of Munich
 Munich, Germany
 brandtf@in.tum.de

Patrick Lederer
 Technical University of Munich
 Munich, Germany
 ledererp@in.tum.de

René Romen
 Technical University of Munich
 Munich, Germany
 rene.romen@tum.de

ABSTRACT

Social decision schemes (SDSs) map the preferences of a group of voters over some set of m alternatives to a probability distribution over the alternatives. A seminal characterization of strategyproof SDSs by Gibbard implies that there are no strategyproof Condorcet extensions and that only random dictatorships satisfy *ex post* efficiency and strategyproofness. The latter is known as the *random dictatorship theorem*. We relax Condorcet-consistency and *ex post* efficiency by introducing a lower bound on the probability of Condorcet winners and an upper bound on the probability of Pareto-dominated alternatives, respectively. We then show that the SDS that assigns probabilities proportional to Copeland scores is the only anonymous, neutral, and strategyproof SDS that can guarantee the Condorcet winner a probability of at least $2/m$. Moreover, no strategyproof SDS can exceed this bound, even when dropping anonymity and neutrality. Secondly, we prove a continuous strengthening of Gibbard’s random dictatorship theorem: the less probability we put on Pareto-dominated alternatives, the closer to a random dictatorship is the resulting SDS. Finally, we show that the only anonymous, neutral, and strategyproof SDSs that maximize the probability of Condorcet winners while minimizing the probability of Pareto-dominated alternatives are mixtures of the uniform random dictatorship and the randomized Copeland rule.

KEYWORDS

Randomized Social Choice; Social Decision Schemes; Strategyproofness; Condorcet-consistency; *ex post* efficiency

ACM Reference Format:

Felix Brandt, Patrick Lederer, and René Romen. 2022. Relaxed Notions of Condorcet-Consistency and Efficiency for Strategyproof Social Decision Schemes. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

Multi-agent systems are often faced with problems of collective decision making: how to find a group decision given the preferences of multiple individual agents. These problems, which have been traditionally studied by economists and mathematicians, are of increasing interest to computer scientists who employ the formalisms of social choice theory to analyze computational multi-agent systems [see, e.g., 8, 9, 26, 30].

A pervasive phenomenon in collective decision making is strategic manipulation: voters may be better off by lying about their preferences than reporting them truthfully. This is problematic since all desirable properties of a voting rule are in doubt when voters act dishonestly. Thus, it is important that voting rules incentivize voters to report their true preferences. Unfortunately, Gibbard [19] and Satterthwaite [28] have shown independently that dictatorships are the only non-imposing voting rules that are immune to strategic manipulations. However, these voting rules are unacceptable for most applications because they invariably return the most preferred alternative of a fixed voter. A natural question is whether more positive results can be obtained when allowing for randomization. Gibbard [20] hence introduced *social decision schemes* (SDSs), which map the preferences of the voters to a lottery over the alternatives and defined SDSs to be *strategyproof* if no voter can obtain more expected utility for any utility representation that is consistent with his ordinal preference relation. He then gave a complete characterization of strategyproof SDSs in terms of convex combinations of two types of restricted SDSs, so-called unilaterals and duples. An important consequence of this result is the *random dictatorship theorem*: random dictatorships are the only *ex post* efficient and strategyproof SDSs. Random dictatorships are convex combinations of dictatorships, i.e., each voter is selected with some fixed probability and the top choice of the chosen voter is returned. In contrast to deterministic dictatorships, the uniform random dictatorship, in which every agent is picked with the same probability, enjoys a high degree of fairness and is in fact used in many subdomains of social choice [see, e.g., 1, 12]. As a consequence of these observations, Gibbard’s theorem has been the point of departure for a lot of follow-up work. In addition to several alternative proofs of the theorem [e.g., 14, 24, 31], there have been extensions with respect to manipulations by groups [4], cardinal preferences [e.g., 16, 23, 25], weaker notions of strategyproofness [e.g., 2, 5, 7, 29], and restricted domains of preference [e.g., 11, 15].

Random dictatorships suffer from the disadvantage that they do not allow for compromise. For instance, suppose that voters strongly disagree on the best alternative, but have a common second best alternative. In such a scenario, it seems reasonable to choose the second best alternative but random dictatorships do not allow for this compromise. On a formal level, this observation is related to the fact that random dictatorships violate *Condorcet-consistency*, which demands that an alternative that beats all other alternatives in pairwise majority comparisons should be selected. Motivated by this observation, we analyze the limitations of strategyproof SDSs by relaxing two classic conditions: Condorcet-consistency and *ex post* efficiency. To this end, we say that an SDS is α -*Condorcet-consistent* if a Condorcet winner always receives a probability of

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

at least α and β -*ex post efficient* if a Pareto-dominated alternative always receives a probability of at most β . Moreover, we say a strategyproof SDS is γ -*randomly dictatorial* if it can be represented as a convex combination of two strategyproof SDSs, one of which is a random dictatorship that will be selected with probability γ . All of these axioms are discussed in more detail in Section 2.2.

Building on an alternative characterization of strategyproof SDSs by Barberà [3], we then show the following results (m is the number of alternatives and n the number of voters):

- Let $m, n \geq 3$. There is no strategyproof SDS that satisfies α -Condorcet-consistency for $\alpha > 2/m$. Moreover, the *randomized Copeland rule*, which assigns probabilities proportional to Copeland scores, is the only strategyproof SDS that satisfies anonymity, neutrality, and $2/m$ -Condorcet-consistency.
- Let $0 \leq \epsilon \leq 1$ and $m \geq 3$. Every strategyproof SDS that is $\frac{1-\epsilon}{m}$ -*ex post efficient* is γ -randomly dictatorial for $\gamma \geq \epsilon$. If we additionally require anonymity, neutrality, and $m \geq 4$, then only mixtures of the uniform random dictatorship and the uniform lottery satisfy this bound tightly.
- Let $m \geq 4$ and $n \geq 5$. No strategyproof SDS that is α -Condorcet-consistent is β -*ex post efficient* for $\beta < \frac{m-2}{m-1}\alpha$. If we additionally require anonymity and neutrality, then only mixtures of the uniform random dictatorship and the randomized Copeland rule satisfy $\beta = \frac{m-2}{m-1}\alpha$.

The first statement characterizes the randomized Copeland rule as the “most Condorcet-consistent” SDS that satisfies strategyproofness, anonymity, and neutrality. In fact, no strategyproof SDS can guarantee more than $2/m$ probability to the Condorcet winner, even when dropping anonymity and neutrality. The second point can be interpreted as a continuous strengthening of Gibbard’s random dictatorship theorem: the less probability we put on Pareto-dominated alternatives, the more randomly dictatorial is the resulting SDS. In particular, this theorem indicates that we cannot find appealing strategyproof SDSs by allowing that Pareto-dominated alternatives gain a small probability since the resulting SDS will be very similar to random dictatorships. The last statement identifies a tradeoff between α -Condorcet-consistency and β -*ex post efficiency*: the more probability a strategyproof SDS guarantees to the Condorcet winner, the less efficient it is. Thus, we can either maximize α for α -Condorcet-consistency or minimize β for β -*ex post efficiency* of a strategyproof SDS, which again highlights the central roles of the randomized Copeland rule and random dictatorships.

2 THE MODEL

Let $N = \{1, 2, \dots, n\}$ be a finite set of voters and let $A = \{a, b, \dots\}$ be a finite set of m alternatives. Every voter i has a *preference relation* $>_i$, which is an anti-symmetric, complete, and transitive binary relation on A . We write $x >_i y$ if voter i prefers x strictly to y and $x \geq_i y$ if $x >_i y$ or $x = y$. The set of all preference relations is denoted by \mathcal{R} . A *preference profile* $R \in \mathcal{R}^n$ contains the preference relation of each voter $i \in N$. We define the *supporting size* for x against y in the preference profile R by $n_{xy}(R) = |\{i \in N : x >_i y\}|$.

Given a preference profile, we are interested in the winning chance of each alternative. We therefore analyze social decision schemes (SDSs), which map each preference profile to a lottery over the alternatives. A *lottery* p is a probability distribution over the

set of alternatives A , i.e., it assigns each alternative x a probability $p(x) \geq 0$ such that $\sum_{x \in A} p(x) = 1$. The set of all lotteries over A is denoted by $\Delta(A)$. Formally, a *social decision scheme (SDS)* is a function $f : \mathcal{R}^n \rightarrow \Delta(A)$. We denote with $f(R, x)$ the probability assigned to alternative x by f for the preference profile R .

Since there is a huge number of SDSs, we now discuss axioms formalizing desirable properties of these functions. Two basic fairness conditions are anonymity and neutrality. Anonymity requires that voters are treated equally. Formally, an SDS f is *anonymous* if $f(R) = f(\pi(R))$ for all preference profiles R and permutations $\pi : N \rightarrow N$. Here, $R' = \pi(R)$ denotes the profile with $>_{\pi(i)} = >_i$ for all voters $i \in N$. *Neutrality* guarantees that alternatives are treated equally and formally requires for an SDS f that $f(R, x) = f(\tau(R), \tau(x))$ for all preference profiles R and permutations $\tau : A \rightarrow A$. This time, $R' = \tau(R)$ is the profile derived by permuting the alternatives in R according to τ , i.e., $\tau(x) >'_i \tau(y)$ if and only if $x >_i y$ for all alternatives $x, y \in A$ and voters $i \in N$.

2.1 Stochastic Dominance and Strategyproofness

This paper is concerned with strategyproof SDSs, i.e., social decision schemes in which voters cannot benefit by lying about their preferences. In order to make this formally precise, we need to specify how voters compare lotteries. To this end, we leverage the well-known notion of stochastic dominance: a voter i (weakly) prefers a lottery p to another lottery q , written as $p \geq_i q$, if $\sum_{y \in A: y >_i x} p(y) \geq \sum_{y \in A: y >_i x} q(y)$ for every alternative $x \in A$. Less formally, a voter prefers a lottery p weakly to a lottery q if, for every alternative $x \in A$, p returns a better alternative than x with at least as much probability as q . Stochastic dominance does not induce a complete order on the set of lotteries, i.e., there are lotteries p and q such that a voter i neither prefers p to q nor q to p .

Based on stochastic dominance, we can now formalize strategyproofness. An SDS f is *strategyproof* if $f(R) \geq_i f(R')$ for all preference profiles R and R' and voters $i \in N$ such that $>_j = >'_j$ for all $j \in N \setminus \{i\}$. Less formally, strategyproofness requires that every voter prefers the lottery obtained by voting truthfully to any lottery that he could obtain by voting dishonestly. Conversely, we call an SDS f *manipulable* if it is not strategyproof. While there are other ways to compare lotteries with each other, stochastic dominance is the most common one [see, e.g., 2, 3, 6, 17, 20]. This is mainly due to the fact that $p \geq_i q$ implies that the expected utility of p is at least as high as the expected utility of q for every vNM utility function that is ordinally consistent with voter i ’s preferences. Hence, if an SDS is strategyproof, no voter can manipulate regardless of his exact utility function [see, e.g., 7, 29]. This observation immediately implies that the *convex combination* $h = \lambda f + (1 - \lambda)g$ (for some $\lambda \in [0, 1]$) of two strategyproof SDSs f and g is again strategyproof: a manipulator who obtains more expected utility from $h(R')$ than $h(R)$ prefers $f(R')$ to $f(R)$ or $g(R')$ to $g(R)$.

Gibbard [20] shows that every strategyproof SDS can be represented as convex combinations of unilaterals and duples.¹ The terms “unilaterals” and “duples” refer here to special classes of SDSs: a *unilateral* is a strategyproof SDS that only depends on the

¹In order to simplify the exposition, we slightly modified Gibbard’s terminology by requiring that duples and unilaterals have to be strategyproof.

preferences of a single voter i , i.e., $f(R) = f(R')$ for all preference profiles R and R' such that $\succ_i = \succ'_i$. A *duplex*, on other hand, is a strategyproof SDS that only chooses between two alternatives x and y , i.e., $f(R, z) = 0$ for all preference profiles R and alternatives $z \in A \setminus \{x, y\}$.

Theorem 1 (Gibbard [20]). *An SDS is strategyproof if and only if it can be represented as a convex combination of unilaterals and duples.*

Since duples and unilaterals are by definition strategyproof, Theorem 1 only states that strategyproof SDSs can be decomposed into a mixture of strategyproof SDSs, each of which must be of a special type. In order to circumvent this restriction, Gibbard proves another characterization of strategyproof SDSs.

Theorem 2 (Gibbard [20]). *An SDS is strategyproof if and only if it is non-pervasive and localized.*

Non-perversity and localizedness are two axioms describing the behavior of an SDS. For defining these axioms, we denote with $R^{i:yx}$ the profile derived from R by only reinforcing y against x in voter i 's preference relation. Note that this requires that $x \succ_i y$ and that there is no alternative $z \in A$ such that $x \succ_i z \succ_i y$. Then, an SDS f is *non-pervasive* if $f(R^{i:yx}, y) \geq f(R, y)$ for all preference profiles R , voters $i \in N$, and alternatives $x, y \in A$. Moreover, an SDS is *localized* if $f(R^{i:yx}, z) = f(R, z)$ for all preference profiles R , voters $i \in N$, and distinct alternatives $x, y, z \in A$. Intuitively, non-perversity—which is now often referred to as monotonicity—requires that the probability of an alternative only increases if it is reinforced, and localizedness that the probability of an alternative does not depend on the order of the other alternatives. Together, Theorem 1 and Theorem 2 show that each strategyproof SDS can be represented as a mixture of unilaterals and duples, each of which is non-pervasive and localized.

Since Gibbard's results can be quite difficult to work with, we now state another characterization of strategyproof SDSs due to Barberà [3]. Barberà has shown that every strategyproof SDS that satisfies anonymity and neutrality can be represented as a convex combination of a supporting size SDS and a point voting SDS. A *point voting SDS* is defined by a scoring vector (a_1, a_2, \dots, a_m) that satisfies $a_1 \geq a_2 \geq \dots \geq a_m \geq 0$ and $\sum_{i \in \{1, \dots, m\}} a_i = \frac{1}{n}$. The probability assigned to an alternative x by a point voting SDS f is $f(R, x) = \sum_{i \in N} a_{|\{y \in A: y \succ_i x\}|}$. Furthermore, *supporting size SDSs* also rely on a scoring vector $(b_n, b_{n-1}, \dots, b_0)$ with $b_n \geq b_{n-1} \geq \dots \geq b_0 \geq 0$ and $b_i + b_{n-i} = \frac{2}{m(m-1)}$ for all $i \in \{0, \dots, n\}$ to compute the outcome. The probability assigned to an alternative x by a supporting size SDS f is then $f(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)}$. Note that point voting SDSs can be seen as a generalization of (deterministic) positional scoring rules and supporting size SDSs can be seen as a variant of Fishburn's C2 functions [18].

Theorem 3 (Barberà [3]). *An SDS is anonymous, neutral, and strategyproof if and only if it can be represented as a convex combination of a point voting SDS and a supporting size SDS.*

Many well-known SDSs can be represented as point voting SDSs or supporting size SDSs. For example, the *uniform random dictatorship* f_{RD} , which chooses one voter uniformly at random and returns his best alternative, is the point voting SDS defined by the scoring vector $(\frac{1}{n}, 0, \dots, 0)$. An instance of a supporting size SDS is the

randomized Copeland rule f_C , which assigns probabilities proportional to the Copeland scores $c(x, R) = |\{y \in A \setminus \{x\}: n_{xy}(R) > n_{yx}(R)\}| + \frac{1}{2}|\{y \in A \setminus \{x\}: n_{xy}(R) = n_{yx}(R)\}|$. This SDS is the supporting size SDS defined by the vector $b = (b_n, b_{n-1}, \dots, b_0)$, where $b_i = \frac{2}{m(m-1)}$ if $i > \frac{n}{2}$, $b_i = \frac{1}{m(m-1)}$ if $i = \frac{n}{2}$, and $b_i = 0$ otherwise. Furthermore, there are SDSs that can be represented both as point voting SDSs and supporting size SDSs. An example is the *randomized Borda rule* f_B , which randomizes proportional to the Borda scores of the alternatives. This SDS is the point voting SDS defined by the vector $(\frac{2(m-1)}{nm(m-1)}, \frac{2(m-2)}{nm(m-1)}, \dots, \frac{2}{nm(m-1)}, 0)$ and equivalently the supporting size SDS defined by the vector $(\frac{2n}{nm(m-1)}, \frac{2(n-1)}{nm(m-1)}, \dots, \frac{2}{nm(m-1)}, 0)$. Both the randomized Copeland rule and the randomized Borda rule were rediscovered several times by authors who were apparently unaware of Barberà's work [see 13, 21, 22, 27].

2.2 Relaxing Classic Axioms

The goal of this paper is to identify attractive strategyproof SDSs other than random dictatorships by relaxing classic axioms from social choice theory. In more detail, we investigate how much probability can be guaranteed to Condorcet winners and how little probability must be assigned to Pareto-dominated alternatives by strategyproof SDSs. In the following we formalize these ideas using α -Condorcet-consistency and β -*ex post* efficiency.

Let us first consider β -*ex post* efficiency, which is based on Pareto-dominance. An alternative x *Pareto-dominates* another alternative y in a preference profile R if $x \succ_i y$ for all $i \in N$. The standard notion of *ex post efficiency* then formalizes that Pareto-dominated alternatives should have no winning chance, i.e., $f(R, x) = 0$ for all preference profiles R and alternatives x that are Pareto-dominated in R . As first shown by Gibbard, random dictatorships are the only strategyproof SDSs that satisfy *ex post* efficiency. These SDSs choose each voter with a fixed probability and return his best alternative as winner. However, this result breaks down once we allow that Pareto-dominated alternatives can have a non-zero chance of winning $\beta > 0$. For illustrating this point, consider a random dictatorship d and another strategyproof SDS g . Then, the SDS $f^* = (1-\beta)d + \beta g$ is strategyproof for every $\beta \in (0, 1]$ and no random dictatorship, but assigns a probability of at most β to Pareto-dominated alternatives. We call the last property β -*ex post* efficiency: an SDS f is β -*ex post* efficient if $f(R, x) \leq \beta$ for all preference profiles R and alternatives x that are Pareto-dominated in R .

A natural generalization of the random dictatorship theorem is to ask which strategyproof SDSs satisfy β -*ex post* efficiency for small values of β . If β is sufficiently small, β -*ex post* efficiency may be quite acceptable. As we show, the random dictatorship theorem is quite robust in the sense that all SDSs that satisfy β -*ex post* efficiency for $\beta < \frac{1}{m}$ are similar to random dictatorships. In order to formalize this observation, we introduce γ -randomly dictatorial SDSs: a strategyproof SDS f is γ -*randomly dictatorial* if $\gamma \in [0, 1]$ is the maximal value such that f can be represented as $f = \gamma d + (1-\gamma)g$, where d is a random dictatorship and g is another strategyproof SDS. In particular, we require that g is strategyproof as otherwise, SDSs that seem “non-randomly dictatorial” are not 0-randomly dictatorial. For instance, the uniform lottery f_U , which

1	1	1		1	1	1
a	b	c		a	b	c
c	c	a		b	c	a
b	a	b		c	a	b
R				R'		

Figure 1: Condorcet-consistent SDSs violate strategyproofness when $m = n = 3$. Due to the symmetry of R' , we may assume without loss of generality that $f(R', a) > 0$. Since f is Condorcet-consistent, it holds that $f(R, c) = 1$. Thus, voter 1 can manipulate by swapping c and b in R .

always assigns probability $\frac{1}{m}$ to all alternatives, is not 0-randomly dictatorial if g is not required to be strategyproof because it can be represented as $f_U = \frac{1}{m}d_i + \frac{m-1}{m}g$, where d_i is the dictatorial SDS of voter i and g is the SDS that randomizes uniformly over all alternatives but voter i 's favorite one. Moreover, it should be mentioned that the maximality of γ implies that g is 0-randomly dictatorial if $\gamma < 1$. Otherwise, we could also represent g as a mixture of a random dictatorship and some other strategyproof SDS h , which means that f is γ' -randomly dictatorial for $\gamma' > \gamma$.

For a better understanding of γ -randomly dictatorial SDSs, we provide next a characterization of these SDSs. Recall for the following lemma that $R^{i:y^x}$ denotes the profile derived from R by only reinforcing y against x in voter i 's preference relation.

Lemma 1. *A strategyproof SDS f is γ -randomly dictatorial if and only if there are non-negative values $\gamma_1, \dots, \gamma_n$ such that:*

- i) $\sum_{i \in N} \gamma_i = \gamma$.
- ii) $f(R^{i:y^x}, y) - f(R, y) \geq \gamma_i$ for all alternatives $x, y \in A$, voters $i \in N$, and preference profiles R in which voter i prefers x the most and y the second most.
- iii) for every voter $i \in N$, there are alternatives $x, y \in A$ and a profile R such that voter i prefers x the most and y the second most in R , and $f(R^{i:y^x}, y) - f(R, y) = \gamma_i$.

The proof of this lemma can be found in the extended version [10]. Lemma 1 gives an intuitive interpretation of γ -randomly dictatorial SDSs: this axiom only requires that there are voters who always increase the winning probability of an alternative by at least γ_i if they reinforce it to the first place. Hence, for small values of γ , this axiom is desirable as it only formulates a variant of strict monotonicity. However, for larger values of γ , γ -randomly dictatorial SDSs become more similar to random dictatorships. Furthermore, the proof of Lemma 1 shows that the decomposition of γ -randomly dictatorial SDSs is completely determined by the values $\gamma_1, \dots, \gamma_n$: given these values for an strategyproof SDS f , it can be represented as $f = \sum_{i \in N} \gamma_i d_i + (1 - \sum_{i \in N} \gamma_i)g$, where g is a strategyproof SDS and d_i the dictatorial SDS of voter i .

Finally, we introduce α -Condorcet-consistency. To this end, we first define the notion of a Condorcet winner. A *Condorcet winner* is an alternative x that wins every majority comparison according to preference profile R , i.e., $n_{xy}(R) > n_{yx}(R)$ for all $y \in A \setminus \{x\}$. *Condorcet-consistency* demands that $f(R, x) = 1$ for all preference profiles R and alternatives x such that x is the Condorcet winner

Table 1: Values of α , β , and γ for which specific SDSs are α -Condorcet-consistent, β -ex post efficient, and γ -randomly dictatorial. Each row shows the values of α , β , and γ for which a specific SDS satisfies the corresponding axioms. f_{RD} abbreviates the uniform random dictatorship, f_U the uniform lottery, f_B the randomized Borda rule, and f_C the randomized Copeland rule.

SDS	α -Condorcet -consistency	β -ex post efficiency	γ -random dictatorship
f_{RD}	0	0	1
f_U	$\frac{1}{m}$	$\frac{1}{m}$	0
f_B	$\frac{1}{m} + \frac{2-(n \bmod 2)}{mn}$	$\frac{2(m-2)}{m(m-1)}$	$\frac{2}{m(m-1)}$
f_C	$\frac{2}{m}$	$\frac{2(m-2)}{m(m-1)}$	0

in R . Unfortunately, Condorcet-consistency is in conflict with strategyproofness, which can easily be derived from Gibbard's random dictatorship theorem. A simple two-profile proof for this fact when $m = n = 3$ is given in Figure 1. To circumvent this impossibility, we relax Condorcet-consistency: instead of requiring that the Condorcet winner always obtains probability 1, we only require that it receives a probability of at least α . This idea leads to α -Condorcet-consistency: an SDS f satisfies this axiom if $f(R, x) \geq \alpha$ for all profiles R and alternatives $x \in A$ such that x is the Condorcet winner in R . For small values of α , this axiom is clearly compatible with strategyproofness and therefore, we are interested in the maximum value of α such that there are α -Condorcet-consistent and strategyproof SDSs.

For a better understanding of α -Condorcet-consistency, β -ex post efficiency, and γ -random dictatorships, we discuss some of the values in Table 1 as examples. The uniform random dictatorship is 1-randomly dictatorial and 0-ex post efficient by definition. Moreover, it is 0-Condorcet-consistent because a Condorcet winner may not be top-ranked by any voter. The randomized Borda rule is $\frac{2(m-2)}{m(m-1)}$ -ex post efficient because it assigns this probability to an alternative that is second-ranked by every voter. Moreover, it is $\frac{2}{m(m-1)}$ -randomly dictatorial as we can represent it as $\frac{2}{m(m-1)}f_{RD} + \left(1 - \frac{2}{m(m-1)}\right)g$, where f_{RD} is the uniform random dictatorship and g is the point voting SDS defined by the scoring vector $\left(\frac{2(m-2)}{n(m(m-1)-2)}, \frac{2(m-2)}{n(m(m-1)-2)}, \frac{2(m-3)}{n(m(m-1)-2)}, \dots, 0\right)$. Finally, the randomized Copeland rule is 0-randomly dictatorial because there is for every voter a profile in which he can swap his two best alternatives without affecting the outcome. Moreover, it is $\frac{2}{m}$ -Condorcet-consistent because a Condorcet winner x satisfies that $n_{xy}(R) > \frac{n}{2}$ for all $y \in A \setminus \{x\}$ and hence, $f_C(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)} = (m-1)\frac{2}{m(m-1)} = \frac{2}{m}$. Note that Table 1 also contains a row corresponding to the uniform lottery. We consider this SDS as a threshold with respect to α -Condorcet-consistency and β -ex post efficiency because we can compute the uniform lottery without knowledge about the voters' preferences. Hence, if an SDS performs worse than the uniform lottery with respect to α -Condorcet-consistency or β -ex post efficiency, we could also dismiss the voters' preferences.

3 RESULTS

In this section, we present our results about the α -Condorcet-consistency and the β -*ex post* efficiency of strategyproof SDSs. First, we prove that no strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$ and that the randomized Copeland rule f_C is the only anonymous, neutral, and strategyproof SDS that satisfies α -Condorcet-consistency for $\alpha = \frac{2}{m}$. Moreover, we show that every $\frac{1-\epsilon}{m}$ -*ex post* efficient and strategyproof SDS is γ -randomly dictatorial for $\gamma \geq \epsilon$. This statement can be seen as a continuous generalization of the random dictatorship theorem and implies, for instance, that every 0-randomly dictatorial and strategyproof SDS can only satisfy β -*ex post* efficiency for $\beta \geq \frac{1}{m}$, i.e., such SDSs are at least as inefficient as the uniform lottery. Even more, when additionally imposing anonymity and neutrality, we prove that only mixtures of the uniform random dictatorship and the uniform lottery satisfy this bound tightly, which shows that relaxing *ex post* efficiency does not allow for appealing SDSs. In the last theorem, we identify a tradeoff between Condorcet-consistency and *ex post* efficiency: no strategyproof SDS that satisfies α -Condorcet consistency is β -*ex post* efficient for $\beta < \frac{m-2}{m-1}\alpha$. We derive these results through a series of lemmas. Because of space restrictions, the proofs of all lemmas and Theorem 5 are deferred to an extended version of this paper [10] and we only present short proof sketches instead.

3.1 α -Condorcet-consistency

As discussed in Section 2.2, Condorcet-consistent SDSs violate strategyproofness. Therefore, we analyze the maximal α such that α -Condorcet-consistency and strategyproofness are compatible. Our results show that strategyproofness only allows for a small degree of Condorcet-consistency: we prove that no strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$. This bound is tight as the randomized Copeland rule f_C is $\frac{2}{m}$ -Condorcet-consistent, which means that it is one of the “most Condorcet-consistent” strategyproof SDSs. Even more, we can turn this observation in a characterization of f_C by additionally requiring anonymity and neutrality: the randomized Copeland rule is the only strategyproof SDS that satisfies $\frac{2}{m}$ -Condorcet-consistency, anonymity, and neutrality.

For proving these results, we derive next a number of lemmas. As first step, we show in Lemma 2 that we can use a strategyproof and α -Condorcet-consistent SDS to construct another strategyproof SDS that satisfies anonymity, neutrality, and α -Condorcet-consistency for the same α .

Lemma 2. *If a strategyproof SDS satisfies α -Condorcet-consistency for some $\alpha \in [0, 1]$, there is also a strategyproof SDS that satisfies anonymity, neutrality, and α -Condorcet-consistency for the same α .*

The central idea in the proof of Lemma 2 is the following: if there is a strategyproof and α -Condorcet-consistent SDS f , then the SDS $f^{\pi\tau}(R, x) = f(\tau(\pi(R)), \tau(x))$ is also strategyproof and α -Condorcet-consistent for all permutations $\pi : N \rightarrow N$ and $\tau : A \rightarrow A$. Since mixtures of strategyproof and α -Condorcet-consistent SDSs are also strategyproof and α -Condorcet-consistent, we can therefore construct an SDS that satisfies all requirements of the lemma by averaging over all permutations on N and A . More formally, the SDS $f^* = \frac{1}{m!n!} \sum_{\pi \in \Pi} \sum_{\tau \in \Gamma} f^{\pi\tau}$ (where Π denotes the

set of all permutations on N and Γ the set of all permutations on A) meets all criteria of the lemma.

Due to Lemma 2, we investigate next the α -Condorcet-consistency of strategyproof SDSs that satisfy anonymity and neutrality. The reason for this is that this lemma turns an upper bound on α for these SDSs into an upper bound for all strategyproof SDSs. Since Theorem 3 shows that every strategyproof, anonymous, and neutral SDS can be decomposed in a point voting SDS and a supporting size SDS, we investigate these two classes separately in the following two lemmas. First, we bound the α -Condorcet-consistency of point voting SDSs.

Lemma 3. *No point voting SDS is α -Condorcet-consistent for $\alpha \geq \frac{2}{m}$ if $n \geq 3$ and $m \geq 3$.*

The proof of this lemma relies on the observation that there can be $\lceil \frac{m}{2} \rceil$ Condorcet winner candidates, i.e., alternatives x that can be made into the Condorcet winner by keeping x at the same position in the preferences of every voter and only reordering the other alternatives. Since reordering the other alternatives does not affect the probability of x in a point voting SDS, it follows that every Condorcet winner candidate has a probability of at least α . Hence, we derive that $\alpha \leq \frac{1}{\lceil \frac{m}{2} \rceil} \leq \frac{2}{m}$ and a slightly more involved argument shows that the inequality is strict.

The last ingredient for the proof of Theorem 4 is that no supporting size SDS can assign a probability of more than $\frac{2}{m}$ to any alternative. This immediately implies that no supporting size SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$.

Lemma 4. *No supporting size SDS can assign more than $\frac{2}{m}$ probability to an alternative.*

The proof of this lemma follows straightforwardly from the definition of supporting size SDSs. Each such SDS is defined by a scoring vector (b_n, \dots, b_0) such that $b_i + b_{n-i} = \frac{2}{m(m-1)}$ for all $i \in \{0, \dots, n\}$ and $b_n \geq b_{n-1} \geq \dots \geq b_0 \geq 0$. The probability of an alternative x in a supporting size SDS f is therefore bounded by $f(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)} \leq (m-1) \frac{2}{m(m-1)} = \frac{2}{m}$.

Finally, we have all necessary lemmas for the proof of our first theorem.

Theorem 4. *The randomized Copeland rule is the only strategyproof SDS that satisfies anonymity, neutrality, and $\frac{2}{m}$ -Condorcet-consistency if $m \geq 3$ and $n \geq 3$. Moreover, no strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$ if $n \geq 3$.*

Proof. The theorem consists of two claims: the characterization of the randomized Condorcet rule f_C and the fact that no other strategyproof SDS can attain α -Condorcet-consistency for a larger α than f_C . We prove these claims separately.

Claim 1: The randomized Copeland rule is the only strategyproof SDS that satisfies $\frac{2}{m}$ -Condorcet-consistency, anonymity, and neutrality if $m, n \geq 3$.

The randomized Copeland rule f_C is a supporting size SDS and satisfies therefore anonymity, neutrality, and strategyproofness. Furthermore, it satisfies also $\frac{2}{m}$ -Condorcet-consistency because a Condorcet winner x wins every pairwise majority comparison in R . Hence, $n_{xy}(R) > \frac{n}{2}$ for all $y \in A \setminus \{x\}$, which implies that $f_C(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)} = (m-1) \frac{2}{m(m-1)} = \frac{2}{m}$.

Next, let f be an SDS satisfying anonymity, neutrality, strategyproofness, and $\frac{2}{m}$ -Condorcet-consistency. We show that f is the randomized Copeland rule. Since f is anonymous, neutral, and strategyproof, we can apply Theorem 3 to represent f as $f = \lambda f_{point} + (1 - \lambda) f_{sup}$, where $\lambda \in [0, 1]$, f_{point} is a point voting SDS, and f_{sup} is a supporting size SDS. Lemma 3 states that there is a profile R with Condorcet winner x such that $f_{point}(R, x) < \frac{2}{m}$, and it follows from Lemma 4 that $f_{sup}(R, x) \leq \frac{2}{m}$. Hence, $f(R, x) = \lambda f_{point}(R, x) + f_{sup}(R, x) < \frac{2}{m}$ if $\lambda > 0$. Therefore, f is a supporting size SDS as it satisfies $\frac{2}{m}$ -Condorcet-consistency.

Next, we show that f has the same scoring vector as the randomized Copeland rule. Since f is a supporting size SDS, there is a scoring vector $b = (b_n, \dots, b_0)$ with $b_n \geq b_{n-1} \geq \dots \geq b_0 \geq 0$ and $b_i + b_{n-i} = \frac{2}{m(m-1)}$ for all $i \in \{1, \dots, n\}$ such that $f(R, x) = \sum_{y \in A \setminus \{x\}} b_{n_{xy}(R)}$. Moreover, $f(R, x) = \frac{2}{m}$ if x is the Condorcet winner in R because of $\frac{2}{m}$ -Condorcet-consistency and Lemma 4. We derive from the definition of supporting size SDSs that the Condorcet winner x can only achieve this probability if $b_{n_{xy}(R)} = \frac{2}{m(m-1)}$ for every other alternatives $y \in A \setminus \{x\}$. Moreover, observe that the Condorcet winner needs to win every majority comparison but is indifferent about the exact supporting sizes. Hence, it follows that $b_i = \frac{2}{m(m-1)}$ for all $i > \frac{n}{2}$ as otherwise, there is a profile in which the Condorcet winner does not receive a probability of $\frac{2}{m}$. We also know that $b_i + b_{n-i} = \frac{2}{m(m-1)}$, so $b_i = 0$ for all $i < \frac{n}{2}$. If n is even, then $b_{\frac{n}{2}} = \frac{1}{m(m-1)}$ is required by the definition of supporting size SDSs as $\frac{n}{2} = n - \frac{n}{2}$. Hence, the scoring vector of f is equivalent to the scoring vector of the randomized Copeland rule, which proves that f is f_C .

Claim 2: No strategyproof SDS satisfies α -Condorcet-consistency for $\alpha > \frac{2}{m}$ if $n \geq 3$.

The claim is trivially true if $m \leq 2$ because α -Condorcet consistency for $\alpha > 1$ is impossible. Hence, let f denote a strategyproof SDS for $m \geq 3$ alternatives. We show in the sequel that f cannot satisfy α -Condorcet-consistency for $\alpha > \frac{2}{m}$. As a first step, we use Lemma 2 to construct a strategyproof SDS f^* that satisfies anonymity, neutrality, and α -Condorcet-consistency for the same α as f . Since f^* is anonymous, neutral, and strategyproof, it follows from Theorem 3 that f^* can be represented as a mixture of a point voting SDS f_{point} and a supporting size SDS f_{sup} , i.e., $f^* = \lambda f_{point} + (1 - \lambda) f_{sup}$ for some $\lambda \in [0, 1]$.

Next, we consider f_{point} and f_{sup} separately. Lemma 3 implies for f_{point} that there is a profile R with a Condorcet winner a such that $f_{point}(R, a) < \frac{2}{m}$. Moreover, Lemma 4 shows that $f_{sup}(R, a) \leq \frac{2}{m}$ because supporting size SDSs never return a larger probability than $\frac{2}{m}$. Thus, we derive the following inequality. w

$$\alpha \leq f^*(R, a) = \lambda f_{point}(R, a) + (1 - \lambda) f_{sup}(R, a) \leq \lambda \frac{2}{m} + (1 - \lambda) \frac{2}{m} = \frac{2}{m}$$

This proves that f^* , and therefore every strategyproof SDS, fails α -Condorcet-consistency for $\alpha \geq \frac{2}{m}$ \square

Remark 1. Lemma 2 can be applied to properties other than α -Condorcet-consistency, too. For example, given a strategyproof and β -*ex post* efficient SDS, we can construct another SDS that satisfies these axioms as well as anonymity and neutrality.

Remark 2. All axioms in the characterization of the randomized Copeland rule are independent of each other. The SDS that picks the Condorcet winner with probability $\frac{2}{m}$ if one exists and distributes the remaining probability uniformly between the other alternatives only violates strategyproofness. The randomized Borda rule satisfies all axioms of Theorem 4 but $\frac{2}{m}$ -Condorcet-consistency. An SDS that satisfies anonymity, strategyproofness, and $\frac{2}{m}$ -Condorcet-consistency can be defined based on an arbitrary order of alternatives x_0, \dots, x_{m-1} . Then, we pick an index $i \in \{0, \dots, m-1\}$ uniformly at random and return the winner of the majority comparison between x_i and $x_{i+1 \bmod m}$ (if there is a majority tie, a fair coin toss decides the winner). Finally, we can use the randomized Copeland rule f_C to construct an SDS that fails only anonymity for even n : we just ignore one voter when computing the outcome of f_C . Note here that for even n , an alternative x is a Condorcet winner in profile R if $n_{xy}(R) \geq \frac{n+2}{2}$ for all $y \in N \setminus \{x\}$, which means that x remains the Condorcet winner after removing a single voter.

Moreover, the impossibility in Theorem 4 does not hold when there are only $n = 2$ voters because random dictatorships are strategyproof and Condorcet-consistent in this case. The reason for this is that a Condorcet winner needs to be the most preferred alternative of both voters and is therefore chosen with probability 1.

Remark 3. The randomized Copeland rule has multiple appealing interpretations. Firstly, it can be defined as a supporting size SDS as shown in Section 2.1. Alternatively, it can be defined as the SDS that picks two alternatives uniformly at random and then picks the majority winner between them; majority ties are broken by a fair coin toss. Next, Theorem 4 shows that the randomized Copeland rule is the SDS that maximizes the value of α for α -Condorcet-consistency among all anonymous, neutral, and strategyproof SDSs. Finally, the randomized Copeland rule is the only strategyproof SDS that satisfies anonymity, neutrality, and assigns 0 probability to a Condorcet loser whenever it exists.

3.2 β -*ex post* Efficiency

According to Gibbard's random dictatorship theorem, random dictatorships are the only strategyproof SDSs that satisfy *ex post* efficiency. In this section, we show that this result is rather robust by identifying a tradeoff between β -*ex post* efficiency and γ -random dictatorships. More formally, we prove that for every $\epsilon \in [0, 1]$, all strategyproof and $\frac{1-\epsilon}{m}$ -*ex post* efficient SDSs are γ -randomly dictatorial for $\gamma \geq \epsilon$. If we set $\epsilon = 1$, we obtain the random dictatorship theorem. On the other hand, we derive from this theorem that every 0-randomly dictatorial and strategyproof SDS is β -*ex post* efficient for $\beta \geq \frac{1}{m}$, i.e., every such SDS is at least as inefficient as the uniform lottery. Moreover, we prove for every $\epsilon \in [0, 1]$ that mixtures of the uniform random dictatorship and the uniform lottery are the only ϵ -randomly dictatorial SDSs that satisfy anonymity, neutrality, strategyproofness, and $\frac{1-\epsilon}{m}$ -*ex post* efficiency. In summary, these results demonstrate that relaxing *ex post* efficiency does not lead to particularly appealing strategyproof SDSs. Furthermore, we also identify a tradeoff between α -Condorcet-consistency and β -*ex post* efficiency: every α -Condorcet consistent and strategyproof SDS fails β -*ex post* efficiency for $\beta < \frac{m-1}{m-2}\alpha$. Under the additional assumption of anonymity and neutrality, we characterize the strategyproof SDSs that maximize the ratio between α and β : all these

SDSs are mixtures of the randomized Copeland rule and the uniform random dictatorship.

For proving the tradeoff between β -*ex post* efficiency and γ -random dictatorships, we first investigate the efficiency of 0-randomly dictatorial strategyproof SDSs. In more detail, we prove next that every such SDS fails β -*ex post* efficiency for $\beta < \frac{1}{m}$.

Lemma 5. *No strategyproof SDS that is 0-randomly dictatorial satisfies β -*ex post* efficiency for $\beta < \frac{1}{m}$ if $m \geq 3$.*

The proof of this result is quite similar to the one for the upper bound on α -Condorcet-consistency in Theorem 4. In particular, we first show that all 0-randomly mixtures of duples and all 0-randomly dictatorial mixtures of unilaterals violate β -*ex post* efficiency for $\beta < \frac{1}{m}$. Next, we consider an arbitrary 0-randomly dictatorial SDS f and aim to show that there are a profile R and a Pareto-dominated alternative $x \in A$ such that $f(R, x) \geq \beta$. Even though Theorem 1 allows us to represent f as the convex combination of a 0-randomly dictatorial mixture of unilaterals f_{uni} and a mixture of duples f_{duple} , our previous observations have unfortunately no direct consequences for the β -*ex post* efficiency of f . The reason for this is that f_{uni} and f_{duple} might violate β -*ex post* efficiency for different profiles or alternatives. We solve this problem by transforming f into a 0-randomly dictatorial SDS f^* that is β -*ex post* efficient for the same β as f and satisfies additional properties. In particular, f^* can be represented as a convex combination of a 0-randomly dictatorial mixture of unilaterals f_{uni}^* and a 0-randomly dictatorial mixture of duples f_{duple}^* such that $f_{uni}^*(R, x) \geq \frac{1}{m}$ and $f_{duple}^*(R, x) \geq \frac{1}{m}$ for some profile R in which alternative x is Pareto-dominated. Consequently, f^* fails β -*ex post* efficiency for $\beta < \frac{1}{m}$, which implies that also f violates this axiom.

Based on Lemma 5, we can now show the tradeoff between *ex post* efficiency and the similarity to a random dictatorship.

Theorem 5. *For every $\epsilon \in [0, 1]$, every strategyproof and $\frac{1-\epsilon}{m}$ -*ex post* efficient SDS is γ -randomly dictatorial for $\gamma \geq \epsilon$ if $m \geq 3$. Moreover, if $\gamma = \epsilon$, $m \geq 4$, and the SDS satisfies additionally anonymity and neutrality, it is a mixture of the uniform random dictatorship and the uniform lottery.*

The proof of the first claim follows easily from Lemma 5: we consider a strategyproof SDS f and use the definition of γ -randomly dictatorial SDSs to represent f as a mixture of a random dictatorship and another strategyproof SDS g . Unless f is a random dictatorship, the maximality of γ entails that g is 0-randomly dictatorial. Hence, Lemma 5 implies that g can only be β -*ex post* efficient for $\beta \geq \frac{1}{m}$. Consequently, $\gamma \geq \epsilon$ must be true if f satisfies $\frac{1-\epsilon}{m}$ -*ex post* efficiency. For the second claim, we observe first that every anonymous, neutral, and strategyproof SDS f can be represented as a mixture of the uniform random dictatorship and another strategyproof, anonymous, and neutral SDS g . Moreover, unless f is 1-randomly dictatorial, g is 0-randomly dictatorial. Thus, Lemma 5 and the assumption that $\gamma = \epsilon$ require that g is exactly $\frac{1}{m}$ -*ex post* efficient. Finally, the claim follows by proving that the uniform lottery is the only 0-randomly dictatorial and strategyproof SDS that satisfies anonymity, neutrality, and $\frac{1}{m}$ -*ex post* efficiency if $m \geq 4$. For $m = 3$ the randomized Copeland rule also satisfies all required axioms and the uniform rule is thus not the unique choice.

Theorem 5 represents a continuous strengthening of Gibbard's random dictatorship theorem: the more *ex post* efficiency is required, the closer a strategyproof SDS gets to a random dictatorship. Conversely, our result also entails that γ -randomly dictatorial SDSs can only satisfy $\frac{1-\epsilon}{m}$ -*ex post* efficiency for $\epsilon \leq \gamma$. Moreover, the second part of the theorem indicates that relaxing *ex post* efficiency does not allow for particularly appealing strategyproof SDSs.

The correlation between β -*ex post* efficiency and γ -randomly dictatorialships also suggests a tradeoff between α -Condorcet-consistency and β -*ex post* efficiency because all random dictatorships are 0-Condorcet-consistent for sufficiently large m and n . Perhaps surprisingly, we show next that α -Condorcet consistency and β -*ex post* efficiency are in relation with each other for strategyproof SDSs. As a consequence of this insight, two strategyproof SDSs are particularly interesting: random dictatorships because they are the most *ex post* efficient SDSs, and the randomized Copeland rule because it is the most Condorcet-consistent SDS.

Theorem 6. *Every strategyproof SDS that satisfies anonymity, neutrality, α -Condorcet consistency, and β -*ex post* efficiency with $\beta = \frac{m-2}{m-1}\alpha$ is a mixture of the uniform random dictatorship and the randomized Copeland rule if $m \geq 4$, $n \geq 5$. Furthermore, there is no strategyproof SDS with $\beta < \frac{m-2}{m-1}\alpha$ if $m \geq 4$, $n \geq 5$.*

Proof. Let f be a strategyproof SDS that satisfies α -Condorcet consistency for some $\alpha \in [0, \frac{2}{m}]$ and let $\beta \in [0, 1]$ denote the minimal value such that f is β -*ex post* efficient. We first show that $\beta \geq \frac{m-2}{m-1}\alpha$ and hence apply Lemma 2 to construct an SDS f' that satisfies strategyproofness, anonymity, neutrality, α' -Condorcet consistency for $\alpha' \geq \alpha$, and β' -*ex post* efficiency for $\beta' \leq \beta$. In particular, if f' is only β' -*ex post* efficient for $\beta' \geq \frac{m-2}{m-1}\alpha'$, then f can only satisfy β -*ex post* efficiency for $\beta \geq \beta' \geq \frac{m-2}{m-1}\alpha' \geq \frac{m-2}{m-1}\alpha$.

Since f' satisfies anonymity, neutrality, and strategyproofness, we can apply Theorem 3 to represent it as a mixture of a supporting size SDS and a point voting SDS, i.e., $f' = \lambda f_{point} + (1 - \lambda) f_{sup}$ for some $\lambda \in [0, 1]$. Let (a_1, \dots, a_m) and (b_0, \dots, b_n) denote the scoring vectors describing f_{point} and f_{sup} , respectively. Next, we derive lower bound for α' and an upper bound for β' by considering specific profiles. First, consider the profile R in which every voter reports a as his best alternative and b as his second best alternative; the remaining alternatives can be ordered arbitrarily. It follows from the definition of point voting SDSs that $f_{point}(R, b) = na_2$ and from the definition of supporting size SDS that $f_{sup}(R, b) = (m - 2)b_n + b_0$. Since a Pareto-dominates b in R , it follows that $\beta' \geq f(R, b) = \lambda na_2 + (1 - \lambda)((m - 2)b_n + b_0)$.

For the upper bound on α , consider the following profile R' where alternative x is never ranked first, but it is the Condorcet winner and wins every pairwise comparison only with minimal margin. We denote for the definition of R' the alternatives as $A = \{x, x_1, \dots, x_{m-1}\}$. In R' , the voters $i \in \{1, 2, 3\}$ ranks alternatives $X_i := \{x_k \in A \setminus \{x\} : k \bmod 3 = i - 1\}$ above x and all other alternatives below. Since $m \geq 4$, none of them ranks x first. If the number of voters n is even, we duplicate voters 1, 2, and 3. As last step, we add pairs of voters with inverse preferences such that no voter prefers x the most until R' consists of n voters. Since alternative x is never top-ranked in R' , it follows that $f_{point}(R', x) \leq na_2$. Furthermore, $n_{xy}(R') = \lceil \frac{n+1}{2} \rceil$ for all $y \in A \setminus \{x\}$ and therefore $f_{sup}(R', x) = (m - 1)b_{\lceil \frac{n+1}{2} \rceil}$. Finally, we

derive that $\alpha' \leq f(R', x) \leq \lambda n a_2 + (1 - \lambda)(m - 1)b_{\lceil \frac{n+1}{2} \rceil}$ because x is by construction the Condorcet winner in R' .

Using these bounds, we show next that f' is only β' -*ex post* efficiency for $\beta' \geq \frac{m-2}{m-1}\alpha'$, which proves the second claim of the theorem. In the subsequent calculation, the first and last inequality follow from our previous analysis. The second inequality is true since $\frac{m-2}{m-1} \leq 1$ and $\frac{m-2}{m-1}(m-1) = (m-2)$. The third inequality uses the definition of supporting size SDSs.

$$\begin{aligned} \beta' &\geq \lambda n a_2 + (1 - \lambda)((m - 2)b_n + b_0) \\ &\geq \frac{m-2}{m-1}\lambda n a_2 + \frac{m-2}{m-1}(1 - \lambda)((m - 1)b_n + b_0) \\ &\geq \frac{m-2}{m-1}\lambda n a_2 + \frac{m-2}{m-1}(1 - \lambda)(m - 1)b_{\lceil \frac{n+1}{2} \rceil} \\ &\geq \frac{m-2}{m-1}\alpha' \end{aligned}$$

Finally, note that, if $\beta' = \frac{m-2}{m-1}\alpha'$, all inequalities must be tight. If the second inequality is tight $a_2 = 0$ and $b_0 = 0$, and when the third inequality is tight $b_n = b_{\lceil \frac{n+1}{2} \rceil}$. These observations fully specify the scoring vectors of f_{point} and f_{sup} . For the point voting SDS, $a_2 = 0$ implies $a_i = 0$ for all $i \geq 2$ and $a_1 = \frac{1}{n}$, i.e., f_{point} is the uniform random dictatorship. Next, $b_0 = 0$ and $b_n = b_{\lceil \frac{n+1}{2} \rceil}$ imply that $b_i = \frac{2}{m(m-1)}$ for all $i \in \{\lceil \frac{n+1}{2} \rceil, \dots, b_n\}$ and $b_i = 0$ for all $i \in \{0, \dots, \lfloor \frac{n-1}{2} \rfloor\}$. Moreover, if n is even, the definition of supporting size SDSs requires that $b_{\frac{n}{2}} = \frac{1}{m(m-1)}$. This shows that f_{sup} is the randomized Copeland rule. Consequently, the SDS f' is a mixture of the uniform random dictatorship and the randomized Copeland rule if $\beta' = \frac{m-2}{m-1}\alpha'$. This proves that every strategyproof SDS that satisfies anonymity, neutrality, α -Condorcet consistency, and β -*ex post* efficiency with $\beta = \frac{m-2}{m-1}\alpha$ is a mixture of the uniform random dictatorship and the randomized Copeland rule. \square

Remark 4. All axioms of the characterization in Theorem 6 are independent of each other. Every mixture of random dictatorships other than the uniform one and the randomized Copeland rule only violates anonymity. An SDS that violates only neutrality can be constructed by using a variant of the randomized Copeland rule that does not split the probability equally if there is a majority tie. Finally, the correlation between α -Condorcet-consistency and β -*ex post* efficiency is required since the uniform lottery satisfies all other axioms. Moreover, all bounds on m and n in Theorem 6 are tight. If there are only $n = 2$ voters, $m = 3$ alternatives, or $m = 4$ alternatives and $n = 4$ voters, the uniform random dictatorship is not 0-Condorcet consistent since a Condorcet winner is always ranked first by at least one voter. Hence, the bound on β does not hold in these cases. In contrast, our proof shows that Theorem 6 is also true when $n = 3$.

4 CONCLUSION

In this paper, we analyzed strategyproof SDSs by considering relaxations of Condorcet-consistency and *ex post* efficiency. Our findings, which are summarized in Figure 2, show that two strategyproof SDSs perform particularly well with respect to these axioms: the uniform random dictatorship (and random dictatorships in general), and the randomized Copeland rule. In more detail, we prove that




Figure 2: Graphical summary of our results. Points in the figures correspond to SDSs and the horizontal axis indicates in both figures the value of β for which the considered SDS is β -*ex post* efficient. In the left figure, the vertical axis states the α for which the considered SDSs are α -Condorcet-consistent, and in the right figure, it shows the γ for which SDSs are γ -randomly dictatorial. Theorems 4 and 6 show that no strategyproof SDS lies in the grey area of the left figure. Theorem 5 shows that no strategyproof SDS lies in the grey area below the diagonal in the right figure. Furthermore, no SDS lies in the grey area above the diagonal since a γ -randomly dictatorial SDS can put no more than $1 - \gamma$ probability on Pareto-dominated alternatives. Finally, the following SDS are marked in the figures: d corresponds to all random dictatorships, c to the randomized Copeland rule, b to the randomized Borda rule, and u to the uniform lottery.

the randomized Copeland rule is the only strategyproof, anonymous, and neutral SDS which guarantees a probability of $\frac{2}{m}$ to the Condorcet winner. Since no other strategyproof SDS can guarantee more probability to the Condorcet winner (even if we drop anonymity and neutrality), this characterization identifies the randomized Copeland rule as one of the most Condorcet-consistent strategyproof SDSs. On the other hand, Gibbard's random dictatorship theorem shows that random dictatorships are the only *ex post* efficient and strategyproof SDSs. We present a continuous generalization of this result: for every $\epsilon \in [0, 1]$, every $\frac{1-\epsilon}{m}$ -*ex post* efficient and strategyproof SDS is γ -randomly dictatorial for $\gamma \geq \epsilon$. This means informally that, even if we allow that Pareto-dominated alternatives can get a small amount of probability, we end up with an SDS similar to a random dictatorship. Finally, we derive a tradeoff between α -Condorcet-consistency and β -*ex post* efficiency for strategyproof SDSs: every strategyproof and α -Condorcet-consistent SDS fails β -*ex post* efficiency for $\beta < \frac{m-2}{m-1}\alpha$. This theorem entails that it is not possible to jointly optimize both notions, which again highlights the special role of the randomized Copeland rule and random dictatorships.

ACKNOWLEDGMENTS

This work was supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1. We thank Dominik Peters for stimulating discussions and the anonymous reviewers for their helpful comments.

REFERENCES

- [1] A. Abdulkadiroğlu and T. Sönmez. 1998. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica* 66, 3 (1998), 689–701.
- [2] H. Aziz, F. Brandl, F. Brandt, and M. Brill. 2018. On the Tradeoff between Efficiency and Strategyproofness. *Games and Economic Behavior* 110 (2018), 1–18.
- [3] S. Barberà. 1979. Majority and Positional Voting in a Probabilistic Framework. *Review of Economic Studies* 46, 2 (1979), 379–389.
- [4] S. Barberà. 1979. A Note on Group Strategy-Proof Decision Schemes. *Econometrica* 47, 3 (1979), 637–640.
- [5] J.-P. Benoît. 2002. Strategic Manipulation in Voting Games When Lotteries and Ties Are Permitted. *Journal of Economic Theory* 102, 2 (2002), 421–436.
- [6] A. Bogomolnaia and H. Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100, 2 (2001), 295–328.
- [7] F. Brandl, F. Brandt, M. Eberl, and C. Geist. 2018. Proving the Incompatibility of Efficiency and Strategyproofness via SMT Solving. *J. ACM* 65, 2 (2018), 1–28.
- [8] F. Brandt, V. Conitzer, and U. Endriss. 2013. Computational Social Choice. In *Multiagent Systems* (2nd ed.), G. Weiß (Ed.), MIT Press, Chapter 6, 213–283.
- [9] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [10] F. Brandt, P. Lederer, and R. Romen. Relaxed notions of Condorcet-consistency and efficiency for strategyproof social decision schemes. Technical report, <https://arxiv.org/abs/2201.10418>, 2022.
- [11] S. Chatterji, A. Sen, and H. Zeng. 2014. Random dictatorship domains. *Games and Economic Behavior* 86 (2014), 212–236.
- [12] Y.-K. Che and F. Kojima. 2010. Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms. *Econometrica* 78, 5 (2010), 1625–1672.
- [13] V. Conitzer and T. Sandholm. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*. 627–634.
- [14] J. Duggan. 1996. A geometric proof of Gibbard’s random dictatorship theorem. *Economic Theory* 7, 2 (1996), 365–369.
- [15] B. Dutta, H. Peters, and A. Sen. 2002. Strategy-Proof Probabilistic Mechanisms in Economies with Pure Public Goods. *Journal of Economic Theory* 106, 2 (2002), 392–416.
- [16] B. Dutta, H. Peters, and A. Sen. 2007. Strategy-proof cardinal decision schemes. *Social Choice and Welfare* 28, 1 (2007), 163–179.
- [17] L. Ehlers, H. Peters, and T. Storcken. 2002. Strategy-Proof Probabilistic Decision Schemes for One-Dimensional Single-Peaked Preferences. *Journal of Economic Theory* 105, 2 (2002), 408–434.
- [18] P. C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489.
- [19] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.
- [20] A. Gibbard. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45, 3 (1977), 665–681.
- [21] J. C. Heckelman. 2003. Probabilistic Borda rule voting. *Social Choice and Welfare* 21 (2003), 455–468.
- [22] J. C. Heckelman and F. H. Chen. 2013. Strategy Proof Scoring Rule Lotteries for Multiple Winners. *Journal of Public Economic Theory* 15, 1 (2013), 103–123.
- [23] A. Hylland. 1980. Strategyproofness of Voting Procedures with Lotteries as Outcomes and Infinite Sets of Strategies. (1980). Mimeo.
- [24] S. Nandeibam. 1997. An alternative proof of Gibbard’s random dictatorship result. *Social Choice and Welfare* 15, 4 (1997), 509–519.
- [25] S. Nandeibam. 2013. The structure of decision schemes with cardinal preferences. *Review of Economic Design* 17, 3 (2013), 205–238.
- [26] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- [27] A. D. Procaccia. 2010. Can approximation circumvent Gibbard-Satterthwaite?. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*. 836–841.
- [28] M. A. Satterthwaite. 1975. Strategy-Proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.
- [29] A. Sen. 2011. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIEs* 2, 4 (2011), 515–527.
- [30] Y. Shoham and K. Leyton-Brown. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- [31] Y. Tanaka. 2003. An alternative proof of Gibbard’s random dictatorship theorem. *Review of Economic Design* 8 (2003), 319–328.