

An Analytical and Experimental Comparison of Maximal Lottery Schemes

Florian Brandl Felix Brandt Christian Stricker

Technische Universität München
{brandlfl,brandtf,stricker}@in.tum.de

Abstract

Randomized voting rules are gaining increasing attention in computational and non-computational social choice. A particularly interesting class of such rules are *maximal lottery (ML) schemes*, which were proposed by Peter Fishburn in 1984 and have been repeatedly recommended for practical use. However, the subtle differences between different *ML* schemes are often ignored. Two canonical subsets of *ML* schemes are *C1-ML* schemes (which only depend on unweighted majority comparisons) and *C2-ML* schemes (which only depend on weighted majority comparisons). We prove that *C2-ML* schemes are the only Pareto efficient—but also among the most manipulable—*ML* schemes. Furthermore, we evaluate the frequency of manipulable preference profiles and the degree of randomization of *ML* schemes via extensive computer simulations. In general, *ML* schemes are rarely manipulable and often do not randomize at all, especially when there are only few alternatives. For up to 21 alternatives, the average support size of *ML* schemes lies below 4 under reasonable assumptions. The average degree of randomization (in terms of Shannon entropy) of *C2-ML* schemes is significantly lower than that of *C1-ML* schemes.

1 Introduction

When aggregating the preferences of multiple agents into one collective choice, it is easily seen that completely symmetric situations call for randomization. Moreover, it has been shown that—apart from guaranteeing impartiality—randomization allows the circumvention of well-known impossibility results that have plagued social choice theory for long [see, e.g., Gibbard, 1977, Brandl et al., 2016, 2018b, Brandl and Brandt, 2018]. Two types of randomized voting rules that have been shown to be attractive from an axiomatic point of view are *random (serial) dictatorships* and *maximal lottery (ML) schemes*. While random dictatorships are renowned for being immune to strategic manipulation, *ML* schemes satisfy desirable consistency conditions (such as Condorcet-consistency, population-consistency, and

composition-consistency) as well as strong notions of Pareto efficiency.

The notion of a maximal lottery was first conceived by Kreweras [1965] and independently proposed and studied in much more detail by Fishburn [1984a]. Interestingly, maximal lotteries have been rediscovered several times by economists [Laffond et al., 1993], mathematicians [Fisher and Ryan, 1995], political scientists [Felsenthal and Machover, 1992], and computer scientists [Rivest and Shen, 2010]. In particular, Laffond et al. [1993, 1996], Dutta and Laslier [1999], Laslier [2000], and Brandt et al. [2018] have extensively studied the *support* of maximal lotteries, called the *bipartisan set* or the *essential set*.¹ Felsenthal and Machover [1992] and Rivest and Shen [2010] also discuss whether *ML* schemes are suitable for real-world political elections. Rivest and Shen conclude that “[the maximal lotteries system] is not only theoretically interesting and optimal, but simple to use in practice; it is probably easier to implement than, say, IRV [instant-runoff voting]. We feel that it can be recommended for practical use.” More recently, Peyre [2013] and Hoang [2017] have popularized maximal lotteries in France under the name *scrutin de Condorcet randomisé* (randomized Condorcet voting system), which alludes to the fact that maximal lotteries are essentially randomized (weak) Condorcet winners (see Section 3).² Easy-to-use voting tools for *C1-ML* and *C2-ML* are available on the websites <https://votation.ovh> (in French) and <https://pnyx.dss.in.tum.de>, respectively.

The literature on maximal lotteries often ignores the fact that there are different variants of maximal lottery schemes, which under certain circumstances may lead to completely contradictory outcomes (e.g., two lotteries with disjoint support). All these (infinitely many) variants are captured by Fishburn’s original classification of maximal lotteries [Fishburn, 1984a]. The two main candidates are *C1-ML schemes* (which only depend on unweighted majority comparisons) and *C2-ML schemes* (which only depend on weighted majority comparisons). *C1-ML* schemes have, for example,

¹The term bipartisan set was proposed by Nobel Prize Laureate Roger Myerson [Laffond et al., 1993].

²Two French YouTube videos about maximal lotteries by Lê Nguyễn Hoang enjoy immense popularity (<https://youtu.be/wKimU8jy2a8> and <https://youtu.be/vAdGZkXh1NM>).

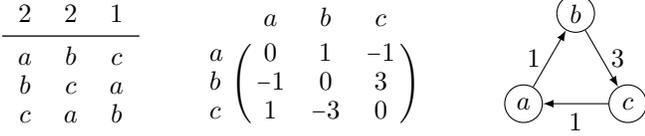


Figure 1: Illustration of the definitions stated in Section 2. The table on the left represents a preference profile R for three alternatives $\{a, b, c\}$ and five agents. The number on top of each column denotes the number of agents with the corresponding preference relation, e.g., there are two agents with the preference relation $\succsim_i: a, b, c$. The corresponding matrix of majority margins and the weighted digraph are displayed on the right. The majority relation of R is cyclic, i.e., $a \succ_R b \succ_R c \succ_R a$. Hence, R does not admit a (weak) Condorcet winner. $C1\text{-}ML(R) = \{1/3 a + 1/3 b + 1/3 c\}$ and $C2\text{-}ML(R) = \{3/5 a + 1/5 b + 1/5 c\}$.

been considered by Kreweras [1965], Felsenthal and Machover [1992], Laffond et al. [1993], Fisher and Ryan [1995], Peyre [2013], and Hoang [2017] while $C2\text{-}ML$ schemes have been considered by Dutta and Laslier [1999], Laslier [2000], Rivest and Shen [2010], Aziz et al. [2013], Brandl et al. [2016], and Brandl and Brandt [2018]. In this paper, we provide a detailed analytical and experimental comparison of all ML schemes.

2 Preliminaries

Let $A = \{a, b, \dots\}$ be a finite set of m alternatives and $N = \{1, \dots, n\}$ be a set of n agents. A (weak) preference relation on A is a complete and transitive binary relation on A . The preference relation reported by agent i is denoted by \succsim_i , and the set of all preference relations on A is denoted by \mathcal{R} . We write \succ_i for the strict part of \succsim_i , i.e., $x \succ_i y$ if $x \succsim_i y$ but not $y \succsim_i x$, and \sim_i for the indifference part of \succsim_i , i.e., $x \sim_i y$ if $x \succsim_i y$ and $y \succsim_i x$. A preference relation \succsim_i is called *strict* if it is anti-symmetric, i.e., either $x \succ_i y$ or $y \succ_i x$ for all distinct alternatives $x, y \in A$. 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. For example, $x \succ_i y \sim_i z$ is represented by $\succsim_i: x, \{y, z\}$. A *preference profile* $R = (\succsim_1, \dots, \succsim_n)$ is an n -tuple containing a preference relation for each agent. The set of all preference profiles is thus given by \mathcal{R}^N . For a preference profile $R \in \mathcal{R}^N$, we denote by $n_{xy} = |\{i \in N: x \succ_i y\}|$ the number of agents who strictly prefer x to y . The *majority margin* of x over y is given by $m_{xy} = n_{xy} - n_{yx}$. The majority margins between all pairs of alternatives can be represented by a skew-symmetric matrix whose rows and columns are indexed by alternatives; the majority margin of x over y is given in the cell indexed by (x, y) . Alternatively, majority margins can be illustrated by a weighted digraph with an edge from x to y with weight m_{xy} if $m_{xy} > 0$. The *majority relation* \succsim_R on alternatives for a given preference profile R can be derived from the majority margins: $x \succsim_R y$ if and only if $m_{xy} \geq 0$. An alternative $x \in A$ is a *Condorcet winner* in R if $x \succ_R y$ for all $y \in A \setminus \{x\}$ and a *weak Condorcet winner* in R if $x \succsim_R y$ for all $y \in A \setminus \{x\}$. The preceding definitions are illustrated in Figure 1.

We consider voting rules that randomize over alternatives. The set of all *lotteries* over A is denoted by $\Delta(A)$, i.e., $\Delta(A) = \{p \in \mathbb{R}_{\geq 0}^A: \sum_{x \in A} p(x) = 1\}$, where $p(x)$ is the probability that p assigns to x . By $\text{supp}(p)$ we denote the *support* of a lottery $p \in \Delta(A)$, i.e., the set of all alternatives to which p assigns positive probability. A lottery p is degenerate if $|\text{supp}(p)| = 1$. We write lotteries as convex combinations of alternatives, e.g., the uniform lottery on $\{a, b\}$ is denoted by $1/2 a + 1/2 b$. A *social decision scheme* (SDS) takes as input a preference profile $R \in \mathcal{R}^N$ and returns a lottery over A . Two common symmetry conditions for SDS s are *anonymity*, i.e., the SDS is invariant under renaming the agents, and *neutrality*, i.e., the SDS is symmetric with respect to the alternatives. Furthermore, an SDS is *majoritarian* if it only depends on the majority relation, i.e., $f(R) = f(\hat{R})$ whenever $\succsim_R = \succsim_{\hat{R}}$.

Defining properties such as efficiency and strategyproofness for SDS s requires to make assumptions about the agents' preferences over lotteries based on their preferences over alternatives. To this end, we consider *lottery extensions*, which map a preference relation on the set of alternatives A to a preference relation on the set of lotteries $\Delta(A)$. The most common lottery extension is *stochastic dominance* (SD), according to which a lottery p is preferred to another lottery q if for every alternative $x \in A$, p is at least as likely to return an alternative that is at least as good as x as q . Formally,

$$p \succsim_i^{SD} q \Leftrightarrow \sum_{y \succsim_i x} p(y) \geq \sum_{y \succsim_i x} q(y) \text{ for all } x \in A. \quad (SD)$$

It is a well-known fact that $p \succsim_i^{SD} q$ if and only if the expected utility of p is at least as large as that of q for every von Neumann Morgenstern utility function consistent with \succsim_i .

Another important lottery extension is based on pairwise comparisons of alternatives [Brandl et al., 2018b, Brandl and Brandt, 2018]. A lottery p is preferred to another lottery q according to *pairwise comparisons* (PC) if p is more likely to return a more preferred alternative than q . Formally,

$$p \succsim_i^{PC} q \Leftrightarrow \sum_{x \succsim_i y} p(x)q(y) \geq \sum_{x \succsim_i y} p(y)q(x). \quad (PC)$$

While the PC extension results in preferences over lotteries that cannot be represented by any von Neumann Morgenstern utility function, it represents a refinement of the SD extension, i.e., $\succsim_i^{SD} \subseteq \succsim_i^{PC}$ [Fishburn, 1984b, Theorem 8]. In contrast to the SD extension, which may result in incomplete preferences over lotteries, the PC extension yields a complete preference relation over lotteries. If $a \succ_i b \succ_i c$, $p = 1/2 a + 1/2 c$, and $q = 2/3 b + 1/3 c$, we have that $p \succ_i^{PC} q$ but neither $p \succsim_i^{SD} q$ nor $q \succsim_i^{SD} p$.

3 Maximal Lottery Schemes

In this paper, we focus on *maximal lottery schemes*, a class of SDS s introduced by Fishburn [1984a]. Every maximal lottery scheme is based on an odd and monotone function $\tau: \mathbb{Z} \rightarrow \mathbb{R}$ with $\tau(1) = 1$ and returns an optimal mixed strategy of the symmetric zero-sum game induced by $(\tau(m_{xy}))_{x, y \in A}$.³ For

³ τ is odd if $\tau(-k) = -\tau(k)$ for all $k \in \mathbb{Z}$. Thus, $\tau(0) = 0$ and τ is completely defined by its values on positive integers.

every such function τ and every $R \in \mathcal{R}^N$, we define the set $ML^\tau(R)$ of maximal lotteries in R with respect to τ as follows.

$$ML^\tau(R) = \{p \in \Delta(A) : \sum_{x,y \in A} p(x)q(y)\tau(m_{xy}) \geq 0 \text{ for all } q \in \Delta(A)\}.$$

When τ is the identity function, maximal lotteries thus correspond to randomized weak Condorcet winners in the following sense: let p be a maximal lottery. Then, for any lottery q , the expected number of agents who prefer the outcome of p to that of q is larger than the expected number of agents who prefer the outcome of q to that of p . Maximal lotteries can be found via linear programming and thus in polynomial time. An SDS f is a maximal lottery scheme based on τ if, for all $R, \hat{R} \in \mathcal{R}^N$, $f(R) \in ML^\tau(R)$ and $f(R) = f(\hat{R})$ whenever $ML^\tau(R) = ML^\tau(\hat{R})$.

The function τ describes how different sizes of majorities are traded off against each other. Roughly speaking, the steeper the function τ , the more emphasize is given to larger majorities. Two notable classes of ML schemes are obtained for particularly natural choices of τ . The class of ML schemes based on the *sign function* is denoted by $C1-ML$, in accordance with Fishburn's [1977] classification of deterministic voting rules. $C1-ML$ schemes are in fact the only ML schemes that are C1 functions, i.e., that only depend on the majority relation and thus treat all sizes of majorities equally. The class of ML schemes based on the *identity function* is referred to as $C2-ML$ for short. When there is an odd number of agents with strict preferences, both $C1-ML$ and $C2-ML$ return a unique lottery, which follows from the fact that optimal mixed strategies are unique in symmetric zero-sum games with odd off-diagonal payoffs; moreover, the support of this lottery contains an odd number of alternatives [Laffond et al., 1997]. For the preference profile R from Figure 1, we have that $ML^\tau(R) = \{1/(\tau(3)+2)(\tau(3)a+b+c)\}$. In particular, $C1-ML(R) = \{1/3a + 1/3b + 1/3c\}$ and $C2-ML(R) = \{3/5a + 1/5b + 1/5c\}$. Dutta and Laslier [1999, Proposition 4.2] gave an example showing that the unique lottery returned by $C1-ML$ can have disjoint support from the unique lottery returned by $C2-ML$.

A potential advantage of $C1-ML$ schemes is that they require less information than other ML schemes. It suffices to only elicit the majority relation in order to compute the election outcome.⁴

In the sequel, we derive analytical results about Pareto efficiency and strategyproofness of ML schemes (cf. Sections 4 and 5) and complement these results by empirical findings about the frequency of manipulation instances and the degree of randomization under the impartial anonymous culture assumption (cf. Section 6).

⁴However, as our simulations show, $C2-ML$ schemes also only depend on the majority relation to a large extent. Precise majority margins are only required for pairwise comparisons between a few crucial alternatives.

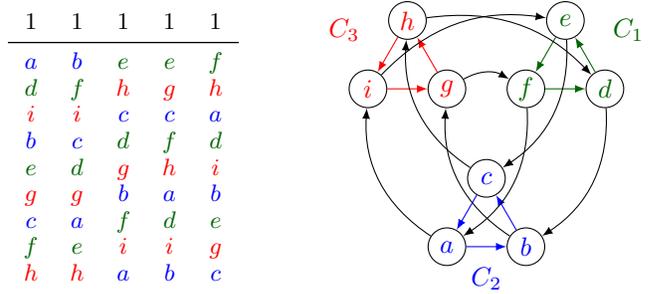


Figure 2: The preference profile R used in the proof of Theorem 4.2 is depicted on the left. The corresponding majority relation is illustrated on the right. Omitted edges point counter-clockwise.

4 Efficiency

A fundamental economic property is (Pareto) efficiency, which prescribes that no agent can be made better off without making another agent worse off. When assuming that lotteries are compared based on stochastic dominance, one obtains the notion of SD -efficiency. Formally, a lottery $p \in \Delta(A)$ is SD -efficient for a preference profile $R \in \mathcal{R}^N$ if there is no lottery $q \in \Delta(A)$ such that $q \succ_i^{SD} p$ for all $i \in N$ and $q \succ_j^{SD} p$ for some $j \in N$. An SDS f is SD -efficient if $f(R)$ is SD -efficient for all $R \in \mathcal{R}^N$. In a similar vein, one can define PC -efficiency, which results in an efficiency notion that is stronger than SD -efficiency. Theorem 4.1 shows two important facts regarding efficiency of ML schemes: first, $C2-ML$ schemes are PC -efficient and, second, no other ML schemes are even SD -efficient. The involved proof of this theorem is omitted due to limited space.

Theorem 4.1. *Every $C2-ML$ scheme is PC -efficient. No other ML scheme is SD -efficient for all values of m and n .*

Theorem 4.1 implies that no $C1-ML$ scheme is SD -efficient. This raises the question whether there is any SD -efficient majoritarian SDS. In Theorem 4.2 we show that the answer to this question is negative when also assuming neutrality. The base case for the following proof was found with the help of a computer.

Theorem 4.2. *Every majoritarian and neutral SDS violates SD -efficiency for $m \geq 9$ (and $n = 5, n = 7, \text{ or } n \geq 9$), even when preferences are strict.*

Proof. We first prove the case $m = 9$ and $n = 5$. Let $A = \{a, b, c, d, e, f, g, h, i\}$ and consider the preference profile $R \in \mathcal{R}^N$ and the corresponding weighted majority graph depicted in Figure 2. Every majoritarian and neutral SDS f yields the uniform lottery $p = 1/9(a + b + c + d + e + f + g + h + i)$ over all alternatives in A for R due to the symmetrical structure of the majority relation of R . However, shifting all the probabilities from the “red” alternatives g, h , and i to the “green” alternatives d, e , and f yields another lottery $q = 1/9(a + b + c) + 2/9(d + e + f)$, which is SD -preferred to p by all agents, i.e., $q \succ_i^{SD} p$ for all $i \in N$. This can be checked by looking at R , because there is always a “green” alternative directly above a “red” alternative for each agent. Additional alternatives can be added at the bottom of each

preference relation. These alternatives will then be Pareto dominated and do not break the symmetry among the first 9 alternatives. If f assigns positive probability to a Pareto dominated alternative, it is SD -inefficient. Otherwise, f is SD -inefficient with the same argument as for the case $m = 9$. For more than 5 agents, analogous proofs can be obtained by duplication and combination of profiles that induce the same majority graph as R . \square

This impossibility is somewhat surprising since *ex post* efficient, majoritarian SDSs do exist. In fact, there is an elegant characterization of such SDSs [see Brandt et al., 2016]. This characterization can be used to show that neutrality is required for the impossibility since every SDS that puts probability 1 on an *ex post* efficient alternative satisfies SD -efficiency. Theorem 4.2 can be used as an argument against majoritarian SDSs (i.e., SDSs based on C1 functions or so-called tournament solutions) in general as the majority relation does not contain enough information to guarantee SD -efficiency.

Theorem 4.2 implies Theorem 1 by Aziz et al. [2013] who showed that no neutral, majoritarian, SD -efficient, and SD -strategyproof SDS exists (although their result already holds when $m = 4$).

5 Strategyproofness

In this section, we derive analytic results about the vulnerability of ML schemes to strategic manipulation. Just like in Section 4, we consider manipulability when preferences over lotteries are based on stochastic dominance. Formally, an SDS f is SD -manipulable for a preference profile $R \in \mathcal{R}^N$ if there is a preference profile $\hat{R} \in \mathcal{R}^N$ and $j \in N$ such that $\succsim_i = \hat{\succsim}_i$ for all $i \neq j$ and $f(\hat{R}) \succ_j^{SD} f(R)$. An SDS f is SD -strategyproof if it is not SD -manipulable for any preference profile $R \in \mathcal{R}^N$. Analogously, one can define PC -manipulability. The resulting notion of PC -strategyproofness is stronger than SD -strategyproofness.

Since only few pairs of lotteries are comparable with respect to stochastic dominance, agents can rarely manipulate and SD -strategyproofness can be seen as a rather weak notion of strategyproofness.⁵ Still, it has proved to be largely prohibitive, even for general SDSs. Brandl et al. [2018a, Theorem 3.1] showed that every anonymous and neutral SDS is either SD -inefficient or SD -manipulable. Since every $C2$ - ML scheme is anonymous and SD -efficient (cf. Theorem 4.1), it follows that every neutral $C2$ - ML scheme violates SD -strategyproofness. The following result shows that this shortcoming in fact pertains to *all* ML schemes.

Theorem 5.1. *All ML schemes violate SD -strategyproofness when $m \geq 3$ and $n \geq 3$. This even holds for strict preferences when $m \geq 4$ and n is odd.*

Proof. We first prove the case $m = 3$ and $n = 3$. Let $A = \{a, b, c\}$ and ML^τ be an arbitrary ML scheme. Consider the following preference profile $R \in \mathcal{R}^N$, where agent i in the last column is indifferent between a and c , and the corresponding matrix of majority margins.

$$\begin{array}{ccc|ccc} 1 & 1 & 1 & & & \\ \hline a & c & b & a & b & c \\ b & a & \{a, c\} & b & \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 1 \end{pmatrix} \\ c & b & & c & \begin{pmatrix} 0 & -1 & 0 \end{pmatrix} \end{array}$$

By the definition of ML schemes, we have that $\tau(1) = 1$ independently of the choice of τ . Then, $ML^\tau(R) = \{\lambda a + (1 - \lambda)c : 1/2 \leq \lambda \leq 1\}$. Now assume that the agent i changes his preference relation to $\hat{\succsim}_i : b, c, a$, i.e., he breaks the tie between a and c . For the resulting preference profile \hat{R} we have $ML^\tau(\hat{R}) = \{q\}$, where $q = 1/3 a + 1/3 b + 1/3 c$. Note that $q \succ_j^{SD} p$ for all $p \in ML^\tau(R)$. Thus, agent i can manipulate ML^τ at R , which shows that no ML scheme based on τ is SD -strategyproof.

Secondly, we prove the case $m = 4$ and $n = 3$ with strict preferences only. Similar to the first case, let $A = \{a, b, c, d\}$ and ML^τ be again an arbitrary ML scheme. Consider the following preference profile $R \in \mathcal{R}^N$ with the corresponding matrix of majority margins.

$$\begin{array}{ccc|cccc} 1 & 1 & 1 & & & & \\ \hline a & c & b & a & b & c & d \\ d & b & d & b & \begin{pmatrix} 0 & -1 & 1 & 1 \\ 1 & 0 & -1 & 1 \end{pmatrix} \\ c & a & a & c & \begin{pmatrix} -1 & 1 & 0 & -1 \end{pmatrix} \\ b & d & c & d & \begin{pmatrix} -1 & -1 & 1 & 0 \end{pmatrix} \end{array}$$

Again independently of the choice of τ , we have $ML^\tau(R) = \{p\}$, where $p = 1/3 a + 1/3 b + 1/3 c$. Now assume that the agent j in the last column changes his preference relation to $\hat{\succsim}_j : b, d, c, a$, i.e., he swaps a and c . The resulting preference profile is denoted by \hat{R} . The majority margins in R and \hat{R} are the same except that $m_{ac} = 1$ and $\hat{m}_{ac} = -1$. Then we have that $ML^\tau(\hat{R}) = \{q\}$, where $q = 1/3 b + 1/3 c + 1/3 d$. Observe that $q \succ_j^{SD} p$. Hence, agent j can manipulate ML^τ at R .

Both cases can be generalized to larger profiles via a construction similar to the one by Brandl et al. [2018a, Lemma 4.1]. \square

For all ML schemes except for $C1$ - ML schemes, Theorem 5.1 also holds when there are only three alternatives and preferences are strict.

The proof of Theorem 5.1 crucially relies on the fact that Condorcet winners may fail to exist. Empirical studies have however observed that almost all preferences profiles that appear in real-world elections admit a Condorcet winner (see Footnote 7), in which case every ML scheme chooses the Condorcet winner with probability 1. As it turns out, no ML scheme can be PC -manipulated whenever there is a Condorcet winner (not even to a profile without a Condorcet winner), which remedies the seemingly severe implications of Theorem 5.1 for most real-world settings. This statement has been shown by Peyre [2013] and Hoang [2017] for $C1$ - ML schemes.⁶ In particular, it implies that choosing the Con-

⁶Hoang [2017] showed *Condorcet-proofness* of the *Randomized Condorcet Voting System (RCVS)*, i.e., of $C1$ - ML schemes, which corresponds to our notion of PC -strategyproofness for profiles with a Condorcet winner.

⁵In fact, this notion is often called *weak* SD -strategyproofness.

dorcet winner on the domain of preference profiles that admit a Condorcet winner is a strategyproof voting rule [see, e.g., Campbell and Kelly, 2003, Moulin, 1988, Lemma 10.3].

Theorem 5.2. *Let R be a preference profile that admits a Condorcet winner. Then, no ML scheme is PC -manipulable at R .*

This positive result for profiles that admit a Condorcet winner is contrasted by the following negative result, which shows that every $C2$ - ML scheme can be manipulated in every sufficiently diverse profile *without* a weak Condorcet winner.

Theorem 5.3. *Let R be a preference profile that does not admit a weak Condorcet winner and in which every preference relation appears at least once. Then, every $C2$ - ML scheme can be manipulated at R .*

The non-trivial proofs of both theorems are omitted because of space restrictions. With minor adaptations in the proof, it can be shown that Theorem 5.3 holds for every ML scheme that is based on a *strictly* monotonic function τ .

6 Experimental Results

We have conducted various experiments concerning the manipulability and the degree of randomization of ML schemes. For these experiments, we confined ourselves to profiles with an odd number of agents with strict preferences. Recall that for every such profile R , both $C1$ - $ML(R)$ and $C2$ - $ML(R)$ contain a unique lottery with odd support size. Hence, on this domain, all $C1$ - ML schemes coincide and all $C2$ - ML schemes coincide. Within this section, we thus refer to the unique $C1$ - ML scheme and the unique $C2$ - ML scheme as $C1$ - ML and $C2$ - ML , respectively.

The stochastic preference model used for our experiments is called *impartial anonymous culture (IAC)*. Under IAC, preference profiles form equivalence classes with two profiles belonging to the same class if they are identical up to permutations of the agents. Every equivalence class is assumed to be equally likely. Impartial culture models are known to significantly underestimate the likelihood of Condorcet winners [see, e.g., Regenwetter et al., 2006]. Our results on the support size, degree of randomization, and manipulability of ML schemes under IAC can thus be interpreted as upper bounds. In real world settings, one would expect lower numbers.⁷

6.1 Support Sizes

For an odd number of agents with strict preferences, an ML scheme returns a degenerate lottery if and only if there is a Condorcet winner. Hence, the frequency of cases where maximal lotteries do not randomize can be directly obtained by looking at the number of profiles that admit a Condorcet winner [see, e.g., Gehrlein, 2002, Regenwetter et al.,

⁷Gehrlein and Lepelley [2011] summarize 37 empirical studies from 1955 to 2009 and concluded that “there is a possibility that Condorcet’s Paradox might be observed, but that it probably is not a widespread phenomenon.” Similar conclusions were drawn by Regenwetter et al. [2007], Mattei et al. [2012], Tideman and Plassmann [2012]. Laslier [2010] and Brandt and Seedig [2014] report concrete probabilities for the existence of Condorcet winners under various distributional assumptions using computer simulations.

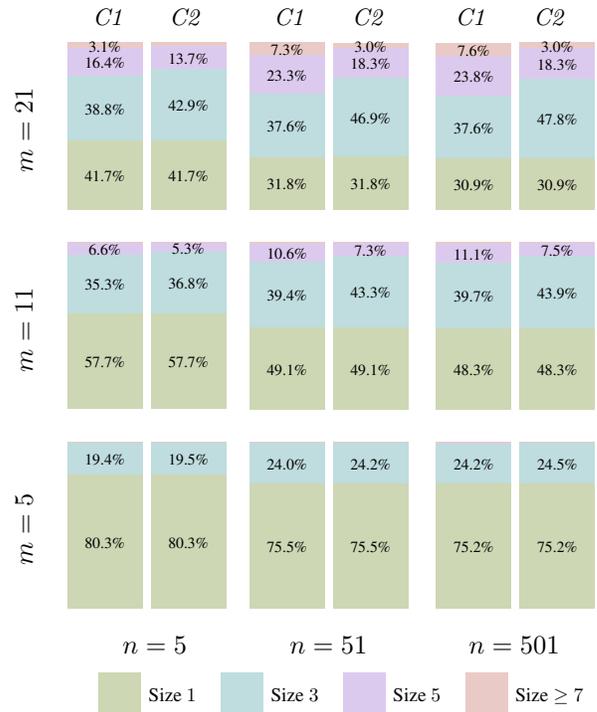


Figure 3: Distributions of the support sizes of $C1$ - ML (left) and $C2$ - ML (right) for $m \in \{5, 11, 21\}$ and $n \in \{5, 51, 501\}$ for 100.000 samples for every combination of parameters according to the IAC model. The bars represent the frequency of a support with size 1, 3, 5, and 7 or more stacked from bottom to top. Frequencies lower than 3% are not labeled.

2006, Laslier, 2010, Gehrlein and Lepelley, 2011, Brandt and Seedig, 2014]

Figure 3 shows the distribution of small supports of the two canonical ML schemes under IAC for different numbers of alternatives and agents, e.g., for $m = 21$ and $n = 501$, $C1$ - ML and $C2$ - ML randomize over 5 or more alternatives in 31.4% and 21.3% of the cases, respectively. Even with these extreme parameters, the average support sizes are still relatively low (3.18 and 2.87, respectively). In general, $C2$ - ML results in slightly lower average support sizes than $C1$ - ML .

The extreme case of disjoint $C1$ - ML and $C2$ - ML supports turns out to be extremely rare. We have found only a handful of these examples in millions of tested profiles. Moreover, the supports of $C1$ - ML and $C2$ - ML almost always coincide (largely due to the fact that Condorcet winners are likely to exist). We have not encountered a single $C1$ - ML SD -efficiency violation during our simulations.

6.2 Degree of Randomization

The support size is a rather crude value to measure the degree of randomization because it ignores the values of non-zero probabilities. We therefore evaluate the degree of randomization of a lottery $p \in \Delta(A)$ in terms of Shannon entropy $H(p) = -\sum_{x \in \text{supp}(p)} p(x) \log_b p(x)$. We set the basis of the logarithm $b = m$ to normalize the maximal entropy to 1, which is attained for the uniform lottery.

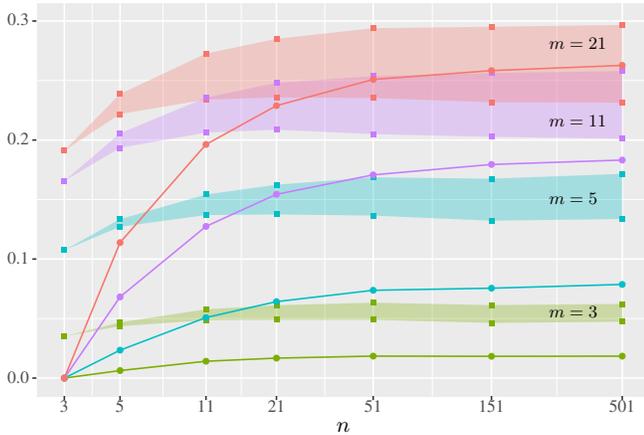


Figure 4: Average normalized Shannon entropies of ML schemes with the lower boundary of each colored area corresponding to $C2-ML$ and the upper boundary corresponding to $C1-ML$. Additionally, the average distances between these two ML schemes according to the L^1 -norm are displayed as solid lines. The same set of sampled profiles as in Figure 3 is used.

Figure 4 shows a considerable disparity between the average degree of randomization for $C1-ML$ and $C2-ML$, which widens as the number of agents or the number alternatives increases. This behavior is also reflected when plotting the average distance between $C1-ML$ and $C2-ML$ according to the L^1 -norm. Somewhat surprisingly, in contrast to $C1-ML$, the average degree of randomization of $C2-ML$ tends to slightly decline for large n .

6.3 Strategyproofness

The analytic results we obtained in Section 5 suggest a strong connection between the existence of Condorcet winners and manipulability of $C2-ML$ schemes. When the number of agents is large compared to the number of alternatives, there is a high probability that every preference relation appears at least once in a randomly chosen profile under the IAC model. Moreover, when there is an odd number of agents with strict preferences, the notions of weak and strict Condorcet winners coincide. Hence, one would expect that the probability that $C2-ML$ is SD -manipulable for a random profile converges to the probability that no Condorcet winner exists as the number of agents grows. Figure 5 shows the empirical frequency that $C2-ML$ is SD -manipulable in a randomly chosen profile under the IAC model for various numbers of alternatives and agents (solid lines). For a fixed number of alternatives, this frequency converges relatively quickly to the probability that no Condorcet winner exists (dashed lines), which confirms the above intuition. In addition, Figure 5 displays the frequency of profiles at which $C1-ML$ is SD -manipulable for 5 alternatives (dashed line). It is significantly lower than for $C2-ML$ for the same number of alternatives and decreases as the number of agents increases. This can be explained by observing that the absolute values of majority margins are unlikely to be 1 for large numbers of agents, which renders it improbable that a single agent can affect the majority relation, and thus the returned lottery, at all.

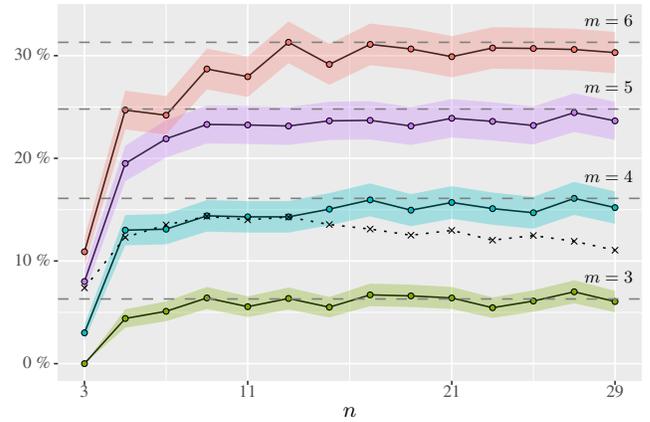


Figure 5: SD -Manipulability of $C2-ML$ (solid lines) with confidence intervals for a confidence level of 95% (in color) and manipulability of $C1-ML$ for $m = 5$ (dotted line). Each point is based on 2000 preference profiles sampled according to the IAC model and on testing all possible deviations for each type of voter. The dashed lines show the limit probabilities that no Condorcet winner exists for $n \rightarrow \infty$ derived by Gehrlein [2002].

7 Conclusion and Discussion

We have provided an extensive comparison of ML schemes. Our results can be used to guide the decision whether to use $C1-ML$ or $C2-ML$ schemes. We showed that $C2-ML$ schemes are the *only* SD -efficient ML schemes (and even satisfy the stronger notion of PC -efficiency). Moreover, we proved that *all* majoritarian and neutral SDSs violate SD -efficiency, a statement that may be of independent interest. $C1-ML$ efficiency failures are, however, extremely rare and may not appear in actual use. We also showed that, while *all* ML schemes are SD -manipulable, they are PC -strategyproof whenever a Condorcet winner exists (which is the case for most real-world preference profiles). Unsurprisingly, $C1-ML$ is less manipulable than $C2-ML$, simply because single agents are unlikely to be able to affect the outcome of majoritarian SDSs. This observation is connected to another potential drawback of $C1-ML$: it is less responsive than $C2-ML$. For $m = 3$ and odd n , $C1-ML$ will randomize with equal probabilities whenever there is no Condorcet winner. While uniform lotteries may be easier to implement in the real world and perhaps be more acceptable to the general public, this rigidity comes at a cost. Consider, for example, a profile for $m = 3$ and $n = 100$ such that 49 agents prefer c to a to b , 48 agents prefer a to b to c , and 3 agents prefer b to c to a . Then, $C1-ML$ selects b with probability $1/3$ even though 97 of 100 agents prefer a to b and no reasonable (deterministic) voting rule would select b . $C2-ML$, on the other hand, puts probability 0.04 on alternative b .⁸

Finally, we evaluated the degree of randomization of ML schemes via computer simulations. When there are only few alternatives, ML schemes usually do not randomize at all.

⁸In this example, Borda's rule would select a while most other rules (Kemeny's rule, maximin, Schulze's rule, ranked pairs, plurality, and all runoff rules) would select c . $C2-ML$ puts probability 0.94 on c .

For up to 21 alternatives, the average support size of ML schemes lies below 4 under reasonable assumptions. The average degree of randomization (in terms of Shannon entropy) of $C2-ML$ schemes is significantly lower than that of $C1-ML$ schemes.

Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/11-1.

References

- H. Aziz, F. Brandt, and M. Brill. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proc. of 12th AAMAS Conference*, pages 455–462. IFAAMAS, 2013.
- F. Brandl and F. Brandt. Arrovian aggregation of convex preferences. 2018. Working Paper.
- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.
- F. Brandl, F. Brandt, M. Eberl, and C. Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, 65(2), 2018a.
- F. Brandl, F. Brandt, and J. Hofbauer. Welfare maximization entices participation. *Games and Economic Behavior*, 2018b. Forthcoming.
- F. Brandt and H. G. Seedig. On the discriminative power of tournament solutions. In *Proc. of 1st AAMAS Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE)*, 2014.
- F. Brandt, C. Geist, and P. Harrenstein. A note on the McKelvey uncovered set and Pareto optimality. *Social Choice and Welfare*, 46(1):81–91, 2016.
- F. Brandt, M. Brill, H. G. Seedig, and W. Suksompong. On the structure of stable tournament solutions. *Economic Theory*, 65(2):483–507, 2018.
- D. E. Campbell and J. S. Kelly. A strategy-proofness characterization of majority rule. *Economic Theory*, 22(3):557–568, 2003.
- B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
- D. S. Felsenthal and M. Machover. After two centuries should Condorcet’s voting procedure be implemented? *Behavioral Science*, 37(4):250–274, 1992.
- P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984a.
- P. C. Fishburn. Dominance in SSB utility theory. *Journal of Economic Theory*, 34(1):130–148, 1984b.
- D. C. Fisher and J. Ryan. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2):217–236, 1995.
- W. V. Gehrlein. Condorcet’s paradox and the likelihood of its occurrence: Different perspectives on balanced preferences. *Theory and Decision*, 52(2):171–199, 2002.
- W. V. Gehrlein and D. Lepelley. *Voting Paradoxes and Group Coherence*. Studies in Choice and Welfare. Springer-Verlag, 2011.
- A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- L. N. Hoang. Strategy-proofness of the randomized Condorcet voting system. *Social Choice and Welfare*, 48(3):679–701, 2017.
- G. Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gössing, Austria (3–27 July 1962)*, pages 73–79, 1965.
- G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993.
- G. Laffond, J. Lainé, and J.-F. Laslier. Composition-consistent tournament solutions and social choice functions. *Social Choice and Welfare*, 13(1):75–93, 1996.
- G. Laffond, J.-F. Laslier, and M. Le Breton. A theorem on symmetric two-player zero-sum games. *Journal of Economic Theory*, 72(2):426–431, 1997.
- J.-F. Laslier. Aggregation of preferences with a variable set of alternatives. *Social Choice and Welfare*, 17(2):269–282, 2000.
- J.-F. Laslier. In silico voting experiments. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, chapter 13, pages 311–335. Springer-Verlag, 2010.
- N. Mattei, J. Forshee, and J. Goldsmith. An empirical study of voting rules and manipulation with large datasets. In *Proc. of 4th COMSOC*, 2012.
- H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- R. Peyre. La quête du graal électoral. *Images des Mathématiques*, CNRS, 2013.
- M. Regenwetter, B. Grofman, A. A. J. Marley, and I. M. Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006.
- M. Regenwetter, A. Kim, A. Kantor, and M.-H. R. Ho. The unexpected empirical consensus among consensus methods. *Psychological Science*, 18(7):629–635, 2007.
- R. L. Rivest and E. Shen. An optimal single-winner preferential voting system based on game theory. In *Proc. of 3rd COMSOC*, pages 399–410, 2010.
- N. Tideman and F. Plassmann. Modeling the outcomes of vote-casting in actual elections. In D. S. Felsenthal and M. Machover, editors, *Electoral Systems: Paradoxes, Assumptions, and Procedures*, Studies in Choice and Welfare. Springer, 2012.