

Consistent Probabilistic Social Choice

Florian Brandl Felix Brandt Hans Georg Seedig
Technische Universität München
{brandlf1,brandtf,seedigh}@in.tum.de

Two fundamental axioms in social choice theory are consistency with respect to a variable electorate and consistency with respect to components of similar alternatives. In the context of traditional non-probabilistic social choice, these axioms are incompatible with each other. We show that in the context of *probabilistic* social choice, these axioms uniquely characterize a function proposed by Fishburn (Rev. Econ. Stud., 51(4), 683–692, 1984). Fishburn’s function returns so-called *maximal lotteries*, i.e., lotteries that correspond to optimal mixed strategies in the symmetric zero-sum game induced by the pairwise majority margins. Maximal lotteries are guaranteed to exist due to von Neumann’s Minimax Theorem, are almost always unique, and can be efficiently computed using linear programming.

1. Introduction

Many important properties in the theory of social choice concern the consistency of aggregation functions under varying parameters. What happens if two electorates are merged? How should an aggregation function deal with components of similar alternatives? How should choices from overlapping agendas be related to each other? These considerations have led to a number of consistency axioms that these functions should ideally satisfy.¹ Unfortunately, social choice theory is rife with impossibility results which have revealed the incompatibility of many of these properties. Young and Levenglick (1978), for example, have pointed out that every social choice function that satisfies Condorcet-consistency violates consistency with respect to variable electorates. On the other hand, it follows from results by Young (1975) and Laslier (1996) that all Pareto-optimal social choice functions

¹Consistency conditions have found widespread acceptance well beyond social choice theory and feature prominently in the characterizations of various concepts in mathematical economics such as proportional representation rules (Balinski and Young, 1978), Nash’s bargaining solution (Lensberg, 1988), the Shapley value (Hart and Mas-Colell, 1989), and Nash equilibrium (Peleg and Tijs, 1996). Young (1994) and Thomson (2014) provide excellent overviews and give further examples.

that are consistent with respect to variable electorates are inconsistent with respect to components of similar alternatives.

The main result of this paper is that, in the context of *probabilistic* social choice, consistency with respect to variable electorates and consistency with respect to components of similar alternatives uniquely characterize an appealing probabilistic social choice function, which furthermore satisfies Condorcet-consistency. Probabilistic social choice functions yield lotteries over alternatives (rather than sets of alternatives) and were first formally studied by Zeckhauser (1969), Fishburn (1972), and Intriligator (1973). Perhaps one of the best known results in this context is Gibbard's characterization of strategyproof probabilistic social choice functions (Gibbard, 1977). An important corollary of Gibbard's characterization, attributed to Hugo Sonnenschein, is that *random dictatorships* are the only strategyproof and *ex post* efficient probabilistic social choice functions. In random dictatorships, one of the voters is picked at random and his most preferred alternative is implemented as the social choice. While Gibbard's theorem might seem as an extension of classic negative results on strategyproof non-probabilistic social choice functions (Gibbard, 1973; Satterthwaite, 1975), it is in fact much more positive (see also Barberà, 1979). In contrast to deterministic dictatorships, the uniform random dictatorship (henceforth, random dictatorship), in which every voter is picked with the same probability, enjoys a high degree of fairness and is in fact used in many subdomains of social choice that are concerned with the fair assignment of objects to agents (see, e.g., Abdulkadiroğlu and Sönmez, 1998; Bogomolnaia and Moulin, 2004; Che and Kojima, 2010; Budish et al., 2013).

Summary of Results

In this paper, we consider two consistency axioms, non-probabilistic versions of which have been widely studied in the literature. The first one, *population-consistency*, requires that, whenever two disjoint electorates agree on a lottery, this lottery should also be chosen by the union of both electorates. The second axiom, *composition-consistency*, requires that the probability that an alternative receives is unaffected by introducing new variants of another alternative. Alternatives are variants of each other if they form a component, i.e., they bear the same relationship to all other alternatives. On top of that, composition-consistency prescribes that the probability of an alternative within a component should be directly proportional to the probability that the alternative receives when the component is considered in isolation. Apart from their intuitive appeal, these axioms can be motivated by the desire to prevent a central planner from strategically partitioning the electorate into subelectorates or by deliberately introducing similar variants of alternatives, respectively. Our first result shows that there is no non-probabilistic social choice function that satisfies both axioms simultaneously. We then move to probabilistic social choice and prove that the only probabilistic social choice function satisfying these properties is the function that returns all *maximal lotteries* for a given preference profile. Maximal lotteries, which were proposed independently by Fishburn (1984) and other authors, are equivalent to mixed maximin strategies of the symmetric zero-sum game given by the pairwise majority mar-

gins. Whenever there is an alternative that is preferred to any other alternative by some majority of voters (a so-called Condorcet winner), the lottery that assigns probability 1 to this alternative is the unique maximal lottery. In other words, maximal lotteries satisfy *Condorcet-consistency*. At the same time, maximal lotteries satisfy population-consistency which has been identified by Young (1974a), Nitzan and Rubinstein (1981), Saari (1990b), and others as *the* defining property of Borda’s scoring rule. As such, the characterization can be seen as one possible resolution of the historic dispute between the founding fathers of social choice theory, the Chevalier de Borda and the Marquis de Condorcet, which dates back to the 18th century.²

Random dictatorship, the most prevalent probabilistic social choice function, satisfies population-consistency, but fails to satisfy composition-consistency. For this reason, we also consider a weaker version of composition-consistency called cloning-consistency, which is satisfied by random dictatorship, and provide an alternative characterization of maximal lotteries using population-consistency, cloning-consistency, and Condorcet-consistency (see Remark 5).

Acceptability of Social Choice Lotteries

Clearly, allowing lotteries as outcomes for high-stakes political elections such as those for the U.S. presidency would be highly controversial and considered by many a failure of deliberative democracy. If, on the other hand, a small group of people repeatedly votes on where to hold their next meeting, randomization would likely be more acceptable and perhaps even desirable. The use of lotteries for the selection of officials interestingly goes back to the world’s first democracy in Athens where it was widely regarded as a principal characteristic of democracy (Headlam, 1933). It has also been early observed in the social choice literature that “unattractive social choices may result whenever lotteries are not allowed to compete. [...] Refusal to entertain lotteries on alternatives can lead to outcomes that to many appear to be inequitable and perhaps even inefficient” (Zeckhauser, 1969).³ In contemporary research, probabilistic social choice has gained increasing interest in both social choice (see, e.g., Ehlers et al., 2002; Bogomolnaia et al., 2005; Chatterji et al., 2014) and political science (see, e.g., Goodwin, 2005; Dowlen, 2009; Stone, 2011).

²In this sense, our main theorem is akin to the characterization of Kemeny’s rule by Young and Levenglick (1978). Young and Levenglick considered *social preference functions*, i.e., functions that return sets of *rankings* of alternatives, and showed that Kemeny’s rule is characterized by strong versions of population-consistency and Condorcet-consistency.

Interestingly, all three rules—Borda’s rule, Kemeny’s rule, and maximal lotteries—maximize aggregate score in a well-defined sense. For maximal lotteries, this is the case because they maximize social welfare according to the canonical skew-symmetric bilinear (SSB) utility functions representing the voters ordinal preferences (Brandl et al., 2017). SSB utility theory goes back to Fishburn (1982) and is a generalization of von Neumann-Morgenstern utility theory.

³It is interesting to note that the “intransitivity difficulties” that Zeckhauser (1969) examines in the context of probabilistic social choice disappear when replacing majority rule with *expected* majority rule. This directly leads to Fishburn’s definition of maximal lotteries.

Whether lotteries are socially acceptable depends on many factors, only some of which are based on formal arguments. In our view, two important factors are (i) *the effective degree of randomness* and (ii) *risk aversion on behalf of the voters*.

As to (i), it is easily seen that certain cases call for randomization or other means of tie-breaking. For example, if there are two alternatives, a and b , and exactly half of the voters prefer a while the other half prefers b , there is no deterministic way of selecting a single alternative without violating basic fairness conditions. There are several possibilities to extend the notion of a tied outcome to three or more alternatives. An important question in this context is whether ties are a rare exception or a common phenomenon. A particularly rigorous and influential extension due to Condorcet declares a tie in the absence of a pairwise majority winner. According to Condorcet, it is the intransitivity of social preferences, as exhibited in the Condorcet paradox, that leads to situations in which there is no unequivocal winner. As it turns out, maximal lotteries are degenerate if and only if there is a Condorcet winner. Our main result thus establishes that the degree of randomness entailed by our axioms is precisely in line with Condorcet’s view of equivocality. Interestingly, there is strong empirical and experimental evidence that most real-world preference profiles for *political* elections do admit a Condorcet winner (see, e.g., Regenwetter et al., 2006; Laslier, 2010; Gehrlein and Lepelley, 2011).⁴ Maximal lotteries only randomize in the less likely case of cyclical majorities. Brandt and Seedig (2014) specifically analyzed the support of maximal lotteries and found that the average support size is less than four under various distributional assumptions and up to 30 alternatives. By contrast, random dictatorship randomizes over all alternatives in almost all elections.

As to (ii), risk aversion is strongly related to the frequency of preference aggregation. If an aggregation procedure is not frequently repeated, the law of large numbers does not apply and risk-averse voters might prefer a sure outcome to a lottery whose expectation they actually prefer to the sure outcome. Hence, probabilistic social choice seems particularly suitable for novel preference aggregation settings that have been made possible by technological advance. The Internet, in particular, allows for much more frequent preference aggregation than traditional paper-and-pencil elections. In recurring randomized elections with a fixed set of alternatives, voters need not resubmit their preferences for

⁴Analytical results for the likelihood of Condorcet winners are typically based on the simplistic “impartial culture” model, which assumes that every preference relation is equally likely. According to this model, a Condorcet winner, for example, exists with a probability of at least 63% when there are seven alternatives (Fishburn, 1973). The impartial culture model is considered highly unrealistic and Regenwetter et al. (2006) argued that it significantly underestimates the probability of Condorcet winners. Gehrlein and Lepelley (2011) summarized 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.” Laslier (2010) and Brandt and Seedig (2014) reported concrete probabilities for the existence of Condorcet winners under various distributional assumptions using computer simulations. A common observation in these studies is that the probability of a Condorcet winner generally decreases with increasing number of alternatives. For example, Brandt and Seedig found that, for 15 voters and a spatial distribution of preferences that is commonly used in political science, the probability of a Condorcet winner ranges from 98% (for three alternatives) to 59% (for 50 alternatives).

every election; rather preferences can be stored and only changed if desired. For example, maximal lotteries could help a group of coworkers with the daily decision of where to have lunch without requiring them to submit their preferences every day. Another example are automatic music broadcasting systems, such as Internet radio stations or software DJs, that decide which song should be played next based on the preferences of the listeners. In contrast to traditional deterministic solutions to these problems such as sequential dictatorships, the repeated execution of lotteries is a memoryless process that guarantees *ex ante* fairness after any number of elections.

Finally, it should be noted that the lotteries returned by probabilistic social choice functions do not necessarily have to be interpreted as probability distributions. They can, for instance, also be seen as fractional allocations of divisible objects such as time shares or monetary budgets. The axioms considered in this paper are equally natural for these interpretations as they are for the probabilistic interpretation.

2. Preliminaries

Let U be an infinite universal set of alternatives. The set of *agendas* from which alternatives are to be chosen is the set of finite and non-empty subsets of U , denoted by $\mathcal{F}(U)$. The set of all linear (i.e., complete, transitive, and antisymmetric) *preference relations* over some set $A \in \mathcal{F}(U)$ will be denoted by $\mathcal{L}(A)$.

For some finite set X , we denote by $\Delta(X)$ the set of all probability distributions with rational values over X . A (fractional) preference profile R for a given agenda A is an element of $\Delta(\mathcal{L}(A))$, which can be associated with the $(|A|! - 1)$ -dimensional (rational) unit simplex. We interpret $R(\succ)$ as the fraction of voters with preference relation $\succ \in \mathcal{L}(A)$. Preference profiles are depicted by tables in which each column represents a preference relation \succ with $R(\succ) > 0$. The table below shows an example profile on three alternatives.⁵

	1/2	1/3	1/6	
a	a	a	b	(Example 1)
b	c	c		
c	b	a		

The set of all preference profiles for a fixed agenda A is denoted by $\mathcal{R}|_A$ and \mathcal{R} is defined as the set of all preference profiles, i.e., $\mathcal{R} = \bigcup_{A \in \mathcal{F}(U)} \mathcal{R}|_A$. For $B \subseteq A$ and $R \in \mathcal{R}|_A$, $R|_B$ is the restriction of R to alternatives in B , i.e., for all $\succ \in \mathcal{L}(B)$,

$$R|_B(\succ) = \sum_{\succ' \in \mathcal{L}(A): \succ \subseteq \succ'} R(\succ').$$

⁵Our representation of preference profiles implicitly assumes that aggregation functions are anonymous (i.e., all voters are treated identically) and homogeneous (i.e., duplication of the electorate does not affect the outcome). Similar models (sometimes even assuming a continuum of voters) have for example been considered by Young (1974b, 1975), Young and Levenglick (1978), Saari (1995), Dasgupta and Maskin (2008), Che and Kojima (2010), and Budish and Cantillon (2012).

For all $x, y \in A$, $R(x, y) = R|_{\{(x, y)\}}$ is the fraction of voters who prefer x to y (the set $\{(x, y)\}$ represents the preference relation on two alternatives with $x \succ y$). In Example 1, $R(a, b) = 5/6$.

Elements of $\Delta(A)$ are called *lotteries* and will be written as convex combinations of alternatives. If p is a lottery, p_x is the probability that p assigns to alternative x .

A *probabilistic social choice function (PSCF)* f is an (upper hemi-) continuous function that, for any agenda $A \in \mathcal{F}(U)$, maps a preference profile $R \in \mathcal{R}|_A$ to a convex subset of $\Delta(A)$.⁶ A PSCF is thus a collection of mappings from high-dimensional simplices to low-dimensional simplices. Two further properties that we demand from any PSCF are *unanimity* and *decisiveness*. Unanimity states that in the case of one voter and two alternatives, the preferred alternative should be chosen with probability 1.⁷ Since we only consider fractional preference profiles, this amounts to for all $x, y \in U$ and $R \in \mathcal{R}|_{\{x, y\}}$,

$$f(R) = \{x\} \text{ whenever } R(x, y) = 1. \quad (\text{unanimity})$$

Decisiveness requires that the set of preference profiles where f is multi-valued is negligible in the sense that for all $A \in \mathcal{F}(U)$,

$$\{R \in \mathcal{R}|_A : |f(R)| = 1\} \text{ is dense in } \mathcal{R}|_A. \quad (\text{decisiveness})$$

In other words, for every preference profile that yields multiple lotteries, there is an arbitrarily close preference profile that only yields a single lottery.

Probabilistic social choice functions considered in the literature usually satisfy these conditions and are therefore well-defined PSCFs. For example, consider *random dictatorship (RD)*, in which one voter is picked uniformly at random and his most-preferred alternative is returned. Formally, RD returns the unique lottery, which is determined by multiplying fractions of voters with their respective top choices, i.e., for all $A \in \mathcal{F}(U)$ and $R \in \mathcal{R}|_A$,

$$RD(R) = \left\{ \sum_{\succ \in \mathcal{L}(A)} R(\succ) \cdot \max_{\succ} (A) \right\}, \quad (\text{random dictatorship})$$

where $\max_{\succ}(A)$ denotes the unique alternative x such that $x \succ y$ for all $y \in A$. For the preference profile R given in Example 1,

$$RD(R) = \{5/6 a + 1/6 b\}.$$

⁶Fishburn (1973, pp. 248–249) argued that the set of lotteries returned by a probabilistic social choice function should be convex because it would be unnatural if two lotteries are socially acceptable while a randomization between them is not (see also Fishburn, 1972, p. 201).

⁷This is the only condition we impose that actually interprets the preference relations. It is equivalent to *ex post* efficiency for agendas of size 2 and is slightly weaker than Young’s faithfulness (Young, 1974a). Our results still hold when replacing unanimity with the even less controversial condition that merely requires that $f(R) \neq \{y\}$ whenever $R(x, y) = 1$.

RD is single-valued and therefore trivially decisive and convex-valued. It is also easily seen that RD satisfies unanimity and continuity and thus constitutes a PSCF.

A useful feature of our definition of PSCFs is that traditional *set-valued* social choice functions (SCFs) (also known as *social choice correspondences*) can be seen as a special case, namely as those PSCFs that map every preference profile $R \in \mathcal{R}|_A$ on some agenda $A \in \mathcal{F}(U)$ to $\Delta(X)$ for some $X \subseteq A$. Such PSCFs will be called (*non-probabilistic*) SCFs.

3. Population-consistency and Composition-consistency

The consistency conditions we consider are generalizations of the corresponding conditions for SCFs, i.e., the axioms coincide with their non-probabilistic counterparts.

The first axiom relates choices from varying electorates to each other. More precisely, it requires that whenever a lottery is chosen simultaneously by two electorates, this lottery is also chosen by the union of both electorates. For example, consider the two preference profiles R' and R'' given below.

$\frac{1}{2}$ $\frac{1}{2}$	$\frac{1}{2}$ $\frac{1}{2}$	$\frac{1}{4}$ $\frac{1}{4}$ $\frac{1}{2}$	
a b	a b	a a b	
b c	c c	b c c	
c a	b a	c b a	(Example 2)
R'	R''	$\frac{1}{2} R' + \frac{1}{2} R''$	

Population-consistency then demands that any lottery that is chosen in both R' and R'' (say, $\frac{1}{2}a + \frac{1}{2}b$) also has to be chosen when both preference profiles are merged. Formally, a PSCF satisfies population-consistency if for all $A \in \mathcal{F}(U)$, $R', R'' \in \mathcal{R}|_A$, and any convex combination R of R' and R'' , i.e., $R = \lambda R' + (1 - \lambda)R''$ for some $\lambda \in [0, 1]$,

$$f(R') \cap f(R'') \subseteq f(R). \quad (\text{population-consistency})$$

Population-consistency is arguably one of the most natural axioms for variable electorates and is usually considered in a slightly stronger version, known as *reinforcement* or simply *consistency*, where the inclusion in the equation above is replaced with equality whenever the left-hand side is non-empty (see also Remark 4). Note that population-consistency is merely a statement about abstract sets of outcomes, which makes no reference to lotteries whatsoever. It was first considered independently by Smith (1973), Young (1974a), and Fine and Fine (1974) and features prominently in the characterization of scoring rules by Smith (1973) and Young (1975). Population-consistency and its variants have found widespread acceptance in the social choice literature (see, e.g., Young, 1974b; Fishburn, 1978; Young and Levenglick, 1978; Saari, 1990a, 1995; Myerson, 1995; Congar and Merlin, 2012).

The second axiom prescribes how PSCFs should deal with *decomposable* preference profiles. For two agendas $A, B \in \mathcal{F}(U)$, $B \subseteq A$ is a component in $R \in \mathcal{R}|_A$ if the alternatives

in B are *adjacent* in all preference relations that appear in R , i.e., for all $a \in A \setminus B$ and $b, b' \in B$, $a \succ b$ if and only if $a \succ b'$ for all $\succ \in \mathcal{L}(A)$ with $R(\succ) > 0$. Intuitively, the alternatives in B can be seen as variants or clones of the same alternative because they have exactly the same relationship to all alternatives that are not in B . For example, consider the following preference profile R in which $B = \{b, b'\}$ constitutes a component.

1/3	1/6	1/2		1/2	1/2		1/3	2/3		
a	a	b		a	b		b'	b		
b'	b	b'		b	a		b	b'		(Example 3)
b	b'	a								
R				$R _{A'}$			$R _B$			

The ‘essence’ of R is captured by $R|_{A'}$, where $A' = \{a, b\}$ contains only one of the cloned alternatives. It seems reasonable to demand that a PSCF should assign the same probability to a (say, $1/2$) independently of the number of clones of b and the internal relationship between these clones. This condition will be called cloning-consistency and was first proposed by Tideman (1987) (see also Zavist and Tideman, 1989). Its origins can be traced back to earlier, more general, decision-theoretic work by Arrow and Hurwicz (1972) and Maskin (1979) where it is called *deletion of repetitious states* as well as early work on majoritarian SCFs by Moulin (1986). For a formal definition of cloning-consistency, let $A', B \in \mathcal{F}(U)$ and $A = A' \cup B$ such that $A' \cap B = \{b\}$. Then, a PSCF f satisfies cloning-consistency if, for all $R \in \mathcal{R}|_A$ such that B is a component in R ,

$$\{(p_x)_{x \in A \setminus B} : p \in f(R)\} = \{(p_x)_{x \in A \setminus B} : p \in f(R|_{A'})\}.$$

When having a second look at Example 3, it may appear strange that cloning-consistency does not impose any restrictions on the probabilities that f assigns to the clones. While clones behave completely identical with respect to uncloned alternatives, they are not indistinguishable from *each other*. It seems that the relationships between clones ($R|_B$) should be taken into account as well. For example, one would expect that f assigns more probability to b than to b' because two thirds of the voters prefer b to b' . An elegant and mathematically appealing way to formalize this intuition is to require that the probabilities of the clones b and b' are directly proportional to the probabilities that f assigns to these alternatives when restricting the preference profile to the component $\{b, b'\}$. This condition, known as *composition-consistency*, is due to Laffond et al. (1996) and was studied in detail for majoritarian SCFs (see, e.g., Laslier, 1996, 1997; Brandt, 2011; Brandt et al., 2011; Horan, 2013).⁸

For a formal definition of composition-consistency, let $p \in \Delta(A')$ and $q \in \Delta(B)$ and

⁸More generally, modular decompositions of discrete structures have found widespread applications in operations research and combinatorial optimization (see, e.g., Möhring, 1985).

define

$$(p \times_b q)_x = \begin{cases} p_x & \text{if } x \in A \setminus B, \\ p_b q_x & \text{if } x \in B. \end{cases}$$

The operator \times_b is extended to sets of lotteries $X \subseteq \Delta(A')$ and $Y \subseteq \Delta(B)$ by applying it to all pairs of lotteries in $X \times Y$, i.e., $X \times_b Y = \{p \times_b q \in \Delta(A): p \in X \text{ and } q \in Y\}$.

Then, a PSCF f satisfies composition-consistency if for all $R \in \mathcal{R}|_A$ such that B is a component in R ,

$$f(R|_{A'}) \times_b f(R|_B) = f(R). \quad (\text{composition-consistency})$$

In Example 3 above, $1/2 a + 1/2 b \in f(R|_{A'})$, $2/3 b + 1/3 b' \in f(R|_B)$, and composition-consistency would imply that $1/2 a + 1/3 b + 1/6 b' \in f(R)$.

4. Non-probabilistic Social Choice

In the context of SCFs (i.e., PSCFs that only return the convex hull of degenerate lotteries), there is some friction between population-consistency and composition-consistency. In fact, the conflict between these notions can be traced back to the well-documented dispute between the pioneers of social choice theory, the Chevalier de Borda and the Marquis de Condorcet (see, e.g., Black, 1958; Young, 1988, 1995; McLean and Hewitt, 1994). Borda proposed a score-based voting rule—*Borda's rule*—that can be axiomatically characterized using population-consistency (Young, 1974a). It then turned out that the entire class of scoring rules (which apart from Borda's rule also includes plurality rule) is characterized by population-consistency (Smith, 1973; Young, 1974b, 1975). Condorcet, on the other hand, advocated *Condorcet-consistency*, which requires that an SCF selects a Condorcet winner whenever one exists. As Condorcet already pointed out, Borda's rule fails to be Condorcet-consistent. Worse, Young and Levenglick (1978) even showed that no Condorcet-consistent SCF satisfies population-consistency (the defining property of scoring rules).⁹ Laslier (1996), on the other hand, showed that no Pareto-optimal rank-based rule—a generalization of scoring rules—satisfies composition-consistency while this property is satisfied by various Condorcet-consistent SCFs (Laffond et al., 1996). One of the few SCFs that satisfies both properties is the rather indecisive Pareto rule (which returns all alternatives that are not Pareto-dominated). Since our definition of PSCFs already incorporates a certain degree of decisiveness, we obtain the following impossibility. (The proofs of all theorems are deferred to the Appendix.)

Theorem 1. *There is no SCF that satisfies population-consistency and composition-consistency.*¹⁰

⁹Theorem 2 by Young and Levenglick (1978) actually uses the strong variant of population-consistency, but their proof also holds for population-consistency as defined in this paper.

¹⁰Theorem 1 still holds when replacing composition-consistency with the weaker condition of cloning-consistency.

In light of this result, it is perhaps surprising that, for probabilistic social choice, both axioms are not only mutually compatible but even *uniquely* characterize a PSCF.

5. Characterization of Maximal Lotteries

Maximal lotteries were first considered by Kreweras (1965) and independently proposed and studied in more detail by Fishburn (1984). Interestingly, maximal lotteries or variants thereof have been rediscovered again by economists (Laffond et al., 1993b),¹¹ mathematicians (Fisher and Ryan, 1995), political scientists (Felsenthal and Machover, 1992), and computer scientists (Rivest and Shen, 2010).¹²

In order to define maximal lotteries, we need some notation. For $A \in \mathcal{F}(U)$, $R \in \mathcal{R}|_A$, $x, y \in A$, the entries $M_R(x, y)$ of the *majority margin matrix* M_R denote the difference between the fraction of voters who prefer x to y and the fraction of voters who prefer y to x , i.e.,

$$M_R(x, y) = R(x, y) - R(y, x).$$

Thus, M_R is skew-symmetric and $M_R \in [-1, 1]^{A \times A}$. A (*weak*) *Condorcet winner* is an alternative x that is *maximal* in A according to M_R in the sense that $M_R(x, y) \geq 0$ for all $y \in A$. If $M_R(x, y) > 0$ for all $y \in A \setminus \{x\}$, x is called a *strict Condorcet winner*. A PSCF is *Condorcet-consistent* if $x \in f(R)$ whenever x is a Condorcet winner in R .

It is well known from the Condorcet paradox that maximal elements may fail to exist. As shown by Kreweras (1965) and Fishburn (1984), this drawback can, however, be remedied by considering lotteries over alternatives. For two lotteries $p, q \in \Delta(A)$, the majority margin can be extended to its bilinear form $p^T M_R q$, the *expected majority margin*. The set of maximal lotteries is then defined as the set of “probabilistic Condorcet winners.”

¹¹Laffond et al. (1993b, 1996), Dutta and Laslier (1999), and others have explored the *support* of maximal lotteries, called the *bipartisan set* or the *essential set*, in some detail. Laslier (2000) has characterized the essential set using monotonicity, Fishburn’s C2, regularity, inclusion-minimality, the strong superset property, and a variant of composition-consistency.

¹²Felsenthal and Machover (1992) and Rivest and Shen (2010) also discussed whether maximal lotteries are suitable for real-world political elections. Rivest and Shen concluded 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.” Felsenthal and Machover wrote that “an inherent special feature of [maximal lotteries] is its extensive and essential reliance on probability in selecting the winner [...] Without sufficient empirical evidence it is impossible to say whether this feature of [maximal lotteries] makes it socially less acceptable than other majoritarian procedures. It is not at all a question of fairness, for nothing could be fairer than the use of lottery as prescribed by [maximal lotteries]. The problem is whether society will accept such an extensive reliance on chance in public decision-making. Different societies may have differing views about this. For example, it is well known that the free men of ancient Athens regarded it as quite acceptable to select holders of public office by lot. Clearly, before [the maximal lotteries system] can be applied in practice, public opinion must first be consulted, and perhaps educated, on this issue.”

Formally, for all $A \in \mathcal{F}(U)$ and $R \in \mathcal{R}|_A$,¹³

$$ML(R) = \{p \in \Delta(A) : p^T M_R q \geq 0 \text{ for all } q \in \Delta(A)\}. \quad (\text{maximal lotteries})$$

As an example, consider the preference profile given in Example 1 of Section 2. Alternative a is a strict Condorcet winner and $ML(R) = \{a\}$. This is in contrast to $RD(R) = \{5/6 a + 1/6 b\}$, which puts positive probability on any first-ranked alternative no matter how small the corresponding fraction of voters.

The Minimax Theorem implies that $ML(R)$ is non-empty for all $R \in \mathcal{R}$ (von Neumann, 1928). In fact, M_R can be interpreted as the payoff matrix of a symmetric zero-sum game and maximal lotteries as the mixed maximin strategies (or Nash equilibrium strategies) of this game. Hence, maximal lotteries can be efficiently computed via linear programming. Interestingly, $ML(R)$ is a singleton in almost all cases. In particular, this holds if there is an odd number of voters (Laffond et al., 1997; Le Breton, 2005). Moreover, we point out in Appendix C.1 that the set of preference profiles that yield a unique maximal lottery is open and dense, which implies that the set of profiles with multiple maximal lotteries is nowhere dense and thus negligible. As a consequence, ML satisfies decisiveness as well as the other properties we demand from a PSCF (such as unanimity, continuity, and convex-valuedness) and therefore constitutes a well-defined PSCF.

In contrast to the non-probabilistic case where majority rule is known as the only reasonable and fair SCF on two alternatives (May, 1952; Dasgupta and Maskin, 2008), there is an infinite number of such PSCFs even when restricting attention to only two alternatives (see, e.g., Saunders, 2010). Within our framework of fractional preference profiles, a PSCF on two alternatives can be seen as a convex-valued continuous correspondence from the unit interval to itself. Unanimity fixes the function values at the endpoints of the unit interval, decisiveness requires that the points where the function is multi-valued are isolated, and population-consistency implies that the function is monotonic. Two natural extreme cases of functions that meet these requirements are a probabilistic version of simple majority rule and the proportional lottery (Figure 1).¹⁴ Interestingly, these two extreme points are taken by maximal lotteries and random dictatorship as for all $x, y \in U$ and $R \in \mathcal{R}|_{\{x,y\}}$,

$$ML(R) = \begin{cases} \{x\} & \text{if } R(x, y) > 1/2, \\ \{y\} & \text{if } R(x, y) < 1/2, \\ \Delta(\{x, y\}) & \text{otherwise,} \end{cases} \quad \text{and} \quad RD(R) = \{R(x, y)x + R(y, x)y\}.$$

¹³Several authors apply the signum function to the entries of M_R before computing maximal lotteries. This is, for example, the case for Kreweras (1965), Felsenthal and Machover (1992), Laffond et al. (1993a), and Fisher and Ryan (1995). Maximal lotteries as defined in this paper were considered by Dutta and Laslier (1999), Laslier (2000), and Rivest and Shen (2010). Fishburn (1984) allowed the application of any odd function to the entries of M_R , which covers both variants as special cases.

¹⁴Fishburn and Gehrlein (1977) compared these two-alternative PSCFs on the basis of expected voter satisfaction and found that the simple majority rule outperforms the proportional rule.




Figure 1: Maximal lotteries and random dictatorship on two-element agendas.

When considering up to three alternatives and additionally taking composition-consistency into account, any such PSCF coincides with majority rule on two-element agendas.¹⁵ When allowing an arbitrary number of alternatives, the axioms completely characterize *ML*.

Theorem 2. *A PSCF f satisfies population-consistency and composition-consistency if and only if $f = ML$.*

The proof of Theorem 2 is rather involved yet quite instructive as it rests on a number of lemmas that might be of independent interest (see Appendix C). The high-level structure is as follows. The fact that *ML* satisfies population-consistency and composition-consistency follows relatively easily from basic linear algebra.

For the converse direction, we first show that population-consistency and composition-consistency characterize *ML* on two-element agendas. For agendas of more than two alternatives, we assume that f is a population-consistent and composition-consistent PSCF and then show that $f \subseteq ML$ and $ML \subseteq f$. The first statement takes up the bulk of the proof and is shown by assuming for contradiction that there is a preference profile for which f yields a lottery that is *not* maximal. We then identify a set of preference profiles with full dimension for which f returns the uniform lottery over a fixed subset of at least two alternatives and which has the uniform profile, i.e., the preference profile in which every preference relation is assigned the same fraction of voters, in its interior. Along the way we show that f has to be Condorcet-consistent for all preference profiles that are close to the uniform profile. It follows that there has to be an ε -ball of profiles around some strict Condorcet profile (close to the uniform profile), for which f returns the uniform lottery over a non-singleton subset of alternatives as well as the lottery with probability 1 on the

¹⁵This also shows that *RD* violates composition-consistency, which can be seen in Example 3 in Section 3. However, *RD* does satisfy cloning-consistency (see Remark 5).

Condorcet winner. This contradicts decisiveness. For the inclusion of ML in f , we take an arbitrary preference profile and an arbitrary vertex of the set of maximal lotteries for this profile and then construct a sequence of preference profiles that converges to the original profile and whose maximal lotteries converge to the specified maximal lottery. From $f \subseteq ML$ and continuity, we obtain that f has to select this lottery in the original preference profile. Finally, convexity implies that $ML \subseteq f$.

6. Remarks

We conclude the paper with a number of remarks.

Remark 1 (Independence of axioms). The axioms used in Theorem 2 are independent from each other. RD satisfies population-consistency, but violates composition-consistency (see also Remark 5). The same is true for Borda's rule. When defining ML^3 via the third power of majority margins $(M_R(x, y))^3$, ML^3 satisfies composition-consistency, but violates population-consistency.¹⁶ Also the properties implicit in the definition of PSCFs are independent. The PSCF that returns all maximal lotteries for the profile in which all preference relations are reversed violates unanimity but satisfies decisiveness, population-consistency, and composition-consistency. When not requiring decisiveness, returning all *ex post* efficient lotteries is consistent with the remaining axioms.¹⁷

Remark 2 (Size of Universe). The proof of Theorem 2 exploits the infinity of the universe. As stated in the previous remark, ML^3 satisfies composition-consistency and violates population-consistency. However, ML^3 does satisfy population-consistency when there are only up to three alternatives. This implies that the statement of Theorem 2 requires a universe that contains at least four alternatives.

Remark 3 (Uniqueness). The set of profiles in which ML is not single-valued is negligible in the sense specified in the definition of PSCFs. When extending the set of fractional profiles to the reals, it can also be shown that maximal lotteries are almost always unique by using an argument similar to that of Harsanyi (1973a).

Remark 4 (Strong population-consistency). ML does not satisfy the stronger version of population-consistency in which the set inclusion is replaced with equality (see Section 3).¹⁸ This can be seen by observing that *every* lottery is maximal for the union of any two electorates whose preferences are completely opposed to each other. When there are at least three alternatives, it is possible to find two such preference profiles which yield the same unique maximal lottery and strong population-consistency is violated. However, whenever ML is single-valued (which is almost always the case), strong

¹⁶Such variants of ML were already considered by Fishburn (1984). See also Footnote 13.

¹⁷Continuity and convexity can also be seen as implicit assumptions. Continuity is needed because the relative interior of ML satisfies all remaining axioms. Whether convexity is required is open.

¹⁸Strong population-consistency is quite demanding. It is for example violated by rather basic functions such as the Pareto rule.

population-consistency is equivalent to population-consistency and therefore satisfied by ML .

Remark 5 (Cloning-consistency and Condorcet-consistency). Requiring cloning-consistency instead of composition-consistency suffices for the proof of Theorem 2 when additionally demanding Condorcet-consistency. It is therefore possible to alternatively characterize ML using population-consistency, cloning-consistency, and Condorcet-consistency. As above, the axioms are independent from each other. ML^3 , as defined in Remark 1, satisfies all axioms except population-consistency. The PSCF that is identical to ML^3 for agendas of size 3 and otherwise identical to ML satisfies all axioms except cloning-consistency. RD satisfies all axioms except Condorcet-consistency.

Remark 6 (Agenda-consistency). ML also satisfies *agenda-consistency*, which requires that the set of all lotteries that are chosen from two overlapping agendas should be identical to the set of lotteries that are chosen from the union of both agendas (and whose support is contained in both agendas). The inclusion from left to right is known as Sen’s γ or *expansion*, whereas the inclusion from right to left is Sen’s α or *contraction* (Sen, 1971).¹⁹ Numerous impossibility results, including Arrow’s well-known theorem, have revealed that agenda-consistency is prohibitive in non-probabilistic social choice when paired with minimal further assumptions such as non-dictatorship and Pareto-optimality (e.g., Sen, 1977, 1986; Campbell and Kelly, 2002).²⁰

Remark 7 (Domains). In contrast to RD , which at least requires that every voter has a unique top choice, ML does not require the asymmetry, completeness, or even transitivity of individual preferences (and still satisfies population-consistency and composition-consistency in these more general domains). In the restricted domains of matching and house allocation, on the other hand, maximal lotteries are known as *popular mixed matchings* (Kavitha et al., 2011) or *popular random assignments* (Aziz et al., 2013b).

Remark 8 (Efficiency). It has already been observed by Fishburn (1984) that ML is *ex post* efficient, i.e., Pareto-dominated alternatives always receive probability zero in all maximal lotteries. Aziz et al. (2013a) strengthened this statement by showing that ML even satisfies SD -efficiency (also known as ordinal efficiency) as well as the even stronger notion of PC -efficiency (see Aziz et al., 2015). While RD also satisfies SD -efficiency, random serial dictatorship (the canonical generalization of RD to weak preferences) violates SD -efficiency (Bogomolnaia and Moulin, 2001; Bogomolnaia et al., 2005).

Remark 9 (Strategyproofness). RD is the only *ex post* efficient strategyproof PSCF (Gibbard, 1977). However, RD and ML violate *group*-strategyproofness (in fact, there exists no *ex post* efficient group-strategyproof PSCF). RD and ML satisfy the significantly weaker notion of ST -group-strategyproofness, which is violated by probabilistic extensions

¹⁹Sen’s α actually goes back to Chernoff (1954) and Nash (1950), where it is called *independence of irrelevant alternatives* (not to be confused with Arrow’s IIA). We refer to Monjardet (2008) for more details.

²⁰Pattanaik and Peleg (1986) obtained a similar impossibility for probabilistic social choice using an interpretation of Sen’s α that is stronger than ours.

of most common voting rules (Aziz et al., 2014). A very useful property of ML is that it cannot be manipulated in all preference profiles that admit a strict Condorcet winner (see Peyre, 2013).

Acknowledgements

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/7-2 and BR 2312/10-1. Early versions of this paper were presented at the 6th Seminar on Ordered Structures in Games and Decisions in Paris (November, 2012), the Workshop on Mathematics of Electoral Systems in Budapest (October, 2013), the 12th Meeting of the Society for Social Choice and Welfare in Boston (June, 2014), the Micro-Theory Seminar at the University of Glasgow (February, 2015), and the Algorithms Seminar at the University of Oxford (November, 2015). The authors thank Bhaskar Dutta, Christian Geist, Paul Harrenstein, Johannes Hofbauer, Olivier Hudry, Vincent Merlin, Bernard Monjardet, Hervé Moulin, Klaus Nehring, Clemens Puppe, Peyton Young, and William S. Zwicker for helpful comments. The authors are furthermore indebted to the editor and anonymous referees for their constructive feedback.

References

- A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- K. J. Arrow and L. Hurwicz. An optimality criterion of decision-making under ignorance. In C. F. Carter and J. L. Ford, editors, *Uncertainty and Expectations in Economics*. Basil Blackwell, 1972.
- H. Aziz, F. Brandt, and M. Brill. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 455–462. IFAAMAS, 2013a.
- H. Aziz, F. Brandt, and P. Stursberg. On popular random assignments. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, volume 8146 of *Lecture Notes in Computer Science (LNCS)*, pages 183–194. Springer-Verlag, 2013b.
- H. Aziz, F. Brandl, and F. Brandt. On the incompatibility of efficiency and strategyproofness in randomized social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 545–551. AAAI Press, 2014.
- H. Aziz, F. Brandl, and F. Brandt. Universal Pareto dominance and welfare for plausible utility functions. *Journal of Mathematical Economics*, 60:123–133, 2015.
- M. L. Balinski and H. P. Young. Stability, coalitions and schisms in proportional representation systems. *American Political Science Review*, 72(3):848–858, 1978.
- S. Barberà. Majority and positional voting in a probabilistic framework. *Review of Economic Studies*, 46(2):379–389, 1979.

- D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
- A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.
- A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- H. F. Bohnenblust, S. Karlin, and L. S. Shapley. Solutions of discrete, two-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 1, pages 51–72. Princeton University Press, 1950.
- F. Brandl, F. Brandt, and J. Hofbauer. Welfare maximization entices participation. 2017. Working paper.
- F. Brandt. Minimal stable sets in tournaments. *Journal of Economic Theory*, 146(4):1481–1499, 2011.
- F. Brandt and H. G. Seedig. On the discriminative power of tournament solutions. In *Proceedings of the 1st AAMAS Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE)*, 2014.
- F. Brandt, M. Brill, and H. G. Seedig. On the fixed-parameter tractability of composition-consistent tournament solutions. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 85–90. AAAI Press, 2011.
- E. Budish and E. Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *American Economic Review*, 102(5):2237–2271, 2012.
- E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- D. E. Campbell and J. S. Kelly. Impossibility theorems in the Arrovian framework. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 1. Elsevier, 2002.
- S. Chatterji, A. Sen, and H. Zeng. Random dictatorship domains. *Games and Economic Behavior*, 86:212–236, 2014.
- Y.-K. Che and F. Kojima. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica*, 78(5):1625–1672, 2010.
- H. Chernoff. Rational selection of decision functions. *Econometrica*, 22(4):422–443, 1954.
- R. Congar and V. Merlin. A characterization of the maximin rule in the context of voting. *Theory and Decision*, 72(1):131–147, 2012.
- P. Dasgupta and E. Maskin. On the robustness of majority rule. *Journal of the European Economic Association*, 6(5):949–973, 2008.

- O. Dowlen. Sorting out sortition: A perspective on the random selection of political officers. *Political Studies*, 57(2):298–315, 2009.
- B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
- L. Ehlers, H. Peters, and T. Storcken. Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences. *Journal of Economic Theory*, 105(2):408–434, 2002.
- D. S. Felsenthal and M. Machover. After two centuries should Condorcet’s voting procedure be implemented? *Behavioral Science*, 37(4):250–274, 1992.
- B. Fine and K. Fine. Social choice and individual ranking I. *Review of Economic Studies*, 41(127):303–323, 1974.
- P. C. Fishburn. Lotteries and social choices. *Journal of Economic Theory*, 5(2):189–207, 1972.
- P. C. Fishburn. *The Theory of Social Choice*. Princeton University Press, 1973.
- P. C. Fishburn. Axioms for approval voting: Direct proof. *Journal of Economic Theory*, 19(1):180–185, 1978.
- P. C. Fishburn. Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26(1):31–67, 1982.
- P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984.
- P. C. Fishburn and W. V. Gehrlein. Towards a theory of elections with probabilistic preferences. *Econometrica*, 45(8):1907–1924, 1977.
- D. C. Fisher and J. Ryan. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2):217–236, 1995.
- W. V. Gehrlein and D. Lepelley. *Voting Paradoxes and Group Coherence*. Studies in Choice and Welfare. Springer-Verlag, 2011.
- A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- B. Goodwin. *Justice by Lottery*. University of Chicago Press, 2005. 1st edition 1992.
- J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory*, 2(1):235–250, 1973a.
- J. C. Harsanyi. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1):1–23, 1973b.
- S. Hart and A. Mas-Colell. Potential, value and consistency. *Econometrica*, 57(3):589–614, 1989.

- J. W. Headlam. *Election by Lot at Athens*. Cambridge University Press, 1933.
- S. Horan. Implementation of majority voting rules. 2013. Working paper.
- M. D. Intriligator. A probabilistic model of social choice. *Review of Economic Studies*, 40(4): 553–560, 1973.
- T. Kavitha, J. Mestre, and M. Nasre. Popular mixed matchings. *Theoretical Computer Science*, 412(24):2679–2690, 2011.
- G. Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3–27 July 1962)*, pages 73–79, 1965.
- G. Laffond, J.-F. Laslier, and M. Le Breton. More on the tournament equilibrium set. *Mathématiques et sciences humaines*, 31(123):37–44, 1993a.
- G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993b.
- 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. Rank-based choice correspondences. *Economics Letters*, 52(3):279–286, 1996.
- J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 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.
- M. Le Breton. On the uniqueness of equilibrium in symmetric two-player zero-sum games with integer payoffs. *Économie publique*, 17(2):187–195, 2005.
- R. Lensberg. Stability and the Nash solution. *Journal of Economic Theory*, 45(2):330–341, 1988.
- E. Maskin. Decision-making under ignorance with implications for social choice. *Theory and Decision*, 11(3):319–337, 1979.
- K. May. A set of independent, necessary and sufficient conditions for simple majority decisions. *Econometrica*, 20(4):680–684, 1952.
- D. C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.
- I. McLean and F. Hewitt. *Condorcet: Foundations of Social Choice and Political Theory*. Edward Elgar Publishing, 1994.

- R. H. Möhring. Algorithmic aspects of the substitution decomposition in optimization over relations, set system and boolean functions. *Annals of Operations Research*, 4(6):195–225, 1985.
- B. Monjardet. Statement of precedence and a comment on IIA terminology. *Games and Economic Behavior*, 62:736–738, 2008.
- H. Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.
- R. B. Myerson. Axiomatic derivation of scoring rules without the ordering assumption. *Social Choice and Welfare*, 12(1):59–74, 1995.
- J. F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- S. Nitzan and A. Rubinstein. A further characterization of Borda ranking method. *Public Choice*, 36(1):153–158, 1981.
- P. K. Pattanaik and B. Peleg. Distribution of power under stochastic social choice rules. *Econometrica*, 54(4):909–921, 1986.
- B. Peleg and S. Tijs. The consistency principle for games in strategic form. *International Journal of Game Theory*, 25(1):13–34, 1996.
- 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.
- R. L. Rivest and E. Shen. An optimal single-winner preferential voting system based on game theory. In *Proceedings of the 3rd International Workshop on Computational Social Choice (COMSOC)*, pages 399–410, 2010.
- D. G. Saari. Consistency of decision processes. *Annals of Operations Research*, 23(1):103–137, 1990a.
- D. G. Saari. The Borda dictionary. *Social Choice and Welfare*, 7(4):279–317, 1990b.
- D. G. Saari. *Basic Geometry of Voting*. Springer, 1995.
- M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- B. Saunders. Democracy, political equality, and majority rule. *Ethics*, 121(1):148–177, 2010.
- A. K. Sen. Choice functions and revealed preference. *Review of Economic Studies*, 38(3):307–317, 1971.
- A. K. Sen. Social choice theory: A re-examination. *Econometrica*, 45(1):53–89, 1977.
- A. K. Sen. Social choice theory. In K. J. Arrow and M. D. Intriligator, editors, *Handbook of Mathematical Economics*, volume 3, chapter 22, pages 1073–1181. Elsevier, 1986.

- J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.
- P. Stone. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press, 2011.
- W. Thomson. Consistent allocation rules. Mimeo, 2014.
- T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- J. von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100(1):295–320, 1928.
- H. P. Young. An axiomatization of Borda’s rule. *Journal of Economic Theory*, 9(1):43–52, 1974a.
- H. P. Young. A note on preference aggregation. *Econometrica*, 42(6):1129–1131, 1974b.
- H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.
- H. P. Young. Condorcet’s theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.
- H. P. Young. *Equity: in Theory and Practice*. Princeton University Press, 1994.
- H. P. Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995.
- H. P. Young and A. Levenglick. A consistent extension of Condorcet’s election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.
- T. M. Zavist and T. N. Tideman. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6(2):167–173, 1989.
- R. Zeckhauser. Majority rule with lotteries on alternatives. *Quarterly Journal of Economics*, 83(4):696–703, 1969.
- W. S. Zwicker. The voter’s paradox, spin, and the Borda count. *Mathematical Social Sciences*, 22(3):187–227, 1991.

APPENDIX

A. Preliminaries

As stated in Section 2, U is an infinite set of alternatives. For convenience we will assume that $\mathbb{N} \subseteq U$. For $n \in \mathbb{N}$, $[n]$ is defined as $[n] = \{1, \dots, n\}$. For two sets A and B , let $\Pi(A, B)$ denote the set of all bijections from A to B (where $\Pi(A, B) = \emptyset$ if $|A| \neq |B|$). Let $\Pi(A) = \Pi(A, A)$ be the set of all permutations of A . We will frequently work with profiles in which alternatives are renamed according to some bijection from one set of alternatives to another. For all $A, B \in \mathcal{F}(U)$, $\succsim \in \mathcal{L}(A)$, and $\pi \in \Pi(A, B)$, let $\pi(\succsim) = \{(\pi(x), \pi(y)) : (x, y) \in \succsim\} \in \mathcal{L}(B)$ and, for $R \in \mathcal{R}|_A$, let $\pi(R) \in \mathcal{R}|_B$ such that $R(\succsim) = (\pi(R))(\pi(\succsim))$. A well-known symmetry condition for PSCFs is *neutrality*, which requires that all alternatives are treated equally in the sense that renaming alternatives is appropriately reflected in the outcome. Formally, a PSCF is neutral if

$$\pi(f(R)) = f(\pi(R)) \text{ for all } A, B \in \mathcal{F}(U), R \in \mathcal{R}|_A, \text{ and } \pi \in \Pi(A, B). \quad (\text{neutrality})$$

We show that composition-consistency implies neutrality by replacing all alternatives with components of size 2.

Lemma 1. *Every composition-consistent PSCF satisfies neutrality.*

Proof. Let f be a composition-consistent PSCF, $A, B \in \mathcal{F}(U)$, $R \in \mathcal{R}|_A$, and $\pi \in \Pi(A, B)$. We have to show that $\pi(f(R)) = f(\pi(R))$. To this end, let $p^A \in f(R)$. First, choose $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ such that $b_i = \pi(a_i)$ for all $i \in [n]$. Since U is infinite, there is $C = \{c_1, \dots, c_n\} \in \mathcal{F}(U)$ such that $C \cap A = \emptyset$ and $C \cap B = \emptyset$. Now, let $R' \in \mathcal{R}|_{A \cup C}$ such that $R'|_A = R$ and $\{a_i, c_i\}$ is a component in R' for all $i \in [n]$. Thus, we have that $p^A \in f(R'|_A)$. We now apply composition-consistency to a_i and the components $\{a_i, c_i\}$ for all $i \in [n]$, which by definition implies that

$$f(R'|_A) \times_{a_1} f(R'|_{\{a_1, c_1\}}) \times_{a_2} f(R'|_{\{a_2, c_2\}}) \cdots \times_{a_n} f(R'|_{\{a_n, c_n\}}) = f(R'|_{A \cup C}).$$

Hence, for $p^{AC} \in f(R'|_{A \cup C})$ we have $p_{a_i}^{AC} + p_{c_i}^{AC} = p_{a_i}^A$ for all $i \in [n]$. Applying composition-consistency analogously to c_i and $\{a_i, c_i\}$ for all $i \in [n]$ yields $p_{c_i}^C = p_{a_i}^{AC} + p_{c_i}^{AC} = p_{a_i}^A$ for all $p^C \in f(R'|_C)$ and $i \in [n]$. Finally, let $R'' \in \mathcal{R}|_{B \cup C}$ such that $R''|_C = R'|_C$ and $\{b_i, c_i\}$ is a component in R'' for all $i \in [n]$. Hence, we have that $p^C \in f(R''|_C)$. As before, it follows from composition-consistency that $p^B \in f(R''|_B)$ where $p_{b_i}^B = p_{c_i}^C$ for all $i \in [n]$. Notice that $p^B = \pi(p^A)$ and $B = \pi(A)$. Since $R''|_B = \pi(R)$ by construction of R'' , we have $p^B \in f(\pi(R))$. Hence, $\pi(f(R)) \subseteq f(\pi(R))$. The fact that $f(\pi(R)) \subseteq \pi(f(R))$ follows from application of the above to $\pi(R)$ and π^{-1} . \square

The following notation is required for our proofs. For some set X , $\text{uni}(X)$ denotes the *uniform distribution* over X . In particular, for $A \in \mathcal{F}(U)$, $\text{uni}(A)$ is the uniform lottery over A , i.e., $\text{uni}(A) = 1/|A| \sum_{x \in A} x$. The *support* of a lottery p is the set of all alternatives

to which p assigns positive probability, i.e., $\text{supp}(p) = \{x \in A: p_x > 0\}$. The 1-norm of $x \in \mathbb{Q}^n$ is denoted by $\|x\|$, i.e., $\|x\| = \sum_{i=1}^n |x_i|$. For $X \subseteq \mathbb{Q}^n$, the *convex hull* $\text{conv}(X)$ is the set of all convex combinations of elements of X , i.e.,

$$\text{conv}(X) = \left\{ \lambda_1 a^1 + \dots + \lambda_k a^k : a^i \in X, \lambda_i \in \mathbb{Q}_{\geq 0}, \sum_{i=1}^k \lambda_i = 1 \right\}.$$

X is *convex* if $X = \text{conv}(X)$. The *affine hull* $\text{aff}(X)$ is the set of all affine combinations of elements of X , i.e.,

$$\text{aff}(X) = \left\{ \lambda_1 a^1 + \dots + \lambda_k a^k : a^i \in X, \lambda_i \in \mathbb{Q}, \sum_{i=1}^k \lambda_i = 1 \right\}.$$

X is an *affine subspace* if $X = \text{aff}(X)$. We say that $a^1, \dots, a^k \in \mathbb{Q}^n$ are *affinely independent* if, for all $\lambda_1, \dots, \lambda_k \in \mathbb{Q}$ with $\sum_{i=1}^k \lambda_i = 0$, $\sum_{i=1}^k \lambda_i a^i = 0$ implies $\lambda_i = 0$ for all $i \in [k]$. The *dimension* of an affine subspace X , $\dim(X)$, is $k - 1$, where k is the maximal number of affinely independent vectors in X . The dimension of a set X is the dimension of $\text{aff}(X)$. The *linear hull* $\text{lin}(X)$ is the set of all linear combinations of elements of X , i.e.,

$$\text{lin}(X) = \left\{ \lambda_1 a^1 + \dots + \lambda_k a^k : a^i \in X, \lambda_i \in \mathbb{Q} \right\}.$$

$B_\varepsilon(x) = \{y \in \mathbb{Q}^n : \|x - y\| < \varepsilon\}$ denotes the ε -ball around $x \in \mathbb{Q}^n$. The *interior* of $X \subseteq \mathbb{Q}^n$ in $Y \subseteq \mathbb{Q}^n$ is $\text{int}_Y(X) = \{x \in X : B_\varepsilon(x) \cap Y \subseteq X \text{ for some } \varepsilon > 0\}$. The *closure* of $X \subseteq \mathbb{Q}^n$ in $Y \subseteq \mathbb{Q}^n$, $\text{cl}_Y(X)$, is the set of all limit points of sequences in X which converge in Y , i.e., $\text{cl}_Y(X) = \{\lim_{k \rightarrow \infty} a^k : (a^k)_{k \in \mathbb{N}} \text{ converges in } Y \text{ and } a^k \in X \text{ for all } k \in \mathbb{N}\}$. X is *dense* in Y if $\text{cl}_Y(X) = Y$. Alternatively, X is dense at $y \in \mathbb{Q}^n$ if for every $\varepsilon > 0$ there is $x \in X$ such that $\|x - y\| < \varepsilon$. X is dense in Y if X is dense at y for every $y \in Y$.

B. Non-Probabilistic Social Choice

Theorem 1. *There is no SCF that satisfies population-consistency and composition-consistency.*

Proof. Assume for contradiction that f is an SCF that satisfies population-consistency and composition-consistency. Let $A = \{a, b, c\}$ and consider the profiles R^1, \dots, R^6 as depicted below.

1/3	1/3	1/3	1/2	1/2	1/2	1/2	1/2	1/2	1/2	1/2	1/2	1/2
a	b	c	a	c	a	b	b	c	a	b	b	a
b	c	a	b	b	c	c	a	a	b	c	a	c
c	a	b	c	a	b	a	c	b	c	a	c	b
R^1			R^2		R^3		R^4		R^5		R^6	

We claim that $\Delta(\{a, b\}) \subseteq f(R^i)$ for all $i \in \{1, \dots, 6\}$. It follows from neutrality that $f(R^1) = \Delta(A)$. Again, by neutrality, $f(R^2|_{\{a, b\}}) = \Delta(\{a, b\})$ and $f(R^2|_{\{a, c\}}) = \Delta(\{a, c\})$. Notice that $\{a, b\}$ is a component in R^2 . Hence, by composition-consistency,

$$f(R^2) = f(R^2|_{\{a, c\}}) \times_a f(R^2|_{\{a, b\}}) = \Delta(\{a, c\}) \times_a \Delta(\{a, b\}) = \Delta(A).$$

A similar argument yields $f(R^i) = \Delta(A)$ for $i = 3, 4$. Unanimity implies that $f(R^5|_{\{b, c\}}) = \{b\}$ and, by neutrality, we have $f(R^5|_{\{a, b\}}) = \Delta(\{a, b\})$. Furthermore, $\{b, c\}$ is a component in R^5 . Hence, by neutrality and composition-consistency,

$$f(R^5) = f(R^5|_{\{a, b\}}) \times_b f(R^5|_{\{b, c\}}) = \Delta(\{a, b\}) \times_b \{b\} = \Delta(\{a, b\}).$$

Similarly, $f(R^6) = \Delta(\{a, b\})$.

Every profile R^i is a vector in the five-dimensional unit simplex $\mathcal{R}|_A$ in \mathbb{Q}^6 . The corresponding vectors are depicted below.

$$\begin{pmatrix} R^1 \\ R^2 \\ R^3 \\ R^4 \\ R^5 \\ R^6 \end{pmatrix} = \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 & 0 & 0 \\ 0 & 0 & 1/2 & 0 & 0 & 1/2 \\ 1/2 & 1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/2 & 0 & 1/2 \end{pmatrix}$$

It can be checked that R^1, \dots, R^6 are affinely independent, i.e., $\dim(\{R^1, \dots, R^6\}) = 5$. It follows from population-consistency that $\Delta(\{a, b\}) \subseteq f(R)$ for every $R \in \text{conv}(\{R^1, \dots, R^6\})$. Hence, $\{R \in \mathcal{R}|_A : |f(R)| = 1\}$ is not dense in $\mathcal{R}|_A$ at $1/6 R^1 + \dots + 1/6 R^6$, which contradicts decisiveness of f . \square

C. Probabilistic Social Choice

In this section we prove that every PSCF that satisfies population-consistency and composition-consistency has to return maximal lotteries. The high-level structure of the proof is described after Theorem 2 in Section 5.

C.1. *ML* Satisfies Population-Consistency and Composition-Consistency

We first show that *ML* is a PSCF that satisfies population-consistency and composition-consistency. This statement is split into two lemmas.

Lemma 2. *ML is a PSCF.*

Proof. *ML* is continuous, since the correspondence that maps a matrix M to the set of vectors x such that $Mx \geq 0$ is (upper-hemi) continuous.

The fact that $f(R)$ is convex for every $R \in \mathcal{R}$ follows from convexity of the set of maximin strategies for all (symmetric) zero-sum games.

ML obviously satisfies unanimity by definition.

ML satisfies decisiveness. Let $A \in \mathcal{F}(U)$ and $R \in \mathcal{R}|_A$. It is easy to see that, for every $\varepsilon > 0$, we can find $R' \in B_\varepsilon(R) \cap \mathcal{R}|_A$ and $k \in \mathbb{N}$ such that $kR'(x, y)$ is an odd integer for all $x, y \in A$ with $x \neq y$. Laffond et al. (1997) have shown that every symmetric zero-sum game whose off-diagonal entries are odd integers admits a unique Nash equilibrium. Hence, $|f(R')| = 1$ and f is decisive. \square

Moreover, the set of symmetric zero-sum games with a unique maximin strategy inherits openness from the set of all zero-sum games with a unique maximin strategy (Bohnenblust et al., 1950, pp. 56–58). Hence, the set of profiles with a unique maximal lottery is open and dense in the set of all profiles and the set of profiles with multiple maximal lotteries is nowhere dense.

Lemma 3. *ML satisfies population-consistency and composition-consistency.*

Proof. To simplify notation, for every $v \in \mathbb{Q}^n$ and $X \subseteq [n]$, we denote by v_X the restriction of v to indices in X , i.e., $v_X = (v_i)_{i \in X}$.

ML satisfies population-consistency. Let $A \in \mathcal{F}(U)$, $R', R'' \in \mathcal{R}|_A$, and $p \in ML(R') \cap ML(R'')$. Then, by definition of ML , $p^T M_{R'} q \geq 0$ and $p^T M_{R''} q \geq 0$ for all $q \in \Delta(A)$. Hence, for all $\lambda \in [0, 1]$,

$$p^T (\lambda M_{R'} + (1 - \lambda) M_{R''}) q = \lambda \underbrace{p^T M_{R'} q}_{\geq 0} + (1 - \lambda) \underbrace{p^T M_{R''} q}_{\geq 0} \geq 0,$$

for all $q \in \Delta(A)$, which implies that $p \in ML(\lambda R' + (1 - \lambda) R'')$.

ML satisfies composition-consistency. Let $A', B \in \mathcal{F}(U)$ such that $A' \cap B = \{b\}$, $A = A' \cup B$, and $R \in \mathcal{R}|_A$ such that B is a component in R . To simplify notation, let $C = A \setminus B$ and $M = M_R$, $M_{A'} = M_{R|_{A'}}$, $M_B = M_{R|_B}$, and $M_C = M_{R|_C}$. Notice first that M and $M_{A'}$ take the following form for some $v \in \mathbb{Q}^{A \setminus B}$:

$$M = \left(\begin{array}{ccc|ccc} & & & | & & | \\ & M_C & & | & v & \dots & v & | \\ \dots & & & | & & & & | \\ \dots & (-v^T) & & | & & & & | \\ & & & | & & & & | \\ & \vdots & & | & & & & | \\ \dots & & & | & & & & | \\ \dots & (-v^T) & & | & & & & | \\ & & & | & & & & | \end{array} \right), \quad M_{A'} = \left(\begin{array}{ccc|c} & & & | \\ & M_C & & | \\ \dots & & & | \\ \dots & (-v^T) & & | \\ & & & | \\ & & & | \\ \dots & & & | \\ & & & | \\ & & & | \\ & & & 0 \end{array} \right).$$

Let $p \in ML(R|_{A'}) \times_b ML(R|_B)$. Then, there are $p^{A'} \in ML(R|_{A'})$ and $p^B \in ML(R|_B)$

such that $p = p^{A'} \times_b p^B$. Let $q \in \Delta(A)$. Then,

$$\begin{aligned} p^T M q &= p_C^T M_C q_C + \|p_B\| (-v)^T q_C + p_C^T v \|q_B\| + p_B^T M_B q_B \\ &= (p_C, \|p_B\|)^T M_{A'}(q_C, \|q_B\|)^T + p_B^T M_B q_B \\ &= \underbrace{(p^{A'})^T M_{A'}(q_C, \|q_B\|)^T}_{\geq 0} + \underbrace{\|p_B\| (p^B)^T M_B q_B}_{\geq 0} \geq 0, \end{aligned}$$

since $p^{A'} \in ML(R|_{A'})$ and $p^B \in ML(R|_B)$, respectively. Hence $p \in ML(R)$.

For the other direction, let $p \in ML(R)$. We have to show that there are $p^{A'} \in ML(R|_{A'})$ and $p^B \in ML(R|_B)$ such that $p = p^{A'} \times_b p^B$.

First, if $\|p_B\| = 0$ let $p^{A'} = p_{A'}$ and $p^B \in ML(R|_B)$ be arbitrary. Let $q \in \Delta(A')$. Then,

$$(p^{A'})^T M_{A'} q = p_C^T M_C q_C + p_C^T v q_b = p^T M(q, 0)^T \geq 0,$$

since $p \in ML(R)$. Hence, $p^{A'} \in ML(R|_{A'})$.

Otherwise, let $p^{A'} = (p_C, \|p_B\|)$ and $p^B = p_B / \|p_B\|$. Let $q \in \Delta(A')$. Then,

$$\begin{aligned} (p^{A'})^T M_{A'} q &= p_C^T M_C q_C + \|p_B\| (-v)^T q_C + p_C^T v q_b \\ &= p_C^T M_C q_C + \|p_B\| (-v)^T q_C + p_C^T v q_b + \frac{q_b}{\|p_B\|} \underbrace{p_B^T M_B p_B}_{=0} \\ &= p^T M(q_C, \frac{q_b}{\|p_B\|} p_B)^T \geq 0. \end{aligned}$$

Hence, $p^{A'} \in ML(R|_{A'})$. Let $q \in \Delta(B)$. Then,

$$\begin{aligned} \|p_B\|^2 (p^B)^T M_B q &= \|p_B\| p_B^T M_B q \\ &= \|p_B\| p_B^T M_B q + \underbrace{p_C^T M_C p_C}_{=0} + \underbrace{\|p_B\| (-v)^T p_C + \|p_B\| p_C^T v}_{=0} \\ &= (p_C, p_B)^T M(p_C, \|p_B\| q) = p^T M(p_C, \|p_B\| q) \geq 0. \end{aligned}$$

Hence, $p^B \in ML(R|_B)$. □

C.2. Binary Choice

The basis of our characterization of ML is the special case for agendas of size 2. The following lemma states that, on two alternatives, whenever a composition-consistent PSCF returns a non-degenerate lottery, it has to return all lotteries. Interestingly, the proof uses composition-consistency on three-element agendas, even though the statement itself only concerns agendas of size 2. In order to simplify notation, define

$$p^\lambda = \lambda a + (1 - \lambda) b.$$

Lemma 4. *Let $A = \{a, b\}$ and f be a PSCF that satisfies composition-consistency. Then, for all $R \in \mathcal{R}|_A$ and $\lambda \in (0, 1)$, $p^\lambda \in f(R)$ implies $f(R) = \Delta(A)$.*

Proof. Let $R \in \mathcal{R}|_A$ and assume $p^\lambda \in f(R)$ for some $\lambda \in (0, 1)$. Define $R' \in \mathcal{R}|_{\{a, b, c\}}$ as depicted below.

$R(a, b)$	$R(b, a)$
a	c
b	b
c	a
R'	

Notice that $R'|_A = R$ and thus, $p^\lambda \in f(R'|_A)$. Neutrality implies that $\lambda a + (1 - \lambda)c \in f(R'|_{\{a, c\}})$. Since A is a component in R' , we have $\lambda p^\lambda + (1 - \lambda)c \in f(R'|_{\{a, c\}}) \times_a f(R'|_A) = f(R')$. Since $\{b, c\}$ is also a component in R' , composition-consistency implies that $\lambda p^\lambda + (1 - \lambda)c \in f(R') = f(R'|_A) \times_b f(R'|_{\{b, c\}})$. Observe that $\lambda p_a^\lambda = p_a^{\lambda^2}$ and hence $p^{\lambda^2} \in f(R'|_A) = f(R)$.

Applying this argument repeatedly, we get $p^{\lambda^{2^k}} \in f(R)$ for all $k \in \mathbb{N}$. Since $\lambda^{2^k} \rightarrow 0$ for $k \rightarrow \infty$ and f is continuous, we get $p^0 = b \in f(R)$. Similarly, it follows that $p^1 = a \in f(R)$. The fact that f is convex-valued implies that $f(R) = \Delta(A)$. \square

The characterization of ML for agendas of size 2 proceeds along the following lines. By unanimity, neutrality, and Lemma 4, we know which lotteries have to be returned by every composition-consistent PSCF for three particular profiles. Then population-consistency implies that every such PSCF has to return all maximal lotteries. Last, we again use population-consistency to show that the function is not decisive if it additionally returns lotteries that are not maximal.

Lemma 5. *Let f be a PSCF that satisfies population-consistency and composition-consistency and $A = \{a, b\}$. Then $f(R) = ML(R)$ for every $R \in \mathcal{R}|_A$.*

Proof. First, note that $R \in \mathcal{R}|_A$ is fully determined by $R(a, b)$. Let $R \in \mathcal{R}|_A$ be the profile such that $R(a, b) = 1/2$. Since $f(R) \neq \emptyset$, there is $\lambda \in [0, 1]$ such that $p^\lambda \in f(R)$. Neutrality implies that $p^{1-\lambda} \in f(R)$ and hence, by convexity of $f(R)$, $p^{1/2} = 1/2(p^\lambda + p^{1-\lambda}) \in f(R)$. It follows from Lemma 4 that $f(R) = \Delta(A)$.

Now, let $R \in \mathcal{R}|_A$ be the profile such that $R(a, b) = 1$. Unanimity implies that $a \in f(R)$. By population-consistency and the first part of the proof, we get $a \in f(R')$ for all $R' \in \mathcal{R}|_A$ with $R'(a, b) \in [1/2, 1]$. Similarly, $b \in f(R')$ for all $R' \in \mathcal{R}|_A$ with $R'(a, b) \in [0, 1/2]$. This already shows that $ML(R) \subseteq f(R)$ for every $R \in \mathcal{R}|_A$.

Finally, let $R \in \mathcal{R}|_A$ be a profile such that $R(a, b) = r > 1/2$. If $f(R) \neq \{a\}$, there is $\lambda \in [0, 1)$ such that $p^\lambda \in f(R)$. We have shown before that $f(R') = \Delta(A)$ if $R'(a, b) = 1/2$. Hence, it follows from population-consistency that $p^\lambda \in f(R')$ for every $R' \in \mathcal{R}|_A$ with

$R'(a, b) \in [1/2, r]$. But then $\{R \in \mathcal{R}|_A : R(a, b) \in [1/2, r]\} \subseteq \{R \in \mathcal{R}|_A : |f(R)| > 1\}$ and hence, $\{R \in \mathcal{R}|_A : |f(R)| = 1\}$ is not dense in $\mathcal{R}|_A$. This contradicts decisiveness of f . An analogous argument shows that $f(R) = \{b\}$ whenever $R(a, b) < 1/2$.

In summary, we have that $f(R) = \{a\}$ if $R(a, b) \in (1/2, 1]$, $f(R) = \{b\}$ if $R(a, b) \in [0, 1/2)$, and $f(R) = \Delta(A)$ if $R(a, b) = 1/2$. Thus, $f = ML$ (as depicted in Figure 1(a)). \square

C.3. $f \subseteq ML$

The first lemma in this section shows that every PSCF that satisfies population-consistency and composition-consistency is Condorcet-consistent for profiles that are close to the uniform profile $\text{uni}(\mathcal{L}(A))$, i.e., the profile in which every preference relation is assigned the same fraction of voters. We prove this statement by induction on the number of alternatives. Every profile close to the uniform profile that admits a Condorcet winner can be written as a convex combination of profiles that have a component and admit the same Condorcet winner. For these profiles we know from the induction hypothesis that the Condorcet winner has to be chosen.

Lemma 6. *Let f be a PSCF that satisfies population-consistency and composition-consistency and $A \in \mathcal{F}(U)$. Then, f satisfies Condorcet-consistency in a neighborhood of the uniform profile $\text{uni}(\mathcal{L}(A))$.*

Proof. Let f be a PSCF that satisfies population-consistency and composition-consistency and $A \in \mathcal{F}(U)$ with $|A| = n$. Let, furthermore, $R \in \mathcal{R}|_A$ be such that $a \in A$ is a Condorcet winner in R and $\|R - \text{uni}(\mathcal{L}(A))\| \leq \varepsilon_n = (4^n \prod_{k=1}^n k!)^{-1}$. We show that $a \in f(R)$ by induction over n . An example for $n = 3$ illustrating the idea is given after the proof. For $n = 2$, the claim follows directly from Lemma 5.

For $n > 2$, fix $b \in A \setminus \{a\}$. First, we introduce some notation. For $\succsim \in \mathcal{L}(A)$, we denote by \succsim^{-1} the preference relation that reverses all pairwise comparisons, i.e., $x \succsim^{-1} y$ iff $y \succsim x$ for all $x, y \in A$. By $\succsim^{b \rightarrow a}$ we denote the preference relation that is identical to \succsim except that b is moved upwards or downwards (depending on whether $a \succsim b$ or $b \succsim a$) until it is next to a in the preference relation. Formally, let

$$X_{\succsim} = \begin{cases} \{x \in A : a \succ x \succ b\} & \text{if } a \succ b, \text{ and} \\ \{x \in A : b \succ x \succ a\} & \text{if } b \succ a, \end{cases}$$

and

$$\succsim^{b \rightarrow a} = \begin{cases} \succsim \setminus (X_{\succsim} \times \{b\}) \cup (\{b\} \times X_{\succsim}) & \text{if } a \succ b, \text{ and} \\ \succsim \setminus (\{b\} \times X_{\succsim}) \cup (X_{\succsim} \times \{b\}) & \text{if } b \succ a. \end{cases}$$

Notice that for every $\succsim' \in \mathcal{L}(A)$, there are at most $n - 1$ distinct preference relations \succsim such that $\succsim' = \succsim^{b \rightarrow a}$. Furthermore, we say that $\{a, b\}$ is a *component* in \succsim if $X_{\succsim} = \emptyset$.

We first show that composition-consistency implies Condorcet-consistency for a particular type of profiles. For $\succsim \in \mathcal{L}(A)$, let $S \in \mathcal{R}|_A$ such that $S(\succsim) + S(\succsim)^{b \rightarrow a} = S(\succsim^{-1}) = 1/2$.

We have that $S(a, x) = 1/2$ for all $x \in A \setminus \{a\}$ and hence, a is a Condorcet winner in S . We prove that $a \in f(S)$ by induction over n . For $n = 2$, this follows from Lemma 5. For $n > 2$, let $x \in A \setminus \{b\}$ such that $x \succ y$ for all $y \in A$ or $y \succ x$ for all $y \in A$. This is always possible since $n > 2$. Notice that $A \setminus \{x\}$ is a component in S and $S(x, y) = 1/2$ for all $y \in A \setminus \{x\}$. If $x = a$, it follows from composition-consistency and Lemma 5 that $a \in f(S)$. If $x \neq a$, it follows from the induction hypothesis that $a \in f(S|_{A \setminus \{x\}})$. Lemma 5 implies that $a \in f(S|_{\{a, x\}})$ as $S(a, x) = 1/2$. Then, it follows from composition-consistency that $a \in f(S|_{\{a, x\}}) \times_a f(S|_{A \setminus \{x\}}) = f(S)$.

Now, for every $\succ \in \mathcal{L}(A)$ such that $\{a, b\}$ is not a component in \succ and $0 < R(\succ) \leq R(\succ^{-1})$, let $S^\succ \in \mathcal{R}|_A$ such that

$$S^\succ(\succ) + S^\succ(\succ^{b \rightarrow a}) = S^\succ(\succ^{-1}) = 1/2 \quad \text{and} \quad S^\succ(\succ)/S^\succ(\succ^{-1}) = R(\succ)/R(\succ^{-1}).$$

From what we have shown before, it follows that $a \in f(S^\succ)$ for all $\succ \in \mathcal{L}(A)$.

The rest of the proof proceeds as follows. We show that R can be written as a convex combination of profiles of the type S^\succ and a profile R' in which $\{a, b\}$ is a component and a is a Condorcet winner. Since R is close to the uniform profile, $R(\succ')$ is almost identical for all preference relations \succ' . Hence $S^\succ(\succ')$ is close to 0 for all preference relations \succ' in which $\{a, b\}$ is a component. As a consequence, $R'(\succ')$ is almost identical for all preference relations \succ' in which $\{a, b\}$ is a component and $R'|_{A \setminus \{b\}}$ is close to the uniform profile for $n-1$ alternatives, i.e., $\text{uni}(\mathcal{L}(A \setminus \{b\}))$. By the induction hypothesis, $a \in f(R'|_{A \setminus \{b\}})$. Since $\{a, b\}$ is a component in R' , it follows from composition-consistency that $a \in f(R')$.

We define

$$S = 2 \sum_{\succ} R(\succ^{-1}) S^\succ,$$

where the sum is taken over all \succ such that $\{a, b\}$ is not a component in \succ and $0 < R(\succ) \leq R(\succ^{-1})$ (in case $R(\succ) = R(\succ^{-1})$ we pick one of \succ and \succ^{-1} arbitrarily). Now, let $R' \in \mathcal{R}|_A$ such that

$$R = (1 - \|S\|)R' + S.$$

Note that, by definition of S , $R'(\succ) = 0$ for all $\succ \in \mathcal{L}(A)$ such that $\{a, b\}$ is not a component in \succ . Hence, $\{a, b\}$ is a component in R' . By the choice of R , we have that

$$\|S\| = \sum_{\succ \in \mathcal{L}(A)} S(\succ) \leq \frac{n! - 2(n-1)!}{n!} + \varepsilon_n = 1 - \frac{2}{n} + \varepsilon_n.$$

Using this fact, a simple calculation shows that

$$R'(\succ) \leq \frac{R(\succ) - S(\succ)}{\frac{2}{n} - \varepsilon_n} \leq \frac{\frac{1}{n!} + \varepsilon_n}{\frac{2}{n} - \varepsilon_n} \leq \frac{1}{2(n-1)!} + \frac{\varepsilon_{n-1}}{4(n-1)!}.$$

Since, for every preference relation \succ where $\{a, b\}$ is a component, there is exactly one other preference relation identical to \succ except that a and b are swapped, we have that

$$R'|_{A \setminus \{b\}}(\succ) \leq \frac{1}{(n-1)!} + \frac{\varepsilon_{n-1}}{2(n-1)!},$$

for every $\succsim \in \mathcal{L}(A \setminus \{b\})$. By the above calculation, we have that

$$\|R'|_{A \setminus \{b\}} - \text{uni}(\mathcal{L}(A \setminus \{b\}))\| \leq \varepsilon_{n-1}.$$

Since $S^\succsim(a, x) = 1/2$ for all $x \in A \setminus \{a\}$ and $\succsim \in \mathcal{L}(A)$, we have that $R'(a, x) \geq 1/2$ for all $x \in A \setminus \{a\}$. Thus, a is a Condorcet winner in $R'|_{A \setminus \{b\}}$. From the induction hypothesis it follows that $a \in f(R'|_{A \setminus \{b\}})$. Using the fact that $R'(a, b) \geq 1/2$, Lemma 5 implies that $a \in f(R'|_{\{a, b\}})$. Finally, composition-consistency entails $a \in f(R'|_{A \setminus \{b\}}) \times_a f(R'|_{\{a, b\}}) = f(R')$.

In summary, $a \in f(S^\succsim)$ for all $\succsim \in \mathcal{L}(A)$ and $a \in f(R')$. Since R is a convex combination of profiles of the type S^\succsim and R' , it follows from population-consistency that $a \in f(R)$. \square

We now give an example for $A = \{a, b, c\}$ which illustrates the proof of Lemma 6. Consider the following profile R , where $0 \leq \varepsilon \leq \varepsilon_3$.

$(1+2\varepsilon)/6$	$1/6$	$1/6$	$(1-\varepsilon)/6$	$(1-\varepsilon)/6$	$1/6$
a	a	b	b	c	c
b	c	a	c	a	b
c	b	c	a	b	a

R

Then, we have that $\|R - \text{uni}(\mathcal{L}(A))\| \leq \varepsilon_3$. Now consider \succsim with $b \succ c \succ a$, which yields S^\succsim as depicted below.

$1/2$	$(1-\varepsilon)/2$	$\varepsilon/2$
a	b	c
c	c	b
b	a	a

S^\succsim

Here, $y \succ a$ for all $y \in A$. Hence, it follows from what we have shown before that $a \in f(S^\succsim)$. No other profiles of this type need to be considered, as \succ and \succ^{-1} are the only preference relations in which $\{a, b\}$ is not a component. Thus $S = 1/3 S^\succsim$.

Then, we have R' , $R'|_{\{a, c\}}$, and $R'|_{\{a, b\}}$ as follows.

$(1+2\varepsilon)/4$	$1/4$	$(1-\varepsilon)/4$	$(1-\varepsilon)/4$	$(1+\varepsilon)/2$	$(1-\varepsilon)/2$	$(2+\varepsilon)/4$	$(2-\varepsilon)/4$
a	b	c	c	a	c	a	b
b	a	a	b	c	a	b	a
c	c	b	a				

R'

$R'|_{\{a, c\}}$

$R'|_{\{a, b\}}$

It follows from Lemma 5 that $a \in f(R'|_{\{a,c\}})$ and $a \in f(R'|_{\{a,b\}})$. Then, composition-consistency implies that $a \in f(R') = f(R'|_{\{a,c\}}) \times_a f(R'|_{\{a,b\}})$.

In summary, we have that

$$R = 2/3 R' + 1/3 S^{\succ},$$

$a \in f(R')$, and $a \in f(S^{\succ})$. Thus, population-consistency implies that $a \in f(R)$.

Lemma 7. *Every PSCF that satisfies population-consistency and composition-consistency returns the uniform lottery over all Condorcet winners for all profiles in a neighborhood of the uniform profile $\text{uni}(\mathcal{L}(A))$.*

Proof. Let f be a PSCF that satisfies population-consistency and composition-consistency and $A \in \mathcal{F}(U)$ with $|A| = n$. Moreover, let $R \in \mathcal{R}|_A$ such that $\|R - \text{uni}(\mathcal{L}(A))\| \leq \varepsilon_n$ and $A' \subseteq A$ be the set of Condorcet winners in R . We actually prove a stronger statement, namely that $\Delta(A') \subseteq f(R)$. Every alternative in A' is a Condorcet winner in R . Thus, it follows from Lemma 6 that $x \in f(R)$ for every $x \in A'$. Since $f(R)$ is convex, $\Delta(A') \subseteq f(R)$ follows. \square

For the remainder of the proof, we need to define two classes of profiles that are based on regularity conditions imposed on the corresponding majority margins. Let $A \in \mathcal{F}(U)$ and $A' \subseteq A$. A profile $R \in \mathcal{R}|_A$ is

$$\text{regular on } A' \text{ if } \sum_{y \in A'} M_R(x, y) = 0 \text{ for all } x \in A', \text{ and}$$

$$\text{strongly regular on } A' \text{ if } M_R(x, y) = 0 \text{ for all } x, y \in A'.$$

By $\mathcal{R}|_A^{A'}$ and $\mathcal{S}|_A^{A'}$ we denote the set of all profiles in $\mathcal{R}|_A$ that are regular or strongly regular on A' , respectively.

In the following five lemmas we show that, for every $A' \subseteq A$, every profile on A can be affinely decomposed into profiles of three different types: profiles that are strongly regular on A' , certain regular profiles, and profiles that admit a strict Condorcet winner in A' .²¹ Lemmas 8, 9, and 10 do not make any reference to population-consistency, composition-consistency, or maximal lotteries and may be of independent interest.

First, we determine the dimension of the space of all profiles that are strongly regular on A' .

Lemma 8. *Let $A' \subseteq A \in \mathcal{F}(U)$. Then, $\dim(\mathcal{S}|_A^{A'}) = |A|! - \binom{|A'|}{2} - 1$.*

Proof. We will characterize $\mathcal{S}|_A^{A'}$ using a set of linear constraints. By definition, $\mathcal{S}|_A^{A'} = \{R \in \mathcal{R}|_A : M_R(x, y) = 0 \text{ for all } x, y \in A'\}$. Recall that $M_R(x, y) = \sum_{\succ: x \succ y} R(\succ) - \sum_{\succ: y \succ x} R(\succ)$. Since $M_R(x, x) = 0$ for all $R \in \mathcal{R}|_A$ and $x \in A$, $\mathcal{S}|_A^{A'}$ can be characterized by $\binom{|A'|}{2}$ homogeneous linear constraints in the $(|A|! - 1)$ -dimensional space $\mathcal{R}|_A$, which

²¹Similar decompositions of majority margin matrices have been explored by Zwicker (1991) and Saari (1995).

implies that $\dim(\mathcal{S}_{|A|}^{A'}) \geq |A|! - \binom{|A|}{2} - 1$. Equality holds but is not required for the following arguments. We therefore omit the proof. \square

Second, we determine the dimension of the space of all skew-symmetric $n \times n$ matrices that correspond to regular profiles and vanish outside their upper left $n' \times n'$ sub-matrix, i.e.,

$$\mathcal{M}_{n'} = \left\{ M \in \mathbb{Q}^{n \times n} : M = -M^T, \sum_{j=1}^n M(i, j) = 0 \text{ if } i \in [n], \text{ and } M(i, j) = 0 \text{ if } \{i, j\} \not\subseteq [n'] \right\}.$$

In Lemma 10, we then proceed to show that every matrix of this type can be decomposed into matrices induced by a subset of regular profiles for which we know that every PSCF has to return the uniform lottery over the first n' alternatives (possibly among other lotteries).

Lemma 9. $\dim(\mathcal{M}_{n'}) = \binom{n'}{2} - (n' - 1)$.

Proof. First note that the space of all $n \times n$ matrices has dimension n^2 . We show that $\mathcal{M}_{n'}$ can be characterized by a set of $(n^2 - (n')^2) + (\binom{n'}{2} + n') + (n' - 1)$ homogeneous linear constraints. Let $M \in \mathbb{Q}^{n \times n}$ and observe that $(n^2 - (n')^2)$ constraints are needed to ensure that M vanishes outside of $[n'] \times [n']$, $\binom{n'}{2} + n'$ constraints are needed to ensure skew-symmetry of M on $[n'] \times [n']$, i.e., $M(i, j) = -M(j, i)$ for all $i, j \in [n']$, $j \geq i$, and $(n' - 1)$ constraints are needed to ensure that the first n' rows (and hence also the columns) of M' sum up to 0, i.e., $\sum_{j=1}^n M(i, j) = 0$ for all $i \in [n' - 1]$. It follows from skew-symmetry and the latter $n' - 1$ constraints that the n' th row of M sums up to 0, since

$$\sum_{j=1}^n M(n', j) = \sum_{i, j=1}^n M(i, j) - \sum_{i=1}^{n'-1} \sum_{j=1}^n M(i, j) = 0.$$

The last $n - n'$ rows of M trivially sum up to 0. Hence, $\dim(\mathcal{M}_{n'}) \geq (n')^2 - (\binom{n'}{2} + n') - (n' - 1) = \binom{n'}{2} - (n' - 1)$. Equality holds but is not required for the following arguments. We therefore omit the proof. \square

Let $\Pi_B^\circ([n])$ be the set of all permutations that are cyclic on B and coincide with the identity permutation outside of B .²² We denote by $\mathcal{M}_{n'}^\circ$ the space of all matrices in $\mathcal{M}_{n'}$ induced by a permutation in $\Pi_B^\circ([n])$ for some $B \subseteq [n']$, i.e.,

$$\mathcal{M}_{n'}^\circ = \left\{ M \in \mathcal{M}_{n'} : M(i, j) = \begin{cases} 1 & \text{if } j = \pi(i), i \in B, \\ -1 & \text{if } i = \pi(j), j \in B, \\ 0 & \text{otherwise,} \end{cases} \text{ for some } \pi \in \Pi_B^\circ([n]), B \subseteq [n'] \right\},$$

with the convention that $\mathcal{M}_2^\circ = \{0\}$. We now show that the linear hull of $\mathcal{M}_{n'}^\circ$ is $\mathcal{M}_{n'}$.

²²A permutation $\pi \in \Pi([n])$ is cyclic on B if $\pi^{|B|}$ is the smallest positive power of π that is the identity function on B .

Lemma 10. $\text{lin}(\mathcal{M}_{n'}^\circ) = \mathcal{M}_{n'}$.

Proof. The idea underlying the proof is as follows: every matrix $M \in \mathcal{M}_{n'}$ corresponds to a weighted directed graph with vertex set $[n]$ where the weight of the edge from i to j is $M(i, j)$. If $M \neq 0$, there exists a cycle along edges with positive weight of length at least 3 in the subgraph induced by $[n']$. We obtain a matrix M' with smaller norm than M by subtracting the matrix in $\mathcal{M}_{n'}^\circ$ from M that corresponds to the cycle identified before.

Let $M \in \mathcal{M}_{n'}$ and $\kappa \in \mathbb{Q} \setminus \{0\}$ such that $\kappa M \in \mathbb{N}^{n \times n}$. We show, by induction over $\kappa \|M\|$, that $M = \sum_{i=1}^{\ell} \lambda_i M^i$ for some $\lambda_i \in \mathbb{Q}$ and $M^i \in \mathcal{M}_{n'}^\circ$ for all $i \in [\ell] \setminus \{1\}$ for some $\ell \in \mathbb{N}$. If $\kappa \|M\| = 0$ then $M = 0$. Hence, the induction hypothesis is trivial.

If $\kappa \|M\| \neq 0$, i.e., $M \neq 0$, we can find $B \subseteq [n']$ with $|B| \geq 3$ and $\pi \in \Pi_B^\circ([n])$ such that $M(i, j) > 0$ if $\pi(i) = j$ and $i \in B$. Note that π defines a cycle of length at least 3 in the graph that corresponds to M . We define $M^1 \in \mathcal{M}_{n'}^\circ$ by letting

$$M^1(i, j) = \begin{cases} 1 & \text{if } \pi(i) = j \text{ and } i \in B, \\ -1 & \text{if } \pi(j) = i \text{ and } j \in B, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\lambda = \min\{M(i, j) : i, j \in [n] \text{ and } M^1(i, j) > 0\}$ and $M' = M - \lambda M^1$. By construction, we have that $M'(i, j) = M(i, j) - \lambda$ if $\pi(i) = j$ and $i \in B$, $M'(i, j) = M(i, j) + \lambda$ if $\pi(j) = i$ and $j \in B$, and $M'(i, j) = M(i, j)$ otherwise. Note that $M(i, j) \geq \lambda$ if $\pi(i) = j$ and $i \in B$ and $M(i, j) \leq -\lambda$ if $\pi(j) = i$ and $j \in B$ by definition of λ . Recall that $\kappa M \in \mathbb{N}^{n \times n}$ and, in particular, $\kappa \lambda \in \mathbb{N}$. Hence, $\kappa M' \in \mathbb{N}^{n \times n}$. Moreover, $\kappa \|M'\| = \kappa \|M\| - 2\kappa \lambda |B| \leq \kappa \|M\| - 1$. From the induction hypothesis we know that $M' = \sum_{i=2}^{\ell} \lambda_i M^i$ with $\lambda_i \in \mathbb{Q}$ and $M^i \in \mathcal{M}_{n'}^\circ$ for all $i \in [\ell]$ for some $\ell \in \mathbb{N}$. By construction of M' , we have that $M = \sum_{i=1}^{\ell} \lambda_i M^i$ with $\lambda_1 = \lambda$. \square

Lemma 11 leverages Lemmas 7, 8, 9, and 10 to show two statements. First, it identifies the dimension of the space of all profiles that are regular on $A' \subseteq A$. Second, it proves that there is a full-dimensional subset of the space of all profiles that are regular on A' for which every PSCF that satisfies population-consistency and composition-consistency returns the uniform lottery over A' .

Lemma 11. *Let f be a PSCF that satisfies population-consistency and composition-consistency and $A' \subseteq A \in \mathcal{F}(U)$. Then, there is $\mathcal{X} \subseteq \mathcal{R}_A^{A'}$ of dimension $|A| - |A'|$ such that $\text{uni}(A') \in f(R)$ for every $R \in \mathcal{X}$.*

Proof. To simplify notation, we assume without loss of generality that $A = [n]$ and $A' = [n']$. For $M \in \mathbb{Q}^{n \times n}$ and $\pi \in \Pi(A)$, let $\pi(M)$ be the matrix that results from M by permuting the rows and columns of M according to π , i.e., $(\pi(M))(i, j) = M(\pi(i), \pi(j))$.

From Lemma 8 we know that we can find a set $\mathcal{S} = \{S^1, \dots, S^{n! - \binom{n}{2}}\} \subseteq \mathcal{R}_{[n]}^{[n']}$ of affinely independent profiles. Since \mathcal{S} can be chosen such that every $S \in \mathcal{S}$ is close to $\text{uni}(\mathcal{L}([n]))$, it follows from Lemma 7 that $\text{uni}([n']) \in f(S)$ for all $S \in \mathcal{S}$. Therefore, it suffices to find

a set of profiles $\mathcal{T} = \{R^1, \dots, R^{\binom{n'}{2} - (n'-1)}\} \subseteq \mathcal{R}^{[n']}$ such that $\text{uni}([n']) \in f(R)$ for every $R \in \mathcal{T}$ and $\mathcal{S} \cup \mathcal{T}$ is a set of affinely independent profiles. If $n' = 2$, we can choose $\mathcal{T} = \emptyset$. For $n' \geq 3$ we construct a suitable set of profiles as follows.

For every $B \subseteq [n']$ with $|B| = k \geq 3$ and $\pi \in \Pi_B^\circ([n])$, let $[n] \setminus B = \{a_1, \dots, a_{n-k}\}$ and R_B^π be defined as follows: $R_B^\pi(\succ) = 1/(2k)$ if

$$\begin{aligned} \pi^0(i) \succ \pi^1(i) \succ \pi^2(i) \succ \dots \succ \pi^{k-1}(i) \succ a_1 \succ \dots \succ a_{n-k} \quad \text{or} \\ a_{n-k} \succ \dots \succ a_1 \succ \pi^{k-1}(i) \succ \dots \succ \pi^2(i) \succ \pi^0(i) \succ \pi^1(i), \end{aligned}$$

for some $i \in B$. Note that R_B^π is regular on $[n']$, since

$$R_B^\pi(i, j) = \begin{cases} \lambda & \text{if } \pi(i) = j \text{ and } i \in B, \\ -\lambda & \text{if } \pi(j) = i \text{ and } j \in B, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

where $\lambda = 1/k > 0$. Hence, for every $M \in \mathcal{M}_{n'}^\circ$, there are $B \subseteq [n']$ and $\pi \in \Pi_B^\circ([n])$ such that $\lambda M = M_{R_B^\pi}$. Notice that B and $[n] \setminus B$ are components in R_B^π . For $j \in B$, we have by construction that $R_B^\pi(j, a_1) = 0$. Hence, it follows from Lemma 5 that $j \in f(R_B^\pi|_{\{j, a_1\}})$ and $a_1 \in f(R_B^\pi|_{\{j, a_1\}})$. Moreover, neutrality, convexity, and composition-consistency imply that $\text{uni}(B) \in f(R_B^\pi)$ by the symmetry of R_B^π with respect to B . Now let $a_i \in \{a_1, \dots, a_{n-k}\}$. Observe that $\{a_1, \dots, a_{i-1}\}$ is a component in R_B^π and $R_B^\pi(a_1, a_i) = 0$. Thus, composition-consistency and Lemma 5 imply that

$$a_i \in f(R_B^\pi|_{\{a_1, a_i\}}) \times_{a_1} f(R_B^\pi|_{\{a_1, \dots, a_{i-1}\}}) = f(R_B^\pi|_{\{1, \dots, i\}}).$$

Furthermore, $\{a_{i+1}, \dots, a_{n-k}\}$ is a component in R_B^π and $R_B^\pi(a_i, a_{n-k}) = 0$. As before, we get

$$a_i \in f(R_B^\pi|_{\{a_i, a_{n-k}\}}) \times_{a_{n-k}} f(R_B^\pi|_{\{a_{i+1}, \dots, a_{n-k}\}}) = f(R_B^\pi|_{\{i, \dots, n-k\}}).$$

Also $\{a_i, \dots, a_{n-k}\}$ is a component in R_B^π and thus,

$$a_i \in f(R_B^\pi|_{\{a_1, \dots, a_i\}}) \times_{a_i} f(R_B^\pi|_{\{a_i, \dots, a_{n-k}\}}) = f(R_B^\pi|_{[n] \setminus B}).$$

As B is a component in S and $R_B^\pi(j, a_1) = 0$, we get

$$a_i \in f(R_B^\pi|_{\{j, a_1, \dots, a_i\}}) \times_j f(R_B^\pi|_B) = f(R_B^\pi).$$

Then, it follows from convexity of $f(R_B^\pi)$ that

$$\text{uni}([n']) = \frac{k}{n'} \text{uni}(B) + \frac{1}{n'} \sum_{a_i \in [n'] \setminus B} a_i \in f(R_B^\pi),$$

since $\text{uni}(B) \in f(R_B^\pi)$ and $a_i \in f(R_B^\pi)$ for every $i \in [n-k]$.

We know from Lemma 10 that $\text{lin}(\mathcal{M}_{n'}^{\circ}) = \mathcal{M}_{n'}$ and, by Lemma 9, $\dim(\mathcal{M}_{n'}^{\circ}) \geq \binom{n'}{2} - (n' - 1)$. Thus, we can find a basis $\{M^1, \dots, M^{\binom{n'}{2} - (n' - 1)}\}$ of $\mathcal{M}_{n'}^{\circ}$ and a set of corresponding profiles

$$\mathcal{T} = \{R^1, \dots, R^{\binom{n'}{2} - (n' - 1)}\} \subseteq \{R_B^{\pi} : B \subseteq [n'] \text{ and } \pi \in \Pi_B^{\circ}\}.$$

We claim that $\mathcal{S} \cup \mathcal{T}$ is a set of affinely independent profiles. Let $S^1, \dots, S^l \in \mathcal{S}$ and $R^1, \dots, R^m \in \mathcal{T}$ be pairwise disjoint. Assume that $\sum_i \lambda_i S^i + \sum_j \mu_j R^j = 0$ for some $\lambda^i, \mu^j \in \mathbb{Q}$ such that $\sum_i \lambda_i + \sum_j \mu_j = 0$. This implies that $\sum_j \mu_j M^j = 0$, which in turn implies $\mu_j = 0$ for all $j \in [m]$, since the M^j 's are linearly independent. Hence, $\sum_i \lambda_i S^i = 0$ and $\sum_i \lambda_i = 0$, which implies that $\lambda_i = 0$ for all $i \in [l]$, since $S^1, \dots, S^{n! - \binom{n'}{2}}$ are affinely independent. Thus, $\mathcal{S} \cup \mathcal{T}$ is a set of affinely independent profiles and $\dim(\mathcal{S} \cup \mathcal{T}) = |\mathcal{S} \cup \mathcal{T}| - 1 = n! - n'$. The above stated fact that $\text{uni}([n']) \in f(R_B^{\pi})$ for every $B \subseteq [n']$ and $\pi \in \Pi_B^{\circ}([n'])$ finishes the proof. \square

We now consider PSCFs that may return a lottery that is *not* maximal. The following lemma shows that for every such PSCF there is a set of profiles with a strict Condorcet winner for which it returns the uniform lottery over a fixed subset of alternatives if we additionally require population-consistency and composition-consistency. Furthermore, this set of profiles has only one regular profile in its linear hull. Later this statement is leveraged to show that every population-consistent and composition-consistent PSCF returns a subset of maximal lotteries.

Lemma 12. *Let f be a PSCF that satisfies population-consistency and composition-consistency. If $f \not\subseteq ML$, there are $A' \subseteq A \in \mathcal{F}(U)$ and $\mathcal{Y} \subseteq \mathcal{R}|_A$ of dimension $|A'| - 1$ such that $\text{uni}(A') \in f(R)$ for every $R \in \mathcal{Y}$ and $\dim(\text{lin}(\mathcal{Y}) \cap \text{lin}(\mathcal{R}|_A^{A'})) = 1$.*

Proof. If $f \not\subseteq ML$, there are $A \in \mathcal{F}(U)$, $R \in \mathcal{R}|_A$, and $p \in f(R)$ such that $p \notin ML(R)$. Since p is not a maximal lottery, by definition, there is $q \in \Delta(A)$ such that $q^T M_R p > 0$. Linearity of matrix multiplication implies that there is $x \in A$ such that $(M_R p)_x > 0$, where $(M_R p)_x$ is the entry of $M_R p$ corresponding to x . We first use composition-consistency to “blow up” alternatives such that the resulting lottery is uniform. Let κ be the greatest common divisor of $\{p_y : y \in A\}$, i.e., $\kappa = \max\{s \in \mathbb{Q} : p_y/s \in \mathbb{N} \text{ for all } y \in A\}$. For every $y \in A$, let $A_y \in \mathcal{F}(U)$ such that $|A_y| = \max\{1, p_y/\kappa\}$, $A_y \cap A = \{y\}$, and all A_y are pairwise disjoint. Moreover, let $A^u = \bigcup_{y \in A} A_y$. Now, choose $R^u \in \mathcal{R}|_{A^u}$ such that $R^u|_A = R$, A_y is a component in R^u for every $y \in A$, and $R^u|_{A_y} = \text{uni}(\mathcal{L}(A_y))$ for every $y \in A_y$. Hence, $\text{uni}(A_y) \in f(R^u|_{A_y})$ for all $y \in A$ as f is neutral and $f(R^u|_{A_y})$ is convex. To simplify notation, let $A^p = \bigcup_{y \in \text{supp}(p)} A_y$. By composition-consistency, it follows that $p' = \text{uni}(A^p) \in f(R^u)$. Observe that

$$(M_{R^u} p')_x = \sum_{y \in \text{supp}(p) \setminus \{x\}} \frac{|A_y|}{|A^p|} M_{R^u}(x, y) = \sum_{y \in A \setminus \{x\}} p_y M_R(x, y) > 0.$$

We now construct a profile $R' \in \mathcal{R}|_{A^u}$ such that x is a strict Condorcet winner in R' and $\text{uni}(A^p) \in f(R')$. To this end, let $R' \in \mathcal{R}|_{A^u}$ be the uniform mixture of all profiles that arise from R^u by permuting all alternatives in $A^p \setminus \{x\}$, i.e.,

$$R' = \frac{1}{|A^p \setminus \{x\}|!} \sum_{\substack{\pi \in \Pi(A^u): \pi(y)=y \\ \text{for all } y \in A^u \setminus A^p \cup \{x\}}} \pi(R^u).$$

Then, $M_{R'}(x, y) = M_{R'}(x, z) > 0$ for all $y, z \in A^p \setminus \{x\}$. Neutrality and population-consistency imply that $p' \in f(R')$.

Let $R^{\text{uni}} = \text{uni}(\mathcal{L}(A^u))$ and define, for $\lambda \in [0, 1]$,

$$R^\lambda = \lambda R' + (1 - \lambda) R^{\text{uni}}.$$

It follows from Lemma 6 that $y \in f(R^{\text{uni}})$ for all $y \in A^u$. Convexity of $f(R^{\text{uni}})$ implies that $f(R^{\text{uni}}) = \Delta(A^u)$. Hence, by population-consistency, $p' \in f(R^\lambda)$ for all $\lambda \in [0, 1]$.

Now, let $S \in \mathcal{R}|_{A^u}$ such that $M_S(y, z) = 0$ for all $y, z \in A^p \cup \{x\}$ and $M_S(y, z) = 1$ for all $y \in A^p \cup \{x\}$, $z \in A^u \setminus (A^p \cup \{x\})$. For $\lambda \in [0, 1]$, let

$$S^\lambda = \lambda S + (1 - \lambda) R^{\text{uni}}.$$

Note that every $y \in A^p \cup \{x\}$ is a Condorcet winner in S^λ . It follows from population-consistency and Lemma 6 that, for small $\lambda > 0$, $y \in f(S^\lambda)$ for all $y \in A^p \cup \{x\}$ and, by convexity of $\Delta(A^p \cup \{x\}) \subseteq f(S^\lambda)$. In particular, $p' \in f(S^\lambda)$ for small $\lambda > 0$.

Finally, let

$$R^x = 1/3 R^\lambda + 2/3 S^\lambda,$$

for some small $\lambda > 0$. Population-consistency implies that $p' \in f(R^x)$. Moreover, $M_{R^x}(x, y) > 0$ for all $y \in A^u \setminus \{x\}$, i.e., x is a strict Condorcet winner in R^x , and hence, it follows from Lemma 6 that $x \in f(R^x)$.

If $p_x > 0$ then, by construction, $p' = \text{uni}(A^p \cup \{x\}) \in f(R^x)$. If $p_x = 0$ then $p' = \text{uni}(A^p) \in f(R^x)$. In this case it follows from convexity of $f(R^x)$ that $\text{uni}(A^p \cup \{x\}) = 1/(|A^p|+1)x + |A^p|/(|A^p|+1)\text{uni}(A^p) \in f(R^x)$.

Hence, in either case, we get a profile R^x such that $\text{uni}(A^p \cup \{x\}) \in f(R^x)$ and $M^x = M_{R^x}$ restricted to $A^p \cup \{x\}$ takes the form

$$M^x = \lambda \cdot \begin{pmatrix} 0 & \dots & 0 & -1 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & -1 & 0 & \dots & 0 \\ 1 & \dots & 1 & 0 & 1 & \dots & 1 \\ 0 & \dots & 0 & -1 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & -1 & 0 & \dots & 0 \end{pmatrix}$$

for some $\lambda > 0$ where all entries except the x th row and column are zero. Let $n' = |A^p \cup \{x\}|$. For every $y \in A^p$, let $M^y \in \mathbb{Q}^{n' \times n'}$ such that $M^y(y, z) = -M^y(z, y) = \lambda$ for all $z \neq y$ and 0 otherwise. Let $\pi^y \in \Pi(A^u)$ such that $\pi^y(x) = y$ and $\pi^y(z) = z$ for all $z \in A^u \setminus (A^p \cup \{x\})$ and $R^y = \pi^y(R^x)$. Then, for every $y \in A^p \cup \{x\}$, $M_{R^y} = M^y$ and, by neutrality, $\text{uni}(A^p \cup \{x\}) \in f(R^y)$.

Let $\mathcal{Y} = \{R^y : y \in A^p \cup \{x\}\}$. We have that $\dim(\mathcal{Y}) = |A^p \cup \{x\}| - 1$ since \mathcal{Y} is a set of affinely independent vectors. Now we determine $\dim(\text{lin}(\mathcal{Y}) \cap \text{lin}(\mathcal{R}|_{A^u}^{A^p \cup \{x\}}))$. To this end, let $\lambda R \in \text{lin}(\mathcal{Y}) \cap \text{lin}(\mathcal{R}|_{A^u}^{A^p \cup \{x\}})$ with $\lambda \in \mathbb{Q}$ and $R \in \mathcal{R}|_{A^u}$. Then $\lambda M_R = \sum_{z \in A^p \cup \{x\}} \lambda_z M^z$ and $\lambda \sum_{z \in A^p \cup \{x\}} M_R(y, z) = 0$ for all $y \in A^p \cup \{x\}$. This implies that $(n' - 1)\lambda^y = \sum_{z \in A^p \cup \{x\} \setminus \{y\}} \lambda^z$ for all $y \in A^p \cup \{x\}$. Hence, $\lambda^y = \lambda^z$ for all $y, z \in A^p \cup \{x\}$ and $\text{lin}(\mathcal{Y}) \cap \text{lin}(\mathcal{R}|_{A^u}^{A^p \cup \{x\}}) = \{\lambda \sum_{y \in A^p \cup \{x\}} R^y : \lambda \in \mathbb{Q}\}$. \square

In Lemma 13, we finally show that every PSCF that satisfies population-consistency and composition-consistency has to yield maximal lotteries. The structure of the proof is as follows. We assume for contradiction that a PSCF satisfies population-consistency and composition-consistency, but returns a lottery that is not maximal. Then we can find a set of profiles with full dimension for which the uniform lottery over a fixed subset of at least two alternatives is returned and the uniform profile is in its interior. Thus, this set contains a profile with a strict Condorcet winner that is close to the uniform profile. For every profile in an ε -ball around this strict Condorcet profile, the function has to return the uniform lottery over a non-singleton subset as well as the lottery with probability 1 on the Condorcet winner, which contradicts decisiveness.

Lemma 13. *Every PSCF f that satisfies population-consistency and composition-consistency has to yield maximal lotteries, i.e., $f \subseteq ML$.*

Proof. Let f be a PSCF that satisfies population-consistency and composition-consistency and $A \in \mathcal{F}(U)$. For $|A| = 2$ the statement follows from Lemma 5. For $|A| > 2$, assume for contradiction that $f \not\subseteq ML$. By Lemma 12, there is $A' \subseteq A$ and $\mathcal{Y} \subseteq \mathcal{R}|_A$ of dimension $|A'| - 1$ such that $\text{uni}(A') \in f(R)$ for every $R \in \mathcal{Y}$ and $\text{lin}(\mathcal{Y}) \cap \text{lin}(\mathcal{R}|_A^{A'})$ has dimension 1. By Lemma 11, there is $\mathcal{X} \subseteq \mathcal{R}|_A^{A'}$ of dimension $|A|! - |A'|$ such that $\text{uni}(A') \in f(R)$ for every $R \in \mathcal{X}$. Since $0 \notin \mathcal{X}$ and $0 \notin \mathcal{Y}$, $\text{lin}(\mathcal{X})$ has dimension $|A|! - |A'| + 1$ and $\text{lin}(\mathcal{Y})$ has dimension $|A'|$. Thus, $\text{lin}(\mathcal{X} \cup \mathcal{Y})$ has dimension $|A|!$. This implies that $\mathcal{X} \cup \mathcal{Y}$ has dimension $|A|! - 1$.

Furthermore, it follows from population-consistency that $\text{uni}(A') \in f(R)$ for every $R \in \text{conv}(\mathcal{X} \cup \mathcal{Y})$. Since $\text{uni}(\mathcal{L}(A))$ is in the interior of $\text{conv}(\mathcal{X} \cup \mathcal{Y})$, there are $x \in A'$ and $R^x \in \text{int}_{\mathcal{R}|_A}(\mathcal{X} \cup \mathcal{Y})$ such that x is a strict Condorcet winner in R^x . Hence, there is $\varepsilon > 0$ such that, for every $R' \in B_\varepsilon(R^x) \cap \mathcal{R}|_A$, $R' \in \text{conv}(\mathcal{X} \cup \mathcal{Y})$ and x is a strict Condorcet winner in R' . Then, we get that $x \in f(R')$ and $\text{uni}(A') \in f(R')$ for every $R' \in B_\varepsilon(R^x) \cap \mathcal{R}|_A$. Thus, $\{R' \in \mathcal{R}|_A : |f(R', A)| = 1\}$ is not dense in $\mathcal{R}|_A$ at R^x . This contradicts decisiveness of f . \square

C.4. $ML \subseteq f$

In this section we show that every PSCF f that satisfies population-consistency and composition-consistency has to yield *all* maximal lotteries. To this end, we first prove an auxiliary lemma. It was shown by McGarvey (1953) that every complete and anti-symmetric relation is the majority relation of some profile with a bounded number of voters. We show an analogous statement for skew-symmetric matrices and fractional preference profiles.

Lemma 14. *Let $M \in \mathbb{Q}^{n \times n}$ be a skew-symmetric matrix. Then, there are $R \in \mathcal{R}|_{[n]}$ and $c \in \mathbb{Q}_{>0}$ such that $cM = M_R$. Furthermore, if there is $\pi \in \Pi([n])$ such that $M(i, j) = M(\pi(i), \pi(j))$ for all $i, j \in [n]$, then $R = \pi(R)$.*

Proof. For all $i, j \in [n]$ with $i \neq j$, let $R^{ij} \in \mathcal{R}|_{[n]}$ be the profile such that $R^{ij}(\succ) = 1/(n-1)!$ if $i \succ j$ and $\{i, j\}$ is a component in R^{ij} and $R^{ij}(\succ) = 0$ otherwise. By construction, we have that $R^{ij}(i, j) = 1$ and $R^{ij}(x, y) = 0$ for all $\{x, y\} \neq \{i, j\}$. Let $c = 1/\sum_{i,j: M(i,j)>0} M(i, j)$ and $R = c \sum_{i,j: M(i,j)>0} M(i, j)R^{ij}$. Then, we have that $M_R = cM$. The second part of the lemma follows from the symmetry of the construction. \square

For profiles which admit a unique maximal lottery, it follows from Lemma 13 that $f = ML$. It turns out that every maximal lottery that is a vertex of the set of maximal lotteries in one of the remaining profiles is the limit point of a sequence of maximal lotteries of a sequence of profiles with a unique maximal lottery converging to the original profile. The proof of Lemma 15 heavily relies on the continuity of f .

Lemma 15. *Let f be a PSCF that satisfies population-consistency and composition-consistency. Then, $ML \subseteq f$.*

Proof. Let f be a PSCF that satisfies population-consistency and composition-consistency, $A \in \mathcal{F}(U)$, and $R \in \mathcal{R}|_A$. It follows from Lemma 13 that $f \subseteq ML$. By neutrality, we can assume without loss of generality that $A = [n]$ and for simplicity $M = M_R$. We want to show that $f(R) = ML(R)$. If $ML(R)$ is a singleton, it follows from $f \subseteq ML$ that $f(R) = ML(R)$. Hence, consider the case where $ML(R)$ is not a singleton. Let $p \in ML(R)$ and assume without loss of generality that $\text{supp}(p) = [k]$.

We first consider the case where k is odd. By Lemma 14, there are $S \in \mathcal{R}|_A$ and $c \in \mathbb{Q}_{>0}$

such that

$$M_S = c \begin{pmatrix} 0 & -\frac{1}{p_1 p_2} & 0 & \dots & 0 & \frac{1}{p_k p_1} & 1 & \dots & 1 \\ \frac{1}{p_1 p_2} & & & & & 0 & \vdots & \ddots & \vdots \\ 0 & & & & & \vdots & \vdots & \ddots & \vdots \\ \vdots & & & & & 0 & \vdots & \ddots & \vdots \\ 0 & & & & & -\frac{1}{p_{k-1} p_k} & 1 & \dots & 1 \\ -\frac{1}{p_k p_1} & 0 & \dots & 0 & \frac{1}{p_{k-1} p_k} & 0 & 1 & \dots & 1 \\ \hline -1 & & \dots & & & -1 & 0 & \dots & 0 \\ \vdots & & \ddots & & & \vdots & \vdots & \ddots & \vdots \\ -1 & & \dots & & & -1 & 0 & \dots & 0 \end{pmatrix}$$

Intuitively, M_S defines a weighted cycle on $[k]$. Note that $(p^T M_S)_i = 0$ for all $i \in \text{supp}(p)$ and $(p^T M_S)_i > 0$ for all $i \in A \setminus \text{supp}(p)$, i.e., p is a quasi-strict maximin strategy in M_S in the sense of Harsanyi (1973b). Since p is a maximin strategy in M_S , it follows that $p \in ML(S)$. For $\varepsilon \in [0, 1]$, we define $R^\varepsilon = (1 - \varepsilon)R + \varepsilon S$ and $M^\varepsilon = M_{R^\varepsilon}$. Population-consistency implies that $p \in ML(R^\varepsilon)$ for all $\varepsilon \in [0, 1]$. Observe that p is a quasi-strict maximin strategy in M^ε for every $\varepsilon \in (0, 1]$. Hence, for every maximin strategy q in M^ε , it follows that $(q^T M^\varepsilon)_i = 0$ for every $i \in [k]$ and $q_i = 0$ for every $i \notin [k]$. It follows from basic linear algebra that

$$\det \left((M_S(i, j))_{i, j \in [k-1]} \right) = c^{k-1} \prod_{i=1}^{k-2} \left(\frac{1}{p_i p_{i+1}} \right)^2 \neq 0,$$

and hence, $(M_S(i, j))_{i, j \in [k]}$ has rank at least $k - 1$. More precisely, $(M_S(i, j))_{i, j \in [k]}$ has rank $k - 1$, since skew-symmetric matrices of odd size cannot have full rank.²³ Furthermore, $\det((M^\varepsilon(i, j))_{i, j \in [k-1]})$ is a polynomial in ε of order at most $k - 1$ and hence, has at most $k - 1$ zeros. Thus, we can find a sequence $(\varepsilon_l)_{l \in \mathbb{N}}$ which converges to zero such that $(M^{\varepsilon_l}(i, j))_{i, j \in [k]}$ has rank $k - 1$ for all $l \in \mathbb{N}$. In particular, if $(q^T M^\varepsilon)_i = 0$ for all $i \in [k]$, then $q = p$. This implies that p is the unique maximin strategy in M^{ε_l} for all $l \in \mathbb{N}$ and hence, $\{p\} = ML(R^{\varepsilon_l}) \subseteq f(R^{\varepsilon_l})$ for all $l \in \mathbb{N}$. It follows from continuity of f that $p \in f(R)$.

Now we consider the case where k is even. $ML(R)$ is a polytope because it is the solution space of a linear feasibility program. Assume that p is a vertex of $ML(R)$. We first show that p is not quasi-strict.²⁴ Assume for contradiction that p is quasi-strict, i.e., $(p^T M)_i > 0$ for all $i \notin [k]$. Then, $\text{supp}(q) \subseteq [k]$ for every maximin strategy q of M . But then $(1 + \varepsilon)p - \varepsilon q$ is also a maximin strategy in M for small $\varepsilon > 0$ as p is a quasi-strict maximin strategy in M . This contradicts the assumption that p is a vertex of $ML(R)$.

²³A skew-symmetric matrix M of odd size cannot have full rank, since $\det(M) = \det(M^T) = \det(-M) = (-1)^n \det(M) = -\det(M)$ and, hence, $\det(M) = 0$.

²⁴The proof of this statement does not make use of the fact that k is even and therefore also holds (but is not needed) for odd k .

Hence, we may assume without loss of generality that $(p^T M)_{k+1} = 0$. Let $e_1 = M(k+1, 1)/p_2$ and $e_i = (M(k+1, i) + p_{i-1}e_{i-1})/p_{i+1}$ for $i \in \{2, \dots, k-1\}$. By Lemma 14, there are $S \in \mathcal{R}|_A$ and $c \in \mathbb{Q}_{>0}$ such that

$$M_S = c \begin{pmatrix} 0 & e_1 & 0 & \dots & 0 & 0 & 1 & \dots & 1 \\ -e_1 & & & \ddots & \vdots & & & & \\ 0 & & \ddots & & 0 & \vdots & \vdots & \ddots & \vdots \\ \vdots & \ddots & & & e_{k-1} & & & & \\ 0 & \dots & 0 & -e_{k-1} & 0 & 0 & & & \\ \hline 0 & & \dots & & 0 & 0 & 1 & \dots & 1 \\ \hline -1 & & & \dots & & -1 & 0 & \dots & 0 \\ \vdots & & & \ddots & & \vdots & \vdots & \ddots & \vdots \\ -1 & & & \dots & & -1 & 0 & \dots & 0 \end{pmatrix}$$

Note that $M_S(1, k) = M_S(k, 1) = 0$. For $\varepsilon > 0$, let $R^\varepsilon = (1 - \varepsilon)R + \varepsilon S$ and $M^\varepsilon = M_{R^\varepsilon}$. We claim that p^ε defined as follows is a maximin strategy in M^ε . To this end, let $s_\varepsilon = \frac{\varepsilon c}{1 - \varepsilon + \varepsilon c}$.

$$p_i^\varepsilon = \begin{cases} (1 - s_\varepsilon)p_i & \text{if } i \in [k], \\ s_\varepsilon & \text{if } i = k+1, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Note that $1/c(p^T M_S)_1 = -p_2 e_1 = -M(k+1, 1)$ and, for $i \in \{2, \dots, k-1\}$,

$$\frac{1}{c}(p^T M_S)_i = p_{i-1}e_{i-1} - p_{i+1}e_i = p_{i-1}e_{i-1} - (M(k+1, i) + p_{i-1}e_{i-1}) = -M(k+1, i).$$

To determine $(p^T M_S)_k$, we first prove inductively that $p_i e_i = 1/p_{i+1} \sum_{j=1}^i M(k+1, j)p_j$ for all $i \in [k-1]$. For $i = 1$, this follows from the definition of e_1 . Now, let $i \in \{2, \dots, k-1\}$. Then,

$$\begin{aligned} p_i e_i &= \frac{p_i}{p_{i+1}} (M(k+1, i) + p_{i-1}e_{i-1}) = \frac{p_i}{p_{i+1}} (M(k+1, i) + \frac{1}{p_i} \sum_{j=1}^{i-1} M(k+1, j)p_j) \\ &= \frac{1}{p_{i+1}} \sum_{j=1}^i M(k+1, j)p_j, \end{aligned}$$

where the second equality follows from the induction hypothesis. Now,

$$\frac{1}{c}(p^T M_S)_k = p_{k-1}e_{k-1} = \frac{1}{p_k} \sum_{j=1}^{k-1} M(k+1, j)p_j = -\frac{1}{p_k} M(k+1, k)p_k = -M(k+1, k),$$

where the third equality follows from the fact that $(p^T M)_{k+1} = 0$.

For $i \in [k]$, it follows from $(p^T M)_i = 0$ that $((p^\varepsilon)^T M)_i = s_\varepsilon M(k+1, i)$. Then, for $i \in [k]$,

$$((p^\varepsilon)^T M^\varepsilon)_i = (1 - \varepsilon)s_\varepsilon M(k+1, i) + \varepsilon c(1 - s_\varepsilon)(-M(k+1, i)) = 0.$$

Furthermore, it follows from $(p^T M)_{k+1} = 0$ that $((p^\varepsilon)^T M^\varepsilon)_{k+1} = 0$ as $M(k+1, k+1) = 0$, and, for $i \in A \setminus [k+1]$,

$$((p^\varepsilon)^T M^\varepsilon)_i \geq (1 - \varepsilon)s_\varepsilon M(k+1, i) + \varepsilon c \geq -(1 - \varepsilon)s_\varepsilon + \varepsilon c > 0.$$

This shows that p^ε is a maximin strategy in M^ε and hence, $p^\varepsilon \in ML(R^\varepsilon)$. Since $|\text{supp}(p^\varepsilon)|$ is odd, it follows from the first case that $p^\varepsilon \in f(R^\varepsilon)$. Note that s_ε goes to 0 as ε goes to 0. Hence, p^ε goes to p as ε goes to 0. It now follows from continuity of f that $p \in f(R)$.

Together, we have that $p \in f(R)$ for every vertex p of $ML(R)$. Since every lottery in $ML(R)$ can be written as a convex combination of vertices, convexity of $f(R)$ implies that $f(R) = ML(R)$. \square

Theorem 2 then follows directly from Lemmas 13 and 15.

Theorem 2. *A PSCF f satisfies population-consistency and composition-consistency if and only if $f = ML$.*