

CHAPTER 1

Rolling the Dice: Recent Results in Probabilistic Social Choice

Felix Brandt

Casting the lot puts an end to disputes and decides between powerful contenders.

— Solomon, c. 900 BC (Proverbs 18:18, RSV)

1.1 Introduction

When aggregating the preferences of multiple agents into one collective choice, 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 two agents such that one prefers a and the other one b , there is no deterministic way of selecting a single alternative without violating one of two basic fairness conditions known as *anonymity* and *neutrality*. Anonymity requires that the collective choice ought to be independent of the agents' identities whereas neutrality requires impartiality towards the alternatives.¹ Allowing lotteries as social outcomes hence seems like a necessity for impartial collective choice. Indeed, most common “deterministic” social choice functions such as plurality rule, Borda's rule, or Copeland's rule are only deterministic as long as there is no tie, which is usually resolved by drawing a lot. 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), and has recently gained increasing attention in political science (see, e.g., Goodwin, 2005; Dowlen, 2009; Stone, 2011; Guerrero, 2014).

It turns out that randomization—apart from guaranteeing impartiality—allows the circumvention of well-known impossibility results such as the Gibbard-Satterthwaite Theorem. Important questions in this context are how much “randomness” is required to achieve positive results and which assumptions are made about the agents' preferences over lotteries. In this chapter, I will survey some recent axiomatic results in the area of probabilistic social choice.

Probabilistic social choice functions (PSCFs) map collections of individual preference relations over alternatives to lotteries over alternatives and were first for-

¹Moulin (1983, pp. 22–25) has provided a complete characterization that shows for which numbers of alternatives and agents there are deterministic single-valued social choice functions that satisfy anonymity and neutrality when individual preferences are strict.

mally 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 (i.e., non-manipulable) PSCFs (Gibbard, 1977). An important corollary of Gibbard's characterization, attributed to Hugo Sonnenschein, concerns the most studied PSCFs: *random dictatorships*. In random dictatorships, one of the agents is picked at random and his most preferred alternative is implemented as the social choice. Gibbard (1977) has shown that *random dictatorships* are the only strategyproof and *ex post* efficient PSCFs. While Gibbard's result 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, 1979b). In contrast to deterministic dictatorships, the uniform random dictatorship (henceforth, *RD*), in which every agent is picked with the same probability, enjoys a high degree of fairness and is in fact used in many subdomains of social choice that are concerned with the fair assignment of objects to agents (see, e.g., Abdulkadiroğlu and Sönmez, 1998; Bogomolnaia and Moulin, 2001; Che and Kojima, 2010; Budish et al., 2013).

One may wonder how Gibbard defined strategyproofness for PSCFs since, in his framework, agents submit their preferences over alternatives, but no preferences over lotteries. Preferences over lotteries are often defined by assuming the existence of a *von Neumann-Morgenstern (vNM) utility function* which assigns cardinal utility values to alternatives. A lottery is preferred to another lottery if the former yields more expected utility than the latter. The notion of strategyproofness considered by Gibbard is a rather strong one. According to his definition, a PSCF is strategyproof if, for *all* vNM utility functions that are compatible with the ordinal preferences, submitting one's true preferences yields at least as much expected utility as submitting any other preference relation. This notion of strategyproofness is sometimes also referred to as *strong SD-strategyproofness* (see Section 1.3.2). According to strong *SD*-strategyproofness, a PSCF may be deemed manipulable just because it can be manipulated for some contrived and highly unlikely vNM utility representations. While it is good to know that *RD* satisfies such a high degree of strategyproofness, an interesting question is whether there are other—perhaps more attractive and “less randomized”—PSCFs that satisfy weaker notions of strategyproofness.

Since there are various problems associated with asking agents to submit their complete preference relations over all lotteries, a common approach to defining axiomatic properties of PSCFs is to systematically extend the agents' preferences over alternatives to (possibly incomplete) preferences over lotteries via so-called *lottery extensions*.² In Section 1.3, I will define a number of lottery extensions, which will in turn lead to varying notions of strategyproofness, efficiency, and participation. On top of that, I will discuss several consistency conditions, which are not based on the individual preferences over lotteries. One such condition is *population-consistency* which requires that whenever a PSCF

²Preference relations over lotteries may, for example, not allow for a concise representation. Moreover, and perhaps more importantly, agents are in many cases not even aware of their complete preferences over lotteries. Even if they *think* they can competently assign vNM utilities to alternatives, these assignments are prone to be based on arbitrary choices.

returns the same lottery for two disjoint electorates, then this lottery should also be returned for the union of both electorates. In Section 1.4, I will review positive and negative axiomatic results for PSCFs. Particular attention will be paid to the case of weak individual preferences, i.e., preferences that may contain ties. Allowing weak preferences can lead to results that significantly differ from those for strict preferences; positive results may turn into impossibilities and easy computational problems may become intractable. In many important subdomains of social choice such as assignment, matching, and coalition formation, ties are unavoidable because agents are indifferent among all outcomes in which their allocation, match, or coalition is the same.

It is impossible to completely cover the topic of probabilistic social choice in this chapter. The selection of results is certainly biased towards work that I was involved in and I apologize in advance for any omissions. In particular, there has been interesting computational work on establishing hardness of manipulation via randomization (Conitzer and Sandholm, 2003; Elkind and Lipmaa, 2005; Walsh and Xia, 2012), approximating deterministic voting rules (Procaccia, 2010; Birrell and Pass, 2011; Service and Adams, 2012), and measuring the worst-case utilitarian performance of randomized voting rules (Anshelevich et al., 2015; Anshelevich and Postl, 2016; Gross et al., 2017).

1.2 Probabilistic Social Choice Functions

Let $N = \{1, \dots, n\}$ be a set of agents and A a finite set of m alternatives. Every agent $i \in N$ is equipped with a complete and transitive *preference relation* $\succsim_i \subseteq A \times A$, the strict (or asymmetric) part of which is denoted by \succ_i . A preference relation \succsim_i is called *strict* if it is antisymmetric, i.e., it is identical to its strict part up to reflexivity. Otherwise, the preference relation is said to be *weak*. A *preference profile* maps each agent $i \in N$ to a preference relation.

The set of all *lotteries* (or *probability distributions*) over A is denoted by $\Delta(A)$, i.e.,

$$\Delta(A) = \left\{ p \in \mathbb{R}^m : p(x) \geq 0 \text{ for all } x \in A \text{ and } \sum_{x \in A} p(x) = 1 \right\}.$$

For convenience, I will also write lotteries as convex combinations of alternatives, e.g., $1/2 a + 1/2 b$ denotes the uniform distribution over $\{a, b\}$.³ A lottery p is *degenerate* if its support is of size 1, i.e., it puts all probability on a single alternative.

Our central object of study are PSCFs, i.e., functions that map a preference profile to a non-empty convex subset of lotteries.⁴ A PSCF is *anonymous* if its outcome is invariant under permutations of the agents. Similarly, a PSCF is *neutral* if permuting alternatives in the preference profile leads to lotteries in which alternatives are permuted accordingly.

³The lotteries returned by PSCFs 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.

⁴We consider set-valued PSCFs because *RSD* and *ML* may return more than one lottery. If there are sufficiently many agents, this is however almost never the case (see also Brandl et al., 2016c). Single-valued PSCFs are called *social decision schemes* (Gibbard, 1977).

In this chapter, we will consider four exemplary PSCFs, all of which are anonymous and neutral: random dictatorship, two probabilistic variants of Borda’s rule, and maximal lotteries.⁵

Random Dictatorship (RD). Perhaps the most-studied PSCF is *random dictatorship*, where one of the agents is picked uniformly at random and this agent’s most-preferred alternative is selected. Thus, the probabilities assigned by *RD* are directly proportional to the number of agents who top-rank a given alternative (or, in other words, the alternative’s plurality score). *RD* is only well-defined for strict preferences. In order to be able to deal with ties in the preferences, *RD* is typically extended to *random serial dictatorship (RSD)*. *RSD* selects a permutation of the agents uniformly at random and then sequentially allows agents in the order of the permutation to narrow down the set of alternatives to their most preferred of the remaining ones. This will always result in a single alternative unless there are two alternatives among which all agents are indifferent.⁶ While implementing *RSD* is straightforward, computing the resulting *RSD* probabilities is #P-complete and therefore intractable (Aziz et al., 2013a). Also, checking whether the *RSD* probability of a given alternative exceeds some fixed value from the interval $(0, 1)$ is NP-complete. Subsequent work has studied the parameterized complexity of these problems (Aziz and Mestre, 2014).

The remaining three PSCFs considered in this chapter are based on pairwise majority comparisons between alternatives. For a given profile of preferences, the $m \times m$ matrix of *majority margins* M is defined by

$$M_{xy} = |\{i \in N : x \succ_i y\}| - |\{i \in N : y \succ_i x\}|.$$

If the output of a neutral PSCF f only depends on M , f is called *pairwise*. Pairwiseness is an informational requirement and is formally defined by demanding that the output for two preference profiles, which give rise to the same majority margin matrix, has to be identical. An advantage of pairwise PSCFs is that they are applicable even when individual preferences are incomplete or intransitive.

Borda’s Rule. Traditionally, Borda’s rule is defined as a scoring rule in which each agent assigns a score of $m - 1$ to his most-preferred alternative, $m - 2$ to his second-most preferred alternative, etc. The alternatives with maximal accumulated score win. Alternatively, Borda scores can be obtained from the majority margin matrix.⁷ The Borda score of alternative x is $\sum_{y \in A} M_{xy}/2 + n$. We will discuss two probabilistic variants of Borda’s rule. The first one, *Borda_{max}* yields all lotteries that randomize over alternatives with *maximal* Borda score. The second one, *Borda_{pro}*, involves much more randomness and assigns probabilities to

⁵Other PSCFs not covered in this chapter include the recently proposed *maximal recursive rule* (Aziz, 2013), *egalitarian simultaneous reservation rule* (Aziz and Stursberg, 2014), and *2-Agree* (Gross et al., 2017).

⁶Simpler extensions of *RD* to weak preferences such as returning a uniform lottery over all first-ranked alternatives of a randomly selected agent typically suffer from a lack of *ex post* efficiency.

⁷This also yields a natural generalization of Borda’s rule for preferences that fail to be antisymmetric, complete, or even transitive. Borda’s rule is the only pairwise scoring rule.

the alternatives that are *proportional* to their Borda scores. Examples are given below.

One of the most influential notions in social choice theory is that of a *Condorcet winner*, i.e., an alternative that is preferred to every other alternative by some majority of agents. Formally, M admits a *Condorcet winner* if it contains a row in which all entries but one are strictly positive. Example 2 below shows that Borda's rule may fail to select a Condorcet winner. It is well-known that Condorcet winners do not exist in general (see Example 3 below). In fact, the absence of Condorcet winners—the so-called *Condorcet paradox*—is the root cause for central impossibility theorems in social choice theory such as Arrow's Theorem or the Gibbard-Satterthwaite Theorem. The essence of Condorcet's paradox is that there are voting situations in which *no matter* which alternative is selected, there will always be another alternative that is preferred by a majority of the agents. In other words, it is impossible to select an outcome that cannot be overturned by an organized majority of agents who all agree with which alternative it should be replaced.


Maximal Lotteries (ML). Maximal lotteries were first considered by Kreweras (1965) and independently rediscovered and studied in detail by Fishburn (1984a).⁸ A lottery p is *maximal* iff $p^T M \geq 0$. A maximal lottery p can thus be seen as a “randomized weak Condorcet winner”, i.e., a lottery that is weakly preferred to every other lottery by an *expected* majority of agents: $p^T M q \geq 0$ for all $q \in \Delta(A)$.⁹ See Example 3 below for a profile with no Condorcet winner, but a unique maximal lottery. Maximal lotteries are equivalent to the mixed maximin strategies (or Nash equilibria) of the symmetric zero-sum game given by M . In contrast to Condorcet winners, maximal lotteries are thus guaranteed to exist by von Neumann's Minimax Theorem. Moreover, most profiles admit a *unique* maximal lottery. This is, for example, the case when there is an odd number of agents with strict preferences (see Laffond et al., 1997; Le Breton, 2005). More generally, if the number of agents goes to infinity, the number of profiles with multiple maximal lotteries goes to zero. Maximal lotteries can be found in polynomial time by solving a linear feasibility problem.¹⁰

⁸Interestingly, maximal lotteries or variants thereof have been rediscovered again by economists (Laffond et al., 1993), mathematicians (Fisher and Ryan, 1995), political scientists (Felsenthal and Machover, 1992), and computer scientists (Rivest and Shen, 2010). In particular, the support of maximal lotteries, called the *bipartisan set* or the *essential set*, has received considerable attention. A number of scholars have recommended maximal lotteries for practical use (Felsenthal and Machover, 1992; Rivest and Shen, 2010; Brandl et al., 2016c; Hoang, 2017). Within the domain of random assignment, maximal lotteries are known as popular mixed matchings (see Chapter 6 of this book).

⁹ $p^T M q > 0$ iff the expected number of agents who prefer the alternative returned by p to that returned by q is at least as large as the expected number of agents who prefer the outcome returned by q to that returned by p . This is reminiscent of the *PC* lottery extension (see Section 1.3.2). However, when not taking the expectation over the number of agents and directly comparing lotteries using lottery extensions such as *SD* or *PC*, all lotteries can be overturned by some majority of agents in the absence of Condorcet winners (see, also Zeckhauser, 1969; Aziz, 2015).

¹⁰Brandt and Fischer (2008, Thm. 5) have shown that deciding whether an alternative receives positive probability in some maximal lottery is P-complete and therefore not amenable to parallelization.

Example 1 In the case of only two alternatives, a and b , the four considered PSCFs break down to two prototypical rules: the proportional lottery (left) and the simple majority rule (right).



It is easily seen that the simple majority rule maximizes the agents' average *ex ante* satisfaction (Fishburn and Gehrlein, 1977). For example, consider three agents, two of which prefer a to b and one of which prefers b to a . Then, under the proportional rule, the former two will be satisfied with probability $2/3$ and the latter one with probability $1/3$. Hence, the average probability of satisfaction is $5/9$, which is lower than that of the simple majority rule ($2/3$). This gap widens when agents are risk-averse.

The proportional rule, on the other hand, steers clear of the “tyranny of the majority” by giving agents with a minority opinion at least the chance of being satisfied. Depending on the concrete setting, this can be very desirable. However, one should be aware that the proportional rule (and, in fact, any rule different from the simple majority rule) can return alternatives that are majority-dominated and therefore subject to strong opposition or even resistance. In other words, there is the possibility of *ex post* majority dissatisfaction.¹¹

Example 2 Consider the following preference profile and its corresponding majority margin matrix.

$$\begin{array}{cc}
 \frac{3}{a} & \frac{2}{b} \\
 \frac{b}{b} & \frac{c}{c} \\
 \frac{c}{c} & \frac{a}{a}
 \end{array}
 \quad
 M = \begin{array}{ccc}
 & a & b & c \\
 a & \begin{pmatrix} 0 & 1 & 1 \\ -1 & 0 & 5 \\ -1 & -5 & 0 \end{pmatrix}
 \end{array}$$

The *RD* lottery is $3/5 a + 2/5 b$. The Borda score of a is $(0 + 1 + 1)/2 + 5 = 6$, that of b is $(-1 + 0 + 5)/2 + 5 = 7$, and that of c is $(-1 - 5 + 0)/2 + 5 = 2$. Hence, *Borda_{max}* returns b and *Borda_{pro}* returns $6/15 a + 7/15 b + 2/15 c$. The profile admits a Condorcet

¹¹Note that this is not possible when the outcomes of PSCFs are implemented as fractional allocations rather than actual lotteries (see Footnote 3).

winner because the first row of M is positive, except for the first entry. ML thus (uniquely) returns a because

$$(1 \ 0 \ 0) \cdot M = (0 \ 1 \ 1) \geq 0.$$

Example 3 Consider the following preference profile and its corresponding majority margin matrix.

$$\begin{array}{ccc}
 2 & 2 & 1 \\
 \hline
 a & b & c \\
 b & c & a \\
 c & a & b
 \end{array}
 \quad
 M = \begin{array}{ccc}
 & a & b & c \\
 a & 0 & 1 & -1 \\
 b & -1 & 0 & 3 \\
 c & 1 & -3 & 0
 \end{array}$$

In this example, RD yields $2/5 a + 2/5 b + 1/5 c$. The Borda scores of a , b , and c are 5, 6, and 3, respectively. Hence, $Borda_{max}$ returns b and $Borda_{pro}$ $5/14 a + 6/14 b + 3/14 c$. The pairwise majority relation is cyclic and there is no Condorcet winner. The unique maximal lottery returned by ML is $3/5 a + 1/5 b + 1/5 c$ because

$$(3/5 \ 1/5 \ 1/5) \cdot M = (0 \ 0 \ 0) \geq 0.$$

1.3 Axioms

The axioms considered in this chapter can be roughly divided into two subgroups: those that are independent of the agents' preferences over lotteries and those that do require preferences over lotteries.¹²

1.3.1 Consistency

We first discuss consistency axioms belonging to the first category. Non-probabilistic versions of these axioms have been widely studied in the literature.

Condorcet-consistency. A PSCF is *Condorcet-consistent* if it uniquely returns a lottery that puts probability 1 on the Condorcet winner whenever a Condorcet winner exists. Condorcet-consistency, which goes back to the 18th century, is one of the oldest formal axioms in social choice theory and considered by many to be desirable (see, e.g., Black, 1958; Fishburn, 1977; Campbell and Kelly, 2003; Dasgupta and Maskin, 2008).

Agenda-consistency. Part of the motivation of Condorcet-consistency is that an alternative that emerges as the unequivocal winner in all pairwise comparisons should also be chosen from the entire set of alternatives. Rational choice theory continues this train of thought by specifying a number of axioms that

¹²Apart from the axioms considered here, some authors have proposed "fairness" conditions for PSCFs such as an axiom that prescribes that every agent should receive positive probability on at least one alternative he does not rank last (Bogomolnaia et al., 2005; Duddy, 2015).

deal with choices from variable subsets of alternatives and postulating whether these choices are consistent with each other. These axioms can be transferred to probabilistic social choice simply by restricting the preference profile in question to a subset of alternatives and observing which lotteries a PSCF returns for the reduced profile. Let p be a lottery and A, B two subsets of alternatives such that p 's support is contained in both A and B . Then, what we call *agenda-consistency* requires that p is returned for A and B iff it is returned for the union of A and B . The implication from left to right is known as Sen's γ or *expansion*, whereas the implication from right to left is Sen's α or *contraction* (see Sen, 1971, 1977, 1986; Schwartz, 1976).

Population-consistency. A PSCF is *population-consistent* if, whenever it returns the same lottery for two preference profiles (defined on disjoint sets of agents), it also returns the same lottery for a profile that results by merging both profiles.¹³ 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 (1974), and Fine and Fine (1974) and features prominently in the characterization of scoring rules by Smith (1973) and Young (1975) as well as the characterization of Kemeny's rule by Young and Levenglick (1978).

Cloning-consistency and composition-consistency. *Cloning-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 and therefore constitute a contiguous interval in each agent's preference ranking. This condition was first considered by Tideman (1987) (see also Zavist and Tideman, 1989). Cloning-consistency imposes no restrictions on the relative probabilities of alternatives within a component. *Composition-consistency* is stronger than cloning-consistency and additionally requires 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. It was first considered by Laffond et al. (1996) and has been analyzed from a computational point of view by Brandt et al. (2011). Cloning-consistency implies neutrality (Brandt et al., 2016c, Lem. 1).

Apart from their intuitive appeal, these axioms can be motivated by the desire to prevent a central planner from strategically tampering with the set of feasible alternatives (e.g., by removing irrelevant alternatives or by introducing variants of alternatives) and the set of agents (e.g., by partitioning the electorate into subelectorates). For formal definitions and examples, the reader is referred to Brandt et al. (2016c).

¹³A slightly stronger variant of this axiom is also known as *reinforcement*.

1.3.2 Efficiency, Strategyproofness, and Participation

Several important axioms require the specification of preferences over lotteries. We will generate these preferences by systematically lifting a preference relation over alternatives to possibly incomplete preferences over lotteries. Formally, for any given preference relation \succsim on A and any pair of lotteries $p, q \in \Delta(A)$, a lottery extension \mathcal{E} prescribes whether $p \succsim^{\mathcal{E}} q$. The strict part $\succ^{\mathcal{E}}$ of $\succsim^{\mathcal{E}}$ is defined by letting $p \succ^{\mathcal{E}} q$ iff $p \succsim^{\mathcal{E}} q$ and not $q \succsim^{\mathcal{E}} p$. We will consider five different lottery extensions in this section. For all examples we assume that the underlying preference relation is $a \succ b \succ c$.

The first, and most conservative, lottery extension we consider is called *deterministic dominance* (DD') and postulates that p is preferable to q iff any alternative possibly returned by p is strictly better than any alternative possibly returned by q . In other words,

$$p \succsim^{DD'} q \quad \text{iff} \quad \forall x, y: [p(x) \cdot q(y) > 0 \Rightarrow x \succ y]. \quad (DD')$$

A variant of this extension can be defined using the weak preference relation rather than the strict one.

$$p \succsim^{DD} q \quad \text{iff} \quad \forall x, y: [p(x) \cdot q(y) > 0 \Rightarrow x \succsim y]. \quad (DD)$$

Hence, $p \succ^{DD} q$ iff every alternative returned by p is at least as good as every alternative returned by q with at least one strict preference. An agent may thus strictly prefer one lottery to another even though he is eventually indifferent between particular instantiations of the lotteries. Clearly, whether $p \succ^{DD} q$ or $p \succ^{DD'} q$ only depends on the supports of p and q .¹⁴ DD' only allows the comparison of lotteries with disjoint supports whereas the supports may overlap for DD as long as the agent is indifferent between all alternatives contained in the intersection of both supports. For example, $2/3a + 1/3b \succ^{DD'} c$ and $2/3a + 1/3b \succ^{DD} 1/2b + 1/2c$. DD' and DD may seem rather crude, but very risk-averse agents who seek to avoid uncertainty under any circumstances may subscribe to these preference extension. Furthermore, many PSCFs based on deterministic social choice functions already violate DD' -strategyproofness.

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

$$p \succsim^{BD} q \quad \text{iff} \quad \forall x, y \in A: [(x \succ y \Rightarrow p(x) \cdot q(y) \geq p(y) \cdot q(x)]. \quad (BD)$$

Apart from its intuitive appeal, the main motivation for BD is that p bilinearly dominates q iff p is preferable to q for every skew-symmetric bilinear (SSB) utility function consistent with \succsim (cf. Fishburn, 1984b; Aziz et al., 2015).¹⁵ For example, $1/2a + 1/2b \succ^{BD} 1/3a + 1/3b + 1/3c$.

¹⁴Within the context of set-valued social choice functions, DD is known as Kelly's preference extension (see, e.g., Kelly, 1977; Brandt, 2015).

¹⁵SSB utility theory is a generalization of von Neumann and Morgenstern's linear expected utility theory, which does not require the controversial independence axiom and transitivity (see, e.g., Fishburn, 1988).

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

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

For example, $1/2 a + 1/2 c \succ^{SD} 1/2 b + 1/2 c$. It is well-known that $p \succsim^{SD} q$ iff, for every vNM utility function compatible with \succsim , the expected utility for p is at least as large as that for q (see, e.g., Brandt et al., 2016a, Lem. 2).

The last lottery extension we consider is called *pairwise comparison (PC)* and postulates that p should be preferred to q iff the probability that p yields a better alternative than q is at least as large as the other way round (Aziz et al., 2015). Formally,

$$p \succsim^{PC} q \quad \text{iff} \quad \sum_{x,y: x \succ y} p(x) \cdot q(y) \geq \sum_{x,y: x \succ y} q(x) \cdot p(y). \quad (PC)$$

For example, $2/3 a + 1/3 c \succ^{PC} b$. The terms in the inequality above can be associated with the probability of *ex ante* regret. Then, a lottery is *PC*-preferred to another lottery if its choice results in less *ex ante* regret. The *PC* extension can alternatively be defined using canonical SSB utility functions. Blavatsky (2006) gave a characterization of the *PC* extension which relies on the axioms that characterize SSB utility functions (cf. Fishburn, 1982, 1988) plus an additional axiom that singles out *PC*. In contrast to the previous three extensions, *PC* yields *complete* preference relations over lotteries.

The five lottery extensions introduced here form a hierarchy, i.e., for any preference relation \succsim ,

$$\succsim^{DD'} \subseteq \succsim^{DD} \subseteq \succsim^{BD} \subseteq \succsim^{SD} \subseteq \succsim^{PC}.$$

The examples mentioned also show that these inclusions are strict if $m \geq 3$.

Other extensions that have been considered in the literature include the downward lexicographic (*DL*), the upward lexicographic (*UL*) (Cho, 2016), and the surething (*ST*) (Aziz et al., 2013b) extensions.

Standard axioms such as efficiency, strategyproofness, and participation can now be defined in varying degrees depending on the underlying lottery extension.

Efficiency. Arguably one of the most fundamental axioms in microeconomic theory, Pareto efficiency prescribes that social outcomes should be “optimal” in a well-defined weak way. For a lottery extension \mathcal{E} , p \mathcal{E} -dominates q if $p \succsim_i^{\mathcal{E}} q$ for all $i \in N$ and $p \succ_i^{\mathcal{E}} q$ for some $i \in N$. A PSCF is \mathcal{E} -efficient if it never returns \mathcal{E} -dominated lotteries. A common efficiency notion that cannot be formalized using lottery extensions is *ex post* efficiency. *Ex post* efficiency requires that whenever $x \succsim_i y$ for all $i \in N$ and $x \succ_i y$ for some $i \in N$ (i.e., y is *Pareto dominated* by x) then y should receive probability 0. It can be shown that *SD*-efficiency implies *ex post* efficiency and *ex post* efficiency implies *BD*-efficiency (Aziz et al., 2015).




Figure 1.1: Logical relationships between varying degrees of efficiency, strategyproofness, and participation. *PC*-efficiency, for example, implies *SD*-efficiency. A PSCF is *ex post* efficient if it puts probability 0 on all Pareto dominated alternatives. Strong *SD*-strategyproofness is equivalent to the strategyproofness notion considered by Gibbard (1977). Very strong *SD*-participation requires that a participating agent is always *strictly* better off (unless he already obtains a most preferred outcome).

Strategyproofness. *Strategyproofness* demands that agents cannot benefit from misrepresenting their preferences. Since most lottery extensions return incomplete preference relations, there are two fundamentally different ways how to define strategyproofness. Consider a preference profile, a resulting lottery p , and a lottery extension \mathcal{E} . The strong notion of strategyproofness, first advocated by Gibbard (1977), requires that every misreported preference relation of an agent will result in a lottery q such that $p \succ^{\mathcal{E}} q$. According to the weaker notion, first used by Postlewaite and Schmeidler (1986) and then popularized by Bogomolnaia and Moulin (2001), no agent can misreport his preferences to obtain a lottery q such that $q \succ^{\mathcal{E}} p$. In other words, the strong version always interprets incomparabilities in the worst possible manner (such that they violate strategyproofness) while the weak version interprets them as actual incomparabilities that cannot be resolved. In the following, strategyproofness (without qualifier) will refer to weak strategyproofness. Note that due to the completeness of the *PC* extension, strong *PC*-strategyproofness and *PC*-strategyproofness coincide. Moreover, strong *SD*-strategyproofness is stronger than *PC*-strategyproofness while (weak) *SD*-strategyproofness is weaker. A PSCF is *group-strategyproof* for some lottery extension if no group of agents can jointly misrepresent their preferences such that all of them are strictly better off.

Participation. Like population-consistency, *participation* is a variable-electorate condition. It requires that no agent is ever better off by *abstaining* from an election or—equivalently—that an agent can never be worse off by participating in an election. Again each preference extension yields a corresponding notion of weak and strong participation. On top of that, we define the notion of *very strong participation*, which demands that a participating agent is always *strictly* better off (unless he already obtains a most preferred lottery). While prohibitive in non-probabilistic social choice, this condition is satisfiable by reasonable PSCFs because incentives can be arbitrarily small. In analogy to group-strategyproofness, a PSCF satisfies *group-participation* if no group of agents is individually strictly better off by abstaining from an election.

In principle, every lottery extension leads to corresponding notions of efficiency, weak and strong strategyproofness, and weak, strong, and very strong participation. The relationships between the most relevant concepts are depicted in Figure 1.1. Some combinations such as *DD*-efficiency or strong *BD*-strategyproofness are omitted because they are extremely weak or prohibitively strong.

The sets of efficient lotteries for the various lottery extensions given above already constitute an interesting research subject (see Aziz et al., 2015). For example, it has been shown that whether a lottery is *BD*-efficient or whether it is *SD*-efficient only depends on its support. Perhaps surprisingly, the set of *SD*-efficient lotteries and the set of *PC*-efficient lotteries may fail to be convex. As a consequence, the convex combination of two *SD*-efficient PSCFs may violate *SD*-efficiency. Finding and verifying *BD*-, *SD*-, and *PC*-efficient lotteries can be achieved in polynomial time.

1.4 Results

A complete overview of which properties are satisfied by which PSCF is given in Table 1.1. Interestingly, some combinations of these axioms are prohibitive in *deterministic* social choice while they can be satisfied by reasonable PSCFs. This is, for example, the case for population-consistency and Condorcet-consistency (Young and Levenglick, 1978), participation and Condorcet-consistency (Moulin, 1988), and population-consistency and cloning-consistency (Brandt et al., 2016c). Each of agenda-consistency and very strong participation is prohibitive on its own when paired with minimal further assumptions.

Gibbard (1977) provided a complete characterization of strongly *SD*-strategyproof PSCFs for strict preferences in terms of convex combinations of so-called *unilaterals* (where only one agent affects the outcome) and *duples* (where only two alternatives may receive positive probability). The most well-known consequence of this result is known as the *Random Dictatorship Theorem*.

Theorem 1.1 (Gibbard, 1977). *RD is the only anonymous, strongly SD-strategyproofness, and ex post efficient PSCF when preferences are strict.*

Subsequent research has provided alternative proofs for this theorem (Duggan, 1996; Nandebam, 1997; Tanaka, 2003) as well as various extensions and

	$RSD (RD)$	$Borda_{max}$	$Borda_{pro}$	ML
efficiency	<i>ex post</i> (SD)	SD	–	PC
strategyproofness	strong SD	–	strong SD	$DD' (DD)$
group-strategyproofness	$DD (SD)$	–	$DD (BD)$	$DD' (DD)$
participation	very strong SD	strong SD	strong SD	PC
group-participation	DD	SD	$DD (BD)$	PC
Condorcet-consistency	–	–	–	✓
population-consistency	– (✓)	✓	✓	✓
agenda-consistency	✓	–	–	✓
composition-consistency	–	–	–	✓
cloning-consistency	✓	–	–	✓
pairwiseness	–	✓	✓	✓
computational complexity	#P-complete (in P)	in P	in P	in P
randomness	<i>a lot</i>	<i>little</i>	<i>a lot</i>	<i>some</i>

Table 1.1: Properties of PSCFs. In general, results hold for weak preferences. A property that is only satisfied for strict preferences is given in parentheses. All results are tight in the sense that each cell contains the strongest version of a satisfied property. The cells of PSCFs that satisfy the strongest version of the corresponding property are highlighted in gray. Non-trivial results are due to Gibbard (1977); Barberà (1979b); Aziz et al. (2013a,b); Brandl et al. (2015b, 2016b,c); Brandt (2015).

variations (e.g., Gibbard, 1978; Barberà, 1979a; Hylland, 1980; Barberà et al., 1998; Benoît, 2002; Dutta et al., 2002, 2007; Nandeibam, 2008; McLennan, 2011; Nandeibam, 2013; Picot and Sen, 2012; Chatterji et al., 2014).¹⁶

Since $Borda_{max}$ and ML are *ex post* efficient, Theorem 1.1 entails that these PSCFs violate strong SD -strategyproofness (in fact, ML fails to satisfy BD -strategyproofness while $Borda_{max}$ does not even satisfy DD' -strategyproofness). Another, less obvious, consequence of Gibbard's characterization is that $Borda_{pro}$ satisfies strong SD -strategyproofness (Barberà, 1979b).¹⁷ While the results on strongly SD -strategyproof PSCFs are encouraging, these PSCFs involve an enormous amount of randomization (it follows from Theorem 1.1 that $Borda_{pro}$ even fails to put probability 0 on Pareto dominated alternatives). In general, there appears to be a pervasive tradeoff between efficiency and strategyproofness. For example, it quickly follows from Theorem 1.1 that PC -efficiency and strong SD -strategyproofness are incompatible, even when preferences are strict: since PC -efficiency is stronger than *ex post* efficiency, the only candidate for such a PSCF would be RD , which is easily seen to violate PC -efficiency. A number of impossibilities illustrating this tradeoff (and other incompatibilities) are given in Table 1.2. Among these, the following result deserves special mention.

¹⁶See also Barberà (2010, Sec. 7).

¹⁷This result has been rediscovered several times (see Heckelman, 2003; Procaccia, 2010; Heckelman and Chen, 2013).

axioms			prefs	$m \geq$	$n \geq$	source
<i>SD</i> -eff.	<i>SD</i> -strategypr.	anon. & neutr.	weak	4	4	<i>c</i>
<i>PC</i> -eff.	<i>PC</i> -strategypr.	anon. & neutr.	weak	3	3	<i>a</i>
<i>ex post</i> -eff.	<i>BD</i> -group-strategypr.	anon. & neutr.	weak	3	3	<i>a</i>
<i>ex post</i> -eff.	<i>DD</i> -strategypr.	pairwise	weak	3	3	<i>f</i>
Cond.-cons.	<i>DD</i> -strategypr.	—	weak	3	$3m$	<i>d</i>
Cond.-cons.	strong <i>SD</i> -part.	—	strict	4	12	<i>e</i>
<i>ex post</i> -eff.	very strong <i>SD</i> -part.	pairwise	strict	2	3	<i>b</i>

a: Aziz et al., 2014, *b*: Brandl et al., 2015b, *c*: Brandl et al., 2016a, *d*: Brandt, 2015, *e*: Brandt et al., 2017a, *f*: Unpublished work with F. Brandl, C. Saile, and C. Stricker

Table 1.2: Impossibility theorems. The first row corresponds to Theorem 1.2.

Theorem 1.2 (Brandl et al., 2016a). *There is no anonymous, neutral, SD-efficient, and SD-strategyproof PSCF when $m, n \geq 4$.*

Alternatively, the theorem can be phrased as follows: let f be an anonymous and neutral PSCF which does not return lotteries that are Pareto dominated for all vNM utility representations compatible with the agents' preferences. Then f can be manipulated for all vNM utility representations compatible with the manipulator's preferences. This sweeping impossibility was obtained with the help of a computer and the proof is long and tedious to verify for humans. It has been verified by the interactive theorem prover Isabelle/HOL (see also Chapter 13 of this book). When preferences are strict, the axioms are compatible (and satisfied by *RD*). Theorem 1.2 implies that *RD* cannot be extended to weak preferences without giving up *SD*-efficiency or *SD*-strategyproofness (Brandl et al., 2016d). When restricting attention to pairwise PSCFs, *SD*-efficiency and *SD*-strategyproofness can be weakened to *ex post* efficiency and *BD*-strategyproofness (see Table 1.2).

Perhaps surprisingly, even the lowest degree of strategyproofness (*DD'*-strategyproofness) is violated by many PSCFs. In particular, $Borda_{max}$ (and PSCFs that randomize over plurality winners, Copeland winners, Nanson winners, etc.) violate *DD'*-strategyproofness. However, a handful of interesting PSCFs are *DD'*-strategyproof. A sufficient condition for *DD'*-strategyproofness is *set-monotonicity*, which requires that weakening alternatives that receive probability 0 does not affect the support of the resulting lottery.

Theorem 1.3 (Brandt, 2015, Brandl et al., 2015a). *Every set-monotonic PSCF satisfies *DD'*-group-strategyproofness and *DD*-group-participation (if completely indifferent agents do not affect the outcome). When preferences are strict, set-monotonicity implies *DD*-group-strategyproofness.*

As a consequence, PSCFs that randomize arbitrarily over the choice sets of some well-known set-valued social choice functions such as the *top cycle*, the *minimal covering set*, or the *bipartisan set* (see, e.g., Brandt et al., 2016) are *DD'*-group-strategyproof and satisfy *DD*-group-participation. It is easily seen that a PSCF is *DD'*-group-strategyproof if, for every preference profile, it returns a lottery whose support contains the support of the lottery returned by

another DD' -group-strategyproof PSCF. This implies that, apart from the PSCFs mentioned above, randomizing over elements of the *uncovered set* is DD' -group-strategyproof (even though the corresponding PSCF violates set-monotonicity). ML was shown to satisfy ST -strategyproofness, a minor strengthening of DD' -strategyproofness (Aziz et al., 2013b).

DD -strategyproofness, on the other hand, is already prohibitive when paired with further assumptions such as Condorcet-consistency or pairwise-ness and *ex post* efficiency (see Table 1.2). Other impossibility theorems involving DD -strategyproofness were given by Kelly (1977) and Barberà (1977).

Let us now turn to consistency conditions. As mentioned above, population-consistency and composition-consistency are incompatible in deterministic social choice. When allowing lotteries as outcomes, these axioms uniquely characterize ML .¹⁸

Theorem 1.4 (Brandl et al., 2016c). *ML is the only anonymous PSCF satisfying population-consistency and composition-consistency when preferences are strict.*

RD satisfies population-consistency, but violates composition-consistency. When replacing composition-consistency with the weaker property of cloning-consistency (which is satisfied by RD) and adding Condorcet-consistency (which is violated by RD), the previous characterization remains intact.

Theorem 1.5 (Brandl et al., 2016c). *ML is the only anonymous PSCF satisfying population-consistency, cloning-consistency, and Condorcet-consistency when preferences are strict.*

Since population-consistency has been identified as the defining property of Borda's scoring rule (Young, 1974; Nitzan and Rubinstein, 1981; Saari, 1990), this theorem can be seen as one possible resolution of the well-documented 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 (see, e.g., Black, 1958; Young, 1988, 1995; McLean and Hewitt, 1994). In this sense, Theorem 1.5 resembles the characterization of Kemeny's rule by Young and Levenglick (1978).¹⁹

On top of population-consistency and composition-consistency, ML also satisfies agenda-consistency. Agenda-consistency, the contraction part of which is at the heart of virtually all choice-theoretic Arrovian impossibility theorems (see, e.g., Sen, 1977, 1986), is also satisfied by RD . Pattanaik and Peleg (1986) considered a significantly stronger version of contraction-consistency, which demands that probabilities cannot decrease when removing arbitrary alternatives (by contrast, we require lotteries to be *unaffected* when removing alternatives that receive probability 0). Together with *ex post* efficiency and an independence

¹⁸The formal statement was shown for a framework using fractional profiles which requires PSCFs to be continuous, decisive, and unanimous (see Brandl et al., 2016c). These are mild technical assumptions that are satisfied by every reasonable PSCF.

¹⁹Interestingly, 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 PC SSB utility functions representing the agents' ordinal preferences (Brandl et al., 2016b).

condition, this stronger contraction-consistency condition characterizes RD in a variable-agenda framework. It is violated by $Borda_{max}$, $Borda_{pro}$, and ML .

Recently, ML has also been characterized using a strengthening of PC -group-participation and additional technical properties (Brandl et al., 2016b).

1.5 Discussion and Future Work

Whether randomization is inadmissible, acceptable, or even desirable strongly depends on the application. While electing a political leader via lottery would probably be controversial, randomly selecting an employee of the day, a restaurant to go to, or background music for a party seems quite natural. Important factors in this context are how frequently elections are repeated and how much randomization is entailed by the voting procedure. The degree of randomization of the PSCFs considered in this chapter greatly differs (see Table 1.1). This can, for example, be illustrated by considering the precise circumstances under which these PSCFs return a degenerate lottery. While $Borda_{pro}$ never returns a degenerate lottery (if $m > 2$), RD and RSD do so only if all agents favor the same alternative, and ML only if there is a weak Condorcet winner. Interestingly, there is strong empirical evidence that most real-world preference profiles for *political* elections do admit a Condorcet winner (see, e.g., Feld and Grofman, 1992; Regenwetter et al., 2006; Laslier, 2010; Gehrlein and Lepelley, 2011). Hence, the actual degree of randomization of ML might be relatively low. For a more comprehensive discussion of the acceptability of randomization in social choice, the reader is referred to Brandl et al. (2016c, pp. 1841–1843).

Many topics in probabilistic social choice deserve further study. For example, to the best of my knowledge, there is no formal analysis of the degree of randomization of specific PSCFs. Furthermore, while strong SD -strategyproofness and weak notions of efficiency such as *ex post* efficiency are well understood, this is not the case for the other two extremes. Only little is known about the structure of the set of PC -efficient lotteries (Aziz et al., 2015) and there is no coherent picture of which PSCFs are DD -strategyproof and which ones are not (see, e.g., Brandt, 2015). There are a number of concrete open problems for strict preferences:

- Are there PC -efficient and BD -strategyproof (or even SD -strategyproof or PC -strategyproof) PSCFs?
- Are there Condorcet-consistent and PC -strategyproof PSCFs?
- Are there PC -efficient PSCFs that satisfy very strong PC -participation?

Similarly, there are challenging questions for weak preferences:

- Are there SD -efficient PSCFs that satisfy very strong SD -participation?
- Are there SD -efficient and DD -strategyproof (or even BD -strategyproof) PSCFs?
- Is neutrality required for the first three impossibilities in Table 1.2?

Just like in non-probabilistic social choice, considering restricted domains of preferences such as dichotomous or single-peaked preferences opens up new avenues for intriguing results (see Ehlers et al., 2002; Bogomolnaia et al., 2005). Also, economic domains such as random assignment, random matching, or random coalition formation may allow for positive results as well as strengthened impossibilities (e.g., Bogomolnaia and Moulin, 2001, 2004; Aziz et al., 2013a,c, 2017; Brandt et al., 2017b; Brandl et al., 2017)

Finally, one of the main appeals of *RD* is its association with a natural voting procedure that implements the *RD* outcome. Apart from its simplicity, this procedure has the advantage of minimal preference elicitation. It would be interesting to study similarly natural procedures or cryptographic protocols that implement other PSCFs. Such procedures and protocols are particularly important in probabilistic social choice because agents not only need to be convinced that the outcome was computed correctly, but also that the randomization was performed faithfully.

Acknowledgments

Thanks to Haris Aziz, Florian Brandl, Markus Brill, Christian Geist, Johannes Hofbauer, and Hans Georg Seedig for many illuminating discussions. I am furthermore grateful to Ulle Endriss and Dominik Peters for helpful feedback on drafts of this chapter. This work is supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/11-1.

Bibliography

- A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 46–59. AAAI Press, 2016.
- E. Anshelevich, O. Bhardwaj, and J. Postl. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 777–783, 2015.
- H. Aziz. Maximal Recursive Rule: A New Social Decision Scheme. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 34–40. AAAI Press, 2013.
- H. Aziz. Condorcet’s paradox and the median voter theorem for randomized social choice. *Economic Theory Bulletin*, 35(1), 2015.
- H. Aziz and J. Mestre. Parametrized algorithms for random serial dictatorship. *Mathematical Social Sciences*, 72:1–6, 2014.

- H. Aziz and P. Stursberg. A generalization of probabilistic serial to randomized social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 559–565. AAAI Press, 2014.
- H. Aziz, F. Brandt, and M. Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013a.
- 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, 2013b.
- 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, 2013c.
- 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.
- H. Aziz, F. Brandt, E. Elkind, and P. Skowron. Computational social choice: The first ten years and beyond. In B. Steffen and G. Woeginger, editors, *Computer Science Today*, volume 10000 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, 2017. Forthcoming.
- S. Barberà. The manipulation of social choice mechanisms that do not leave “too much” to chance. *Econometrica*, 45(7):1573–1588, 1977.
- S. Barberà. A note on group strategy-proof decision schemes. *Econometrica*, 47(3):637–640, 1979a.
- S. Barberà. Majority and positional voting in a probabilistic framework. *Review of Economic Studies*, 46(2):379–389, 1979b.
- S. Barberà. Strategy-proof social choice. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, chapter 25, pages 731–832. Elsevier, 2010.
- S. Barberà, A. Bogomolnaia, and H. van der Stel. Strategy-proof probabilistic rules for expected utility maximizers. *Mathematical Social Sciences*, 35(2):89–103, 1998.
- J.-P. Benoît. Strategic manipulation in voting games when lotteries and ties are permitted. *Journal of Economic Theory*, 102(2):421–436, 2002.

- E. Birrell and R. Pass. Approximately strategy-proof voting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 67–72. AAAI Press, 2011.
- D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- P. R. Blavatsky. Axiomatization of a preference for most probable winner. *Theory and Decision*, 60(1):17–33, 2006.
- A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
- A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.
- A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 18–24. AAAI Press, 2015a.
- F. Brandl, F. Brandt, and J. Hofbauer. Incentives for participation and abstention in probabilistic social choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1411–1419. IFAAMAS, 2015b.
- F. Brandl, F. Brandt, and C. Geist. Proving the incompatibility of efficiency and strategyproofness via SMT solving. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 116–122. AAAI Press, 2016a.
- F. Brandl, F. Brandt, and J. Hofbauer. Welfare maximization entices participation. 2016b. Working paper.
- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016c.
- F. Brandl, F. Brandt, and W. Suksompong. The impossibility of extending random dictatorship to weak preferences. *Economics Letters*, 141:44–47, 2016d.
- F. Brandl, F. Brandt, and J. Hofbauer. Random assignment with optional participation. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 326–334. IFAAMAS, 2017.
- F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4):793–804, 2015.
- F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.

- 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.
- F. Brandt, M. Brill, and P. Harrenstein. Tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2016.
- F. Brandt, C. Geist, and D. Peters. Optimal bounds for the no-show paradox via SAT solving. *Mathematical Social Sciences*, 2017a. Special Issue in Honor of Hervé Moulin. Forthcoming.
- F. Brandt, J. Hofbauer, and M. Suderland. Majority graphs of assignment problems and properties of popular random assignments. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 335–343. IFAAMAS, 2017b.
- 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. A strategy-proofness characterization of majority rule. *Economic Theory*, 22(3):557–568, 2003.
- 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.
- W. J. Cho. Incentive properties for ordinal mechanisms. *Games and Economic Behavior*, 95:168–177, 2016.
- V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 781–788. Morgan Kaufmann, 2003.
- 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.
- C. Duddy. Fair sharing under dichotomous preferences. *Mathematical Social Sciences*, 73:1–5, 2015.
- J. Duggan. A geometric proof of Gibbard’s random dictatorship theorem. *Economic Theory*, 7(2):365–369, 1996.
- B. Dutta, H. Peters, and A. Sen. Strategy-proof probabilistic mechanisms in economies with pure public goods. *Journal of Economic Theory*, 106(2):392–416, 2002.

- B. Dutta, H. Peters, and A. Sen. Strategy-proof cardinal decision schemes. *Social Choice and Welfare*, 28(1):163–179, 2007.
- 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.
- E. Elkind and H. Lipmaa. Hybrid voting protocols and hardness of manipulation. In *Proceedings of the 16th International Symposium on Algorithms and Computation (ISAAC)*, volume 3827 of *Lecture Notes in Computer Science (LNCS)*, pages 206–215. Springer-Verlag, 2005.
- S. L. Feld and B. Grofman. Who's afraid of the big bad cycle? Evidence from 36 elections. *Journal of Theoretical Politics*, 4(2):231–237, 1992.
- 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. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- P. C. Fishburn. Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26(1):31–67, 1982.
- P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984a.
- P. C. Fishburn. Dominance in SSB utility theory. *Journal of Economic Theory*, 34(1):130–148, 1984b.
- P. C. Fishburn. *Nonlinear preference and utility theory*. The Johns Hopkins University Press, 1988.
- 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.

- A. Gibbard. Straightforwardness of game forms with lotteries as outcomes. *Econometrica*, 46(3):595–614, 1978.
- B. Goodwin. *Justice by Lottery*. University of Chicago Press, 2005. 1st edition 1992.
- S. Gross, E. Anshelevich, and L. Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 544–550, 2017.
- A. Guerrero. Against elections: The lottocratic alternative. *Philosophy and Public Affairs*, 42(2):135–178, 2014.
- J. W. Headlam. *Election by Lot at Athens*. Cambridge University Press, 1933.
- J. C. Heckelman. Probabilistic Borda rule voting. *Social Choice and Welfare*, 21: 455–468, 2003.
- J. C. Heckelman and F. H. Chen. Strategy proof scoring rule lotteries for multiple winners. *Journal of Public Economic Theory*, 15(1):103–123, 2013.
- L. N. Hoang. Strategy-proofness of the randomized Condorcet voting system. *Social Choice and Welfare*, 48(3):679–701, 2017.
- A. Hylland. Strategyproofness of voting procedures with lotteries as outcomes and infinite sets of strategies. Mimeo, 1980.
- M. D. Intriligator. A probabilistic model of social choice. *Review of Economic Studies*, 40(4):553–560, 1973.
- J. S. Kelly. Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2):439–446, 1977.
- G. Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gössing, Austria (3–27 July 1962)*, pages 73–79, 1965.
- G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993.
- G. Laffond, J. Lainé, and J.-F. Laslier. Composition-consistent tournament solutions and social choice functions. *Social Choice and Welfare*, 13(1):75–93, 1996.
- G. Laffond, J.-F. Laslier, and M. Le Breton. A theorem on symmetric two-player zero-sum games. *Journal of Economic Theory*, 72(2):426–431, 1997.
- J.-F. Laslier. 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.

- I. McLean and F. Hewitt. *Condorcet: Foundations of Social Choice and Political Theory*. Edward Elgar Publishing, 1994.
- A. McLennan. Manipulation in elections with uncertain preferences. *Journal of Mathematical Economics*, 47(3):370–375, 2011.
- H. Moulin. *The Strategy of Social Choice*. North-Holland, 1983.
- H. Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45:53–64, 1988.
- S. Nandeibam. An alternative proof of Gibbard's random dictatorship result. *Social Choice and Welfare*, 15(4):509–519, 1997.
- S. Nandeibam. A note on the structure of stochastic social choice functions. *Social Choice and Welfare*, 30(3):447–455, 2008.
- S. Nandeibam. The structure of decision schemes with cardinal preferences. *Review of Economic Design*, 17(3):205–238, 2013.
- 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.
- J. Picot and A. Sen. An extreme point characterization of random strategy-proof social choice functions: The two alternative case. *Economics Letters*, 115(1):49–52, 2012.
- A. Postlewaite and D. Schmeidler. Strategic behaviour and a notion of ex ante efficiency in a voting model. *Social Choice and Welfare*, 3(1):37–49, 1986.
- A. D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 836–841. AAAI Press, 2010.
- 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. The Borda dictionary. *Social Choice and Welfare*, 7(4):279–317, 1990.
- 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.
- T. Schwartz. Choice functions, "rationality" conditions, and variations of the weak axiom of revealed preference. *Journal of Economic Theory*, 14:414–427, 1976.

- 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.
- T. C. Service and J. A. Adams. Strategyproof approximations of distance rationalizable voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 569–576. IFAAMAS, 2012.
- 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.
- Y. Tanaka. An alternative proof of Gibbard’s random dictatorship theorem. *Review of Economic Design*, 8:319–328, 2003.
- T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- T. Walsh and L. Xia. Lot-based voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 603–610. IFAAMAS, 2012.
- H. P. Young. An axiomatization of Borda’s rule. *Journal of Economic Theory*, 9(1):43–52, 1974.
- 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. 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.