

Collective Choice Lotteries*

Dealing with Randomization in Economic Design

Felix Brandt[†]

Randomization is playing an ever increasing role in economic design with examples ranging from fair allocation to matching markets to voting. I propose and briefly discuss three interdisciplinary and interrelated research questions that deserve further attention: *(i)* when are collective choice lotteries acceptable, *(ii)* how do agents compare lotteries, and *(iii)* how can randomized rules be implemented.

1 Introduction

Economic design has seen the emergence of a number of attractive randomized rules for allocation, matching, and even voting. In this essay, I would like to raise three fundamental questions in the context of randomization in economic design, which I think have not been sufficiently addressed in the literature.

I will take the perspective of social choice theory and refer to randomized collective choice rules, which map a collection of individual preference relations to a socially most-preferred, representative, or otherwise adequate lottery over the alternatives. However, since alternatives may represent allocations (with or without payments), matchings, coalition structures, or any other type of economic outcome, the following observations should be equally relevant for mechanism design, market design, auction theory, matching markets, random assignment, fair division, and so forth.

When aggregating the preferences of multiple agents into a single 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 basic fairness conditions (referred to as *anonymity* and *neutrality*). However, *ex ante* fairness can easily be restored by returning an even chance lottery over a and b . When allowing for more randomization than is necessary to break ties, classic impossibilities such as *Arrow's theorem*, the *Gibbard-Satterthwaite theorem*, the *no-show paradox*, or the *incompatibility of Condorcet-consistency with population-consistency* can be circumvented under suitable assumptions about the agents' preferences over lotteries (e.g., Gibbard, 1977; Brandl et al., 2016, 2018; Brandl and Brandt, 2018). This is reminiscent of game theory, where the availability of randomized strategies is a crucial requirement of fundamental results like the minimax theorem and the Nash equilibrium existence theorem. In the context of voting, two classes of rules that turned out to be particularly desirable from an axiomatic perspective are *random (serial) dictatorships* and rules that return *maximal lotteries* (whose existence follows from the minimax theorem) (see, e.g., Fishburn, 1984a; Brandl et al., 2016; Aziz et al., 2018). These rules have also

*To appear in J.-F. Laslier, H. Moulin, R. Sanver, W.S. Zwicker (Eds.), *The Future of Economic Design*, Springer Verlag.

[†]Institut für Informatik, Technische Universität München, Germany

been considered in matching and allocation subdomains of the general voting domain where random serial dictatorship is known as *random priority* and maximal lotteries as *mixed popular matchings* or *popular random assignments* (see, e.g., Bogomolnaia and Moulin, 2001; Brandt et al., 2017). Further interesting possibilities emerge in the voting domain when assuming that the agents' preferences adhere to certain structural restrictions (such as single-peakedness or dichotomousness) (e.g., Ehlers et al., 2002; Bogomolnaia et al., 2005).

Curiously, the use of lotteries for the selection of officials 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).¹ Randomization has also turned out to be a valuable tool to achieve fairness in matching markets, in particular when individual preferences may contain ties. Bogomolnaia and Moulin (2004) have identified attractive randomized matching rules for the important case of dichotomous preferences. Randomization is perhaps most common when assigning objects to agents. When objects are *indivisible*, it is impossible to deterministically assign objects such that agents with identical preferences receive the same objects (*equal treatment of equals*). This problem is usually avoided by randomization, i.e., by assigning lotteries over objects to the agents. Randomization is particularly natural in the unit-demand (aka house allocation) case, where each agent receives at most one object, because it is not possible to compensate agents via bundles of objects. Besides random priority, the *probabilistic serial rule* by Bogomolnaia and Moulin (2001) has gathered a lot of interest. Even when objects are *divisible* (such as in *cake cutting* problems), randomization has been exploited to achieve *ex ante* fairness and strategyproofness for piecewise valuation functions (e.g., Chen et al., 2013). In settings that involve the approximation of some global measure, such as social welfare in combinatorial auctions, it is well-known that randomized rules can outperform deterministic ones (see, e.g., Nisan and Ronen, 2001; Dobzinski and Dughmi, 2013; Fischer and Klimm, 2015).²

In summary, randomized collective choice rules have emerged in various areas of economic design. At the same time, a number of pressing interdisciplinary research questions remain unresolved.

- (i) When are collective choice lotteries acceptable?
- (ii) How do agents compare lotteries?
- (iii) How can randomized rules be implemented?

In the remainder of this essay, I will comment on these questions.

When are collective choice lotteries acceptable? Whether randomization is inadmissible, acceptable, or even desirable obviously depends on the application. While electing a political leader via lottery would be controversial and perhaps considered by some a failure of deliberative democracy, 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 collective choice rules are executed and how much randomization is entailed by the rule. The interplay between these properties can be explained by risk aversion on behalf of the agents. For example, most people would probably be more content with randomization for daily collective decisions than for annual ones.

¹For further philosophical discussions of using lotteries to achieve fairness, see, for example, Broome (1991) and Nissan-Rozen (2012).

²For a related discussion of the benefits and challenges of randomization, see Haris Aziz's piece in this volume (Aziz, 2018)

Interestingly, humans appear to have less reservations against randomization in matching and allocation than in voting. In fact, randomized voting rules are rarely used in the real world while randomized matching and allocation rules are widely applied. This is partly due to the difference of the *public good* nature of voting versus the *private good* nature of matching and allocation. In private good settings, extensive randomization is often accepted in order to satisfy fairness conditions such as *envy-freeness* (which have no equivalent in the public good setting). There also seem to be cultural and psychological factors influencing an agent’s stance towards lotteries, which are largely unexplored.

To the best of my knowledge, there is no formal analysis of the *degree of randomization* of collective choice rules. While randomization is likely more acceptable if lotteries are only invoked in rare cases to break otherwise unresolvable ties, the degree of randomization can also be considered for a single lottery. Intuitively, an even chance lottery over two alternatives entails more uncertainty than a lottery with almost all probability on the first alternative and negligible probability on the second alternative. Suitable metrics for the degree of randomization may include Shannon entropy, the distance to the nearest degenerate lottery, the variance of the lottery, or just the size of the support.

How do agents compare lotteries? Central concepts such as Pareto efficiency (no agent can be made better off without making another one worse off) or strategyproofness (no agent can obtain a more preferred outcome by misrepresenting his preferences) can only be meaningfully defined by making reference to the agents’ preferences over lotteries. These preferences are typically defined by assuming the existence of a *von Neumann-Morgenstern (vNM) utility function* which assigns cardinal utility values to alternatives and asserting that a lottery is preferred to another lottery if the former yields more expected utility than the latter.

There are at least two problems with this approach. First, various experimental studies have shown that vNM utility theory is systematically violated by human decision makers. An alternative model that, in my view, is much under-appreciated is Fishburn’s *skew-symmetric bilinear utility (SSB)* theory, a significant generalization of vNM utility theory which assigns a utility value to each *pair* of alternatives and dispenses with the controversial independence and transitivity axioms (see, e.g., Fishburn, 1984b, 1988; Aziz et al., 2015; Brandl et al., 2018). The second problem is that asking agents to submit their vNM utility functions—or, equivalently, their complete preferences over lotteries—is usually impractical. I would even argue that most agents are not even aware of these preferences in the first place. (Even if agents *think* they can competently assign vNM utilities to alternatives, these assignments are prone to be based on arbitrary choices because of missing information and the inability to fully grasp the consequences of these choices.) One approach to bypass this problem is to systematically extend the agents’ preferences over alternatives to (possibly incomplete) preferences over lotteries via so-called *lottery extensions*.³ The most influential preference extension is *first-order stochastic dominance*, which is obtained by quantifying over all consistent vNM utility functions. However, there are also other sensible ways to extend preferences to lotteries. For example, by quantifying over all consistent SSB utility function, one obtains the *bilinear dominance* extension, which is coarser than stochastic dominance and therefore leads to weaker notions of efficiency and strategyproofness (Aziz et al., 2015, 2018). Another very intuitive, but little studied, lottery extension is given by postulating that lottery p is preferred to lottery q if and only if p is more likely to return a better alternative than q (see, e.g., Brandl and Brandt, 2018). This extension corresponds to the canonical SSB utility representation consistent with the ordinal preference over alternatives and has been supported by recent experimental evidence (Butler et al., 2016).

³For other approaches to come up with coherent collective decisions when individual decision-makers violate traditional decision-theoretic assumptions, see Regenwetter et al. (2009) and Danan et al. (2016).

How can randomized rules be implemented? An often neglected problem in collective choice, especially with more sophisticated rules, is that agents need to be convinced that their preferences were taken into account properly and that the outcome was computed correctly. These concerns are of particular importance for *randomized* rules because the randomization itself needs to be implemented in a verifiable way. This issue can for example be tackled using distributed cryptographic protocols that allow agents to jointly generate random numbers and to verify the correctness of the rule's outcome (including the randomization) (see, e.g., Brandt and Sandholm, 2005).

As an alternative to the use of intricate protocols from cryptography, physical, publicly verifiable randomization procedures remain of great interest. This can be based on the simple observation that, even in 2018, most randomization procedures of public interest such as drawing lottery numbers on live television or deciding kick-offs and penalty kick orders in soccer matches are still realized via simple physical devices such as urns, dice, or coins. In a similar vein, random priority is a popular allocation rule because it is associated with a natural allocation procedure: each of the agents is asked for his most preferred of the remaining objects in random order. Apart from its simplicity, the random priority procedure has the advantage of only asking agents to explore those parts of their preferences that are required to compute the outcome. Curiously, this procedure runs in polynomial time even though computing the probabilities resulting from random priority is $\#P$ -complete (Aziz et al., 2013; Saban and Sethuraman, 2015). An interesting question is whether similar procedures (for example, adaptive urn processes) exist for other randomized rules such as the probabilistic serial rule.

Acknowledgements

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1. I am grateful to Florian Brandl, Ashley Piggins, Marcus Pivato, and Bill Zwicker for helpful comments.

References

- H. Aziz. A probabilistic approach to voting, allocation, matching, and coalition formation. In J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker, editors, *The Future of Economic Design*. Springer-Verlag, 2018.
- H. Aziz, F. Brandt, and M. Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013.
- 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. Brandl, F. Brandt, and M. Brill. On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110:1–18, 2018.
- 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 and F. Brandt. Arroviaan aggregation of convex preferences. 2018. Working paper.
- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.
- F. Brandl, F. Brandt, and J. Hofbauer. Welfare maximization entices participation. *Games and Economic Behavior*, 2018. Forthcoming.
- F. Brandt and T. Sandholm. On correctness and privacy in distributed mechanisms. In *Revised selected papers from the 7th AAMAS Workshop on Agent-Mediated Electronic Commerce (AMEC)*, volume 3937 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 212–225, 2005.
- 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, 2017.
- J. Broome. Fairness. *Proceedings of the Aristotelian Society, New Series*, 91:81–101, 1991.
- D. Butler, G. Pogrebna, and P. R. Blavatsky. Predictably intransitive preferences. 2016. Working paper.
- Y. Chen, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.
- E. Danan, T. Gajdos, B. Hill, and J.-M. Tallon. Robust social decisions. *American Economic Review*, 106(9):2407–2425, 2016.
- S. Dobzinski and S. Dughmi. On the power of randomization in algorithmic mechanism design. *SIAM Journal on Computing*, 42(6):2287–2304, 2013.
- O. Dowlen. Sorting out sortition: A perspective on the random selection of political officers. *Political Studies*, 57(2):298–315, 2009.
- 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.
- F. Fischer and M. Klimm. Optimal impartial selection. *SIAM Journal on Computing*, 44(5):1263–1285, 2015.
- P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984a.
- P. C. Fishburn. SSB utility theory: An economic perspective. *Mathematical Social Sciences*, 8(1):63–94, 1984b.
- P. C. Fishburn. *Nonlinear preference and utility theory*. The Johns Hopkins University Press, 1988.
- 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.
- 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.
- N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1):166–196, 2001.
- I. Nisan-Rozen. Doing the best one can: a new justification for the use of lotteries. *Erasmus Journal for Philosophy and Economics*, 5(1):45–72, 2012.
- M. Regenwetter, B. Grofman, A. Popova, W. Messner, C. P. Davis-Stober, and D. R. Cavanaugh. Behavioural social choice: a status report. *Philosophical Transactions of the Royal Society*, 364(1518):833–843, 2009.
- D. Saban and J. Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005–1014, 2015.
- P. Stone. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press, 2011.