

A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation



Haris Aziz

Abstract Randomisation and time-sharing are some of the oldest methods to achieve fairness. I make a case that applying these approaches to social choice settings constitutes a powerful paradigm that deserves an extensive and thorough examination. I discuss challenges and opportunities in applying these approaches to settings including voting, allocation, matching, and coalition formation.

1 Introduction

Suppose two agents have opposite preferences over the two possible social outcomes. What should be a fair resolution for this problem?

If the outcome is required to be deterministic, then it is patently unfair to one of the agents. However, one can regain fairness by at least three different approaches: (1) resort to randomisation so that each of the social outcomes has equal probability, (2) treat the outcomes as divisible and resort to time sharing where each social outcome has half of the time share or, (3) use a uniform frequency distribution if there will be multiple occurrences of the discrete outcomes. Mathematically, all three resolutions towards fairness are equivalent because the outcomes have equal probability, time-share, or frequency. In the rest of the article, when I will describe a probabilistic approach to social choice, I will use it abstractly so as to model approaches (1), (2), and (3).

I argue that although a probabilistic approach has been applied in several social choice settings in both theory and practice (see e.g. Stone 2011), there is potential to revisit fundamental social choice settings such as voting, allocation, matching, and coalition formation with this powerful paradigm. Considering the natural aversion of many people to important decisions being based on the toss of the coin, such an approach may be especially useful for time and budget sharing scenarios. A related chapter in this book is by Brandt (2019).

H. Aziz (✉)

UNSW Sydney and Data61 CSIRO, Sydney, NSW, Australia
e-mail: haris.aziz@data61.csiro.au

2 A Case for Probabilistic Social Choice

I first list some of the compelling reasons why probabilistic social choice is a powerful and useful approach.

- (i) **Modeling time-sharing** Some of the settings and their corresponding results in the literature ignore the possibility of time-sharing. For example, several results in voting suppose that only one alternative is selected. However, the voting could be about deciding the fraction of time different genres of music are played on radio. Similarly most of the results in matching and coalition formation suppose that agents form exclusively one coalition or set of partnerships (Manlove 2013). Many of these results need to be re-examined when we allow the flexibility of time-sharing.
- (ii) **Participatory Budgeting** Voting can also be used to decide on which projects get how much budget (see e.g. Fain et al. 2016; Aziz et al. 2017). The approach is getting traction as grassroots participatory budgeting movements grow in stature (Cabannes 2004). In this context, a probabilistic social choice view is useful because the probability of an alternative can represent the fraction of the budget allocated to it.
- (iii) **Achieving fairness** As explained in the example above, a probabilistic or time-sharing approach to social choice is geared towards achieving fairness. In the example above, no deterministic mechanism can simultaneously be anonymous (treat agents symmetrically) and neutral (treat social outcomes symmetrically). On the other hand, a probabilistic approach easily overcomes this impossibility. Suppose that in the example, the two social outcomes are allocating one item each to the agents where one item is valuable to both and the second item is useless to both. Then the only way to avoid envy is to use a probabilistic approach in a broad sense.
- (iv) **Incentivizing participation** Another reason to take a probabilistic approach is to provide participation incentives (Brandl et al. 2015; Aziz et al. 2017, 2018). For the example described above, at least one of the agents appears to have no strict incentive to participate if the decision is made deterministically. On the other hand, probabilistic rules can be designed that give each voter the ability to at least make an epsilon difference (in expectation) to the outcome.
- (v) **Achieving strategyproofness without resorting to dictatorship** Some of the most striking results in social choice give the message that if one wants agents to have incentives to report truthfully, then one has to resort to dictatorship. In our running example it would mean selecting the preferred social outcome of one pre-specified agent. However, with a probabilistic or time-sharing approach, one can still achieve strategyproofness and also circumvent the prospect of a single agent with over-riding power (Aziz et al. 2014; Chen et al. 2013; Procaccia 2010; Procaccia and Tennenholtz 2013).
- (vi) **Achieving stability** In much of the social choice and cooperative game theory literature, a theme of results involves the lack of outcome which satisfies an appropriate notion of stability. In voting, the most prominent result within

this theme is Condorcet’s theorem says that it can be possible that for any given social outcome, a majority of people prefer another outcome. However, Condorcet’s cycles vanish when the probabilistic ‘maximal lottery’ rule is used (Aziz et al. 2013; Brandl et al. 2016).¹ Similarly, core stable outcomes may not exist for general settings such as hedonic coalition formation in which agents have preferences over coalitions they are members of. However, if we allow for probabilistic outcomes or for time-sharing arrangements, there exist stable outcomes (Aharoni and Fleiner 2003).

- (vii) **Circumventing impossibility results** Social choice is at times notorious for some of the bleak impossibility results in its literature. Several results showing that no apportionment method simultaneously satisfies basic monotonicity axioms (Young 1994). However, these problems disappear if a bit of randomisation is used. Similarly, there are results pointing out that no deterministic voting rule satisfies some basic consistency properties. However, this is not anymore the case if one uses the *maximal lotteries* randomised voting rule (Brandl et al. 2016).²
- (viii) **Better welfare guarantees** When considering cardinal preferences over outcomes, a probabilistic approach may achieve better approximation welfare guarantees while simultaneously achieving other axiomatic properties (Anshelevich et al. 2015; Anshelevich and Postl 2016; Procaccia 2010). In some cases, randomization may allow for better welfare or ex ante Pareto improvement while satisfying stability constraints (Manjunath 2016).

3 Research Challenges and Opportunities

I outlined several advantages of considering a probabilistic approach. At an abstract level, resorting to randomisation means that one can consider the full continuous space of outcomes. This can both be a challenge as well as an opportunity for new and exciting research.

- (i) **Formalizing and exploring a range of solution concepts and axioms** When considering probabilistic or time sharing approaches, there are several ways in which a solution concept for deterministic settings can be extended to probabilistic settings. Take for example pairwise stability for the classic stable matching problem in which we want to pair men and women in a way so that no man and woman not paired with each other want to elope. When considering probabilistic outcomes, there is a hierarchy of stability concepts that are all generalisations of deterministic stability (Aziz and Klaus 2019; Dogan and Yildiz 2016; Kesten and Unver 2015; Manjunath 2016). Understanding the nature and structure of these properties is already a significant research direction. More

¹The argument for the existence of such a lottery invokes von Neumann’s minimax theorem.

²For further discussion on probabilistic approaches to circumvent impossibility results in voting, we refer to the survey by Brandt (2017).

importantly, the potential to generalize important axioms based on stability, Pareto optimality, and fairness in several different ways gives useful creative leeway for institution designers for exploring the tradeoffs and compatibility between different levels of properties.

Similar to the potential of defining a several levels of axiomatic properties, one can explore different levels of properties of mechanisms. A case in point is strategyproofness and participation incentives.

- (ii) **Eliciting, representing, and reasoning about preferences** In most voting or matching settings, agents express preferences over deterministic outcomes. As we move from deterministic to probabilistic settings, there is an interesting challenge to elicit and represent agents' risk attitudes towards different lotteries. One possible approach that involves compact preferences is to extend the agents' preferences over discrete outcomes to preferences over probabilistic outcomes by *lottery extension* methods such as first order stochastic dominance (Brandl 2013; Brandt 2017; Bogomolnaia and Moulin 2001; Aziz et al. 2013; Cho 2016).
- (iii) **Designing time-sharing mechanisms** Since modeling time-sharing is one of the most important motivations of probabilistic social choice, it is important to come up with compelling time-sharing mechanisms. Although voting has been studied for decades, a probabilistic perspective leads to interesting and meaningful new voting rules (see e.g., Aziz and Stursberg 2014). When allowing for fractional outcomes, several well-known mechanisms such as Gale-Shapley Deferred Acceptance or Gale's Top Trading Cycles need to be generalized (Kesten and Unver 2015; Athanassoglou and Sethuraman 2011; Bogomolnaia and Moulin 2004).
- (vi) **Efficiency issues** When trying to achieve fairness via randomization, a straightforward approach is to uniformly randomize over reasonable deterministic outcomes or reasonable deterministic mechanisms. However, such a naive approach such as randomizing over Pareto optimal alternatives can lead to loss of ex ante efficiency. This phenomenon is starkly highlighted by the fact that random serial dictatorship that involves uniform randomization over the class of serial dictatorships can lead to unambiguous loss of welfare (Bogomolnaia and Moulin 2001; Aziz et al. 2013). This issue motivates the need to design interesting new mechanisms that are not victim to such a phenomenon.
- (v) **Computational complexity** Generally speaking, optimising in continuous environments is computationally more tractable than in discrete environments. However, when considering time-sharing outcomes that are implicitly convex combinations of a potentially exponential number of discrete outcomes, computing the time shares can be a computationally arduous task (Aziz et al. 2014). Therefore, when formulating time-sharing mechanisms for different social choice settings, computational complexity rears its head as a potential challenge as well an opportunity for innovative algorithmic research.
- (vi) **Instantiating a lottery** As mentioned earlier, uniformly randomizing over desirable outcomes can result in loss of efficiency or computational intractabil-

ity. Therefore, mechanisms may first use an alternative way to find an *expected* ‘fractional’ outcome—say a weighted matching signifying probabilities for partnerships. If approach (1) is being used, then the expected outcome needs to be instantiated via a concrete lottery. Finding a suitable lottery is trivial in single-winner voting and easy for simple assignment settings³ but can become a challenge for richer settings with more complex constraints (Budish et al. 2013). When instantiating a lottery, an interesting challenge that arises is to instantiate over deterministic outcomes that *also* satisfy some weaker notions of stability, fairness, or other properties (see e.g. Teo et al. 2001; Akbarpour and Nikzad 2017).

To conclude, a probabilistic approach to social choice in particular voting, allocation, matching, and coalition formation leads to several interesting research questions and directions.

Acknowledgements The author is supported by a Julius Career Award and a UNSW Scientia Fellowship. He thanks Gabrielle Demange, Jörg Rothe, Nisarg Shah, Paul Stursberg, Etienne Billette De Villemeur, and Bill Zwicker for useful feedback. He thanks all of his collaborators on this topic in particular Florian Brandl, Felix Brandt, and Bettina Klaus for several insightful discussions. Finally he thanks Hervé Moulin and Bill Zwicker for encouraging him to write the chapter.

References

- Aharoni, R., & Fleiner, T. (2003). On a lemma of scarf. *Journal of Combinatorial Theory Series B*, 87, 72–80.
- Akbarpour, M., & Nikzad, A. (2017). Approximate random allocation mechanisms. Technical report 422777, SSRN.
- Anshelevich, E., & Postl, J. (2016). Randomized social choice functions under metric preferences. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 46–59). AAAI Press.
- Anshelevich, E., Bhardwaj, O., & Postl, J. (2015). Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)* (pp. 777–783).
- Athanassoglou, S., & Sethuraman, J. (2011). House allocation with fractional endowments. *International Journal of Game Theory*, 40(3), 481–513.
- Aziz, H., & Klaus, B. (2019). Random matching under priorities: stability and no envy concepts. *Social Choice and Welfare*.
- Aziz, H., & Stursberg, P. (2014). A generalization of probabilistic serial to randomized social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)* (pp. 559–565). AAAI Press.
- Aziz, H., Brandt, F., & Brill, M. (2013). On the tradeoff between economic efficiency and strategy proofness in randomized social choice. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)* (pp. 455–462). IFAAMAS.
- Aziz, H., Brandt, F., Brill, M., & Mestre, J. (2014). Computational aspects of random serial dictatorship. *ACM SIGecom Exchanges*, 13(2), 26–30.

³Any fractional bipartite matching can be represented as a convex combination of discrete matchings via Birkhoff’s algorithm.

- Aziz, H., Bogomolnaia, A., & Moulin, H. (2017). Fair mixing: the case of dichotomous preferences. Working paper.
- Aziz, H., Luo, P., & Rizkallah, C. (2018). Rank maximal equal contribution: A probabilistic social choice function. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)* (pp. 910–916).
- Bogomolnaia, A., & Moulin, H. (2001). A new solution to the random assignment problem. *Journal of Economic Theory*, *100*(2), 295–328.
- Bogomolnaia, A., & Moulin, H. (2004). Random matching under dichotomous preferences. *Econometrica*, *72*(1), 257–279.
- Brandl, F. (2013). Efficiency and incentives in randomized social choice. Master's thesis, Technische Universität München.
- Brandl, F., Brandt, F., & Hofbauer, J. (2015). Incentives for participation and abstention in probabilistic social choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)* (pp. 1411–1419). IFAAMAS.
- Brandl, F., Brandt, F., & Seedig, H. G. (2016). Consistent probabilistic social choice. *Econometrica*, *84*(5), 1839–1880.
- Brandt, F. (2017). Rolling the dice: Recent results in probabilistic social choice. In U. Endriss (Ed.), *Trends in computational social choice* (chap. 1, pp. 3–26). AI Access.
- Brandt, F. (2019). Collective choice lotteries: Dealing with randomization in economic design. In J.-F. Laslier, H. Moulin, R. Sanver, & W. S. Zwicker (Eds.), *The future of economic design*. Springer-Verlag. Forthcoming.
- Budish, E., Che, Y.-K., Kojima, F., & Milgrom, P. (2013). Designing random allocation mechanisms: Theory and applications. *American Economic Review*, *103*(2), 585–623.
- Cabannes, Y. (2004). Participatory budgeting: A significant contribution to participatory democracy. *Environment and Urbanization*, *16*(1), 27–46.
- Chen, Y., Lai, J. K., Parkes, D. C., & Procaccia, A. D. (2013). Truth, justice, and cake cutting. *Games and Economic Behavior*, *77*(1), 284–297.
- Cho, W. J. (2016). Incentive properties for ordinal mechanisms. *Games and Economic Behavior*, *95*, 168–177.
- Dogan, B., & Yildiz, K. (2016). Efficiency and stability of probabilistic assignments in marriage problems. *Games and Economic Behavior*, *95*, 47–58.
- Fain, B., Goel, A., & Munagala, K. (2016). The core of the participatory budgeting problem. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, Proceedings* (pp. 384–399). 11–14 Dec 2016.
- Kesten, O., & Unver, U. (2015). A theory of school choice lotteries. *Theoretical Economics*, 543–595.
- Manjunath, V. (2016). Fractional matching markets. *Games and Economic Behavior*, *100*, 321–336.
- Manlove, D. F. (2013). *Algorithmics of matching under preferences*. Hackensack: World Scientific Publishing Company.
- Procaccia, A. D. (2010). Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)* (pp. 836–841). AAAI Press.
- Procaccia, A. D., & Tennenholtz, M. (2013). Approximate mechanism design without money. *ACM Transactions on Economics and Computation*, *1*(4), 1–26.
- Stone, P. (2011). *The luck of the draw: The role of lotteries in decision making*. Oxford: Oxford University Press.
- Teo, C.-P., Sethuraman, J., & Tan, W.-P. (2001). Gale-shapley stable marriage problem revisited: Strategic issues and applications. *Management Science*, *49*(9), 1252–1267.
- Young, H. P. (1994). *Equity: In theory and practice*. Princeton: Princeton University Press.