

Reverting to Simplicity in Social Choice



Nisarg Shah

Abstract Social choice theory studies the design of rules for eliciting individual preferences and making collective decisions in various domains. The domain of interest—be it election of a candidate based on citizens' votes, allocation of goods to agents, or matching of agents to other agents—dictates the space of possible decisions as well as the criteria that make certain decisions socially more desirable. This often leads to the design of domain-specific solution concepts. In contrast, certain conceptually simple solutions from microeconomics (e.g., maximization of a social welfare function) seem to naturally apply to a broad range of domains, after suitable modeling of agent preferences, and achieve compelling properties. In this article, I point to recent advances in the understanding of such solution concepts in different domains and advocate for reverting to simplicity in social choice.

The past few decades have seen an accelerating shift from analysis of elegant theoretical models to treatment of important real-world problems, which often bear complexity in the forms of constraints, priorities, or endowments. This has paved the way for the design of complex algorithmic solutions. In microeconomics, complex solutions also stem from normative economics, where the goal is to find *some* solution that satisfies a combination of axioms, and no emphasis is placed on the simplicity or the intuitive appeal of the solution itself. On the one hand, having too few axioms could permit a solution to make arbitrary choices from a wide range of possibilities, and on the other hand, having nonessential axioms could unnecessarily complicate the resulting solution. For instance, with the rise of computational economics, polynomial time computability became a popular desideratum. But for many problems, it has become dispensable due to the availability of fast integer programming solvers.

In this article, I examine the appeal of reverting to conceptually simple solutions, and its practical implications on the future of economic design. I draw on my experience of research on social choice theory, which studies societal decision making based on individual preferences. I begin with the positive case of *fair division*, where the power of conceptually simple solutions is relatively well understood, and then discuss *voting*, where the appeal is less clear.

N. Shah (✉)
University of Toronto, Toronto, ON, Canada
e-mail: nisarg@cs.toronto.edu

First, what do I mean by “conceptually simple solutions”? While there is no clear definition, such solutions are typically easy to describe, intuitively appealing, and importantly, defined for a wide range of domains. In theoretical computer science, greedy algorithms and dynamic programming are recognized to be conceptually simple. In this article, I focus on an approach that stems from welfare economics: maximizing a collective utility function (or simply, welfare maximization). In this approach, one first defines, for each participant, a utility function that maps each possible outcome to a real number, then defines a collective utility function (CUF) that aggregates individual utilities, and finally chooses the outcome maximizing the collective utility. Three popular choices for the CUF are utilitarian (sum of utilities), Nash (product of utilities), and egalitarian (minimum utility).¹

Let me discuss their efficacy in fair division applications, where the goal is to fairly divide a common pool of resources among participants. The study of fair division begins by answering what *fair* means. One of the most compelling notions of fairness is envy-freeness (Foley 1967), which requires that no participant prefer what another participant receives to what she receives. Take the classic setting of rent division, where n rooms and a total rent are to be divided among n roommates. Svensson (1983) showed that envy-free outcomes are guaranteed to exist when individuals have quasi-linear utilities, i.e., if their utility is their value for the assigned room minus the rent paid. When participants can be charged payments for receiving goods, utilitarian CUF is typically preferred. Indeed, for rent division the First Welfare Theorem implies that in any envy-free outcome, the allocation of rooms must maximize utilitarian CUF. There still exist multiple divisions of rent that guarantee envy-freeness; Gal et al. (2016) showed that choosing the rent division that maximizes egalitarian CUF (subject to envy-freeness) provides additional guarantees. It is worthwhile remarking that this solution concept has the added benefit of being polynomial time computable.

Let us now turn our attention to goods allocation *without* money. The lack of money makes interpersonal comparison of utilities less meaningful, which in turn makes Nash CUF compelling because, under mild conditions, it is uniquely independent of individual utility scale (Moulin 2003).² Indeed, consider the cake-cutting setting (Steinhaus 1948), where a divisible heterogeneous good (“cake”) is to be allocated. It is commonly assumed that participants have additive utilities; in fact, cake-cutting is the quintessential fair division setting with additive utilities. It has been shown that cutting the cake by maximizing Nash CUF satisfies most desiderata considered in the literature:

- It is equivalent to a market equilibrium approach (strong competitive equilibrium from equal incomes, or s-CEEI) (Sziklai and Segal-Halevi 2015).
- It satisfies group envy-freeness (Berliant et al. 1992), which generalizes envy-freeness and Pareto optimality (Weller 1985).

¹The leximin method is a refinement, in which, after maximizing the minimum utility, ties are broken in favor of higher second minimum utility, remaining ties are broken in favor of higher third minimum utility, and so on.

²That is, scaling the utilities of an individual does not alter the outcome maximizing Nash CUF.

- It produces an outcome in the core,³ which generalizes a different fairness notion called proportionality, and Pareto optimality.
- It satisfies intuitive properties such as resource monotonicity (dividing more cake cannot be worse for anyone) and population monotonicity (dividing between more participants cannot be better for anyone) (Sziklai and Segal-Halevi 2015).

The intractability of computing this outcome, except in special cases Aziz and Ye (2014), has led researchers to explore finite, bounded, and polynomial time computable solutions.⁴ Such solutions are often conceptually more intricate than simply maximizing Nash CUF; take, for example, the polynomial-time Even-Paz protocol (1984) for proportional cake-cutting or the (surprisingly complex) bounded-time Aziz-Mackenzie protocol (2016) for envy-free cake-cutting. It is unclear if one can instead compute the outcome maximizing Nash CUF up to an accuracy sufficient for real-world problems.

For allocating *indivisible* goods without money, under additive utilities, there is no outstanding method due to the impossibility of achieving strong fairness notions like envy-freeness. Recently, Caragiannis et al. (2016a) showed that maximizing Nash CUF achieves envy-freeness up to one good: no participant would envy another participant if the former got to remove at most one good from the latter's bundle. It remains to be seen whether maximizing Nash CUF leads to (relaxations of) stronger guarantees for this setting as it does for cake-cutting. Nonetheless, we have deployed this approach to our fair division website, <http://www.spliddit.org>, due to the simplicity of the solution concept and the usefulness of envy-freeness up to one good in explaining fairness of the chosen outcome to the participants.

In the literature on fair division with non-additive utilities, one prominent method is the leximin method, which maximizes egalitarian CUF. Kurokawa et al. (2015) studied the leximin method, motivated by the real-world problem of allocating unused classroom space to charter schools. Under mild conditions, they showed that when utility functions satisfy an “optimal utilization” requirement,⁵ the leximin method satisfies proportionality, envy-freeness, and Pareto optimality, along with a strong game-theoretic desideratum called group strategyproofness.

Welfare maximization, while prevalent in fair division applications, is largely overshadowed by the axiomatic approach in voting applications, although it has had its advocates (Harsanyi 1955; Hillinger 2005). One potential reason is that while fair division deals with preferences over an exponential outcome space, and thus imposes restricted utilities for tractability, voting applications typically deal with a small outcome space, and thus ask voters to report rankings over possible outcomes. Unfortunately, in this ordinal model, no conceptually simple, or intricate, solution dominates due to celebrated impossibility results like Arrow's (1951) impossibility.

The situation improves a little if one adheres to cardinal utility theory, which posits that voter preferences have intensities, which can be represented by cardi-

³This is an easy derivation given the equivalence to s-CEEI.

⁴Geometric requirements such as contiguity of allocated pieces of cake have also inspired a significant body of research.

⁵Leontief and dichotomous utilities are examples that satisfy this requirement.

nal utilities. Nonetheless, in most voting applications, such as political elections, it is cognitively difficult for voters to specify a real-valued utility for an alternative. Procaccia and Rosenschein (2006) reconciled this conflict by introducing the *implicit utilitarian voting* framework, where voters still report ranked preferences, but these are treated as proxies for underlying cardinal utilities. To address the lack of full information and derive a unique solution concept, they combined welfare maximization with an elegant solution concept from theoretical computer science: optimization of worst-case approximation, where the worst case is over all possible full information (cardinal utilities) consistent with the reported ordinal information. Boutilier et al. (2015) showed the promise of this framework for selecting a single alternative. Caragiannis et al. (2016b) extended the approach to selecting a subset of alternatives and replacing worst-case approximation by another simple concept from learning theory, *minimax regret*. While they showed that implicit utilitarian rules perform well on real data, the framework currently lacks the strong axiomatic justification that welfare maximization has in fair division. One direction for the future is to study complex voting problems with exponential outcome space, such as participatory budgeting (Cabannes 2004; Benade et al. 2017), where restricted utility forms may again be imposed. Another direction is to bridge the gap between voting and fair division by treating the outcomes of a voting process as public goods. Recent work has shown that some of the fairness notions considered in the fair division literature are well-defined for voting problems, provide non-trivial guarantees, and can be achieved (Conitzer et al. 2017; Fain et al. 2018; Aziz et al. 2017).

Let me now cast a wider net. In this article, I examined the appeal of using conceptually simple solutions in social choice, surveyed their success in fair division applications, and contrasted with the relative lack thereof in voting applications. There are also practical implications of gravitating towards such solution concepts.

Perhaps the most obvious implication is that we need to invest more effort to understand the limits of their capabilities—the domains for which different solution concepts are suitable, and the practical considerations they can incorporate. For instance, it would be interesting to study if the attractiveness of the leximin method or Nash CUF extends beyond the “optimal utilization” domain (Kurokawa et al. 2015) and the additive utility domain (Weller 1985; Varian 1974; Berliant et al. 1992; Sziklai and Segal-Halevi 2015; Caragiannis et al. 2016a), respectively. Kurokawa et al. (2015) show that the leximin method can incorporate arbitrary external constraints on feasible outcomes as long as the outcome space remains convex. Would a similar result hold for Nash or utilitarian CUF? Also, most CUFs can easily incorporate priorities for participants in the form of real-valued weights.⁶ Are there other forms of priorities they can incorporate? Finally, implicit utilitarian voting uses worst-case approximation and minimax regret to deal with partial information. Would these be useful to deal with partial information in fair division?

⁶Utilitarian, Nash, and egalitarian CUFs admit weighted variants that use sum of utilities multiplied by weights, product of utilities to the power weights, and minimum of utilities divided by weights, respectively.

Going one step further, there is also a subtle methodological implication for social choice researchers. Instead of starting from a set of axioms and designing *some* solution concept that satisfies them, one may want to examine fundamental solutions that emerge (often uniquely) from simple concepts such as welfare maximization, worst-case approximation, or minimax regret. Budish (2012) contrasts the prevalence of a similar optimization-based approach in the mechanism design literature to the prevalence of the axiomatic approach in the applied matching literature. Note that axioms can still be useful for justifying the choice of one solution concept over another, and for explaining appropriateness of the chosen outcome to the participants (cf. the article in this volume by Ariel Procaccia). Also, the axioms could be setting-dependent even if the solution concept is more generally defined.⁷ It is also worthwhile remarking that even if a solution concept is simple, the algorithm for computing its outcome may be complicated; see, e.g., the algorithm for maximizing Nash CUF (Aziz and Ye 2014; Caragiannis et al. 2016a) or the algorithms for implicit utilitarian voting (Boutilier et al. 2015; Caragiannis et al. 2016b).

Conceptually simple solutions are easy to convey to participants and intuitively appealing, and therefore have a practical advantage over complex solutions. They have thus been advocated in other areas of computational economics as well. For instance, in algorithmic mechanism design literature, simple auctions have been shown to be approximately optimal for many complex settings (Hartline and Roughgarden 2009; Daskalakis and Pierrakos 2011; Greenwald et al. 2017), and have been advocated for their robustness (Hartline 2013). I believe such solution concepts have potential for significant real-world impact in social choice applications.

References

- Arrow, K. (1951). *Social choice and individual values*. New York: Wiley.
- Aziz, H., & Mackenzie, S. (2016). A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)* (pp. 416–427).
- Aziz, H., & Ye, C. (2014). Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)* (pp. 1–14).
- Aziz, H., Bogomolnaia, A., & Moulin, H. (2017). Fair mixing: the case of dichotomous preferences. [arXiv:1712.02542](https://arxiv.org/abs/1712.02542).
- Benade, G., Nath, S., Procaccia, A. D., & Shah, N. (2017). Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)* (pp. 376–382).
- Berliant, M., Thomson, W., & Dunz, K. (1992). On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3), 201–216.

⁷Conversely, sometimes an axiom is broadly defined, but is achieved in different settings by different solution concepts. For instance, the core is achieved by maximizing the Nash CUF in allocation of public goods Fain et al. (2018); Aziz et al. (2017), by the top trading cycles mechanism in housing markets (Shapley and Scarf 1974), and through stable matching algorithms in two-sided matching markets (Gale and Shapley 1962).

- Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A. D., & Sheffet, O. (2015). Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227, 190–213.
- Budish, E. (2012). Matching versus mechanism design. *ACM SIGecom Exchanges*, 11(2), 4–15.
- Cabannes, Y. (2004). Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization*, 16(1), 27–46.
- Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., & Wang, J. (2016a). The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)* (pp. 305–322).
- Caragiannis, I., Nath, S., Procaccia, A. D., & Shah, N. (2016b). Subset selection via implicit utilitarian voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 151–157).
- Conitzer, V., Freeman, R., & Shah, N. (2017). Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)* (pp. 629–646).
- Daskalakis, C., & Pierrakos, G. (2011). Simple, optimal and efficient auctions. In *Proceedings of the 7th Conference on Web and Internet Economics (WINE)* (pp. 109–121).
- Even, S., & Paz, A. (1984). A note on cake-cutting. *Discrete Applied Mathematics*, 7, 285–296.
- Fain, B., Munagala, K., & Shah, N. (2018). Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. Forthcoming.
- Foley, D. (1967). Resource allocation and the public sector. *Yale Economics Essays*, 7, 45–98.
- Gal, Y., Mash, M., Procaccia, A. D., & Zick, Y. (2016). Which is the fairest (rent division) of them all? Manuscript.
- Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1), 9–15.
- Greenwald, A., Oyakawa, T., & Syrgkanis, V. (2017). Simple vs optimal mechanisms in auctions with convex payments. [arXiv:1702.06062](https://arxiv.org/abs/1702.06062).
- Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4), 309–321.
- Hartline, J. D. (2013). *Mechanism Design and Approximation*. Book Draft. <http://jasonhartline.com/MDnA>.
- Hartline, J. D., & Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Economics and Computation (EC)* (pp. 225–234).
- Hillinger, C. (2005). The case for utilitarian voting. Discussion papers in economics, University of Munich, Department of Economics.
- Kurokawa, D., Procaccia, A. D., & Shah, N. (2015). Leximin allocations in the real world. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)* (pp. 345–362).
- Moulin, H. (2003). *Fair division and collective welfare*. Cambridge: MIT Press.
- Procaccia, A. D., & Rosenschein, J. S. (2006). The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)* (pp. 317–331).
- Shapley, L., & Scarf, H. (1974). On cores and indivisibility. *Journal of mathematical economics*, 1(1), 23–37.
- Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16, 101–104.
- Svensson, L.-G. (1983). Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica*, 51(4), 939–954.
- Szicklai, B., & Segal-Halevi, E. (2015). Resource-monotonicity and population-monotonicity in cake-cutting. [arXiv:1510.05229](https://arxiv.org/abs/1510.05229).
- Varian, H. (1974). Equity, envy and efficiency. *Journal of Economic Theory*, 9, 63–91.
- Weller, D. (1985). Fair division of a measurable space. *Journal of Mathematical Economics*, 14(1), 5–17.