Deliberation for Social Choice

Brandon Fain[*] Ashish Goel[†] Kamesh Munagala[‡]

Abstract

In large scale collective decision making, social choice is a normative study of how one ought to design a protocol for reaching consensus. However, in instances where the underlying decision space is too large or complex for ordinal voting, standard voting methods of social choice may be impractical. In such circumstances, one would like to have a general way of discovering socially preferable outcomes without relying on an understanding of the fundamental structure of the space. We propose deliberation, modeled by many bargaining scenarios among small sets of agents, as a natural solution to this problem. We describe the general method and analyze its outcomes in canonical spaces; our main result is that deliberation finds a 1.207-approximation to the social welfare optimum in median graphs.

Introduction. Suppose you are a city council member in Urbanopolis, a large city offering a complex variety of public services. You are in charge of transportation services and have been tasked to (i) choose a set of 5 new projects out of 10 proposed for the city to initiate before the next fiscal year and (ii) determine the allocation of existing transportation funds to 20 ongoing transportation initiatives and services. Mindful of the public interest, you decide to elicit the best way forward based on the consensus of the citizens of Urbanopolis. This example highlights two problems in social choice theory. In (i), the decision space is combinatorial; there are 252 possible sets of projects, far too many to put to any ordinal vote. Furthermore, projects likely exhibit substitutions or complements, so any assumption of additive utility for cardinal utility elicitation and subsequent welfare optimization seems ill founded. In (ii), the decision space is continuous, at least within some range, and multidimensional. Assuming that utility is some metric over the space seems reasonable enough, but the imposition of any particular model of utility seems arbitrary.

Intuitively, neither of these problems are hard for humans. That is to say, if any small set of interested agents look at the proposals, think for a bit, and discuss with one another, we suspect that agents can determine reasonable socially preferable outcomes. We seek to develop practical mechanisms based on this insight. We propose a mechanism that we term *sequential deliberation*.

Sequential Deliberation. There is a decision space S of feasible outcomes (these may be projects, sets of projects, or continuous allocations) called decision points or alternatives and a set N of agents. We assume each agent has a hidden cardinal utility for each outcome. Sequential deliberation proceeds for T rounds and each round can be viewed as bargaining between the agents with a disagreement outcome.

---

[*]Department of Computer Science, Duke University, 308 Research Drive, Durham, NC 27708. `btfain@cs.duke.edu`. Supported by NSF grants CCF-1408784 and IIS-1447554.

[†]Supported by the Army Research Office Grant No. 116388, the Office of Naval Research Grant No. 11904718, and by the Stanford Cyber Initiative. Author's Address: Management Science and Engineering Department, Stanford University, Stanford CA 94305. Email: `ashishg@stanford.edu`

[‡]Department of Computer Science, Duke University, Durham NC 27708-0129. `kamesh@cs.duke.edu`. Supported by NSF grants CCF-1408784 and IIS-1447554.

1. In each round $t = 1, 2, \ldots, T$:
   (a) A pair (or more generally, a subset) of agents u^t and v^t are chosen at random.
   (b) These agents are presented with a disagreement alternative o^t.
   (c) They are asked to output a consensus alternative, and they are told that if they fail to reach a consensus, the alternative output will be o^t.
   (d) Let a^t denote the alternative that is output in round t. We set $o^{t+1} = a^t$, where we assume o^1 is the bliss point of a random agent.
2. The final social choice is a^T.

Figure 1: A framework for sequential deliberation.

Sequential deliberation as outlined in Figure 1 is easily implementable. Clearly, this framework is an instantiation of a more general framework that leads to alternative design choices. For example, the size of the subset that is chosen at each round can vary. Further, the disagreement outcome o^t can be an arbitrary function of the history up to that point, including simply choosing a random alternative. Finally, the ultimate social choice could be some other function of $\{a^1, a^2, \ldots, a^T\}$, for instance, putting them to vote.

Analytic Model. In order to analyze sequential deliberation, we need a model for preferences as well as a model for bargaining. We assume that the set S of alternatives are vertices of a *median graph*. A median graph has the following property: For each triplet of vertices u, v, w, there is a unique point that is common to the three shortest paths, those between u, v, between v, w, and between u, w. This point is the *median* of u, v, w. Several natural graphs are median graphs, including grid graphs in arbitrary dimensions, trees, and hypercubes [5].

Each agent u has a bliss point $p_u \in S$, and his disutility for an alternative $a \in S$ is simply $d(p_u, a)$, where $d(\cdot)$ is the distance function on the median graph. (Note that this disutility can have an agent-dependent scale factor.) The model for bargaining is simply Nash bargaining. Given a disagreement alternative o, agents u and v will choose that alternative a that maximizes:

$$\text{Nash product } = (d(p_u, o) - d(p_u, a)) \times (d(p_v, o) - d(p_v, a))$$

The Nash product maximizer need not be unique; in the case of ties we postulate that agents select the outcome that is closest to the disagreement outcome. Note that our model is fairly general. First, the bliss points of the agents in N form an arbitrary subset of S. Further, the alternative chosen by bargaining need not correspond to any bliss point, so that agents are exploring the space of alternatives when they deliberate. The social cost of an alternative $a \in S$ is given by $SC(a) = \sum_{u \in N} d(p_u, a)$. Let $a^* \in S$ be the minimizer of social cost, that is, the social optimum. We measure the quality of outcome a as

$$\text{Quality}(a) = \frac{SC(a)}{SC(a^*)}$$

where we use the expected social cost if a is the outcome of a randomized algorithm.

Our Results. In our analytic model (Nash bargaining on a median graph), sequential deliberation (see Figure 1) is superior to one-shot deliberation (a^1 where o^1 is the bliss point of a random agent), which in turn is superior to any mechanism restricted to choosing the bliss point of some agent.

Nash Bargaining. On a median graph, Nash bargaining between agents u and v using disagreement outcome o finds the median of p_u, p_v, o.

Baseline. A baseline algorithm is to choose a bliss point p_u of a random agent $u \in N$. There exist median graphs on which this has quality at least 2. It is easy to check that the quality is always at most 2, so the bound of 2 is tight. A stronger statement holds: Any mechanism restricted to choosing the bliss point of some agent cannot have quality better than 2.

Sequential Deliberation. For sequential Nash bargaining on a median graph, the expected quality of outcome a^T has an upper bound approaching 1.207 as $T \to \infty$.

One-shot Deliberation. If $T = 1$, the expected quality of bargaining on a median graph is upper bounded by 1.316 and this bound is tight.

It is important to note that while we present analytic results for deliberation in specific decision spaces, the process of deliberation is well defined regardless of the underlying decision space and the mediator's understanding of the space. At a high level, this flexibility and generality in practice are key advantages of sequential deliberation.

Related Work. While the real world complexities of the model are beyond the analytic confines of this work, deliberation as an important component of collective decision making and democracy is studied in political science. For examples (by no means exhaustive), see [3, 9]. The combinatorial version of the problem is intimately related to the combinatorial public projects problem on which there is a long line of work [8, 7, 1, 2]. However, these results focus on truthful mechanism design and the winner determination problem, whereas we emphasize preference elicitation and the discovery of preferable outcomes for social choice in a broader setting.

Two-person bargaining has been a classical game theoretic problem since its framing by Nash in 1950 [6] as a two-person game wherein there is a disagreement outcome and two agents must cooperate to reach a decision; failure to cooperate results in the adoption of the disagreement outcome. Nash postulates four axioms that a bargaining solution ought to satisfy: Pareto optimality, symmetry (between the agents), invariance with respect to affine transformations of utility (that is, the outcome should not depend on the scale an agent uses to report his utility), and independence of irrelevant alternatives (informally that the presence of a feasible outcome that agents do not select does not influence their decision). Nash proved that the solution maximizing the Nash product (essentially maximizing the geometric mean of utility, normalized by the disagreement outcome) is the unique solution to the bargaining problem satisfying these axioms. We use this classical result to characterize our deliberation process.

Our paper is inspired by the *triadic consensus* results of Goel and Lee [4]. In that work, the authors focus on small group interactions with the goal of reaching consensus. In their model, three people deliberate at the same time, and they choose a fourth individual to whom they grant their votes. This individual takes these votes and participates in future rounds, until all votes accumulate with one individual, who is the consensus outcome. The analysis proceeds through a median graph, on which the authors show that the quality of the consensus approaches 1. However, the protocol crucially assumes individuals know the positions of other individuals, and requires the space of alternatives to coincide with the space of individuals. We make neither of these assumptions – this makes our protocol more practical, but at the same time, restricts our quality to be bounded away from 1.

References

[1] D. Buchfuhrer, M. Schapira, and Y. Singer. Computation and incentives in combinatorial public projects. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, pages 33–42, New York, NY, USA, 2010. ACM.

[2] S. Dughmi. A truthful randomized mechanism for combinatorial public projects via convex optimization. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, pages 263–272, New York, NY, USA, 2011. ACM.

[3] J. Fishkin and R. Luskin. Experimenting with a democratic ideal: Deliberative polling and public opinion. *Acta Politica*, 40(3):284–298, 2005.

[4] A. Goel and J. Lee. Towards large-scale deliberative decision-making: Small groups and the importance of triads. In *ACM EC*, 2016.

[5] D. E. Knuth. *The Art of Computer Programming: Combinatorial Algorithms, Part 1.* Addison-Wesley Professional, 1st edition, 2011.

[6] J. F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

[7] C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Foundations of Computer Science, 2008. FOCS '08. IEEE 49th Annual IEEE Symposium on*, pages 250–259, Oct 2008.

[8] M. Schapira and Y. Singer. *Inapproximability of Combinatorial Public Projects*, pages 351–361. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[9] D. F. Thompson. Deliberative democratic theory and empirical political science. *Annual Review of Political Science*, 11(1):497–520, 2008.