Egalitarianism of Random Assignment Mechanisms

(Extended Abstract)

Haris Aziz
Data61 and UNSW
Sydney, Australia
haris.aziz@nicta.com.au

Aris Filos-Ratsikas
University of Oxford
Oxford, UK
Aris.Filos-Ratsikas@cs.ox.ac.uk

Jiashu Chen
UNSW
Sydney, Australia
jiashu.chen@student.unsw.edu.au

Simon Mackenzie Nicholas Mattei
Data61 and UNSW
Sydney, Australia
{simon.mackenzie,nicholas.mattei}@nicta.com.au

ABSTRACT

We consider the egalitarian welfare of random assignment mechanisms when agents have unrestricted cardinal utilities over the objects. We define and give bounds on how well different random assignment mechanisms approximate the *optimal egalitarian value (OEV)* and investigate the effect that different well-known properties like ordinality, envy-freeness, and truthfulness have on the achievable egalitarian value. Finally, we conduct detailed experiments analyzing the tradeoffs between efficiency with envy-freeness or truthfulness using two prominent random assignment mechanisms — random serial dictatorship and the probabilistic serial mechanism — for different classes of utility functions and distributions.

Keywords

Algorithms; Economics; Social Choice; Game Theory

1. INTRODUCTION

We explore the tradeoffs between fairness and efficiency for randomized mechanisms for the *assignment problem*. Specifically, we consider settings where n agents express preferences (cardinal or ordinal) over a set of m indivisible objects. The objective is to assign the objects to agents in a fair and mutually beneficial manner [1, 3]. This general setting has a number of important and significant applications including the assignment of tasks to cores in cloud computing, kidneys to patients in organ exchanges, runways to airplanes in transportation, and students to seats in schools. We consider the classic assumption that irrespective of whether agents are asked to report ordinal or cardinal preferences, they have an underlying utility structure, where each agent assigns real or cardinal valuations to the objects.

A well-established criterion for fairness is the Rawlsian concept of maximizing the happiness of the least satisfied agent [4]. Following the spirit of this idea, we quantify the fairness of an allocation in terms of its *egalitarian value*: the minimum ratio of the value of objects assigned to an agent to his total valuation for all the objects. The *optimal egalitarian value (OEV)* for a valuation profile of all agents is the best egalitarian value achievable over all assignments. The optimal egalitarian value is well-grounded for a number of reasons. If each agent has a total utility of one for the set of all objects (a standard assumption in the literature [5]), it is equivalent to the maximum egalitarian welfare. The OEV does not change if agents scale their relative values for the objects, unlike other metrics. Furthermore, if the egalitarian value of each agent is $1/n$, then the well-known proportionality requirement [4] is satisfied.

The egalitarian value is not the only criterion for desirable allocation mechanisms. Allocation mechanisms may have other goals and requirements such as envy-freeness or truthfulness. Crucially, both these properties are incompatible with optimizing the egalitarian value.Thus, it is natural to examine the *tradeoffs* between optimizing the egalitarian value and achieving other desirable properties. Evaluating these tradeoffs also motivates the study of how established mechanisms with other desiderata perform in terms of the egalitarian value. For a given mechanism J, we examine the approximation ratio $guar(J)$, which is the minimum ratio (among all valuation profiles) of the egalitarian value of an allocation returned by the mechanism to the OEV.

We study randomized assignment mechanisms for which achieving ex ante fairness is easier compared to deterministic mechanisms. Thus, to evaluate the performance of the mechanisms, we compare their egalitarian value with the OEV achieved by any randomized allocation. Note that computing the allocation with the optimal egalitarian value is an NP-hard problem when we restrict ourselves to deterministic allocations [4]. On the other hand, when we consider randomized allocations, the optimal egalitarian value can be computed in polynomial time via a linear program.

We give extra consideration to two randomized assignment mechanisms — *random serial dictatorship (RSD)* and *probabilistic serial (PS)*, which are probably the best-known and most-studied mechanisms in the random assignment literature [3].In RSD,[1] a permutation over the agents is se-

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

[1]The original definition of RSD is for n agents and n objects; the definition here is an adaptation for $n < m$.

lected uniformly at random and each agent in the permutation picks the most preferred m/n units of object that are not yet allocated [3]. In PS, each object is considered to have an infinitely divisible probability weight of one. To compute an allocation, agents simultaneously and with the same speed eat the probability weight of their most preferred object which has not been completely consumed. Once an object has been completely eaten by a subset of agents, each of these agents moves on to eat their next most preferred object that has not been completely eaten. The procedure terminates after all the objects have been eaten. The random allocation of an agent by PS is the amount of each object he has eaten [3]. PS satisfies stochastic dominance (SD) envy-freeness (envy-freeness with respect to all cardinal utilities consistent with the ordinal preferences). We also define a mechanism which we refer to as *Optimal Egalitarian and Envy-Free Mechanism (OEEF)*, which maximizes the egalitarian value of an allocation under the constraint that the allocation is envy-free. Allocations under this mechanism can be computed in polynomial time via linear programming since envy-freeness can be captured by linear constraints.

2. THEORETICAL RESULTS

We present novel theoretical and empirical results regarding fairness in randomized mechanisms. Our main theoretical contributions are: (1) For *any* SD envy-free mechanism J: $guar(J) = O(n^{-1})$. (2) For *any* envy-free mechanism J: $guar(J) = \Omega(n^{-1})$ and $guar(J) = O(n^{-1/5})$. (3) For *any* truthful-in-expectation mechanism J: $guar(J) = O(n^{-1/5})$. (4) For *any* ordinal mechanism J: $guar(J) = O(n^{-1})$.

The first three results apply to mechanisms that may be cardinal mechanisms. As a result of our general bounds, we also get asymptotically tight bounds of $\Theta(n^{-1})$ for RSD and PS. As a result of our general bounds for envy-free mechanisms, we obtain bounds for well-known envy-free mechanisms such as *competitive equilibrium with equal incomes (CEEI)* [7] and the *pseudo-market* mechanism [6]. Since a random assignment of indivisible objects can also be interpreted as a fractional assignment of divisible objects, *our results apply as well to fair allocation of divisible objects*. The constructions that provide the upper bounds for the *guar* values can be considered as extreme examples that may not be common in real-life scenarios.

3. EMPIRICAL RESULTS

In order to better understand how the mechanisms may perform in practice, we consider the approximation ratio achieved by RSD and PS. We also examine the effect of imposing the envy-freeness constraint on the overall allocation. To study this, we generate ordinal profiles via a Mallows model for different levels of dispersion ϕ from a common reference ranking of objects, assigning cardinal utilities via the Borda and exponential scoring functions. Sweeping ϕ from 0, where all agents have the same preference, to 1.0, where all preference orders are equally likely (the Impartial Culture), allows us to make statements regarding situations where agent preferences are more or less correlated.

Overall, there is a negligible difference between the minimum and average achievable approximation ratios for PS and RSD under Borda utilities. While PS performs slightly better than RSD when agents have more extreme (exponential) utilities, both mechanisms perform strictly worse when agents' valuations are more similar, as they are under Borda utilities. When we require envy-freeness (as in OEEF) with exponential utilities, as ϕ increases towards 1.0 (i.e. Impartial Culture) the achievable approximation ratio first decreases slightly and then increases. Hence, as agents value more disparate objects highly, satisfying envy-freeness does not impose as stiff a penalty on the achievable approximation ratio. In our experiments, the requirement of envy-freeness as a constraint in itself (as in the OEEF mechanism) does not have a large impact on the OEV. However, since PS returns an SD envy-free (envy-free for all cardinal utilities consistent with the ordinal preferences) allocation, its achievable approximation ratio is strictly less than OEEF.

4. CONCLUSION

We present theoretical and experimental results concerning how well different randomized mechanisms approximate the optimal egalitarian value. It has been well-known that egalitarianism can be incompatible with envy-free or truthfulness. In this paper, we quantified how much egalitarianism is affected by such properties. For details including full proofs and comprehensive coverage of the experiments, please see the our full technical report [2].

Acknowledgments

Data61 is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. Aris Filos-Ratsikas was supported by the ERC Advanced Grant 321171 (ALGAME).

REFERENCES

[1] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. In *Proc. of 13th AAMAS Conference*, pages 1305–1312, 2014.

[2] H. Aziz, J. Chen, A. Filos-Ratsikas, S. Mackenzie, and N. Mattei. Egalitarianism of random assignment mechanisms. Technical Report arXiv:1507.06827, arXiv.org, 2015.

[3] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

[4] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proc. of 13th AAMAS Conference*, pages 1321–1328. IFAAMAS, 2014.

[5] A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Proc. of 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.

[6] A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, 87(2):293–314, 1979.

[7] V. V. Vazirani. Combinatorial algorithms for market equilibria. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 5, pages 103–134. Cambridge University Press, 2007.