Pareto Optimal Allocation under Uncertain Preferences (Extended Abstract)

Haris Aziz
Data61, CSIRO and UNSW
Sydney, Australia
haris.aziz@data61.csiro.au

Ronald de Haan
Technische Universität Wien
Vienna, Austria
dehaan@ac.tuwien.ac.at

Baharak Rastegari
University of Glasgow
Glasgow, UK
baharakr@gmail.com

ABSTRACT

The assignment problem is one of the most well-studied settings in social choice, matching, and discrete allocation. We consider this problem with the additional feature that agents' preferences involve uncertainty. The setting with uncertainty leads to a number of interesting questions including the following ones. How to compute an assignment with the highest probability of being Pareto optimal? What is the complexity of computing the probability that a given assignment is Pareto optimal? Does there exist an assignment that is Pareto optimal with probability one? We consider these problems under two natural uncertainty models: (1) the lottery model in which each agent has an independent probability distribution over linear orders and (2) the joint probability model that involves a joint probability distribution over preference profiles. For both of these models, we present a number of algorithmic and complexity results highlighting the differences and similarities in the complexity of the two models.

Keywords

Pareto optimality; uncertain preferences; house allocation; matching under preferences

1. INTRODUCTION

When preferences of agents are aggregated to identify a desirable social outcome, *Pareto optimality* is a minimal requirement. Pareto optimality stipulates that there should not be another outcome that is at least as good for all agents and better for at least one agent. We take Pareto optimality as a central concern and consider a richer version of the classical *assignment problem* where the twist is that agents may express uncertainty in their preferences. The assignment problem is a fundamental setting in which n agents express preferences over n items and each agent is to be allocated one item. The setting is a classical one in discrete allocation and its axiomatic and computational aspects have been well-studied [2, 3, 6, 8, 10, 15, 21, 22]. Our motivation for studying assignment with uncertain preferences is that agents' preferences may not be completely known due to lack of information or communication.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

Our work is inspired by the recent work of Aziz et al. [7] who examined the stable marriage problem under uncertain preferences. Uncertainty in preferences has already been studied in voting [16]. Similarly, in auction theory, it is standard to examine Bayesian settings in which there is probability distribution over the types of the agents. Although computational aspects of Pareto optimal outcomes have been intensely studied in various settings such as assignment, matching, housing markets, and committee voting [3, 5, 6, 9, 12, 17, 18, 19], there has not been much work on Pareto optimal under uncertain preferences. In the presence of uncertainty, one can relax the goal of computing a Pareto optimal outcome and focus on computing outcomes that have the *highest probability* of being Pareto optimal. We will abbreviate Pareto optimal as PO. If an assignment is Pareto optimal with probability one, we will call it certainly PO.

We consider the following uncertainty models:

- **Lottery Model**: For each agent, we are given a probability distribution over linear preferences.
- **Joint Probability Model**: A probability distribution over linear preference profiles is specified.

Note that both the lottery model and the joint probability model representation can be exponential in the number of agents but if the support of the probability distributions is small, then the representation is compact. Also note that the product of the independent uncertain preferences in the lottery model results in a probability distribution over preference profiles and hence can be represented in the joint probability model. However, the change in representation can result in a blowup. Thus whereas the joint probability model is more general than the lottery model, it is not as compact. In view of this, complexity results for one model do not directly carry over to the other model.

The most natural computational problems that we will consider are as follows.

- PO-Probability: what is the probability that a given assignment is PO?
- AssignmentWithHighestPO-Probability: compute an assignment with the highest probability of being PO.

We also consider simpler problems than PO-Probability:

- IsPO-ProbabilityNon-Zero: for a given assignment, is its probability of being PO non-zero?

Problems	Lottery Model	Joint Probability Model
PO-Probability	#P-complete but in FPT w.r.t. parameter: # uncertain agents	in P
IsPO-ProbabilityNon-Zero IsPO-ProbabilityOne	in P in P	in P in P
ExistsPossiblyPO-Assignment	in P (trivially)	in P (trivially)
ExistsCertainlyPO-Assignment	NP-complete	NP-complete
AssignmentWithHighestPO-Prob	NP-hard	NP-hard

Table 1: Summary of results.

- IsPO-ProbabilityOne: for a given assignment, is its probability of being PO one?

We also consider a problem connected to AssignmentWithHighestPO-Probability: ExistsCertainlyPO-Assignment asks whether there exists an assignment that is PO with probability one. Note that ExistsPossiblyPO-Assignment—the problem of checking whether there exists some PO assignment with non-zero probability—is trivial for all uncertainty models in which the induced *certainly preferred* relation is acyclic. An agent certainly prefers an item to another if the preference is with probability 1. The reason for the triviality is that the certainly preferred relation can be completed in a way so that it is transitive, and then for the completed deterministic preferences there exists at least one PO assignment.

We say that a given uncertainty model is *independent* if any uncertain preference profile L under the model can be written as a product of uncertain preferences L_a for all agents a, where all L_a's are independent [7]. Note that the lottery model is independent but the joint probability model is not.

Results.

We show that for both the lottery model and the joint probability model, ExistsCertainlyPO-Assignment is NP-complete. We also prove that AssignmentWithHighestPO-Probability is NP-hard for both models. In view of these results, we see that as we move from deterministic preferences to uncertain preferences, the complexity of computing Pareto optimal assignments jumps significantly. On the other hand, we show that for a general class of independent uncertainty models, both problems IsPO-ProbabilityNon-Zero and IsPO-ProbabilityOne can be solved in linear time. Whereas PO-Probability is polynomial-time solvable for the joint probability model, we prove that it is #P-complete for the lottery model. This problem becomes polynomial-time solvable for the lottery model if there is a constant number of uncertain agents. Moreover, we show that the problem PO-Probability for the lottery model can be solved in fixed-parameter tractable time when parameterized by the number of uncertain agents.

Our results are summarized in Table 1.

2. CONCLUSIONS

Computing Pareto optimal outcomes is an active line of research in economics and computer science. In this paper, we examined the problem for an assignment setting where the preferences of the agents are uncertain. Our central technical results are computational hardness results. We see that as we move from deterministic preferences to uncertain preferences, the complexity of computing Pareto optimal outcomes jumps significantly. The computational hardness results carry over to more complex models in which there may be more items than agents, agents may have capacities, and items may have copies. For future work, we have started considering other uncertainty models [7]. If we allow for intransitive preferences, even a possibly Pareto optimal assignment may not exist and the problem of checking whether a possible Pareto optimal assignment exists becomes interesting. An orthogonal but equally interesting direction will be to consider other fairness, stability, or efficiency desiderata [4].

Acknowledgments.

Aziz is supported by a Julius Career Award. De Haan is supported by the Austrian Science Fund (FWF), project P26200.

REFERENCES

[1] A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3): 689–701, 1998.

[2] A. Abdulkadiroğlu and T. Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88 (2):233–260, 1999.

[3] D. J. Abraham, K. Cechlárová, D. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In *Proceedings of the 16th International Symposium on Algorithms and Computation (ISAAC)*, volume 3341 of *Lecture Notes in Computer Science (LNCS)*, pages 1163–1175, 2005.

[4] H. Aziz. Computational social choice: Some current and new directions. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4054–4057, 2016.

[5] H. Aziz and B. de Keijzer. Housing markets with indifferences: a tale of two mechanisms. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1249–1255, 2012.

[6] H. Aziz, S. Mackenzie, L. Xia, and C. Ye. Ex post efficiency of random assignments. In *Proceedings of the 14th International Conference on Autonomous*

Agents and Multi-Agent Systems (AAMAS), pages 1639–1640, 2015.

[7] H. Aziz, P. Biro, S. Gaspers, R. de Haan, N. Mattei, and B. Rastegari. Stable matching with uncertain linear preferences. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, 2016.

[8] H. Aziz, P. Biro, J. Lang, J. Lesca, and J. Monnot. Optimal reallocation under additive and ordinal preferences. In *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2016.

[9] H. Aziz, J. Lang, and J. Monnot. Computing Pareto Optimal Committees. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

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

[11] S. J. Brams and D. L. King. Efficient fair division: Help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421, 2005.

[12] A. Erdil and H. Ergin. Two-sided matching with indifferences. June 2015.

[13] G. Flum and M. Grohe. The parameterized complexity of counting problems. *SIAM Journal on Computing*, 33(4):892–922, 2004.

[14] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. Springer-Verlag, 2006.

[15] P. Gärdenfors. Assignment problem based on ordinal preferences. *Management Science*, 20:331–340, 1973.

[16] N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.

[17] P. Krysta, D. Manlove, B. Rastegari, and J. Zhang. Size versus truthfulness in the house allocation problemâĹŮ. In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*, pages 453–470. ACM Press, 2014.

[18] D. Manlove. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company, 2013.

[19] D. Saban and J. Sethuraman. House allocation with indifferences: a generalization and a unified view. In *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*, pages 803–820. ACM Press, 2013.

[20] D. Saban and J. Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005 –1014, 2015.

[21] L.-G. Svensson. Queue allocation of indivisible goods. *Social Choice and Welfare*, 11:323–330, 1994.

[22] L.-G. Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.

[23] L. G. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8 (3):410–421, 1979.