

# Social Welfare in One-Sided Matching Mechanisms

## (Extended Abstract)

George Christodoulou  
University of Liverpool  
Ashton Street, Liverpool, U.K.  
gchristo@liv.ac.uk

Aris Filos-Ratsikas  
University of Oxford  
Parks Road, Oxford, UK  
aris.filos-  
ratsikas@cs.ox.ac.uk

Søren Kristoffer Stil  
Frederiksen  
Aarhus University  
Abogade 34, Aarhus, DK  
sorensf@gmail.com

Paul W. Goldberg  
University of Oxford  
Parks Road, Oxford, UK  
paul.goldberg@cs.ox.ac.uk

Jie Zhang  
University of Oxford  
Parks Road, Oxford, UK  
jie.zhang@cs.ox.ac.uk

Jinshan Zhang  
University of Oxford  
Ashton Street, Liverpool, UK  
jinshan.zhang@liv.ac.uk

### ABSTRACT

We study the Price of Anarchy of mechanisms for the well-known problem of one-sided matching, or house allocation, with respect to the social welfare objective. We consider both ordinal mechanisms, where agents submit preference lists over the items, and cardinal mechanisms, where agents may submit numerical values for the items being allocated. We present a general lower bound of  $\Omega(\sqrt{n})$  on the Price of Anarchy, which applies to *all* mechanisms and we show that a very well-known mechanism, Probabilistic Serial achieves a matching upper bound. We extend our lower bound to the Price of Stability of a large class of mechanisms that satisfy a common proportionality property.

### Keywords

One-sided matching, probabilistic serial, truthfulness, price of anarchy, Nash equilibrium

### 1. INTRODUCTION

*One-sided matching* (also called the house allocation problem) is the fundamental problem of assigning items to agents, such that each agent receives exactly one item, which has numerous applications. In this setting, agents are often asked to provide *ordinal preferences*, i.e. preference lists, or rankings of the items. We assume that underlying these ordinal preferences, agents have numerical values specifying how much they value each item [7]. In game-theoretic terms, these are the agents' von Neumann-Morgenstern utility functions [12] and the associated preferences are often referred to as *cardinal preferences*.

A *mechanism* is a function that maps agents' valuations to matchings. However, agents are rational strategic entities that might not always report their valuations truthfully; they may misreport their values if that results in a better matching (from their own perspective). Assuming the agents report their valuations strategically to maximize

their utilities, it is of interest to study the *Nash equilibria* of the induced game, i.e. strategy profiles from which no agent wishes to unilaterally deviate.

A natural objective for the designer is to choose the matching that maximizes the *social welfare*, i.e. the sum of agents' valuations for the items they are matched with, which is the most prominent measure of aggregate utility in the literature. Given the strategic nature of the agents, we are interested in mechanisms that maximize the social welfare *in the equilibrium*. We use the standard measure of equilibrium inefficiency, the *Price of Anarchy* [8], that compares the maximum social welfare attainable in any matching with the *worst-case* social welfare that can be achieved in an equilibrium.

We evaluate the efficiency of a mechanism with respect to the Price of Anarchy of the induced game. We study both deterministic and randomized mechanisms: in the latter case the output is a probability mixture over matchings, instead of a single matching. We are interested in the class of *cardinal* mechanisms, which use cardinal preferences, and generalize the ordinal mechanisms.

Note that our setting involves no monetary transfers and generally falls under the umbrella of *approximate mechanism design without money* [10]. In general settings without money, one has to fix a canonical representation of the valuations. A common approach in the literature is to consider the *unit-sum* normalization, i.e. each agent has a total value of 1 for all the items.

### Our main contributions

Our main contribution can be summarized by the following theorem.

**THEOREM 1.** *The Price of Anarchy of Probabilistic Serial is  $O(\sqrt{n})$  and the Price of Anarchy of any mechanism is  $\Omega(\sqrt{n})$ . This implies that Probabilistic Serial is (asymptotically) optimal among all mechanisms for the problem.*

The upper bound of the theorem holds with respect to the very general equilibrium notion, the *coarse correlated equilibrium*, as well as the *Bayes-Nash equilibrium*, when agents have incomplete information about the preference orderings of other agents.

We extend the lower bound of Theorem 1 to the *Price of Stability*, a more optimistic measure of efficiency, for all

**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.

Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

mechanisms that satisfy a simple *proportionality* property, which requires that any agent can report some preference list that guarantees her an allocation at least as good as an equal split (in terms of probability) of each item, regardless of the reports of the other agents. This property is satisfied by most mechanisms in literature, including Probabilistic Serial.

Finally, for deterministic mechanisms, we prove that no mechanism can achieve a Price of Anarchy smaller than  $\Omega(n)$ , proving that randomization is essential for non-trivial guarantees to be achievable. The proofs to our results can be found in the full version of the paper [5].

## Discussion

The one-sided matching problem was introduced in [7] and has been studied extensively ever since, with various desirable objectives related to truthfulness, fairness and economic efficiency. The objective of social welfare maximization for one-sided matching problems has been studied before in the literature, for mechanisms that are truthful [6]. The authors in [6] prove that Random Priority, a well-known truthful mechanism, achieves an approximation ratio of  $O(\sqrt{n})$ , which is asymptotically optimal among all *truthful* mechanisms. Our lower bounds are more general, since they apply to *all* mechanisms, not just truthful ones. In particular, our lower bound on the Price of Anarchy of all mechanisms generalizes the corresponding bound for truthful mechanisms in [6].

We establish the tightness of the bounds in Theorem 1, by bounding the inefficiency of perhaps the most popular one-sided matching mechanism, Probabilistic Serial. The mechanism was introduced in [3] and since then, it has been in the center of attention of the matching literature, with related work on characterizations, strategic aspects and hardness of manipulation. Somewhat surprisingly, the Nash equilibria of the mechanism were only recently studied [1].

A somewhat different recent branch of study considers ordinal measures of efficiency instead of social welfare maximization, under the assumption that agents' preferences are only expressed through preference orderings over items [2, 9, 4]. While interesting, these measures of efficiency do not accurately encapsulate the socially desired outcome the way that social welfare does, especially since an underlying cardinal utility structure is part of the setting [3, 7, 12]. Our results actually suggest that in order to achieve the optimal welfare guarantees, one does not even need to elicit this utility structure; agents can only be asked to report preference orderings, which is arguably more appealing.

Finally, we point out that our work is in a sense analogous to the literature that studies the Price of Anarchy in item-bidding auctions for settings without money. Furthermore, the extension of our results to very general solution concepts (coarse correlated equilibria) and settings of incomplete information (Bayes-Nash equilibria) is somehow reminiscent of the *smoothness* framework [11] for games. While our results are not proven using the smoothness condition, our extension technique is similar in spirit.

## Acknowledgements

G. Christodoulou, P.W. Goldberg and Jinshan Zhang were supported by the EPSRC grant EP/K01000X/1. P.W. Goldberg was supported by COST Action IC1205. A. Filos-

Ratsikas and Jie Zhang were supported by the ERC Advanced Grant 321171 (ALGAME). A. Filos-Ratsikas and S.K.S. Frederiksen acknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which part of this work was performed and support from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. The authors would like to thank Piotr Krysta for useful discussion.

## REFERENCES

- [1] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Equilibria under the probabilistic serial rule. *arXiv preprint arXiv:1502.04888*, 2015.
- [2] A. Bhargat, D. Chakrabarty, and S. Khanna. Social Welfare in One-Sided Matching Markets without Money. In *APPROX-RANDOM*, pages 87–98, 2011.
- [3] A. Bogomolnaia and H. Moulin. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100:295–328, 2001.
- [4] D. Chakrabarty and C. Swamy. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 105–120, 2014.
- [5] G. Christodoulou, A. Filos-Ratsikas, S. K. S. Frederiksen, P. W. Goldberg, J. Zhang, and J. Zhang. Social Welfare in One-sided Matching Mechanisms. *arXiv preprint arXiv:1502.03849*, 2015.
- [6] A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.
- [7] A. Hylland and R. Zeckhauser. The Efficient Allocation of Individuals to Positions. *The Journal of Political Economy*, 87(2):293–314, 1979.
- [8] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999.
- [9] P. Krysta, D. Manlove, B. Rastegari, and J. Zhang. Size versus truthfulness in the House Allocation problem. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 453–470. ACM, 2014.
- [10] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 177–186. ACM, 2009.
- [11] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 513–522. ACM, 2009.
- [12] J. Von Neumann and O. Morgenstern. *Theory of games and economic behavior (60th Anniversary Commemorative Edition)*. Princeton university press, 2007.