One-Sided Matching with Dynamic Preferences

(Doctoral Consortium)

Hadi Hosseini
Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada
h5hossei@uwaterloo.ca

ABSTRACT

One-sided matching is concerned with the allocation of indivisible objects to a set of agents, in the absence of monetary transfers. In contrast to many real-life scenarios such as shift scheduling or course assignment, traditional matching mechanisms assume that agents precisely know their preferences. However, preferences over alternatives may change over time due to idiosyncratic reasons or as a function of earlier outcomes. My research focuses on the theoretical investigation of dynamic matching markets, along with the experimental study of various matching mechanisms in such dynamic environments. In particular, I design and analyze truthful sequential mechanisms in settings with dynamic ordinal preferences to further elucidate a principled approach to decision making in dynamic models of resource allocation and fill the current gap between stochastic reasoning models and the game-theoretical elements of the matching theory.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Economics, Algorithms, Theory

Keywords

Mechanism Design; Matching; Dynamic Preferences

1. INTRODUCTION

In recent years, the problem of allocating indivisible resources to self-interested agents has generated an interesting, rich interplay at the interface of computer science, economics, and game theory. Most theoretical frameworks and empirical studies for resource allocation encompass a wide range of desirable economic properties such as efficiency, truthfulness, and fairness. A broad array of real-life allocation problems, in fact, prohibit the use of transferable currencies (such as money): allocating human organs to patients (e.g. kidney exchange), assigning campus housing or college courses to students, allocating precious medical resources to patients, assigning faculty members or legislators to committees, and allocating teaching load among faculty are real-life examples of the myriad types of application domains [10, 11, 12].

The absence of monetary transfers has promoted new types of markets, mainly matching and assignment markets. In the context of mechanism design, researchers have been extensively studying these problems and their game-theoretical properties [2, 5, 11]. Although these markets are quintessential to many realistic scenarios, a large number of economic and game-theoretical properties remain unstudied in environments where agents' preferences are dynamically changing. Similar to dynamic mechanism design for auctions [6, 4, 3], the advent of online platforms and web-based applications (e.g. shift scheduling or reservation systems) has imposed a more dynamic nature on the market design problem. Mainly, decisions must be made when agents' underlying private preferences over objects evolve based on what is perceived and learned whilst interacting with the outside world. These dynamics present a few new challenges when seeking to sustain desirable decision policies in multiagent systems with self-interested agents. While matching with dynamic arrival or departure of agents has recently attracted attention (e.g. [1, 9]), none of these works consider the decision problem and strategic behavior of agents with dynamic preferences over a desired planning horizon.

My research combines algorithmic and computational aspects of mechanism design and game theory with insights from theoretical and empirical economics. In particular, my research focuses on analyzing the game-theoretical and incentive properties of matching problems in the presence of dynamic ordinal preferences. My work significantly contributes to the analysis of dynamic matching markets and advancing the theoretical foundations of sequential matching under dynamic preferences.

2. MATCHING WITH DYNAMIC ORDINAL PREFERENCES

Inspired by the seminal random matching mechanisms for fair assignment of goods and alternatives to selfish agents (e.g. Random Serial Dictatorship [2] and Probabilistic Serial mechanisms [5]), I initiated the study of matching problems where agents' preference dynamically evolve and change over time [7]. These mechanisms explicitly consider the incentive problems associated with the revelation of agents' private information in a sequence of matching decisions.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

I proposed a generic sequential stochastic decision process, where each decision represents a matching from the set of alternatives to the set of agents. In static environments, the Random Serial Dictatorship (RSD) mechanism guarantees strategyproofness, fairness (in terms of equal treatment of equals), and ex post efficiency [2]. The first result of this work shows that, in contrast to static settings, RSD is prone to manipulation under dynamic preferences. These findings show that an agent may strategically misreport her true idiosyncratic preference and sacrifice her immediate random assignment to alter the decision trajectory and the evolution of preferences, and subsequently, benefit in the future. Nevertheless, I showed that a nicely crafted *history-dependent* matching policy, namely RSD with adjusted priorities (ARSD), guarantees the global truthfulness while sustaining the local properties of fairness and ex post efficiency in each round. The key idea behind the proposed strategyproof mechanism, ARSD, is to use the future matching decisions to incentivize truthful behavior in the current period. Moreover, as a result of balancing the agents' priority orderings, this elegant random mechanism provides some guarantee on the maximum ex post envy among the participating agents.

2.1 Restricting Preference Dynamics

Global strategyproofness in dynamic settings is a strong requirement, which requires truthfulness given any transition of agent preferences, and for the space of all utility models consistent with agents' preferences. Therefore, inspired by the mechanism design literature on quasi-linear utilities and single-peaked preferences, I restrict attention to a particular class of utility functions from the literature that allows us to formulate the matching problem as a planning problem and leverage Markov Decision Process (MDP) models. In particular, I considered the game-theoretic aspects of dynamic matching problems, where each agent's best-response strategy is conditioned upon the truthful revelation of all other agents [8], which only requires a weaker concept of *within-period ex post Nash equilibrium* [4, 3]. I first argued that even under these assumptions, no optimal matching policy (deterministic or stochastic) satisfies the truthful notion of *within-period ex post incentive compatibility*. Moreover, I showed that the sequential RSD policy does not satisfy within-period ex post incentive compatibility for more than two agents even when agents are endowed with linear positional scoring utilities. Nonetheless, I was able to overcome this impossibility result by examining some additional mild restrictions on the dynamics of the agents' preferences, for example when agents are myopic or when players have rotational preferences and lose their desires after receiving an object (e.g. book reading club). In addition, I showed that when players are single-minded, that is they are only interested in their top choice in each period, even if manipulation of sequential RSD is possible, the manipulation will always benefit all agents, resulting in a more efficient matching policy.

3. FUTURE DIRECTIONS

There are many compelling open problems in the area of dynamic matching without monetary transfers. To date, the literature on matching and random assignment has been mostly focused on studying the random mechanisms in the context of economics and mechanism design. Yet, the advancement in computational methods of decision theory and the recent interest in game-theoretical approaches have provided a vibrant framework for analyzing matching problems under various premises. One immediate future direction is to study dynamic matching models where agents are capable of learning the preferences and reporting strategies of other participating agents in the market. The design of such decision policies and the analysis of their properties is a promising future direction which broadens the applicability of this work, both from theoretical and experimental point of view.

The incompatibility of efficiency and strategyproofness in static matching [5], and its counterpart in dynamic environments, poses a few intriguing questions. Despite the results on the existence of strategyproof history-dependent mechanisms [7], it is still unclear whether there exist some efficient policies in the policy space that satisfy truthfulness in expectation. Thus, I am investigating the existence of truthful policies wherein given a common-knowledge transition kernel, each agent's best response, looking forward into the future, is to report its truthful preferences. Certainly such efficient policies are not required to satisfy local properties of ex post efficiency and strategyproofness; that is, at some steps the policy may yield a Pareto dominated matching decision while maximizing the expected matching outcomes.

REFERENCES

[1] A. Abdulkadiroğlu and S. Loerscher. Dynamic house allocations. *Working papar*, 2007.

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

[3] S. Athey and I. Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

[4] D. Bergemann and J. Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.

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

[6] R. Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 220–229. ACM, 2008.

[7] H. Hosseini, K. Larson, and R. Cohen. Matching with dynamic ordinal preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015.

[8] H. Hosseini, K. Larson, and R. Cohen. Sequential matching with restricted preference dynamics. 2015. Working paper.

[9] M. Kurino. House allocation with overlapping generations. *American Economic Journal: Microeconomics*, 6(1):258–289, 2014.

[10] A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.

[11] T. Sönmez and M. U. Ünver. Course bidding at business schools. *International Economic Review*, 51(1):99–123, 2010.

[12] M. U. Ünver. Dynamic kidney exchange. *The Review of Economic Studies*, 77(1):372–414, 2010.