Games and Economic Behavior 110 (2018) 295–317

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Matching in the large: An experimental study ☆

Yan Chen [a,b,*], Ming Jiang [c], Onur Kesten [d], Stéphane Robin [e], Min Zhu [f]

[a] *School of Information, University of Michigan, 105 South State Street, Ann Arbor, MI 48109-2112, USA*
[b] *Department of Economics, School of Economics and Management, Tsinghua University, Beijing, 100084, China*
[c] *Antai College of Economics and Management, Shanghai Jiao Tong University, 1954 Huashan Road, Shanghai, 200030, China*
[d] *Tepper School of Business, Carnegie Mellon University, PA 15213, USA*
[e] *Université de Lyon, CNRS, GATE UMR 5824, F-69130 Ecully, France*
[f] *School of Business, Beijing Normal University, No. 19 Xinjiekouwai Street, Beijing, 100875, China*

ARTICLE INFO

Article history:
Received 2 January 2017
Available online xxxx

JEL classification:
C78
C92
D47
D82

Keywords:
Matching
School choice
Experiment
Scale

ABSTRACT

We compare the performance of the Boston Immediate Acceptance (IA) and Gale–Shapley Deferred Acceptance (DA) mechanisms in a laboratory setting where we increase the number of participants per match. In our experiment, we first increase the number of students per match from 4 to 40; when we do so, participant truth-telling increases under DA but decreases under IA, leading to a decrease in efficiency under both mechanisms. Furthermore, we find that DA remains more stable than IA, regardless of scale. We then further increase the number of participants per match to 4,000 through the introduction of robots. When robots report their preferences truthfully, we find that scale has no effect on human best response behavior. By contrast, when we program the robots to draw their strategies from the distribution of empirical human strategies, we find that our increase in scale increases human ex-post best responses under both mechanisms.

1. Introduction

School choice has been a widely-debated education policy within the U.S. and across the world (Hoxby, 2003; He, 2014), affecting the education experiences and labor market outcomes for millions of students each year. In practical implementations of school choice plans, game theory has played an instrumental role. Within these streams of research, two mechanisms that have received significant scholarly attention are the Gale–Shapley Deferred Acceptance (DA) mechanism

☆ We thank Doug Bernheim, Fuhito Kojima, Nancy Kotzian, Erin Krupka, Steve Leider, Sera Linardi, Yusufcan Masatlioglu, Alvin Roth, Marie Claire Villeval, and Alistair Wilson for helpful discussions, an Advisory Editor and three anonymous referees for their constructive comments which significantly improved the paper. Comments from audiences at Carnegie Mellon, Guelph, Lyon, Michigan, Shanghai Jiao Tong, Stanford, Tsinghua, the 2014 International Economic Science Association Meetings (Honolulu, HI), and the 2014 European Economic Science Association Meethings (Prague) are gratefully acknowledged. We thank Yefeng Chen and Tingting He for excellent research assistance. The financial support from the National Science Foundation through grant no. SES-0962492 to Chen and Kesten, from Shanghai Pujiang Program through grant no. 17PJC059 to Jiang, from the Agence Nationale de la Recherche under contract ANR-13-BSHS1-0010 to Robin, and from the National Natural Science Foundation of China through grant no. 71704013 to Zhu are gratefully acknowledged. This work was performed within the framework of the LABEX CORTEX (ANR-11-LABX-0042) of Université de Lyon, within the program Investissements d'Avenir (ANR-11-IDEX-007) operated by the French National Research Agency (ANR). This research was approved by the University of Michigan IRB.

* Corresponding author.
E-mail addresses: yanchen@umich.edu (Y. Chen), mjiang@sjtu.edu.cn (M. Jiang), okesten@andrew.cmu.edu (O. Kesten), stephane.robin@univ-grenoble-alpes.fr (S. Robin), zhu@bnu.edu.cn (M. Zhu).

https://doi.org/10.1016/j.geb.2018.04.004
0899-8256/

(Gale and Shapley, 1962) and the Boston Immediate Acceptance (IA) mechanism (Abdulkadiroğlu and Sönmez, 2003). Indeed, the question of which mechanism best meets the goals of a school choice plan has been at the center of intensive research as well as ongoing policy discussions (Abdulkadiroğlu and Sönmez, 2003; Ergin and Sönmez, 2006; Abdulkadiroğlu et al., 2011; Pathak and Sönmez, 2013).

Of these two mechanisms, IA is widely used in practice, yet is theoretically concerning due to the incentives it creates for gaming, which in turn jeopardize its welfare and stability properties. Another question that arises in the theoretical literature is how the number of participants impacts the performance of a school choice mechanism. To address these concerns, our goal in this paper is to provide an experimental assessment of the respective performances of DA and IA in large markets by assessing how the performance of the two mechanisms changes with the number of participants, or market size.

Market size has been predicted to play an influential role in a broad class of economic environments. Theoretically, in a sufficiently large market, impossibilities related to incentives, welfare, or stability may be overturned. Likewise, existing possibilities may be sharpened to unique solutions. This becomes particularly important in the design of practical applications of choice mechanisms such as combinatorial auctions, school choice plans, labor market clearinghouses, course allocations, and kidney exchange, where the market size can range from hundreds to millions of participants. For example, school assignments in New York City match nearly 100,000 students per year, while the centralized college admissions systems in China and Turkey match millions of students, respectively.

Motivated by practical concerns, there has been a surge of interest in the theoretical study of large matching markets within the last decade. This research has tended to focus on either investigating the asymptotic properties of finite discrete markets or modeling the market as a continuum mass of agents. Within this literature, a significant number of papers have examined the question of whether the good incentive and stability properties found in DA hold as the market size grows, as conjectured in Roth and Peranson (1999). Indeed, this conjecture is supported by a number of studies that show that a sufficient increase in market size eliminates partner incentives for preference misrepresentations in marriage problems (Immorlica and Mahdian, 2005), college incentives for capacity and preference misrepresentations in college admissions problems (Kojima and Pathak, 2009), and school incentives to disrespect quality improvements in school choice problems (Hatfield et al., 2012). Similarly, another line of research demonstrates that, in a large market satisfying certain regularity assumptions, DA always produces a stable matching in a discrete two-sided matching model. This result is found in the context of the entry-level labor market for U.S. doctors (Kojima et al., 2013; Ashlagi et al., 2011), as well as in continuum models of many-to-one and many-to-many matching (Azevedo and Leshno, 2011; Azevedo and Hatfield, 2013; Che et al., 2013).

Theoretical insights as well as computation and controlled laboratory experiments have all been used in designing real-life school choice plans (Roth, 2002). For example, matching theorists played an integral role in advocating for the adoption of DA in the New York City (Abdulkadiroğlu et al., 2005a) and Boston (Abdulkadiroğlu et al., 2005b) school systems. Experimental data helped make the case for DA in Boston's decision to switch from IA in 2005 (Abdulkadiroğlu et al., 2005b). Finally, laboratory results underscored the vulnerability of IA to strategic manipulation (Chen and Sönmez, 2006). In both New York City and Boston, this combination of support has been used to make the case for DA. In addition to providing data to support institutional choice, an experimental study of how DA and IA perform in a large market can shed light on how large is "large" in understanding how the mechanisms perform in a large market setting.[1]

Despite the theoretical interest in understanding the effects of market size, this attention has not been transferred to the experimental setting. Laboratory experiments continue to be small in scale compared to the corresponding real-world implementations. For example, the largest school choice experiments have 36 subjects per match (Chen and Sönmez, 2006; Calsamiglia et al., 2010), a far cry from the hundreds of thousands of students in Beijing and Shanghai assigned to various high schools each year (He, 2014).

In our study, we bridge the gap between laboratory experiments and real-world implementations by comparing the performance of two school choice mechanisms in the laboratory setting where market size increases. Studying large matching markets in the laboratory is of interest for two main reasons. First, large markets often have different theoretical properties than their smaller counterparts. Second, large markets make the decision process more complex and thus may impact participant behavior. In our experiment, we compare and contrast the large market characteristics of DA to IA. Unlike DA, which is known to be strategy-proof regardless of market size, IA has been shown to be prone to strategic play in both small and large markets (Kojima and Pathak, 2009; Azevedo and Budish, 2013). In our experimental design, we use a one-sided choice setting. Thus, our experiment can be seen as a complement to existing theoretical studies that focus on the impact of market size in a two-sided matching setting.

Creating a large market in the laboratory is challenging because of both physical and financial constraints. To address these challenges, we conduct two distinct experiments. In our first experiment, we conduct a series of all-human sessions varying the number of participants from 4 to 40 students per match. Note that random re-matching in our 40 student setting requires 80-subject sessions. In our second experiment, we conduct human–robot sessions where the robot strategies are

[1] A notable exception is Rustichini et al. (1994) who show in the context of double auctions six traders of each type are sufficient to obtain efficiency to within one percent. In the context of assignment problems, for a fixed set of object types and a given agent's utility function, Kojima and Manea (2010) calculate the number of copies of each object type needed in order for the probabilistic serial mechanism to become strategy-proof.