Games and Economic Behavior 100 (2016) 353–375

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Boston versus deferred acceptance in an interim setting: An experimental investigation [☆]

Clayton R. Featherstone [a,*], Muriel Niederle [b,c]

[a] University of Pennsylvania, The Wharton School, 3620 Locust Walk – SHDH 1352, Philadelphia, PA 19104, United States
[b] Stanford University, Department of Economics, United States
[c] NBER, United States

ARTICLE INFO

Article history:
Received 16 July 2013
Available online 14 October 2016

JEL classification:
C78
C90
D47
D82

Keywords:
School choice mechanisms
Market design
Experimental economics
Matching theory

ABSTRACT

While much of the school choice literature advocates strategyproofness, recent research has aimed to improve efficiency using mechanisms that rely on non-truthtelling equilibria. We address two issues that arise from this approach. We first show that even in simple environments with ample feedback and repetition, agents fail to reach non-truthtelling equilibria. We offer another way forward: implementing truth-telling as an ordinal Bayes–Nash equilibrium rather than as a dominant strategy equilibrium. We show that this weaker solution concept can allow for more efficient mechanisms in theory and provide experimental evidence that this is also the case in practice. In fact, truth-telling rates are basically the same whether truthtelling is implemented as an ordinal Bayes–Nash equilibrium or a dominant strategy equilibrium. This provides a proof-of-concept that ordinal Bayes–Nash design might provide a middle path, achieving efficiency gains over strategy-proof mechanisms without relying on real-life agents playing a non-truth-telling equilibrium.

1. Introduction

Former Boston Public Schools superintendent Thomas Payzant described the rationale for switching away from a manipulable choice mechanism quite succinctly: "A strategy-proof algorithm 'levels the playing field' by diminishing the harm done to parents who do not strategize or do not strategize well."[1,2] This quote highlights two main assumptions implicit in much of the school choice literature. The first is that a significant fraction of parents will strategize poorly if faced with a situation where truth-telling is sub-optimal. The second is that if truth-telling is not a dominant strategy, then parents might fail to submit preferences truthfully. In this paper, we present an experiment that investigates both of these assump-

[☆] This paper has also circulated under the titles "Improving on Strategy-proof School Choice Mechanisms: An Experimental Investigation", "School Choice Mechanisms under Incomplete Information: An Experimental Investigation", and "Ex Ante Efficiency in School Choice Mechanisms: An Experimental Investigation".

* Corresponding author.
E-mail address: cfeather@gmail.com (C.R. Featherstone).

[1] A school choice mechanism maps students' reported preferences over schools to lotteries over possible allocations of students to schools. A strategy-proof school choice mechanism admits truthful preference revelation as a weakly dominant strategy. A non-strategy-proof mechanism is also called manipulable.

[2] For more about Boston Public Schools' decision to adopt a deferred acceptance assignment system, see Abdulkadiroğlu et al. (2005b) and Abdulkadiroğlu et al. (2006).

http://dx.doi.org/10.1016/j.geb.2016.10.005
0899-8256/

tions in the context of two common school choice mechanisms: the strategy-proof Deferred Acceptance mechanism ("DA") and the non-strategy-proof, "immediate acceptance" mechanism once used in Boston ("Boston").[3]

The Boston mechanism is quite popular, presumably because it maximizes the number of students that receive their reported first choice.[4] However, it also provides participants with incentives to manipulate their preference reports, and in fact, a large number of experiments have confirmed that students do indeed deviate from truth-telling (Chen and Sönmez, 2006; Pais and Pintér, 2008; Calsamiglia et al., 2010).[5] These papers mainly focus on determining what rules-of-thumb participants use when choosing their strategically manipulated preference reports. While Boston Public Schools has replaced the Boston mechanism with strategy-proof DA, several recent papers (Abdulkadiroğlu et al., 2011; Miralles, 2009; Abdulkadiroğlu et al., 2015) have demonstrated that non-truth-telling equilibria under the old Boston mechanism can yield significant efficiency gains. For example, Abdulkadiroğlu et al. (2011) shows in a stylized environment that the Boston mechanism can sustain a generically non-truth-telling equilibrium that dominates the truth-telling DA equilibrium by allowing students to better express their cardinal preferences.[6] This raises a new question that current experiments fail to address, namely whether students would play close enough to the equilibrium prediction in a Boston mechanism for these gains to be realized.[7] Existing experiments cannot address this question, as they generally did not give participants sufficient information to calculate the equilibrium.[8]

The first part of our experiment tries to close this gap by considering a very simple incomplete information environment that has a unique, non-truth-telling equilibrium under the Boston mechanism.[9] As theory predicts, students truth-tell under DA. Under Boston, we confirm former results that students do not play truth-telling strategies. However, even in our simple environment with feedback and repetition, students fail to play the theoretical equilibrium and in general are far from playing a best response. This lines up well with the general finding of the experimental literature on Bayesian games: that subjects have a hard time finding the equilibrium.[10] As such, we confirm the worry of Boston Public Schools superintendent Thomas Payzant that equilibria that entail untruthful reporting may be difficult to successfully implement in actual school choice problems.

Payzant's second assumption is that strategy-proofness may be necessary to "level the playing field". However, dominant-strategy incentive compatibility, while appealing, imposes a strong constraint on design, and it has been shown empirically that this constraint can come with costs.[11] If we could successfully implement truth-telling in a weaker way, then we would open the door to the possibility of leveling the playing field with non-strategy-proof designs that are customized to induce truth-telling in particular environments.

In the second part of the experiment, we try to show that this could work by considering an environment in which truth-telling is an ordinal Bayes–Nash equilibrium under the non-strategy-proof Boston (and of course, a dominant-strategy equilibrium under DA).[12] In this environment, we show that participants truth-tell when it is an equilibrium to do so, whether it is implemented as an ordinal Bayes–Nash equilibrium or a dominant strategy equilibrium.[13] Indeed, truth-telling rates turn out to be high and not statistically different across the two mechanisms. Furthermore, we empirically observe

---

[3] DA is currently used in New York City (Abdulkadiroğlu et al., 2005a, 2009) and Boston (Abdulkadiroğlu et al., 2005b, 2006); Boston, or a close variant, is or has been used in Charlotte-Mecklenburg (North Carolina), Lee County (Florida), Minneapolis, and Seattle (Abdulkadiroğlu and Sönmez, 2003). These are, of course, non-exhaustive lists.

[4] Mechanisms like Boston are also popular in two-sided markets, see e.g. Roth (1991). While the Boston mechanism maximizes the number of students that receive the school they ranked as their first choice, it does not necessarily maximize the number of students that receive their first or second choice. For an investigation on mechanisms that seek to directly optimize the distribution of ranks received by students, see Featherstone (2016).

[5] See Calsamiglia et al. (2011) and Chen and Sönmez (2011) for comments on the data analysis in Chen and Sönmez (2006).

[6] Troyan (2012) suggests an extension that deals with more complex priority structures.

[7] Abdulkadiroğlu et al. (2006) provide suggestive evidence that some participants in Boston may not have manipulated their preference list in an optimal way.

[8] In Pais and Pintér (2008), students can calculate the equilibrium under the full information treatment, but not under their treatments with less information. The paper does not however address the extent to which students play equilibrium.

[9] In our experiments, students know their own preferences, but only the distribution from which other students' preferences are drawn. In one of the experimental environments, this distribution is degenerate; hence, each student actually knows the true preferences of the other students. Students also know the distribution of lottery tie-breaker draws, but not the realization.

[10] The general finding in the experimental literature on the implementation of Bayes–Nash equilibria that involve misreporting of preferences is that the behavior of participants is not close to equilibrium play. See, for example, Featherstone and Mayefsky (2015), Fragiadakis and Troyan (2015), and the summary of work in auctions in the old experimental handbook (Kagel and Roth, 1995, Chapter 8), the new experimental handbook (Kagel and Roth, 2016, Chapter 10 by Kagel and Levin), and in Kagel and Levin's book on the winner's curse (Kagel and Levin, 2009).

[11] Theoretically, dominant-strategy incentive compatibility imposes an incentive compatibility constraint at every possible strategy profile. This is a larger set of constraints than, say, those required by Bayes–Nash incentive compatibility, which only requires a constraint for each unilateral deviation from the proposed equilibrium. Empirically, Abdulkadiroğlu et al. (2009) show that gains from the improvements suggested by Erdil and Ergin (2008) can be significant, but come at the cost of strategy-proofness. See also Kesten and Ünver (2013).

[12] An ordinal Bayes–Nash equilibrium is a Bayes–Nash equilibrium in which each player's equilibrium strategy yields an outcome distribution that first-order stochastically dominates the outcome distributions of all other potential strategies (when all other players are playing equilibrium strategies). See Ehlers and Massó (2007).

[13] While we believe that only truth-telling ordinal Bayes–Nash equilibria will predict actual behavior, we cannot rule out that even non-truth-telling ordinal Bayes–Nash equilibria will predict behavior. It this were so, the case for considering non-strategy-proof mechanisms would be stronger, as a weakly larger set of outcomes can be implemented as an ordinal Bayes–Nash equilibrium without the constraint that the equilibrium entails truth-telling. We suspect, however, that non-truth-telling ordinal Bayes–Nash equilibria will not predict actual behavior; for instance, Fragiadakis and Troyan (2015) consider an assignment game with a unique non-truth-telling ordinal Bayes–Nash equilibrium in which subjects play the equilibrium only about 1% of the time.