Games and Economic Behavior 98 (2016) 180–196

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Truncation strategies in two-sided matching markets: Theory and experiment

Marco Castillo [a], Ahrash Dianat [b,*]

[a] Department of Economics, Texas A&M University, United States
[b] Division of the Humanities and Social Sciences, California Institute of Technology, United States

ARTICLE INFO

Article history:
Received 26 April 2015
Available online 23 June 2016

JEL classification:
C78
C92
D47

Keywords:
Two-sided matching
Truncation strategies
Experiments

ABSTRACT

We investigate strategic behavior in a centralized matching clearinghouse based on the Gale–Shapley deferred acceptance algorithm. To do so, we conduct a laboratory experiment to test the degree to which agents strategically misrepresent their preferences by submitting a "truncation" of their true preferences. Our experimental design uses a restricted environment in which a particular form of truncation is always a best response. We find that subjects do not truncate their preferences more often when truncation is profitable. They do, however, truncate their preferences less often when truncation is dangerous – that is, when there is a risk of "over-truncating" and remaining unmatched. Our findings suggest that behavioral insights can play an important role in the field of market design.

1. Introduction

Two-sided matching theory has informed the design of institutions in areas as diverse as kidney exchange (Roth et al., 2004), entry-level labor markets (Roth and Peranson, 1999), and school choice (Abdulkadiroğlu and Sönmez, 2003). These institutions often operate as centralized clearinghouses, in which participants submit rank-order lists of their preferences and then a particular algorithm selects the final outcome (i.e., who is paired with whom). In this context, a widely-used matching algorithm is the Gale–Shapley deferred acceptance algorithm.[1,2] In the DA algorithm, the market is divided into "proposers" and "receivers". This algorithm has an important property: if all agents submit their true preferences, then the resulting outcome is stable and is also the most preferred stable outcome for the proposing side of the market.[3] In the DA algorithm, it is well-known that the proposers have a dominant strategy of truth-telling (Dubins and Freedman, 1981). The receivers, on the other hand, might have incentives to misrepresent their preferences to produce a more favorable outcome for themselves (Gale and Sotomayor, 1985).

We investigate whether – and under what conditions – receivers behave strategically in the preference-revelation game induced by the DA algorithm. We focus attention on a particular class of strategic behavior: truncation strategies (i.e., submitting a shortened preference list that otherwise maintains the order of the true preferences). This emphasis arises for

* Corresponding author.
E-mail addresses: marco.castillo@tamu.edu (M. Castillo), adianat@caltech.edu (A. Dianat).

[1] The algorithm was first introduced by Gale and Shapley (1962).

[2] Henceforth, DA algorithm.

[3] A matching is said to be stable if no agent prefers remaining unmatched to her current allocation and no pair of agents mutually prefer each other to their current allocations.

http://dx.doi.org/10.1016/j.geb.2016.06.006
0899-8256/

two reasons. First, existing theoretical work suggests that truncation strategies are intuitively appealing (Coles and Shorrer, 2014; Roth and Rothblum, 1999). Second, when evaluating an agent's profitable misrepresentation opportunities in the DA algorithm, it suffices to restrict attention to truncation strategies. In other words, misrepresenting one's true preferences in a manner other than truncation can do no better than what can be achieved via truncation (Roth and Rothblum, 1999).

To make progress on this question, we focus attention on two-sided markets with an equal number of agents on both sides of the market and in which each agent prefers being matched to remaining single. These assumptions guarantee that all agents are matched in equilibrium (i.e., at all stable matchings). We then characterize the conditions under which a receiver acting on her own can secure a match that is no worse than her most preferred achievable partner.[4] Our proposition is a straightforward extension of a classic result: in markets with more than one stable matching, there will be an incentive for some receiver to truncate her preferences whenever all other agents report their preferences truthfully (Gale and Sotomayor, 1985). Although an agent's optimal truncation is a function of other agents' reported preferences, we show that no direct knowledge of the strategies of other agents is required. Rather, it merely suffices for other receivers to be constrained to truncation strategies to be able to calculate the best response. In general, the optimal truncation strategy will change if other receivers are allowed to play non-truncation strategies since these strategies can substantially alter the set of stable matchings. In those situations, it is possible to be optimally truncating with respect to the true preferences but sub-optimally truncating with respect to the reported preferences. While not breaking new ground, this result is methodologically important for our experimental design. It allows us to construct environments that maintain the key interactive features of matching markets while essentially reducing optimal truncation to a decision-theoretic problem.

Even after removing this aspect of strategic uncertainty, there are two practical difficulties that present themselves with respect to optimal truncation. First, an agent might be unable to identify the existence of a profitable opportunity to misrepresent her preferences. Second, an agent might "over-truncate" her preferences and remain unmatched (her worst possible outcome).[5] In a laboratory experiment, we investigate whether the degree of truncation depends on the ordinal distance between the most preferred and least preferred achievable partners in an agent's preference list (i.e., the "profitability" of truncation) as well as the rank of the most preferred achievable partner in an agent's preference list (i.e., the "riskiness" of truncation). The first measure is important since it represents the magnitude of the pecuniary gain from optimal truncation compared to the benchmark of truth-telling. The second measure is important since the rank of the most preferred achievable partner determines the likelihood of remaining unmatched by mistakenly over-truncating.

There are two points worth emphasizing about the notions of profitability and riskiness that we define. First, since both of these features are functions of the set of achievable partners, they are only well-defined in the context of our experimental design. If agents are allowed to play non-truncation strategies, then the set of achievable partners under the reported preferences can differ substantially from the set of achievable partners under the true preferences. Second, we should only expect behavioral agents to be responsive to these features of the market. A sophisticated agent would only be responsive to the *existence* of a profitable truncation opportunity and not the *magnitude* of the pecuniary gain from playing an optimal truncation strategy. Similarly, a sophisticated agent would have the ability to calculate the set of stable outcomes and could thus implement the optimal truncation strategy without the risk of over-truncation.

To mirror the theoretical conditions, the experiment is conducted in an environment with complete information about other agents' preferences. The proposing side of the market is automated to play its dominant strategy of truthful preference revelation. The experimental subjects play in the role of the receivers and they are restricted to either truth-telling or truncation strategies. Importantly, since we have removed the element of strategic uncertainty over other players' actions, the *only* risk associated with truncation in our environment comes from over-truncating. Our experiment tests the extent to which agents truncate their preferences in situations that are particularly conducive to truncation behavior.[6] Ideally, the simplicity of our environment will provide insight into the reasons why market participants choose to behave either straightforwardly or strategically.

We find truth-telling to be the most common strategy in our experimental markets: 56% (511/920) of submitted rank-order lists are identical to subjects' true preferences. We also find that truncation is not sensitive to considerations of profitability, but is sensitive to the rank of the most preferred achievable partner. This result is robust to alternative specifications. We consider this to be remarkable given the difficulty in identifying achievable match partners even in small markets. Regarding aggregate outcomes, 88% (203/230) of our experimental markets culminate in stable outcomes. This is not only due to the fact that truth-telling is common, but also because over-truncation is rare.[7] We also find that final outcomes are closer to the receiver-optimal stable matching than to the proposer-optimal stable matching. However, this result is not entirely surprising. Since strategic behavior has positive spillover effects in our environment, the receiver-optimal stable outcome can still be attained when only a subset of agents truncates its preferences optimally.

[4] Two individuals are said to be achievable for each other if they are paired at some stable matching.

[5] Over-truncation refers to the situation where an agent truncates "too much" and leaves her most preferred achievable partner off her submitted rank-order list.

[6] Another factor that makes truncation more attractive in our experiment is the linearity in subject payoffs. In most applications, it is reasonable to expect a utility discontinuity between matching with one's least preferred partner and remaining unmatched. If the risk of remaining unmatched is important in our environment, then it is likely to be even more important in field settings where remaining unmatched is likely to be very costly.

[7] Due to the constrained nature of our strategy space, over-truncation is the only way to observe instability in final outcomes.