Matching and Learning – An Experimental Study

Guillaume Haeringer*
Lan Nguyen†
Silvio Ravaioli‡

December 3, 2019

Extended Abstract

We use a lab experiment to study the patterns and effects of learning in centralized matching mechanisms widely used in school choice and other real-world settings. Theoretical results on the properties of these mechanisms rely on strong assumptions about agents' behavior, and there is evidence of their violations in the existing experimental literature. Our contribution is two-fold. First, we provide a robustness test of existing experimental results by adding a previously-unexplored component: learning. Second, our results have a clear policy implication. If a mechanism performs better when participants have accumulated more experience, the market designer should take this into account by creating or encouraging an explicit learning process.

In our laboratory experiment, we adopt a 2x2 design, with two matching mechanisms (between-subjects) and two treatments (within-subjects). The two mechanisms are Gale-Shapley Deferred Acceptance (DA) and Boston Immediate Acceptance (IA) algorithms. In each experimental session, we use only one of the mechanisms, and at the beginning of the session, the rules of the relevant mechanism are clearly explained, accompanied by an example. Each session is divided into two parts for the two different treatments: part one, without priority zones, and part two, with priority zones (where students are informed of their higher chance of being admitted in two out of five schools). Within each part, there are 8 rounds, each consisting of 5 periods. In every period, each participant plays the role of a student who applies to multiple schools sequentially. There are 20 students and 5 schools. Participants are informed of their own payoffs for being admitted into the schools and asked for an order in which they would apply to those schools. The distinction between rounds and periods allows for two types of learning. Within a round, the environment (schools' priorities and students' preferences) is unchanged, and periods differ only in the rankings submitted by participants. Thus, agents learn about the environment and their opponents. Across

*Baruch College. Email: guillaume.haeringer@baruch.cuny.edu. [Link to webpage.]
†Department of Economics, Columbia University. Email: tn2304@columbia.edu. [Link to webpage.]
‡Department of Economics, Columbia University. Email: sr3300@columbia.edu. [Link to webpage.]

the rounds, the environment changes, but the mechanism does not. Agents who adopt a model-based learning process can become more knowledgeable about the mechanism itself and extrapolate the best strategy beyond specific environments.

We have now run a pilot experiment with 4 sessions and 74 subjects. Our preliminary analysis shows that (i) a higher fraction of participants are truthful under the DA mechanism, and (ii) average payoffs are higher under the IA mechanism. This is expected, given the strategy-proofness of DA and efficiency of IA. (iii) For DA, the fraction of participants who satisfy truncated truthfulness, that is, whether each participant's rank-order list is according to her true preference above her own match, increases over time across the different rounds within each session, indicating some learning with respect to the strategy-proofness of this mechanism. We do not observe the same trend for IA, for which truth-telling is not necessarily a good strategy. (iv) Across periods within a round, truncated truthfulness decreases over time for both mechanisms, although payoffs remain relatively stable. This is suggestive evidence that in later periods, participants become more knowledgeable about the environment of the round and, therefore, can guarantee themselves the same payoffs with different strategies that may be less truthful. This across-period pattern is robust to both treatments with and without priority zones. These results confirm our conjectures that learning may play an important role in matching problems, and mechanisms may display persistent deviations from optimal behavior even after extensive experience. We are in the process of testing the robustness of the results under weaker definitions of truthfulness, as well as comparing measures of strategic optimality that are more relevant for IA.