Equilibrium Selection and the Role of Information in Repeated Matching Markets

Ernan Haruvy
University of Texas at Dallas

M. Utku Ünver
University of Pittsburgh

Abstract

We examine a repeated one-to-one matching environment. In such environments, when a partnership forms between a worker and a firm, the previous partners of these parties may find themselves displaced, and may in turn displace other agents when they find new suitable partners. Using an experimental strategic framework, we seek here to characterize the convergence outcome of this process and its sensitivity to environmental factors such as the information structure.

JEL Classification Numbers: C78, C99, D83

Introduction

A significant body of research has applied insights from matching theory to economic settings (see Roth and Sotomayor, 1990). However, little attention has been given to repeated matching environments which are settings characterized by buyers and sellers repeatedly seeking and matching with each other (as in spot labor markets or business-to-business interactions). In this work, we show that the deferred acceptance algorithm (Gale and Shapley, 1962) is a useful predictor for repeated matching markets since it can be thought of as a description of dynamic reactions by naive agents in a decentralized environment (e.g., Roth and Vande Vate, 1990; Roth and Xing, 1997). However, in high information environments, one might expect agents to behave less naively and therefore the algorithm to be less useful as a predictor. In a series of experiments, we change the information environment to examine this issue.

The experimental treatments we study involve four firms and four workers. Each worker has preferences over firms and each firm has preferences over workers. Each firm makes a single offer per period and each worker can accept a single offer in each period. These matching decisions are repeated many times. The experimental conditions vary along two dimensions, information and worker volition, resulting in a 2x2 design.

We find that with low information and non-strategic workers, the firm-optimal outcome predicted by the firm-proposing deferred acceptance algorithm is almost always reached. In the high information and strategic worker conditions, repeated interactions occasionally culminate in worker-optimal stable matching, but the firm-optimal stable matching is nevertheless reached in the vast majority of instances.

Theory

The deferred acceptance algorithm (hereafter DAA) (Gale and Shapley, 1962) is used to find stable outcomes. Workers and firms reveal their preferences over each other, and the algorithm is executed to find a matching. The firm proposing DAA works as follows: Each firm ranks all workers in order of preference. In each step, each firm makes an offer to his highest ranked worker to whom it has not yet extended an offer. Each worker who has received at least one offer conditionally accepts the best offer and rejects all others. A firm whose offer is rejected makes an offer in the next step to the next worker on his list. The algorithm terminates when no offers are rejected in a step. The conditionally accepted offers by workers then become permanently accepted and are realized as matches. The worker-proposing algorithm is symmetrically defined.

A matching is said to be "stable" if there is no worker-firm pair, each of whom prefers the other to his current partner and there is no agent who prefers being unmatched to his current partner. There is a Pareto-dominant stable matching for all firms (workers), and this is the worst stable matching for all workers (firms). The outcome of the DAA is a "stable" matching, which is "optimal" for the proposing side.

This algorithm has many applications (see Roth and Sotomayor, 1990) and is most famous for the matching of residents to hospitals. However, it has historically been used purely for centralized matching based on stated preferences. It has not hitherto been used for prediction in a repeated decentralized environment.

Why should the DAA have any predictive use in repeated decentralized environments? Consider the following repeated game strategy for a firm: In the first period, firm initially makes an offer to his highest ranked worker. In the k^{th} period, if firm did not get matched in the previous period, he makes an offer to his highest ranked worker that he has not yet made an offer to. If firm was matched in the previous period, he makes an offer to the choice that he was matched in $(k-1)^{th}$ period. We name this repeated game strategy the "going-down-the-list" (hereafter GDL) strategy. Consider any preference profile where the DAA converges to the firm-optimal stable outcome in S steps. Let the game be played $K \geq S$ times using the same preference profile.

We first argue that when workers myopically accept best offers, this strategy profile will result in the firm-optimal stable matching:

Lemma 1: In the last K-S+1 periods of a finitely repeated matching game, the outcome of the "GDL" strategy profile coincides with the firm-optimal stable matching when workers are myopic.

The proof is straightforward due to the parallel between the GDL strategy and DAA.

We also define a delayed version of the GDL strategy: In the first period, a firm makes an offer to his highest ranked worker. In the k^{th} period, if firm did not get matched in the previous period, he makes an offer to his highest ranked worker to whom he has not yet made an offer if he got rejected by the same choice more than ℓ times. He makes an offer to the same choice to whom he made an offer in $(k-1)^{th}$ period if he got rejected by that choice less than or equal to ℓ times. If he was matched in the previous period, he makes an offer to the choice to whom he was matched in $(k-1)^{th}$ period. In this strategy, ℓ is the delay and it can be different for every choice. As before, workers are assumed to myopically best respond. Lemma 1 can be trivially generalized for the delayed GDL strategies.

Corollary 1: If the number of periods in a repeated game is sufficiently large then the outcome of delayed GDL strategies will converge to the firm-optimal stable matching when workers are myopic.

When information about firms' preferences is available, workers can profitably deviate from myopic best response. From a worker's perspective, the firm-optimal stable matching is dominated by the worker-optimal stable matching. Any worker who strictly prefers the worker-optimal stable matching can reject any offer from a firm ranked lower than the worker-optimal stable match against GDL firms *to guarantee* convergence to a matching in which he gets matched to his worker-optimal stable partner.

Given the possibility of strategic manipulation by workers, we change the information available to subjects in our experiments to examine the impact of information on the predictive value of the DAA.

Experimental Design[1]

In each condition, three preference profiles are used such that the first profile converges in the firm-proposing DAA after five iterations, the second profile converges

[1] Instructions at www.utdallas.edu/~eeh017200/B2B/Instructions_all.doc

in six to seven iterations, and the last profile converges in eight iterations. That is, convergence becomes increasingly difficult as the experiment proceeds. Each profile is characterized by two stable matching outcomes—a worker-optimal and a firm-optimal outcome.

The payoff to a firm from matching to a worker can be 1, 2, 3, or 4 tokens. Staying unmatched is costly and results in the loss of 1 token for each unmatched party. Each subject begins with a 10 token initial endowment to prevent bankruptcies early on.

The experimental conditions vary along two dimensions, resulting in a 2x2 design. One dimension is the strategic behavior of workers. Workers can be computerized or human. A computerized worker is automated to select the best offer among all offers it has in a given period. Human workers are free to select any offer among the offers made to them. The other dimension is the amount of information available. In high information conditions, each subject has complete information about all other preference profiles and all past offers. In low information conditions, subjects see only their own preference profiles and only the offers that pertain to them.

In each period there are two stages. In the first stage, each firm can extend an offer to one and only one worker. After offers are simultaneously submitted by firms, workers simultaneously decide which offer to accept in the second stage. Each worker can accept only one offer in a period. When workers are automated, each worker accepts the best incoming offer. Each subject then observes his payoff for that market. Then the market game is repeated with the same preferences and same players in it. The first preference profile is used in the first 30 periods, the second preference profile in the next 60 periods, and the last preference profile in the last 60 periods (each session lasts for 150 periods).

Eight to ten cohorts were run for each condition. The subjects were Boston area undergraduate students. Subjects earned a participation fee and their token earnings were converted into $US. The exchange rate was 30 tokens per $1.

Results

We observe that delayed GDL strategy profile is not a bad approximation to firms' behavior. Table 1 shows that a significant portion of firms applied strategies strictly consistent with a delayed GDL strategy. In this table, rows titled "n^{th} offer" give the number of firms whose actions were strictly consistent with a delayed GDL strategy by the n^{th} distinct offer made in the profile. If a firm did not make n distinct offers, but it was consistent by $(n-1)^{th}$ offer with the delayed GDL then it is also consistent by n^{th} offer. The majority of firms initially offer to their first choice in all conditions. While the proportion of delayed GDL players increases in low-information conditions, it falls in high-information conditions after profile 1. Hence the lack of information forces firms to use delayed GDL more frequently. However, agents do not purely use this strategy; rather they experiment. Hence, the ratio of firms who behave consistently with a delayed GDL strategy gets smaller over time within each profile's periods.

In Table 2, we examine what the converged outcomes are. We say a *strong convergence* occurred if at least six out of the last 10 periods of the profile were that particular outcome. We also define a *weak convergence* if at least three periods were a particular outcome and the majority of the remaining in the last 10 periods had at most one player deviating from this outcome. We report the sum of the strongly and weakly converged outcomes. The numbers of weak convergences are reported in parentheses.

In the human worker conditions, we observe occasional convergence to the worker-optimal stable outcome. The number of convergences to the worker optimal outcome is 4/27 with low information versus 4/24 with high information. For the computerized worker conditions, we get similar results: 19/30 instances of firm optimal outcomes with low information and 20/30 instances with high information. The main difference is in the emergence of the worker-optimal stable matching. The worker-optimal stable matching emerges only once in the low information condition with computerized workers, although it emerges in 4/30 instances with high information and computerized workers.

Lastly, we examine whether the amount of information improves the welfare of workers, as one would expect if workers can behave strategically rather than myopically. We compare the conditions using two-sample t-tests in blocks of 10 periods. Though it appears that occasionally in the first block information makes a significant difference, and workers do obtain higher welfare levels, by the last block of the conditions with strategic workers, welfare levels converge and are frequently no longer significantly different between information conditions. The speed of convergence seems to be the main difference between the treatments. In computerized worker conditions, information makes a statistically significant difference at the end, especially in profile 3. However, the greater differences are still early on, indicating some convergence.

Conclusions

The first contribution of this work is in showing the DAA to be a useful predictor in repeated matching environments.

A second contribution pertains to the effect of information in repeated matching environments. First, high information may speed up convergence and improve welfare by allowing agents to avoid costly mistakes. Second, high information can alter the outcomes in the workers' favor since workers can unilaterally deviate, forcing convergence to a matching with their worker-optimal stable partner. We find evidence to support the first assertion: Mostly adding information made a statistically significant difference in average welfare for both workers and firms. Comparing settings, under both information conditions, with strategic workers as opposed to computerized workers, we find little or no significant difference in agents' welfare or in the number of instances of convergence to worker-optimal outcomes, suggesting that workers did not behave strategically. Instead, information improves welfare through more informed firm proposals.

References

Gale, D. and L. Shapley (1962), "College admissions and the stability of marriage," *American Mathematical Monthly*, 69, 9-15.

Roth, A. E. and M. Sotomayor (1990), "Two-sided matching: A study in game-theoretic modeling and analysis," Cambridge, Cambridge University Press.

Roth, A. E. and J. H. Vande Vate (1990), "Random Paths to Stability in Two-Sided Matching", *Econometrica*, 58, 1475-1480.

Roth, A. E. and X. Xing (1997), "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists," *Journal of Political Economy*, 105, 284-329.

Table 1: Proportion of firms who play consistent with a delayed GDL strategy

		Low Information-Human Workers			Low Information-Computerized Workers		
		Profiles			Profiles		
		1	2	3	1	2	3
Offers	1st	0.64	0.86	0.97	0.6	0.83	0.83
	2nd	0.64	0.81	0.94	0.58	0.68	0.75
	3rd	0.5	0.36	0.47	0.45	0.25	0.38
	4th	0.39	0.14	0.39	0.45	0.2	0.2
		High Information-Human Workers			**High Information-Computerized Workers**		
		Profiles			**Profiles**		
		1	**2**	**3**	**1**	**2**	**3**
Offers	1st	0.72	0.47	0.38	0.73	0.55	0.43
	2nd	0.53	0.22	0.28	0.58	0.4	0.35
	3rd	0.44	0.19	0.22	0.43	0.18	0.23
	4th	0.44	0.16	0.19	0.4	0.13	0.1

Table 2: Convergence results

		1. Low Information-Human Workers (9 cohorts)			1A. Low Information-Computerized Workers (10 cohorts)		
		Profiles			Profiles		
		1	2	3	1	2	3
Outcomes	Firm Optimal	8 (1)	8	6(2)	5(2)	7	7(4)
	Worker Optimal	1	1	2	1(1)	0	0
	Other	0	0	1(1)	1	3(1)	2
	None	0	0	0	3	0	1
		2. High Information-Human Workers (8 cohorts)			**2A. High Information-Computerized Workers (10 cohorts)**		
		Profiles			**Profiles**		
		1	**2**	**3**	**1**	**2**	**3**
Outcomes	Firm Optimal	6	4	5	7(1)	5	8
	Worker Optimal	0	2	2	0	4	0
	Other	2	1	0	2	1	2(1)
	None	0	1	1	1	0	0

Appendix A for Referees Only: Value profiles used in the experiment.

Matching (X,Y,Z,W)- The numbers in the parentheses are the workers matched to firms 1, 2, 3 and 4, respectively in the depicted matching.

Profile 1

Firm Payoffs

	Worker 1	Worker 2	Worker 3	Worker 4
Firm 1	3	1	4	2
Firm 2	4	1	3	2
Firm 3	4	3	2	1
Firm 4	1	4	3	2

Firm-Optimal Stable Matching: **(4,1,3,2)**

Worker Payoffs

	Firm 1	Firm 2	Firm 3	Firm 4
Worker 1	2	4	3	1
Worker 2	4	1	2	3
Worker 3	2	1	4	3
Worker 4	3	1	2	4

Worker-Optimal Stable Matching: **(2,1,3,4)**

Profile 2

Firm Payoffs

	Worker 1	Worker 2	Worker 3	Worker 4
Firm 1	2	4	3	1
Firm 2	4	3	1	3
Firm 3	3	1	2	4
Firm 4	4	2	1	3

Firm-Optimal Stable Matching: **(3,4,1,2)**

Worker Payoffs

	Firm 1	Firm 2	Firm 3	Firm 4
Worker 1	4	1	3	2
Worker 2	3	1	2	4
Worker 3	1	4	2	3
Worker 4	4	3	2	1

Worker-Optimal Stable Matching: **(1,4,3,2)**

Profile 3

Firm Payoffs

	Worker 1	Worker 2	Worker 3	Worker 4
Firm 1	1	3	4	2
Firm 2	1	3	2	4
Firm 3	1	4	2	3
Firm 4	2	4	1	3

Firm-Optimal Stable Matching: **(4,2,3,1)**

Worker Payoffs

	Firm 1	Firm 2	Firm 3	Firm 4
Worker 1	1	2	4	3
Worker 2	3	4	1	2
Worker 3	1	4	2	3
Worker 4	4	2	1	3

Worker-Optimal Stable Matching: **(4,2,1,3)**