Signaling in Matching Markets

Peter Coles[*] — *Harvard Business School*
Muriel Niederle[†] — *Stanford University and NBER*

February 2006 (preliminary and incomplete)

Abstract

We evaluate the effect of preference signaling in two sided matching markets. Firms and workers have strict preferences over members of the other side of the market. Each firm makes an offer to exactly one worker. Workers select the best offer from those available to them. The short time frame produces congestion and the market fails to reach a stable outcome. But if workers are able to signal their preferences, (i.e. their top choice firm,) firms may use this information as guidance for their offer choices. We find that in this signaling setting, it is optimal for firms to make use of these signals in the form of cutoff strategies. However, making use of signals imposes a negative externality on other firms. We find that on average, introducing a signaling technology increases the average number of matches, one possible measure of social welfare.

1 Introduction

Many entry-level labor markets, as well as many markets for educational positions share the feature that applicants become available at the same time and search for positions in the near future. An outcome in these markets is stable if no applicant-employer pair would prefer a match with each other to their current match. To ensure stability, significant information about preferences needs to be processed. In many settings, market frictions limit the amount of information that can be processed, so that stable outcomes are an unrealistic hope. However, these frictions may be mitigated by allowing applicants to signal preference information to potential employers.

[*]Harvard Business School, pcoles@hbs.edu

[†]Department of Economics, Stanford University, www.stanford.edu/~niederle, email: niederle@stanford.edu.

By introducing a signaling mechanism, where signals are (by construction) scarce, and hence credible, we may achieve a second best outcome. In this paper, we model such a signaling mechanism, analyzing behavior and welfare implications. Ultimately we hope the model will provide insight and policy recommendations as to when such a mechanism might be useful, and what form a signaling mechanism should take.

In two-sided matching markets (e.g with firms matched to workers, or men to women), the amount of information about preferences that must be transmitted to guarantee that a particular matching is stable can be significant. In a market where firms make costless proposals to workers, in a stable matching, each firm must be convinced that no worker better than its current match would prefer to renege on its match and pair with the firm instead. Segal [5] points out that the minimal informational requirements can be achieved via a firm proposing deferred acceptance algorithm (see e.g. Roth Sotomayor [3].) This means that every firm must have made an offer, and been rejected by, every worker it prefers to its current match. To get an idea of the number of offers that must be made, observe that with N workers and N firms with uncorrelated preferences over each other, the expected number of offers is approximately $N \cdot H_N$, where H_N is the Nth harmonic number (Knuth) [2]. When $N = 100$, $100 \times H_{100} \approx 500$ offers must be made. Since offers are often made sequentially, the time to reach a stable matching can be significant (Roth and Xing) [4].

There are many reasons we may not expect markets to solve this problem, that is to fully extract all the information necessary to achieve a stable matching. A market simply may not have enough time to clear because, for example, classes start in college, or jobs are about to start. Without a fixed endpoint, employers may still experience time constraints, as applicants may begin accepting offers elsewhere on which they cannot costless renege. Even when a market experiences no time restrictions, an employer may find itself subject to frictions. For example, offers may be costly, or in the extreme case a firm may be restricted in the number of offers it may make.[1]

In markets where frictions are an issue, employers often face a tradeoff between making an offer to their most desirable candidate (who is still available) and a less desirable candidate who is more likely to accept an offer. In such markets, we might expect applicants to convey their willingness to accept an offer. In a market with no frictions, and where the value of each

[1]For example, departments may face deans who argue that after four rejections, making an offer to a fifth choice candidate should be ruled out, because 'we are not a fifth choice candidate institution.'

match is known to the firm, applicants' indications of preferences are not useful to the firm.[2] Furthermore, if applicants may costlessly convey such preference signals to all employers, these signals may again be useless.

Signaling in Matching Markets

In practice, we find a multitude of markets that are prone to congestion, and in many of these markets, some form of preference signaling takes place. In some markets, formal signaling mechanisms have been introduced. In others, we witness informal signaling, where due to reputation or other considerations, the signals may still be credible.

An example of a market where preference signals are unambiguously interpreted by employers as a measure of likelihood of acceptance (and do not convey information about the value of a match) is the market for clinical psychologists as described by Roth and Xing [4]. From 1973 to 1998 the market operated under very specific rules that essentially emulated a deferred acceptance algorithm with offers and acceptances made over telephone. In this market, program directors for internships in clinical psychology compete for doctoral students. The rules impose a uniform time regime in which offers can be made, with fixed start and end times.[3] Because the market operates in a decentralized manner, and offers take time, the outcome of the market is not stable. Indeed, some programs find themselves rejected close to the end, and do not have enough time to make offers during the operation of the market. The aftermarket is also heavily regulated, and includes only applicants who have not accepted an offer elsewhere. A program that finds itself rejected close to the end of the market may not be able to fill its slot with a desirable applicant. This illustrates the market congestion; time can run out before programs have a chance to make offers to all the candidates they are interested in. On a site visit in 1993, Roth and Xing describe the behavior of one program and its program directors on selection day, which then lasted from 9 am to 4 pm. The program had 5 positions to fill, and a rank order list of 20 acceptable candidates. The co-directors said that their general strategy in the market was to not "tie up offers with people who will hold them the whole day." At the beginning of the market they made offers to candidates 1, 2, 3, 5 and 12, where 3, 5 and 12 had indicated

[2] For the base model in this paper, we assume that the information gathering stage has taken place, so that firms valuations of candidates are known. Hence, applicant signals of preferences are only valuable in that they affect the likelihood of offers being accepted. Alternative formulations could include signaling as part of an information gathering stage, so that signals convey information about the valuation of a match, or even yield considerations, where rejections themselves affect firm utility.

[3] In its later stages, the interval was a seven hour period within a single day.

they would accept an offer immediately. Candidates 1 and 2 were deemed attractive enough to be worth taking a chance on. Later in the day, other candidates called to report that the program was now the highest ranked program on their list, and indeed, the co-directors of the program decided to make the offers that were turned down by candidates 1 and 2 to the best of those candidates. The program eventually hired candidates 3, 5, 8, 10 and 12, all of which had indicated they would accept an offer immediately. Only candidates 1 and 2 received offers without having signaled their intent. Candidates 4, 6, 7, 9 and 11 (who were all preferred to some of the applicants eventually hired) were never made an offer. Roth and Xing summarize this episode by remarking on the program directors' concern over making offers which ran the risk of being rejected late in the day, the consequent attention paid to candidates who indicated that they would accept immediately, and the willingness of candidates to convey such information.

An important example of preference signaling is the college admissions market, in which early application to a college can be interpreted as a preference signal. For many colleges this is a market with two application periods, one early and one 'regular.' Although early application rules vary by college, many schools require that early applicants not send early applications to other competing schools, so that students are often faced with choosing exactly one school to which they can apply early. A distinction across schools is the commitment implications of early application. Some colleges used 'early action,' in which applicants may apply early but without any commitment to accept an offer from the college should they receive one. Other colleges use an 'early decision' plan that included students' commitment to attend the college should they be admitted. These plans, although varying by school, represent a formal way of sending a credible preference signal.[4] Avery et. al [1] examine this market in detail, and show that both early action and early decision applications result in a higher chance of admission than do regular applications.

In this paper we provide an explicit model of a labor market with frictions, and consider the effects of introducing a signaling mechanism. We examine the strategic behavior of firms and workers in the presence of such a mechanism, as well as welfare effects and comparative statics. We find that when preferences are uncorrelated it is optimal for firms to make use of these signals in the form of cutoff strategies. However, making use of signals imposes a negative externality on other firms. We find that on average, in-

[4]We also observe signaling in other forms during the regular admissions process via college visits, contact with the admissions office, etc.

troducing a signaling technology increases the average number of matches, one possible measure of social welfare. We hope that the model may provide insight into policy recommendations for second best solutions in congested markets.

2 No Signaling Structure

We first set up the market in the absence of a signaling mechanism. Let F = set of firms, W = set of workers, $|F| = |W| = N$. Each firm f has preferences over the workers chosen uniformly and randomly from the set of all strict preference orderings. Worker preferences are analogously chosen.

Each firm f makes an offer to a worker, and offers are made simultaneously. Workers choose an offer from those available to them.

Let Θ_{f_i} be the set of all preference lists (or rank order lists or ROLs) for firm i, and let Θ_{w_j} be the set of all preference lists for worker j. Lists $\theta_{f_i} \in \Theta_{f_i}$ and $\theta_{f_i} \in \Theta_{f_i}$ are vectors of length N. Let $\Theta = \Theta_{f_1} \times \ldots \times \Theta_{f_N} \times \Theta_{w_1} \times \ldots \times \Theta_{w_N}$ indicate preference list profiles.

We will consider Bayesian Nash Equilibria. The vector of agent types θ_F[5] is drawn from the space $\Theta_F = \Theta_{f_1} \times \ldots \times \Theta_{f_N}$ according to the uniform distribution. That is, the type of firm i is simply his preference list θ_{f_i}, and from i's perspective, other firms' preferences are uncorrelated and uniformly distributed.

Firm i with preferences θ_{f_i} values a match with worker w_j as $g(\theta_{f_i}, w_j)$, where $g(\theta_{f_i}, \cdot)$ is a von-Neumann Morgenstern utility function. We assume that utility of a match depends on worker rank. That is, for any permutation σ, we have $g(\sigma(\theta_{f_i}), \sigma(w_j)) = g(\theta_{f_i}, w_j)$. [6] We will sometimes abuse notation and write $g(\theta_{f_i}, \theta^j_{f_i})$ as $g(j)$, firm utility from matching with its jth ranked worker.

Definition 1. The **offer game with no signals** is given by

$$\Gamma^0 = [F, \{S_i\}, \{\pi_i(\cdot)\}, \Theta_F, U(\cdot)]$$

[5]Sequential rationality ensures that workers will always select the best offer from those available to them. Hence, we assume this behavior for the workers and focus on the reduced game with only firms as players.

[6]Let $\sigma : \{1, \ldots, N\} \to \{1, \ldots, N\}$ be a permutation. Abusing notation, we apply σ to preference lists, workers, and sets of workers such that the permutation applies to the worker indices. For example, suppose $N = 3, \sigma(1) = 2, \sigma(2) = 3$ and $\sigma(3) = 1$. Then we have $\theta_f = (w_1, w_2, w_3) \Rightarrow \sigma(\theta_f) = (w_2, w_3, w_1)$ and $\sigma(w_1) = w_2$.

In this definition, F is the set of firms. Message space $S_i = W$ for each firm i. Function $\pi_i : S \times \Theta_{f_i} \to \mathbb{R}$ yields the payoff to firm i as a function of the message profile and firm i's type. Note that this function will incorporate an expectation over all possible preferences of workers. Set Θ_F is the type space, the space of firm preference profiles. Finally, $U(\cdot)$ is the uniform distribution over the type space Θ_F.

Recall that a strategy for a firm i in this Bayesian game is a mapping $s_i : \Theta_{f_i} \to W$.

Definition 2. Strategy profile $(s_1(\cdot), \ldots, s_N(\cdot))$ is a Bayesian Nash Equilibrium if for all i and θ_{f_i} we have

$$s_i(\theta_{f_i}) \in \arg\max_{w \in W} \mathbb{E}_{\theta_{-f_i}}[\pi(w, s_{-i}(\theta_{-f_i}), \theta_{f_i}) \mid \theta_{f_i}]$$

We focus on strategies that depend on workers' rank within a firm's preference list, rather than index.

Definition 3. Firm i's strategy s_i is **anonymous** if $\forall\ \sigma \in \Sigma, \theta_{f_i} \in \Theta_{f_i}$ we have $s_i(\sigma(\theta_{f_i})) = \sigma(s_i(\theta_{f_i}))$.

For example, 'always make an offer to my second ranked worker' is an anonymous strategy, whereas 'always make an offer to w_2 is not.

Due to the uniform distribution over the type space, each firm optimally makes an offer to the highest ranked worker on its preference list. Hence, the unique Bayesian Nash Equilibrium of the offer game with no signals is

$$s_i(\theta_{f_i}) = \theta^1_{f_i} \text{ for all } i, \theta_{f_i}.$$

3 Model: Signaling Structure

We now modify the game so that each worker sends a signal to exactly one firm on its preference list. Rather than observing all signals from workers who have sent it a signal, firm i only sees its most preferred worker signal.[7]

Define $b_i(\theta) \in W \cup \{\emptyset\}$ to be the highest ranked worker according to θ_{f_i} who has ranked f_i first. We refer to this worker as f_i's 'top guaranteed worker' (TGW) or as f_i's 'top signal.'

Firm i's type space is now $\Theta_{f_i} \times (W \cup \{\emptyset\}) \equiv \Phi_{F_i}$. That is, f_i's type is its preference profile combined with its TGW, if it exists.

[7]One way to think about it is that the number of signals is a second order effect, whereas the value of this top worker is a first order effect. Otherwise, we can think of the firm as naively focusing on its best guaranteed option and forgetting about the others, as it would never make offers to them anyway.

Definition 4. The **offer game with signals** is given by

$$\Gamma = [F, \{S_i\}, \{\pi_i(\cdot)\}, \Phi_F, p(\cdot)]$$

In this definition, F is the set of firms. Message space $S_i = W$ for each firm i. Function $\pi_i : S \times \Phi_{f_i} \to \mathbb{R}$ yields the payoff to firm i as a function of the message profile and firm i's type. Note that this function will incorporate an expectation over all possible preferences of workers *consistent with f_i's type*. Distribution $p(\cdot)$ over types is derived from the uniform distribution over preference profiles.

A strategy for firm i consists of a function $s_i : \Theta_{f_i} \times (W \cup \{\emptyset\}) \to W$. Since we can infer f_i's TGW w from the complete firm/worker profile θ, we abuse notation and may also describe f_i's action as $s_i(\theta)$. Similarly, we sometimes will write i's payoff function as $\pi_i(s_i(\theta), s_{-i}(\theta), \theta)$.

Once again, we require strategies to be anonymous.

Definition 5. Firm i's strategy s_i is **anonymous** if $\forall\, \sigma \in \Sigma, \theta_{f_i} \in \Theta_{f_i}, w \subseteq (W \cup \{\emptyset\})$, we have $s_i(\sigma(\theta_{f_i}), \sigma(w)) = \sigma(s_i(\theta_{f_i}, w))$.

In the offer game with signals, we again focus on Bayesian Nash Equilibria.

Definition 6. Strategy profile $(s_1(\cdot), \dots, s_N(\cdot))$ is a Bayesian Nash Equilibrium if for all i and $\phi_{f_i} \in \Phi_{f_i}$ we have

$$s_i(\phi_{f_i}) \in \arg\max_{w \in W} \mathbb{E}_{\phi_{-f_i}}[\pi(w, s_{-i}(\phi_{-f_i}), \phi_{f_i}) \mid \phi_{f_i}]$$

We will be interested in a special type of strategy, the *cutoff* strategy.

Definition 7. Firm f_i's strategy $s_i(\cdot, \cdot)$ is a **cutoff strategy** if for all θ_{f_i}, and all w, w' with $w \preceq_{\theta_{f_i}} w'$, we have $s_i(\theta_{f_i}, w) = w \Rightarrow s_i(\theta_{f_i}, w') = w'$.

Observe that when strategies are anonymous, a cutoff strategy can be summarized by a single value. A firm using a *cutoff strategy with cutoff j* will make an offer to its TGW, provided its TGW is ranked j or higher. Otherwise it will make an offer to its TRW.

Cutoff strategies seem reasonable; if it is optimal for firm i to make an offer to his TGW, it seems that in a scenario where i has a better TGW, that worker too should receive an offer. The caveat is that worker signals provide information about the signals other firms receive, affecting their behavior, and hence the optimal decision for firm i.

Example 1. Non-optimal cutoff strategies. Suppose firms $-i$ all use the strategy of making an offer to the most preferred worker who has not signaled to them. From firm i's perspective, a higher TGW indicates that in expectation, other firms now have fewer signals from workers other than i's TGW, signals which firms $-i$ shy away from. Hence a higher TGW *reduces* riskiness of an offer to TRW. For carefully constructed utility functions, a cutoff strategy is not optimal. (For example, a utility function for firm i such that i is nearly indifferent between workers ranked j and $j+1$.)

The following proposition states that so long as other firms use cutoff strategies, it is optimal for firm i to use a cutoff strategy.

Proposition 1. *If firms $-i$ are using cutoff strategies $s_{-i}(\cdot)$, then it is optimal for firm i to use a cutoff strategy.*

To prove this proposition, we will need a 'junk signal' lemma.

Define a *junk signal* for a firm as a signal from a worker to whom the firm does not make an offer. The following lemma states that if other firms use cutoff strategies, then firms prefer not to receive junk signals.

Lemma 2. *Let firms $-i$ use cutoff strategies $s_{-i}(\cdot)$ and let firm i use strategy $s_i(\cdot)$. Then for $\forall\ \overline{\theta_{f_i}} \in \Theta_{f_i}, \forall\ w, w' \in W$ with $w \succeq_{\overline{\theta_{f_i}}} w'$, we have*

$$\begin{array}{l} \mathbb{E}_\theta[\quad \pi_i(s_i(\theta), s_{-i}(\theta), \theta) \quad | \quad \theta \in \Theta^w \quad] \ \leq \\ \mathbb{E}_\theta[\quad \pi_i(s_i(\theta), s_{-i}(\theta), \theta) \quad | \quad \theta \in \Theta^{ww'} \quad] \end{array}$$

where

$$\begin{array}{lll} \Theta^w & = \{ \quad \theta \in \Theta \quad | \quad \theta_{f_i} = \overline{\theta_{f_i}} \quad \textit{and} \quad b_i(\theta) = w & \} \\ \Theta^{ww'} & = \{ \quad \theta \in \Theta \quad | \quad \theta_{f_i} = \overline{\theta_{f_i}} \quad \textit{and} \quad b_i(\theta) = w \quad \textit{and} \quad \theta^1_{w'} = f_i & \}. \end{array}$$

Proof. Consider arbitrary $\overline{\theta_{f_i}} \in \Theta_{f_i}$ and $w, w' \in W$ with $w \succeq_{\overline{\theta_{f_i}}} w'$. If $s_i(\overline{\theta_{f_i}}, w) = w$, then the inequality holds trivially, because in both cases the payoff is f_i's certain payoff from matching with w. Assume $s_i(\overline{\theta_{f_i}}, w) = \overline{\theta^1_{f_i}}$, f_i's TRW. Define

$$\begin{array}{llllll} \Theta^w_+ & = & \{\theta \in \Theta^w \quad | \quad \pi_i(s_i(\theta), s_{-i}(\theta), \theta) > 0\} & \Theta^w_- & = & \Theta^w \backslash \Theta^w_+ \\ \Theta^{ww'}_+ & = & \{\theta \in \Theta^{ww'} \quad | \quad \pi_i(s_i(\theta), s_{-i}(\theta), \theta) > 0\} & \Theta^{ww'}_- & = & \Theta^{ww'} \backslash \Theta^{ww'}_+ \end{array}$$

We will now argue that

$$|\Theta^w| \ = \ N \cdot |\Theta^{ww'}| \quad \text{and} \tag{3.0.1}$$

$$|\Theta^w_+| \ \geq \ N \cdot |\Theta^{ww'}_+|. \tag{3.0.2}$$

This will be enough to show the result, as (3.0.1) and (3.0.2) combine to give

$$\frac{|\Theta_+^w|}{|\Theta^w|} \geq \frac{|\Theta_+^{ww'}|}{|\Theta^{ww'}|}$$

or

$$\frac{|\Theta_+^w|}{|\Theta^w|} \cdot g(\overline{\theta_{f_i}}, \overline{\theta_{f_i}^1}) \geq \frac{|\Theta_+^{ww'}|}{|\Theta^{ww'}|} \cdot g(\overline{\theta_{f_i}}, \overline{\theta_{f_i}^1})$$

which is equivalent to the inequality in the lemma. Define function $\psi : \Theta \times \{1, \ldots, N\} \to \Theta$ by

i. $(\psi(\theta, j))_{w'} = \theta_{w'}$, except that w' removes the highest ranked firm on its preference list and reinserts it in position j.

ii. $(\psi(\theta, j))_{\tilde{w}} = \theta_{\tilde{w}} \quad \forall \tilde{w} \neq w'$.

iii. $(\psi(\theta, j))_{\tilde{f}} = \theta_{\tilde{f}} \quad \forall \tilde{f}$.

To show (3.0.1), observe that $\psi : \Theta^{ww'} \times \{1, \ldots, N\} \to \Theta^w$ is a bijection.

To show (3.0.2), observe first that $\psi(\Theta_+^{ww'} \times \{1, \ldots, N\}) \subseteq \Theta_+^w$. To see this, consider any $\theta \in \Theta_+^{ww'}$ and $j \in \{1, \ldots, N\}$. For $j = 1$ clearly $\psi(\theta, j) \in \Theta_+^w$. For $j > 1$, $\psi(\theta, j)$ differs from θ in that w' moves f_i to position j and hence w' now signals to the second firm (call this firm k) on its list under θ. All firms other than k will make offers as before. Firm k, however, has a new signal under $\psi(\theta, j)$, namely w'. If this is a junk signal, then firm k will offer as under θ. If this signal (w')is now firm k's TGW, then because firm k is using cutoff strategies, the only way it may change its offer is to switch from its TRW to its new TGW w'. But from f_i's perspective, this only reduces competition for its TRW, as its TRW cannot be w'. Hence, $\psi(\theta, j) \in \Theta_+^w$. Furthermore, observe that $\psi : \Theta_+^{ww'} \times \{1, \ldots, N\} \to \Theta_+^w$ is a one-to-one mapping. □

To summarize, when junk signaler w' switches away from sending f_i a signal, this only be beneficial for f_i. The only impact it could have is to draw another firm to signal to w' and away from f_i's top choice.

We are now ready to demonstrate the optimality of cutoff strategies

Proof of Proposition 1. Suppose firms $-i$ are using cutoff strategies $s_{-i}(\cdot)$, and that $s_i(\cdot)$ is a best response to $s_{-i}(\cdot)$. Let $\overline{\theta_{f_i}}$ and w be any pair such that $s_i(\overline{\theta_{f_i}}, w) = w$. Let w' be the worker ranked just above w

according to $\overline{\theta_{f_i}}$. We will show that $s_i(\overline{\theta_{f_i}}, w') = w'$, from which the result follows. Define

$$\begin{array}{lcll} \Theta^w & = & \{\theta \in \Theta & \mid \; \theta_{f_i} = \overline{\theta_{f_i}} \;\text{ and }\; b_i(\theta) = w\} \\ \Theta^{w'} & = & \{\theta \in \Theta & \mid \; \theta_{f_i} = \overline{\theta_{f_i}} \;\text{ and }\; b_i(\theta) = w'\} \\ \Theta^{w'}_{+} & = & \{\theta \in \Theta^{w'} & \mid \; \theta^1_w = f_i\} \\ \Theta^{w'}_{-} & = & \{\theta \in \Theta^{w'} & \mid \; \theta^1_w \neq f_i\} \end{array}.$$

We calculate the returns to f_i naming its TRW when its TGW is w'. Begin by writing

$$\mathbb{E}_\theta\left[\pi_i(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\middle|\; \Theta^{w'}\right] = \begin{array}{lcl} p(\Theta^{w'}_{+}|\Theta^{w'}) & \cdot & \mathbb{E}_\theta\left[\pi_i(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\middle|\; \Theta^{w'}_{+}\right] \quad + \\ p(\Theta^{w'}_{-}|\Theta^{w'}) & \cdot & \mathbb{E}_\theta\left[\pi_i(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\middle|\; \Theta^{w'}_{-}\right] \end{array}.$$

Observe that

$$\mathbb{E}_\theta\left[\pi(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\middle|\; \Theta^{w'}_{-}\right] = \mathbb{E}_\theta\left[\pi(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\mid \Theta^{w}\right],$$

because under each condition, exactly one of $\{w, w'\}$ has signalled to f_i, and no worker ranked higher than w' has. By the junk signal lemma,

$$\mathbb{E}_\theta\left[\pi(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\middle|\; \Theta^{w'}_{+}\right] < \mathbb{E}_\theta\left[\pi(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\mid \Theta^{w}\right].$$

Hence, we have

$$\mathbb{E}_\theta\left[\pi(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\middle|\; \Theta^{w'}\right] < \mathbb{E}_\theta\left[\pi(\theta^1_{f_i}, s_{-i}(\theta), \theta) \;\mid \Theta^{w}\right]$$

and the result follows. □

The next proposition states that firms prefer that other firms not make use of signals. That is, using signaling strategies imposes a negative spillover on other firms.

Proposition 3. *Suppose that $s_{-i}(\cdot)$ and $s'_{-i}(\cdot)$ vary only in that firm $j \neq i$ has $s_j(\overline{\theta_{f_j}}, w) = \overline{\theta^1_{f_j}}$ while $s'_j(\overline{\theta_{f_j}}, w) = w$ for some pair $(\overline{\theta_{f_j}}, w)$. Then for all $i, s_i(\cdot)$, we have*

$$\mathbb{E}_\theta[\pi_i(s_i(\theta), s_{-i}(\theta), \theta)] \geq \mathbb{E}_\theta[\pi_i(s_i(\theta), s'_{-i}(\theta), \theta)].$$

The intuition is that when firms $-i$ make offers to their guaranteed workers, these workers will certainly be unavailable to firm i. Furthermore, when firm i is making a risky offer to its top ranked worker, this worker has not signaled to i, and hence has signaled to someone else. Therefore, offers from firms $-i$ to their TGWs are more likely to create competition for i than offers from $-i$ to their TRWs.

Proof. By the law of iterated expectations, we have

$$\mathbb{E}_\theta[\pi_i(s_i(\theta), s_{-i}(\theta), \theta] = \mathbb{E}\mathbb{E}[\pi_i(s_i(\theta), s_{-i}(\theta), \theta) \mid \theta_{f_i}, b_i(\theta)].$$

Hence, it is enough to show that conditional on any type θ_{f_i}, w, firm i prefers that firm j use the 'signal-ignoring' strategy $s'(\cdot)$. That is, $\forall \theta_{f_i}, w$,

$$\mathbb{E}[\pi_i(s_i(\theta), s_{-i}(\theta), \theta) \mid \theta_{f_i}, b(\theta)] \leq \mathbb{E}[\pi_i(s_i(\theta), s'_{-i}(\theta), \theta) \mid \theta_{f_i}, b_i(\theta)]. \quad (3.0.3)$$

Let $s_{-i}(\cdot)$ and $s'_{-i}(\cdot)$ satisfy the assumptions, and let $s_i(\cdot), \overline{\theta_{f_j}}$ and w be arbitrary. Define

$$\begin{array}{llclcl} \Theta^w \equiv \{ & \theta \in \Theta & \mid & \theta_{fi} = \overline{\theta_{fi}} & \text{and} & b_i(\theta) = w & \} \\ \Theta^w_+ \equiv \{ & \theta \in \Theta^w & \mid & \pi_i(s_i(\theta), s_{-i}(\theta), \theta) & < & \pi_i(s_i(\theta), s'_{-i}(\theta), \theta) & \} \\ \Theta^w_- \equiv \{ & \theta \in \Theta^w & \mid & \pi_i(s_i(\theta), s_{-i}(\theta), \theta) & > & \pi_i(s_i(\theta), s'_{-i}(\theta), \theta) & \} \end{array} .$$

and $\Theta^w_0 = \Theta^w - \Theta^w_+ - \Theta^w_-$.

In words, Θ^w_+ and Θ^w_- are the states of the world where f_j's strategy is critical to f_i. For states in Θ^w_+, f_i's offer is successful when opponents use s'_{-i}, but unsuccessful when they play s_{-i}. For states Θ^w_-, firm j's strategy is again critical to for firm i, but in the opposite direction.

We will show that

$$|\Theta^w_+| \leq |\Theta^w_-| \quad (3.0.4)$$

by constructing a one-to-one function $\psi : \Theta^w_+ \to \Theta^w_-$.

If $\Theta^w_+ = \emptyset$, we are done. Otherwise, consider any $\theta \in \Theta^w_+$. Since f_j's strategy is critical for f_i, we must have $s_i(\theta) = \overline{\theta^1_{f_i}} \equiv w_1$. Under $s'_{-i}(\cdot)$, f_i's offer is successful but under $s_{-i}(\cdot)$ it is not. Hence, we must have

$$\begin{array}{lclclcl} s_j(\theta) & = & \theta^1_j & & & = & w_1 \quad \text{but} \\ s'_j(\theta) & = & b_j(\theta) & \equiv & w_x & \neq & w_1. \end{array}$$

Note that w_1 has signaled to neither i nor j.

Define $\psi(\theta) \equiv \theta'$ as follows.[8] Beginning with θ, in f_j's ROL θ_{f_j}, reverse the position of w_1 and w_x. Then switch the top two firms of w_1 and w_x's ROLs (call the old top firm of w_1 f_y). Each list will now have this new top firm ranked twice. Replace the second appearance with the old top firm that was just swapped out. That is,

$$\begin{array}{rlcll} f_j: & (w_1, \ldots, w_x, \ldots) & \rightarrow & (w_x, \ldots, w_1, \ldots) & \text{and} \\ w_1: & (f_y, \ldots, f_j, \ldots) & \rightarrow & (f_j, \ldots, f_y, \ldots) & \\ w_x: & (f_j, \ldots, f_y, \ldots) & \rightarrow & (f_y, \ldots, f_j, \ldots) & . \end{array}$$

We now show that θ' is in Θ^{w}_{-}:

1. θ' satisfies $\theta'_{f_i} = \overline{\theta_{fi}}$ because f_i's ROL is unchanged.
2. θ' satisfies $b_i(\theta') = w$ because only w_1 and w_x's top candidates have been exchanged, and in neither case was this top candidate f_i.[9]

Therefore, f_i again makes offer to w_1, i.e. $s_i(\theta') = w_1$.

3. Under θ' and s'_j, f_i firm i's offer to w_1 is now un successful because

$$s'_j(\theta') = b_j(\theta') = w_1.$$

4. Under θ' and s_j, f_i's offer to w_1 is now successful:
   (a) Firm j makes its offer to $(\theta')^1_j = w_x \neq w_1$.
   (b) Firm y does not make an offer to w_1, as $(\theta')^1_y \neq w_1$ and $b_y(\theta') \neq w_1$. To see this, observe that
      i. If $w_1 = (\theta')^1_y = \theta^1_y$, then under θ, firm j's strategy could not be critical; firm y would always make a successful offer to its TRW w_1, who is also its TGW.
      ii. By construction of θ', $b_y(\theta') \neq w_1$.
   (c) Offers by other firms are the same under θ and θ', and hence, non pivotal.

[8] An alternative means of constructing $\psi(\cdot)$ would be to adjust the instance of w_x and w_i everywhere except in f_i's list, and then to swap the lists of w_i and w_x. This formulation may be the only one that works when number of offers k>1.

[9] This is the step that fails if we try to construct the one-to-one function in the other direction.

Hence, we have constructed a one-to-one mapping given by $\psi(\theta) = \theta'$, which proves (3.0.4). Finally, since the distribution $P_{\theta|\theta\in\Theta^w}(\cdot)$ is uniform, (3.0.3) follows and the proposition is proved. □

We now examine how a firm should adjust its behavior in response to changes in the behavior of opponents. We find that responding to signals is a case of strategic complements. One consequence of this result is that strategic complements can very well allow for multiple equilibria, which by Proposition 7 can be Pareto ranked.

Proposition 4. *Responding to worker signals is a case of* ***strategic complements****. That is, if firm j reduces its cutoff point (responds more to signals), firm i will optimally also reduce its cutoff.*

Proof. By proposition 3, when an opponent of firm i reduces its cutoff, the return to i's risky strategy of making an offer to its TRW is lowered. The safe payoff of going for a TGW remains the same. Hence, any tradeoff previously in favor of the TGW remains so, while some additional decisions may now favor the TRW - a lowering of i's cutoff. □

When your opponent firms are making offers to workers who have signaled to them, making an offer to a worker who has not signaled to you is particularly risky. She has signaled to someone else, who is then very much inclined to make her an offer. The greater this inclination on the part of your opponents, the riskier it is for you to make offers to your TRW, and hence, the more inclined you are to make an offer to your TGW as well. As we shall see in the next section, this feature may potentially lead to perverse welfare consequences.

4 Equilibria and Welfare

We now turn to equilibrium predictions of the signaling game. Continuity properties guarantee the existence of equilibria in which players all use the same cutoff. Interestingly, the strategic complementarity in cutoffs ensures that we can find a *pure strategy* symmetric cutoff strategy. The results in section 3 allow us to rank the equilibria from the perspectives of both workers and firms. We also consider another measure of welfare, the ex-ante expected number of matches.

Before proving equilibrium existence, a few observations about mixed strategies. If a firm finds two distinct cutoffs i, j $(i < j)$ to be optimal, we have

i. $|i - j| = 1$

ii. When TGW has rank j, it is indifferent between its TRW and its TGW.

iii. Mixing between cutoffs i, j $(i < j)$ is the same as randomizing between TRW and TGW when its TGW is ranked j and making an offer to TGW when TGW is higher ranked than j.

In light of observations, we can interpret a mixed cutoff strategy as a value in $[1, N]$. That is, cutoff $j + \frac{a}{b}$ is mixed strategy $(1\text{-}\frac{a}{b}) \cdot j + \frac{a}{b} \cdot (j + 1)$.

Proposition 5. *There exists a symmetric cutoff equilibrium.*

Proof: Observe that best response correspondences are

1. Convex
2. Upper hemicontinuous (from the theorem of the maximum)
3. map from $[1, N]$ to $[1, N]$

The fixed point follows from Kakutani. □

However, by using the conjecture that signaling strategies are strategic complements, we can show that there is a pure strategy equilibrium using the lattice point strategy sets $\{1, 2, \ldots, N\}$.

Proposition 6. *When strategy sets are* $\{1, 2, \ldots, N\}$, *there exists a pure strategy symmetric equilibrium.*

Proof: The proof is an application of Tarski's fixed point theorem on lattices. Consider the best response function $\text{br}(j)$ which yields the optimal cutoff for f when other firms all use cutoff j.[10] By the strategic complements assumption, $\text{br}(j)$ maps $\{k, k + 1, \ldots, N\}$ to itself and is (weakly) monotonic increasing. By Tarski's fixed point thereom , there exists a fixed point, which is a pure strategy symmetric equilibrium. □

In general, cutoff equilibria are not unique. It is possible, for example, to construct a setting in which two equilibria exist, one of which involves an interior cutoff, and one which involves a cutoff of 1; that is, firms ignore signals entirely.

The next proposition states that n the case of multiple symmetric cutoff equilibria, firms prefer the equilibrium with the highest cutoff; that is, the cutoff that involves the least use of signals.

[10] If ever two cutoffs are optimal, choose the lower cutoff.

Proposition 7. *Suppose there exist two symmetric pure strategy equilibria with cutoffs j, k where $j < k$. Then all firms have higher (ex-ante) expected payoffs in the equilibrium with cutoff k than in the equilibrium with cutoff j.*

Proof. Let s' and s be symmetric pure strategy cutoff equilibria with cutoffs j and k, respectively, with $j < k$.

Since s is an equilibrium, for any firm i, we have

$$\mathbb{E}[\pi_i(s_i(\theta), s_{-i}(\theta), \theta)] \geq \mathbb{E}[\pi_i(s'_i(\theta), s_{-i}(\theta), \theta)].$$

By proposition 3, we have

$$\mathbb{E}[\pi_i(s'_i(\theta), s_{-i}(\theta), \theta)] \geq \mathbb{E}[\pi_i(s'_i(\theta), s'_{-i}(\theta), \theta)].$$

Combining the two inequalities proves the result. □

This result may seem unusual, in that it appears to indicate that firms prefer less signaling to more. However, note that this comparison only holds for equilibria. It may very well be that a cutoff equilibrium is preferred by all firms to a non-equilibrium strategy profile with higher cutoff.

When a firm optimally chooses to respond to signals, it increases its payoff and lowers that of others. Hence, the overall effect on welfare is ambiguous. However, by another welfare standard, the welfare effect is unambiguous. Proposition 8 states that in expectation, firms responding to signals increases the number of matches.

Proposition 8. *(Number of Matches.) Let firms $-i$ use strategies $s_{-i}(\cdot)$. Let $s_i(\cdot)$ differ from $s'_i(\cdot)$ only in that for some pair $(\overline{\theta_{f_i}}, w)$, we have $s_i(\overline{\theta_{f_i}}, w) = \overline{\theta^1_{f_i}}$ while $s'_i(\overline{\theta_{f_i}}, w) = w$. Then*

$$\mathbb{E}_\theta[\ |\mu(s_i(\theta), s_{-i}(\theta), \theta)|\ ] \leq \mathbb{E}_\theta[\ |\mu(s'_i(\theta), s_{-i}(\theta), \theta)|\ ].$$

Proof. Define

$$\begin{aligned} \Theta^+ &= \{\ \theta \in \Theta\ :\ |\mu(s_i(\theta), s_{-i}(\theta), \theta)| < |\mu(s_i(\theta), s'_{-i}(\theta), \theta)|\ \} \quad \text{and} \\ \Theta^- &= \{\ \theta \in \Theta\ :\ |\mu(s_i(\theta), s_{-i}(\theta), \theta)| > |\mu(s_i(\theta), s'_{-i}(\theta), \theta)|\ \}. \end{aligned}$$

Sets Θ^+ and Θ^- are the profiles for which firm i's strategy swap (s to s') affects the number of matches. In particular, this means

$$\theta \in \Theta^+ \cup \Theta^- \Rightarrow \theta_{f_i} = \overline{\theta_{f_i}} \ \text{ and } \ b_i(\theta) = w.$$

For any $\theta \in \Theta^+$, it must be that

$$|\mu(s_i(\theta), s'_{-i}(\theta), \theta)| - |\mu(s_i(\theta), s_{-i}(\theta), \theta)| = 1. \tag{4.0.5}$$

For $\theta \in \Theta^+$, it must be the case that without firm i's offer, $\theta^1_{f_i}$ has an offer from another firm, and w does not. Similarly, for $\theta \in \Theta^-$, it is w who has an outside offer, and hence

$$|\mu(s_i(\theta), s_{-i}(\theta), \theta)| - |\mu(s_i(\theta), s'_{-i}(\theta), \theta)| = 1. \tag{4.0.6}$$

We will now show that $|\Theta^+| \geq |\Theta^-|$. Equations (4.0.5) and (4.0.6) along with the uniformity of the distribution over preferences will then be enough to prove the result.

Construct function $\psi : \Theta \to \Theta$ as follows: Let $\psi(\theta)$ be the profile in which workers have preferences as in θ, but firms $-i$ all swap the positions of workers $\overline{\theta^1_{f_i}}$ and w in their preference lists. For any $\theta \in \Theta^-$, without firm i's offer, w has an offer from another firm, and $\overline{\theta^1_{f_i}}$ does not. But then when preferences are $\psi(\theta)$, without firm i's offer

(i) Worker $\overline{\theta^1_{f_i}}$ ***must*** have another offer.

(ii) Worker w ***cannot*** have another offer.

To see (i), note that under θ, worker w signals to f_i, so his outside offer must come from a firm j who has ranked him first. Under $\psi(\theta)$, this firm ranks $\overline{\theta^1_{f_i}}$ first. If $\overline{\theta^1_{f_i}}$ has not signaled to firm j, then by anonymity, $\overline{\theta^1_{f_i}}$ gets f_j's offer. If $\overline{\theta^1_{f_i}}$ has signaled to firm j, then by the positive interpretation of signals assumption, $\overline{\theta^1_{f_i}}$ again gets f_j's offer.

To see (ii), suppose to the contrary that under $\psi(\theta)$, worker w does in fact receive an offer from some firm $j \neq i$. Since w signals to f_i, w must be f_j's TRW under $\psi(\theta)$, so that $\overline{\theta^1_{f_i}}$ is f_j's TRW under θ. But then by anonymity and the positive interpretation of signals assumption, $\overline{\theta^1_{f_i}}$ receives f_j's offer under θ, a contradiction.

From (i) and (ii), we have

$$\theta \in \Theta^- \Rightarrow \psi(\theta) \in \Theta^+.$$

Since $\psi(\cdot)$ is clearly injective, we have $|\Theta^+| \geq |\Theta^-|$, which is enough to prove the result. □

To summarize the intuition in the proof, observe that when a firm switches his offer his TRW to his TGW, it is the outside offers of these workers that determine the impact on the total number of matches. If both firms have outside offers, or if neither has an outside offer, the number of matches is unaffected. When exactly one firm has an outside offer, it is more likely to be the TRW, as this worker has signaled to another firm, while the TGW has not. Hence, by making an offer to its TGW, the firm maximizes the expected total number of matches.

In addition to increasing the expected number of matches, response to signals also unambiguously increases worker welfare.

Proposition 9. *(Worker Welfare.) Let firms $-i$ use strategies $s_{-i}(\cdot)$. Let $s_i(\cdot)$ differ from $s'_i(\cdot)$ only in that for some type class $\sigma \in \Sigma(\widetilde{\theta_f}, \tilde{w})$, we have $s_i(\sigma) = TRW(\sigma)$ while $s'_i(\sigma) = TGW(\sigma)$. Then*

$$\mathbb{E}_\theta[\pi^W(s_i(\theta), s_{-i}(\theta), \theta)] \geq \mathbb{E}_\theta[\pi^W(s_i(\theta), s'_{-i}(\theta), \theta)].$$

Proof. For every profile θ such that the switch $(s \to s')$ causes a net worker welfare decrease, we can find a profile such that the the change causes a net worker welfare increase of greater magnitude as follow. Begin with θ. Note that firm i's type is $(\theta_i, b_i(\theta)) \equiv (\theta_i, w)$. Now swap the positions of θ_i^1 and w in the lists of all firms $-i$. Swap the lists of θ_i^1 and w, except for f_i, which maintains its ranking in the two lists.

This switch has the effect of swapping the outside (non f_i) offers of θ_i^1 and w, except possibly adding some offers for θ_i^1 (who may not signaled to f_i) and possibly removing some offers to w (who has signaled to f_i). Under this new profile θ', net worker welfare is greater from f_i's offer to w, and because of the extra outside offers for θ_i^1 (and fewer for w), the welfare difference is greater than under θ.

Finally, since under anonymous strategies, workers all receive ex ante identical utility, this change must improve welfare for all workers. □

5 Market Size

The value of a signal to a firm depends very much on market conditions. If there are many more workers than firms, a firm may recognize its position of power and ignore signals from workers. Rather, it will make an offer to its top ranked worker, knowing this worker will likely be pleased to receive an offer at all. On the other hand, if workers are in short supply, then a firm will be ecstatic to receive any worker signal and will likely respond with an offer to avoid remaining unmatched. This idea is expressed in the following proposition.

Proposition 10. *Let s^l, s^h be the equilibria with lowest and highest cutoffs in the market with M firms and N workers. Let $\tilde{s}^l, \tilde{s}^h$ be the equilibria with lowest and highest cutoffs in the market with $M+1$ firms and N workers. Let $k^l, k^h, \tilde{k}^l,$ and $\tilde{k}^h$ be the respective cutoffs of these equilibria. Then*

$$\tilde{k}^l > k^l \text{ and } \tilde{k}^h > k^h.$$

Proof. Define best response function $\text{br}_M(j)$ to be the optimal cutoff for firm f when the other $M-1$ firms all use cutoff j[11] Define $\text{br}_{M+1}(j)$ similarly for the $M+1$ firm market.

Observe first that $\text{br}_M(k_h) = k_h$. If otherwise $\text{br}_M(k_h) > k_h$, then we can find a cutoff equilibrium with cutoff $> k_h$. Next observe that $\forall j, \text{br}_{M+1}(j) > \text{br}_M(j)$. That is, when there is an additional firm, if other firms continue to behave in the same manner, it is optimal for f to pay more attention to signals. But now $\text{br}_{M+1}(\cdot)$ maps the interval $[k_h, M+1]$ onto itself and is weakly monotonic increasing. Hence by Tarski's fixed point theorem, there is an (integer) cutoff $j^* \geq k_h$ with $\text{br}_{M+1}(j^*) = j^*$, from which we have $\tilde{k}^h > k^h$.

□

6 Extensions

There are a number of important problems that this baseline model can be used to explore, including robustness to changes in preferences and market size, as well as questions involving optimal mechanism choice.

1. **Perfect Correlation.**

   (a) *Perfect correlation of worker preferences.* In this case, a signaling mechanism adds no value. To see this, observe that the #1 firm will ignore signals, as he is the top choice of all workers. The #2 firm doesn't care what the firms below him are doing, and has no information about what firm #1 is planning on doing. Hence, firm #2 will ignore signals. And so on....

[11] If ever two cutoffs are optimal, choose the higher cutoff.

(b) *Perfect correlation of firm preferences.* Even in the no signaling case, the equilibrium is tricky. There may be an asymmetric 'coordination equilibrium' where firm 1 makes an offer to worker 1, etc.[12] There may also be a symmetric equilibrium where firms mix over workers.

Here signaling can add value. Since workers have no information about the preferences of their fellow workers, they will signal their top choices truthfully. Firms, after observing their TGW, may still conceivably make an offer to a worker other than their TGW. For example, there is a chance that #2 ranked worker may have the same top choice firm as the #1 ranked worker, in which case it may be worth making the #2 ranked worker an offer. We postulate a symmetric mixed strategy equilibrium where firms, upon seeing their TGW, mix offers over workers ranked $2 \ldots TGW - 1$. Under some regularity conditions (e.g linear utility of rank), this mixture will be over an interval.

At best, we might hope to compare the (unique symmetric?) equilibria in the no signaling and signaling cases in a simple 3x3 model.

2. **Correlation: Two tiers of school.** There are two tiers of firms, A and B, and two tiers of workers, A and B. All firms in A are preferred by all workers to firms in B, and all workers in A are preferred by all firms to workers in B. Within tiers, preferences are uncorrelated. If workers have one signal to send, top tier workers may choose to send the signal to their top choice in tier A or B. Firms may make offers to workers in tier A or B. Assume there is an imbalance in that there are more top tier workers than top tier firms. Then the only cross signaling we would expect would be for top tier firms to signal to second tier firms. Similarly, we would only expect cross tier offers to come from second tier firms to top tier workers.

   Unlike the single tier models, there is now an important difference between early decision and early action. Under early action, a top tier worker may be willing to signal to a second tier school, knowing that he can always accept offers from top tier schools who send him offers.

[12] Anonymity is no longer a requirement, as workers are naturally indexed by rank.

Under early decision, we would expect a smaller fraction of top tier workers to send signals to second tier firms, meaning that more top tier workers would fall through the cracks.

One possible modeling simplification is that firms value all workers in A at u_A and all workers in B at U_B. However, we then lose the within tier TRW-TGW trade off.

3. **Size of Market.** Increasing the number of firms shifts the set of equilibria so that workers respond more to signals. We conjecture that increasing the number of firms also reduces worker welfare. We further conjecture that the value of a signaling mechanism is greatest in balanced markets.

4. **Additional Signals.** Concavity of introducing additional signals (i.e workers send top two signals). Conjecture: when firms only have one offer to make, the optimal number of signals is 1. When firms have k offers to make, the optimal number of signals cannot be $> k$.

5. **Endogenous Signal Credibility.** Suppose that workers were in fact allowed to signal to multiple firms, but that all signals were made public. In this setting, despite the costless nature of the signals, workers may be hesitant to signal to all firms, as their signals may then be disregarded. A worker who signals to just a few firms may be more likely to be take seriously. Such a mechanism has the advantage of providing flexibility in how a worker signals; a single "strong" signal versus multiple "weak" signals, and credibility is endogenous. Furthermore, one could envision worlds in which the publicizing is itself a costly act on the part of the signaler, so that the need for a signaling mechanism itself is obviated.

7 Conclusion

Market congestion is a reality, and the abundance of signaling in congested markets suggests that signaling deserves consideration as a means of achieving second-best outcomes.

We have examined a natural signaling mechanism; allowing workers to signal to a sole firm. In a setting of uncorrelated preferences, workers will send this signal to their top choice firm, and the firms use this information as guidance for their offers. We find that in this signaling setting, it is optimal for firms to make use of these signals in the form of cutoff strategies. However, making use of signals imposes a negative externality on other firms, so that while beneficial for workers, the overall welfare effect of a signaling technology on firms is ambiguous. We find that on average, introducing a signaling technology increases the average number of matches, one possible measure of social welfare.

Ultimately the goal is to provide policy advice for a broad set of environments as to when an organized signaling mechanism might improve outcomes, and should this be the case, what forms the signaling mechanism might take. Our hope is that the approach in this paper will serve as a tool and as a benchmark; a framework for examining settings with alternative preference/market assumptions, and a point of comparison for alternative signaling mechanisms.

References

[1] Christopher Avery, Andrew Fairbanks, and Richard Zeckhauser. *The Early Admissions Game: Joining the Elite.* Harvard University Press, 2003.

[2] Donald E. Knuth. Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms. *CRM Proceedings and Lecture Notes, American Mathematical Society*, 10, 1997.

[3] Alvin E. Roth and Marilda A. Oliveira Sotomayor. *Two-sided Matching; a Study in Game-theoretic Modeling and Analysis.* Cambridge University Press, 1990.

[4] Alvin E. Roth and Xiaolin Xing. Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. *Journal of Political Economy*, 105, April.

[5] Ilya Segal. The communication requirements of social choice rules and supporting budget sets. 2004.