Signalling Preferences in Interviewing Markets*

Robin S. Lee† Michael Schwarz‡

September 5, 2007

Abstract

The process of match formation in matching markets can be divided into three parts: information sharing, investments in information acquisition, and the formation of matches based on available information. The last stage where agents are assumed to know their preferences has been studied in seminal work of Gale and Shapley (1962), and a model of second stage costly information acquisition is introduced and studied in Lee and Schwarz (2007). This paper focuses on the first stage – information sharing – and examines mechanisms which allow workers to signal their preferences over matching partners prior to the assignment of interviews. The incentives of firms and workers vis-a-vis information revelation are partially aligned – all other things being equal, a worker prefers to have an interview with a firm that is high in his preference ranking and a firm prefers to invest in interviewing a worker who ranks a firm highly because such worker is more likely to accept a job if offered. However, the incentives are far from being perfectly aligned. For instance, if firms pay the full cost of interviewing, each worker would prefer to have as many interviews as possible, and in a world with bilateral communication no information is revealed as each workers would want to tell each firm that it is his first choice. But if communication is moderated through an intermediary or there is a restriction on the number of messages a worker can send, then cheap talk becomes informative. Currently existing market institutions that facilitate information exchange prior to interviewing are discussed.

KEYWORDS: cheap talk, job search, labor market, matching, interview assignment

*We thank David McAdams for helpful comments.

†Harvard University and Harvard Business School, contact: lee54@fas.harvard.edu. Part of this research was conducted during an internship at Yahoo! Research, Berkeley.

‡Yahoo! Research and NBER, contact: mschwarz@yahoo-inc.com.

Dagstuhl Seminar Proceedings 07271
Computational Social Systems and the Internet
http://drops.dagstuhl.de/opus/volltexte/2007/1163

1 Introduction

The process of match formation in markets such as the marriage and labor markets can be loosely divided into three parts: information sharing, investments in information acquisition, and the formation of matches based on available information. The last stage where agents are assumed to know their preferences has been studied in seminal work of Gale and Shapley (1962) and in flourishing literature that followed (see Roth and Sotomayor (1990) for a survey). A model of second stage costly information acquisition, which we label the "interview assignment problem," is introduced and studied in Lee and Schwarz (2007). This paper focuses on the first stage (information sharing). Information sharing includes the worker revealing information about his ability and/or his preferences over potential employers. We examines mechanisms which allow workers to signal their preferences over matching partners prior to the assignment of interviews.

Considerable resources are devoted to information acquisition (e.g., interviewing or dating) in matching markets. With regards to labor markets, a firm typically can only interview a limited number of candidates due to the often significant costs per interview; however, since these interviews affect the formation of preferences and determine which partners are feasible matches, their allocation is as important to efficiency as the design and execution of the matching process itself. Lee and Schwarz (2007) tackles this interviewing assignment problem "in isolation" assuming that agents ex-ante do not have private information about their preferences, and demonstrate that even in a simple version of the interview problem it may be difficult to characterize the equilibrium or solve for the efficient outcome. However, in many instances, the preferences of workers and firms regarding whom to interview are partially aligned. If there are two workers who are identical except for the fact that one strongly prefers the firm and the other does not, the firm will prefer to interview the worker who has the stronger preference. When interviews are costly, interviewing a worker who has a strong preference for the firm indicates a greater likelihood that the worker will accept a job offer if one is made; similarly, interviewing a worker who does not like the firm is a bad investment since the worker would not be likely to accept a job.

When interviewing costs are primarily borne by the firm – which can be significant once the time of employees are accounted for – the more important it is for the firm to only interview those workers who are predisposed to accepting an offer.

Nonetheless, without a credible means of signalling, workers have little incentive to truthfully report their preferences and instead will exaggerate their affinity for certain firms: when the cost of unemployment are sufficiently high, workers will want to maximize the number of interviews in order to have the greatest chance of being employed. In turn, this inability to meaningfully communicate can lead to a great deal of inefficiency and friction in the market. To illustrate, consider the allocation of interviews in marriage markets – i.e., dates. A recent article profiled on Yahoo!'s frontpage mentions "mistakes" that women often make in online dating markets, including the following:

> Mistake #7: Expecting Him to Tell the Truth in His Profile You don't like to be lied to. Nobody does. And once you've gone out with a man who claimed to be 5'9" but is really 5'5", it's hard to keep dating. But haven't you ever done the same thing? The typical woman exaggerates her height by one inch and lowers her weight by 20 pounds. And it's not just a coincidence that the most popular ages for women on dating sites are 29, 39, 44 and 49. You want to be given a chance. You don't want to be judged before you meet. And you're insecure that telling the truth won't get you in the door against younger, thinner women. So if there are good reasons why an honest woman might be tempted to misrepresent herself, wouldn't it make sense that an honest man might be tempted to do the same thing?[1]

This paper explores ways to allow for meaningful communication between parties before interviews are assigned. In a stylized setting, we show that private bilateral communication between workers and firms is uninformative and is ignored. However, either utilizing an intermediary – which prevents workers from sending conflicting messages – or a restriction on the number of messages a worker can send can be sufficient to sustain an non-babbling equilibrium.

The communication between agents prior to the interviewing stage can be viewed as a form of cheap talk. Workers may communicate to a firm information about their ability and skills as well as information about their preferences. Note that to maximize the probability of an interview, a worker may want to communicate to each firm that he is a good match (a low quality firm would prefer to interview a

[1] http://dating.personals.yahoo.com/singles/gettingstarted/616/10-classic-online-dating-mistakes, accessed on August 23, 2007.

medium or low ability worker because high quality workers are unlikely to accept offers). In a world where firms pay the cost of interviewing, workers want to get as many interviews as possible. Thus, if communication is bilateral, a worker will have an incentive to tell each firm that he is a perfect match both based on his ability and preferences making informative communication impossible. But if communication is public or has some public component, information transmission becomes possible. The information transmission may be imperfect even with public communication for reasons similar to the classic model of cheap talk (e.g., Crawford and Sobel (1982)).[2] However, this paper emphasizes a novel aspect of cheap talk that is essentially unrelated to classic cheap talk models: unlike the classic model of cheap talk where there is a single receiver of information, the present game has multiple receivers with different preferences. The ability to send messages publicly via a centralized authority or some central communication channel is critical for making information transmission possible.[3,4]

These insights are indeed borne out in reality. In certain markets, an intermediary that can monitor or restrict the form of communication between workers and firms does exist. Such is the case with on-campus recruiting at many professional schools. For example, at Harvard Law School, a mechanism very similar to that described in this section is employed for assigning summer associate interviews at law firms to second-year law students: each student is allowed to submit a preference rank-ordered list of up to 35 firms; firms then assign interviews giving priority to those students who rank those firms the highest. Importantly, employers are prohibited from pre-screening candidates – they cannot view student records, resumes or transcripts – and consequently applicants truly are ex ante homogenous.

[2]For instance, if a man would like to date women who are a few years younger than them and if women prefer to date men their same age, perfect information revelation is not possible; however, it is possible to reveal some information about age because interests are partially aligned – although a man might view an ideal date to be a few years younger than him, he would not want a date someone much younger because it is very unlikely to lead to a match (again, given that women prefer men close to their age).

[3]Politicians face a similar problem. To be elected a politician may want to tell each constituency what it wants to hear. If all messages are broadcasted to all constituencies, communication may become more truthful if a politician values getting the support of like minded people more than the support of other voters.

[4]In a model that focuses on unravelling in matching markets, Ostrovsky and Schwarz (2006) allow for an interpretation of a transcript as a cheap talk about student's ability. That paper shows that if transcripts of all (or a random sample of students) are publicly observable than schools can credibly reveal to employers some information about student ability even when schools cannot commit to honest grading in a one period game.

On the other hand, some business schools (such as Michigan and UCLA) provide students with a limited budget to bid on interview spots that certain companies have left open.[5] And in the job market for junior economists, candidates are allowed to send up to two signals to potential departments expressing their interest in being interviewed.[6] In the latter two cases, providing applicants with a fixed and limited budget effectively accomplishes the same goal as monitoring and limiting the number of (conflicting) signals that can be sent.

2 Signalling Preferences

For discussion, we adopt a stylized model of communication, interviewing, and matching. We assume there are an equal number N of firms and workers (where the set of firms is given by F and the set of workers by W), each firm hires at most one worker, and market participants will eventually be matched via a matching process which we assume to be a firm-proposing deferred acceptance algorithm (Gale and Shapley (1962)). Before engaging in this process, firms must discover their preferences over workers by interviewing them; they do so by simultaneously assigning interviews (and conducting them) to a subset of workers, and each firm bears a strictly positive cost c per interview. On the other hand, workers already know their preferences before interviewing, and they are uncorrelated and distributed uniformly over firms.

As shown in Lee and Schwarz (2007), for any number of interviews x, there exists a cost c such that in equilibrium, each firm will interview exactly x workers and each worker will obtain exactly x interviews. In this equilibrium, firms will discover their ranking of workers during their interviews, and participants will report their true preferences over partners during the matching process. Since workers cannot influence interview assignment in this model, their preferences in this stage are ignored.

Nonetheless, it may be socially desirable for firms to know and utilize worker preferences when assigning interviews. For example, if firms ignored worker preferences, often there may be two workers who would wish to swap interviews if each

[5] See also Sönmez and Ünver (2004). They show that although using a single instrument (bids) to communicate both preferences and claims can have efficiency costs in the case of course bidding, under certain assumptions these problems do not extend to interview-bidding since agents there do not have binding quotas.

[6] http://www.aeaweb.org/joe/signal/

preferred the other's firm without adversely affecting the utility of any firm; in fact, if given the choice, firms would wish for such trades to happen since each firm (weakly) prefers interviewing the worker who prefers that firm more strongly than the other. Doing so would increase the probability that an offer, if made, would be accepted, and increase the returns from interviewing a particular worker.

We now consider the possibility of allowing for a communication stage between workers and firms prior to the interview stage. We first show that private, bilateral communication between each firm and worker cannot yield any improvement for the workers: all communication is either ignored or uninformative in any equilibrium. We then construct a mechanism whereby workers "publicly" communicate their complete preferences to an intermediary, and this intermediary recommends to each firm a subset of workers to interview. Such a mechanism has an equilibrium whereby workers report their true preferences, and firms obey its recommendations. Finally, in cases where such public communication is infeasible or the communication of complete rank order lists infeasible, we show that merely being able to monitor the number of interviews each worker has during the interview assignment stage can substantially improve upon the no-communication outcome.

For the purposes of this discussion, we abstract away from the coordination issues introduced in Lee and Schwarz (2007), and assume the following:

Assumption 2.1. *Fix c such that there exists a symmetric pure strategy equilibrium interview assignment $\hat{\eta}$ in which each firm and each worker conducts x interviews.*

An interview assignment η is simply a correspondence from the set of workers and firms $F \cup W$ into itself such that $f \in \eta(w)$ (which means w interviews with f) if and only if $w \in \eta(f)$, where f and w are elements of F and W, respectively.

2.1 Bilateral Communication

We first address the case where prior to the interview selection stage, each worker simultaneously is able to privately send a message m_{wf} to each firm indicating how much that worker prefers that firm. Though the message space M can be defined generally, one natural candidate would be $M \equiv \{1, ..., N\}$, where m_{wf} can be understood to be the rank that firm f is on w's preference list.

Our first result shows that if workers communicate privately with each firm, then it turns out that such communication is "cheap" and cannot achieve a better outcome for workers than in the case without worker communication.

Proposition 2.1. *If communication between workers and firms is private and non-verifiable, then communication is uninformative or ignored in equilibrium.*

This result should not be too surprising – with private bilateral communication, there is nothing to prevent workers from sending conflicting messages to each firm, and thus every worker will tell each firm what it needs to say in order to maximize his chance of obtaining an interview. Indeed, we often observe in practice that applicants often send cover letters to firms stating each firm is their "top choice"; in turn, firms have learned to discount such cheap talk accordingly to the extent that such signals are ignored.

2.2 "Public" Communication

One natural solution to the bilateral communication problem would be to restrict workers from sending multiple (and possibly conflicting) messages by allowing them only to communicate once with a central intermediary. By requiring workers to communicate only once with one party, they cannot simultaneously inform different firms that it is their top choice. Consider the following interview-assignment mechanism which requires each worker to submit a rank-ordered list of their preferences $\tilde{\mathcal{P}}_w$ over firms, and then subsequently provides each firm privately with a subset of workers to interview $\eta(\tilde{\mathcal{P}}_W)$ where $\eta(\cdot)$ is generated as follows:[7]

- In the first round, each firm is assigned all workers who have ranked that firm their first choice for an interview according to $\tilde{\mathcal{P}}_w$. If more than x workers rank a particular firm as their top choice, then x workers are chosen at random from those that do.

- In general, in round t, any firm who has not yet received x interviews accepts all workers who also have not received x interviews and ranks that firm as their t-th highest choice. In the case that accepting all worker who rank that firm their t-th choice results in a firm interviewing more than x candidates, the firm chooses at random enough workers such that the firm ends up with exactly x total candidates.

This procedure terminates with an interview assignment that allocates each worker and each firm exactly x interviews. This mechanism can be seen as a worker-proposing deferred acceptance algorithm for interviews, where firms are indifferent

[7] $\tilde{\mathcal{P}}_W \equiv \{\tilde{\mathcal{P}}_w\}_{\forall w \in W}$ is simply the set of preferences for all workers.

among all workers (thus will accept any interview offer up to quota), and each worker and firm has a quota of x.

Proposition 2.2. *There exists an equilibrium whereby workers report preferences truthfully ($\tilde{\mathcal{P}}_w = \mathcal{P}_w \; \forall \; w$), and firms allocate interviews according to the assignment proposed by the intermediary, $\eta(\tilde{\mathcal{P}}_w)$.*

By construction, this outcome is Pareto optimal for workers among all outcomes that assign each worker and firm x interviews – no two workers could swap firm interviews and each be better off. Furthermore, no firm would not wish to switch a worker w with a worker w' from any other firm: a firm does not know how highly it ranked on w's or w''s preferences, nor does it know the identities of the other firms interviewing the workers; all it knows is that w will receive exactly x interviews in the equilibrium described, and if it did swap, it would receive an interview with a worker w' who prefers that firm less than any of its current interviewers.[8] Thus, the resultant allocation is pairwise stable.[9]

We refer to this particular mechanism as "public" communication because workers must commit to making only one announcement of their rank-ordered preference list and cannot provide conflicting messages to different parties. This is even despite the fact that firms at no point are privy to the announcements, an assumption we make only for technical reasons.[10] Note also that an intermediary in this setting functions also as a coordination device, enabling firms to evenly distribute interviews by giving each worker exactly x.

2.3 Monitoring the Number of Interviews

In many cases however, public communication or a centralized intermediary may be infeasible – the former may be unobservable or unverifiable, and the latter may be too costly to implement. Furthermore, requiring workers to submit complete preference listings may be both burdensome for them to compute or too complex to

[8] If w ranked firm f as its kth ranked firm, any w' who did not get an interview with f must have ranked f as its kth or lower choice (otherwise, w' would have been assigned f as an interview in addition to or in place of w). Since f does not know the identities of the other firms that interview w or w', a firm would weakly prefer interviewing the worker who preferred the firm higher.

[9] A *pairwise stable* match is a matching in which there is no firm and worker pair who are not matched that would prefer to be matched to each other than to their existing partners.

[10] This condition is to avoid "overlap" concerns and ensure the existence of a symmetric pure strategy interview assignment among firms, as we have assumed. If firms could observe the announcements, then the equilibrium assignment $\eta(\mathcal{P}_W)$ may no longer be an equilibrium for firms to follow due to asymmetric overlap conditions for certain announcements of $\mathcal{P}_W$.

communicate when N is large.[11] In this setting, is it still possible to improve upon the no-communication outcome? As long as firms are able to observe the number of interviews any worker obtains during the interview assignment stage – even with private bilateral communication – the answer is affirmative.

Assume each worker w can initially send a signal to any firm f; i.e., $m_{wf} \in \{0, 1\}$, where $m_{wf} = 0$ indicates that w did not send f a signal. Next, during the interview assignment stage, nature selects a random ordering of firms; firms then take turns selecting a worker to interview, only observing the number of interviews each worker has received up to that point and not the identities of the firms interviewing him. Importantly, we assume that only after every firm has stopped assigning interviews do interviews actually occur.

Proposition 2.3. *For some N and $(k_1, \ldots, k_N) \in \{1, \ldots, N\}^N$, there exists an equilibrium in which each worker w_i sends a message to a firm f if and only if f is one of his k_ith highest ranked forms, and a firm assigns an interview to a worker only if he has signalled to it.*

This mechanism only requires that firms be able to observe the number of interviews a worker receives during the interview assignment stage; it places no restrictions on the number of messages a worker receives. Instead, the decision to moderate the number of signals sent is borne by the worker, induced by the following tradeoff from sending an additional signal: an extra signal decreases the probability that a worker will receive less than x interviews; on the other hand, an extra signal increases the probability that a less desirable firm will interview the worker in place of a more desirable firm.

The proposition follows immediately by noting that it holds for $k_i = N \; \forall \; i$ and that for any fixed k_i, if worker w_i signals to a firm f', it will signal to all firms $f \succ f'$ – otherwise he could do better by signalling to the higher ranked firm instead of f'. In general, we cannot say anything specific about the number of signals sent in equilibrium unless we translate worker preferences over firms into cardinal utilities. Nonetheless, in terms of comparative statics, it follows that if workers are almost indifferent between their top and lowest firm, then they will choose to signal more than if the differences between firms were greater.[12]

[11] See Segal (2005) for a discussion on the costs of full preference revelation in two-sided matching markets.

[12] Coles and Niederle (2006) explore the use of limited signals in matching markets, whereby participants can preferences prior to a match. In this setting, preferences are known ex ante, and

Though the message space could be larger without substantially changing the outcome, any informative message (i.e., one that a firm receives and uses to influence its behavior) must be interpreted the same way in order for there to be any notion of truthful revelation – if firms differentially treated the receipt of two different messages, then workers would have incentive only to report the message that maximized its chances of obtaining an interview. Thus, the simple communication protocol used here not only is sufficient, but minimizes the complexity of communication and possible scope for strategic misrepresentation.

A Proofs

Proof of Proposition 2.1. Clearly, there exists an equilibrium where every firm ignores any communication that occurs and plays, and workers send meaningless messages.

We now show that there any equilibrium where communication is utilized is uninformative. Assume that there exists an equilibrium where a firm f "listens" to a signal received by a worker w in that a higher message m_{wf} corresponds to a higher probability that worker w receives an interview. Since a worker cannot be worse off by receiving more interviews (holding the firms with which he does interview fixed) and can in fact be better off (being more likely to be hired by a more preferred firm), all workers will have an incentive to tell each firm that the firm is their highest choice. In turn, if there is a signal that provides a higher probability a worker is interviewed then another signal, then every worker will have an incentive to provide the signal that provides the highest probability of being hired to each firm. The same "message" thus will be sent to every firm by every worker. To fully describe the resultant equilibrium, we let firms randomly select y candidates from all workers who provided the highest signal (e.g., the firm is their top choice), and all workers tell every firm this message. This resultant equilibrium is outcome equivalent to a symmetric mixed strategy equilibrium without communication whereby firms randomly selected y workers to interview. □

Proof of Proposition 2.2. Assume workers report truthfully. Then given the mechanism assigns each firm exactly x interviews with workers who also have x interviews,

thus the interview stage of information acquisition is ignored. However, they are able to obtain some comparative statics indicating that welfare may improve with the introduction of a signalling mechanism.

by assumption 2.1, there exists a Nash equilibrium for the interview assignment stage whereby firms will interview according to $\eta(\mathcal{P}_W)$.

Given firms interview according to $\eta(\mathcal{P}_W)$ and all other workers report truthfully, we now show that any individual worker will not wish to deviate. First, a worker cannot truncate his preferences or omit to include any firm because the mechanism requires a full rank ordering; even if submitting an incomplete list of preferences were to be allowed, a worker still would not wish to do so since truncation would create a positive probability of obtaining less than x interviews but yield no benefit. Second, if a worker prefers firm i to j, he will not report that j is ranked higher than i: since all other workers have preferences symmetrically distributed across firms, and since each firm gives priority for interviews to workers who apply in earlier rounds, the probability of obtaining an interview for a given firm is strictly increasing in the rank it was submitted; thus, a worker would do strictly better by reporting i as higher ranked than j, since he values an interview with firm i more than one with firm j. □

References

Coles, P., and M. Niederle (2006): "Signaling in Matching Markets," mimeo.

Crawford, V. P., and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431–1451.

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

Lee, R. S., and M. Schwarz (2007): "Interviewing in Two-Sided Matching Markets," mimeo.

Ostrovsky, M., and M. Schwarz (2006): "Synchronization under Uncertainty," *International Journal of Economic Theory*, 2(1), 1–16.

Roth, A. E., and M. A. O. Sotomayor (1990): *Two-Sided Matching: A study in game-theoretic modeling and analysis*. Econometric Society Monographs, Cambridge, MA.

Segal, I. (2005): "Communication in Econonomic Mechanisms," in *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*,

ed. by R. Blundell, W. K. Newey, and T. Persson, vol. 1, chap. 6. Cambridge University Press, Cambridge, U.K.

Sönmez, T., and M. U. Ünver (2004): "Course Bidding at Business Schools," *mimeo*.