Interviewing in Two-Sided Matching Markets*

Robin S. Lee† Michael Schwarz‡

November 23, 2007

Abstract

We introduce the *interview assignment problem*, which generalizes the one-to-one matching model of Gale and Shapley (1962) by including a stage of costly information acquisition. Agents do not know their preferences over potential partners unless they choose to conduct costly interviews. Although there may exist many equilibria in which all agents are assigned the same number of interviews, we show the efficiency of the resultant match can vary significantly depending on the degree of *overlap* – the number of common interview partners among agents – exhibited by the interview assignment. Among all such equilibria, the one with the highest degree of overlap yields the highest probability of being matched for any agent. Our analysis is used to motivate new and explain existing coordinating mechanisms prevalent in markets with interviewing.

1 Introduction

The theory of two-sided matching generally assumes that agents know their true preferences over potential partners prior to engaging in a match.[1,2] However, in matching markets ranging from labor markets to marriage markets, information acquisition plays an important role: interviews and dates to learn these preferences are often costly and thus scarce. Since these interviews affect the formation of preferences, the efficiency of the match depends not only on the matching mechanism but also the procedure for assigning interviews.

*We thank Kyna Fong, David McAdams, Michael Ostrovsky and Al Roth for helpful comments. First draft: October 2006.

†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.

[1]For a survey, see Roth and Sotomayor (1990).

[2]Chakraborty, Citanna, and Ostrovsky (2007) is a notable exception in which agents do not know their preferences; unlike the present paper which focuses on learning via costly interviewing, Chakraborty et al. investigates the stability of matching mechanisms with interdependent values over partners.

We generalize the one-to-one matching model of Gale and Shapley (1962) to allow for a stage of costly information acquisition. To our knowledge, this paper is the first to analyze the interview assignment problem in the context of two-sided matching. Throughout this paper we will refer to agents as "firms" and "workers," but note that this label can be changed to men and women, colleges and students, hospitals and doctors, and so forth. Firms and workers do not ex ante know their idiosyncratic preferences over potential matching partners, but instead must discover them through a costly interviewing process. We analyze a two-stage game: in the first stage, firms simultaneously choose a subset of workers to interview and learn their preferences over these workers, and in the second stage firms and workers participate in a one-to-one match using a firm-proposing deferred acceptance algorithm in which firms make "job offers" to workers.[3]

We utilize results pioneered in the one-to-one matching literature for the deferred acceptance subgame, and primarily focus on the first stage of interviewing. Even so, the interview assignment problem is generally difficult and possibly intractable. To allow for analysis while still maintaining a model rich enough to yield meaningful results, we make the following assumptions: firms bear the full cost of interviewing; a firm and worker must interview in order to be matched;[4] workers prefer being matched to any firm than be unemployed; firms may find some workers undesirable and choose to remain unmatched; and workers and firms are ex ante homogenous, with preferences over partners independent and idiosyncratic to each agent.

Even if all firms and workers are ex ante identical (prior to the realization of their idiosyncratic preferences), agents are not indifferent over whom they interview with. Since interviews are costly, firms care about how many interviews a potential interviewee has: as the number of interviews a worker has increases, the probability a job offer being accepted declines as the worker might obtain and accept an offer from another firm. Thus, all else being equal, workers who have few interviews are more attractive to interview because they are more likely to accept if an offer is made.

However, we also investigate a more subtle form of coordination that also is

[3] In a companion paper, Lee and Schwarz (2007) consider the possibility that workers initially know their own preferences, and examine mechanisms which allow workers to signal their preferences prior to the assignment of interviews.

[4] For example, the National Residency Matching Program is a prominent example of a market between hospitals and medical school graduates which utilizes a centralized match (Roth (1984)). Hospitals rarely if ever rank students whom they do not interview.

important. Although there may exist many equilibria in which all agents conduct the same number of interviews, the efficiency of the match can be very different depending on the degree of *overlap* – the number of common interviewees among firms – exhibited by the interview assignment. Consider two firms f and f' who are the only firms that interview workers w and w': if firm f has an offer rejected by worker w, it must be that the worker accepted an offer from firm f'; consequently, firm f will then face no "competition" for worker w' and obtain him for certain if it made him an offer. If firms f and f' did not interview the same set of workers, then firm f could possibly be rejected by both w and w' and not be matched despite making offers to both workers (since being rejected by w no longer implies obtaining w' for certain). Thus, a firm's expected payoffs depends not only on the number of interviews its workers receive, but also the identities of the firms interviewing those workers. In general, this paper shows that among equilibrium interview assignments in which all workers and firms obtain exactly the same number of interviews, the assignment which exhibits the highest degree of overlap yields the highest probability of employment for any agent.

The interview assignment problem can be seen as a many-to-many assignment problem since firms may be assigned to many workers and workers to many firms in the interview stage. However, as firms care about the identities of other firms who interview its candidates, there are externalities imposed on agents not directly involved in a particular pairwise match. Our setting thus does not fit into the standard many-to-many matching framework; instead of relying on non-equilibrium concepts such as pairwise stability often employed in that literature, we utilize standard non-cooperative equilibrium conditions when analyzing interview assignments.

Our paper is closely related to the simultaneous search model in Chade and Smith (2006), which considers a problem faced by a single decision maker who must choose a portfolio of ranked stochastic options. Whereas in their model the probability of obtaining a particular option is assumed to be given, our paper endogenizes the probability that selecting a particular worker for an interview leads to a match, both as a function of other firms' actions and the outcome of the second-stage deferred acceptance algorithm.[5]

[5]Both our paper and Chade and Smith (2006) are significantly different from the literature on costly *sequential* search. E.g., Shimer and Smith (2000) and Atakan (2006) have added frictions to decentralized sequential search and matching economies such as the one proposed in Becker (1973) in order to test the robustness of assortative matching; Lien (2006) provides an example in which the assignment of interviews in sequential search markets may be non-assortative.

There are also parallels to the literature on information acquisition in mechanism design: firms (bidders) here must interview (invest) in workers (objects) in order to learn their private values over workers, and an firm's incentive to learn its valuation for a particular worker is reduced when others choose also to interview.[6] However, instead of focusing on environments with a single seller and multiple buyers as would be the case in an auction environment, we consider a matching market between many buyers and sellers. In a sense, the interview stage becomes similar to a bipartite network formation model in which one side of the market (the firms) unilaterally decide which links (interviews) to form, and total payoffs depend on the total network which is created (i.e., there are significant externalities across agents).[7]

Furthermore, we wish to emphasize that the use of the second stage "match" is only an approximation for the dynamics of hiring processes in a variety of industries and settings.[8] In some situations that utilize a centralized match such as the National Residency Matching Program, the relationship is quite exact; in others, a decentralized matching market may still be modelled as a deferred acceptance procedure. Whenever preferences are ex ante unknown and need be revealed through a costly interview, due diligence, or even dating process, our analysis remains relevant.

2 Model

2.1 Setup and Definitions

There are N workers and N firms, represented by the sets $W = \{w_1, \ldots, w_N\}$ and $F = \{f_1, \ldots, f_N\}$. Each worker w has a strict preference orderings over firms P_w. If firm f hires worker w, it realizes a firm specific surplus $\delta_{w,f} \in \mathbb{R}$. If a firm does not hire a worker, it receives a reservation utility of $\underline{\delta}$ which we will assume to be 0. A worker can only work for one firm and a firm can only hire one worker; we refer to this hiring decision as a *match* between a firm and a worker.

The main innovation of our model is that $\{P_w\}_{w \in W}$ and $\{\delta_{w,f}\}_{w \in W, f \in F}$ are

[6]See Bergemann and Välimäki (2005) for a survey.

[7]See Jackson (2004) for a survey. Kranton and Minehart (2001) study a similar network formation game between buyers and sellers in which sellers have only one good to sell, may only trade with buyers with whom they have formed a link, and buyers receive a random draw over their valuation of a good.

[8]We choose to abstract away from wage negotiations and assume that such wages are fixed or already embedded in user preferences across firms, or there are no wages as in dating markets. Indeed, many jobs provide the same salary to all workers in entry level positions, despite relative differences in quality of workers.

unknown ex ante prior to a match, and can only be revealed through a costly interview process. Firms and workers are allowed to conduct multiple interviews, but each interview costs a fixed amount $c \in \mathbb{R}^+$ and all costs are borne by the firm conducting it. When a firm f interviews worker w, it learns the value of $\delta_{w,f}$. We assume $\{\delta_{w,f}\}_{w \in W, f \in F}$ comprises i.i.d draws from the distribution H, where H has finite first and second moments (so that all order statistics have finite expectations), continuous density h, and $\int x dH(x) \leq \underline{\delta}$. This last condition ensures that a firm would (weakly) prefer not to hire a worker it has not interviewed. Finally, we impose one further condition

$$E_\delta[\delta - y|\delta > y] - E_\delta[\delta - y'|\delta > y'] \leq 0 \; \forall \; y > y' \geq 0 \tag{2.1}$$

which states that if δ is distributed according to H, then the expected value of $\delta - y$ given $\delta > y$ (weakly) falls as y increases.

Worker preferences are distributed uniformly over firms – i.e., for any two firms f and f', a given worker has as likely a chance of preferring f to f', and vice versa – and workers always prefer working for any firm than remaining unemployed. If a worker interviews with a subset of firms $F_w \subset F$, then the worker will realize his relative rankings over only those firms $f \in F_w$. Finally, since a firm will never make a job offer to a worker whom it never interviewed, how a worker w ranks a firm $f' \notin F_w$ is irrelevant.

2.2 Timing and Description of Game

The timing of the interview and matching game is as follows:

(1) In the first stage, each firm f chooses a set of workers $W_f \subset W$ to interview and bears an interview cost $c|W_f|$.[9] These choices define an *interview assignment* η, a correspondence from the set $F \cup W$ into itself such that $f \in \eta(w)$ if and only if $w \in \eta(f)$. Thus $\eta(f) \equiv W_f \subset W$ represents the workers interviewed by firm f under η, and $\eta(w) \equiv F_w \subset F$ represents the set of firms that interview worker w. Each firm privately realizes $\{\delta_{w,f}\}_{w \in W_f}$ and each worker privately forms preferences over the firms it interviews with. Although each

[9] Since interviewing is costless from a worker's perspective, it is strictly in his best interest to maximize the number of interviews he receives. To see this, note that a worker will not receive a job offer unless he is interviewed. The more interviews a worker has, the more likely firms will receive favorable draws on his quality, and thus the more job offers that worker will receive.

firm observes the entire interview assignment η, each worker only observes the set of firms with whom he interviews, $\eta(w)$.

(2) In the second stage, firms and workers engage in a firm-proposing deferred acceptance algorithm for employment, analyzed by Gale Shapley (1962). In this algorithm, each firm f reports preferences $\tilde{P}_f$ and each worker w reports preferences $\tilde{P}_w$. The algorithm proceeds as follows:

   – Step 1: Each firm makes a job offer to its first choice worker (or, if all interviews yielded negative draws on δ, does not make any offers). Each worker who receives an offer "holds" onto its most preferred offer and rejects the rest.

   – In general, at step t: Each firm who was rejected in step $t-1$ makes a job offer to the most preferred and acceptable worker who has not yet rejected it. Each worker who receives an offer compares all offers received (including an offer he may be holding from a previous round), holds onto his most preferred offer, and rejects the rest.

   The algorithm stops after a step when no firm's offer is rejected; at this point all firms have either a job offer that is currently being held or has no workers it wishes to make an offer to that has not already rejected it. Any worker who is holding a job offer from a firm is hired by that firm (an event we also refer to as the worker *accepting* an offer), and any worker who does not have a job offer remains unemployed. This algorithm yields a one-to-one matching μ which is a one-to-one correspondence from $F \cup W$ onto itself such that (i) $\mu^2(x) = x$, (ii) if $\mu(f) \neq f$ then $\mu(f) \in W$, and (iii) if $\mu(w) \neq w$ then $\mu(w) \in F$. We say worker w is hired by firm f if $\mu(w) = f$, and worker w is unemployed if $\mu(w) = w$. Similarly, we say firm f hires worker w if $\mu(f) = w$ and firm f does not hire anyone if $\mu(f) = f$.

The firm-proposing deferred acceptance algorithm utilized in the final job match results in what is referred to as the *firm optimal stable matching* (FOSM) for utilized preferences.[10] We utilize this particular procedure and outcome as a reasonable

[10] See Gale and Shapley (1962), Roth and Sotomayor (1990). A *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. *Firm optimal* means that no firm can do better (match with a more preferred worker) in another stable matching than in the FOSM, according to the preferences used.

approximation for the outcome of the hiring procedure, even in decentralized hiring markets.

These types of matching mechanisms can be susceptible to "gaming" in that participants may find it preferable to misrepresent their true preferences. However, as long as workers prefer being employed over being unemployed strongly enough, in an equilibrium both sides will use their preferences realized during the interview stage for the job match: for each f, $\tilde{P}_f$ will rank workers in descending order according to the realized values of $\{\delta_{w,f}\}_{w \in W_f}$, and any worker who was not interviewed or was found to have a negative $\delta_{w,f}$ are considered unacceptable matches; for each w, $\tilde{P}_w$ will truthfully rank any two firms it interviewed with according with true preferences P_w.

Lemma 2.1. *Let $f_{i(k)}$ represent worker i's k-th ranked firm, and let his utility from being employed by a firm given by $u_i(f)$. Denote $u_i(\emptyset)$ the utility to worker i from being unemployed. There exists a $\beta > 0$ such that if*

$$u_i(f_{i(N)}) - u_i(\emptyset) > \beta(u_i(f_{i(1)}) - u_i(f_{i(N)}) \; \forall \, i \in W \tag{2.2}$$

it is an equilibrium for both workers and firms to use their true preferences when conducting the firm-proposing deferred acceptance algorithm.

(All proofs are located in the appendix.) Since we utilize a firm proposing deferred acceptance algorithm, it is a dominant strategy for firms to use their true preferences (Dubins and Freedman (1981), Roth (1982)). The fact that workers also use their true preferences may seem surprising in light of the negative existence results in the two-sided matching literature of a mechanism which elicits truthful reporting from both sides. However, since preferences are independently drawn and workers do not observe the entire interview assignment, each worker perceives the probability of receiving a job offer to be the same for any firm. Thus, a worker will not wish to swap the ordering of any two firms in his reported preferences. Furthermore, as long as each worker places a high enough disutility of being unemployed (condition (2.2)), no worker will reject any firm that makes him an offer (i.e., rank a firm as unacceptable in his reported preferences). For our analysis, we assume (2.2) holds.[11]

[11]This is the only part of our analysis that relies on cardinal utilities; as long as remaining unemployed is sufficiently unattractive for any worker, the analysis proceeds relying only on ordinal utilities.

3 Interview Assignment

Since the behavior of agents in the matching stage is well-characterized, we now turn to analyzing the decisions of firms during the interview stage. We are interested in "symmetric" equilibria in which all firms interview the same number of workers. However, even with this restriction, there are still several equilibrium outcomes which differ in the total number and distribution of interviews conducted. The expected number of unemployed workers or the costs expended on interviewing can vary and depend on the equilibrium chosen.

3.1 Firm's Expected Utility

Since we have shown that firms and workers report preferences honestly in an equilibrium of the second stage matching process, a firm's expected utility from interviewing any subset of workers W_f given the actions of other firms W_{-f} can be computed.

For illustrative purposes, consider the expected utility of a firm from interviewing one worker w:

$$EU_f(\{w\}, W_{-f}) = \underbrace{Pr(\delta_{w,f} \geq 0)}_{(1)} \underbrace{E[\delta_{w,f} | \delta_{w,f} \geq 0]}_{(2)} \underbrace{Pr(f \succ_w f' \; \forall \; f' \subset \overline{F}_w | f \in \overline{F}_w)}_{(3)} - c$$

where $\overline{F}_w$ denotes the set of firms that make a job offer to worker w given all other firms interview the subsets of workers W_{-f}.[12] The expected utility can be separated into three parts: (1) the probability that a job offer is made to the worker at some stage of the job matching process (which here, due to only interviewing one worker, is equivalent to the probability that the firm receives a positive draw on $\delta_{w,f}$), (2) the expected surplus this worker will provide contingent on being hired, and (3) the probability the worker accepts this offer from the firm given that the firm makes an offer (equivalent to the probability the worker prefers the firm to all other firms who make him an offer). Notice conditional on being made a job offer from firm f, a worker's $\delta_{w,f}$ is independent of his probability of actually accepting the offer – the latter is a function of his other $\delta_{w\cdot}$ draws with other firms and his own preferences,

[12] When we say a firm f makes a job offer to a worker w, we are referring to the event that during *any* stage of the deferred acceptance algorithm, firm f finds itself proposing to worker w; this definition is independent of whether worker w rejects the offer, holds onto it, or ultimately accepts it.

both of which are independent of $\delta_{w,f}$. Thus, the expected value of $\delta_{w,f}$ conditional on being hired is simply the expected value of $\delta_{w,f}$ conditional on being made an offer, which corresponds to (2).

If a firm decides to interview K workers, it is equivalent to taking K "draws" on δ. The realization of the jth highest δ draw is itself a random variable, known as the jth order statistic which we denote by $\delta_{j:K}$. By the logic of the deferred acceptance algorithm, we can then construct the expected utility from interviewing K workers as the expected surplus from hiring the top worker of K interviews times the probability of hiring him, plus the expected utility from hiring the 2nd highest worker times the probability of losing the highest worker times the probability of hiring the 2nd highest worker, and so forth. Formally then, a firm's expected utility from interviewing the subset W_f:

$$\begin{aligned} EU(W_f, W_{-f}) &= \Lambda_{K,K}\overline{P}_{(K)} + \Lambda_{K-1,K}(1-\overline{P}_{(K)})\overline{P}_{(K-1)} + \\ &\quad \ldots + \Lambda_{2,K}\overline{P}_{(2)}\prod_{i=3}^{K}(1-\overline{P}_{(i)}) + \Lambda_{1,K}\overline{P}_{(1)}\prod_{i=2}^{K}(1-\overline{P}_{(i)}) - cK \end{aligned} \tag{3.1}$$

where $K = |W_f|$, $\Lambda_{j,K} = Pr(\delta_{j:K} \geq 0)E[\delta_{j:K}|\delta_{j:K} \geq 0]$ is the expected value of the jth highest worker interviewed conditional on him being a desirable hire times the probability he is a desirable hire (equivalent to (1) and (2) in the single worker example), and $\overline{P}_{(j)}$ represents the probability that firm f "wins" its jth highest worker conditional on making him an offer – i.e., firm f was rejected by all workers which would yield higher surplus, and the worker prefers f over any other firm that makes him an offer (equivalent to (3) in the single worker example). The probability a firm eventually is matched to any worker is simply equation (3.1) with $Pr(\delta_{j:K} \geq 0)$ replacing $\Lambda_{j,K}$:

$$Pr(\mu(f) \neq f|W_f, W_{-f}) = \sum_{j=1}^{K} Pr(\delta_{j:K} \geq 0)\overline{P}_{(j)} \prod_{i=j+1}^{K}(1-\overline{P}_{(i)}) \tag{3.2}$$

The probabilities $\overline{P}_{(j)}$ are a function of the other firms' actions W_{-f}, and may be difficult to compute. However, one observation that aids analysis is that from a firm's perspective, any worker's preferences are randomly generated uniformly over all the firms that interview him; consequently, if n firms make a job offer to a worker at any point during the deferred acceptance stage, each firm considers itself to have a $\frac{1}{n}$ probability of being the firm that the worker accepts (i.e., of being the highest ranked firm for that worker). Thus, sufficient for determining $\overline{P}_{(j)}$ is simply the probability distribution over the number of firms that "compete" by making an offer to the jth ranked worker.

Let P_i^j indicates the probability that when a firm makes an offer to its jth highest worker, i other firms also make that worker a job offer. Then it follows:

$$\overline{P}_{(j)} = \sum_{i=0}^{N} \frac{1}{i+1} P_i^j$$

The following example illustrates how this symmetry can be used to compute expected utilities for firms:

Example 3.1. *Let $N = 4$, and index firms by $\{A, B, C, D\}$ and workers by $\{1, 2, 3, 4\}$. Consider the following interview assignment η:*

$$\eta(A) = \{1, 2\} \qquad \eta(B) = \{2, 3\} \tag{3.3}$$

$$\eta(C) = \{3, 4\} \qquad \eta(D) = \{1, 4\} \tag{3.4}$$

Assume $\delta = 1$ with probability .9 and $\delta = -10$ with probability .1. This corresponds to the case where a worker is most likely to generate positive surplus, but there is a slight chance that he may be very costly to a firm.

Since all firms have a symmetric interview assignments, any firm's profits can be expressed using (3.1) *with the same values for each P_i^j:*

$$\pi = \Lambda_{2,2} \underbrace{(P_0^2 + \frac{1}{2}P_1^2)}_{\overline{P}_{(2)}} + \Lambda_{1,2} \underbrace{(1 - P_0^2 - \frac{1}{2}P_1^2)}_{1-\overline{P}_{(2)}} \underbrace{(P_0^1 + \frac{1}{2}P_1^1)}_{\overline{P}_{(1)}} - 2c \tag{3.5}$$

where the first component is the expected gain times the probability of hiring the most preferred worker, and the second component is the expected gain times the probability of hiring the second most preferred worker (given it lost the first choice worker). Since $E[\delta|\delta \geq 0] = 1$, we have $\Lambda_{2,2} = .99$ and $\Lambda_{1,2} = .81$.

Consider firm A. Without loss of generality, assume firm A's top worker is *worker 1. Then the probability firm A faces competition for worker 1 from D is:*

$$P_1^2 = \underbrace{\frac{1}{2}.99}_{(1)} + \underbrace{\frac{1}{2}.81[\frac{P_1^2}{2}]}_{(2)}$$

where (1) is the probability that D's top worker is also worker 1, and it receives a positive draw on worker 1's quality, and (2) is the probability that D's top choice worker is worker 4 but it loses out to firm C, and then subsequently makes an offer to worker 1. Firm D can only lose worker 4 if C competes for the same worker, which in turn is the very same probability P_1^2.

Next, assume firm A lost its top worker 1 and now is evaluating its competition for its next best worker 2. Again, similar logic allows us to calculate the probability of competition:

$$P_1^1 = \underbrace{\frac{1}{2}.99}_{(1)} + \underbrace{\frac{1}{2}.81[\frac{1}{2}\frac{.99}{2}]}_{(2)}$$

where (1) is the same as before, but (2) is now slightly different. Now if B's top worker is worker 3, then B would only lose worker 3 if C also competed for worker 3. However, since A could only have lost 1 if D employed 1, C faces no competition for worker 4 (or must have already won 4), and thus B will lose worker 3 only if worker 3 is C's top choice and C receives a positive draw on 3.

Noting $P_0^j = 1 - P_1^j$ *for* $j = 1, 2$, $\Lambda_{2,2} = .99$ *and* $\Lambda_{1,2} = .81$, *we can solve* (3.5) *and find a firm's expected profits* $\pi \approx .86 - 2c$. *Thus, if a worker can generate \$100K surplus for a firm or lose \$1M, a firm will obtain in expectation approximately \$86K minus the cost of two interviews.*

Furthermore, the probability that a firm remains unmatched is

$$Pr(\delta_{2:2} < 0) + Pr(\delta_{1:2} > 0)[(P_1^2\frac{1}{2})(P_1^1\frac{1}{2})] + Pr(\delta_{2:2} > 0 \& \delta_{1:2} < 0)[P_1^2\frac{1}{2}] \approx .14$$

In Appendix A, we show how this example's intuition generalizes to calculate expected utilities for other interview assignments.

3.2 Equilibrium Analysis

Having defined each firm's expected utility from an interview assignment η, we now turn to defining what it means for η to be an equilibrium interview assignment.

Formally, a firm's strategy during the interview assignment stage is a probability measure ν_f over the powerset of all workers $\mathcal{P}(W)$. A strategy profile $\nu \equiv \{\nu_f\}_{f \in F}$ is a Nash Equilibrium of this game iff

$$\int_{\nu_f} \int_{\nu_{-f}} EU_f(W_f, W_{-f}) \geq \int_{\nu'_f} \int_{\nu_{-f}} EU_f(W_f, W_{-f}) \; \forall \; \nu'_f, f$$

Any mention of equilibrium refers to the solution concept of *subgame perfect Nash equilibrium.*

We say firm f interviews x workers if $\nu_f(W_f) > 0$ iff $|W_f| = x$; we say firm f interviews y workers at random if $\nu_f(W_f) > 0$ iff $|W_f| = y$, and $|W_f| = |W'_f|$ for any f, f' implies $\nu_f(W_f) = \nu_f(W'_f)$. A pure strategy for a firm simply assigned probability 1 to one particular element $W_f \in \mathcal{P}(W)$. Finally, if there is a pure strategy equilibrium in which each firm f interviews the subset of workers W_f, we say the correspondence η is an equilibrium interview assignment if $\eta(f) = W_f \; \forall \; f$.

A natural candidate for a symmetric equilibrium would be if each firm randomly selects y workers to interview. For certain values of c, an equilibrium in which firms randomize exists:

Proposition 3.1. *For any $y \in \{0, ..., N\}$, there exists $c > 0$ such there is an equilibrium in which each firm interviews y workers at random.*

A mixed strategy equilibrium seems a reasonable outcome if firms are unable to monitor how many interviews a worker receives, and if they are unable to coordinate with other firms on which workers to interview. Indeed, since the outcome of this mixed-strategy equilibrium is a distribution of interview assignments across workers, certain firms ex post would have been better off had they been able to coordinate and not compete excessively for the over-popular (but no better) candidates.

An alternative would be if firms could coordinate and select a single subset of workers such that every worker and firm received the same number of interviews. Example 3.1 illustrated such an assignment for $N = 4$. Again, this too may be an equilibrium:

Proposition 3.2. *For any $x \in \{0, ..., N\}$, there exists a $c > 0$ such that there exists a symmetric pure-strategy equilibrium interview assignment η in which each worker and each firm receives exactly x interviews.*

This and the previous existence proof relies on the result established in lemma B.1 that conditional on other firms utilizing a particular strategy, a given firm's utility from interviewing an additional worker is decreasing in the number of workers it is already interviewing. That is, a firm gains more from the xth interview it conducts (holding everyone else's actions fixed at interviewing x workers) than it gains from the $x+1$th. Thus, if the cost of interviewing is less than the gain from interviewing the xth worker but greater than the gain from interviewing the $x+1$th worker for a firm, every firm interviewing x workers will be an equilibrium as it will not wish to add, remove, or replace any workers in its set of interviewees.[13]

Unlike in the mixed strategy case, implicit in the construction of a pure strategy equilibrium is a means for firms to somehow distinguish subsets of workers when they are of the same size – i.e., a firm must be able to differentiate W_f from W'_f whenever $|W_f| = |W'_f|$. Furthermore, it also requires a great deal of coordination among firms in terms of exactly how to partition the space of workers or which particular equilibrium to play; for any $x < N$, there are at least $N!$ different symmetric equilibrium in which x interviews are conducted by each firm and x interviews are received by each worker. As a consequence, firms need not only to be able to identify which workers to interview in a particular pure strategy equilibrium, but also need to coordinate with all other firms which particular pure strategy equilibrium to play. If firms are able to coordinate, the following example shows they can achieve a better outcome in a pure strategy equilibrium than mixed:

Example 3.2. *Consider $N = 3$ and index firms by $\{A, B, C\}$ and workers by $\{1, 2, 3\}$. Consider the following interview assignment:*

$$\eta(A) = \{1, 2\} \qquad \eta(B) = \{2, 3\}$$
$$\eta(C) = \{3, 1\}$$

Following the same type of calculations as in example 3.1, each firm's expected

[13]Due to integer constraints, there may exist values of c for which no symmetric equilibrium exists. To see why, consider the mixed-strategy case. Assume that no firm interviews any worker, and let G denote the gain from a firm deviating and randomly interviewing one worker. Let G' represent the gain from interviewing one worker when every other firm also interviews one worker at random. Clearly $G' < G$ since the gain to interviewing a worker falls when other firms may also interview that worker. Thus, as long as $c \in (G', G)$, no symmetric mixed-strategy equilibrium exists – neither everyone interviewing no workers nor everyone interviewing one worker is an equilibrium (and as arguments in the proof of the previous proposition can show, everyone interviewing more than one worker in not an equilibrium either).

profits is $\pi \approx .88 - 2c$ *and the probability of being unmatched is approximately* $.12$.

However, now consider the case where each firm now randomly select 2 *workers to interview. From a given firm's perspective, there are now several possible interview assignments – e.g., the most preferred worker has 3 interviews, and the least preferred has 0 other interviews; the most preferred worker has 2 interviews, and the least preferred has 2; and so on. For each case, it is possible to compute precisely the expected profits and probabilities of being unmatched. We find that* $\pi = .84 - 2c$ *and the probability of being unmatched is approximately* $.16$.

In the following section, we compare outcomes of the different equilibria described here.

3.3 Overlap

It is not surprising that the inability to coordinate on a pure-strategy equilibrium as opposed to playing a mixed strategy equilibrium can lead to efficiency losses. However, this is not the only form of coordination that can be achieved by firms in order to improve outcomes. It turns out that a firm cares not only about the number of interviews its interviewees are already receiving, but the identities of those firms that its interviewees are interviewing with.[14] Indeed, the construction of the previous pure-strategy equilibria took this into account: each firm received a symmetric subset of workers – symmetric not only in the number of interviews each worker received, but also the type of firms that were already interviewing the worker.

Why does the identity of other firms matter? Consider the decision of firm f choosing to interview an additional candidate when it is already interviewing worker w. Firm f can choose between workers w' and worker w'' who each already have the same number of interviews, except w' also happens to be interviewing with the same firms interviewing worker w, whereas worker w'' is not – and thus we say worker w' exhibits *overlap* with worker w since they have interviewers in common. It turns out, the distinction between worker w' and w'' is not trivial – a firm f will strictly prefer to interview worker w'. This is due to the fact that if firm f loses its first choice worker (be it w or w') to a firm f', then firm f will face less "competition" among firms for its second choice worker since f' no longer needs to match. This

[14] In addition, a firm cares about the identities of the firms who interview the workers who are interviewed by the firms who interview the same set of workers, and so on and so forth.

generalizes naturally as well: if firm f's candidates all overlap with the same other firms, then it means that for every worker who rejects f's job offer, effectively one less firm is then "competing" for its next highest ranked worker.

For the purposes of our analysis, there one specific type of overlap that is focal:

Definition 1. *An interview assignment η that assigns x interviews to each firm and worker exhibits perfect overlap if if $\eta(f) \cap \eta(f') \neq \emptyset$ implies $\eta(f) = \eta(f')$ $\forall$ f, f'.*

Although perhaps subtle, the existence of greater overlap can have dramatic effects.

Example 3.3. *Recall in example 3.1 that the probability a firm is unmatched was approximately $.14$, and a firms' expected profits was approximately $.86 - 2c$.*

Now take the setup of example 3.1, but we now assume that

$$\eta(A) = \{1, 2\} \qquad \eta(B) = \{1, 2\} \qquad (3.6)$$
$$\eta(C) = \{3, 4\} \qquad \eta(D) = \{3, 4\}$$

such that there is perfect overlap. Now if both 1 and 2 are acceptable workers for A, then A is guaranteed to hire at least one of them with certainty: if A loses its top choice worker, it means B hired 1 and there no longer is competition for worker 2. It then follows that $P_0^1 = 1$ (the probability of facing no competition for the second best worker, given the first best worker rejected the firm). Additionally, the probability A's top worker receives another job offer is simply $P_1^2 = \frac{1}{2}Pr(\delta_{2:2} > 0)$, which is the probability that a B's top choice worker coincides with A's top choice. We thus find

$$\pi = \Lambda_{2,2}(P_0^2 + \frac{1}{2}P_1^2) + \Lambda_{1,2}(1 - P_0^2 - \frac{1}{2}P_1^2) - 2c \approx .95 - 2c$$

Furthemore, the probability of remaining unmatched is now

$$Pr(\delta_{2:2} < 0) + Pr(\delta_{2:2} > 0 \& \delta_{1:2} < 0)(\frac{1}{2}P1) \approx .05$$

Hence, we see that with overlap, the probability that any worker or firm is unmatched is drastically reduced, and that a firm generates in expectation greater surplus from the same number of interviews – an increase of over 10%.

Indeed, an interview assignment with no overlap as depicted in example 3.1 is an equilibrium for firms to follow for $c \in (.14, .23)$.[15] *On the other hand, as long as*

[15] To see why, observe that interviewing an additional worker for any firm can yield at most a gain

$c \in (.5, .26)$, *the interview assignment depicted here where firms interview* 2 *workers with perfect overlap is an equilibrium. Consequently, for any value of* $c \in (.14, .23)$*, both interview assignments* (3.3) *and* (3.6) *are equilibria, but the latter equilibrium dominates.*

Thus, there may be many different pure strategy equilibria that still assign each firm and each worker x interviews, but exhibit different degrees of overlap among firms. However with higher degrees of overlap, (1) the greater the probability that at any stage of the deferred acceptance algorithm a worker will accept a firm's offer, (2) the more likely a firm will receive a higher-ranked worker in the match, and (3) each firm is less likely to remain unmatched and have all of its acceptable candidates reject its offers. We can show that a symmetric perfect-overlap equilibrium will outperform any other symmetric equilibrium (including one in mixed strategies) if firms interview the same number of workers.

Proposition 3.3. *Consider an interview assignment* η *which assigns each firm and each worker exactly* x *interviews with perfect overlap. There is no other interview assignment* η' *in which firms each receive exactly* x *interviews such that either every firm strictly receives higher utility or every firm is matched with strictly higher probability. Furthermore, if all firms instead chose* x *workers at random, they would also be strictly worse off.*

Again, due to integer constraints, for a given c and N, a symmetric pure-strategy with perfect overlap may not exist: the construction of the equilibria is sensitive to the relationship of N versus x, where x is the number of interviews per firm in equilibrium. However, as we show in Appendix C, as N grows large there exists a correlated equilibrium in which each firm achieves perfect overlap of interviews with probability close to 1.

Finally, we have only demonstrated that a perfect overlap pure strategy equilibrium is more efficient than any equilibrium in which firms conduct the same number of interviews. However, we have not discussed whether the inability to coordinate on a pure strategy equilibrium as opposed to a mixed equilibrium would result in a greater or lower quantity of interviewing. It turns out, such a comparison is not

in expected utility of $1 - .86 = .14$, and thus a firm will not interview an additional worker if $c > .14$. Furthermore, if a firm drops a worker, its expected utility is now $EU = \Lambda_{1,1}(P_0^2 + \frac{1}{2}P_1^2) \approx .62$. Thus, interviewing 2 workers instead of 1 yields an expected gain of approximately .24; if the cost of interviewing a worker is less than .24, no firm will choose to drop any worker. It is also straightforward to see why a firm would not want to switch which workers it interviews.

possible in general. In Appendix D, we provide an example in which a mixed strategy equilibrium can result in more or less interviews conducted than a pure strategy equilibrium.

4 Discussion

The ability to limit interviews for candidates who already have several and grant more to others who have few can assist firms in coordinating the allocation of interviews. It is thus not surprising that many institutions observed in practice aid in this regard. With on-campus recruiting at colleges, interviews are conducted on only a limited number of days thereby making time considerations a limiting factor on the number of interviews any given candidate can feasibly conduct; in academic job markets, placement officers aid in identifying candidates who have not received many interviews. Additionally, as studied in Mongell and Roth (1991), the "rush" system by which sororities on college campuses recruit new members can be seen as a way of limiting the number of interviews a potential candidate may receive, and equalizing the number of interviews conducted by each sorority.[16]

Note also that if the costs of interviewing are sufficiently low such that the number of interviews x in any equilibrium is close to N, the differences between equilibria – mixing versus pure strategy, or varying degrees of overlap within pure strategy – become less and less pronounced. Indeed, at the extreme if $x = N$, all equilibria coincide with the same interview assignment in which every firm interviews every worker. Thus, if the population of N can be divided into smaller subgroups in which agents in each group can only interview other agents in that group, the inability to coordinate on a pure strategy equilibrium becomes less problematic – i.e., it is equivalent to mixing in an environment where x is close to N. Consequently, by partitioning the population, the probability of overlap is increased and the variance of interviews each worker receives is reduced.

In this light, some of the institutions for improving overlap are the creation of specialized fields or job divisions for certain positions, even if responsibilities and requirements do not differ. For example, universities may choose to interview candidates only within a specific field of a discipline as opposed to across fields within a given year in order to maximize overlap. Furthermore, segmentation via

[16] E.g., "[a] rushee who receives more invitations than the number of parties permitted in a given round must decline, or 'regret,' the excess invitations," (Mongell and Roth, 1991).

geography (e.g., firms interviewing only local candidates) contribute to encouraging overlap as well.

5 Concluding Remarks

Our paper introduced the interview assignment problem and provided a model for analysis when information acquisition is costly. We illustrated two distinct forms of miscoordination in the assignment of interviews – workers may receive varying numbers of interviews and firms may not efficiently overlap their interviews – and explored results in a simple stylized environment. There are a number of directions for future research that are beyond the scope of the paper. Possibilities include extending the model to include additional features of real world interview environments,[17] the social planner's problem and the calculation of the first best optimal assignment of interviews, and the assignment of interviews if they are allocated in a sequential process.

A Equilibrium Analysis

A.1 Symmetric Pure Strategy

In any symmetric pure strategy interview assignment, firms not only interview the same number of workers, but also the same "types" of workers – i.e., all workers have the same probability of having any number of total interviews, and all workers share the same degree of overlap. For this section, we consider a symmetric interview assignment η in which each firm and each worker conducts exactly K interviews.

Let P_i^j indicate the probability that when a firm makes an offer to its jth highest ranked worker, i other firms also make that worker an offer. Let $\overline{P}_{(j)}$ indicate the probability that a firm obtains its jth highest worker given it makes that worker an offer. (These probabilities are all conditional on having been rejected by all workers ranked higher than j.) Since worker preferences over firms are random but uniform and symmetric, for any set of firms that make an offer to a worker, each firm has an equal chance of being a particular worker's highest ranked firm. Thus:

$$\overline{P}_{(j)} = \sum_{i=0}^{K-1} \frac{1}{i+1} P_i^j$$

[17] E.g., wages, heterogenous agents, allowing firms to hire more than one worker, and the sharing of interviewing costs.

Recall equation (3.1) which we restate here:

$$EU(W_f, W_{-f}) = \Lambda_{K,K}\overline{P}_{(K)} + \Lambda_{K-1,K}(1-\overline{P}_{(K)})\overline{P}_{(K-1)} + \ldots + \Lambda_{2,K}\overline{P}_{(2)}\prod_{i=3}^{K}(1-\overline{P}_{(i)}) + \Lambda_{1,K}\overline{P}_{(1)}\prod_{i=2}^{K}(1-\overline{P}_{(i)}) - cK$$

where $K = |W_f|$, $\Lambda_{j,K} = Pr(\delta_{j:K} \geq 0)E[\delta_{j:K}|\delta_{j:K} \geq 0]$

A.1.1 Perfect Overlap

In the special case of perfect overlap, we can explicitly calculate the expected utilities and probabilities that a firm will remain unmatched. Consider the assignment η in which each firm and worker receives exactly K interviews with perfect overlap. We wish to characterize $\overline{P}_{(k)}$ for all $k \leq K$, where $\overline{P}_{(k)}$ is the probability a firm's kth highest ranked worker accepts a job offer *conditional on the firm having been rejected by all higher ranked workers.* Consider firm f and denote the workers it interviews under η as $w_K, \ldots, w_1$ in decreasing order of preference.

Consider $\overline{P}_{(1)}$. Clearly $\overline{P}_{(1)} = 1$, since if a firm was rejected by all $K-1$ higher ranked workers, this means that there is no other firm with a job offer extended to w_1 and firm f obtains him with certainty (conditional on making him an offer).

Now consider $\overline{P}_{(2)}$. If a firm f is considering making an offer to its 2nd least ranked worker w_2, it means that $K-2$ firms and workers have already been matched.[18] Consequently, there is at most one other firm f' who has not yet been matched and whose only attainable workers are w_1 and w_2. Since all draws on δ are i.i.d., it now follows that f will face competition for w_2 if and only if f''s highest ranked worker *of those remaining* is desirable and is w_2 – i.e., $P_1^2 = \frac{1}{2}Pr(\delta_{2:2} \geq 0)$. If not, then firm f would obtain w_2 upon making him an offer.

In general, it is easily shown that conditional on firm f having been rejected by its top k workers, the resultant expected utility is identical to that of interviewing the remaining $K-k$ workers with $K-k$ other firms with perfect overlap.

We thus can generalize this logic and note if firm f makes an offer to any w_l, then the probability that the $l-1$ other remaining firms who have not yet been matched make an offer to w_l (or to any of the remaining l workers) can be defined recursively as

$$\rho^l = \frac{1}{l}\sum_{i=1}^{l} Pr(\delta_{i:l} \geq 0) \prod_{j=i+1}^{l} (1 - \hat{P}_{j-1}(\rho^j)) \tag{A.1}$$

where $\hat{P}_k(\rho)$ is the probability a worker who has k interviews accepts a job offer from a firm given the other $k-1$ firms submit a job offer with probability ρ. If we let $R_{i,k}(\rho)$ denote the probability a worker receives i offers out of k interviews, given he receives at least one offer and each firm submits an offer with probability ρ, then

$$R_{i,k} = \binom{k-1}{i-1}(\rho)^{i-1}(1-\rho)^{k-i}$$

[18]Otherwise, f would have been matched with a higher ranked worker.

and we see that

$$\hat{P}_k(\rho) = \sum_{i=1}^{k} \frac{1}{i} R_{i,k}(\rho) = \sum_{i=1}^{k} \frac{\binom{k-1}{i-1}}{i} (\rho)^{i-1}(1-\rho)^{k-i} = \sum_{i=1}^{k} \frac{\binom{k}{i}}{k} (\rho)^{i-1}(1-\rho)^{k-i}$$

Thus, with perfect overlap, any firm's probability of obtaining its kth highest ranked worker conditional on making an offer is

$$\overline{P}_{(k)} = \hat{P}_k(\rho^k) \prod_{j=k+1}^{K} (1 - \hat{P}_{j-1}(\rho^j))$$

A.1.2 Lower Bound

In general, it is difficult to explicitly characterize $\overline{P}_{(k)}$ for general forms of symmetric overlap. The reasoning is as follows: consider a firm f. Unlike with perfect overlap, following the rejection of f by its top ranked worker w_K, f no longer faces identical competition from its remaining firms who interview w_{K-1}. Indeed, there may exist a firm f' who also interviewed w_K and w_{K-1}, but a firm f'' that only interviewed w_{K-1} and not w_K. Consequently, firm f having lost w_K now expects f' to have a different probability of making an offer to w_{K-1} than firm f''. As a result, the ability to treat firms symmetrically disappears in all states following a worker's rejection in non-perfect overlap cases.

Nonetheless, we still can explicitly compute a lower bound on these probabilities, and hence characterize the lower bound of utility achievable under any pure strategy equilibrium by making assumptions to restore this symmetry. Recall that with any symmetric interview assignment η, $\overline{P}_{(k)} \leq \overline{P}_{(k-1)}\ \forall\ k$.[19] But if we assume that contingent on having lost a previous worker, a firm faces the same competition as before (i.e., $\overline{P}_{(k)} = \overline{P}_{(k-1)}$), then we can provide a lower bound on the utility achievable in any pure strategy symmetric assignment.

The reason for this particular exercise is two-fold: (1) for $K << N$, such an approximation is close to that achievable with a pure-strategy interview assignment with low overlap, and (2) the tractable closed form expressions help elucidate the intuition for some of the dynamics of the interview assignment game.

First, by assuming a firm's previous rejections do not influence his future competition means that the probability a competitor makes an offer to a given worker of the job-matching deferred acceptance algorithm does not change from round to round. We denote this probability ρ.

Again, let $R_{i,K}$ denote the probability a worker receives i offers out of K interviews,

[19] This is because contingent on having been rejected by a kth ranked worker w_k, firm f now faces less competition for its $k-1$ ranked worker w_{k-1} *even if* no other firm interviewed both workers: once w_k is matched, there is one less firm who is now competing for any other worker; the loss of that firm reduces competition for some other worker w', which in turn makes it more likely for some firm f' to hire that worker, which in turn reduces competition for worker w'', and so forth. This chain of worker-firm interview assignments thus influences the probability of facing competition for worker w_{k-1}, and increases the probability of obtaining him. In the extreme case of perfect overlap, this benefit manifested itself explicitly – having lost w_k directly implied that one less firm could possibly compete for worker w_{k-1}.

given he receives at least one offer.

$$R_{i,K} = \binom{K-1}{i-1}(\rho)^{i-1}(1-\rho)^{K-i}$$

Let $\overline{P}_K(\rho)$ denote the probability a worker who has K interviews accepts a job offer from a firm. This is also a function of ρ, since a worker's acceptance depends on the likelihood of receiving offers from *other* firms.

$$\overline{P}_K(\rho) = \sum_{i=1}^{K} \frac{1}{i} R_{i,K} = \sum_{i=1}^{K} \frac{\binom{K-1}{i-1}}{i}(\rho)^{i-1}(1-\rho)^{K-i} = \sum_{i=1}^{K} \frac{\binom{K}{i}}{K}(\rho)^{i-1}(1-\rho)^{K-i}$$

Lemma A.1. *For any value of $\rho \in (0,1)$, $\overline{P}_K(\rho)$ is decreasing in K. For any value of $K \in \{1,\dots,N\}$, $\overline{P}_K(\rho)$ is decreasing in ρ.*

Proof. $\overline{P}_K(\rho)$ is simply the expected value of $\frac{1}{x+1}$ where x is distributed according to the binomial distribution with $K-1$ trials and probability of success ρ. Consider what happens when K increases by 1: for each state of the world where there had previously been r successes, there are now two possible states with either r or $r+1$ successes, depending on the outcome of the new trial. Thus, in each state of the world, $\frac{1}{x+1}$ is now weakly decreasing and consequently the expected value of $\frac{1}{x+1}$ decreases as K increases. The second part of the lemma follows via similar reasoning: the expected value of $\frac{1}{x+1}$ with a fixed number of trials is decreasing as the probability of success (ρ) increases. □

Finally, we can define ρ:

$$\rho = \frac{1}{K}\sum_{i=0}^{K-1} Pr(\delta_{K-i:K} \geq 0)(1-\overline{P}_K(\rho))^i \tag{A.2}$$

Since the right hand sides of this equation is continuous and a function from $[0,1] \to [0,1]$, by Brouwer's Fixed Point Theorem there exists a fixed point ρ. Since $\overline{P}_K$ is decreasing in ρ, $(1-\overline{P}_K)$ is increasing in ρ and consequently the RHS of this equation is strictly decreasing in ρ. Thus, any fixed point must be unique.

We stress once again that our simplifying assumption – that contingent on being rejected, competition for the next best worker does not change – allows us to write equation (A.2) and hence solve explicitly for a firm's expected utility. Additionally, comparing (A.1) to this expression allows on to see how fiercer competition for higher ranked workers without overlap leads to more competition for lower ranked workers.

A.2 Mixed Strategy Analysis

A symmetric mixed strategy equilibrium where everyone interviews K candidates is more computationally involved to characterize, since not only are the number of interviews that any particular candidate expects to receive is random, but also is the degree of overlap. If a firm interviews a subset of K workers with equal probability, then the probability that any given worker receives k interviews given he receives at least 1 can be computed – letting $q = \frac{K}{N}$, we can compute this probability as $g_K(k) = \binom{N-1}{k}q^i(1-q)^{N-1-i}$. But this is inadequate, since the distribution across firm identities matters as well.

We focus on a given firm f which interviews a sequence of K workers $W_f \equiv \{w_1, \ldots, w_K\}$ in order of rank. If every other firm interviews a random subset of K workers, it induces a distribution over the space of interview assignments $\Omega \equiv \{\eta | \eta(f) = W_f, |\eta(f')| = K \ \forall \ f'\}$. We now have the expected utility of a firm defined as

$$EU_f(W_f, \nu_{-f}) = \frac{1}{|\Omega|} \sum_{\eta \in \Omega} \left[\Lambda_{K:K} \overline{P}_K(\eta) + \ldots + \Lambda_{1:K} \overline{P}_1(\eta) \prod_{j=2}^{K} (1 - \overline{P}_j(\eta)) \right]$$

where the probabilities $\overline{P}_j(\cdot)$ are functions of the realized interview assignment.

B Proofs

Proof of Lemma 2.1. Since this is equivalent to the marriage problem that yields the M-optimal stable matching (with firms as men), firms have a dominant strategy to report their preferences truthfully (Dubins and Freedman (1981), Roth (1982)). For workers, it is sufficient to rule out two types of deviations: (i) a worker may rank some firm as "unacceptable" and reject any offer from that firm; (ii) a worker may rank firm j' higher than j in his reported preferences despite preferring j to j' in his true preferences.

To see why deviation (i) may be effective, note that declaring a firm as unacceptable can lead to the following "chain" of events: a worker rejects some firm j's offer (instead of holding onto it or accepting it), which leads to that firm to offer a job to another worker who then rejects another firm who he prefers less, and so on, until a firm j' who was rejected by another worker makes an offer to the original worker, and this worker prefers j' to j. As long as the gain to such a deviation is never greater than the potential loss from employing it, a worker will never choose to reject any firm.

Let $L \leq N$ be the maximum number of interviews any worker receives in any equilibrium. Assume worker i is considering ranking firm j as unacceptable. In order for this to be profitable, firm j upon being rejected (conditional on making i an offer) must propose to a worker that already has an existing offer from another firm j', and that worker must prefer j to j'. The probability that this firm j is preferred to any j' by another worker is exactly $\frac{1}{2}$, and consequently the probability that rejecting a firm leads to a profitable manipulation is at most $\frac{1}{2}$. Thus the gain to rejecting a firm is bounded by $\frac{1}{2}(u_i(f_{i(1)}) - u_i(f_{i(N)}))$, where the term in parenthesis is the maximum gain possible to i by obtaining a more preferred firm. However, if a worker receives L interviews and rejects an offer, the probability that he receives no other offer is at least $(\frac{1}{2})^{L-1}$, since he receives an offer with probability at most $\frac{1}{2}$ (the probability that his δ for a firm is positive). Consequently, by rejecting firm j, he risks losing at least $(\frac{1}{2})^{L-1}(u_i(f_{i(N)}) - u_i(\emptyset))$. Clearly as long as $\beta > \frac{1}{(\frac{1}{2})^L}$ and the inequality (2.2) holds, no worker will find it profitable to reject any firm's offer.

To rule out deviation (ii), we first establish the following claim: prior to engaging in the match, the expected probability of being hired by a firm is strictly decreasing in the rank a worker orders that firm in his reported preferences. First recall preferences are independently drawn for all agents and privately realized and workers do not observe the complete interview assignment. Thus, a worker perceives the probability of receiving a job offer is the same for any firm. If this probability is denoted by p, then the expected probability of being hired by a firm ranked in nth position is $(1-p)^{n-1} \times p$ (since in order to be hired by the nth firm, all firms that were ranked higher must not have made a job offer). This expression is decreasing in n.

Having established the claim, it is straightforward to show that if any worker ranked f' higher than f despite preferring f to f', he would be better off not doing so and instead reporting truthfully. □

Proof of Proposition 3.1. Assume each firm randomly selects y workers to interview: e.g., each firm plays a strategy ν_f which assigns equal positive probability to only those subsets of workers of size y. We show that there exists a c such that no firm will wish to deviate.

Consider now firm f. Let

$$g_f(k,\nu_{-f}) = \max_{W_f s.t. |W_f|=k} \int_{\nu_{-f}} EU_f(W_f, W_{-f}) - \max_{W_f s.t. |W_f|=(k-1)} \int_{\nu_{-f}} EU_f(W_f, W_{-f}) \tag{B.1}$$

denote the expected gain to interviewing an additional kth worker (not including costs). In this particular case, since every firm is randomizing uniformly, firm f is indifferent over each worker. Thus, any choice of k workers is optimal.

We first prove the following lemma:

Lemma B.1. *$g_f(k,\nu_{-f})$ is decreasing in k.*

Proof. Denote $\mu(f)$ as the worker matched to firm f if it interviews an additional candidate w, and denote $\mu'(f)$ as the worker it is matched to if it does not interview w. We can decompose $g_f(k,\nu_{-f})$ as follows:

$$\begin{aligned} g_f(k,\nu_{-f}) \quad = \quad & \underbrace{Pr(\mu(f)=w)}_{(1)} \Bigg[\underbrace{Pr(\mu'(f)=f)}_{(2)} \underbrace{(E[\delta_w|\mu(f)=w \& \mu'(f)=f])}_{(3)} + \\ & + \underbrace{Pr(\mu'(f)=w')}_{(4)} \underbrace{(E[\delta_w - \delta_{w'}|\mu(f)=w \& \mu'(f)=w'])}_{(5)} \Bigg] \end{aligned}$$

(1) *Probability that interviewing w results in hiring w*: w is only hired if firm f makes it an offer, which occurs only if $\delta_{w,f} \geq 0$ if $\mu'(f) = f$, and if $\delta_{w,f} \geq \delta_{w',f}$ if $\mu'(f) = w'$. If $\delta_{w,f} = 0$, then $Pr(\delta_{w,f} \geq 0)$ does not change with k, but if $\delta_{w,f} = w'$, then $Pr(\delta_{w,f} \geq \delta_{w',f})$ is decreasing in k. To see why, $[\delta_{w',f}|\mu'(f) = w']$ is simply the utility of interviewing $k-1$ candidates contingent on making a hire without accounting for interviewing costs. This amount is clearly increasing in k: in every state of the world (i.e., for any realization of δ for all firms), $\delta_{w',f}$ given w' is hired is weakly increasing in k as interviewing an additional worker cannot hurt the expected surplus realized from the eventual hire; the more workers that are interviewed, the greater the expected utility of the worker that is eventually hired (keeping the actions of other firms fixed).

Increasing k does not affect the probability that w accepts or rejects an offer made by firm f – indeed, this probability is only influenced by ν_{-f}, which is held fixed. Thus, $Pr(\mu(f) = w)$ is decreasing in k.

(2) *Probability that without interviewing w, firm f would have been unmatched*: Clearly this is decreasing in k – the more workers f interviews, the less likely it will remain unmatched.

(3) *Expected gain from hiring worker w given alternative under μ' was being unmatched:* Again, $E[\delta_{w,f}|\mu(f) = w\&\mu'(f) = f] = E[\delta_{w,f}|\delta_{w,f} > 0]$ – $\mu(f) = w$ and $\mu'(f) = f$ implies only that $\delta_{w,f} > 0$, since (1) $\delta_{w,f}$ is a necessary and sufficient condition for f to have made a job offer to worker w (since no other worker f interviewed accepted its offers), and (2) the decision of whether or not w accepts f's offer is independent of $\delta_{w,f}$ and is only a function of w's preferences. Thus this is independent of k.

(4) *Probability that without interviewing w, some other worker w' was hired*: Since (2) is decreasing in k and $(4) = 1 - (2)$, this is increasing in k.

(5) *Expected gain from hiring worker w given alternative under μ' was being matched to w':* As in (1) and (3), note $E[\delta_{w,f} - \delta_{w',f}|\mu(f) = w\&\mu'(f) = w'] = E[\delta_{w,f} - \delta_{w',f}|\delta_{w,f} > \delta_{w',f}\&\mu'(f) = w']$; i.e., we know $\delta_{w,f} > \delta_{w',f}$ or otherwise f would not have made an offer to worker w after interviewing him. Consequently, since $E[\delta_{w,f} - y|\delta_{w,f} \geq y]$ falls as y increases by our regularity condition imposed on $H(\cdot)$ (see equation (2.1)), $[\delta_{w',f}|\mu' = w']$ is increasing in every state of the world as k increases, and $\delta_{w',f}$ and $\delta_{w,f}$ are independent, it follows that (5) is decreasing in κ.

Since (1) is decreasing in κ, (2) is decreasing, (3) does not change, (5) is decreasing, and $(2) + (5) = 1$ while $(3) > (5)$ since $\delta_{w',f} > 0$, the lemma is proved. □

By the previous lemma, we can find $c \in (g_f(y, \nu_{-f}), g_f(y-1, \nu_{-f}))$. For such c, given every other firm is interviewing a subset of y workers at random, no individual firm will wish to interview more than y candidates (since doing so earns an expected gain of less than c per additional candidate) or less than y candidates (since doing so gives up an expected gain greater than c per candidate). Furthermore, every firm is indifferent over all subsets of y workers, so a mixed strategy is an equilibrium. □

Proof of Proposition 3.2. For any x, we construct the following interview assignment η: assign each firm and worker a number $0, \ldots, N-1$. For each firm i, assign that firm the set of workers $\{i, [i+1]_N, \ldots, [i+x]_N\}$, where $[j]_N$ denotes the worker who corresponds to the index $j\,\mathbf{mod}\,N$. Clearly this assignment generates symmetric "overlap" and symmetric probability that a firm will make an offer to a worker.

As in the previous proof, we define the gain $g_f(k, W_{-f})$ of firm f to interviewing an additional kth worker as in (B.1), except now the other firms do not randomize. The twist now is that workers faced by firm f are no longer ex ante symmetric – indeed, for $k \leq x$, the optimal choice of workers to interview for f is any subset of k workers in $\eta(f)$ (each worker in $\eta(f)$ only has $x-1$ interviews from other firms and any other worker $w \notin \eta(f)$ already has x interviews); for $k > x$, it must interview workers who already have x interviews in addition to those workers in $\eta(f)$. Nonetheless, it is still straightforward to extend lemma B.1 to this setting, and that the gain to interviewing an additional worker is once again decreasing in k.

Let $c \in (g_f(x, W_{-f}), g_f(x-1, W_{-f}))$. For such c, not only will no firm choose to add or drop workers to interview, no firm will wish to change the composition of its candidates – any firm can only swap a worker with x interviews for one that will have $x+1$ interviews, and thus such a deviation leaves the firm strictly worse off. Thus η is an equilibrium interview assignment. □

Proof of Proposition 3.3. From (3.1), it is clear that a perfect overlap equilibrium yields higher utility than any other *pure strategy* symmetric equilibrium in which firms interview x workers: since $\bar{P}_{(k)}$ is strictly higher under perfect overlap than without (i.e., a firm has a higher probability of obtaining it's kth ranked worker contingent on making him an offer under perfect overlap than without), both a firm's utility and probability of being matched is also strictly higher.

To show that a perfect overlap equilibrium outperforms a mixed strategy equilibrium, first note that the number of competing firms for firm f's kth ranked worker (contingent on making him an offer) will always be $k-1$ under perfect overlap. However, under a mixed strategy equilibrium, the *expected* number of competitors i (denoted $E(i,k)$) will be strictly greater than $k-1$. This follows because there are $N-(x-k)$ firms and workers are still unmatched when a firm is making an offer to its kth worker; in order to have an expectation of achieving fewer than $k-1$ competitors for the kth worker, over half of the firms who have not been matched must also have made offers to each of firm f's $k+1,...,K$ ranked workers, which in turn occurs with probability less than 1/2. Thus, since $\bar{P}_k(\text{mixed strategy}) = E_i[\frac{1}{i+1}] < \frac{1}{E(i,k)+1}$ by Jensen's inequality, and since $\frac{1}{E(i,k)+1} < \frac{1}{k} = \bar{P}_k(\text{perfect overlap})$ (which follows because $k-1$ is less than $E(i,k)$), $\bar{P}_{(k)}$ is strictly higher under perfect overlap than a mixed strategy equilibrium and a firm's utility and probability of being matched is also strictly higher. □

C Existence of Correlated Equilibrium with Almost Perfect Overlap

In small markets, a given c might require an x such that the integer constraints N precludes perfect overlap. Nonetheless, with a large enough market, this issue is not a problem: with a correlated device or an intermediary, there still will exist a symmetric equilibrium where each firm and worker receives x interviews, and each firm in expectation receives the same degree of overlap. For a given firm, as the market size grows, the probability of receiving perfect overlap approaches 1.

Proposition C.1. *If there exists N, c such that a symmetric pure-strategy equilibrium where each firm and each worker receives x interviews with perfect overlap exists, then for any $\epsilon > 0$, there exists an $\overline{N}$ such that $\forall N > \overline{N}$ a correlated equilibrium exists in which each firm interviews x workers, each worker receives x interviews, and with probability $1-\epsilon$ each firm achieves perfect overlap.*

Proof of Proposition C.1. For any N, we can partition the population into $\lfloor \frac{N}{x} \rfloor - 1$ groups of exactly x workers and firms, and 1 group of $x + (N - \lfloor \frac{N}{x} \rfloor)$ workers and firms.[20] For any such partition π, associate an interview assignment $\eta(\pi)$ whereby in each of the groups with exactly x workers and firms, every firm interviews every worker in that group, and in the group with slightly more than x workers and firms, the interview assignment among workers and firms assigns each agent x symmetric interviews as in the proof of proposition 3.2. Thus, $\eta(\pi)$ gives each firm and each worker x interviews, and for but only $x+(N-\lfloor \frac{N}{x} \rfloor)$ workers and firms, there is perfect overlap.

Consider the space of all possible π and associated $\eta(\pi)$. For any $\varepsilon > 0$, there exists an $\overline{N}$ such that for any $N > \overline{N}$, if a π is chosen at random, the probability that a given

[20] $\lfloor x \rfloor$ represents the greatest integer less than or equal to x

firm achieves perfect overlap in the interview assignment $\eta(\pi)$ is at least $1-\varepsilon$. Thus, for sufficiently large N and small ε, we can construct a correlated equilibrium in which firms "coordinate" on a given $\eta(\pi)$ at random, and achieve perfect overlap with probability $1-\varepsilon$. □

Even though N may not be a multiple of x, as long as N is sufficiently large, the perfect overlap quantity of interviews can still be achieved for almost all firms. This outcome is strictly preferable to that achieved in a true symmetric pure strategy equilibrium.

D Quantity of Interviews in Equilibrium

For a fixed c, the marginal contribution of an extra worker in a mixed strategy equilibrium can be either larger or smaller than in a pure strategy equilibrium for the xth worker. Thus, either may result in more interviewing, as the following example illustrates:

Example D.1. *Let $N=2$ and let δ be drawn from the same distribution as in the previous examples. If $c \in (.0405, .9)$, a pure strategy equilibrium in which each firm interviews 1 worker exists. However, within this range, if $c \in (.6525, .9)$, an asymmetric mixed strategy equilibrium in which one firm mixes and the other firm interviews no worker exists. On the other hand, if $c \in (.1, .29)$, the only mixed strategy equilibrium that exists is one in which each firm interviews both workers.*

References

Atakan, A. E. (2006): "Assortative Matching with Explicit Search Costs," *Econometrica*, 74, 667–680.

Becker, G. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81, 813–846.

Bergemann, D., and J. Välimäki (2005): "Information in Mechanism Design," 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. 5. Cambridge University Press, Cambridge, U.K.

Chade, H., and L. Smith (2006): "Simultaneous Search," *Econometrica*, 74(5), 1293–1307.

Chakraborty, A., A. Citanna, and M. Ostrovsky (2007): "Two-Sided Matching with Interdependent Values," mimeo.

Dubins, L. E., and D. A. Freedman (1981): "Machiavelli and the Gale-Shapley Algorithm," *American Mathematical Monthly*, 88, 485–494.

Echenique, F., and J. Oviedo (2004): "A Theory of Stability in Many-to-Many Matching Markets," *mimeo.*

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

Jackson, M. (2004): "A Survey of Models of Network Formation: Stability and Efficiency," in *Group Formation in Economics; Networks, Clubs and Coalitions*, ed. by G. Demange, and M. Wooders, chap. 1. Cambridge University Press, Cambridge, U.K.

Kranton, R. E., and D. F. Minehart (2001): "A Theory of Buyer-Seller Networks," *American Economic Review*, 91(3), 485–508.

Lee, R. S., and M. Schwarz (2007): "Signalling Preferences in Interviewing Markets," in *Computational Social Systems and the Internet*, ed. by P. Cramton, R. Müller, E. Tardos, and M. Tennenholtz, no. 07271 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany.

Lien, Y.-C. (2006): "Costly Interviews and Non-Assortative Matchings," mimeo, Stanford University.

Mongell, S., and A. E. Roth (1991): "Sorority Rush as a Two-Sided Matching Mechanism," *American Economic Review*, 81(3), 441–464.

Roth, A. E. (1982): "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 7, 617–628.

——— (1984): "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, 92, 991–1016.

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.

Shimer, R., and L. Smith (2000): "Assortative Matching and Search," *Econometrica*, 68, 343–369.