

NBER WORKING PAPER SERIES

INTERVIEWING IN TWO-SIDED MATCHING MARKETS

Robin S. Lee
Michael Schwarz

Working Paper 14922
<http://www.nber.org/papers/w14922>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2009

We thank Itay Feinmesser, Kyna Fong, Rachel Kranton, David McAdams, Michael Ostrovsky, Al Roth, and Lones Smith for helpful comments. All errors are our own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Robin S. Lee and Michael Schwarz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Interviewing in Two-Sided Matching Markets

Robin S. Lee and Michael Schwarz

NBER Working Paper No. 14922

April 2009

JEL No. C78,D85,J01

ABSTRACT

We introduce the interview assignment problem, which generalizes the one-to-one matching model of Gale and Shapley (1962) by introducing a stage of costly information acquisition. Agents may learn preferences over partners via costly interviews. Although there exist multiple equilibria where all agents receive the same number of interviews, efficiency depends on overlap – the number of common interview partners among agents. We prove the equilibria with the highest degree of overlap yields the highest probability of being matched. The analysis suggests that institutions which ration interviews or create labor market segmentation may lead to greater efficiency in information acquisition activities.

Robin S. Lee

NYU Stern, Department of Economics

44 West 4th Street, 7th Floor

New York, NY 10012

rslee@stern.nyu.edu

Michael Schwarz

Yahoo! Research

2397 Shattuck Ave

Berkeley, CA 94704

and NBER

mschwarz@yahoo-inc.com

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 However, in many matching markets, information acquisition plays an important role: in labor markets, firms interview workers; in marriage markets, men date women; and in real estate markets, buyers attend open-houses. Such interviews, dates, and meetings are often costly and scarce – e.g., firms pay up to a third of an employee’s annual salary as commission to headhunters just to narrow down the potential field of candidates to interview² – and since these interviews affect the formation of preferences, their assignment can crucially affect the efficiency of the eventual matching process.

Our paper generalizes 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 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 and 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 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. We primarily focus on the first stage interviewing decisions, and use standard results from the one-to-one matching literature to analyze the second stage assignment. 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; 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.³

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

²Coles, Edelman, Hall, and Bennett (2008).

³Although homogeneity is a strong assumption to apply for an entire market, it is potentially plausible if a market is partitioned into subsets of workers and firms of similar quality. For example, although workers in general may have a greater likelihood of preferring a “top-tier” firm in the entire market over a “bottom-tier” firm, within a tier worker preferences may be fairly evenly distributed; this also may be true with firm preferences over workers within a tier. Thus, in the sense that workers and firms of similar quality are the only ones that interview and match with each other, and – conditional on all available information – worker and firm preferences are generally uniform and i.i.d among a particular subset of agents, our model is less

Even if all firms and workers are *ex ante* identical (prior to the realization of their idiosyncratic preferences), our model emphasizes that 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. 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.

We also investigate a more subtle form of coordination which proves just as important. Although there may exist many equilibria in which all agents conduct the same number of interviews, the efficiency of the match (defined as the number of expected hires) 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 firm or worker. This result is also not fragile to the i.i.d. assumption over preferences: e.g., we also note that overlap minimizes unemployment when firm preferences over workers are perfectly correlated.

The analysis suggests that institutions which either limit the number of interviews candidates receive or segment the market such that firms only interview workers within well defined categories can improve coordination, reduce interviewing costs, and increase the efficiency of the matching process. For example, 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; and the “rush” system by which sororities on college campuses recruit new members limits and equalizes interviews: “[a] rushee who receives more invitations than the number of parties permitted in a given round must decline, or ‘regret,’ the excess invitations”

restrictive than it might appear.

(Mongell and Roth, 1991).

Examples of improving overlap include segmentation via geography (e.g., firms interviewing only local candidates) and other (potentially arbitrary) dimensions. Note that if the number of interviews each firm conducts in equilibrium is close to the total number of workers *or* if the population could be partitioned such that agents in each group can only interview other agents in that group, interviews would be more evenly dispersed and the degree of overlap drastically increased. Thus, even if academic departments have no preference over the field in which to hire, they may still benefit from restricting the hiring to a particular field.

Finally, 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. 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 centralized mechanism such as a deferred acceptance procedure given certain assumptions (Niederle and Yariv, 2008; Schweinzer, 2008; Haeringer and Wooders, 2009). Whenever preferences are *ex ante* unknown and need be revealed through a costly interview, due diligence, or even dating process, our analysis remains relevant.

Related Literature

The interview assignment problem is a variant of the many-to-many matching problem since firms may be assigned to many workers and workers to many firms in the interview stage (Roth, 1984; Blair, 1988; Echenique and Oviedo, 2006; Konishi and Ünver, 2006). In contrast to this literature, our model allows for externalities imposed on agents not directly involved in a particular pairwise match as firms care about the identities of other firms who interview its candidates. Additionally, instead of relying on non-equilibrium concepts such as pairwise stability, we utilize standard non-cooperative equilibrium conditions when analyzing interview assignments. The special structure of the interview assignment problem considered herein enables us to solve for the equilibrium allocations.

The matching literature generally also assumes agents know their preferences before entering a match. However, there are notable exceptions: in settings different than ours, Das and Kamenica (2005), Chakraborty, Citanna, and Ostrovsky (2007), Josephson and Shapiro (2008), and Hoppe, Moldovanu, and Sela (2009) explore the role of information acquisition in matching markets. Das and Kamenica consider sequential learning in the context of dating markets where men and women repeatedly go on dates to learn their preferences; Chakraborty et al. investigate the stability of matching mechanisms with interdependent

values over partners; Hoppe et. al examine the effectiveness of signalling in an assortative matching environment; and Josephson and Shapiro study the role of adverse selection when firms sequentially interview workers in a decentralized matching process.

Although the literature on search seems related since it also explores the role of frictions in matching markets, it is quite dissimilar: the sequential and directed search literature (e.g., Shimer and Smith (2000); Atakan (2006); Eeckhout and Kircher (2009)) primarily investigates frictions due to delay and the impact on assortative matching; the model of simultaneous search in Chade and Smith (2006) focuses on the actions of a single decision maker who must choose a portfolio of ranked stochastic options, where the probability of obtaining a particular option is assumed to be exogenous.⁴ In contrast, we abstract away from notions of assortative matching and explicitly explore the miscoordination aspects of information acquisition by focusing on the externalities that firms impose on other firms via their interview decisions, thereby endogenizing the probability that a firm hires a worker. As a result, the interview assignment problem instead exhibits strong parallels to the literature on information acquisition in mechanism design: here firms (bidders) must interview (invest) to learn their private values over workers (objects), and firms' incentives to learn valuations over workers are reduced the more other firms choose to do so (c.f. Bergemann and Välimäki (2005)).

In a sense, our interview stage is also similar to the bipartite network formation models of Kranton and Minehart (2000, 2001), which explore the formation of trading links between buyers and sellers and where prices are influenced by these connections to potential trading partners. Although the economic environment appears to be very different from a matching market, there is a close mathematical connection. Indeed, their papers can be reinterpreted as an interviewing problem in a market with transferable utility: a seller (firm) may only trade with (hire) a buyer (worker) with whom it has formed a link (interviewed). A key mathematical difference, however, is that Kranton and Minehart (2000, 2001), consider models with transferable utility (prices) while our paper considers a traditional matching market model without transferable utility; furthermore, they do not allow for buyers to vary in their valuations over sellers, whereas idiosyncratic preferences are central to our analysis.

⁴Lien (2006) provides an example in which the assignment of interviews may be non-assortative.

2 Model

2.1 Setup and Definitions

There are N workers and N firms, represented by the sets $W = \{w_1, \dots, w_N\}$ and $F = \{f_1, \dots, f_N\}$. 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 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.⁵

The main innovation of our model is that $\{P_w\}_{w \in W}$ and $\{\delta_{w,f}\}_{w \in W, f \in F}$ are unknown 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 $c \in \mathbb{R}^+$. Interviewing costs are borne by firms, and firms are assumed to be risk-neutral expected utility maximizers. 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 $\int x dH(x) < 0$. This condition ensures that a firm would not hire a worker it has not interviewed.⁶ Finally, we impose one further condition:

$$(E_\delta[\delta | \delta \geq y] - y) \text{ is decreasing in } y \forall y \geq 0 \quad (2.1)$$

which implies that a firm's expected gains to interviewing an additional worker falls as it interviews more candidates. A sufficient condition for this is that the density h is log-concave, which is satisfied by the normal, exponential, uniform, and Weibull (for $\gamma > 1$) densities (c.f. Burdett (1996)).

Worker preferences are strict, 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 rank working for any firm over being 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 $f \in F_w$.⁷

⁵We abstract away from wage negotiations and assume that there are no wages (as in dating markets), or that wages are fixed or already embedded in user preferences across firms (as may be the case in markets for some entry level jobs).

⁶The 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.

⁷Since workers cannot communicate preferences to firms, when they learn their preferences does not matter in this model. Coles and Niederle (2007) and Lee and Schwarz (2007) consider settings where workers initially know their preferences over firms, and examine mechanisms which allow them to signal to firms prior to the assignment of interviews.

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 $c|W_f|$.⁸ 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 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 over acceptable workers 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

⁸Since 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.

yields a one-to-one matching μ . 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.⁹ This particular procedure can be viewed as an approximation for the outcome of the hiring process in decentralized 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 in accordance 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*

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

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

For the rest of the paper, all proofs are located in the appendix; a sketch of the proof of this lemma is as follows: 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

⁹See 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.

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 and the truth-telling equilibrium is the one played in any second stage subgame.

3 Interview Assignment

Having characterized the behavior of agents in the second stage (deferred acceptance algorithm), we now focus on the equilibrium first stage assignment of interviews. We are interested in “symmetric” equilibria in which all firms interview the same number of workers. Even with this restriction, there are still many different equilibria, and the expected number of unemployed workers or the costs expended on interviewing can vary depending on the equilibrium chosen.

3.1 Equilibrium Analysis

Formally, a firm’s strategy during the interview assignment stage is a probability measure ν_f over the powerset of all workers $\mathcal{P}(W)$; i.e., for a set of workers W_f , ν_f assigns a probability that f interviews those and only those workers. Note firm f interviews x workers if for any $W_f \subset W$, $\nu_f(W_f) > 0$ implies $|W_f| = x$; and firm f interviews y workers at random if $\nu_f(W_f) > 0$ if and only if $|W_f| = y$, and $|W_f| = |W'_f|$ implies $\nu_f(W_f) = \nu_f(W'_f)$.

A strategy profile $\nu \equiv \{\nu_f\}_{f \in F}$ is a subgame perfect Nash Equilibrium of this game if and only if:

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

where $EU_f(W_f, W_{-f})$ is the expected utility for firm f given it interviews the set of workers W_f , and all other firms interview the workers $W_{-f} \equiv \{W_{f'}\}_{\forall f' \in F, f' \neq f}$. We will explicitly define each firm’s expected utility in the next subsection, but for now it is sufficient to note that it includes the expected value of δ for the worker that a firm expects to hire, minus the costs of interviewing $|W_f|$ workers. The reason for subgame perfection is that we assume firms anticipate equilibrium play in the second stage matching process when computing first stage interviewing profits.

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

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

Such a mixed strategy equilibrium may be 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 concentrate interviews on only a few 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. This is equivalent to a pure strategy equilibrium, where each firm assigns probability 1 to one particular element $W_f \in \mathcal{P}(W)$. 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 $\eta(f) = W_f \forall f$. This type of equilibrium may also exist:

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

This and the previous existence proof relies on a result established in the Appendix (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 x th 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 x th 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.¹⁰

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

¹⁰Due 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 is not an equilibrium either).

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

$$\begin{aligned}\eta(A) &= \{1, 2\} & \eta(B) &= \{2, 3\} \\ \eta(C) &= \{3, 1\}\end{aligned}$$

where each firm interviews two workers, and each worker obtains two interviews. Assume $\delta = 1$ with probability .9 and $\delta = -10$ with probability .1, which corresponds to the case where a worker is likely to generate positive surplus to a firm, but may be very costly.

In the next subsection, we show how one can explicitly compute each firm's expected utility. For now, we note that each firm's expected profits are $\pi \approx .88 - 2c$ and the probability of being unmatched is approximately .12.

However, 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 – some workers may receive 3 interviews and others 1 or 0; or, each worker gets exactly 2 interviews; 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.

3.2 Firm's Expected Utility

We now describe some of the details required to solve example 3.1, as well as provide insights into the proofs of this paper. A reader interested in only the main results may wish to skip directly to section 4.

In section 2, we proved firms and workers report preferences honestly in an equilibrium of the second stage matching process; as a result, each 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 :

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

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} . 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. 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 receiving 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 δ_w . draws with other firms and his own preferences, 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 $K > 0$ workers, it is equivalent to taking K “draws” on δ . The realization of the j th highest δ draw is itself a random variable, known as the j th 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 :

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

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 j th 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 j th 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). We stress that $\overline{P}_{(j)}$ is a function of the entire interview assignment $\eta \equiv \{W_f, W_{-f}\}$, but omit this notation when it is unambiguous. Finally, the probability a firm eventually is matched to any worker is simply equation (3.1) with $Pr(\delta_{j:K} \geq 0)$ replacing $\Lambda_{j,K}$:

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

Since the probabilities $\overline{P}_{(j)}$ are a function of the other firms’ actions W_{-f} , they 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 j th ranked worker.

Let P_i^j indicate the probability that when a firm makes an offer to its j th 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 3.2. Let $N = 4$, and index firms by $\{A, B, C, D\}$ and workers by $\{1, 2, 3, 4\}$. Consider the following interview assignment η :

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

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

As in example 3.1, let $\delta = 1$ with probability .9 and $\delta = -10$ with probability .1.

Since all firms have symmetric interview assignments, any firm’s profits can be expressed

using (3.1) with the same values for each P_i^j :

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

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 100 thousand surplus for a firm or lose 1 million, a firm will obtain in expectation approximately 86 thousand 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.

4 Main Result

It is not surprising that coordinating on a pure-strategy equilibrium as opposed to playing a mixed strategy equilibrium can lead to efficiency gains, as demonstrated by example 3.1. 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.¹¹ 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 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 is one specific type of overlap that is focal:

Definition 4.1. An interview assignment η that assigns x interviews to each firm and worker exhibits *perfect overlap* if and only 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.

¹¹In 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.

Example 4.1. Recall in example 3.2 that the probability a firm is unmatched was approximately .14, and a firm's expected profits were approximately $.86 - 2c$.

Now take the setup of example 3.2, but we now assume that

$$\begin{aligned}\eta(A) &= \{1, 2\} & \eta(B) &= \{1, 2\} \\ \eta(C) &= \{3, 4\} & \eta(D) &= \{3, 4\}\end{aligned}\tag{4.1}$$

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 B 's top choice worker coincides with A 's top choice. We thus find

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

Furthermore, the probability of remaining unmatched is now

$$Pr(\delta_{2:2} < 0) + Pr(\delta_{2:2} > 0 \& \delta_{1:2} < 0)(\frac{1}{2}P_1) \approx .05$$

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 3.2 is an equilibrium for firms to follow for $c \in (.14, .23)$.¹² On the other hand, as long as $c \in (.5, .26)$, 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 (4.1) 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, (i) the greater the probability that at any stage of the deferred

¹²To see why, observe that interviewing an additional worker for any firm can yield at most a gain 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 change the set of workers it interviews.

acceptance algorithm a worker will accept a firm's offer, (ii) the more likely a firm will receive a higher-ranked worker in the match, and (iii) 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 minimize unemployment over any other pure-strategy symmetric equilibrium where firms interview the same number of workers.

Theorem 4.1. *Consider an equilibrium interview assignment η which assigns each firm and each worker exactly x interviews with perfect overlap. There is no other pure-strategy equilibrium interview assignment η' in which each firm interviews exactly x workers such that every firm and worker is matched with strictly higher probability.*

It is worthwhile to stress that our overlap result is not necessarily sensitive to the assumption that firm preferences over workers are i.i.d.: e.g., consider the modification that firm preferences over workers are perfectly correlated such that all firms share the same common value for each worker (but would have to still interview to learn these values, and could not share information). In this setting, an equilibrium with perfect overlap would also maximize expected employment among all other symmetric equilibria, including those in mixed-strategies. To see this, note that with perfect overlap, any worker that generates positive surplus will be employed as long as every firm conducts at least 1 interview; thus the expected probability of unemployment will simply be $Pr(\delta < 0)$. On the other hand, either imperfect overlap or random interviewing introduces matching frictions and strictly increases unemployment. E.g., consider the assignment in (3.3): if all workers were desirable and ranked according to their number, firm B and worker 4 could remain unemployed if D hired worker 1, A hired worker 2, and C hired worker 3; under the perfect overlap assignment in (4.1), no unemployment would occur.

Also note that due to integer constraints, for a given c and N a symmetric pure-strategy equilibrium with perfect overlap may not exist: the construction of 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, although we have focused on the probability of being matched, we have not discussed whether coordination – e.g., 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 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.

5 Concluding Remarks

The cost of interviewing and the importance of how interviews are allocated has remained outside of the scope of the matching literature. This paper takes a step towards incorporating interviewing into matching models. We have studied the interview assignment problem and provided a model for analysis when information acquisition is costly. We identified two distinct forms of potential miscoordination in the assignment of interviews: (i) workers may receive varying numbers of interviews; (ii) more subtly, firms may not efficiently overlap their interviews. Consequently, institutions which limit the number of interviews workers can receive or which artificially create labor market segmentation to encourage overlap may lead to better coordination in information acquisition activities.

As a first attempt to study interviewing in matching markets, our model is deliberately simple and hence relatively tractable. It is worth mentioning several aspects of interviewing markets that are beyond the scope of this paper but are interesting for future work. For example, our paper assumes that firms bear the full cost of interviewing; this may be a reasonable approximation in some markets, but less so in others. For example, economic PhDs do not pay for travel and accommodations for academic interviews whereas medical residents do. Can this division – or even the interviewing costs themselves – be endogenized? There is also a normative question of what division of interviewing costs promotes efficiency and how it differs from market equilibrium.

Other potentially important extensions include allowing for (ex-ante) heterogeneous agents or correlated preferences, and determining the socially efficient allocation of interviews in this environment; this analysis becomes even more complex when interviews can be assigned dynamically (i.e., sequentially) based on information revealed in earlier interviews. Finally, although the present model considered a matching market without transfers (or rigid wages that cannot be tailored to an individual candidate), generalizing the model to environments with endogenous transfers is another challenging and important problem.

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.

Recall P_i^k is the probability that when a firm makes an offer to its k th highest ranked worker,

i other firms also make that worker an offer. Let $\bar{P}_{(k)}$ indicate the probability that a firm obtains its k th highest worker given it makes that worker an offer. These probabilities are all conditional on having been rejected by all workers ranked higher than k . 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:

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

For reference, it is also useful to restate equation (3.1):

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

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 be matched. Consider the assignment η in which each firm and worker receives exactly K interviews with perfect overlap. We wish to characterize $\bar{P}_{(k)}$ for all $k \leq K$, where $\bar{P}_{(k)}$ is the probability a firm's k th 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, \dots, w_1 in decreasing order of preference.

Consider $\bar{P}_{(1)}$. Clearly $\bar{P}_{(1)} = 1$, 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 $\bar{P}_{(2)}$. 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 (otherwise, f would have been matched with a higher ranked worker). 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 also 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 resulting 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_k , then the probability that the $k - 1$ other remaining firms who have not yet been matched make an offer to w_k (or to any of the remaining k workers) can be defined recursively as

$$\rho_k = \frac{1}{k} \sum_{i=1}^k Pr(\delta_{i:k} \geq 0) \prod_{j=i+1}^k (1 - \hat{P}_j(\rho_j)) \quad (\text{A.1})$$

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 . 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 . For the case of perfect overlap, note that $P_i^k = R_{i,k}(\rho)$. Then

$$R_{i,k}(\rho) = \binom{k-1}{i-1} (\rho)^{i-1} (1-\rho)^{k-i} \quad (\text{A.2})$$

and we see that

$$\begin{aligned} \hat{P}_k(\rho) &= \sum_{i=1}^k \frac{1}{i} R_{i,k}(\rho) = \sum_{i=1}^k \frac{\binom{k-1}{i-1}}{i} (\rho)^{i-1} (1-\rho)^{k-i} = \sum_{i=0}^{k-1} \frac{\binom{k-1}{i}}{i+1} (\rho)^i (1-\rho)^{k-i-1} \\ &= \frac{1}{k} \sum_{i=0}^{k-1} \binom{k}{i+1} (\rho)^i (1-\rho)^{k-i-1} = \frac{1}{\rho k} \sum_{i=1}^k \binom{k}{i} (\rho)^i (1-\rho)^{k-i} \\ &= \frac{1}{\rho k} \left[\left(\sum_{i=0}^k \binom{k}{i} (\rho)^i (1-\rho)^{k-i} \right) - \binom{k}{0} (\rho)^0 (1-\rho)^k \right] \\ &= \frac{1}{\rho k} [1 - (1-\rho)^k] \end{aligned} \quad (\text{A.3})$$

where the first equality of the second line follows from $\frac{1}{i+1} \binom{k-1}{i} = \frac{1}{k} \binom{k}{i+1}$.¹³

Thus any firm's probability of obtaining its k th highest ranked worker conditional on making an offer is simply

$$\bar{P}_{(k)} = \hat{P}_k(\rho_k)$$

where ρ_k can be computed from (A.1) and by noting $\hat{P}_1(\rho_1) = 1$.

A.1.2 Lower Bound

In general, it is difficult to explicitly compute $\bar{P}_{(k)}$ 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 , $\bar{P}_{(k)} \leq \bar{P}_{(k-1)} \forall k$.¹⁴ But if we assume that contingent on having lost a previous worker, a firm

¹³We thank Itay Feinmesser aiding in these simplifications.

¹⁴This is because contingent on having been rejected by a k th 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} .

faces the same competition as before (i.e., $\bar{P}_{(k)} = \bar{P}_{(k-1)}$), 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 $K \ll N$, 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 given he receives at least one offer, and let $\hat{P}_k(\rho)$ be 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 . They are the same as defined in (A.2) and (A.3).

The simplifying assumption that contingent on being rejected, competition for the next best worker does not change, allows us to write the following fixed-point definition for ρ :

$$\rho = \frac{1}{K} \sum_{i=1}^K Pr(\delta_{i:K} \geq 0)(1 - \hat{P}_K(\rho))^{K-i} \quad (\text{A.4})$$

and solve explicitly for the lower bound on a firm's expected utility. Additionally, comparing (A.1) to (A.4) allows one 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 $g_K(k) = \binom{N-1}{k} q^k (1-q)^{N-1-k}$. 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 $W_f \equiv \{w_1, \dots, w_K\}$ 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

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

where the probabilities $\bar{P}_j(\cdot)$ 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 and privately realized for all agents, and workers do not observe the complete interview assignment. Thus, a worker perceives the probability of receiving a job offer as the same for any firm. If this probability is denoted by p , then the expected probability of being hired by a firm ranked in the n th position is $(1-p)^{n-1} \times p$ (since in order to be hired by the n th ranked 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 reporting truthfully. \square

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

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

denote the expected gain to interviewing an additional k th 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:

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

- (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 since $\delta_{w',f}$ and $\delta_{w,f}$ are independent, it follows that (5) is decreasing in k .

Since (1) is decreasing in k , (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. \square

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. \square

Proof of Proposition 3.2. For any x , we construct the following interview assignment η : assign each firm and worker a number $0, \dots, N-1$. For each firm i , assign that firm the set of workers $\{i, [i+1]_N, \dots, [i+x-1]_N\}$, where $[j]_N$ denotes the worker who corresponds to the index $N \bmod j$.¹⁵ 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 k th 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. \square

Proof of Theorem 4.1. First, recall the discussion and notation in Appendix A: $\hat{P}_k(\rho)$ is the probability a worker who has k interviews accepts a job offer given the other $k-1$ firms submit a job offer with probability ρ , and is defined in (A.3). The following lemma will be useful:

Lemma B.2. *For any value of $\rho \in (0, 1)$, $\hat{P}_k(\rho)$ is strictly decreasing in k . For any value of $k \in \{1, \dots, N\}$, $\hat{P}_k(\rho)$ is strictly decreasing in ρ .*

Proof. Follows directly from (A.3).¹⁶ \square

To prove the theorem, note equation (3.1) implies a perfect overlap equilibrium yields higher utility and probability of being matched than any other pure strategy symmetric equilibrium in which firms interview x workers if it has a higher value of $\bar{P}_{(k)}$ for all k . Conditional on being rejected by $x-k$ workers the expected number of competitors for its k th highest worker is strictly greater than k without perfect overlap, and equal to k with perfect overlap. By lemma B.2, this means $\hat{P}_k(\rho)$ is strictly higher under perfect overlap than without it for a fixed ρ . Furthermore, by (A.1), ρ will strictly be higher without perfect overlap since the greater the probability of rejection by a worker (as firms face stiffer competition), the more likely that a firm will make an offer

¹⁵ $a \bmod b$ is the remainder when a is divided by b .

¹⁶Intuitively, note $\bar{P}_K(\rho)$ 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 decreasing and consequently the expected value of $\frac{1}{x+1}$ strictly 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 strictly decreasing as the probability of success (ρ) increases.

conditional on interviewing a worker. Hence, $\hat{P}_k(\rho)$ at the equilibrium ρ is strictly higher with perfect overlap than without it by the same lemma, and thus $\bar{P}_{(k)}$ is strictly higher with perfect overlap as well. \square

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; i.e., 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 every firm.

Proposition C.1. *If there exist N, c such that a symmetric pure-strategy equilibrium where each firm and each worker conducts x interviews with perfect overlap exists, then for any $\epsilon > 0$, there exists an \bar{N} such that $\forall N > \bar{N}$ 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 $x + (N - \lfloor \frac{N}{x} \rfloor x)$ workers and firms.¹⁷ 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 $x + (N - \lfloor \frac{N}{x} \rfloor x)$ workers and firms, there is perfect overlap.

Consider the space of all possible π and associated $\eta(\pi)$. For any $\epsilon > 0$, there exists an \bar{N} such that for any $N > \bar{N}$, if a π is chosen at random, the probability that a given firm achieves perfect overlap in the interview assignment $\eta(\pi)$ is at least $1 - \epsilon$. 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 $1 - \epsilon$. \square

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 x th 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

¹⁷ $\lfloor x \rfloor$ represents the greatest integer less than or equal to x

$c \in (.1, .29)$, 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.
- 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.
- BLAIR, C. (1988): “The Lattice Structure of the Set of Stable Matchings with Multiple Partners,” *Mathematics of Operations Research*, 13(4), 619–628.
- BURDETT, K. (1996): “Truncated Means and Variances,” *Economics Letters*, 52, 263–267.
- 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.
- COLES, P., B. EDELMAN, B. HALL, AND N. BENNETT (2008): “TheLadders,” Harvard Business School Case 9-908-061.
- COLES, P., AND M. NIEDERLE (2007): “Signaling in Matching Markets,” mimeo.
- DAS, S., AND E. KAMENICA (2005): “Two-Sided Bandits and the Dating Market,” in *Proceedings of the Nineteenth International Joint Conferences on Artificial Intelligence*, pp. 947–952.
- 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 (2006): “A Theory of Stability in Many-to-Many Matching Markets,” *Theoretical Economics*, 1(2), 233–273.
- ECKHOUT, J., AND P. KIRCHER (2009): “Sorting and Decentralized Price Competition,” mimeo.
- GALE, D., AND L. SHAPLEY (1962): “College admissions and the stability of marriage,” *American Mathematical Monthly*, 69, 9–15.
- HAERINGER, G., AND M. WOODERS (2009): “Decentralized Job Matching,” mimeo.
- HOPPE, H. C., B. MOLDOVANU, AND A. SELA (2009): “The Theory of Assortative Matching Based on Costly Signals,” *Review of Economic Studies*, 76(1), 253–281.

- JOSEPHSON, J., AND J. SHAPIRO (2008): “Interviews and Adverse Selection,” mimeo.
- KONISHI, H., AND M. U. ÜNVER (2006): “Credible Group Stability in Many-to-Many Matching Problems,” *Journal of Economic Theory*, 129(1), 57–80.
- KRANTON, R. E., AND D. F. MINEHART (2000): “Networks Versus Vertical Integration,” *RAND Journal of Economics*, 31(3), 570–601.
- (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.
- NIEDERLE, M., AND L. YARIV (2008): “Matching Through Decentralized Markets,” mimeo.
- 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.
- SCHWEINZER, P. (2008): “Labour Market Recruiting with Intermediaries,” *Review of Economic Design*, 12(2), 119–127.
- SHIMER, R., AND L. SMITH (2000): “Assortative Matching and Search,” *Econometrica*, 68, 343–369.