

Interviewing in Two-Sided Matching Markets*

Robin S. Lee[†]

Michael Schwarz[‡]

March 2017

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. Firms learn preferences over workers via costly interviews. Even if all firms and workers conduct the same number of interviews, realized unemployment depends also on the extent to which agents share common interviewing partners. We introduce the concept of *overlap* that captures this notion, and prove that unemployment is minimized with *perfect overlap*: i.e., if two firms interview any common worker, they interview the exact same set of workers.

KEYWORDS: interviewing, two-sided matching, market design

*We thank Neil Thakral for exceptional research assistance, and Itay Fainmesser, Kyna Fong, Rachel Kranton, David McAdams, Michael Ostrovsky, Al Roth, Lones Smith, two anonymous referees and the editor for helpful comments. All errors are our own.

[†]Harvard and NBER, robinlee@fas.harvard.edu.

[‡]Google and NBER, schwarz.m@gmail.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 and women date; 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 interview (cf. Coles et al., 2008)—and because these interviews affect the formation of preferences, their assignment can crucially affect the outcome of the eventual matching process.

Our article 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 article is one of the first to analyze the interview assignment problem in the context of two-sided matching. Throughout this article 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 are matched one-to-one 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 to being unemployed; firms may find some workers unacceptable and choose to remain unmatched; and workers and firms are *ex ante* homogeneous, with preferences over partners independent and idiosyncratic to each agent.²

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. As interviews are costly, firms care about how many interviews a potential interviewee has: as the number of interviews a worker has increases, the probability of 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

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

²Although 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 restrictive than it might appear.

to accept if an offer is made. As a result, even though there are equilibria in which firms all conduct the same number of interviews, firms would prefer not to randomize over workers but instead coordinate so that all workers receive the same number of interviews. We prove that for any number of interviews conducted by all firms, there are values of interviewing costs where firms randomly choosing workers to interview comprise an equilibrium of our game; additionally, there are (potentially different) values of interviewing costs so that all firms and workers conducting the same number of interviews is also an equilibrium.

We then investigate a more subtle form of coordination which also proves to have significant effects on equilibrium employment. For our analysis, we focus on *symmetric* interview assignments, which we define to be interview assignments in which each firm and worker has the same number of interviews, and even if two firms or two workers swap their interview assignments, the original interview assignment can be obtained via a renaming of other agents. Although there may exist many equilibria in which all agents conduct the same number of interviews, the number of expected hires in the match can be very different depending on whether or not the interview assignment exhibits *greater overlap*—a partial ordering that we define formally which captures the number of common interview partners among agents.

To illustrate why overlap matters, consider two firms, f and f' , who are the only firms that interview workers w and w' : if f has an offer rejected by w , it must be that the worker accepted an offer from f' ; consequently, f will then face no “competition” for w' and obtain him for certain if it makes him an offer. If f and f' did not interview the same set of workers, then f could possibly be rejected by both w and w' and not be matched despite making offers to both workers (because being rejected by w no longer implies obtaining w' for certain). Thus, a firm’s expected payoff depends not only on the number of interviews its workers receive, but also the identities of the firms interviewing those workers.

The main result of this article is that if the ex ante probability that a worker is found to be unacceptable to a firm is sufficiently small, then among all symmetric interview assignments in which every firm and worker conducts the same number of interviews, any assignment with *perfect overlap*—i.e., if any two firms interview the same worker, they interview exactly the same set of workers—minimizes unemployment for any firm or worker. We note that this result is also not fragile to the i.i.d. assumption over preferences: e.g., an interview assignment with perfect overlap also minimizes unemployment when firm preferences over workers are perfectly correlated.

We also use simulation results to quantify the impact on unemployment from two forms of coordination if all firms interview the same number of workers: (i) allowing firms to ensure that all workers receive the same number of interviews (as opposed to randomly assigning interviews), and (ii) conditional on all workers receiving the same number of interviews, increasing the amount of overlap in the interview assignment. We find that unemployment reductions from both forms of coordination can be substantial.

Our analysis suggests that institutions which, conditional on the total number of interviews that are conducted, limit the number of interviews that 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 number of successful matches. 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” (Mongell and Roth, 1991). Our model’s reliance on a single simultaneous round of interviewing can also be motivated by time constraints in these and other examples. With the possibility of delay, firms may benefit from engaging in sequential rounds of interviews; this is outside of the scope of the current analysis.

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 assignment would be closer to achieving perfect overlap. Thus, even if academic departments have no preference over the field in which to hire, each may still benefit from restricting its candidate search to a particular field if all departments did the same.

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 modeled as a centralized mechanism such as a deferred acceptance procedure given certain assumptions (Niederle and Yarov, 2009; Schweinzer, 2008; Haeringer and Wooders, 2011). Whenever preferences are *ex ante* unknown and need be revealed through a costly interview, due diligence, or dating process, our analysis remains relevant.

Related Literature. The interview assignment problem is a variant of the many-to-many matching problem because 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, although we assume our second-stage assignment satisfies non-equilibrium concepts such as pairwise stability, if we interpret this assignment as being generated from a centralized mechanism employing the deferred-acceptance algorithm, we are also able to utilize standard non-cooperative equilibrium conditions when analyzing interview assignments. Indeed, the special structure of the interview assignment problem considered herein enables us to

solve for the equilibrium allocations.

The matching literature also generally assumes agents know their preferences before entering a match. There are notable exceptions: Das and Kamenica (2005) consider sequential learning in the context of dating markets where men and women repeatedly go on dates to learn their preferences; Chade (2006) explores marriage markets where participants observe a noisy signal about the quality of their potential mates; Chakraborty, Citanna and Ostrovsky (2010) investigate the stability of matching mechanisms with interdependent values over partners; Josephson and Shapiro (2016) study the role of adverse selection when firms sequentially interview workers in a decentralized matching process; and Hoppe, Moldovanu and Sela (2009) examine the effectiveness of signaling in an assortative matching environment. Furthermore, the literature on college admissions such as Nagypal (2004) and Chade, Lewis and Smith (2014) can be reinterpreted as matching with costly information acquisition, where the applicants bear the cost of providing colleges with a noisy signal of their own quality.

A related issue not addressed in this article but examined in a companion piece (Lee and Schwarz, 2007) is the communication or signaling of preferences between agents once they are known, but prior to making offers. Coles, Kushnir and Niederle (2013) 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 in decentralized labor markets; they find that allowing for signaling can improve employment and worker welfare.

Although the literature on search seems related because 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 (2010)) 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.³ 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 (cf. 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

³Lien (2013) provides an example in which the assignment of interviews may be non-assortative.

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) whereas our article 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.

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 $F = \{f_1, \dots, f_N\}$. 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.⁴ The main innovation of our model is that firm preferences for workers 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 profit maximizers.

When a firm f interviews a worker w , it learns the value of $\delta_{w,f} \in \mathbb{R}$ which represents the firm-specific profits it realizes if it hires w . If a firm does not hire a worker, it receives a reservation surplus which we will assume to be 0. We assume $\{\delta_{w,f}\}_{w \in W, f \in F}$ comprises private i.i.d. draws from the continuous distribution H , where H has density h , finite first moments so that all order statistics have finite expectations (David and Nagaraja, 2003), and $\int x dH(x) < \infty$ which ensures that a firm would not hire a worker it has not interviewed.⁵ We impose the following condition:

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

This assumption is used in later proofs to guarantee that a firm's expected gain from interviewing an additional worker does not increase as it interviews more candidates.⁶ A sufficient condition for this is that the density h is weakly log-concave, which is satisfied by the normal, exponential, uniform, and Weibull (for $\gamma > 1$) densities (cf. Burdett (1996)). As H is continuous, the values of $\{\delta_{w,f}\}_{w \in W}$ for firm f can be used to generate strict preferences \mathcal{P}_f over workers that are interviewed, as ties occur with zero probability.

If a worker w is hired by firm f , he realizes surplus $\zeta_{w,f}$, where $\{\zeta_{w,f}\}_{w \in W, f \in F}$ comprise i.i.d. draws from the continuous distribution G with finite support $[\underline{\zeta}, \bar{\zeta}]$. If a worker is not matched to

⁴We 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).

⁵The 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.

⁶See Lemma A.2 in the appendix.

a firm, he receives $\zeta_0 < \underline{\zeta}$, and thus a worker always prefers being employed to being unemployed. Worker preferences are bounded to ensure workers report preferences truthfully in the second stage match of our game. For each worker w , $\{\zeta_{w,f}\}_f$ can be used to generate strict preferences \mathcal{P}_w over firms. In our model, workers cannot communicate preferences to firms, and thus whether they know their preferences prior to being interviewed or learn them afterwards does not matter in this model.⁷ For expositional purposes, we assume workers know their preferences over all firms prior to interviewing; as a firm will never make a job offer to a worker whom it never interviewed, how a worker ranks firms that do not interview him 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 $c \times |W_f|$, where $|W_f|$ is the number of interviews f conducts. We assume that workers do not decline interviews.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 $\eta(w) \subset F$ represents the set of firms that interview w . Each firm privately realizes $\{\delta_{w,f}\}_{w \in W_f}$ and forms ordinal preferences \mathcal{P}_f over workers.⁹
- (2) In the second stage, we assume that firms and workers match according to the *firm optimal stable matching* (FOSM; cf. Gale and Shapley (1962); Roth and Sotomayor (1990)) given preferences $\{\mathcal{P}_f, \mathcal{P}_w\}_{f \in F, w \in W}$ learned in stage 1.¹⁰ The realized matching is represented by μ , where we say w is hired by f if $\mu(w) = f$, and w is unemployed if $\mu(w) = w$. Similarly, we say f hires w if $\mu(f) = w$ and firm f does not hire anyone if $\mu(f) = f$.

We note that the FOSM used in the second stage of our game is the equilibrium outcome when firms and workers engage in a firm-proposing *deferred acceptance algorithm* (DAA) for employment (Gale and Shapley, 1962), and we utilize this particular procedure as an approximation for the outcome of the hiring process. In Appendix A.1, we describe the algorithm (required for the construction of later proofs), and we also prove that—as long as workers sufficiently dislike being

⁷Coles, Kushnir and Niederle (2013) 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.

⁸Interviewing is costless from a worker's perspective and a worker can elect to rank a firm as unacceptable in the second stage deferred acceptance algorithm. If we assume that workers' dis-utility from unemployment is sufficiently large (which, as discussed below, is sufficient to ensure that workers report their true preferences in a firm-proposing deferred acceptance algorithm), it can be shown that a worker will not wish to decline any interview.

⁹Although a firm f only learns $\{\delta_{w,f}\}$ for those workers it interviews, as the expected value of hiring any $w \notin \eta(f) < 0$, how each firm ranks workers it does not interview is irrelevant.

¹⁰A *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.




Figure 1: Examples of Interview Assignments between firms $\{f_1, f_2, f_3, \dots\}$ and workers $\{w_1, w_2, w_3, \dots\}$.

unemployed—the only equilibrium of the DAA is for both firms and workers to utilize their true preferences learned from the first-stage interviewing process.

3 The Interview Assignment Problem

Having specified the outcome of the second stage given any realization of preferences, we now focus on the equilibrium first stage assignment of interviews.

3.1 Interview Assignments

Note any interview assignment η can be represented by a bi-partite network or graph $g(\eta)$ on nodes $\{F, W\}$, where we say $fw \in g(\eta)$ if f interviews w under interview assignment η (i.e., $w \in \eta(f)$). It will occasionally be convenient to represent interview assignments as networks; we provide some examples of interview assignments in network form in Figure 1.

3.2 Equilibrium Analysis

We now turn to the equilibrium determination of an interview assignment. Formally, each 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. We say that firm f (always) interviews K workers if for any $W_f \subset W$, $\nu_f(W_f) > 0$ implies $|W_f| = K$; and we say that firm f interviews K workers “(uniformly) at random” if $\nu_f(W_f) > 0$ if and only if $|W_f| = K$, 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:

$$\int \int \pi_f(W_f, W_{-f}) d\nu_f d\nu_{-f} \geq \int \int \pi_f(W_f, W_{-f}) d\nu'_f d\nu_{-f} \quad \forall \nu'_f, f,$$

where $\pi_f(W_f, W_{-f})$ is the expected profit 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 profit 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. Imposing subgame perfection implies firms anticipate the outcome of the equilibrium play in the second stage match when conducting interviews in the first stage.

We are first interested in examining equilibria in which all firms interview the same number of workers. We show that even among this set of equilibria, the expected number of unemployed workers or the costs expended on interviewing can vary significantly.

First, consider an assignment where every firm randomly selects y workers to interview. For certain values of c , this assignment is an equilibrium:

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

(All proofs are located in the appendix.) 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; however, this may lead some workers to receive more interviews than others.

Next, consider an assignment in which every firm selects a single subset of workers so that every firm *and every worker* receives the same number of interviews. This is equivalent to firms employing a strategy in which each firm f assigns probability 1 to one particular element $W_f \in \mathcal{P}(W)$; if these strategies constitute an equilibrium, we say the correspondence $\eta : \eta(f) = W_f \forall f$ is an pure-strategy equilibrium interview assignment. Such an equilibrium may also exist:

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

Both existence proofs in this section rely on a result established in in the Appendix (Lemma A.2) 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 k th interview it conducts (holding everyone else's actions fixed) than it gains from the $k + 1$ th. Thus, if the cost of interviewing is less than the gain from interviewing the k th worker but greater than the gain from interviewing the $k + 1$ th worker for a firm, every firm interviewing k workers will be an equilibrium as it will not wish to add, remove, or replace any workers in its set of interviewees.¹¹

¹¹Due to integer constraints, there may exist values of c for which no equilibrium where all firms interview the same number of workers 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 $G' < G$ as 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 where all firms interview the same number of workers exists—neither everyone

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; there are a large number of different symmetric pure-strategy equilibria in which k interviews are conducted by each firm and 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, however, the following example shows they can achieve a better outcome in a pure-strategy equilibrium than mixed:

Example 3.1. Consider $N = 3$ and the following symmetric interview assignment in which each firm interviews two workers and each worker obtains two interviews:

$$\begin{array}{ccc}
 f_1 & f_2 & f_3 \\
 | & | & | \\
 w_1 & w_2 & w_3
 \end{array}
 \quad
 \eta(f_1) = \{w_1, w_2\}, \quad
 \eta(f_2) = \{w_2, w_3\}, \quad
 \eta(f_3) = \{w_3, w_1\} \quad (3.1)$$

Assume $\delta = 1 + \epsilon$ with probability .9 and $\delta = -10 + \epsilon$ with probability 0.1, where ϵ is drawn iid from a continuous distribution with arbitrarily small bounded support $[-\bar{\epsilon}, \bar{\epsilon}]$.¹² This corresponds to the case where a worker is likely to generate positive surplus to a firm, but may be very costly if the worker does not.

In the next subsection, we show how one can explicitly compute each firm’s expected profit from an interview assignment. For now, we note that each firm’s expected profit is $\pi \approx 0.88 - 2c$ and the probability of being unmatched is approximately 0.12. Such an assignment is an equilibrium for $c \in (0.10, 0.24)$.

However, consider the case where each firm now randomly selects 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. In this case it is also possible to compute the expected profits ($\pi \approx 0.85 - 2c$) and probabilities of being unmatched (0.15), indicating firms (and workers) are strictly worse off ex ante by mixing. These are equilibrium strategies for $c \in (0.11, 0.27)$. ■

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). However, there may exist equilibria in which firms interview different numbers of workers (as in a symmetric mixed strategy equilibrium where firms randomize over the number of interviews that they conduct).

¹²The use of ϵ is to ensure the distribution of δ is continuous with no point masses.

3.3 Firms' Expected Profits

We now describe some of the details required to solve example 3.1 and provide insights into the results of this article.

In section 2, we noted that the second stage FOSM between firms and workers can be generated from a firm-proposing DAA in which firms and workers use their true preferences. As a result, each firm's expected profit from interviewing any subset of workers W_f given the actions of other firms W_{-f} can be computed using the logic of the DAA. For illustrative purposes, consider the expected profit of a firm f from interviewing only w :

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

where \bar{F}_w denotes the set of firms that make a job offer to w given all other firms interview the subsets of workers W_{-f} . When we say a firm f makes a job offer to w , we are referring to the event that during *any* stage of the DAA, firm f finds itself proposing to w ; this definition is independent of whether w rejects the offer, holds onto it, or ultimately accepts it. The expected profit 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 smallest δ draw is itself a random variable, known as the j th order statistic which we denote by $\delta_{j:k}$.¹³ By the logic of the DAA, we can then construct the expected profit 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 profit 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 profit from interviewing the subset W_f (where $|W_f| = k$):

$$\begin{aligned} \pi_f(W_f, W_{-f}) = & \Lambda_{k,k} \bar{P}_{(k)}^\eta + \Lambda_{k-1,k} \bar{P}_{(k-1)}^\eta \left(1 - \bar{P}_{(k)}^\eta\right) + \\ & \cdots + \Lambda_{2,k} \bar{P}_{(2)}^\eta \prod_{i=3}^k \left(1 - \bar{P}_{(i)}^\eta\right) + \Lambda_{1,k} \bar{P}_{(1)}^\eta \prod_{i=2}^k \left(1 - \bar{P}_{(i)}^\eta\right) - ck \end{aligned} \quad (3.2)$$

¹³Thus $\delta_{1:k}$ is the smallest and $\delta_{k:k}$ is the highest draw.

where $\Lambda_{j,k} = Pr(\delta_{j:k} \geq 0)E[\delta_{j:k} | \delta_{j:k} \geq 0]$ is the expected value of the j th lowest worker interviewed conditional on him being acceptable times the probability he is acceptable (equivalent to (1) and (2) in the single worker example), and $\bar{P}_{(j)}^\eta$ represents the probability that firm f “wins” its j th lowest worker conditional on making him an offer as a function of the entire interview assignment η —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). Finally, the probability a firm eventually is matched to any worker is simply equation (3.2) with $Pr(\delta_{j:k} \geq 0)$ replacing $\Lambda_{j,k}$:

$$\Pr(\mu(f) \neq f | W_f, W_{-f}) = \sum_{j=1}^k \left[\Pr(\delta_{j:k} \geq 0) \bar{P}_{(j)}^\eta \prod_{i=j+1}^k (1 - \bar{P}_{(i)}^\eta) \right] \quad (3.3)$$

As the probabilities $\bar{P}_{(j)}^\eta$ are a function of the entire interview assignment, 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 round of the DAA, each firm considers itself to have a $1/n$ probability of being the firm that the worker accepts (i.e., of being the highest ranked firm for that worker). Thus, sufficient for determining $\bar{P}_{(j)}^\eta$ is simply the probability distribution over the number of firms that “compete” by making an offer to the j th ranked worker.

Let $P_{(j),i}^\eta$ indicate the probability, conditioned on the event that a firm makes an offer to its j th lowest worker, that i other firms also make that worker a job offer. Then it follows:

$$\bar{P}_{(j)}^\eta = \sum_{i=0}^N \frac{1}{i+1} P_{(j),i}^\eta$$

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

Example 3.2. Let $N = 4$, and consider the interview assignment η depicted in Figure 1(a). As in example 3.1, let $\delta = 1 + \epsilon$ with probability .9 and $\delta = -10 + \epsilon$ with probability 0.1.

As all firms have symmetric interview assignments, any firm’s profits can be expressed using (3.2) with the same values for each $P_{(j),i}^\eta$:

$$\pi = \Lambda_{2,2} \underbrace{\left(P_{(2),0}^\eta + \frac{1}{2} P_{(2),1}^\eta \right)}_{\bar{P}_{(2)}^\eta} + \Lambda_{1,2} \underbrace{\left(1 - P_{(2),0}^\eta - \frac{1}{2} P_{(2),1}^\eta \right)}_{1 - \bar{P}_{(2)}^\eta} \underbrace{\left(P_{(1),0}^\eta + \frac{1}{2} P_{(1),1}^\eta \right)}_{\bar{P}_{(1)}^\eta} - 2c \quad (3.4)$$

where the first term is the expected gain times the probability of hiring the most preferred worker, and the second term is the expected gain times the probability of hiring the second most preferred worker (given it lost the first choice worker).

Consider f_1 . Without loss of generality, assume f_1 prefers w_1 to w_2 . Assume now that f_1

was rejected by w_1 , and now makes an offer to w_2 . The probability f_1 faces competition for w_2 conditional on having been rejected by w_1 is given by:

$$P_{(1),1}^\eta = \underbrace{\frac{1}{2}0.99}_{(1)} + \underbrace{\frac{1}{2}0.81 \left[\frac{1}{2} \frac{0.99}{2} \right]}_{(2)}$$

where (1) represents the probability that f_2 's most preferred acceptable worker is w_2 , and (2) represents the probability that w_2 is f_2 's second most preferred worker, is also acceptable, and f_2 is rejected by w_3 . As f_1 could only have been rejected by w_1 if f_4 hired w_1 , f_3 faces no competition for w_4 , and thus f_2 can only be rejected by w_3 if w_3 is f_3 's top acceptable choice.

Next, examine the probability f_1 faces competition for w_1 conditional on making him an offer; this is given by:

$$P_{(2),1}^\eta = \underbrace{\frac{1}{2}0.99}_{(1)} + \underbrace{\frac{1}{2}0.81 \left[\frac{\tilde{P}_{(2),1}^\eta}{2} \right]}_{(2)}$$

where (1) is the probability that f_4 's top acceptable worker is also w_1 , and (2) is the probability that f_4 's top acceptable worker is w_4 , is rejected by w_4 , and then subsequently makes an offer to w_1 . Firm f_4 can only lose to w_4 if f_3 competes for the same worker, which occurs with probability $\tilde{P}_{(2),1}^\eta$. This will be different than the probability f_1 faces competition for w_1 , because $\tilde{P}_{(2),1}^\eta$ is conditional on f_1 making an offer to w_1 .

To determine $\tilde{P}_{(2),1}^\eta$, note if w_4 is f_4 's top acceptable choice, f_4 will face competition only if f_3 's top acceptable choice is also w_4 , or if f_3 is rejected by w_3 first and then offers to w_4 . However, because w_3 will only reject f_3 if f_2 makes w_3 an offer, and f_2 will do so only if w_3 is its top acceptable choice (otherwise, f_2 would have hired w_2 as f_1 is not competing for w_2 given it made an offer to w_1 first), we see $\tilde{P}_{(2),1}^\eta = P_{(1),1}^\eta$.

As $E[\delta|\delta \geq 0] = 1$, $\Lambda_{2,2} = 0.99$ and $\Lambda_{1,2} = 0.81$; also, noting $P_{(j),0}^\eta = 1 - P_{(j),1}^\eta$ for $j = 1, 2$ allows us to solve for a firm's expected profit $\pi \approx 0.86 - 2c$ from (3.4). Thus, if a worker can generate 100 thousand surplus for a firm with 90% probability or lose 1 million with 10% probability, a firm will obtain in expectation approximately 86 thousand minus the cost of two interviews.

Furthermore, the probability that a firm remains unmatched is

$$\Pr(\delta_{2:2} < 0) + \Pr(\delta_{1:2} > 0) \left[\left(\frac{P_{(2),1}^\eta}{2} \right) \left(\frac{P_{(1),1}^\eta}{2} \right) \right] + \Pr(\delta_{2:2} > 0, \delta_{1:2} < 0) \left[\frac{P_{(2),1}^\eta}{2} \right] \approx 0.14$$

■

4 Overlap

It is not surprising that coordinating on a pure-strategy equilibrium as opposed to playing a mixed-strategy equilibrium can increase the probability of employment, as demonstrated by example 3.1

and supported by simulation results in the next section. 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

Why does the identity of other firms matter? Consider the decision of f choosing to interview an additional candidate when it is already interviewing w . Firm f can choose between w' and w'' who each already have the same number of interviews, except w' also happens to be interviewing with the same firms interviewing w , whereas w'' is not. It turns out, the distinction between w' and w'' is not trivial— f will strictly prefer to interview w' . This is due to the fact that if f loses its first choice worker (be it w or w') to f' , then f will face less “competition” among firms for its second choice worker because 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 fewer firm is then “competing” for its next highest ranked worker.

We discuss our main results by focusing on a particular set of interview assignments that we call *symmetric*:

Definition 4.1. An interview assignment η is *symmetric* if and only if for any $f', f'' \in F$, there exist bijections $\gamma^F : F \rightarrow F$ and $\gamma^W : W \rightarrow W$ where $\gamma^F(f') = f''$ such that for all $f \in F$, $w \in W$, $f'w \in g(\eta)$ if and only if $\gamma^F(f)\gamma^W(w) \in g(\eta)$.

In other words, in a symmetric interview assignment, any firm can switch its position in a network with any other firm and an “identical” or *isomorphic* network can be obtained via a relabeling of all other nodes. Not only does a symmetric network imply that each firm and worker receives the same number of interviews (given there are the same number of firms and workers), but also (because agents are ex ante symmetric) every agent will be indifferent between keeping and swapping all their interviews with any other agent.

Consider the interview assignments depicted in Figure 1. Both (a) and (b) are symmetric assignments where each firm and worker has exactly 2 interviews. However, although (c) and (d) are assignments where each firm and worker has exactly 3 interviews, (c) is symmetric and (d) is not. To see why, note that in (d), f_1 and f_2 interview the exact same candidates, but, neither firms f_3 , f_4 , or f_5 share the exact same set of candidates with any other firm. Thus, if f_1 switched positions with f_3 in the network, there exists no relabeling of all other nodes such that f_1 interviews the same three candidates with any other firm.

Finally, we note that a symmetric interview assignment implies not only that each *component* of the network contains the same number of firms and workers, but within each component the interview assignment is also symmetric.^{15,16} In Figure 1, all assignments have 1 component except

¹⁴In 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.

¹⁵A *component* of a network is a distinct subgraph in which all nodes are connected via some path in the network (cf. Jackson, 2008).

¹⁶It is straightforward to show that there must be the same number of firms as workers within any component

for (b) which has 2.

We introduce a partial ordering over different symmetric interview assignments to capture this idea:

Definition 4.2. Consider two symmetric interview assignments η and η' in which both firms and workers receive k interviews. We say η exhibits *greater overlap* than η' if and only if $C(g(\eta)) < C(g(\eta'))$, where $C(g(\cdot))$ is the number of firms (or workers) in any component of $g(\cdot)$. We define the *degree of non-overlap* of assignment η to be $d(\eta, k) \equiv (C(g(\eta)) - k)$. We say that η exhibits *perfect overlap* if $d(\eta, k) = 0$ and *minimal overlap* if $d(\eta, k) = N - k$.

Recall for any symmetric interview assignment where each firm and worker receives k interviews, the associated graph will have the same number of agents across each component. The smaller the size of each component across different assignments, the greater the sense in which firms “share” the same set of interviewees with other firms. Indeed, the smallest number of firms a component can contain is k , which occurs under *perfect overlap*; this implies if any two firms interview the same worker, those two firms interview the same exact set of workers: i.e., $\eta(f) \cap \eta(f') \neq \emptyset$ implies $\eta(f) = \eta(f') \forall f, f'$.

Although perhaps subtle, the concept of greater overlap can have significant effects as demonstrated by the following example.

Example 4.1. Recall in example 3.2 that a firm’s expected profits was approximately $0.86 - 2c$, and the probability a firm is unmatched was approximately 0.14. Note this interview assignment, given by Figure 1(a), exhibited minimal overlap.

Now take the same setup of example 3.2 except now consider the interview assignment η depicted in Figure 1(b) which exhibits perfect overlap. In this case, if both w_1 and w_2 are acceptable for f_1 , then f_1 is guaranteed to hire at least one of them with certainty: if f_1 loses its top choice worker, it means f_2 hired w_1 and there no longer is competition for w_2 . It then follows that $P_{(1),0}^\eta = 1$ (the probability of facing no competition for the second best worker, given the first best worker rejected the firm). Additionally, the probability f_1 ’s top worker receives another job offer is simply $P_{(2),1}^\eta = \frac{1}{2} \Pr(\delta_{2:2} > 0)$, which is the probability that f_2 ’s top choice worker coincides with f_1 ’s top choice. We thus find a firm’s expected profit is:

$$\pi_f = \Lambda_{2,2} \left(P_{(2),0}^\eta + \frac{1}{2} P_{(2),1}^\eta \right) + \Lambda_{1,2} \left(1 - P_{(2),0}^\eta - \frac{1}{2} P_{(2),1}^\eta \right) - 2c \approx 0.95 - 2c \quad ,$$

of a symmetric interview assignment (as there are N firms and N workers). To see why there must be the same number of firms (and workers) across components, consider an assignment η with two components, the first with exactly K firms and the other with exactly K' firms, $K > K'$. Let f be in the larger component and f' in the other. If the first component contains K firms, this means there exists an alternating sequence of unique firms and workers $\{f_{(1)}, w_{(1)}, f_{(2)}, w_{(2)}, \dots, w_{(K-1)}, f_{(K)}\}$ which contains f such that $f_{(i)} w_{(i)} \in g(\eta) \forall i \in \{1, \dots, K\}$ and $w_{(i)} f_{(i+1)} \in g(\eta) \forall i \in \{1, \dots, K-1\}$. If η is symmetric, then there must exist a bijection $\gamma^F : F \rightarrow F$ such that $\gamma^F(f) = f'$ and a bijection $\gamma^W : W \rightarrow W$ such that $\gamma^F(f_{(i)}) \gamma^W(w_{(i)}) \in g(\eta) \forall i \in \{1, \dots, K\}$ and $\gamma^W(w_{(i)}) \gamma^W(f_{(i+1)}) \in g(\eta) \forall i \in \{1, \dots, K-1\}$. But this implies f' is in a component with at least K firms, which is a contradiction.

and the probability of remaining unmatched is

$$\Pr(\delta_{2:2} < 0) + \Pr(\delta_{2:2} > 0 \& \delta_{1:2} < 0) \left(\frac{1}{2} P_{(2),1}^\eta \right) \approx 0.05 \quad .$$

Hence, we see that with perfect 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 minimal overlap given by Figure 1(a) is an equilibrium for firms to follow for $c \in (0.09, 0.23)$. On the other hand, as long as $c \in (0.05, 0.26)$, the interview assignment depicted in 1(b) with perfect overlap is an equilibrium. Consequently, for any value of $c \in (0.09, 0.23)$, both interview assignments shown in Figure 1(a) and (b) are equilibria, but the latter equilibrium dominates. ■

Thus, there may be many different pure-strategy equilibria that still assign each firm and each worker the same number interviews, but exhibit different degrees to which workers coincide among firms.

We prove that if the probability that a worker is found to be undesirable is sufficiently small, then among the set of symmetric interview assignments in which firms and workers receive the same number of interviews, the assignments that are characterized by perfect overlap minimize the (ex ante) probability that any firm or worker will remain unemployed:

Theorem 4.1. *For any x such that N is divisible by x , there exists $\epsilon > 0$ such that if $\Pr(\delta < 0) < \epsilon$, an interview assignment with perfect overlap with x interviews per firm and worker minimizes expected unemployment over any other symmetric interview assignment in which each firm and worker receives x interviews.*

The intuition for the proof is as follows. Note first that the probability that any firm is unmatched (and hence a worker is unemployed) under perfect overlap is bounded above by the probability that a firm finds at least one worker to be undesirable. Second, note that without perfect overlap, the probability that any firm is unmatched is bounded below by the probability that a set of preferences are realized in which all of a firm’s interviewees are matched to other firms; by an application of Hall’s marriage theorem, we can show that a set of preferences resulting in such a match will always exist in a symmetric interview assignment without perfect overlap. If the probability that a worker is found to be undesirable is sufficiently small, the upper bound on the probability of unemployment under perfect overlap is smaller than the lower bound on the probability of unemployment under “non-perfect” overlap, and the result follows.

Remarks. It is worth noting that this 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

Table 1: Comparison Between Pure and Mixed-Strategy Equilibria

		Pure-Strategy (Perfect Overlap)			Mixed-Strategy (N=10)			Mixed-Strategy (N=100)		
		Prob.	Exp.	Equil.	Prob.	Exp.	Equil.	Prob.	Exp.	Equil.
		Unemp.	Profit	Costs	Unemp.	Profit	Costs	Unemp.	Profit	Costs
# Interviews Per Agent	1	0.10	0.45	(0.10,0.45)	0.39	0.30	(0.18,0.30)	0.40	0.30	(0.18,0.30)
	2	0.05	0.54	(0.06,0.20)	0.23	0.43	(0.11,0.17)	0.25	0.42	(0.12,0.17)
	3	0.03	0.60	(0.04,0.12)	0.15	0.50	(0.09,0.12)	0.18	0.49	(0.09,0.11)
	4	0.02	0.64	(0.03,0.09)	0.10	0.56	(0.06,0.09)	0.13	0.54	(0.06,0.08)
	5	0.02	0.67	(0.03,0.06)	0.07	0.60	(0.05,0.06)	0.10	0.58	(0.05,0.06)
	6	0.01	0.69	(0.02,0.05)	0.05	0.64	(0.04,0.05)	0.08	0.61	(0.04,0.05)
	7	0.01	0.71	(0.02,0.04)	0.03	0.68	(0.03,0.04)	0.06	0.63	(0.04,0.04)
	8	0.01	0.73	(0.02,0.03)	0.02	0.71	(0.03,0.03)	0.05	0.65	(0.03,0.04)
	9	0.01	0.74	(0.01,0.03)	0.01	0.73	(0.02,0.03)	0.05	0.67	(0.03,0.03)
	10	0.00	0.76	(0.01,0.02)	0.00	0.76	(0.00,0.02)	0.04	0.69	(0.02,0.03)

Notes: 10,000 simulation runs. $\delta_{w,f}$ is distributed iid uniformly on $[0, 1]$ with 90% probability, and $\delta_{w,f} \leq -10$ with 10% probability. “Prob. Unemp.” is the probability that a firm or worker remains unemployed, “Exp. Profit” does not include interviewing costs, and “Equil. Costs” is the range of c such that the assignment in question is an equilibrium.

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 non-perfect overlap or random interviewing introduces matching frictions and strictly increases unemployment.

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 A.3, if there exists an equilibrium with perfect overlap for $N = x$ for a given value of c , then as N grows large there exists a correlated equilibrium in which each firm interviews x workers and achieves perfect overlap with probability close to 1.

5 Simulation Results

This article has examined two forms of coordination in the assignment of interviews: (i) pure vs. mixed-strategy (i.e., even though firms all conduct the same number of interviews, some workers may receive different numbers); and (ii) the degree of non-overlap across assignments. In this section, we simulate interview markets to evaluate the impact of these forms of coordination on employment.

5.1 Pure vs. Mixed-Strategy Equilibria

To compare outcomes under pure and mixed-strategy equilibria (i.e., when firms conduct the same number of interviews, but may not be able to coordinate on the set of workers that they interview), we simulate markets where $\delta_{w,f}$ is distributed iid uniformly on $[0, 1]$ with 90% probability, and $\delta \leq -10$ with 10% probability; this also ensures the expected value of hiring a random worker without interviewing him is negative. We allow number of interviews each firm conducts to vary from 1 to 10, where firms coordinate on a perfect overlap equilibrium, or play a mixed-strategy equilibrium in which the number of total agents in the market varies from 10 to 100.¹⁷

Table 1 reports the probability that any given firm is unemployed, its expected profit not including interviewing costs, and the range of interviewing costs such that the given assignment is an equilibrium. When firms can coordinate on a pure-strategy equilibrium, the probability of unemployment is bounded above by 10% (the probability a worker is unacceptable); however, under a mixed-strategy equilibrium it is much higher, particularly when firms interview a small number of workers. Furthermore, as the number of agents in the market grows larger, a mixed-strategy equilibrium performs worse; a pure-strategy equilibrium (because the size of each component remains fixed) is not affected.

What is also evident from this exercise is that a pure-strategy equilibrium can result in either more or fewer interviews being conducted than a mixed-strategy equilibrium.¹⁸ Although we do not prove this result generally, in this example, for any c under which both a symmetric pure- and mixed-strategy equilibria exist, firms receive strictly higher expected profit (and higher probability of being matched) under the pure-strategy equilibrium.

5.2 Comparing Overlap

We now focus on pure-strategy equilibria with symmetric interviewing assignments, and vary the degree of non-overlap across assignments and the probability a worker is determined to be unacceptable (but firms still do not hire workers they do not interview).¹⁹

Table 2 reports the probability that a firm is unmatched across different specifications. Simulations suggest that as the degree of non-overlap increases for a given number of interviews or probability that a worker is found undesirable, the probability of unemployment increases. Furthermore, the increase is the largest when moving from perfect overlap to an assignment with degree of non-overlap of 1; as the degree of non-overlap grows, the rate at which unemployment increases

¹⁷For all simulations involving pure-strategy equilibria, we assume that N is equal to the number of firms in each component (i.e., the number of interviews per worker plus the degree of non-overlap of the interview assignment); the number of components in the market does not affect results.

¹⁸For example, if $c \in (0.30, 0.45)$, a pure-strategy equilibrium exists where every firm exactly interviews 1 worker; however, no mixed-strategy equilibrium exists where all firms randomly interview at least 1 worker. However, if $c = 0.06$, a pure-strategy exists where firms interview 4 workers; a mixed-strategy equilibrium exists where firms randomly select 5 workers.

¹⁹There may be many symmetric interview assignments with the same degree of non-overlap that differ in expected profits and unemployment. For our simulations, we construct the symmetric interview assignments according to the procedure used in the proof of Proposition 3.2.

Table 2: Probability of Being Unemployed

		Degree of non-overlap ($ C(g(\eta)) - k$)										
		0	1	2	3	4	5	10	25	50	100	
$(Pr(\delta_{w,f}) < 0) = 0\%$	# Interviews Per Agent	1	2	3	4	5	6	7	8	9	10	
	1	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	2	0.0%	6.3%	9.5%	10.5%	10.9%	11.1%	11.1%	11.1%	11.1%	11.1%	11.1%
	3	0.0%	2.2%	3.9%	4.8%	5.3%	5.5%	5.7%	5.7%	5.7%	5.7%	5.7%
	4	0.0%	1.4%	2.5%	3.4%	4.1%	4.5%	5.3%	5.5%	5.4%	5.4%	5.4%
	5	0.0%	0.9%	1.6%	2.2%	2.7%	3.1%	3.9%	4.2%	4.2%	4.1%	4.1%
	6	0.0%	0.6%	1.0%	1.5%	1.9%	2.2%	3.2%	3.7%	3.7%	3.7%	3.7%
	7	0.0%	0.4%	0.7%	1.0%	1.3%	1.6%	2.5%	3.1%	3.1%	3.1%	3.1%
	8	0.0%	0.3%	0.5%	0.7%	0.9%	1.2%	1.9%	2.7%	2.7%	2.7%	2.7%
	9	0.0%	0.2%	0.4%	0.5%	0.7%	0.9%	1.5%	2.3%	2.4%	2.4%	2.4%
10	0.0%	0.1%	0.3%	0.4%	0.5%	0.6%	1.2%	1.9%	2.1%	2.1%	2.1%	
$(Pr(\delta_{w,f}) < 0) = 10\%$	# Interviews Per Agent	1	2	3	4	5	6	7	8	9	10	
	1	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%
	2	5.5%	11.4%	13.6%	14.1%	14.4%	14.4%	14.4%	14.4%	14.4%	14.4%	14.4%
	3	3.4%	5.7%	7.3%	8.0%	8.3%	8.4%	8.5%	8.5%	8.5%	8.5%	8.5%
	4	2.3%	3.9%	4.9%	5.7%	6.3%	6.6%	7.1%	7.2%	7.2%	7.2%	7.2%
	5	1.6%	2.6%	3.4%	4.0%	4.5%	4.8%	5.4%	5.6%	5.6%	5.6%	5.6%
	6	1.2%	1.8%	2.4%	2.9%	3.3%	3.6%	4.4%	4.7%	4.7%	4.7%	4.7%
	7	0.9%	1.4%	1.8%	2.1%	2.4%	2.7%	3.5%	3.9%	4.0%	3.9%	3.9%
	8	0.7%	1.0%	1.3%	1.6%	1.8%	2.1%	2.8%	3.4%	3.4%	3.4%	3.4%
	9	0.5%	0.8%	1.0%	1.2%	1.4%	1.6%	2.3%	2.9%	3.0%	3.0%	3.0%
10	0.4%	0.6%	0.8%	0.9%	1.1%	1.3%	1.8%	2.5%	2.6%	2.6%	2.6%	
$(Pr(\delta_{w,f}) < 0) = 25\%$	# Interviews Per Agent	1	2	3	4	5	6	7	8	9	10	
	1	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%
	2	15.1%	19.4%	20.4%	20.6%	20.6%	20.7%	20.7%	20.7%	20.7%	20.7%	20.7%
	3	10.1%	12.0%	13.2%	13.6%	13.7%	13.8%	13.9%	13.9%	13.9%	13.9%	13.9%
	4	7.2%	8.7%	9.6%	10.2%	10.5%	10.6%	10.9%	10.9%	10.9%	10.9%	10.9%
	5	5.4%	6.4%	7.2%	7.6%	8.0%	8.2%	8.5%	8.6%	8.6%	8.6%	8.6%
	6	4.2%	4.9%	5.5%	5.9%	6.2%	6.5%	7.0%	7.1%	7.1%	7.1%	7.1%
	7	3.3%	3.8%	4.3%	4.6%	4.9%	5.1%	5.7%	5.9%	5.9%	5.9%	5.9%
	8	2.6%	3.1%	3.4%	3.7%	4.0%	4.2%	4.7%	5.1%	5.1%	5.1%	5.1%
	9	2.1%	2.5%	2.8%	3.0%	3.2%	3.4%	4.0%	4.4%	4.4%	4.4%	4.4%
10	1.7%	2.0%	2.2%	2.4%	2.6%	2.7%	3.3%	3.8%	3.8%	3.8%	3.9%	

Notes: 10,000 simulation runs. Each cell indicates the probability a firm is unmatched or a worker is unemployed given every firm and worker receives 1-10 interviews, and the degree of non-overlap increases. Each panel of the table varies the probability of any worker being unacceptable from 0% to 10% to 25%.

falls, and once the degree of non-overlap reaches 10, incremental reductions are marginal.

Some intuition for this result can be gleaned from Figure 2: in panels (a)-(c), the first components of three different symmetric interview assignments $\{\eta^0, \eta^1, \eta^2\}$ are depicted in which each firm and worker receives 2 interviews; η^0 has perfect overlap, whereas the other two have less overlap. If all workers are acceptable to f_1 , then under η^0 there is no possible stable match which results in firm f_1 being unmatched. However, there is one potential stable match under η^1 and two under η^2 in which f_1 is unmatched; as overlap decreases further, even more stable matches (supportable by some set of preferences) can be constructed.

Overall, the gains to overlap can be substantial: focusing on settings where firms interview more than one worker and where the probability of a worker being unacceptable is 10%, the chance that a firm or worker is unmatched is approximately 3-4 times smaller under a perfect overlap assignment




Figure 2: (a)-(c) Symmetric interview assignments where each firm and worker receive 2 interviews, each with different levels of overlap. (d)-(f) Given interview assignments η^1 and η^2 , only μ^1, μ_1^2 , and μ_2^2 are stable matches in which no worker is unacceptable and firm f_1 is unmatched.

than under an assignment with a degree of non-overlap larger than 10.

6 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 article 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 optimally overlap their interviews. Consequently, conditional on a certain number of interviews being conducted, institutions that limit the number of interviews workers can receive or 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. Within this model, there are several open questions, including: characterizing equilibrium existence when firms may interview different numbers of workers; whether for fixed costs of interviewing, pure-strategy equilibria always outperform mixed-strategy equilibria; and whether for certain types of symmetric interview assignments, greater overlap guarantees lower expected unemployment.

It also is worth mentioning several aspects of interviewing markets that are beyond the scope

of this article but are interesting for future work. For example, our article assumes that firms bear the full cost of interviewing; this may be a reasonable approximation in some markets, but less so in others (e.g., 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 reduces unemployment 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 Appendix

A.1 Second Stage Matching

The *firm optimal stable matching* (FOSM) used in the second stage of our game is the outcome of a firm-proposing deferred acceptance algorithm (DAA) for employment when firms and workers utilize their true preferences (Gale and Shapley, 1962). We first describe the algorithm (whose features are leveraged in proving this article’s main results), and then prove that—as long as workers place sufficient disutility on being unemployed—firms and workers will utilize their true preferences in equilibrium when the firm-proposing DAA is used.

Before proceeding, note that in the main text, we have assumed that workers have only ordinal preferences over firms; this was sufficient as long as we assumed the outcome of the second stage was the FOSM. However, if we wish to interpret the FOSM as arising from a firm-proposing DAA, we will need to make assumptions on the relative utility of being employed versus unemployed (see Lemma A.1 to follow); this is the only instance in which we assume workers have cardinal preferences.

The Deferred Acceptance Algorithm (DAA). In the firm-proposing DAA, each firm f utilizes preferences \mathcal{P}_f over acceptable workers and each worker w utilizes preferences \mathcal{P}_w . The algorithm proceeds as follows:

- Step 1: Each firm makes a job offer to its highest-ranked worker (or, if all workers are unacceptable, 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 any step in which no firm’s offer is rejected; at this point all firms have either a worker holding its job offer, or has no workers it wishes to make an offer to that has not already rejected it. At this point, 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.

Truthful Revelation of Preferences. A matching mechanism such as the DAA can be susceptible to “gaming” in that participants may find it preferable to misrepresent their true preferences; indeed, there are several negative results in the two-sided matching literature on the existence of mechanisms that elicit truth-telling from participants. Nonetheless in this section, we prove that as long as workers place a high enough disutility of being unemployed, firms and workers will utilize their true preferences realized during the interview stage under the firm-proposing DAA. This is formally stated and proven in the following lemma:

Lemma A.1. *There exists a finite $\beta > 0$ such that if $\zeta - \zeta_0 > \beta(\bar{\zeta} - \zeta)$, the only equilibrium of the firm-proposing deferred acceptance algorithm is for both workers and firms to report their true preferences.*

Proof of Lemma A.1. As 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: some worker w rejects some firm f ’s offer (instead of holding onto it or accepting it); f then makes an offer to another worker w' who accepts and rejects some other firm f' ; this f' in turn proposes to another worker, who then may reject another firm; and so on, until a firm f'' , who was rejected by another worker, makes an offer to the original worker w , and w prefers f'' to f . 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 w is considering rejecting a firm’s offer that it finds preferable to the one w is already holding. In order for this to be profitable, f upon being rejected (conditional on making w an offer) must propose to another worker that already has an existing offer from another firm f' , and that worker must prefer f to f' . The probability that f is preferred to any f' by another worker is at most $1/2$, which is an upper bound on the probability that rejecting a firm leads to a profitable manipulation. Thus the gain to rejecting a firm’s offer is bounded by $(\bar{\zeta} - \zeta)/2$, where the term in parenthesis is the maximum gain possible to w 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 $(\lambda)^{L-1}$, where $\lambda = \Pr(\delta_{w,f} < 0) > 0$ is the probability that a worker is unacceptable to a firm. Consequently, by rejecting firm j , he risks losing at least

$(\lambda)^{L-1}(\zeta - \zeta_0)$. Clearly as long as $\beta > (\lambda)^{L-1}/2$, the inequality holds and no worker will find it profitable to reject any firm's offer (unless it is holding an offer from a more preferred firm).

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 $(1-p)^{n-1} \times p$ (as 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

A.2 Proofs

Proof of Proposition 3.1. Assume each firm selects $x \leq N$ workers to interview uniformly at random: i.e., each firm plays a strategy ν_f which assigns equal positive probability to all distinct subsets of workers of size x . We show that there exists a c such that this is an equilibrium.

Consider firm f . If every other firm is randomizing uniformly, firm f is indifferent over interviewing any particular worker, so for any given number of interviews k that f conducts, any set of workers is optimal. Let

$$g_f(k, \nu_{-f}; x) = E_{\nu_{-f}}[\pi_f(W_f, W_{-f}) - \pi_f(W'_f, W_{-f}) \mid |W_f| = k, |W'_f| = k-1, |W_{-f}| = x] \quad (\text{A.1})$$

denote f 's expected gain to interviewing an additional k th worker (not including costs) given all other firms are interviewing x workers uniformly at random, and $\pi_f(\cdot)$ is defined in (3.2).

We first prove the following lemma:

Lemma A.2. $g_f(k, \nu_{-f}; x)$ is strictly decreasing in k .

Proof. Let $w' = \mu^{k-1}(f)$ be the worker f would have matched to if it only conducted $k-1$ interviews. Note f only benefits from interviewing the k th worker (denoted w_k) if $\mu^k(f) = w_k$, as otherwise f would still be matched to w' . We decompose $g_f(k, \nu_{-f})$ as follows:

$$g_f(k, \nu_{-f}; x) = \underbrace{\Pr(\mu^k(f) = w_k)}_{(1)} \times \underbrace{(E[\delta_{w_k, f} - \delta_{w', f} \mid \mu^k(f) = w_k \ \& \ \mu^{k-1}(f) = w'])}_{(2)} \quad (\text{A.2})$$

- (1) *Probability that interviewing w_k results in being matched to w_k :* As f is indifferent between interviewing any two workers, it must be that $\Pr(\mu^k(f) = w_k) = \Pr(\mu^k(f) = w)$ for any w that f is already interviewing. Intuitively, it is clear that interviewing an additional worker strictly reduces the probability that f is eventually matched to a given worker it already is interviewing. Formally, consider any w , and assume that $\mu^{k-1}(f) = w$. The probability that f is still matched to w after k interviews is:

$$\Pr(\mu^k(f) = w \mid \mu^{k-1}(f) = w) = \Pr(\delta_{w, f} \geq \delta_{w_k, f}) + \Pr(\delta_{w, f} < \delta_{w_k, f}) \times \Pr(\mu^{k-1}(w_k) P_{w_k} f)$$

where the first term is the probability that f does not make an offer to w_k after interviewing him, and the second term is that f makes an offer to w_k but is rejected. This is strictly less than 1 because $\Pr(\mu^{k-1}(w_k) P_{w_k} f) < 1$. Thus (1) is strictly decreasing in k .

- (2) *Expected gain from being matched to w_k given f was matched to w' :* Note (2) is equivalent to $E[\delta_{w_k, f} - \delta_{w', f} \mid \delta_{w_k, f} > \delta_{w', f}]$: i.e., f would only make an offer to w_k after interviewing him if $\delta_{w_k, f} > \delta_{w', f}$. As $\delta_{w', f}$ is weakly increasing in k (i.e., the value of f 's existing match cannot fall by conducting more interviews), and because Assumption 2.1 implies that $[E[\delta_{w_k, f} - y \mid \delta_{w_k, f} \geq y]]$ weakly falls as y increases, it follows that (2) is at least weakly decreasing in k .

As (1) is strictly and (2) is at least weakly decreasing in k , the lemma is proved. \square

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

Proof of Proposition 3.2. Assume N is divisible by x . Construct a symmetric perfect-overlap interview assignment η as follows: divide firms and workers into $M \equiv N/x$ components, where component $m \in \{1, \dots, M\}$ contains firm and worker indexed by $\{(m-1)x+1, \dots, (m-1)x+x\}$. In the first component ($m=1$), let $\eta(f_i) = \{w_i, w_{[i+1]_x}, \dots, w_{[i+x-1]_x}\}$ for each $i \in \{1, \dots, x\}$, where we define $[a]_b \equiv \text{mod}(a-1, b) + 1$;²⁰ repeat the same assignment within each component for each $m \geq 2$.

W.l.o.g, consider $f \equiv f_1$, and let

$$g_f(k, \eta_{-f}) = [\pi_f(\{w_1, \dots, w_k\}, W_{-f}) - \pi_f(\{w_1, \dots, w_{k-1}\}, W_{-f}) \mid W_{-f} = \{\eta(f')\}_{\forall f' \in F, f' \neq f}]$$

be f 's expected gain to interviewing w_k when it is already interviewing $\{w_1, \dots, w_{k-1}\}$ and all other firms are interviewing according to η .

Lemma A.3. $g_f(k, \eta_{-f})$ is strictly decreasing in k .

Proof. The proof is similar to Lemma A.2: decompose $g_f(k, \eta_{-f})$ as follows:

$$g_f(k, \eta_{-f}) = \underbrace{\Pr(\mu^k(f) = w_k)}_{(1)} \times \underbrace{(E[\delta_{w_k, f} - \delta_{w', f} \mid \mu^k(f) = w_k \ \& \ \mu^{k-1}(f) = w'])}_{(2)} \quad (\text{A.3})$$

For $k \in \{1, \dots, x\}$: Within the first component, (1) is strictly decreasing in k by the same logic as in Lemma A.2: as $\{w_1, \dots, w_x\}$ are identical from the perspective of f (they are all receiving $x-1$ interviews from $\{f_2, \dots, f_x\}$), the probability of being matched to the k th worker (given the first k workers are interviewed) is the same as being matched to any $w \in \{w_j\}_{j < k}$. As before, because interviewing an additional worker must reduce the probability f is matched to a given worker, (1) is strictly decreasing in k . Also, as before, (2) is weakly decreasing in k .

For $k = x+1$: For the $k+1$ interview, f must interview a worker who already has x interviews, and is in a different component of the graph associated with η . To see that the gain to interviewing w_{x+1} is less than the gain to interviewing the w_x , note that (1) can be further decomposed as:

$$\Pr(\mu^k(f) = w_k) = \underbrace{\Pr(\delta_{w_k, f} > \delta_{w', f})}_{(1a)} \times \underbrace{\Pr(f P_{w_k} \mu^{k-1}(w_k) \mid \delta_{w_k, f} > \delta_{w', f})}_{(1b)} \quad (\text{A.4})$$

where (1a) is the probability that f makes an offer to the k th worker, and (1b) is the probability that w_k accepts an offer from f given f proposes to w_k (where $\mu^{k-1}(w_k)$ is the firm that w_k would match to if f had conducted $k-1$ interviews). (1a) is decreasing with k because the value of $\delta_{w', f}$ is weakly increasing as more interviews are made. Note that (1b) is strictly increasing for $k \in \{1, \dots, x\}$ (i.e., within a component). To see why, consider the probability w_1 accepts f 's offer conditional on f having made an offer, and examine the difference if f interviews the first $k-1$ versus k workers within the first component. If f interviews k candidates and $w_k P_f w_1$, then f will propose to w_1 if and only if w_k rejected f ; however, if f only interviewed $k-1$ candidates, it would have proposed to w_1 regardless of where it ranked w_k (because it didn't interview w_k). As w_k rejecting f raises the probability w_1 accepts f 's offer (as there are now fewer firms that w_1 can match to), the probability of w_1 accepting an offer conditional on one being made strictly increases in k . Furthermore, as (1b) is higher for $k=1$ than $k=x+1$ (as w_1 receives fewer interviews than w_{k+1}), we see that (1b) is higher for $k=x$ compared to $k=x+1$. Thus, given (2) is still decreasing, the gain to interviewing w_{k+1} is less than the gain to interviewing w_k .

For $k = \{x+2, \dots, N\}$: As f is indifferent between interviewing any set of workers *within* the same component, the same logic as with $k \in \{1, \dots, x\}$ can be used to show that $g_f(k, \eta_{-f})$ is strictly decreasing in k for $k \in \{mx+2, \dots, mx+x\}$, $m = \{1, \dots, M-1\}$. Finally, the same logic for $k = x+1$ can be applied to show that $g_f(k, \eta_{-f})$ is strictly decreasing between components (i.e., $k = (m+1)x+1$, $m \in \{1, \dots, M-2\}$). \square

Note that Lemma A.3 implies that if f were to interview k workers where $k \leq x$, it is optimal for f to interview only workers contained within $\{w_1, \dots, w_x\}$; otherwise, f would have a profitable deviation by swapping any worker that it is interviewing not in that set with one within it (as doing so would strictly increase its profits). Furthermore, by Lemma A.3, we can find $c \in (g_f(x, \eta_{-f}), g_f(x-1, \eta_{-f}))$ such that η is an equilibrium interview assignment: i.e., no individual firm will wish to interview more or fewer than x candidates or change the set of workers that it is interviewing. \square

Proof of Theorem 4.1. Consider a perfect overlap interview assignment η where every firm and worker receives $k = x$ interviews. Under this assignment, the ex ante probability that any firm is unmatched is bounded above by $\Pr(\delta_{1:k} < 0)$: i.e., the probability that the first order statistic out of k draws is negative. To see this, note that if every worker that a firm interviews yields non-negative surplus (which occurs with probability $\Pr(\delta_{1:k} \geq 0)$), then that firm is

²⁰ $\text{mod}(a, b) = a - \lfloor \frac{a}{b} \rfloor b$, where $\lfloor x \rfloor$ represents the greatest integer less than or equal to x .

guaranteed to be matched during the second stage deferred acceptance algorithm: if a firm is not matched by the k th round, it must have been because the previous $k - 1$ workers have rejected that firm and accepted an offer from another firm; thus, by the k th round, no other firms are “competing” for the k th worker and the firm will be able to hire him. Consequently, a firm can only remain unmatched under a perfect overlap interview assignment if at least one worker is found by that firm to be unacceptable; the probability of this occurrence can be no greater than $1 - \Pr(\delta_{1:k} \geq 0)$.

Consider now any symmetric interview assignment η' where every firm and worker receives k interviews and $d(\eta', k) > 0$ (i.e., η' does not have perfect overlap). As all firms and workers are ex ante identical, consider any particular firm f . Let the set of workers interviewed by f be given by $W_f \equiv \eta'(f)$, where $|W_f| = k$ (by construction). If $\eta'(W_f) \setminus f$ represents the set of all firms excluding f that interview workers in W_f , note that $|\eta'(W_f) \setminus f| \geq k$: i.e., there must be at least k other firms that share at least one worker with f under the interviewing assignment η' (otherwise, if $|\eta'(W_f) \setminus f| = k - 1$, η' would have perfect overlap). Furthermore, note that for any subset of workers $W' \subset W_f$, $|\eta'(W') \setminus f| \geq k - 1 \geq |W'|$: i.e., there are at least $k - 1$ other firms than f that interview the workers in W' . Thus, by Hall’s marriage theorem, there exists a one-to-one matching μ such that every worker in W_f is matched to a firm that interviews that worker under η' and is not f . In such a matching, f_1 is unmatched (as all of the workers it interviews are matched with another firm). Consequently, the probability that f is unmatched given η' can be no less than the probability that preferences for all agents are such that each worker in W_f prefers its matched partner under μ to f , and each firm other than f matched to a worker in W_f finds that worker to be desirable and its most preferable.

The upper bound on the probability that any firm is unmatched under perfect overlap assignment η is $\Pr(\delta_{1:k} < 0)$, which is decreasing as $\Pr(\delta < 0)$ decreases, and $\lim_{\Pr(\delta < 0) \rightarrow 0} \Pr(\delta_{1:k} < 0) = 0$. The lower bound on the probability that any firm is unmatched under η' is strictly positive (the realization of a non-empty set of preferences). Thus, there exists an $\epsilon > 0$ such that if $\Pr(\delta < 0) < \epsilon$, this upper bound is below the lower bound, and η has a lower probability of unemployment than η' . \square

A.3 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 an equilibrium where each firm and worker receives x interviews, and the probability each firm is located in a component with exactly x workers and x firms 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. Note the induced interview assignment may not be symmetric, as exactly one component may have more than x firms or workers.

Proposition A.4. *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 $\forall \tilde{N} > \bar{N}$ a correlated equilibrium exists in which each firm interviews x workers, each worker receives x interviews, and with probability $1 - \epsilon$ each firm obtains the same expected profit as under a perfect overlap equilibrium.*

Proof of Proposition A.4. For any \tilde{N} , we can partition the population into $\lfloor \frac{\tilde{N}}{x} \rfloor - 1$ groups of exactly x workers and x firms, and 1 group of $x + (\tilde{N} - \lfloor \frac{\tilde{N}}{x} \rfloor x)$ workers and firms. 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 interviews as in the proof of proposition 3.2. Thus, $\eta(\pi)$ gives each firm and each worker x interviews, and for but only $x + (\tilde{N} - \lfloor \frac{\tilde{N}}{x} \rfloor x)$ workers and firms, firms and workers are located components with only $2x$ total agents (i.e., 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 $\tilde{N} > \bar{N}$, 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 \tilde{N} and small ϵ , we can construct a correlated equilibrium in which firms “coordinate” on a given $\eta(\pi)$ at random, and achieves the utility of a perfect overlap equilibrium with probability $1 - \epsilon$. \square

References

- Atakan, Alp E.** 2006. "Assortative Matching with Explicit Search Costs." *Econometrica*, 74: 667–680.
- Bergemann, Dirk, and Jusso Välimäki.** 2005. "Information in Mechanism Design." In *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*. Vol. 1, , ed. Richard Blundell, Whitney K. Newey and Torsten Persson, Chapter 5. Cambridge, U.K.:Cambridge University Press.
- Blair, Charles.** 1988. "The Lattice Structure of the Set of Stable Matchings with Multiple Partners." *Mathematics of Operations Research*, 13(4): 619–628.
- Burdett, Kenneth.** 1996. "Truncated Means and Variances." *Economics Letters*, 52: 263–267.
- Chade, Hector.** 2006. "Matching with noise and the acceptance curse." *Journal of Economic Theory*.
- Chade, Hector, and Lones Smith.** 2006. "Simultaneous Search." *Econometrica*, 74(5): 1293–1307.
- Chade, Hector, Gregory Lewis, and Lones Smith.** 2014. "Student Portfolios and the College Admissions Problem." *Review of Economic Studies*, 81(3): 971–1002.
- Chakraborty, Archishman, Alessandro Citanna, and Michael Ostrovsky.** 2010. "Two-Sided Matching with Interdependent Values." *Journal of Economic Theory*, 145(1): 85–105.
- Coles, Peter, Alexey Kushnir, and Muriel Niederle.** 2013. "Preference Signaling in Matching Markets." *American Economic Journal: Microeconomics*, 5(2): 99–134.
- Coles, Peter, Benjamin Edelman, Brian Hall, and Nicole Bennett.** 2008. "TheLadders." Harvard Business School Case 9-908-061.
- Das, Sanmay, and Emir Kamenica.** 2005. "Two-Sided Bandits and the Dating Market." *Proceedings of the Nineteenth International Joint Conferences on Artificial Intelligence*, 947–952.
- David, Herbert Aron, and Hakidy Navada Nagaraja.** 2003. *Order Statistics*. . 3 ed., Wiley.
- Dubins, Lester E., and David A. Freedman.** 1981. "Machiavelli and the Gale-Shapley Algorithm." *American Mathematical Monthly*, 88: 485–494.
- Echenique, Federico, and Jorge Oviedo.** 2006. "A Theory of Stability in Many-to-Many Matching Markets." *Theoretical Economics*, 1(2): 233–273.
- Eckhout, Jan, and Philipp Kircher.** 2010. "Sorting and Decentralized Price Competition." *Econometrica*, 78(2): 539–574.
- Gale, David, and Lloyd Shapley.** 1962. "College admissions and the stability of marriage." *American Mathematical Monthly*, 69: 9–15.
- Haeringer, Guillaume, and Myrna Wooders.** 2011. "Decentralized Job Matching." *International Journal of Game Theory*, 40(1): 1–28.

- Hoppe, Heidrun C., Benny Moldovanu, and Aner Sela.** 2009. "The Theory of Assortative Matching Based on Costly Signals." *Review of Economic Studies*, 76(1): 253–281.
- Jackson, Matthew O.** 2008. *Social and Economic Networks*. Princeton University Press.
- Josephson, Jens, and Joel Shapiro.** 2016. "Costly Interviews." *International Journal of Industrial Organization*, 45: 10–15.
- Konishi, Hideo, and M. Utku Ünver.** 2006. "Credible Group Stability in Many-to-Many Matching Problems." *Journal of Economic Theory*, 129(1): 57–80.
- Kranton, Rachel E., and Deborah F. Minehart.** 2000. "Networks Versus Vertical Integration." *RAND Journal of Economics*, 31(3): 570–601.
- Kranton, Rachel E., and Deborah F. Minehart.** 2001. "A Theory of Buyer-Seller Networks." *American Economic Review*, 91(3): 485–508.
- Lee, Robin S., and Michael Schwarz.** 2007. "Signalling Preferences in Interviewing Markets." *Dagstuhl Seminar Proceedings*. Dagstuhl, Germany.
- Lien, Yuan-Chuan.** 2013. "Costly Interviews and Non-Assortative Matchings." mimeo, Hong Kong University.
- Mongell, Susan, and Alvin E. Roth.** 1991. "Sorority Rush as a Two-Sided Matching Mechanism." *American Economic Review*, 81(3): 441–464.
- Nagypal, Eva.** 2004. "Optimal Application Behavior with Incomplete Information." mimeo.
- Niederle, Muriel, and Leeat Yariv.** 2009. "Decentralized Matching with Aligned Preferences." mimeo, Stanford University.
- Roth, Alvin E.** 1982. "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research*, 7: 617–628.
- Roth, Alvin E.** 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, Alvin E., and Marilda A. Oliveira Sotomayor.** 1990. *Two-Sided Matching: A study in game-theoretic modeling and analysis*. Cambridge, MA: Econometric Society Monographs.
- Schweinzer, Paul.** 2008. "Labour Market Recruiting with Intermediaries." *Review of Economic Design*, 12(2): 119–127.
- Shimer, Robert, and Lones Smith.** 2000. "Assortative Matching and Search." *Econometrica*, 68: 343–369.