DEPARTMENT OF ECONOMICS AND FINANCE

Course in Economics and Business

Bachelor degree thesis in IO & Competition Theory

Thickness, congestion and signaling in the job market for new economic PhDs

Candidate: Sofia de Caprariis, 187541

Thesis supervisor: Prof. Gian Luigi Albano

Academic year: 2016/17

Index

Introduction

The aim of this dissertation is analyzing the functioning and efficacy of the signaling mechanism in ameliorating job market matching, both in term of quantity and "quality" of matches. Particular focus is given to the usage of the instrument in a particular job market, the job market for new PhDs in economics. The complex mechanism regulating this market, which matches junior economists to academic and institutional employers that have vacant positions, will be further explained and analyzed throughout the dissertation, it is useful to anticipate that the market has long been organized around the centralization actions of the American Economic Association (AEA).

The first centralization mechanism provided by the AEA is the Job Openings for Economists (JOE) that the association issues periodically and that allow job market candidates to learn who is hiring.

The second such action is the organization of the January ASSA meetings; each year, during the seminars organized for the AEA members, a job fair is organized that allows job market candidates to meet with recruiting committees of universities and other perspective employers in pre-scheduled interview sessions. The fair lasts three days, during which the bulk of the interviewing for the job market takes place; in almost the totality of cases, no offers are made or accepted prior to the meetings.

In 2006, the AEA established a standing Ad Hoc Committee tasked with ameliorating the functioning of the JOE and the other AEA job market services. Noble prize Al Roth presides on the Committee as chair, and several of the members are experts in games theory.

One of the actions of the committee was to introduce, in the 2007-8 job market, a signaling mechanism. It is precisely the signaling mechanism that makes the job market for new PhDs in economics so interesting from a market design point of view.

The mechanism consists of two signals of particular interest that candidates may send to perspective employers through the AEA. In this case, the AEA acts as a centralizing institution, which guarantees that each candidate does not send more than his allotted two signals. The signals are facultative and non binding (upon receiving an offer from an employer he signaled, the candidate does not have to accept it). The only thing each perspective employer can observe about signals are the signals it received.

The aim of the signaling mechanism is to alleviate the problems of congestion tied to thickness in the market; to work well, a market must be thick, that is, it must bring into contact a large number of buyers and sellers, but a thick market is at risk of becoming a congested one. If the congestion is severe enough, the market may not clear in the allotted time, or the matchings may be suboptimal.

The job market for new PhDs in economics is thick in that it brings together thousands of candidates and hundreds of employers, and congested because candidates apply widely, the cost per application being negligible, and recruitment committees cannot, over the three days of the interviewing fair, interview all those who sent them an application. The decision of which subset of applicants to interview is a fundamental and a strategic one, as perspective employers must take into account how likely a candidate would be to accept their job offer. It is for this reason that the signaling mechanism is crucial in dealing with issues of congestion in this market.

In this dissertation, the evolution of the signaling mechanism is reconstructed through a literature review, going from the introduction of the concept of signaling itself by Spence in 1973 to the design of the market for new PhDs in economics.

This evolution is analyzed and described by retracing the steps taken by the models that have characterized it and led to the construction of the instrument through an evolutionary process characterized by elements of critique and continuous refinement.

1. Market Design and Stable Matching

1.1 The characteristics of a well-functioning market

As Nobel Prize Alvin Roth systematized in his fundamental paper "What have we learned from market design?" [Roth; 2008] there exist a number of characteristic that a market must display in order to work well, the first such characteristic being that it must provide thickness; this means that they must attract a sufficient proportion of potential market participants to join the market, willing to transact with one another. As more and more participants (buyers and sellers) come together to the transact in the market, prices reflect true market levels better.

Thickness, however, leads to congestion, so the second characteristic of a well-functioning market is to provide enough time or to make transactions fast enough to enable market participants to consider a large enough number of possible different transactions, in order to settle on satisfactory ones. That is, a well-functioning market must be able to overcome congestion.

Congestion is especially a problem in markets in which transactions are heterogeneous and offers cannot be made to the whole market. If transactions take even a short time to complete but offers must be addressed to particular participants (as in offers of a job), then someone who makes an offer runs the risk that other opportunities may disappear while the offer is being considered. And even financial markets (in which offers can be addressed to the whole market) experience congestion on days with unusually heavy trading and large price movements, when prices may change significantly while an order is being processed, and some orders may not be able to be processed at all. When individual participants are faced with congestion they may react in ways that damage other properties of the market, for example, if they try to gain time by transacting before others.

The third characteristic of a well-functioning market is safety. It must be safe for buyers and sellers to participate in the market as simply as possible, so as to avoid incentives to transact outside of the market or engage in strategic behavior that causes overall welfare to decrease.

1.2 Unraveling issues and subsequent centralization

One function of a market is to establish a time at which large numbers of buyers and sellers can plan to make transactions [Roth and Xing; 1994]. However, establishing such a time can be difficult. In the past few decades unraveling of transactions times became a problem in several annual entry-level professional labor markets. Year after year, transactions in this market were made earlier and earlier, often in a way that made the market at any moment very thin. In some cases this resulted in such extremes new employees being hired up to 2 years before they would complete their education and start working.

Over time, many of the markets that had gone through this unraveling process, developed market institutions designed to alleviate the problem, for instance, by introducing centralized market-clearing institutions.

The first phase of the new markets is organized in the same way as the old markets: applicants and employers contact and interview one another in a decentralized fashion. The centralization comes into play in a second phase, when each applicant then submits to a centralized clearinghouse a rank ordering (in order of preference) of each employer with which he or she has interviewed and each employer submits a rank ordering of all the applicants they have interviewed. Leaving an applicant or employer off the preference list means that the worker or job is unacceptable, that, is the employer/applicant would prefer remaining unmatched at the end of the game to being matched with that worker/job. An algorithm is then utilized to obtain a matching of applicants to positions, through the preference lists.

The algorithms used in several effective centralized market-clearing mechanisms are approximately the same as the deferred acceptance algorithm first formally studied by Gale and Shapley in 1962 [Roth; 1984, 1991]. The deferred acceptance algorithm produces a stable matching of applicants to employers: there exists no worker and firm pair that is not matched to each other that would prefer being so to their current matches.

1.2.1 Deferred acceptance algorithm

The deferred acceptance algorithm may be firm-proposing or worker-proposing. In the firm-proposing deferred acceptance algorithm, each employer begins by offering each of its available positions to its most preferred candidates. This means, that if it has k identical positions, it offers them to the candidates in the top k spots of its preference list. Each candidate, after receiving the offers, rejects any unacceptable ones, and should he have received more than one offer, rejects them all but the most preferred one, the one from the firm ranked highest in his preference list, which is held without commitment. If any of their offers has been rejected, each firm offers the position to its next-highest-ranked candidate who has not yet rejected it. This continues as long as there acceptable candidates left on their lists. If a candidate gets new offers, he compares them with any offer he may be holding and, as he did before, rejects all but the most preferred. The procedure continues until there is no firm left which wishes to make any further offers. At this point, each candidate accepts the position (if any) he is currently holding, and he becomes matched to it.

This procedure results in stable matching. So long as all agents have strict preferences, it is firm-optimal among all the possible stable matchings. This means that there is no firm which prefers any of the other possible stable matching to this one.

The worker-proposing deferred acceptance algorithm works in much the same way, save that the roles are reversed: it is the candidates that make the offers and the employers which accept or reject them. The worker-proposing deferred acceptance algorithm, like the firm-proposing one, also results in stable matching. However, it is worker-optimal among all the possible stable matchings, rather than firm-optional. [Gale and Shapley; 1962. Roth and Xing;1997].

2. Introducing Signals

In most cases, the employer cannot be sure of the productive capabilities of a candidate during the hiring procedure, and, even after hiring, it will probably take some time for the employer to learn them. This implies that hiring an employee is an investment decision. Since the candidate's productive capabilities are not known beforehand, it is a decision taken under uncertainty. For this reason, hiring someone is in a way equivalent to the purchase of a 1ottery. Spence (1973) develops this analogy by assuming that the employer pays the certain monetary equivalent of the lottery to the worker as wage. If the employer is risk-neutral, the wage is taken to be the individual's marginal contribution to the firm.

The wages the employer is willing to pay are determined by the perceptions the employer has about the lottery. While the employer cannot directly observe the marginal product prior to hiring, he has available to him a vast number of data on the candidate, in the form of his personal characteristics and attributes. These observable data includes, amongst other things, education, previous work experience, ethnicity, gender, date of birth, criminal and service records.

Some of these characteristics are alterable, while others are not. For instance, an individual may invest in education, which is costly in terms of money and time. On the other hand, aspects such as date of birth are not alterable. Of the different observable attributes and characteristics, the unalterable one are referred to as indices, while those subject to manipulation by the individual are called "signals".

For each set of signals and indices, the employer has an expected marginal product for the individual displaying them. This expected marginal product is taken to be the wage offered to applicants with those characteristics.

As indices are unalterable, the applicant can do very little about them. Signals, on the other hand, being alterable, are susceptible to manipulation by the candidate. Of course, there may be costs tied to these manipulations, called "signaling costs". Education, for example, is costly.

The individual, when deciding to acquire an education, does not necessarily think of himself as signaling. His decision to invest in education is determined by whether or not there are sufficient returns to education. Individuals are, then, assumed to select signals for the purpose of maximizing the difference between the wages they are offered and signaling costs [Spence,M.; 1973]. Signaling costs play a key role in this type of signaling situation, for they play the function of the less direct costs and benefits that come with a reputation for signaling reliably, which can acquired by participants who are more prominent in their markets than candidates in the job market.

3. Cheap Talk and Matching Markets

The expression 'cheap talk', when used in the context of games of incomplete information, refers to communication between the players that is both direct and costless [Crawford and Sobel; 1982]. Cheap-talk models are markedly different from the more standard signaling models. In standard signaling models, the communication of private information by the informed agents takes place indirectly, through the choices the agents make, concerning, for instance, the level of education attained. As mentioned above, this choices cause the agent to bear a cost, the "signaling costs", and the fact that choices are differentially costly is exactly what makes signaling credible. E.g. the distinction between high-productivity and low-productivity workers may be made apparent by the decision of which level of education to attain, as some levels of education are too costly for low-productivity workers.

Cheap-talk models address one central question: if communication is costless and direct, can any information be credibly transmitted from the informed agent to the uninformed one? If so, how much?

There is always a 'babbling' equilibrium in cheap talk models. A babbling equilibrium is one in which the agents consider communication to transmit no credible information, so that there are no incentives for them to communicate anything meaningful [Krishna, V. and Morgan, J. 2004].

In the case of job markets, the communication that takes place between agents before the interviewing stage can be considered as a type of cheap talk.

In many instances, in job markets, candidates have incentive to apply for a several positions, since the cost of applying to a position is low and the value of being employed is high. For this reason, many employers find themselves having to evaluate hundreds of applications, which is an almost impossible task. Additionally, often there is a cost tied to the pursuing of candidates. For this reason, employers often are not only concerned with assessing the quality of the candidate, but also whether or not he is likely to accept a job offer if the employer makes one.

This implies that in many cases, the preferences of workers and firms concerning to whom to assign interview slots are only partially aligned.

Ceteris paribus, when interviews are costly, a firm will prefer to interview a worker who has a strong preference for the firm over one who does not, as, he would be likelier to accept a job offer should he receive one, while interviewing a worker who has no or little interest in the firm would be a bad investment, as he would probably decline a job offer.

And the more the interviewing costs, which can be significant, are borne by the firm, the more important it is to grant interview slots only to candidates who are likely to accept an offer.

However, when applicants don't have a mean to credibly signal their preferences, the have an incentive to misrepresent their preferences, exaggerating their compatibility with some firms. In fact, when the cost of remaining unmatched (i.e. the cost of unemployment) is sufficiently high, workers will attempt to maximize the number of interviews, so as to have increase their chances of finding a match.

The information workers communicate to firms concerns not only their preferences, but also their skills. So, in order to maximize the number of interviews he receives, a worker has an incentive to communicate to each firm that he is a good match, also where his abilities are concerned; e.g. a low quality firm has a preference for interviewing medium or low ability workers, as high quality candidates would be likely to decline their job offer

The inability of agents to exchange meaningful information can cause of issue of inefficiency and friction in the market.

In the case of direct, bilateral communication, meaningful communication is impossible, as job candidates will have an incentive not to be truthful in order to maximize the number of interviews. This changes if communication is public, or has some public component.

As shown in the classic model of cheap talk, information transmission is not necessarily perfect even in the case of public communication [Crawford and Sobel; 1982], and this may be the case in the present game as well, which, however presents a fundamental difference with respect to the classic model of cheap talk. In the classic model there is a single receiver of information, while in the present game there are multiple receivers, who have also have different preferences. For this reason, the possibility to transmit information publicly, either through a centralized authority or a central communication channel is fundamental to make meaningful communication possible [Lee and Schwarz; 2007].

3.1 Lee and Schwarz model: how to make information credible

Lee and Schwarz (2007) adopt a stylized model of communication, interviewing, and matching, and assume there are an equal number N of firms and workers (where the set of firms is given by F and the set of workers by W). Each firm hires at most one worker, and market participants will be matched via a firm-proposing deferred acceptance algorithm. Firms do not know ex-ante their preferences over workers, so, before the matching process, they must discover them through interviews. Firms simultaneously assign (and conduct) interviews to a subset of workers. Each firm is subject to a strictly positive cost per interview, c.

Worker's preferences, on the other hand, are uncorrelated and distributed uniformly over the firms and are known by the workers prior to the interviewing stage.

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

In some cases, however, firms knowing and taking into account worker preferences when assigning interviews may lead to Pareto improvements. For instance, there may be two workers who each prefer to interview with the other worker's firm and so, if possible, would switch interview spots. Allowing them to do so would not negatively affect the utility of any firm, and in fact, may even increase the firm's utility, as each firm weakly prefers interviewing a worker who has a strong preference for their firm to one who does not. As mentioned above, it would increase the probability of a job offer being accepted, so that interviewing him would increase the returns to the firm from interviewing a specific worker.

There are two alternative ways of allowing for a communication stage between workers and firms prior to the interview stage. Either private, bilateral communication between firm and worker must be allowed or a mechanism must be developed, allowing workers to "publicly" communicate their complete preferences to a intermediary, and the intermediary then gives recommendations to each firm regarding which subset of workers to interview.

Assumption 3.1. *Fix c such that there exists a symmetric pure strategy equilibrium interview assignment* $\hat{\eta}$ *in which each firm and each worker conducts x interviews.*

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

3.1.1 Private, bilateral communication

In the case of private, bilateral communication, before the firms allocate interview spots, each worker has the possibility to send a message (m_{wf}), simultaneously and privately, to each of the firms, indicating the worker's preferences where the firm is concerned. The message space (M) can be defined generally, but one natural candidate is $M \equiv \{1,...,N\}$, where m_{wf} can be understood to indicate how a given firm f is ranked on worker w's preference list.

If workers communicate privately with each firm, then the information shared is non-verifiable, and for this reason, the market outcome it can achieve for workers is not different from the one in which no communication between workers and firms is permitted.

Proposition 3.1. *If communication between workers and firms is private and non- verifiable, then communication is uninformative or ignored in equilibrium.*

If communication is private and bilateral, nothing is keeping workers from sending conflicting messages to each firm, and so each worker will send to each firm a message designed to maximize his chances of being scheduled for an interview.

Indeed, it has often been observed in practice that candidates on the job market will take the time to craft different cover letters to send to each firm, each stating that that particular firm is at the top of the candidate's preference lists. Over time, firms have learned to discount such cheap talk, and so these signals are ignored.

3.1.2 The case of public communication

One solution to the issues tied to private, bilateral communication problem would be to keep workers from directly communicating with the firms and sending multiple (and possibly conflicting) messages, allowing them instead to communicate only once, to a central intermediary. This would keep workers from simultaneously telling each firm that it is their top choice.

Assuming the following interview-assignment mechanism: each worker submits to a central authority a rank-ordered preference list over the firms ($\tilde{P}_w$), and subsequently, the mechanism privately communicates to each firm which subset of workers to interview $\eta(\tilde{P}_W)$, where $\eta(\cdot)$ is generated using the following algorithm, with $\tilde{P}_W \equiv \{\tilde{P}_w\}_{\forall w \in W}$ being the set of preferences for all workers:

- In the first round, each firm is assigned all workers who have ranked it as their top choice for an interview in their preference list $\tilde{P}_w$. If the workers that ranked a particular firm as the first in their list are more than x, x workers are selected at random among them.
- In general, in round t, any firm who has not yet allocated x interviews spots offers an interview to all workers who have not yet been offered x interviews and rank that firm as their *t-th* highest choice. If doing so would lead the firm to offer interviews to more than x candidates, the firm chooses randomly a number of workers such that it will interview, in total, exactly x workers.

This mechanism terminates with an assignment of interview slots that allocates exactly x interviews to each firm and each worker. This mechanism can be considered a worker-proposing deferred acceptance algorithm for interviews, where firms are indifferent among workers and for this reason will accept all interview offers up to quota, the quota being, for each worker and each firm, x.

Proposition 3.2. *There exists an equilibrium whereby workers report preferences truthfully* ($\tilde{P}_W = P_w \forall w$)*, and firms allocate interviews according to the assignment proposed by the intermediary,* $\eta(\tilde{P}_W)$.

The resultant allocation is also pairwise stable, as, in addition to the outcome being, by construction, Pareto-optimal for workers among all the possible outcomes assigning each worker and firm x interviews, no firm has an incentive to swap a worker w with a worker w′, who is interviewing with any other firm. This is due to the fact that the firm has no way of knowing how highly it is ranked in *w*'s or *w″*s preferences, nor does it know the identities of the other firms interviewing workers in the job market; it knows only that, in the equilibrium described above, w will receive exactly x interviews, and if the firm switched workers with another firm, it would interview a worker w′ who prefers that firm less than any of the current workers currently scheduled for an interview.

This mechanism is referred to as "public" communication because, when it is utilized, workers must commit to making a single announcement of their rank-ordered preference list and, thus, are kept from giving conflicting messages to different firms. This is even despite the fact that firms at no point are privy to the announcements, an assumption made only for technical reasons. In this setting, a centralized intermediary also takes on the function of a coordination device, as it enables firms to distribute interviews evenly, giving exactly x interviews to each worker.

However, implementing public communication or a centralized intermediary is infeasible in several cases, as, for instance, public communication may be impossible to observe or verify, and the costs associated with the implementation of a centralized intermediary may be too high. Additionally, there are issues associated with requiring workers to submit a complete rank-order preference list, as, when N is sufficiently large, it may be onerous for them to compute a ranking, which may also be too complex for them to communicate.

However, even in this setting, it is still possible to improve upon the no-communication outcome, as long as firms may monitor the number of interviews any worker obtains during the interview assignment stage, even if private, bilateral communication is allowed.

3.1.3 Monitoring the number of interviews

Assuming each worker w can, prior to the interview assignment stage, send a signal to any firm f; i.e., $m_{wf} \in \{0,1\}$, where $m_{wf} = 0$ indicates that w did not send f a signal and $m_{wf} = 1$ indicates that he did. During the interview assignment stage, nature selects a random ordering of firms; which then take turns selecting a worker to interview. Doing this, they are able to observe only the number of interviews the worker has received up to that point and not the identities of the firms who have offered them to him. It is assumed that the interviewing stage occurs after each firm has stopped assigning interviews.

Proposition 3.3. *For some N and* $(k_1,\ldots,k_N) \in \{1,\ldots,N\}^N$, *there exists an equilibrium in which each worker* w_i *sends a message to a firm f if and only if f is one of his* k_i*th highest ranked forms, and a firm assigns an interview to a worker only if he has signaled to it.*

This mechanism only requires that firms be able to observe the number of interviews a worker receives during the interview assignment stage; it places no restrictions on the number of messages a worker receives. The worker decides how to moderate the number of signals, as he is confronted with the following tradeoff from sending one additional signal: sending an extra signal both reduces the probability that the number of interviews the worker receives is less than x, but it also increases the probability of him being interviewed by a less desirable firm in place of a more desirable one.

The above proposition is an immediate logical consequence of noting that it holds for $k_i = N\ \forall i$ and that for any fixed k_i, if worker w_i signals to a firm f', he will signal to all firms $f \succ f'$, as, if he didn't, he could do better by signaling a more preferred firm instead of f'. In general, nothing specific can be said about the equilibrium number of signals sent, unless worker preferences over firms are translated into cardinal utilities. In terms of comparative statics, however, it is apparent that, if workers are almost indifferent between their top choice and the lowest-ranked firm in their preference list, then they will choose to send a higher number of signals than if they had a strong preference for their top firm with respect to the others.

Even if the size of the message space could increase without the outcome being substantially changed, for there to be any kind of truthful revelation, all informative messages (i.e., those received by a firm receives and used to influence its behavior) must be interpreted in the same manner, as, if firms differentiated in their behavior upon receiving two different messages, then workers would have incentive only to send the message that maximized its chances of obtaining an interview. Thus, the simple communication protocol used here not only is sufficient, but minimizes the complexity of communication and possible scope for strategic misrepresentation [Lee and Schwarz; 2007].

3.2 The value of a signaling mechanism: the Coles et al (2013) model.

A further step in the study of the role of signaling in this type of markets was taken by Coles et al. (2013). This paper studies a mechanism that aids employers in the evaluation process by allowing applicants to credibly signal information about their preferences for positions.

The structure of the model used in Coles et al (2013) presents also some key differences to the one used in Lee and Schwarz 2007;

First of all, the model of Coles et al (2013) focuses on the strategic question of offer choice and abstracts away the question of acquiring information that determines preferences. It is assumed that each agent knows her own preferences over agents on the other side of the market, but is uncertain of the preferences of other agents, Lee and Schwarz (2007), on the other hand, focused on the assignment of interviews in a job market in which the workers knew their preferences over the firms and the firms could find out their preferences about the workers only through interviews, but bore a strictly positive cost *c* per interview, so that assigning interview slots to workers was a strategic choice.

The model of Coles et al (2013) depicts the third and final part of the labor market described in Lee and Schwarz (2007): the formation of matches on the basis of available information. In the model, firms may make a limited number of offers to workers, so that firms must carefully select the workers to whom they make offers.

A class of block-correlated worker preferences is considered. Firms can be partitioned into blocks, so that all workers agree about ranking of firm blocks, but their preferences within a block of firms are idiosyncratic. The preferences of firms over workers are both idiosyncratic and uniformly distributed. Preferences were modeled in this way to allow for correlation of workers' preferences and to keep the model tractable.

The paper focuses on a simpler environment with a single block of firms where agents care about getting a match, but not the quality of the match, so as to make the comparison of performances across market settings as simple as possible, as it is necessary to understand when a signaling mechanism might be most helpful

Each worker is allowed to send one single signal, binary in nature, to one of the firms. The signal may not transmit any further information. Firms observe their signals, but cannot observe the signals received by the other firms. After this, each firm makes, simultaneously, exactly one offer to a worker. Finally, each worker chooses to accept one offer among the ones available to him. The paper focuses on equilibria in anonymous strategies, in order to eliminate any coordination devices beyond the signaling mechanism.

3.2.1 The basic intuition of the model: a simplified example

Consider a market with two firms $\{f_1, f_2\}$ and two workers $\{w_1, w_2\}$. For each agent, a match with one's most preferred partner from the other side of the market yields payoff 1, while a match with one's second choice partner yields $x \in (0, 1)$. Remaining unmatched yields payoff 0.

Agent preferences are random, uniform and independent. For each firm f, the probability that f prefers worker w_1 to worker w_2 is one half, as is the probability that f prefers w_2 to w_1. Worker preferences over firms are similarly symmetric. Agents know their own preferences, but not the preferences of other agents.

The game without signals

Once agent preferences are realized, each firm may make a single offer to a worker and workers then accept at most one of their available offers. Sequential equilibria will be examined, which guarantees that workers accept their best available offer.

In the unique equilibrium of this game where firm strategies do not depend on the name of the worker, each firm makes an offer to its most preferred worker, as firms do not know which worker is more likely to accept an offer. This market is congested and there is a fifty percent chance that both firms make an offer to the same worker, in which case there will only be one match; on average there are 1.5 matches, and the expected payoff for each firm is $\frac{3}{4}1 + \frac{1}{3}0 = 0.75$.

For workers, if they receive exactly one offer, that offer is equally likely to be from their first or second choice firm. Additionally, there is a fifty percent chance that one worker receives two offers, attaining a payoff of one while the other worker receives no offers and has a payoff of zero. The expected payoff for each worker is then (2 + x)/4.

The game with signals

A signaling mechanism is introduced: before the offering stage, each worker is allowed to send a signal, binary in nature and containing no additional information, to one of the firms. The focus is on non-babbling equilibria, where firms consider the signals to be informative, that is, a sign of being the most preferred firm of that worker, and workers send a signal to their most preferred firm.

Three different situations are possible:

- A firm may receive a signal from its most preferred worker. It this case it will make this worker an offer, since it will certainly be accepted.
- A firm may receive no signals. In this case, since each worker is equally likely to accept an offer from the firm, it optimally makes an offer to its most preferred worker.
- A firm may receive a signal only from its second ranked worker. In this case, the firm may "respond" to the signal, by making an offer to the signaling worker, or "ignore" the signal, by instead making the offer to its most preferred worker.

Suppose f_1 prefers w_1 to w_2 and only w_2 sent a signal to f_1, which implies w_1 sent a signal to f_2. The payoffs of f_1, conditional on receiving a signal from its second ranked worker, and the strategies of f_2, are summarized below, in Table 1.

Table 1: Firm f_1's payoffs conditional on receiving a signal from its second ranked worker and the strategies of f_2

f_1 / f_2	Respond	Ignore
Respond	x	x
Ignore	0	1/2

Table 1 shows that strategies of firms are strategic complements, that is, if a firm responds to signals, then the other firm is weakly better off from responding to signals as well. So, if f_2 switches from the action ignore (not making an offer to a second choice worker who has signaled) to the safe action of responding (making an offer to a second choice worker who has signaled), then f_1 optimally also takes the safe action of responding.

For what concerns equilibrium analysis, if $x > 0.5$ there is a unique equilibrium in which both firms respond to signals. When the value of the first choice worker is much higher than that of the second ranked worker $x < 0.5$, there exist two equilibria in pure strategies. In the first, both firms respond to signals (Respond-Respond) and in the second both firm ignore signals (Ignore-Ignore).

Table 2 summarizes welfare properties of these equilibria.

The expected firm and worker payoffs, as well as the expected number of matches, when signals are ignored are the same the ones in the game with no signaling, since agent actions are identical across these two situations. This makes it possible to study the effects of introducing a signaling mechanism, as outcomes in the offer game without signals are identical to those when both firms ignore signals. [1]

Table 2: Firm payoffs, worker payoffs and number of marches when both firms use the same strategy

	Firm payoffs	Worker payoffs	Number of matches
Respond-respond	(5+2x)/8	(6+x)/8	7/4
Ignore-ignore	3/4	(2+x)/4	3/2

3.2.2. The model: the offer game with and without signals

Let $\mathcal{F} = \{f_1, \ldots, f_F\}$ be the set of firms, and $\mathcal{W} = \{w_1, \ldots, w_W\}$ be the set of workers, with $|\mathcal{F}| = F$ and $|\mathcal{W}| = W$. In this model, only markets with at least two firms and two workers are considered. Firms and workers have preferences over each other.

[1] For details [Coles et al. 2013], pp. 105-107.

For each firm f, Θ_f is taken to be the set of all possible preference lists over workers, where $\theta_f \in \Theta_f$ is a vector of length W. The worker of rank one is the most preferred worker, while the worker of rank W is the least preferred worker.

The set of all firm preference profiles is $\Theta_F = (\Theta_f)^F$. The value that firm f with preference list θ_f assigns to a match with worker w is $u(\theta f, w)$, where $u(\theta f, \cdot)$ is a von-Neumann Morgenstern utility function.

Firms are symmetric in the sense that a firm's utility for a match depends only on a worker's rank; for any permutation ρ of worker indices, $u(\rho(\theta f\ ), \rho(w)) = u(\theta f\ , w)$. Additionally, all firms have the same utility function $u(\cdot,\cdot)$.

$\theta w, \Theta w$ and ΘW for workers are defined similarly. Worker w with preference list θw values a match with firm f as $v(\theta_w, f)$, where match utility again depends only on rank, and all workers share the same utility function.

It is assumed that workers and firms derive zero utility from being unmatched, and that any match is preferable to remaining unmatched for all participants. The set of all agent preference list profiles is denoted as $\Theta \equiv \Theta_F \times \Theta_W$ and let $t(\cdot)$ be the distribution over this set. A market is given by the 5-tuple $\langle \mathcal{F}, \mathcal{W}, t, u, v \rangle$.

In this model, block-correlated distributions of preferences are considered: firms can be partitioned in blocks, so that all workers agree which block contains the most desirable firms, which block the second most desirable set of firms and so on. However, within a block, workers have idiosyncratic preferences over firms. Each firm has preferences over the workers chosen uniformly, randomly and independently from the set of all strict preference orderings over all workers.

Definition 3.1. *The distribution of agent preferences $t(\cdot)$ is block-correlated if there exists a partition $\mathcal{F}_1, \ldots, \mathcal{F}_B$ of the firms into blocks are such that:*

1. *For any $b < b'$ where $b, b' \in \{1, \ldots, B\}$,each worker prefers every firm in block $\mathcal{F}_b$ to any firm in block $\mathcal{F}_b{}'$;*
2. *Each worker's preferences within each block $\mathcal{F}_b$ are uniform and independent; and*
3. *Each firm's preferences over workers are uniform and independent.*

3.2.2.1 The Offer Game with No Signals

In the absence of a signaling mechanism the game takes place as follows:

1. After preferences of firms and workers are realized, each firm simultaneously makes an offer to at most one worker.
2. Workers choose to accept at most one offer from those available to them. Sequential rationality ensures that workers will always select the best available offer. For this reason,

the behavior in the last stage is taken as given and, the focus is on the reduced game with only firms as strategic players.

Once its preference list θ_f (f 's type) is realized, firm f decides whether and to whom to make an offer. It may use a mixed strategy, σ_f, mapping the set of preferences lists to the set of distributions over the union of workers with the no-offer option, denoted by N; that is $\sigma_f : \Theta_f \rightarrow \Delta(W \cup N)$.

A profile of all firms' strategies is denoted as $\sigma_F = (\sigma_{f_1}, \ldots \sigma_{f_F})$, and the set of firm f's possible strategies is denoted as Σ_f.

The function $\pi_f : (\Sigma_f)^F \times \Theta \rightarrow \mathbb{R}$ denotes the payoff of firm f as a function of firm strategies and realized agent types. From this, follows the Bayesian Nash equilibrium of the offer game with no signals.

Definition 3.2. *Strategy profile $\hat{\sigma}_F$ is a Bayesian Nash equilibrium in the offer game with no signals if for all $f \in \mathcal{F}$ and $\bar{\theta} \in \Theta$, the strategy $\hat{\sigma}$ maximizes the profit of firm f of type $\bar{\theta}$. That is,*

$$\hat{\sigma}_F(\bar{\theta}_F) \in argmax_{\sigma_f \in \Sigma_f} \mathbb{E}_{\theta_{-f}} (\pi(\sigma_f, \hat{\sigma}_f, \theta)|\bar{\theta}).$$

The focus is on equilibria in which firm strategies are anonymous: they depend only on workers' ranks within a firm's preference list. This excludes any strategy relying on worker indices, thus eliminating any coordination linked to the identity of workers.

Definition 3.3. *Firm f' s strategy σf is anonymous if for any permutation ρ, and for any preference profile $\theta_f \in \Theta_f$, it follows that $\sigma_f(\rho(\theta_f)) = \rho(\sigma_f(\theta_f))$*

3.2.2.2. The Offer Game with Signals

The game is modified to allow each worker to send a "signal" to exactly one firm. The signal is binary in nature: the only decision of workers is whether and to whom to send a signal, and the signal may transmit no additional information.

The signal does not directly affect the utility a firm derives from a worker, as the firm's utility from hiring a worker is determined by how high the firm ranks that worker. What the signal of a worker may affect is a firm's beliefs over whether that worker is likely to accept an offer.

Since the market is congested and firms can only make one offer, these beliefs may affect the firm's decision of whom to make an offer.

The offer game with signals takes place as follows:

1. Agents' preferences are realized and each worker decides whether to send a signal, and to which firm. Signals are sent simultaneously, and are observed only by firms who have received them.
2. Each firm simultaneously makes an offer to at most one worker;
3. Each worker accepts at most one offer from the set of offers he receives.

Once again, since sequential rationality ensures that workers will always select the best available offer, this behavior for workers as given and the focus is on the reduced game consisting of the first two stages.

In the first stage, each worker sends a signal to a firm, or else chooses not to send a signal. A mixed strategy for worker *w* is a map from the set of all possible preference lists to the set of distributions over the union of firms and the no-signal option, denoted by N; that is, $\sigma_f : \Theta_f \rightarrow \Delta(W \cup N)$. In the second stage, each firm observes the set of workers that sent it a signal, $W^S \subset W \cup N$, and based on these signals forms beliefs $\mu f(\cdot \mid W^S)$ about the preferences of workers. Each firm, based on these beliefs as well as its preferences, decides whether and to whom to make an offer. A mixed strategy of firm *f* is a map from the set of all possible preference lists, Θ_f, and the set of all possible combinations of received signals, 2^W, which is the set of all subsets of workers, to the set of distributions over the union of workers and the no-offer option. That is, $\sigma_f : \Theta_f \times 2^W \rightarrow \Delta(W \cup N)$. A profile of all worker and firm strategies is denoted as $\sigma_W = (\sigma_{w_1}, \ldots \sigma_{w_W})$ and $\sigma_F = (\sigma_{f_1}, \ldots \sigma_{f_F})$ respectively.

The payoff to firm *f* is a function of firm and worker strategies and realized agent types, which is denoted as $\pi_f : (\Sigma_w)^W \times (\Sigma_f)^F \times \Theta \rightarrow \mathbb{R}$. Similarly, the payoff of workers is defined as $\pi_w : (\Sigma_w)^W \times (\Sigma_f)^F \times \Theta \rightarrow \mathbb{R}$. As the offer game with signals is a multi-stage game of incomplete information, sequential equilibrium is taken to be the solution concept.

Definition 4: *The strategy profile* $\hat{\sigma} = (\hat{\sigma}_W, \hat{\sigma}_F)$ *and posterior beliefs* $\hat{\mu}_f(\cdot \mid W^S)$ *for each firm f and each subset of workers* $W^S \subset W \cup N$ *are a sequential equilibrium if*

- *for any* $w \in W, \bar{\theta}_w \in \Theta_W : \hat{\sigma}_w(\bar{\theta}_w) \in argmax_{\sigma_w \in \Sigma_w} E_{\theta_{-w}}(\pi_w(\sigma_w, \hat{\sigma}_{-W}, \theta) \mid \bar{\theta}_w)$,
- *for any* $f \in \mathcal{F}, \bar{\theta}_f \in \Theta_f, W^S \subset W \cup N$:

  $\hat{\sigma}_f(\bar{\theta}_f, W^S) \in argmax_{\sigma_f \in \Sigma_f} E_{\theta_{-f}}(\pi_f(\sigma_f, \hat{\sigma}_{-f}, \theta) | \bar{\theta}_f, W^S, \hat{u}_f)$,

where $\hat{\sigma}_{-a}$ *denotes the strategies of all agents except* a, *for* $a = w, f$, *and beliefs are defined using Bayes' rule.*

As in the game without signaling, the focus is on equilibria where agents use anonymous strategies, thereby eliminating unrealistic sources of coordination.

Definition 5. *Firm f's strategy* σ_f *is anonymous if for any permutation* ρ, *preference profile* $\sigma_f \in \Theta_f$, *and subset of workers* $W^S \subset W \cup N$ *who send f a signal, it follows that* $\sigma_f(\rho(\theta_f), \rho(W^S)) = \rho(\sigma_f(\theta_f, W^S))$. *Worker w's strategy* σ_w *is anonymous if for any permutation*

ρ that permutes only firm orderings within blocks and any preference profile $\theta_w \in \Theta_w$, it follows that $\sigma_w(\rho(\theta_w)) = \rho(\sigma_w(\theta_w))$.[2]

3.2.3. Equilibrium Analysis

As will be shown below, there are marked differences between the equilibrium in the offer game with no signals and the offer game with signals: in the offer game with signals, the expected number of matches and the welfare of workers in any non-babbling equilibrium are greater than in the offer game with no signals. The welfare of firms changes ambiguously with the introduction of a signaling mechanism.

3.2.3.1 The Offer Game with No Signals

In the case of the offer game with no signals, when deciding whom to make an offer to, firms must consider both the utility from hiring a specific worker and the likelihood that this worker will accept their offer. Because preferences of firms are independently and uniformly chosen from all possible preference orderings, and since firms use anonymous strategies, an offer to any worker will be accepted with equal probability. Hence, each firm optimally makes an offer to the highest-ranked worker on its preference list. This is the unique equilibrium when firms use anonymous strategies.

Proposition 1. *The unique equilibrium of the offer game with no signals when firms use anonymous strategies and workers accept the best available offer is $\sigma_f(\theta_f) = \theta_f^1$ for all $f \in \mathcal{F}$ and $\theta_f \in \Theta_f$.*

For the above statement to hold, it is only necessary that firm strategies be anonymous in equilibrium, in all other cases, firm deviations that do not satisfy the anonymity assumption are still allowed.

As seen in the example in Section 3.2.1, in this equilibrium there might be considerable lack of coordination, leaving many firms and workers unmatched.

3.2.3.2. The Offer Game with Signals

In the case of the offer game with signals, the focus is on equilibria where firms within each block play symmetric, anonymous strategies.

That is, if firm f and firm f' belong to the same block $\mathcal{F}_b$, for some $b \in \{1, \dots, B\}$, they play the same anonymous strategies and have the same beliefs. Such firm strategies and firm beliefs are called block-symmetric.

Equilibria where firm strategies and firm beliefs are block-symmetric and worker strategies are anonymous and symmetric are denoted as block-symmetric equilibria.

---

[2] For details [Coles et al. 2013], pp. 107-111.

The model's first step in characterizing the set of block-symmetric equilibria is to pin down the strategies of workers, who must choose whether to send a signal, and if so, to which firm.

In block-symmetric equilibria, firms within each block use the same anonymous strategies. Hence, the ex-ante probability of a worker w receiving an offer from a firm in block $\mathcal{F}_b$, conditional on w sending and not sending a signal to it can be denoted as p_b^s and p_b^{ns} correspondingly. The equilibrium probability that a worker sends his signal to a firm in block $\mathcal{F}_b$ is denoted as α_b, where $\alpha_b \in [0,1]$ and $\sum_{b=1}^{B} \alpha_b \leq 1$.

The following proposition characterizes worker strategies in all block-symmetric sequential equilibria that satisfy a multiplayer analog of Criterion D1 of Cho and Kreps (1987), that is, If, under some ongoing equilibrium, a non-equilibrium signal is received which is equilibrium dominated for some types but not others, then beliefs cannot place positive probability weight on the former set of types.[3]

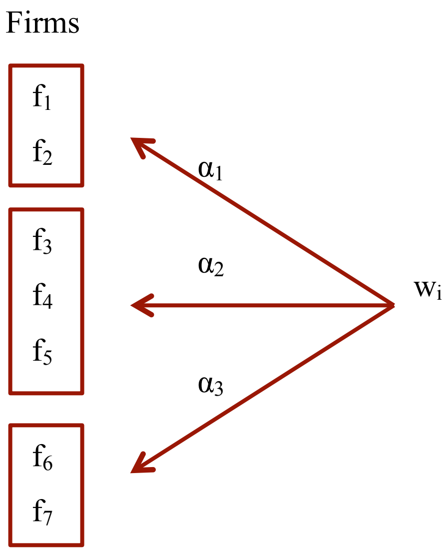

Figure 1: Signals in Block Symmetric equilibria.

Proposition 2 (Worker Strategies). *Consider a block-symmetric sequential equilibrium that satisfies Criterion D1. Then either*

1. *Signals do not influence offers: for every b ∈ {1, ..., B},* $p_b^s = p_b^{ns}$ *or*
2. *Signals sent in equilibrium increase the chances of receiving an offer: there exists* $b_0 \in \{1, \ldots, B\}$, *such that* $p_{b_0}^s > p_{b_0}^{ns}$ *and*

[3] See [Cho and Kreps; 1987] page 205 ff

(a) for any $b \in \{1, \dots, B\}$, such that $\alpha_b > 0$, it follows that $p_b^s > p_b^{ns}$, and if a worker sends her signal to block $\mathcal{F}_b$, she sends her signal to her most preferred firm within $\mathcal{F}_b$, and

(b) for any $b' \in \{1, \dots, B\}$, such that $\alpha_{b'} = 0$, workers' strategies are optimal for any off- equilibrium beliefs of firms from block $\mathcal{F}_{b'}$.

Proposition 2 states that there are two types of block-symmetric equilibria that satisfy Criterion D1. Equilibria of the first type are babbling, where firms ignore signals. The outcomes of these equilibria coincide with the outcome in the offer games with no signals. Consequently, the signaling mechanism adds no value in this case.

In equilibria of the second type, workers send signals only to their most preferred firm in each block, possibly mixing across these top firms (see Figure 1). It is quite natural to expect that in equilibrium, workers may signal to multiple blocks with positive probability. If all workers were signaling to the same block, the benefits to a single worker from signaling to a different block could be quite high.

In equilibrium workers only send signals to blocks in which firms respond to signals; that is, the chances of receiving an offer from the firm they signaled must be higher than if they had not sent that signal. Moreover, if in equilibrium worker w is not prescribed to signal to some block $\mathcal{F}_{b'}$, then w's choice of $\alpha_{b'} = 0$ is optimal for any beliefs of firms in block $\mathcal{F}_{b'}$. In particular, this strategy would be optimal even if firms in block $\mathcal{F}_{b'}$ interpreted unexpected signals in the most favorable way; i.e., upon receiving a signal from worker w, each firm f in $\mathcal{F}_{b'}$ believes that it is w's most preferred firm within block $\mathcal{F}_{b'}$.

Strategies where workers send signals only to their most preferred firm in any block (or mix over such firms) are denoted best-in-block strategies. Beliefs where a firm interprets a signal from a worker w as indicating it is the most preferred firm of w in that block are denoted as best-in-block beliefs.

In the model, it is assumed that workers use symmetric best-in-block strategies and that firms have best-in-block beliefs, and firm offers in the second stage of the game are examined.

f's most preferred worker is denoted as T_f (f's top-ranked worker). A firm f that has received signals from a subset of workers WS $\subset$ W $\cup$ N is considered. |WS| is denoted to be the number of received signals and it is assumed that |N| = 0. f's most preferred worker in this subset is denoted to be S_f (f's most preferred signaling worker). The expected payoff to f from making an offer to T_f or S_f (whichever yields greater payoff) is strictly greater than the payoff from making an offer to any other worker.

This results from symmetry of worker strategies and block-symmetry and anonymity of firm strategies: for any two workers who sent a signal, f's expectation that these workers will accept an offer is identical. Hence, if f makes an offer to a worker who sent a signal, it should make that offer to the worker it prefers the most among them. The same logic holds for any two workers who have not sent a signal.

This suggests a special kind of strategy for firms, called a cutoff strategy.

Definition 6 (Cutoff Strategies). *Strategy σ_f is a cutoff strategy for firm f if $\exists j_1, \ldots, j_W \in \{1, \ldots, W\}$, such that for any $\theta_f \in \Theta_f$ and any set WS of workers who sent a signal,*

$$\sigma_f(\theta_f, W^S) = \begin{cases} S_f \; if \; rank_{\theta_f}(S_f) \leq j_{|W^s|} \\ \quad T_f \; otherwise \end{cases}$$

($j_1, \ldots, j_W$) f's is called cutoff vector, which has as its components cutoffs for each positive number $|W^S|$ of received signals.

A firm f which employs a cutoff strategy need only look at the rank of the most preferred worker who sent it a signal, conditional on the number of signals f has received. If the rank of this worker is below a certain cutoff (lower ranks are better since one is the most preferred rank), then the firm makes an offer to this most preferred signaling worker S_f. Otherwise the firm makes an offer to its overall top ranked worker T_f. Cutoffs may in general depend on the number of signals the firm receives. This is because the number of signals received provides information about the signals the other firms received. This in turn affects the behavior of other firms and hence the optimal decision for firm *f*. Any cutoff strategy is, by definition, an anonymous strategy.

While, above, cutoffs were initially defined as integers, in this stage of the model. the definition is extended to include all real numbers in the range (1, W). This is done by letting a cutoff $j + \lambda$, where $\lambda \in (0, 1)$, correspond to mixing between cutoff j and cutoff $j + 1$ with probabilities $1 - \lambda$ and λ respectively.

Cutoff strategies are both intuitive and optimal strategies for firms. Whenever other firms use anonymous strategies and workers signal to their most preferred firms within blocks, for any strategy of firm *f* there exists a cutoff strategy that provides firm *f* with a weakly higher expected payoff.

In effect, since firm and worker strategies are anonymous, and the probability that firm *f*'s offer to T_f or S_f will be accepted depends only on the number of signals firm f receives, and not on the identity of the signaling workers, if *f* finds it optimal to make an offer to S_f, it will certainly make an offer to a more preferred S_f, provided the number of signals it receives is the same.

All the equilibrium results in this paper involve firms using cutoff strategies.

Since cutoff strategies can be represented by cutoff vectors, a natural partial order can be imposed on them: firm *f*'s cutoff strategy σ'_f is greater than cutoff strategy σ_f if all cutoffs of σ'_f are weakly greater than all cutoffs of σ_f and at least one of them is strictly greater: firm *f* responds more to signals than firm *f'* when σ_f is greater than σ'_f.

The next step of the model is to examine how a firm should adjust its behavior in response to changes in the behavior of opponents. The result is the same as the one pictured in the intuition in section 3.2.1: responding to signals is a case of strategic complements.

Proposition 3.4 (Strategic Complements). *Suppose workers play symmetric best-in-block strategies, all firms use cutoff strategies, and firm f uses a cutoff strategy that is a best response. If one of the other firms increases its cutoffs (responds more to signals), then the best response for firm f is also to increase its cutoffs.*

When other firms make offers to workers who have signaled to them, it is risky for firm *f* to make an offer to a worker who has not signaled to it. Such a worker has signaled to another firm, which is more inclined to make her an offer. The greater this inclination on the part of the firm's opponents, the riskier it is for firm *f* to make an offer to its most preferred overall worker T_f. Hence as a response, firm *f* is also more inclined to make an offer to its most preferred worker among those who sent a signal, namely S_f.

The next result establishes the existence of equilibria in block correlated settings in the offer game with signals. To prove the theorem, the authors first demonstrate equilibrium existence while requiring firms to use only cutoff strategies and then invoke the optimality of cutoffs result to show that this step is not restrictive.

Theorem 3.1 (Equilibrium Existence). *There exists a block-symmetric equilibrium where 1) workers play symmetric best-in-block strategies, and 2) firms play block-symmetric cutoff strategies.*

In the case in which there is only one firm block, the characterization of equilibria is even more distinct. When there is only one block of firms, an optimal strategy for each worker, for any anonymous firm cutoff strategies, is to send a signal to his most preferred firm.

Taking this behavior to be fixed, the strategic complements property stated in proposition 3.4 is used by the authors to apply Theorem 5 from Milgrom and Roberts (1990), that is, that all serially undominated strategies (and thus all rationalizable and equilibrium strategies) lie in an interval [$\bar{x}$, $\underline{x}$] whose maximum and minimum points are the largest and smallest Nash equilibria [P. Milgrom and J. Roberts; 1990].

From applying this concept, it follows that, in the case in which there is a single firm block, there exists a symmetric equilibrium in pure cutoff strategies where workers signal to their most preferred firms and accept their best available offer and firms use symmetric cutoff strategies and there also exist pure symmetric equilibria with smallest and largest cutoffs.[4]

3.2.3. Welfare Effects of Introducing a Signaling Mechanism

In order to analyze the effect of introducing a signaling mechanism on the market outcome, the authors take into consideration three outcome measures: the number of matches in the market, the welfare of firms, and the welfare of workers.

For agent welfare comparisons, the criterion utilized is Pareto ex-ante expected utility. The expected welfare for a firm f and a worker w are captured by π_f and π_w respectively, where $\pi_f, \pi_w\colon (\Sigma_w)^W \times (\Sigma_f)^F \times \Theta \to \mathbb{R}$.

The function $m\colon (\Sigma_w)^W \times (\Sigma_f)^F \times \Theta \to \mathbb{R}$ denotes the expected total number of matches in the market as a function of agent strategies and types.

So that it is possible for the comparisons between the outcomes of an offer game with and without signals to be strict, the authors take into consideration a block with at least two firms where in equilibrium workers send signals with positive probability to that block. As, otherwise, only weak comparisons may be made.

Theorem 3.2 (Welfare). *Consider any non-babbling block-symmetric equilibrium of the offer game with signals, in which there is a block $\mathcal{F}_b$ with at least two firms such that $\alpha_b > 0$. Then,*

- *i. The expected number of matches is strictly greater than in the unique equilibrium of the offer game with no signals.*
- *ii. The expected welfare of workers is strictly greater than in the unique equilibrium of the offer game with no signals.*
- *iii. The welfare of firms may be greater or smaller than in the unique equilibrium of the offer game with no signals.*

The introduction of a signaling mechanism increases the expected number of matches because firms, by responding to signals, make offers to workers who are more likely to accept them. Additionally, as firms responding to signals means that the probability of workers receiving an offer, and receiving one from their top-ranked firm, increases, the introduction of the signaling mechanism leads to an unambiguous increase in the expected welfare of workers.

However, if a firm responds to signals, this has a negative spillover on the welfare of other firms in the same block. This follows from the fact that when a firm decides to respond to a signal, by making an offer to the signaling worker, any other firm in the same block planning to make an offer

[4] "Preference signalling in matching markets", Coles et al. 2012, pages 12-16

to the same worker (despite not having received a signal) has no possibility of having its offer accepted.

While there is an "information effect", as firms derive positive value from using signals to better assess which workers will accept offers, there is also a potentially harmful "competition effect" where firm response to signals has a negative spillover on the welfare of other firms.

As either effect can dominate over the other, that the impact of introducing a signaling mechanism on firm welfare is ambiguous.

In the case of a model with a single block of firms, the authors make, again, sharper predictions. Changes in welfare are compared not only in the case of the game with and without signals, but also over signaling equilibria. Firms and workers have opposing preferences when there are multiple symmetric equilibria, as workers prefer the equilibrium with the greatest cutoffs, in which firms respond the most to signals and firms have a preference for the equilibrium with the smallest cutoffs, in which firms respond the least.

Theorem 3 (Welfare: Single Block of Firms). *Consider any two symmetric cutoff strategy equilibria in the offer game with signals with one block of firms where in one equilibrium firms have greater cutoffs (respond more to signals). Compared to the equilibrium with lower cutoffs, in the equilibrium with greater cutoffs: (i) the expected number of matches is weakly greater, (ii) workers have weakly higher expected payoffs, and (iii) firms have weakly lower expected payoffs.*

This result cannot be generalized to encompass an entire block-correlated market, as, in this case, there is no longer a purely negative spillover on other firms when a firm responds more to signaling. A low ranked firm making an offer to a worker who has signaled it will face less competition when a higher ranked firm focuses on a (different) worker who has signaled it, so firms in lower ranked blocks may draw benefit from a firm in a higher ranked block increasing its responses to signaling.

3.3 Kushnir (2013) and the internal critique of the Coles et al. model

One of the authors of Coles et al. (2013), A. Kushnir, showed in his paper "Harmful signaling in matching markets" (2013) that introducing the signaling mechanism impedes match formation in some markets, proving that the conclusions of the Coles et al 2013 paper are not universal, but rely on the assumption that preferences are block-correlated.

In order to do so, Kushnir considered a model similar to the one of Coles et al., but which displays one difference; in Kushnir's model, preferences are not block-correlated, but almost complete: each worker has either the same "typical" commonly known preferences with a probability close to one or "atypical" preferences taken from some distribution with the complementary probability close to zero. The preferences of workers are ex-ante independently distributed. Firms have some fixed and commonly known preferences over workers that need not to be the same.

Kushnir considers a decentralized matching game with three stages, as in Coles et al. (2013):

1. Agents' preferences are realized and each worker chooses a firm to which to sends one private costless signal, binary in nature and containing no further information. All workers send their allotted signal simultaneously.
2. Each firm decides to whom to extend a job offer for the single position they have available, and in doing so consider also the signals they received.
3. Each worker chooses, among all the offers they received, which one to accept and declines all other offers.

Kushnir shows that information asymmetry among firms is caused by the signaling mechanism, producing coordination failures during the second stage and, thus, causing the expected number of matches to decrease. A relevant result of the model is that, if at least three firms respond to some worker's signal, the introduction of signaling strictly decreases the expected number of matches.

The author finds the overall effect of the introduction of signaling on the firm and worker welfare to be ambiguous.

The signaling mechanism has a positive effect on both firm and worker welfare in that it makes it easier to obtain stable matchings between some workers and firms. However, the introduction of the signaling mechanism leads to some workers and firms remaining unmatched at the end of the game, negatively impacting both firm and worker welfare.

3.3.1 The Model

The model depicts a two-sided matching market with workers and firms. The set of workers and the set of firms are denoted as $\mathcal{W}$ and $\mathcal{F}$ respectively with $|\mathcal{W}| = W$ and $|\mathcal{F}| = F$. It is assumed that $W \geq F$.

Each firm has only one position available, and may offer it to at most one worker and each worker can accept at most one offer.

Worker w ranks firms according to some strict preference list θ_w. Θ_w denotes the set of possible worker's preference lists.

By convention, the firm of rank one is the most preferred firm, while the firm of rank F is the least preferred firm. The set of all workers' preference profiles is denoted as $\Theta_{\mathcal{W}} = (\Theta_w)^W$.

θ_f, Θ_f and $\Theta_{\mathcal{F}}$ for firms are defined similarly.

Each agent a has cardinal utility compatible with preference list θ_a. For simplicity, the author assumes that firms and workers have the same utility function. Additionally, utility depends only on the rank of the matching partner, meaning that the utility of an firm or worker from being matched with a worker or firm on the *kth* position in its/his preference list equals $u(k)$. An agent's cardinal utility from being unmatched is normalized to zero.

All agents strictly prefer being matched at the end of the game to remaining unmatched: for any k, $u(k) > 0$.

Each firm f has some fixed publicly known preference list θ_f. Preference lists are not necessarily the same across firms.

Each worker can be one of two types: "typical" or "atypical":

- All workers of typical type have the same commonly known preference list θ_0
- Atypical workers have preferences that are identically and independently distributed according to some distribution $A(\Theta_w)$ with positive probability of realization for each possible worker preference profile.

Ex-ante, each worker is typical with probability $1 - \varepsilon$ and atypical with complementary probability ε, for some $\varepsilon \in (0,1)$. The analysis focuses on the case when ε is small.

For convenience, firms are given their names according to the their rank in the preference list of workers of the typical type, $\theta_0 = (f_1, ..., f_F)$.

The notation for workers is organized so that, if all firms had the same preferences θ^*, workers would be named according to this preference list $\theta^* = \{w_1, ..., w_W\}$. Worker w_1 is the best worker among all workers $\mathcal{W}$ according to firm f_1's preferences, and for each $i=2,...,F$ worker w_i is the best worker among $\mathcal{W} \backslash \{w_1, \dots, w_{i-1}\}$ according to firm f_i's preferences. The rest of the workers, $\mathcal{W} \backslash \{w_1, \dots, w_F\}$, are named according to some prespecified order.

Kushnir analyzes two settings:

- The game without signals, which is made up of two stages.
  - o Agents' preferences are realized, and each firm makes at most one offer to some worker.
  - o Workers choose among the offers they received which one to accept and decline all the others.
- The game with signals, which consists of three stages.
  - o Agents' preferences are realized and each worker chooses a firm to which to sends one private costless signal, binary in nature and containing no further information. All workers send their allotted signal simultaneously.
  - o Each firm decides to whom to extend a job offer for the single position they have available, and in doing so consider also the signals they received.
  - o Each worker chooses, among all the offers they received, which one to accept and declines all other offers.

As seen in section 3.2, sequential rationality ensures that workers will always select the best available offer at the last stage of the offering game with and without signals, so that workers' behavior at the final stage of the game may be taken as given and the focus shifted on the reduced games with one and two stages respectively.

The notions of the game without signals can be adapted to describe the agents' strategies and the equilibrium concept for the game with signals.

In the game with signals:

- In the first stage, a strategy of worker *w* is a map from the set of all possible preference lists to the union of the set of firms and no-firm option, denoted by $\mathcal{N}$, $\sigma_w : \Theta_w \to \mathcal{F} \cup \mathcal{N}$.
- In the second stage, a strategy of firm *f* is a map from the set of all possible combinations of received signals, 2^W, which is the set of all subsets of workers, to the union of workers and the no-worker option, denoted by denoted $\mathcal{N}$, $\sigma_f : 2^W \to \mathcal{W} \cup \mathcal{N}$.

As it is assumed that each firm has some fixed, publicly known preferences, the dependence of firm strategy on preferences is omitted.

$\sigma_W = (\sigma_{w_1}, \dots \sigma_{w_W})$ is denoted to be a profile of all workers' strategies as, and the set of worker's strategies is denoted as Σ_w. σ_F and Σ_f are defined similarly as a profile of all firms' strategies and the set of firm's strategies respectively.

The final matching and agents' utilities for a given strategy profile of agents $\sigma = (\sigma_W \sigma_f)$ and realized agent types $\theta \in \Theta_{\mathcal{F}} \times \Theta_{\mathcal{W}}$ can be determined.

The utility of agent *a* given a strategy profile σ and a profile of types θ is denoted as $\pi_a(\sigma, \theta)$. The interim expected payoff of worker *w* with preferences θ_w from strategy σ_w when the other agents follow a strategy profile σ_{-w} equals $u_w(\sigma_w | \sigma_{-w}, \theta_w) = \sum_{\theta-w} t(\theta_{-w}) \pi_w((\sigma_w, \sigma_{-w}), (\theta_w, \theta_{-w}))$, where $t(\theta_{-w})$ denotes the joint distribution of all agents except worker w preferences.

Since signals are private and a firm does not observe a worker's action unless it receives a signal from the worker, the game with signals is the game of incomplete information and unobserved actions [Fudenberg and Tirole; 1991].

For each possible set of received signals, $\mathcal{W}_f^S$, firm *f* forms beliefs about the distribution of both workers' types and actions. Namely, $\mu_f(\theta, \mathcal{W}_{-f}^S \,| \mathcal{W}_f^S)$ specifies the probability firm *f* assigns to outcome $(\theta, \{\mathcal{W}_f^S\}_{f \in F})$ conditional on receiving signals from set $\mathcal{W}_f^S$ of workers.

The interim expected payoff of firm f given a subset of received signals $\mathcal{W}_f^S \subset \mathcal{W}$, beliefs $\mu_f(\cdot \,| \mathcal{W}_f^S)$, and other agents' strategy profile σ_{-f} is given by $\mu_f(\sigma_f \,| \sigma_{-f}, \mathcal{W}_f^S, \mu_f) =$ $\sum_\theta \sum_{\mathcal{W}_{-f}^S \in (2^W)^{F-1}} \mu_f(\theta, \mathcal{W}_{-f}^S \,| \mathcal{W}_f^S) \, \pi_f f(\sigma_f, \sigma_{-f}, \theta)$. To conduct theformal analysis the notion of sequential equilibrium is used.

Definition 1 *A strategy profile* $(\hat{\sigma}_w, \hat{\sigma}_f)$ *and firm beliefs* $\{\hat{\mu}_f\}_{f\in F}$ *form a sequential equilibrium if*

- *for any* $w \in \mathcal{W}, \theta_w \in \Theta_{\mathcal{W}}: \hat{\sigma}_w(\theta_w) \in arg\, max_{\sigma_w \in \Sigma_w} u_w(\sigma_w | \hat{\sigma}_{-w}, \theta_w)$ *and*
- *for any* $f \in \mathcal{F}, \mathcal{W}_f^S \subset \mathcal{W}: \hat{\sigma}_f(\mathcal{W}_f^S) \in arg\, max_{\sigma_f \in f} u_f(\sigma_f | \hat{\sigma}_{-f}, \mathcal{W}_f^S, \hat{\mu}_f)$,

where beliefs are defined using Bayes' rule.[5]

3.3.2 Equilibrium analysis

The game without signaling serves the purpose of a benchmark, so its equilibrium is described first. There exists a unique equilibrium in the game without signals, so long as ε *is* sufficiently small. The top firm, firm f_1, makes an offer to its most preferred worker, i.e. worker w_1.

The second top firm, firm f_2, anticipates that worker w_1 is likely to accept firm f_1's offer, so that its optimal strategy is to offer a position to worker w_2, its favorite worker among $\mathcal{W}\backslash\{w_1\}$.

The third top firm, firm f_3, anticipates that workers w_1 and w_2 are likely to accept firm f_1's and f_2's offers, respectively, so that its optimal strategy is to make an offer to worker w_3, its most preferred worker among $\mathcal{W}\backslash\{w_1, w_2\}$. This process continues like this until all the available positions have been filled.

Theorem 1 *For any* $\varepsilon \in (0, \bar{\varepsilon})$,[6] *there is the unique equilibrium in the game without signals. Firm* f_j, $j = 1,\ldots,F$, *makes an offer to worker* w_j.

The number of matches in the unique equilibrium in the game without signals is maximal and equals to F (since $W \geq F$). This match is defined as the "no signaling" one.

In the case of the game with signals, a subset of workers $\mathcal{W}_f^S \subset \mathcal{W}$ is said to be reached for firm f when workers follow strategy profile σ_W if firm f receives signals from exactly set $\mathcal{W}_f^S$ with positive probability. When workers follow strategy profile σ_W, firm f responds to worker w's signal, if there is a subset of workers $\mathcal{W}_f^S \subset \mathcal{W}, w \notin \mathcal{W}_f^S$, such that both $\mathcal{W}_f^S$ and $\mathcal{W}_f^S \cup w$ are reached for firm f when workers follow strategy profile σ_W, and $\sigma_f(\mathcal{W}_f^S) \neq \sigma_f(\mathcal{W}_f^S \cup w)$.

Firm f responding to worker w's signal means that worker w's signal changes the firm's strategy with positive probability. This happens only if the signal transmits useful information about worker w's preferences, and firm f is ready to take this information into account and act upon it.

The phrase "firm responds to a worker's signal in equilibrium" indicates that the firm responds to the worker's signal when the worker follows the equilibrium strategies.

[5] For details, see [Kushnir A., 2013], pp. 210-212.

[6] $\bar{\varepsilon} = \min\left(\min_j \left(\frac{u(j)-u(j+1)}{u(j)}, \frac{u(F)}{u(F)+u(1)}\right)\right.$

There always exists a babbling equilibrium in which firms do not respond to signals, in this case, the only possible outcome is the no signaling match, since signals are not taken into considered by the firms.

On the other hand, if firms respond to signals in equilibrium they make their offers based on the set of signals they receive and this might change the matching outcome. The following theorem establishes the existence of such an equilibrium. Moreover, it provides a condition when there exists an equilibrium when at least three firms responds to some worker signals.

Theorem 2 *Fix any $\varepsilon \in (0, \bar{\varepsilon})$. If there are at least three firms that weakly prefer some worker wi to their no signaling matches, i.e. $|(f_j \in \mathcal{F} : w_i \succcurlyeq_{f_j} w_j)| \geq 3$, then there exists an equilibrium of the game with signals in which at least three firms respond to worker w_i signals.*

The above result may be illustrated by way of an example with three firms and three workers. For simplicity, it is assumed that all firms rank the workers in the same way (w_1, w_2, w_3), i.e. each firm strictly prefers worker w_1 to worker w_2 to worker w_3. Hence, all three firms weakly prefer worker w_1 to their no signaling matches.

The following set of strategies constitutes an equilibrium of the game with signals with the property that all three firms respond to worker w_1's signal. Worker w_1 sends her signal to the best firm according to her realized preference list. Worker w_2 (w_3) always sends her signal to her no signaling match, i.e. firm f_2 (f_3). Each firm makes its offer to worker w_1 only if it receives a signal from her. If firm f_1 does not receive a signal from worker w_1 it makes its offer to worker w_2. If firm f_2 (f_3) does not receive a signal from worker w_1 it makes an offer to worker w_2 (w_3).

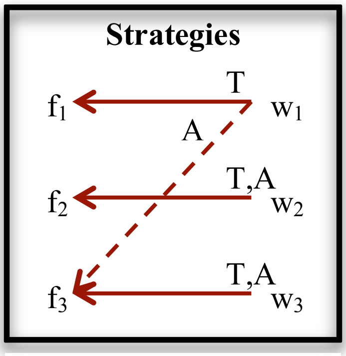

Figure 2: Equilibrium Behavior

Each firm believes that it is worker w_1 most preferred firm only if it receives a signal from him, so that each firm optimally makes an offer to worker w_1 only when it receives his signal.

If firm f_1 does not receive worker w_1's signal, it believes that worker w_1 has sent a signal to his top preferred firm, and that that firm that will make an offer to w_1 and be accepted.

In this case, firm f_1 can secure a match with worker w_2 with probability at least equal $1 - \varepsilon$. Likewise, if firm f_2 (f_3) does not receive a signal from worker w_1 it can secure at best worker w_2 (w_3) with probability at least equal $1 - \varepsilon$.

All firms ignore worker w_2 and worker w_3 signals, which, additionally, do not provide any information about the worker realized preferences.[7]

---

[7] For details, see [Kushnir A., 2013], pp. 212-213.

3.3.3 Number of matches

The expected number of matches is different in the case of the offer game with signals and the offer game without signals.

As shown in section 3.3.2, as long as the probability of atypical preferences in sufficiently small, the number of matches in the unique equilibrium of the game without signals is maximal.

This implies that introducing a signaling mechanism cannot increase the expected number of matches.

Theorem 3 *If at least three firms respond to some worker signals in an equilibrium of the game with signals the expected number of matches in this equilibrium is strictly smaller than in the unique equilibrium of the game without signals.*

To illustrate the theorem, the same example employed in section 3.3.2 is used.

Since $A(\Theta_w)$ has the full support the preference profile in which firm f_3 is the most favorite firm of worker w_1 realizes with positive probability. In this cae, w_1 sends a signal to firm f_3, and workers w_2 and w_3 send their signals to firm f_2 and f_3 correspondingly. Firm f_3 makes an offer to worker w_1, and firm f_1 anticipates that worker w_1 is atypical and makes an offer to worker w_2. Firm f_2 also makes its offer to worker w_2, leading, most likely, to firm f_2 remaining unmatched, as "typical" type worker w_2 would prefer firm f_1 to firm f_2.

This is illustrated in figure 4.

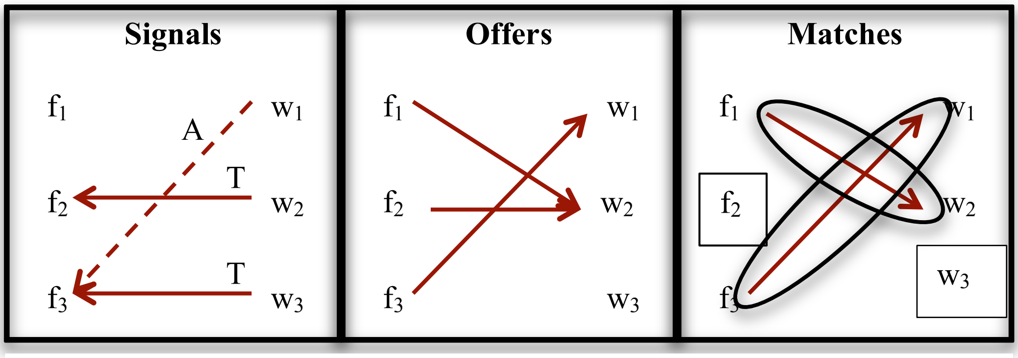

Figure 3: Mismatches

The coordination failure arises because, firm f_2 has no way of anticipating firm f_1's behavior, as it has no information about worker w_1's type. In this case, the number of matches for some realization of preferences is strictly smaller than the maximum match; the expected number of matches in this equilibrium is strictly smaller than the maximum match.

This holds only if at least three firms respond to signals in equilbrium. In fact, if only two firms respond to some worker signals in an equilibrium of the game with signals, the expected number of matches in this equilibrium might be equal to maximum match. In such an equilibrium if some firm

f_j secures a better match with some atypical worker w_i, $i < j$, firm f_i always makes its offer to worker w_j in an equilibrium. This does not decrease the number of matches, since firms exchange their partners.[8]

3.2.3 Welfare

The effect from the introduction of signaling on welfare depends on the relative magnitudes of firms' and workers' cardinal utilities.

Signals play two roles in equilibria when firms respond to signals.

- Signals help to secure "better" matches between some atypical workers and firms, which positively affects the welfare of agents on both sides of the market.
- The introduction of signals also leaves some workers and firms unmatched, negatively affecting the welfare of agents on both sides of the market.

Example 1 illustrates that the introduction of signals is beneficial for a matching market according to egalitarian welfare criterion if and only if the decrease in the number of matches is offset by better matches of atypical workers. A similar calculation shows that the total welfare of firms changes ambiguously.

Example 1 *The example considered is the one discussed in section 3.3.2. Additionally, it is assumed that preferences of atypical workers are independently uniformly distributed among all possible preference order lists. Workers' cardinal utilities from being matched to first, second, and third choice are* $\delta + \lambda, \delta$, *and* $\delta - \lambda$ $(\delta > \lambda)$ *respectively. The expected total welfare of workers in no signaling match equals*

$E[W_{worker}^{nosignals}] = \sum_{i=1}^{3}\left[(1-\varepsilon)u(i) + \varepsilon\frac{1}{3}\sum_{l=1}^{3}u(l)\right] = 3\delta.$

One may check that the expected total welfare of workers in the equilibrium discussed after Theorem 2 equals (terms of the order of ε^2 *and* ε^3 *are omitted)*

$E\left[W_{worker}^{nosignals}\right] = 3\delta + \left(-\frac{1}{3}\delta + 2\lambda\right)\varepsilon$

Hence, the expected total welfare of workers increases only if the difference in utilities between adjacent firms is large enough, $\lambda > \frac{1}{6}\delta$.[9]

[8] For details, see [Kushnir A., 2013], pp. 213-214.
[9] For details, see [Kushnir A., 2013], pp. 214-215.

3.4 Role of signals in matching markets: a comparison between the two models, Coles et al. (2013) and Kushnir (2013)

Coles et al. (2013) show that the introduction of signals increases the expected number of matches and the welfare of workers in a model similar to the one discussed in Kushnir (2013), save for the fact that in it, preferences of workers over firms are block-correlated.

Preferences of workers over firms in Kushnir (2013), on the other hand, are almost complete, and the introduction to the signaling mechanism decreases the number of matches with respect to the no signaling equilibrium and ambiguously affects welfare.

The introduction of signaling, then, has markedly differing effects depending on the environment in which it is introduced. These effects for block-correlated and almost-complete preferences are summarized in table 3.

Table 3: Almost complete and block-correlated distribution preferences

Preferences	No signals	Matches	$E[W_{firm}]$	$E[W_{worker}]$
Block-correlated	0	+	±	+
Almost complete	0	−	±	±

Why does the introduction of a signaling mechanism affect matching markets in differing ways? Kushnir (2013) argues that this is related to the two different roles signals play in matching markets, as they simultaneously transmit information and cause information asymmetry, leading to coordination failures. Because of this, the signaling tool might be less powerful than it was previously believed to be.

When agents' preferences are block-correlated as in Coles et al. (2013), the amount of information regarding agents' preferences available before the game is small, so the aspect of information transmission is more important in match formation.

When, on the other hand, the information about agents' preferences is almost complete, as in Kushnir (2013), the information asymmetry introduced by usage of the signaling tool dominates over its role as a transmitter of information, so that introducing it leads to coordination failures.

This comparison is summarized in Table 4. [10]

Table 4: The roles of signals

Preferences	Transmit Information	Introduce information asymmetry

[10] For details, see [Kushnir A., 2013], pp. 215-216.

Almost complete	Small	Large
Block-correlated	Large	Small

4. AEA Job Market

4.1 How the Job Market for new PhDs in economics works

The North American market for new PhDs in economics was, until the mid-seventies, a fairly decentralized, thin market. In fact, several economics departments filled assistant professors positions through words of mouth and letters of inquiry instead of advertising for them. This changed in 1974, when the AEA (American Economics Association) started to create more thickness in the market through the publication of JOE (Job Openings for Economists).

Since then, other centralized marketplace institutions have been established in this market, most of them by the AEA. Over time, the annual ASSA meetings, organized AEA, have become a fundamental centralization mechanism for the conduction of job interviews, leading to a thick market.

In 2005, the AEA established a permanent Ad Hoc Committee on the Job Market, composed by Al Roth (chair), John Cawley, Peter Coles, Phillip Levine, Muriel Niederle, and John Siegfried. The Committee is tasked with overseeing policy on issues concerning the JOE and the AEA job market services. It has introduced two new mechanisms in the market: market signaling, which takes place each year in early December, and the scramble, which takes place in March, during the final part of the market.

In the early fall, economics departments start advertising for positions. They do so through various routes, but almost all departments with positions to fill advertise (also) on JOE. This makes JOE a very important centralization tool, used by departments to advertise openings and for candidates to see which positions are available.

For this reason, JOE presents a fairly complete and accurate depiction of the academic part of the job market for new PhDs in economics.

PhDs candidates who are set to complete their graduate studies within the academic year then answer the ads by sending applications to the departments they are interested in. The cost the candidates incur by applying to an additional department is negligible (in fact, since the introduction of market clearinghouses allowing candidates to send all their applications simultaneously by uploading them once on the platform, it is equal to zero), so each candidate applies to several departments.

Because of this departments most often receive hundreds of different applications and the recruiting committee does not have enough time to meet with them all. In fact, out of the hundreds of candidates, only a handful will be offered an interview at the AEA meetings in January. Thus, as a consequence of the actions taken to ensure thickness, this market suffers heavily from problems of congestion.

Through the AEA January meetings, in a short time frame, perspective employer can interview several applicants and, at the same time, applicants can meet with several perspective employers. Only a few interviews take place before the meetings, and, even in those cases, it's almost unheard of for candidates and departments to be matched before January, when they have the chance to conduct other interviews and meet with other perspective employers, respectively.

So, the AEA January meeting provides thickness for the market, by providing it with a centralized location and a common starting time for all agents, but the thickness provided by the meetings is associated with congestion. As mentioned above, it is impossible for the recruiting committee of a department to interview, in the span of a few days, all the candidates that sent them an application, as departments often receive hundreds of applications.

It is apparent that the choice of which applicants to interview is a crucial one, as universities rarely, if ever, extend job offers to candidates they have not interviewed at the AEA meetings. Additionally, the decisions has some fundamental strategic aspects to it, as well: departments need not only be concerned with how much they like a candidate when it comes to offering him an interview, but also with how likely that candidate is to accept a job offer were he to receive one. Only elite schools may be reasonably confident that any job offer they make will be accepted, so for all other departments, it may be worthwhile to offer interviews to candidates they prefer less, but who are more likely to accept a position.

Since the cost of applying is negligible, candidates apply for several positions in order to minimize risk of remaining unmatched at the end of the market, so that the fact that a student applied to a department is not necessarily a signal of intense interest.

This issue is even more pressing in the case of liberal arts colleges, non-academic employers and schools outside of North America. Since these types of employers present marked differences with respect to North American research universities, they have need of a way to distinguish between candidates who are genuinely interested in them and those who simply applied everywhere.

After the January meetings come to a close, the departments start inviting candidates they are interested in for so-called "fly-outs", or "campus visits", in which the candidate visits the campus, meets the faculty and gives a presentation.

In this phase, the market is much less organized: different departments schedule fly-outs at different times, make offers at different times and demand answers for those offers at different times, so that candidates may find themselves having to accept or refuse an offer without knowing if they will receive other offers from competing departments later on.

Market congestion is an issue in this case, as well, as departments can invite only a small percentage of the applicants they met with at the January meetings to fly-outs, as campus visits are costly to organize.

Regardless of the differing schedules, the fly-outs and offering stage that make up the latter part of the primary market are concluded by the end of March, at the latest. At this point, the market is no longer congested, but it is also no longer thick. Candidates and departments remaining unmatched at this stage find it difficult to get information about who, on the other side of the market, also remained unmatched. This leads to difficulties in coordinating a match.

To alleviate this issue, the Ad Hoc Committee on the Job Market devised a "scramble", which is a market mechanisms designed to make the secondary market thicker.

The scramble was first introduced in the 2005-6 Job Market and it is organized through the JOE Network. The precise dates of the scramble vary from one year to the next, but, to make the description as clear as possible, the dates for the 2017 scramble will be used in the following paragraphs.

- On March 15th, registrations to the scramble open. Only unmatched candidates (departments with positions left unfilled and candidates who have yet not accepted an offer) are eligible for the scramble.
- On March 22, registrations for the scramble close.
- On March 24th, the scramble webpage goes online for viewing. Only those who registered are given access to the page, and may only see which agents are left unmatched on the other side of the market.
- Scramble viewing remains open until April 10th, but no changes are made to the page from March 24th to April 10th. The page does not include messaging or matching services. It is only meant to provide agents a "snapshot" of the secondary market, which is why it is kept open for such a brief period of time.
- 615 candidates and 101 jobs, listed by 89 different employers participated in the 2016 scramble.

4.2 The signaling mechanism in the job market for new economics PHDs: why and how

It is apparent from the description above that a well-functioning first stage it is crucial for optimal market outcome, and thus it is fundamental to alleviate the congestion issues affecting this phase.

In the 2006-7 job market, the Ad Hoc Committee introduced a new mechanism to facilitate match formation in the primary market: signaling.

The aim of the signaling mechanism is to deal with issues of congestion during the decision of which subset of applicants to interview at the January meetings. It allows candidates, in early

December, to send up to two signals of interest to perspective employers. The AEA acts as a centralization institution, as it transmits the signals from the applicants to the departments.
While sending a signal to one perspective employer is costless in terms of money, it does bear an opportunity cost in the sense that it keeps the candidate from sending the signal to another department [Chakrabortya and Harbaugh, 2007]. So, by limiting the number of available signals to two, each signal transmits meaningful information, making the signaling mechanism helpful in limiting coordination failures.
In particular, signaling can help alleviate the problem of candidates "falling through the cracks", that is, of an employer believing that the candidate was too desirable to accept its job offer, despite this not being the case. Without the aid of coordination mechanisms, this happens with relative frequency in this market and leads to unstable matchings, as both the candidate and the employer in question would prefer to be matched to each other rather than to the agent they are currently matched with.
Without the signaling mechanism, it is difficult for job market candidates to credibly signal their interest in a particular department: any attempt to contact the perspective employer directly would probably be considered non-credible cheap talk, as there is nothing stopping the applicant from telling the exact same thing to all other perspective employers, so as to maximize his chances of obtaining an interview.
Listed in Table 5 is data on the usage of signaling over the years. Precisely, it lists the number of candidates signaling, the number of signals, number of JOE listings overall and the number of JOE listings receiving signals, as well as the percentage of JOE listings signaled from 2009 to 2016.

Table 5: Usage of signaling over the years, 2009-2015

Year	*Nr of candidates signaling*	*Nr of signals*	*Nr of JOE listings*	*Nr of JOE listings receiving signals*	*Percentage of JOE listings signaled*
2009	979	1924	2285	667	29,2
2010	1095	2145	2842	601	21,1
2011	1284	2520	2836	673	23,7
2012	1285	2513	2915	697	23,9
2013	1260	2474	2790	715	25,6
2014	1475	2891	3051	737	24,2
2015	1770	3425	3304	707	21,4

Source: AEA, Report of the Director, several years

The number of candidates using signaling as almost doubled in time period considered, but the same cannot be said for the number of JOE listings receiving signals, which have increased less than proportionally than the number of JOE listings overall. So, it's relevant to find out which employers receive the most signals and what their unifying characteristics are.

4.2. Usage of Signaling in the Job Market, evidence from 2006-2009

The 2010 paper "The Job Market for New Economists: A Market Design Perspective" by Coles et al. analyzes the usage of signals by job market candidates, using statistics from the AEA archives on job markets in the years from 2006 through 2009 and data from surveys conducted by the Committee for the same period.

Statistics on how signals were used in the 2006-2009 job market are displayed in table 6.

Table 6: Number of signal senders, signals sent, and signals recipients, 2006-2009

Year	Number of signalers	Number of signals	Number of employers signaled	Number of Joe listings signaled	Number of Joe listings (entire year)	Fraction of Joe listings signaled
2006	971	1890	519	674	2634	25.5%
2007	1022	2010	489	672	2914	23.1%
2008	979	1926	461	687	2881	23,8%
2009	978	1922	449	666	2285	29,1%

Source: Coles et al. 2010

In the years considered by the paper, unlike what happened later on, the number of job market candidates who made use of the signaling mechanism remained roughly constant over the years, ranging from a low point of 971 in 2006 (the introduction year) to a peak of 1022 in 2007.

Data from the survey conducted in December 2008 shows that one third of participants doesn't use the signaling mechanism and, less then half of them does so out of strategic concerns (41% believed signaling wouldn't be helpful in securing interviews and 5% believed it would outright hurt their chances). The rest of them did not signal out of misinformation (26% missed the deadline and 21% weren't aware of the existence of the signaling mechanism). Unfortunately, similar data are not available for more recent years, so there is no certain explanation for the subsequent rise in the number of signalers, which started after 2010 (see table 5). It may be due to the component linked to the spread of information (a reduction in the number of misinformed candidates) or to the component linked to strategic choices (a decrease in the number of agents skeptical of the usefulness of the tool).

Concerning the composition of the group of agents using the signaling mechanism, the results of surveys conducted by the Committee in years 2006–2009 indicate that, among those sending signals, the biggest homogenous group is represented by PhD candidates (roughly 67% of signalers) - roughly seven times larger than the second largest homogenous group, made up by assistant professors (9% of signalers).

Despite the fact that anyone is allowed access to the AEA job market, and thus also the signaling mechanism may be used by anyone, almost all (99%) of the signalers either had a PhD or had nearly completed the process of earning one. Of those, the overwhelming majority had economics as their primary field (92%) and received (or about to receive) their PhD from a US institution.

Usage of signaling by the candidates was roughly the same regardless of the rank of their home institution, so that there does not seem to be any correlation between sending signals and home institution rank. This result was obtained by analyzing the participation to the signaling mechanism of job market candidates listed on web pages of the 100 highest-ranked economics departments, (ranking data was obtained from econphd.net).

However, if candidates signal regardless of the rank of their home institution, there appears to be a certain degree of uniformity where the recipients of their signals are concerned. In fact, number of signals sent each year total roughly three times the number of JOE listings receiving signals, so that, each year, only roughly between 23% and 29% of JOE listings and only roughly between 450 and 500 employers received at least one signal. At this point, a relevant question is if there some characteristics common to the most frequently signaled employers, and if so, what these characteristics are.

Figure 4 Signals Received by Employer Rank, 2006-2009

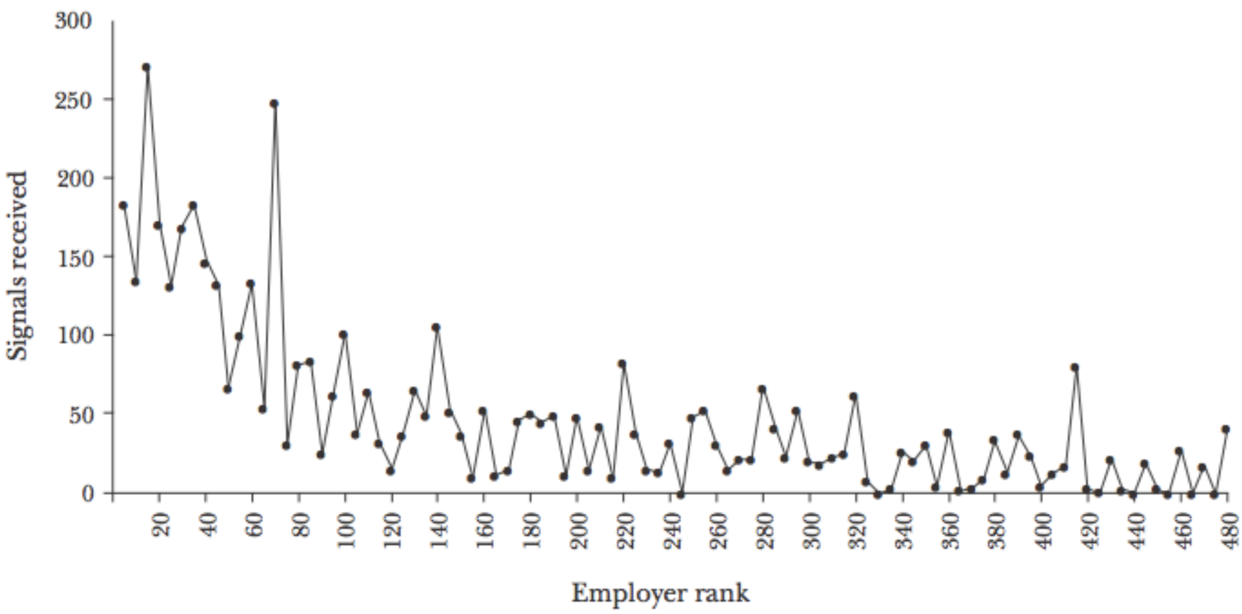

Source: Coles et al. 2010.

As seen in section 3.4, block-correlated worker's preferences are fundamental for the effective functioning of the signaling mechanism. This is a fair assumption to make in the job market for new economic PhDs, in that candidates generally agree about the partitioning of firms in different tiers, based on institution rank, while maintaining idiosyncratic preferences within each tier.

As such, it is relevant to first investigate the correlation between number of signals received and employers' rank. This correlation is depicted in figure 5.

In order to account for imprecisions in the ranking of departments and to protect confidentiality, employers were grouped into blocks of five employers, so that the first point on Figure 5 shows the total number of signals received by the top five employers in rank over years 2006–2009. Over that period of time, the top five ranked employers together received roughly 180 signals, averaging at roughly 9 signals per year for each employer.

In general, better-ranked employers received more signals for the top 100 ranked departments, at which point signals received by rank are fairly evenly distributed. The spikes in the graph in figure 5 can be explained by the fact that they are, in most cases, related to departments situated in popular geographic areas, such as large cities.

Despite the fact that the number of signals received appears to be correlated to employer's rank, institutions of all ranks, even those ranked very low, receive signals. This is a sign that signals may be fulfilling one of their roles, that of facilitating coordination, as low-ranked departments have an especially hard time distinguishing which applicants are truly interested in them and which one have applied only out of a desire of maximizing the number of interview offers they receive.

Table 7: Signals inflow, 2006-2009

	Signals to…					
Signals from…	Tier 1	Tier 2	Tier 3	Tier 4	Unranked	Total
Tier 1	101	368	274	105	159	1007
Tier 2	111	478	521	385	646	2141
Tier 3	40	165	339	385	706	1635
Tier 4	15	84	116	179	603	997
Unranked	13	51	54	62	232	412
Total	280	1146	1304	1116	2346	6192

Source: Coles et al. 2010.

Table 7 display statistics about signal inflow. Signal inflow is how signals from candidates with a given home department rank are allocated across departments of different rank. This is relevant as an attempt to reconstruct how candidates preferences are distributed over institutions.

For the purpose of this analysis, only signals sent to economics departments were considered and the departments were divided in 5 tiers. The tiers all have different sizes, and departments are divided in the following way:

- Tier 1: departments ranked 1-10
- Tier 2: departments ranked 11-50
- Tier 3: departments ranked 51-150
- Tier 4: departments ranked 151-480
- Tier 5: unranked departments

It is apparent from the table that candidates who come from institutions in the top tiers tend to send signals to departments in lower tiers. For instance, in the years considered, only 1426 signals were sent to the departments in the first two tiers, but candidates studying in those departments sent 3148.

Those which received the highest number of signals in the years considered (2346 signals) are departments in tier 5, i.e. unranked departments. Thus, unranked departments appear to be a fairly common destination for signals.

Taking into consideration the top 21 most signaled employers in the years 2006-2009 (both academic and nonacademic employers were taken into account when determining their identity) is useful in determining which characteristics are most taken into account by candidates when using the signaling mechanism.

First of all, while, as established above, there is a correlation between rank and number of signals received, only 1/3 of these employers are economics department with rank equal or higher than 21.

A much more important unifying feature than rank, in this group, is geography: roughly a half of these employers are situated in metropolitan area of either Boston, New York, or Washington, D.C. and of the eleven remaining employers, three are situated in California.

Additionally, within the geographic areas mentioned above, signals tend not to be concentrated in the area's most highly ranked departments.

For instance, despite Boston being both a very popular geographic area for signaling (four of the top 21 signaled departments are situated in Boston) and the place in which two of the top five economic departments are situated, only one of those two departments is part of top 21 most signaled employers, and, on average, the remaining three employers situated in Boston that are part of the top 21 most signaled employers were ranked slightly above 60.

The reason for this distribution of signals may be related to candidates having strong preference for specific locations: big cities that facilitate dual job searches.

An explanation for the fact that candidates have a tendency to send their signals to lower ranked institutions may be found in the fact that departments at the top of the pecking order do not need to worry about discerning which of the candidates that sent them applications are truly interested in them when offering interview, as they can be reasonably sure that all of them are. Thus, top ranked departments are likely not to consider signals when deciding who to interview.

4.2 The effectiveness of signaling in the AEA job market

In Coles et al. 2010, the authors focused their analysis on how effective signals actually are as a means of obtaining interviews with perspective employers.

Ideally, in order to measure this, one would need to know the probability of being offered an interview with a given employer both in the case in which a signal was sent to him and in the case it wasn't sent to him.

The best way to obtain this knowledge is, arguably, conducting a randomized experiment, for instance by randomly deleting one of the signals each candidate submits and calculating interview rates both for sent and unsent signals. This was impossible for two reasons:

- Such an action, in this market, would be frowned upon in that it may negatively affect the chances of some candidates, and was thus deemed to be inappropriate
- It may also negatively affect the credibility of signals, as, an employer, if a candidate claims that he sent them a signal, has no way of knowing (if it did not receive a signal from that candidate) if he is lying or if his signal was randomly deleted.

An alternative to randomized experiments is surveys, so the authors conducted the following surveys on candidates in the 2007-8 and 2008-9 job markets:

- In December, as the candidates sent their two allotted signals they were asked where they would have sent a third signal, had they had one. In the 2008-9 market, also candidates who had not made use of the signaling mechanism where interviewed, with the following question: “If you had used the signaling mechanism, where would you have sent your signals?”
- After the end of the market, respondents received follow-up questions about the outcome of their applications (interviews, fly-outs, job offers…) concerning both the employers they actually signaled and those to which they would have sent their hypothetical signals. Questions regarding the outcome of their signals were posed also to candidates who had not submitted a third hypothetical signal.

These data were used to estimate a random-effects, ordinary least squares, linear probability model in which the unit of observation is a pairing of a candidate and an employer.

The sample includes only employers that were signaled or would have been signaled by the candidate.

The dependent variable is an indicator variable for whether the candidate received an interview from the employer.

The regressor of interest is an indicator variable that equals one if the candidate signaled that employer, and zero if the candidate would have signaled the employer, either if the candidate had had a third signal (if they sent two) or the candidate had sent any signals (if they sent none).

In all cases it can be assumed that the candidate had a genuine interest for the employer, believing it to be a good match.

For candidates who sent two signals, one can think of this model as exploiting something akin to a regression discontinuity: it can be assumed that the third-choice employer was roughly an equally good match for the candidate as the first- and second-choice employers, but only the top two employers were sent signals.

For candidates who did not send any signals, one can think of this model as utilizing a situation similar to a natural experiment, in which the treatment group sent their two signals, and the control group did not. This is a meaningful comparison on that, while all candidates considered the employers to be good matched, signals were sent only in a subset of cases.

Table 8: The association of sending a signal with receiving an interview

	(A)	(B)	(C)	(D)	(E)
Sample	Economics departments	Economics departments	Economics departments	Economics departments	All employers
Panel data grouping	Candidate	n/a	Candidate	Candidate	Candidate
Signal	0.0680**	0.158***	0.0321	0.0580*	0.0634**
	(0.0287)	(0.0443)	(0.0328)	(0.0317)	(0.0271)
Is graduate student		0.176***			
		(0.0481)			
(Is graduate student) *Signal		-0.101*			
		(0.0556)			
Applied to liberal arts school			-0.0482		
			(0.0587)		
(Applied to liberal arts school) *Signal			0.177***		
			(0.0675)		
Applied to international school				-0.0295	
				(0.0654)	
(Applied to international school) *Signal				0.0583	
				(0.0739)	
Applied to nonacademic employer					0.132*
					(0.0765)
(Applied to nonacademic employer) *Signal					0.130
					(0.0827)

Constant	0.247***	0.114***	0.259***	0.252***	0.254***
	(0.0256)	(0.0376)	(0.0296)	(0.0283)	(0.0242)
R^2	0.003	0.012	0.015	0.004	0.003
Observations	1529	1520	1529	1529	1868
Number of candidate groups	728	n/a	728	728	785

Source: Coles et al. 2010

Table 8 displays the main results of the regression[11].

Column A:

- The intercept shows that applications to economics departments that were unaccompanied by signals resulted in interviews with a 24.7 percent success rate.
- The coefficient on the indicator variable for having sent a signal indicates that sending a signal is associated with a 6.8 percentage point higher probability of receiving an interview, meaning that applications accompanied by signals resulted in interviews with a 31.5 percent success rate. This difference is statistically significant at the 5 percent level, suggesting that signals are helpful in securing interviews.

Having established this, it is relevant to understand whether the correlation between sending a signal with being offered an interview is stronger for some groups than for others.

In column B of Table 8, signaling outcomes of current graduate students to those for non–graduate students are compared. In order to do so, two indicator variables are added to the regression: one for whether the applicant is a graduate student, and the other an interaction term for being a graduate student who signaled.

Column B:

- For current graduate students, signals are associated with a 5.7 percentage point increase in the probability of receiving an interview (adding the coefficient of .158 on sending a signal to the coefficient of –.101 on the interaction of sending a signal and graduate student status).
- Absent a signal, the probability of an interview for a graduate student is 29.0 percent (adding .114 and .176); with a signal, the probability is 34.7 percent.
- For non–graduate students, signals are associated with a 15.8 percentage point increase in the probability of receiving an interview.

This difference may be explained by the fact that non–graduate students might be more easily be over-looked than graduate students, as, for instance, they are less likely to appear on department web-pages of job candidates. Signals are, then, for them, a very useful mean to attract the attention of perspective employers.

[11] For details, see [Coles et al., 2010] pp 195-196.

In column C of Table 8, signaling outcomes of signals sent to liberal arts colleges are compared to those for non–liberal arts colleges.

- A signal to an economics department in a liberal arts college was associated with a 20.9 percentage point higher probability of receiving an interview
- A signal sent to an economics department not in liberal arts colleges was associated with a 3.2 percentage point higher probability of an interview, which is not statistically significant.

The reason why liberal arts colleges may take signals into consideration more than other perspective employers do may be that, when it comes to the role played by faculty members, they have different preferences than research universities. Without signals, liberal arts colleges have a hard time discerning which candidates are genuinely interested in a position and which ones have applied to them simply to minimize their chances of remaining unemployed but have a strong preferences for working in a research university.

Another group of employers that might find signals particularly informative is economics departments outside the United States, as they face a problem similar to the one faced by liberal arts colleges: discerning applicants who are genuinely interested in them to those who have a strong preference for a job in a North American institution. This hypothesis is tested in column D of Table 8. It is not possible to reject the null hypothesis that the association of signaling with probability of interview is the same for international and for North American employers.

Column E of Table 8 compares outcomes for signals sent to academic and to nonacademic positions. It is not possible to reject the null hypothesis that the association of signaling with probability of interview is the same for academic and nonacademic employers, though applications to nonacademic employers are associated with higher interview rates regardless of signals.

Table 8, continued

	(F)	(G)	(H)	(I)	(J)
Sample	U.S. economics departments	U.S. economics departments	Economics departments	Economics departments	Economics departments
Panel data grouping	Candidate	Candidate	Candidate	Candidate	Candidate
Signal	0.0191	-0.00325	0.00712	0.00122	0.104***
	(0.0378)	(0.0395)	(0.0370)	(0.0373)	(0.0345)
Town (pop<50000)	-0.0665	-0.0419			
	(0.0593)	(0.0656)			
Town *Signal	0.133**	0.0601			
	(0.0661)	(0.0735)			
Applied to liberal arts school		-0.0433		-0.0314	
		(0.0676)		(0.0669)	
(Applied to liberal arts school) *		0.178**		0.155	
		(0.0778)		(0.0767)	

Signal					
Applied to unranked school			-0.0370	-0.0221	
			(0.0514)	(0.0588)	
(Applied to unranked school) *Signal			0.152**	0.109	
			(0.0591)	(0.0674)	
Received many signals (≥ than 7)					0.00227
					0.0307
(Received many signals) *Signal					-0.0648*
					(0.0347)
Constant	0.271***	0.278***	0.265***	0.266***	0.246***
	(0.0344)	(0.0363)	(0.0340)	(0.0344)	(0.0274)
R^2	0.015	0.015	0.012	0.015	0.006
Observations	1266	1266	1529	1529	1529
Number of candidate groups	638	638	728	728	728

Source: Coles et al. 2010

Schools in small towns might face a problem similar to the one faced by liberal arts colleges and international institutions: that is, distinguishing genuinely interest candidates from those who have applied everywhere out of risk-aversion.

Column F of Table 8 displays outcomes for economics departments in towns with population under 50,000.

- Signals to these departments were associated with a 15.2 percentage point higher probability of an interview
- Signals sent to department in larger cities resulted in 1.9 percentage point higher probability of an interview. This difference is significant at the 1 percent level.

When controlling for whether the school is a liberal arts college, as shown in column G of Table 8, there is still a 6.0 percentage point difference based on city size, but the result is no longer statistically significant. (Signaling a liberal arts college that is not in a small town is associated with a 17.8 percentage point increase in the probability of interview, which is statistically significant.)

The final employers taken into account for the purpose of the regression are departments that do not appear in econphd.net's rankings of economics departments. These 361 schools received 1,188 signals over the two-year period 2007–8, 37% of all signals to economics departments during that time. Column H in Table 8 shows that signals to unranked departments are associated with a 15 percentage point higher probability of interview relative to signals sent to ranked schools. When controlling for whether the institution is a liberal arts college (Column I of Table 8), the difference drops to 10.9 percentage points, which is no longer statistically significant.

Column J of Table 8 examines the role of the total number of signals that an employer receives. Employers are divided into two groups: those that received seven or more signals and those that received six or fewer. A signal to an employer who received six or fewer signals was associated with a 10.4 percentage point higher probability of an interview. In contrast, sending a signal to an employer who received seven or more signals was associated with just a 3.9 percentage point increase in the probability of interview.

These results are only an indicate of correlation, and do not indicate causality, but it is possible that signals become less effective when sent to schools that receive many of them.

In general, there may be issues with the interpretation of these results; in particular, the assumption that candidates may be equally good matches for institutions they actually signaled and those they only indicated for a hypothetical signal may not be a true one.

If candidates are better fits for the institutions they actually signaled, it is impossible to tell whether the perspective employer may have been able to pick up on that even without receiving a signal, so that it is difficult to estimate the true effect of sending a signal on the probability of receiving an interview.

Other issues are that, if an applicant sends a signal, it is possible that he will not have to make any attempts to communicate his genuine interest in that perspective employer along other avenues, something he would have done had he not sent the signal.

Additionally, while the authors had data on sent signal, they had no way of knowing how many of signals sent were actually received, in the sense that the communications of the signals sent by the AEA to the employers may have been, in part, overlooked. If this was the case, the results of this model, are indicative on the effect of sending a signal, but understate the effect of receiving a signal.[12]

[12] For details, see [Coles et al., 2010] pp 197-200.

Conclusions

A well-functioning market must be thick, must be able to overcome the congestion caused by thickness and must be safe. Congestion, if severe enough, may lead to the market not being cleared within the allotted trading period. In matching markets, cheap talk makes it impossible to solve the problem of congestion through informal, decentralized means. Thus, the signaling mechanism, in the context of matching markets is a tool constructed to overcome cheap talk and congestion issues.

Nevertheless, as seen in sections 3.3 and 3.4, the effectiveness of the signaling mechanism cannot be ascertained for the generality of cases. This is due to the fact that it produces asymmetric information that, depending on the structure of agents' preferences, may actually lead to a worse market outcome than the no-signaling case, as it reduces the number of matches in equilibrium.

This stems from the fact that a firm may only observe the signals it received, and has no information about the signals received by the other firms, and, if workers' preferences are almost complete, the equilibrium outcome contains mismatches in which agents that would prefer being matched to each other than being unmatched at the end of the game remain unmatched.

On the other hand, the signaling mechanism works effectively if workers' preferences are block-correlated.

Additionally, in terms of welfare, in the game with signals, in equilibrium, almost complete preferences are associated with ambiguous effects on welfare of both firms and workers, whereas, block-correlated preferences are associated with ambiguous effects on welfare of firms but surely positive effect on welfare of workers.

In the AEA job market, preferences can be fairly assumed to be block-correlated, which idiosyncratic preferences existing within each tier. Since 2007, when it was first introduced, usage of the signaling mechanism in this market has increased over time and there is empirical evidence pointing to a correlation between sending a signal to an employer and higher probability of obtaining an interview with it, implying that employers take signals into account when allotting interview slots.

However, in this field, data are scant and there is no specific empirical evidence pointing to the introduction of the signaling mechanism in the AEA job market having increased the number of matches.

So even if there is reason to suppose that the introduction of signals has handled the problem of cheap talk and alleviated the issue of congestion, there is no indisputable evidence that signals have lead to better market outcomes in terms of number of matches.

This is tied to the fact that, while the assumption that preferences in this market are block-correlated is a fair one, it remains an assumption and not a proved fact.

As proving it is very difficult, if not impossible, a relevant further step of the research would be to seek a viable route towards relaxing the assumption, by analyzing the hypothetical consequences of altering the signaling mechanism by making the content of the signals publicly available, so as to remove the problem of asymmetric information.

Biblography

American Economic Association. Several years. Report of the director: Job Opening for Economists. American Economic Review, also posted on the AEA website https://www.aeaweb.org/articles?id=10.1257/aer.106.5.740

American Economic Association. 2011. Report of Ad Hoc Committee on the Job Market. American Economic Review: Papers and Proceedings. Vol. 101 Nr. 3, pp. 744–746.

Cawley, J. Several years. A guide and advice for economists in the US junior academic job market.. Posted on the American Economic Association (AEA) Job Openings for Economists (JOE) webpage, http://www.aeaweb.org/joe

Chakraborty,A, Harbaugh, R. Persuasion by Cheap Talk. 2010. American Economic Review, Vol. 100, 5, pp. 2361-82.

Cho I. and Kreps D., 1987, "Signaling Games and Stable Equilibria". In The Quarterly Journal of Economics, Vol. 102, No. 2, (May, 1987), pp. 179-221

Coles, **P., Cawley,J , Levine, P.B., Niederle, M, Roth, A.E., Siegfried, J.J.**. 2010. The Job Market for New Economists: A Market Design Perspective: Journal of Economic Perspectives—Volume 24, Number 4, Pages 187–206

Coles, **P., Kushnir, A., Niederle, M.** 2013. Preference Signaling in Matching Markets. American Economic Journal: Microeconomics 5:2, 99-134

Crawford, V.P., Sobel, J.. 1982 Strategic Information Transmission. Econometrica, Vol. 50, 6, pag. 1431 – 1451.

Fudenberg, D., Tirole, J.1991. Game Theory. MIT Press.

Gale, D. and Shapley, L. 1962 College Admissions and the Stability of Marriage. American Mathematical Monthly, January, 59, pp. 9-15.

Krishna, V. and Morgan, J. 2004. The art of conversation: eliciting information from experts through multi-stage communication. Journal of Economic Theory 117,147–79

Kushnir, A. 2013. Harmful signaling in matching markets. Games and Economic Behavior 80, 209-218.

Lee, R., Schwarz, M. 2007. Signaling Preferences in Interviewing Markets. In Computational Social Systems and the Internet. , ed. P. Cramton, P. Muller, E.Tardos and M.T. Bisin.

Milgrom,P.,Roberts, J. 1990. Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. Econometrica, Vol. 58, No., pp. 1255-1277

Spence, M. 1973. Job Market Signaling.The Quarterly Journal of Economics, Vol. 87, No. 3., pp. 355-374.

Bandyopadhyay, S. , Iskhakov, F., Johnson, T., Lee, S., McArthur, D., Rust, J., Watson, J. 2013. Can the Job Market for Economists be Improved?. pp. 189-221 in Handbook of Market Design, Neeman, Z., Niederle,M., Roth, A.E., and Vulkan, N.(eds.), Oxford: Oxford University Press.

Roth, A.E. 2008. What Have We Learned from Market Design? The Economic Journal, Volume 118, Issue 527, Pages 285–310

Roth, A.E. The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory.1984. Journal of Political Economy, December, 92, pp. 991-1016.

Roth, A.E. A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom.1991. American Economic Review, June, 810, pag. 415-40

Roth, A.E., Xiaoling Xing. 1994. Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions. American Economic Review, vol. 84 n. 4, pag. 992-1043

Roth, A.E., Xiaoling Xing. 1997. Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists. Journal of Political Economy, Vol. 105, pag. 284-329.

University of Arkansas, Sam M. Walton College of Business. Several Years .Survey of the labor market for new p.h.d. hires in economics. Posted on http://cber.uark.edu/files/2015-16_New_PhD_Labor_Market_Survey_Report.pdf