I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17052

Sorting through Cheap Talk: Theory and Evidence from a Labor Market

John J. Horton
Ramesh Johari
Philipp Kircher

JUNE 2024

IZA Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17052

Sorting through Cheap Talk: Theory and Evidence from a Labor Market

John J. Horton
MIT and NBER

Ramesh Johari
Stanford University

Philipp Kircher
Cornell University, Universite catholique de Louvain and IZA

JUNE 2024

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

IZA DP No. 17052 JUNE 2024

ABSTRACT

Sorting through Cheap Talk: Theory and Evidence from a Labor Market*

In a labor market model with cheap talk, employers can send messages about their willingness to pay for higher-ability workers, which job-seekers can use to direct their search and tailor their wage bid. Introducing such messages leads – under certain conditions – to an informative separating equilibrium that affects the number of applications, types of applications, and wage bids across rms. This model is used to interpret an experiment conducted in a large online labor market: employers were given the opportunity to state their relative willingness to pay for more experienced workers, and workers can easily condition their search on this information. Preferences were collected for all employers but only treated employers had their signal revealed to job-seekers. In response to revelation of the cheap talk signal, job-seekers targeted their applications to employers of the right "type," and they tailored their wage bids, affecting who was matched to whom and at what wage. The treatment increased measures of match quality through better sorting, illustrating the power of cheap talk for talent matching.

JEL Classification: J64, D83, C87

Keywords: sorting, cheap-talk, gig-economy, freelancer, field-experiment, online job search platform

Corresponding author:
Philipp Kircher
Dept of Econ
ILR
Cornell University
Ives Hall, Ithaca
NY 14853
USA

E-mail: pk532@cornell.edu

* We thank Andrey Fradkin, Dave Holtz, Ed Lazear, Paul Oyer, Isaac Sorkin, Liran Einav, Mike Luca, Meng Liu, Steve Tadelis, Steven Davis, Dean Eckles, and Richard Zeckhauser for helpful comments and suggestions. We also thank participants at a number of seminar presentations and at the Marketplace Innovation Workshop at Stanford, the ACM-EC Conference, and the NBER Summer Institute for IT & Digitization. Thanks go to Ada Yerkes Horton for help in preparing the manuscript, and to Lucia Casal for valuable research assistance. Kircher gratefully acknowledges support through European Research Council Grant ERC-818859. This paper supersedes earlier working papers under the titles "Cheap Talk Messages for Market Design: Theory and Evidence from a Labor Market with Directed Search" and "Buyer Signalling Improves Matching: Evidence from a Field Experiment". Code (eventually): http://john-joseph-horton.com/papers/hot towel.code.html.

1 Introduction

Today's workforce repeatedly needs to look for jobs in a changing marketplace. This is most evident for freelancers, who make up nearly 40% of the US workforce in 2022 and who frequently change whom they are working for (Lowe (2022)). But even workers in traditional employment often need to look for new employment at sub-annual frequency.[1] Finding the right job commensurate to the skills of each worker can become a challenge.

This paper investigates whether cheap-talk in the sense of Crawford and Sobel (1982) is a conducive tool to help navigate this challenge and improve labor market matching. Recent theory papers have pointed out that labor markets could become much more efficient when firms can use cheap talk to signal how important it is for them to hire (Menzio, 2007; Kim and Kircher, 2015). There has been no attempt to investigate the usefulness of such cheap talk empirically. We extend the theory to allow for the empirically relevant setting where firms are not only interested in hiring *per se*, but also differentially care about worker ability. Then we investigate the implications in a field experiment where firms are provided with additional cheap-talk opportunities to express these preferences. We find that cheap talk is informative: job applications and hires exhibit more sorting, and wage demands and resulting wage payments adjusted to the posted messages. These findings are broadly in line with predictions from theory.

Consider an online matching platform that brings together workers with short-term jobs. The market designer understands that employers prefer workers with more experience (Pallais, 2014; Stanton and Thomas, 2015) and are often willing to accept higher wage bids over lower wage bids (Horton and Vasserman, 2020). But firms likely differ in how strong this preference is, even within a narrow category of work: Some employers will pay more for more experienced, expert workers because of the importance or complexity of their tasks. This is important information for the workers, but it is hard to obtain objective information on this. So the planner contemplates facilitating cheap talk by introducing an easy way to express such preferences for firms and to search across this for workers.

Since the use of cheap talk to improve market efficiency has not received much empirical attention, we rely on a simple theory to highlight potential channels through which it might operate. The market under consideration for our empirical analysis (Horton, 2010) exhibits a market structure that resembles the theoretical structure in Kim and Kircher (2015): workers who choose to enter the market apply to job openings and include a wage demand, and firms

[1] A recent set of papers have explored heterogeneity in lob-loss rates among employed workers (Hall and Kudlyak (2022); Ahn et al. (2023); Gregory et al. (2021)). Gregory et al. (2021) explicitly consider employment in the traditional sense that excludes self-employment, contracting work or any other form of employment that is not covered by Unemployment Insurance. They find that 17% of that workforce has twice the likelyhood of loosing employment in the first year than to continue beyond the first year. For another 26% that chance is roughly even.

typically pick the worker they like and pay the demanded wage. The novel feature is the high premium on the experience and ability of applicants—a feature likely shared with most labor markets. To capture this, we extend the theory to accommodate (observable) ability differences between workers, and we allow more complex jobs to have higher ability requirements in order to produce positive output but to gain relatively more from matching with higher-ability workers. Firms can provide a cheap talk message about their "type"—whether they are primarily interested in high or low-ability workers. This message can help to resolve workers' uncertainty about precisely what kind of firm they face.

We characterize when it is an equilibrium for firms that can send a message to reveal their type truthfully, and show that this requires sufficiently heterogeneous preferences among employers. We then compare outcomes for such firms relative to identical counterparts that cannot send messages. Fife insights emerge here: First, cheap talk leads to more sorting in the sense that high-type firms see their applicant pool shifted towards higher worker types, whereas the opposite arises for low-type firms. Second, the effects of cheap talk on the number of applicants is ambiguous, but if low types value all workers less disclosing their type truthfully attracts less applicants. Third, cheap talk increases wages at high-type firms even when holding worker ability constant since high-ability workers know that they are valued more and extract parts of the surplus (and low types require higher wages to be compensated for more competition by high types). The opposite holds for low-type firms. This means that cheap talk not only allocates talent, but also shifts the distribution of surplus. Fourth, firms with access to cheap talk messages generate more social surplus than their non-message counterparts, and increases in the number of firms with access to messages improves efficiency. As a matter of robustness, we show in a fifth insight that these points carry over even if we compare markets where all firms can send messages compared with markets where none can, and the comparative statics apply more generally when comparing any informative cheap talk with no cheap talk.

The obvious caveat applies: cheap talk need not induce informative messages, as is the case in the well-known "babbling" equilibria. Whether an actual marketplace could benefit from introducing signaling therefore remains an empirical question. We approach this empirical question with our experiment. We structure the analysis along the lines of the previous insights. Our experiment introduced a clear and searchable language for employers to express their vertical preferences. In the experiment, employers self-reported a message about their willingness to pay for worker productivity, as proxied by worker experience. Workers could easily search for particular messages and direct their applications toward these. All employers were asked this question but only treated employers had their messages shown to job-seekers. This novel design allowed us to isolate the effects of the signal revelation on sorting.

We find evidence that messages are truthful: a randomly selected set of employers are asked about their type knowing that we will not show their answers to job seekers. A likely focal point for them is to simply report their preferences truthfully. We see that the distribution of answers varies across occupations in ways one might predict: for example, firms seeking administrative support ask less often for “expert” workers than those looking for assistance with Networking and Information Systems. Another randomly selected set of firms face the same question, but are informed that job seekers can see their answers. If there are strategic incentives to manipulate the answer to affect the application pool or the wage demands, one would expect a different distribution of answers. There is no detectable difference in any of our occupational groups.

Comparing firms with the same message across these groups allows us to analyze the effects of cheap talk. We also do this for a third randomly selected group of firms, which provide their answer without knowing whether we show it to job seekers or not. This way we can randomize whether messages are visible, ensuring that firms with the same message are indeed otherwise identical. Either way, we find four insights: First, the revelation of the cheap talk message induced substantial additional sorting of job-seekers, with firms requesting “experts” increasing the ability proxy of their applicant pool, while those requesting ”rookies” decreasing it. Second, the latter see a sharp decline in the number of applications. Third, job-seekers bid up against high-type firms and bid down against low-type firms if messages are revealed. Forth, this sorting and bidding strongly affected who was matched to whom and at what price without changing the number of matches formed. This lack of a decrease in matches formed came despite somewhat fewer applications being sent overall, suggesting an improvement in matching efficiency. Regarding match outcomes, we estimated that the revelation of employer preferences increased the total transaction volume on the platform by about 3%. This increase came from an increase in the quality of matches (but not the quantity), leading to larger within-relationship expenditure and hours worked. This increase in hours worked even occurred among employers selecting “high” vertical preferences who hired workers at higher wages. As employers decide on hours worked, this is strong evidence of match quality improvements. Regarding subjective evaluations, there is marginally significant evidence that employers rated the platform more positively post-transaction. Regarding robustness, we do not have much evidence on full roll-out towards the entire market, but at least the shares of messages remain rather stable among those firms who can send messages.

One novel contribution of this paper is to explicitly explore the use of cheap talk messages to improve the search process in matching markets, underpinning theoretical insights with experimental evidence. Our empirical evidence supports the idea that markets reach beneficial separating equilibria, which seem focal relative to uninformative ”babbling” equilibria. Although several papers have considered the effect of cheap talk in matching markets, these all have focused on 1:1 scenarios in which participants are privately signaling their interest in one particular counter-party (Lee and Niederle, 2015; Coles et al., 2013; Kushnir, 2013). Our focus on searchable cheap talk messages about general preferences is novel. Perhaps the closest related paper is Tadelis and Zettelmeyer (2015), but in this case, it was the platform that invested in the costly collection of additional objective information about quality to thicken markets and encourage sorting, as opposed to participants themselves using cheap talk for organization. Belot et al. (2018) study information provision about occupational fit for job-seekers, but it is their platform that determines the definition of "fit."

Another contribution of our paper is to show that a market-designing platform can substantively improve market information with low-cost market interventions. In contrast to conventional models that take information limitations as essentially a fixed feature, our paper shows these limitations are mutable. Furthermore, the exact intervention could be implemented by any other computer-mediated job board that controls how job posts are initially posted and displayed to job seekers.[2] As more of economic life and the job search process becomes computer-mediated (Kuhn and Skuterud, 2004; Kuhn and Mansour, 2014; Varian, 2010; Marinescu and Wolthoff, 2016), the opportunity to shape matching markets through purely informational interventions is likely to grow (Horton, 2017; Belot et al., 2018; Gee, 2019; Bhole et al., 2021).

A final contribution of the paper is providing strong evidence in favor of a directed search characterization of the matching process.[3] As we can observe applications in our empirical setting, we can observe how much of the search is already directed, with workers responsive to the observable attributes of jobs, not simply in where they apply, but in how they bid. Yet the explicit messages introduce a significant additional effect broadly in line with our model predictions. Although the sorting between workers and jobs that can both be ordered along one "quality" dimension has long been studied both in random search models (e.g., Shimer and Smith (2000), Teulings and Gautier (2004); Gautier and Teulings (2006), Eeckhout and Kircher (2010), (Bagger and Lentz, 2019)) as well as in directed search models (e.g., Shi (2001, 2002), Shimer (2005), Eeckhout and Kircher (2011), Cai et al. (2021)), no paper that we are aware of has explored how cheap talk could be embedded in a job search platform to facilitate sorting. Our theory borrows the basic multilateral meeting and one-on-one production structure from Shi (2002) and Shimer (2005), but instead of wage posting our firms only post messages. Workers

[2] E.g., Monster, Indeed, SimplyHired, LinkedIn, Craigslist, Facebook Jobs, Upwork, etc.

[3] For an overview of directed or competitive search models, see Wright et al. (2021), with fundamental contributions going back to, e.g., Peters (1991), Peters (1997), Moen (1997), Acemoglu and Shimer (1999), and Burdett et al. (2001). Apart from Menzio (2007) and Kim and Kircher (2015) discussed above, noteworthy recent contributions include Kim (2012) who studies market segmentation with a focus on adverse selection, and Albrecht et al. (2016) who study signaling in the housing market where messages entail some real costs.
Sorting based on costly signals has also been studied in matching games without search frictions where individuals can "burn money" to signal their types (see, e.g., Hoppe et al. (2009)). Clearly such signalling games where different announcements carry different costs are very different from the pure cheap-talk setting that we analyze and that characterizes many online markets.

instead ask for particular wages as we document for the empirical setting that we analyze and as modelled in Kim and Kircher (2015) (who abstracted from sorting considerations). Although our empirical context is an online labor market, the search and matching process is quite similar in its fundamentals to the matching process in more conventional markets, and as such, it is likely that our characterization generalizes.

The rest of the paper is organized as follows: Section 2 presents the model. Section 3 explains the empirical context and design. Section 4 presents the results of the experiment on the matching process and outcomes. We discuss directions for further research and conclude in Section 5.

2 A model of cheap talk in a directed search labor market

We first present a stylized model of a labor market that captures some of the salient features in our experimental setting. The salient features we have in mind are: workers can see a large number of employment options on an online platform; they have to select a small number (in the model, just one) to apply for work and include a wage demand. Workers have an outside option of not participating on the platform. Employers on the platform select the applicant they like best and pay the wage demanded. Workers differ in their ability, and jobs differ in their returns to worker ability. Job openings have text that entails a signal about the firm and its returns to worker ability, but is not codified and hard to search. Cheap talk is explicitly introduced by asking employers to choose one from a discrete set of options in the spirit of "looking for the highest ability even at high wages" or "looking to pay lowest wages, with less concern for ability", which are codified and easy for workers to search over.

The theory is intended to capture these features in a simple way but does not attempt to provide the largest theoretical generality. Rather, we aim for tractability to help organize our thoughts regarding the kind of analysis and results we might expect to find in the experimental setting where cheap talk is or is not facilitated. The key insights are collected in our main proposition (Proposition 1). For our derivations, we assume that the researcher can observe the true type of the firm, even though in the model this is the firms' private information that workers need to infer through cheap talk or signals. The subsequent experimental setting uses several treatment arms to achieve this.[4]

[4] In one treatment arm the firms first choose their message, and then the experimenter randomly decides whether to reveal that message to prospective workers or not. This ensures that the experimenter still observes the message even when workers do not observe it. If messages are truthful, this means that the experimenter observes the firm type. In a different arm firms know whether their message is shown or not before they choose their message. Assuming that firms have no incentive to lie when the message is not shown, we can study truth-telling for those for whom it matters, i.e., where the message will be shown to workers.

2.1 Setup

Players and Production: There is an endogeneous mass of workers with heterogeneous abilities $a \in \mathbb{R}_{++}$. Worker ability is known and publicly observable. There exists a unit mass of jobs that differ in their marginal valuation v_T of workers' ability: Each job can hire at most one worker, and a match between job type T and worker ability a produces output $v_T(a - c_T)$, where c_T allows more complex jobs to have higher ability requirements (i.e., "costs") to achieve positive output. For simplicity we focus on two types of jobs: $T \in \{H, L\}$. High types value ability more: $v_H > v_L > 0$, and their jobs might have higher requirements: $c_H \geq c_L \geq 0$. The fraction of high firm types is δ_H, and $\delta_L = 1 - \delta_H$ is the fraction of low firm types. The type of the job is *private* information of the firm that posts the job. We will use the words "job" and "firm" interchangeably. There is an infinite supply of workers of all types whose decision to use this platform is determined endogenously via an entry condition.

Payoffs for market participants: Both workers and jobs are risk neutral. If a worker is hired for a job, the firm's profit is the output minus the wage it pays to the worker. Therefore, a worker of type a who gets paid wage b at job T generates firm profits:

$$\pi(a, b|T) = v_T(a - c_T) - b. \tag{1}$$

A worker who gets hired has utility equal to his wage b. The payoff of jobs and workers who remain unmatched is normalized to zero.

Market entry: Workers who participate in the market forgo their non-negative outside option $u(a)$, which we assume to be increasing and strictly convex, with Inada conditions $u'(0) = 0$ and $u'(a)$ unbounded. For some proofs, it will be convenient to assume that the elasticity of u' is bounded away from zero and infinity. We further assume that for each job type T there is at least one worker type a for which surplus of entering is strictly positive ($i.e.$, $v_T(a - c_T) - u\,(a) > 0$). It will also be convenient to assume that the difference in requirements between the types of firms is not too large, and in particular that $c_H - c_L < u(c_H)/u'(c_H)$.

Market Interaction, Information and Timing: Workers who enter the market can apply to one job (for simplicity) and propose a wage. Workers can obtain two types of information about jobs to decide where to apply and what wage to ask for: they receive some exogenous signals about the type of the job, and some firms can also actively choose a message to send to the workers. For the latter, with i.i.d. probability ρ firm i is asked to choose a message $m_i \in \{L, H\}$, where the message space is without loss of generality given that we only have two types of firms. With probability $1 - \rho$ the firm cannot send a message, which we denote by $m_i = \emptyset$. For notational convenience, let $M = \{\emptyset, L, H\}$.

Each firm also generates two signals that workers can see: $s_{1,i}$ and $s_{2,i}$, each of which is drawn from set $\{L, H\}$. The role of these signals is to provide some information to workers even in the absence of cheap talk, which workers can use to choose the firm (based on first signal) and their wage bid (based on both signal to represent that workers have more information after they went through the process of preparing a full bid). We assume that signals are i.i.d. conditional on the true firm type T_i. The signal coincides with the true type with probability ψ_1 and ψ_2, respectively, with the obvious requirement that these probabilities exceed one-half.

The first signal can be easily observed by workers prior to choosing where to apply. From the joint distribution of messages and first signals, each worker chooses a job announcement with a desired message and first signal $(m, s_1) \in M \times \{L, H\}$. Since the worker lacks other information at this point, all firms with the same (m, s_1) have equal chances of being selected by the worker. In a large market, this randomness leads to a Poisson distribution of workers at a given firm, with Poisson parameter equal to the ratio of workers to firms at a given (m, s_1) combination.[5] The worker then prepares a bid for this job, and in this process, learns its second signal $s_{2,i}$. We assume that the second signal is very precise (ψ_2 close to 1). At that point, the worker already sunk his preparation time and cannot choose another firm, but s/he can adjust the bid b to the message and both signals, or chooses to remain unmatched. We assume that workers only submit a bid if s/he believes that the bid has a strictly positive probability of being accepted.

Each firm observes all its bids and the ability of the associated workers, hires the most profitable among these workers, and pays a wage equal to this worker's bid. It can reject all bidders.

Equilibrium Concept: Our equilibrium concept follows the literature on large games where each individual's payoff depends on his own actions and the distribution of actions across all other players (e.g., Mas-Colell (1984) and its applications to directed search in Eeckhout and Kircher (2010) and Peters (2010)). To summarize the distribution of actions by firms, let $\mu(m, T)$ denote the equilibrium probability that firms of type T send message $m \in \{H, L\}$ conditional on being selected to send a message. This uniquely determines the probability $\Psi(T|m, s_1)$ of facing firm T after observation (m, s_1) whenever m is played with positive probability. If m is not played with positive probability, $\Psi(.)$ is a belief that agents hold in case a firm sends such a message off-equilibrium. Given such beliefs, Bayes' rule generates the joint probability $\Psi(s_2, T|m, s_1)$ of the second signal and type, and a conditional belief $\Psi(T|m, s_1, s_2)$ after observing all signals.

The easiest way to sum up the workers' behavior is through the ratio $\Lambda(a, b_L, b_H|m, s_1)$ of workers with ability weakly above a who will bid weakly above b_L after a low second signal and above b_H after a high second signal, per job with initial observable (m, s_1). This is also called the

[5] This way of introducing search underlies most of the micro-founded work on directed search, and careful derivations that start with a finite population and takes limits are presented in Peters (1991), Peters (1997), Burdett et al. (2001) or the summary in Wright et al. (2021).

queue length of workers. Sometimes it is useful to focus on the marginal $\Lambda_a(a|m, s_1)$ of workers with ability above a per firm with observable (m, s_1), or on the conditional $\Lambda(a, b|m, s_1, s_2)$ that provides the ratio of workers with ability-bid above (a, b) to firms with observables (m, s_1, s_2).

From the perspective of a firm with type T and observables (m, s_1, s_2) this generates a queue of workers who generate a profit weakly higher than π according to

$$\Gamma(\pi|m, s_1, s_2, T) = \int_{\{a,b|\pi(a,b|T)\geq\pi\}} d\Lambda(\cdot|m, s_1, s_2).$$

We denote by $\Gamma_+(\pi|.)$ the queue of workers who generate a profit strictly higher than π, and the difference $\triangle(\pi|m, s_1, s_2, T)$ describes the mass of workers who generate exactly the profit π.[6] We will later rule out the existence of such mass points in equilibrium. The chance that no worker offers a profit weakly above π under a Poisson distribution has probability $\exp(-\Gamma(\pi|.))$, so the chance of having at least one worker who offers such profit is the complement. A firm of type T then obtains the following expected profit from sending signal m:

$$\Pi(m|T) = \sum_{s_1,s_2} \left[\psi(s_1, s_2|T) \int_{\pi\geq 0,s_1,s_2} \pi d(1 - e^{-\Gamma(\pi|m,s_1,s_2,T)}) \right], \tag{2}$$

where $\psi(s_1, s_2|T)$ denotes the probability of signals given the firm type. Consider now a worker who applies to a job with observables (m, s_1) and intends to bid $b(m, s_1, s_2)$. For simplicity denote as $b_L := b(m, s_1, L)$ and as $b_H := b(m, s_1, H)$ the bids after receiving the second signal. This worker's expected utility of applying to this firm is

$$V(m, s_1, a, b_L, b_H) = \sum_{s_2,T} \Psi(s_2, T|m, s_1) b_{s_2} \left(e^{-\Gamma_+(\pi(a,b_{s_2}|T)|m,s_1,s_2,T)} \right), \tag{3}$$

as long as the worker generates weakly positive profits (i.e., $\pi(a, b_{s_2}|T) \geq 0$) and there are no mass points of workers that generate the same profit (i.e., $\triangle(.|m, s_1, s_2, T) = 0$). The first factor inside the summation is the probability of facing signal s_2 and firm type T. The second factor is the worker's bid which he collects as a wage if hired. The final term captures the probability that the worker is hired, which equals the Poisson probability that no other worker is present who generates a strictly higher profit for the firm.

If the worker generates negative profits (i.e., $\pi(a, b_{s_2}|T) < 0$) firms of type T would not hire this worker. If there is a mass point of workers that generate exactly the same profit (i.e., $\triangle(\pi(a, b_{s_2}|T)|m, s_1, s_2, T) = 0$ for some s_2 and T) profits would be strictly lower than in (3) because the firm would randomize its hiring among these equally-attractive workers. Appendix

[6] The definition is $\Gamma_+(\pi|m, s_1, s_2, T) = \int_{\{a,b|\pi(a,b|T)>\pi} d\Lambda(\cdot|m, s_1, s_2)$.

A derives the exact expression that captures these cases. We can then define an equilibrium.

Definition (Equilibrium Definition)**.** *An equilibrium consists of firm behavior, worker behavior, and beliefs* $(\mu(\cdot), \Lambda(\cdot), \Psi(\cdot))$ *such that*

[1] beliefs $\Psi(\cdot)$ *are consistent with Bayes' rule according to firm behavior* $\mu(\cdot)$*,*

[2] firms maximize profits such that $\Pi(m|T) = \max_{m'} \Pi(m', T)$ *when* $\mu(m, T) > 0$*, and*

[3] workers' decisions to enter and bid are optimal such that: $V(m, s_1, a, b_L, b_H) \leq u(a)$ *for all* (m, s_1, a, b_L, b_H)*, and this holds with equality if* (a, b_L, b_H) *is in the support of* $\Lambda(\cdot|m, s_1)$*.*

The last condition ensures that there are no jobs that offer strictly higher returns to workers than their outside option, and those that do enter obtain exactly their outside option.

We conclude with the immediate observation that there cannot be meaningful mass points in equilibrium, i.e., mass points at firm types to which workers assign a non-zero probability. Workers at the mass point have strictly positive probability that another worker that generates the same profit for the firm gets hired instead of them. A marginally lower wage bid (i.e., a marginally higher profit for the firm) would have ensured being hired, inducing a profitable upward jump in the hiring probability. This rules out that mass points can be optimal for workers, simplifying the notation going forward.

Lemma. *There are no meaningful mass points in equilibrium, i.e.,* $\triangle(\pi|m, s_1, s_2, T) = 0$ *for all* $\pi \geq 0$ *and all* (m, s_1, s_2, T) *with* $\Psi(s_2, T|m, s_1) > 0$*.*

Proof: See Appendix A.

2.2 Empirical Prediction

Our main proposition compares firms that can send a cheap talk message (according to probability ρ) with those who are not asked to send a message (according to probability $1 - \rho$). We focus on a setting where firms who send a message reveal this message truthfully. We will later prove that this is an equilibrium when valuations are sufficiently far apart, but will also discuss other possible equilibria (see Subsection 2.5).

Proposition 1. *The following compares firms that engage in truthful cheap-talk with otherwise identical counterparts that cannot.*

1. *Ability sorting: Truthful cheap talk induces more assortative matching, in the sense that the distribution of ability for a randomly selected applicant as well as for the actual hire increases (resp., decreases) in first order stochastic dominance for high (resp., low) firm types.*

2. *Number of applications: In general the effect on the number of applications is theoretically ambigous. One can characterize limiting cases: All else equal, if v_H becomes sufficiently large then truthful cheap talk reduces the applications at low-type firms relative to their "no-message" counterparts, with opposite effects on high-message firms. For given $c_H > c_L$ and all else equal, if v_H is sufficiently low and therefore close to v_L then truthful cheap talk increases the applications at low-type firms relative to their "no-message" counterparts, and the opposite holds for high type firms.*

3. *Wage bidding: Any given worker type bids less at low firm types under truthful cheap talk compared to the absence of cheap talk. The average wage demand at such firms and the average wage of the selected worker fall under truthful cheap talk. The opposite holds for high firm types.*

4. *Efficiency: Social surplus per firm is higher at firms who send truthful messages than at their counterparts who cannot send messages.*

5. *Market-wide rollout: Incentive to use different messages do not dependent on the fraction of firms that use cheap-talk (since equilibrium wage bids and queue lengths conditional on m, s_1, s_2 do not depend on the fraction ρ of firms that can message).*

Proof: Developed in Section 2.3. □

The intuition behind the proposition is that workers have better information about the firm type when they see the messages. Higher ability worker find it more worthwhile to enter as they can better target the high firm types that value them. Lower ability workers are also willing to enter because they find it easier to avoid firms with complex (high cost) jobs where their output would be too low. And workers with intermediate ability find it easier to target the right firms, where those who target high firms ask for higher wages but face tougher competition from above while those who target lower firms find it easier to match despite the need to lower their wage demands.

While intuitive, the structure of the proof is slightly involved. Figure 1 provides more details. Consider first Panel A on the left. It characterizes the job application behavior under truthful messages. Workers split into three groups: the lowest ability workers who enter approach only low firms (though even lower ability workers do not enter the market at all). Workers with middle ability look as if they choose firms at random. Workers with high abilities focus exclusively on high firms (and even higher ability workers stay out of the market because of superior outside options).

Compare this with Panel B on the right. Now workers have to make their choice based on the first signal, and sometimes mistake a low firm for a high one and vice versa. From the

Figure 1: Illustration how workers approach firms with and without messages.

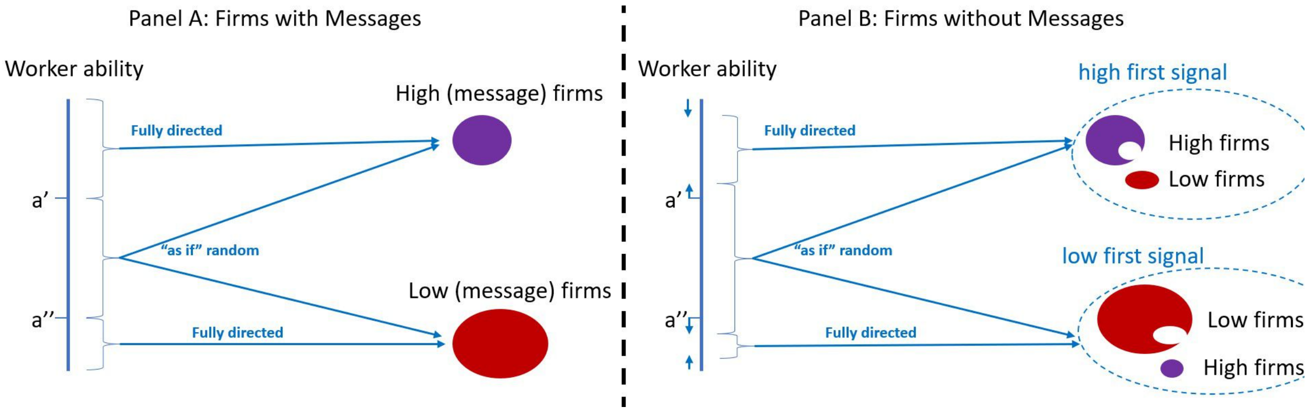

Notes: Panel A of this figures shows the application behavior per worker-ability type a to firms that truthfully send a high or low message. Panel B shows the application behavior per worker-ability type a to firms that have a high or low first signal and no message. Note that in Panel B the set of worker types that are directed are more compressed while the set of types that apply at random expand, relative to Panel A.

workers' perspective this turns out to be rather similar to the setting before, only that workers now consider some notion of average valuation and average cost within the firms that share the same signal. Some high ability workers stay out since the valuation of the firms they meet at the high signal on average has gone down. Similarly, some low ability workers now stay out because at the low signal the costs have gone up. The intermediate set of worker types that select at random expands. So for a given high type firm the distribution of workers it meets deteriorates relative to Panel A for two reasons: even if it emits a high signal it looses out on some of the best workers and attract some of the least able ones instead, and sometimes it emits the low signal and gets the associated set of low worker types. Wage demands decline because workers are less sure if they are highly valued (not depicted). For low firm types the opposite happens.

The next subsections provide some of the theoretical details. They start by considering truthful cheap talk to establish the results depicted in Panel A of Figure 1, and study the associated wages and establish constraint efficiency. They then show that the interaction without messages in Panel B has a similar structure to that in Panel A, with associated consequences for wages. Finally, the incentives for truthful cheap talk are studied, and the section concludes with several insights on robustness.

Empirically inclined readers that find the intuition in this subsection sufficient are invited to jump to Section 3. That sections provides more details on the empirical setups, and the subsequent section first studies indicators of truth-telling, and then investigates empirical analogues to the points in Proposition 1. But for more theory-inclined readers the next sections provide the main steps of the proof, which exploit the beauty of this intentional simplistic setup which enables theoretical insights through a sequence of relatively easy first-order conditions and comparative statics. We then return to questions of incentive-compatibility and robustness.

2.3 Theoretical Analysis (including proof of Proposition 1)

Firms with messages: To understand this market, consider first the firms that can send a cheap-talk message. Their message space is $\{L, H\}$. As we will see, the following analysis is valid whether these firms constitute the entire market (i.e., $\rho = 1$) or just a part of the market (i.e., $0 < \rho < 1$). Assume for now that these firms declare their type truthfully in equilibrium (i.e., $\mu(T|T) = 1$)), and we impose equilibrium conditions [1] and [3].[7]

Under the assumption that firms report truthfully, workers do not rely on signals to infer the firm type and we ease notation by suppressing the reference to signals.[8] As a simplification of (3), we can now express the expected value for a worker of ability a who bids b at a job with message $m \in \{L, H\}$ (no matter what the signals are) by

$$V(m, a, b) = e^{-\Gamma(\pi(a,b|m)|m)} b, \tag{4}$$

as long as the workers generates weakly positive profits $\pi(a, b|m) \geq 0$. Since the worker perfectly infers the firm type on the equilibrium path, it is possible to reformulate the wage bid b into the profit $\pi \geq 0$ that the worker leaves to the firm. The worker then keeps the output minus this profit. For worker-firm pairs that have strictly positive output so that $v_m(a - c_m) \geq \pi \geq 0$ this delivers expected utility

$$W(m, a, \pi) = e^{-\Gamma(\pi|m)} \left[v_m(a - c_m) - \pi \right]. \tag{5}$$

Clearly $V(m, a, b) = W(m, a, \pi(a, b|m))$. By equilibrium condition [3], if workers with ability a and bid b enter at message m — i.e. if (a, b) is in the support of $\Lambda(\cdot|m)$ — then b must maximize (4). But then $\pi(a, b|m)$ must be weakly positive and maximizes (5).

We can therefore study the problem where the worker's choice variable is π, which is chosen to maximize (5). This objective is supermodular in (a, π), indicating that it is less costly for high ability types to offer high profits to firms. It implies that higher ability workers at a given message leave higher profits in equilibrium. And since there are no mass points, this means that they leave strictly higher profits almost everywhere.

[7] In the next subsection, we additionally impose equilibrium condition [2] on firm optimality to study the conditions under which firms would indeed be willing to report truthfully.

[8] Workers not relying on signals simply means that $\Psi(T|T, s_1, s_2) = 1$ for all s_1 and s_2, and it is easy to show that there is no benefit from conditioning the workers' actions on the the signals. Therefore, to ease notation, we suppress references to the signals, i.e., $b = b_H = b_L$ and we write $\Lambda(a, b|m)$ instead of $\Lambda(a, b, b|m, s_1)$.

For ease of notation, let

$$U(a|m) := \arg\max_{\pi} W(m, a, \pi) \tag{6}$$

denote the equilibrium utility for worker type a that approaches message m, and free entry then requires that $U(a|m) \leq u(a)$, with equality if a is in the support of $\Lambda(a|m)$.

We can use this together with the definition of (5) to prove that at a given message m the worker abilities that enter in equilibrium form a non-empty interval $[\underline{a}_m, \bar{a}_m]$, and no single type enters with positive mass (see proof in Appendix A). The reason is that there is no mass point in terms of profits that are offered (see earlier Lemma) and for each profit level, there is a different ability type that finds it optimal to enter.

So the task is to characterize the non-empty interval $[\underline{a}_m, \bar{a}_m]$. Within this interval, consider a worker with ability a that leaves profit π at message m in equilibrium.[9] Then, the marginal increase in equilibrium utility with respect to ability is

$$e^{-\Lambda(a|m)} v_m = u'(a). \tag{7}$$

The left side follows from the differentiation of (5) with respect to ability, where any change in profits can be neglected by the envelope condition, and then substituting $\Lambda(a|m)$ for $\Gamma(\pi|m)$ which follows since higher ability workers leave higher profits. The right hand side has to hold because the equilibrium utility equals the forgone entry cost. Differentiating once more with respect to ability leads to

$$\lambda(a|m) e^{-\Lambda(a|m)} v_m = u''(a) \tag{8}$$

where $\lambda(a|m) = -d\Lambda(a|m)/da$ is akin to a "density" of type a workers queueing at this message. Dividing this by (7) yields

$$\lambda(a|m) = \frac{u''(a)}{u'(a)}, \tag{9}$$

which ties down the type density at any ability level. The key question concerns the end-points of the interval. Clearly at the highest type (7) reduces to the end-point condition

$$v_m = u'(\bar{a}_m), \tag{10}$$

which uniquely determines $\bar{a}_m$. One can then recover the lowest ability type that approaches message m since at any candidate $\underline{a}_m$ the queue length $\Lambda(\underline{a}_m|m) = \int_{a'}^{\bar{a}_m} \lambda(a|m)da$ is uniquely determined by (9) and (10). Moreover, the lowest ability type also leaves the lowest profits, which implies that he leaves zero profits (assume not, then this type could lower the bid so he

[9] Formally this means that there exists (a, b) in the support of $\Lambda(a, b|m)$ such that $\pi = \pi(a, b|m)$.

leaves zero profits and still win with the same probability). Then it has to hold that

$$e^{-\Lambda(\underline{a}_m|m)} v_m(\underline{a}_m - c_m) = u(\underline{a}_m), \tag{11}$$

which uniquely ties down $\underline{a}_m$.

Finally, we briefly consider how wage bids $b_m(a)$ are determined for worker $a \in [\underline{a}_m, \bar{a}_m]$. In equilibrium, the utility of such a worker is $e^{-\Lambda(a|m)} b_m(a)$. We can totally differentiate this with respect to the worker type, and set it equal to $u'(a)$ which needs to hold due to free entry. Using (7) this reduces after cancellation to the following differential equation:

$$\begin{aligned} b'_m(a) &= v_m - b_m(a)\lambda(a|m) \\ &= v_m - b_m(a)u''(a)/u'(a) \end{aligned} \tag{12}$$

which together with end-point condition that the lowest type extracts all suprlus:

$$b_m(\underline{a}_m) = v_m[\underline{a}_m - c_m] \tag{13}$$

determines the wage bids uniquely. For the following summary of these insights, it is useful to define function $B(a)$ on the positive reals uniquely via differential equation $B'(a) = 1 - B(a)u''(a)/u'(a)$ with $B(0) = 0$, which can be thought of as the "normalized" wage at a firm with unit valuation and no costs. We then obtain:

Proposition 2. *Under truthful cheap talk, the interval of worker types $[\underline{a}_m, \bar{a}_m]$ that approach message $m \in \{L, H\}$ and the queue $\Lambda(a|m)$ are fully determined by (9) - (11), and the wages they ask for are fully determined by (12) and (13). On $[\underline{a}_m, \bar{a}_m]$ the wages for such a firm with arbitrary v_m and c_m can be expressed relative to the "normalized" wages according to*

$$b_m(a) = v_m[B(a) - B(\underline{a}_m) + \underline{a}_m - c_m]. \tag{14}$$

It is worth noting that Proposition (2) fully characterizes the application and bidding behavior that firms face when they send messages, and that holds irrespective of whether all firms send such messages ($\rho = 1$) or only some fraction of the firms sends such messages ($\rho < 1$). The characterization makes no reference to the number of firms that have the ability to message, nor to the behavior at firms that do not. This property of directed search models with entry is known in the literature as "block-recursiveness", and comes about since workers can see the message and enter for this message exactly as to equalize their outside option. If many firms with this message are present, many workers enter, and the opposite happens if few firms are present. A separate queue forms for each combination of messages or signals. Since workers can choose which firms to approach, for those who approach firms that send messages it is not important what happens at firms that do not send messages. We will now turn to those firms.

Firms without messages: The firms that cannot send a message could be in a different market where cheap-talk is impossible (i.e., $\rho = 0$), or they could be in the same market as the firms that can send messages (i.e., $0 < \rho < 1$). Workers can only rely on the signals about these firms when they decide which firm to approach and which wage bid to place. For ease of exposition and to be able to leverage previous derivations, assume that the second signal is so informative as to eliminate uncertainty about the firm type (i.e., $\psi_2 = 1$). We provide the appropriate analysis for $\psi_2 < 1$ in the following subsection, and show convergence so that our results apply for high enough precision. Assuming perfect precision of the second signal allows us to consider the consequences of the first message, and enables us to draw heavily on the insights already developed earlier.

Now consider workers who see signal s_1 and believe that the fraction of low firms at this signal is $\Psi(L|s_1)$ where we suppressed the message because it is absent ($m = \emptyset$). The fraction with high type is evidently the complement. Note that these fractions are exogeneously given; for example, $\Psi(L|L) = \psi_1\delta_L/(\psi_1\delta_L + (1-\psi_1)\delta_H)$, and do not depend on the equilibrium interaction.

It turns out that workers approach these firms similar to firms whose type is the average valuation given their beliefs. That is, they approach firms as if they know that their productivity is averaged according to

$$v_\Psi = \Psi v_L + (1-\Psi)v_H,$$

where $\Psi = \Psi(L|s_1)$ represents the belief that the firm is of low type given its signal. Workers also approach firms as if they had the average cost, though the exact expression for this requires some re-arranging. We used to write production as $v(a-c)$, which evidently equals $va - vc$. So in this new setting worker expect the average of vc as the second term. If we want to write production in expectation as $v_\Psi(a - c_\Psi)$, we have to define

$$c_\Psi = \frac{\Psi v_L c_L + (1-\Psi)v_H c_H}{\Psi v_L + (1-\Psi)v_H},$$

where c_Ψ is in $[c_L, c_H]$ and monotonically increases from c_L to c_H when Ψ decreases from one to zero. Workers with types in $[\underline{a}_\Psi, \bar{a}_\Psi]$ then approach such firms, and their bids are then conditional on the second signals that indicates whether the firm is of low or high type. Bids are still driven by (14), except that the lowest ability type that extracts all surplus is determined by the first signal and not the second. We therefore have:

Proposition 3. *Firms with $m = \emptyset$ and signal s_1 induce belief $\Psi = \Psi(L|s_1)$. The interval of worker types $[\underline{a}_{s_1}, \bar{a}_{s_1}]$ that approach them and the queue $\Lambda(a|s_1)$ are fully determined by (9) - (11) when using for the valuation the average given by v_Ψ and for the cost c_Ψ. After the second signal s_2 these workers bid according to*

$$b_{s_2}(a) = v_{s_2}[B(a) - B(\underline{a}_{s_1}) + \underline{a}_{s_1} - c_{s_2}]. \tag{15}$$

Proof idea: The key challenge is to prove that workers who approach a firm will actually always end up bidding there, which is not obvious as they approach a firm under uncertainty and learn additional information from the second signal. Under our assumption that cost differences between firms are not too large this turns out to be true. The averaging then obtains because the second signal is so precise (the role of the second signal is discussed in detail in Appendix B). For the full proof see Appendix A.□

For the wage equation it is worth noting that the wage conditions on the true valuation and cost (v_{s_2}, c_{s_2}) that are revealed through the second signal. But at the time where the worker decides which job to bid for, he only has the first signal as he has not spent enough time on preparing his bid to know all the information about the job. So lowest type $\underline{a}_\Psi$ who extracts all surplus makes his decision where to apply based on the average quality of the firm.

Comparing firms with and without messages: To compare firms, it is therefore useful to study the comparative statics on the types of workers who approach these firms, and their bidding. The worker types are governed by the valuation v and costs c, which are equal to the true type for firms with truthful messages and equal to the average (interim) type given after workers have observed the first signal for firms without messages.

Corollary 1. *Consider a firm with interim expected valuation v and cost c.*

- *The maximal worker type $\bar{a}$ strictly increases in valuation v (but does not depend on c).*
- *The minimal worker type $\underline{a}$ strictly increases in cost c (but does not depend on v).*
- *For interior worker types in $(\underline{a}, \bar{a})$ the density $\lambda(a|.)$ is independent of v or c.*

Proof: The first point follows directly from (10), which depends on the firm's valuation but not its threshold for ability. The third point follows from (9), which depends on neither valuation nor threshold. The second point states that the lowest type $\underline{a}_m$ is only a function of the threshold, but not of the valuation. This is not directly obvious from (11), but in Appendix A we show that for a given $\underline{a}_m$ the utility at this ability stays constant when v_m changes when one adjusts for the associated changes in the queue $\Lambda(\underline{a}_m|m)$: When the valuation increases more higher types are added to leave the utility at the low end exactly unchanged. Because of this feature one can find a different way of constructing the lowest type. First, define $\underline{v}_m(a) = u(a)/[a - c_m]$. This is the lowest valuation at which type a workers could enter and recover their outside option. Then let $\underline{a}_m = \arg\min \underline{v}_m(a)$, i.e., it is the ability level that needs the lowest valuation to make entry possible. Using the first order condition for this minimization reveals that the lowest type can be solved for as the unique solution to $\underline{a}_m = c_m + u(\underline{a}_m)/u'(\underline{a}_m)$. The lowest type $\underline{a}_m$ is increasing in c_m, which follows by implicitly differentiating. $\square$

One can directly apply these comparative statics noting that v and c are lowest for firms with a truthful low message, higher for firms without message but a low first signal, even higher for firms without a message but a high first signal, and highest for firms with a truthful high message. That means that both the lower and the upper bound of the interval $[\underline{a}, \bar{a}]$ of workers who approach these firms is increasing in this order. Whenever these interval overlap across firms with different messages/signals, those workers type that overlap generate the same queue at different firms according to (9). This means that they can implement this queue length by choosing among these firms at random. An individual worker might not choose at random, but the outcome is indistinguishable from one where they do.

We can now return to Figure 1 that we used in preceeding subsection to illustrate graphically how workers choose between firms. Panel A illustrates firms that can send messages, and do so truthfully. Among those firms, the lowest type workers only approach low message firms because the ability threshold c_H at high firms is too high. Workers with intermediate ability choose "as if" at random between firms at both messages. Here $a' = \bar{a}_{m=L}$ and $a'' = \underline{a}_{m=H}$, so that the interval between indicates the overlap between $[\underline{a}_{m=L}, \bar{a}_{m=L}]$ and $[\underline{a}_{m=H}, \bar{a}_{m=H}]$. High worker types again are fully directed, this time towards high message firms which have sufficiently high valuation to warrant their entry cost. Whether high or low-message firms attract more applicants overall depends on parameters, but it is clear that the distribution of worker abilities at high message firms first order stochastically dominates that at low message firms.

Panel B considers the firms cannot send messages, and workers rely on the first signal to determine where to go. At the high signal there are mostly high type firms, but some low type firms also send this signal. Similarly, at the low signal there are mostly low firms, but some high firms also send this signal. Between the high and low signal firms, the workers choose in a similar fashion as in Panel A. But relative to Panel A, in Panel B the lowest and highest worker types stay out of the market, as the average costs are too high at the low end and the average valuation is too low at the upper end. Moreover, the range of worker types that act as if they choose randomly expands both at the upper and lower end. So in Panel B the types that are directed are a strict subset of the types in Panel A, while the types that visit at random are a strict superset.

Firms with and without messages might interact in the same mearket (which arises for interior message probability ρ, but by block-recursivity the insights also apply if one switches the market from no messages where $\rho = 0$ to enabling messages for all firms where $\rho = 1$). Within a market the worker types that approach both message and no-message firms do so as if choosing at random. For example, workers with ability slightly above a' apply to high message as well as to no-message firms, and among no-message firms they apply both those those with high and low first signal. Their queue length is the same at all of them, so they can choose between them at random to implement this. These workers do not apply to low message firms, though.

What are the consequences for wages if the applicant pool changes at a given firm type? At a given firm, workers of a particular ability will adjust their wages depending on the queue of other applicants, which changes depending on whether the firm can send a (truthful) message or not, and in the absence of a message whether the first signal is high or low. They do this even though the second signal allows them a clear idea of the firm's identity. The following corollary summarizes comparative statics for the wage at a particular ability level. It allows comparisons across firms, but the most notable insight arises when holding the firm type constant: a better applicant pool leads any given worker to increase his/her wage. The reason is that there is less competition from lower-ranked workers:

Corollary 2. *Consider a firm with type T and associated valuation v_T and cost c_T that attracts workers with abilities in $[\underline{a}, \bar{a}]$. For a given worker with ability a, his wage bid at this firm increases in v, decreases in c, and increases in $\underline{a}$.*

Proof: Follows directly from (15).□

If one compares a high type firm that sends a truthful message with a high type firm that cannot send a message, for any given worker the wage demand is slightly lower in the no-message firm in the event of a high first signal and substantially lower in the event of a low first signal. On the contrary, if one compares a low type firm that can send a truthful message with a high type firm that cannot, for any given worker his/her wage demand is slightly higher in the no-message firm in the event of a low first signal and substantially higher in the event of a high first signal. So for high types firms wages increase if they can send a message, while they decrease at low type firms.

We can now complete the proof of our main Proposition 1: Part 1 and 3 follow directly from Corollary 1 and 2. For parts 2 and 4 see Appendix A. Intuitively, the number of applicants is ambigous as cheap-talk moves both the upper and lower bound of the types that approach a firm, and it depends which shift is more significant. For the surplus, note that the equilibrium implements the planner's allocation given the information frictions, but a planner can assign worker types in a more targeted way if he has more information about the firm type. Part 5 follows directly from the "block-recursivity" property. □

These insights build intuition for the type of differences we expect to see with the introduction of cheap talk, if messages are indeed truthful. The next subsection provides some important steps that were omitted and discuss robustness, followed by a discussion of incentive compatibility that leads over to the empirical section.

2.4 Omissions, Extensions and Robustness

Two features of the previous analysis warrant further discussion. First, in the absence of messages the derivations so far have considered the limit where the second signal is perfect ($\psi_2 = 1$). This simplified the exposition because workers either chose one set of firms or the other, or they choose at random, as illustrated in Figure 1. In Appendix B.1 we discuss that our theory is best viewed as an approximation for the case where the second signal is not fully informative, and we show that equilibria converge to the limit equilibrium when precision increases (i.e., when $\psi_2 \nearrow 1$).

Second, we have focused on the particular production function $v_T(a - c_T)$. In Appendix B.2 we consider the broader class $\phi(a, T)$ of increasing production functions that are concave in ability, and ask which general properties underlie our main insights illustrated in Figure 1. For our particular environment we showed that the highest ability at H-firms is higher than the highest ability at L-firms, and this is ensured under more general production functions when output ϕ is super-modular. Our proofs were simple because workers that approach both types of firms form queues as if they choose at random. This arises withing the broader set of production functions when the elasticity of ϕ_a with respect to ability is constant in firm type. Finally, among such production functions only those that are log-supermodular imply that the lowest ability at H-firms is higher than the lowest ability at L-firms. If these conditions hold, the insights from Figure 1 continue to apply to this larger class of production functions.

2.5 Truth-telling

Subsection 2.3 assumed that firms report their type truthfully through their messages, and imposed only equilibium conditions [1] and [3]. Now we additionally impose condition [2] to investigate whether truth-telling is indeed incentive-compatible and therefore an equilibrium.

The basic trade-off is the following: a firm with low type who deviates from the candidate equilibrium and sends an H message is taken for a high firm, because in equilibrium only high firms send such messages on the equilibrium path. Such a deviant firm would attract better workers, but each of these workers demands a higher wage. The opposite happens for high firm types that pretend to be low types. Since high firm types care more about worker types and low firm types more about the wage, we can show that sufficient differences in valuations indeed imply the existence of an equilibrium with truthful cheap talk. As in many cheap-talk games, if firm types are very similar full type revelation generically fails.

Proposition 4. *For given $u(a)$, v_L and c_H, thruthtelling is incentive compatible if v_H is sufficiently large. For $c_H = c_L = 0$, when v_H approaches v_L, incentive compatability for thruth-telling generically fails.*

Proof: See Appendix C.1.□

This highlights that incentives for truth-telling are there in some situations, and we outline in Appendix C.2 that this carries over to the somewhat broader class of production functions discussed in the preceding section. But Proposition 4 also indicates that such incentives are not guaranteed in all situation. Beyond sufficient differences in firm types, incentives for truth-telling also rely on the presence of the second signal, as we discuss in Appendix C.3. Comfortingly, we do not just have to speculate on truth-telling but can investigate proxies for truth-telling in the empirical exercise. Moreover, one can show that our insights do not rely on perfect type revelation, but apply even to partially revealing equilibria:

Corollary 3. *For any partially revealing equilibrium where messages carry some information (i.e., where $\Psi(T|L, s_1) \neq \Psi(T|H, s_1)$) the comparison of firms with and without messages is characterized by points (1)-(5) of Proposition 1.*

Proof: The first three points in Proposition 1 are based on comparative statics that continue to apply - see Appendix C.4. Point (5) remains valid because block-recursivity persists, and the efficiency proofs proceeds in analogous steps as in the full-information case. □

Finally, to build the bridge to the ensuing empirical section, we conclude by discussing how we interpret "truth-telling" in an empirical setting: For this, we ask firms without a message a fictitious question: Do you prefer ("rookie", $) or ("expert", $$$)? Firms arguably interpret this as choosing between two combinations of ability and wage-bid: (a_1, b_1) or (a_2, b_2) with $a_1 < a_2$ and $b_1 < b_2$. Since there is no consequences to the question, firms might simply select the option that would generate more profits. Under this assumption, if two firms choose differently, the one that chooses the first must be the L type and the one that chose the second must be the H type. This holds not only for our specific production function, but for any supermodular production function, and one can extend this idea to more than two types.[10] Emirically we check if the fraction who chooses the first answer varies meaningfully across different occupations, and then consider the other firms that can send messages by answering the same question. They have a strategic incentive since their answer will be shown to workers, but if they continue to answer truthfully their distribution of answers should be similar to that of the other firms. We will next describe the empirical setup, followed by the empirical results on truth-telling, and subsequently we return to the empirical counterparts of the five points in Proposition 1.

3 Empirical context

The empirical setting for our analysis is a large online labor market. In these markets, employers hire workers to perform tasks that can be done remotely. Markets differ in their scope and focus, but common services provided by the platform include soliciting and promulgating job openings, hosting user profile pages, processing payments, arbitrating disputes, certifying worker skills, and maintaining a reputation system (Horton, 2010; Filippas et al., 2018).

In the online labor market we use as our empirical setting, would-be employers write job descriptions, self-categorize the nature of the work and required skills, and then post the job openings to the platform website. Job openings are learned about by workers via electronic searches or email notifications. Employers can also search worker "profiles" and invite workers to apply for their openings (Horton, 2017). Worker "profiles" are similar to resumes, containing the details of past jobs completed by the worker, education history, skills, and so on. For both workers and employers, some of the information available to the other side of the market is "hard" in the sense that it is verified by the platform. Examples of verified, public information include hours-worked, hourly wage rates, total earnings, and feedback ratings from past trading partners.

If a worker chooses to apply to a particular job opening, they submit an application, which includes a wage bid (for hourly jobs) or a total project bid (for fixed-price jobs) and a cover letter. In our analysis, we only make use of hourly job openings, as the preference revelation opportunity was only available for hourly job openings.

After a worker submits an application, the employer can choose to interview the applicant. They can also hire an applicant at the terms proposed in the application, or make a counteroffer, which the worker can counter, and so on. The process is not an auction and neither the employer

[10] One can easily extend this logic to more messages: In the empirical setting we also ask additionally about ("intermediate", $$). Assume firms interpret this message as ability-wage combination $(a_{\text{medium}}, b_{\text{medium}})$ with $a_{\text{medium}} \in \{a_1, a_2\}$ and $b_{\text{medium}} \in \{b_1, b_2\}$. If there is a third type of firms with $v_{\text{medium}} \in \{v_L, v_H\}$ and some firms choose each of the messages, then L firms choose the low ability, $Medium$ firms choose the medium ability and H firms choose the high ability due to super-modularity between the firms valuation and ability.

nor the worker are bound to accept any offer. Despite the possibility of back-and-forth bargaining, it is fairly rare, with about 90% of hired workers being hired at the wage they initially proposed (Barach and Horton, 2021). In this market, employers typically collect a more or less complete pool of applicants and then select a subset to interview and ultimately hire (which also seems to characterize the process in conventional markets (Davis and de la Parra, 2021)).

To work on hourly contracts, workers must install custom tracking software on their computers. The tracking software essentially serves as a digital punch clock. The software records not only the time spent working (to the second), but also the count of keystrokes and mouse movements. The software also captures an image of the worker's computer screen at random intervals. All of this captured data is sent to the platform's servers and then made available to the employer for inspection, in real time. These features give employers tools to precisely monitor hours-worked, and to an extent, effort. As employers can end contracts at will, the employer can be thought of as the party choosing hours-worked.

The marketplace we study is not the only market for online work, and so it is important to keep in mind the "market" versus "marketplace" distinction made by Roth (2018). Relatedly, a concern with treating job openings as our primary unit of analysis is that every job opening we see on the platform could be simultaneously posted on several other online labor market sites and in the conventional market. However, survey evidence suggests that online and offline hiring are only very weak substitutes and that multi-homing of job openings is relatively rare. When asked what they would have done with their most recent project if the platform were not available, only 15% of employers responded that they would have made a local hire. Online employers report that they are generally deciding among (a) getting the work done online, (b) doing the work themselves, and (c) not having the work done at all. The survey also found that 83% of employers said that they listed their last job opening only on the platform in question.

3.1 Experimental design

During the experiment, employers posting job openings were asked for their vertical preference, using the interface shown in Figure 2. The choice was mutually exclusive and was mandatory. Employers selecting "Entry Level ($)" are referred to as "low" throughout the paper, those selecting "Intermediate ($$)" as "medium," and those selecting "Expert ($$$)" as "high." The use of varying dollar symbols to indicate an option's relative position in some vertical price/quality space is commonplace, particularly in online settings (e.g., Diamond and Moretti (2018)). Recall the discussion of "truth-telling" in such a message space at the very end of Section 2.

The platform's goal for the intervention was to give market participants more information and encourage better matches. There are several papers that explore the effects of a "platform"

Figure 2: Vertical preference signaling opportunity presented to employers when posting their job openings

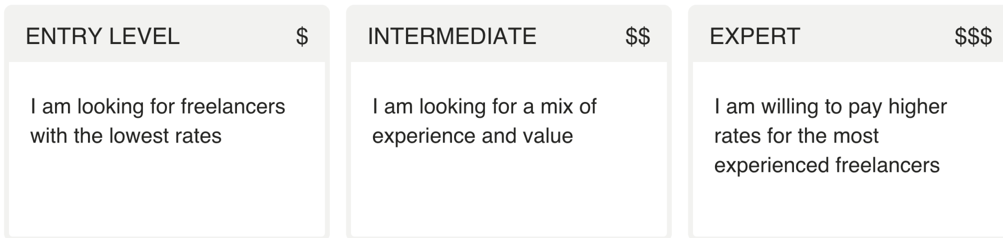

Notes: This figure shows the vertical preference signaling interface and language presented to employers when they posted a job.

changing the information available, which is typically about sellers, such as their quality (Luca, 2016; Jin and Leslie, 2003), past experience, (Barach and Horton, 2021) and capacity to take on more work (Horton, 2019). The stylized fact of these information disclosures is that they redirect buyers to "better" sellers, and, in the shadow of this effect, improve seller quality.[11] This is distinct from our approach which is not trying to change incentives for quality per se, but rather simple improve matching.

The experiment was run by the platform from 2013-07-18 to 2013-12-05. A total of 50,877 employers were allocated to the experiment. These employers collectively posted 220,510 job openings.[12] Upon posting a job opening, employers were randomized to one of two experimental "arms," with each arm having two groups. The two arms of the experiment and their component experimental cells with their allocations are listed in Table 1. These allocated job posts were all "normal" job posts in the market, by real employers trying to complete actual tasks and spending their own money, creating a true "field" context for a market design intervention (Harrison and List, 2004).

In the two cell "explicit arm," employers knew for certain, *ex ante*, whether their tier choice would be revealed. We use an indicator variable, ShownPref, to indicate whether preferences were revealed. Because the value of ShownPref was known by employers *ex ante* in the explicit arm, tier choice cannot be considered exogenous: an employer might claim "high" preferences when they know the choice will not be shown, but "medium" when they know the choice will be shown. This conditioning is not a concern in our other experimental arm, the two cell "ambiguous arm," in which employers were told that their choice *might* be shown to job-seekers. In this arm,

[11] As an example of how this works in another online market, Lewis (2011) shows that on eBay, the revelation of information about quality (through descriptions and prices) and the contracts created by these disclosures largely overcome the adverse selection problem.

[12] The duration of the experiment was chosen *ex ante* by the platform to detect a 1 percentage point change in the fill rate with 80% power, but the experiment was ultimately run substantially longer than this for unrelated business reasons, i.e., one author was traveling and neglected to turn the experiment off at the agreed-upon date.

Table 1: Description of the arms of the experiment and the experimental groups

	Allocation	Vertical Preference Shown to Job-Seekers? (ShownPref)	Employer knows *ex ante* whether signal will be revealed:
Explicit Arm			
ShownPref = 1	16,011; 32.8%	Yes	Yes
ShownPref = 0	15,767; 32.3%	No	Yes
Ambiguous Arm			
ShownPref = 1	11,344; 23.3%	Yes	No
ShownPref = 0	5,649; 11.6%	No	No

Notes: This table lists the cells of the experiment and the number of assigned employers. The fraction in each cell is also reported. Employers made the vertical preference signaling choice when they posted their opening. See Figure 2 for the actual interface. Employers in the two-cell explicit arm were told *ex ante* that the platform would reveal or would not reveal their vertical preferences to workers. Employers in the ambiguous arm were told that the platform *might* reveal their preferences to workers; whether workers were shown employer vertical preferences was randomly determined *ex post*. If ShownPref = 1, would-be applicants could observe the employer's vertical preference before applying, otherwise they could not, ShownPref = 0.

employers were then randomized to either have their choice revealed or not. For these employers, tier choice can be regarded as exogenous, as it is chosen before ShownPref is determined. If the employer's preferences were to be revealed, their job opening was labeled with the employer's vertical preference in the interface shown to workers. The labeling was prominently displayed to make it salient to applying workers.[13] Randomization was effective—see Appendix D.

It is important to note that with our experimental design, workers could simultaneously see and interact with job openings by employers in different cells. As such, the SUTVA condition is inherently—and intentionally—violated. This kind of violation is a typical concern in marketplace experiments (Blake and Coey, 2014). However, we *want* "interference" both in our experiment and in equilibrium, as a goal of the signaling opportunity is to induce workers to sort, by applying to some job openings and not applying to others. This feature of our experimental design does require care when generalizing the results to a market equilibrium.

Employers can and do post multiple job openings, though they are not allowed to have multiple listings for the same position. During the five month experimental period, all subsequent job postings received the same treatment assignment as the original posting, to prevent employer "hunting" for a better cell. This feature of our data can potentially give us more statistical power,

[13] In the ambiguous arm, among those employers shown preferences, employers were further split to have a notice about whether the worker was able to condition upon their signal. The idea motivating this treatment was that employers might infer that bids were more shaded up/down if they knew the worker knew the signal. However, we find no evidence this was the case, and so for simplicity, we pool these observations together, ignoring this feature of the design. As it is, it appeared to have no effect on any outcome it could have affected.

though as experimental group assignment could affect the probability an employer posts a follow-on opening—or the attributes of that opening—we generally restrict our analysis to the first job opening by an employer after the start of the experiment. However, when assessing the effects of the signaling feature on match outcomes, we will use all the job openings to gain more statistical power.

4 Results

4.1 Informative messages

At the end of the theory section we showed that truthful cheap talk can be sustained in equilibrium in some circumstances, but that this is not always possible (see Proposition 4). Here we investigate whether truthful cheap talk seems to be present in this market. We exploit that employers with ShownPref$= 0$ in the explicit arm of the experiment have no strategic incentive since they know that their choice is not shown to potential applicants, and we assume they tell their true preferences. We will see that these employers clearly do make conscious choices as their responses display substantial variation across different types of jobs, in line with intuition. We compare their distribution of choices with those employers where ShownPref$= 1$ in the explicit arm of the experiment, so that a strategic motive is present as they know that their choices will be shown to the other side of the market. If they are not telling the truth, their distribution should be tilted relative to the previous group.

Figure 3 plots the fraction of employers selecting each of the three tiers, by category of work and by whether their choice was to be revealed. For this figure, we only use data from the explicit arm of the experiment. For each fraction, a 95% confidence interval is reported. The number of openings in that category is reported at the top of each facet ("$n = \ldots$").

Figure 3 shows that vertical preferences vary both within and between categories of work. Between categories, if we look only at the ShownPref $= 0$ fractions, we can see that in "Administrative Support," about 59% of employers selected the low tier. In contrast, in "Networking and Information Systems" only about 20% of employers selected the low tier. Although vertical preferences clearly vary between categories, the relationship is far from deterministic; within categories, there is substantial variation, though the medium tier is the most common selection in all categories except for "Administrative Support" and "Customer Service." Because of this within-category variation in tier choice, workers cannot fully learn an employer's vertical preferences simply by knowing the category of work.

When the experiment was designed, it was expected that employers might condition their tier choice on whether their choice would be shown to would-be applicants or just to the platform,

Figure 3: Employer tier choice by category of work in the explicit arm of the experiment, by whether their choice would be shown to would-be applicants

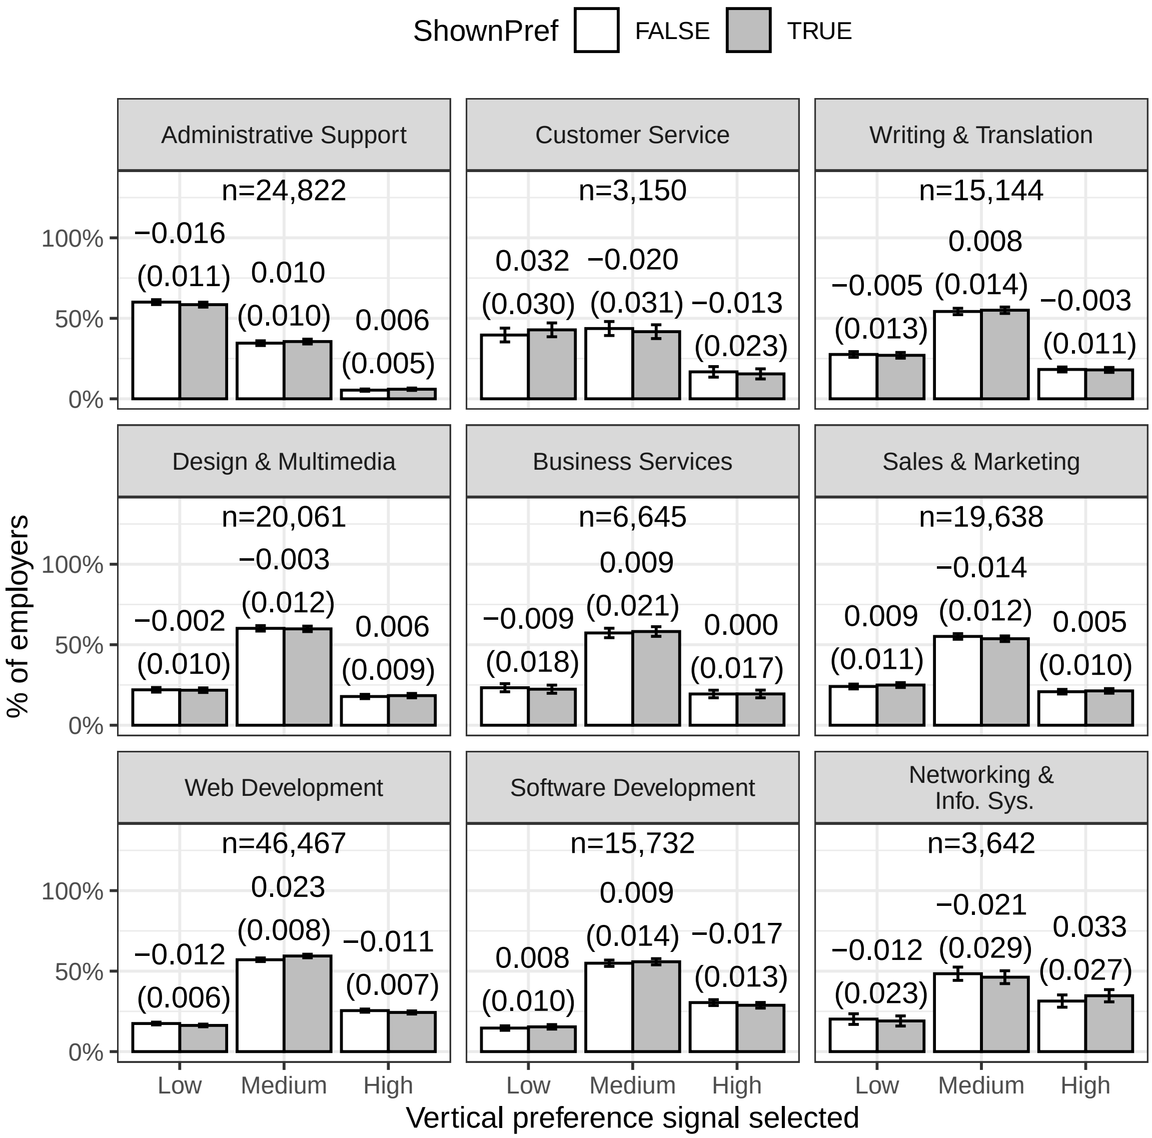

Notes: This figure shows the fraction of employers in the explicit arm of the experiment selecting the various vertical preference tiers, by ShownPref, and by the category of work of the associated job opening. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with the lowest rates."; (2) Intermediate: "I am looking for a mix of experience and value."; (3) Expert: "I am willing to pay higher rates for the most experienced [workers]." We refer to these tiers as "low," "medium," and "high," respectively. If ShownPref = 1, would-be applicants could observe the employer's vertical preference before applying, otherwise they could not, ShownPref = 0. Employers in the two-cell explicit arm were told *ex ante* that the platform would reveal or would not reveal their vertical preferences to workers. A 95% confidence interval is shown for each point estimate. Above each tier fraction in a category of work, the difference between the ShownPref = 1 and ShownPref = 0 fractions is shown, as well as the standard error for the difference.

as outlined before. The design intent of the explicit arm was to test this "endogenous tier" hypothesis.

Despite this possibility of a gap between what employers would message to the platform versus job-seekers, there is no visual evidence in Figure 3 that tier selection depended on revelation: within each category, the fractions choosing the different tiers do not seem to depend on ShownPref. In none of the categories of work is the difference in fractions (shown between bars, with the standard error) conventionally significant, and furthermore, a χ^2-test of ShownPref versus tier selection has a p-value of 0.17.

Despite no evidence of a difference in the fractions of employers picking the various tiers by ShownPref, there could be some hidden compositional shift that leaves the fractions unchanged. As such, we will primarily make use of data from ambiguous arm of the experiment. Despite this caution, the simplest explanation is that employers did not—at least during the experimental period—believe that revelation to workers would be harmful, and so tier choices reflected preferences they were willing to share to both would-be applicants and with the platform. As such, the ambiguous and explicit arms can be safely pooled.

4.2 Sorting

The first implication of Proposition 1 in our theory is that applicant ability should increase at high tier job openings and decrease at low tier job openings when messages are revealed. The same should hold for actual hires. We now empirically investigate how how message revelation affected applicant pool composition, when prior experience is used as a proxy for experience. For this analysis, we use the two cell ambiguous arm of the experiment. Recall that in the ambiguous arm, the tier was chosen *ex ante* by the employer, without knowing whether it would be revealed to job-seekers. As such, differences in the applicant pool composition are causally attributable to revelation.

To measure changes in the applicant pool composition, we estimate the *application level* regression

$$\log y_{ij} = \sum_k \beta_k^s \cdot \textsc{Tier}_{kj} + \epsilon_j \big| s = \textsc{ShownPref}_j, \tag{16}$$

where y_{ij} is some outcome of interest for worker i applying to job opening j, $\textsc{Tier}_{kj}$ is an indicator for whether employer j selected signal tier k and $\textsc{ShownPref}_j$ is an indicator for the treatment status of employer j. We estimate this model separately for $\textsc{ShownPref} = 1$ and $\textsc{ShownPref} = 0$, giving coefficients β_k^1 and β_k^0 for these two samples, respectively. In both regressions, we use weighted least squares, weighting each observation by the inverse of the total number of applicants to the associated job opening. This weighting ensures that all job openings count equally towards the point estimate. We cluster standard errors by the job opening.

The collection of point estimates of β_k^s from (16) are plotted in Figure 4, illustrating the sorting of job-seekers with and without access to the message. The left panel of Figure 4 plots both sets of $\hat{\beta}_k^s$ coefficients from (16) where the outcome is the applicant's log total prior earnings at the time of application. Workers with no experience at the time of application are dropped from the sample. For each of the three tiers, the difference between the two coefficients for the two regressions i.e., $\hat{\beta}_k^1 - \hat{\beta}_k^0$, is labeled, with the standard errors reported below the point estimate.[14] This difference is the effect of message revelation on applicant pool composition.

We can see from the pattern of $\hat{\beta}_k^0$—experience levels by tier when preferences are not revealed—that there is already substantial sorting without the message. High tier employers get more experienced applicants and low tier employers get less experienced applicants, with medium tier employers getting applicants in the middle. This is unsurprising, but note that this is technically not the prediction of the model in Section 2 in which we model workers as approaching firms at random. This simplification of the model follows from our focus on the changes induced by allowing cheap talk messages rather than levels.

If we look at the ShownPref $= 1$ coefficients, $\hat{\beta}_k^1$, we can see that revelation increases sorting in the expected directions according to Proposition 1: revealing the message increases applicant quality—proxied by prior earnings (Pallais, 2014)—at firms in the high tier and decreases it for those selecting the low tier. The effects of revelation are substantial. Revealing the employer's vertical preference raised past average hourly earnings by 7.4% in the high tier and 5.3% in the medium tier. In the low tier, revelation lowered applicant prior earnings by -18.4%.

Proposition 1 establishes that ability should change in the sense of first order stochastic dominance. Figure 4 shows the average effects of message-revelation on applicant ability. This is much weaker than the statement in Proposition 1. If one measures the n-th decile of the ability of applicants at a given opening, information revelation should move ability in the same direction no matter which decile we consider. So for high jobs the median ability (5th decile) should weakly increase when messages are revealed, but also the 10th percentile or the 90th percentile of applicant ability.

Moreover, Corollary 1 predicts that applicant pool differences should be driven either high ability workers (if valuation differences are important) and/or low ability workers (if cost differences are important). Truthful cheap talk as investigated in the previous subsection relied theoretically on large valuation differences, so we expect to see difference for high ability workers, and therefore at high deciles in the ability distribution per vacancy.

[14] The standard error for the difference is calculated directly from the point estimates for the two tiers, without considering the covariance, which should be mechanically zero because of the randomization of ShownPref.

Figure 4: Comparison of applying and hired worker experience by employer vertical preference tier and message revelation in the ambiguous arm

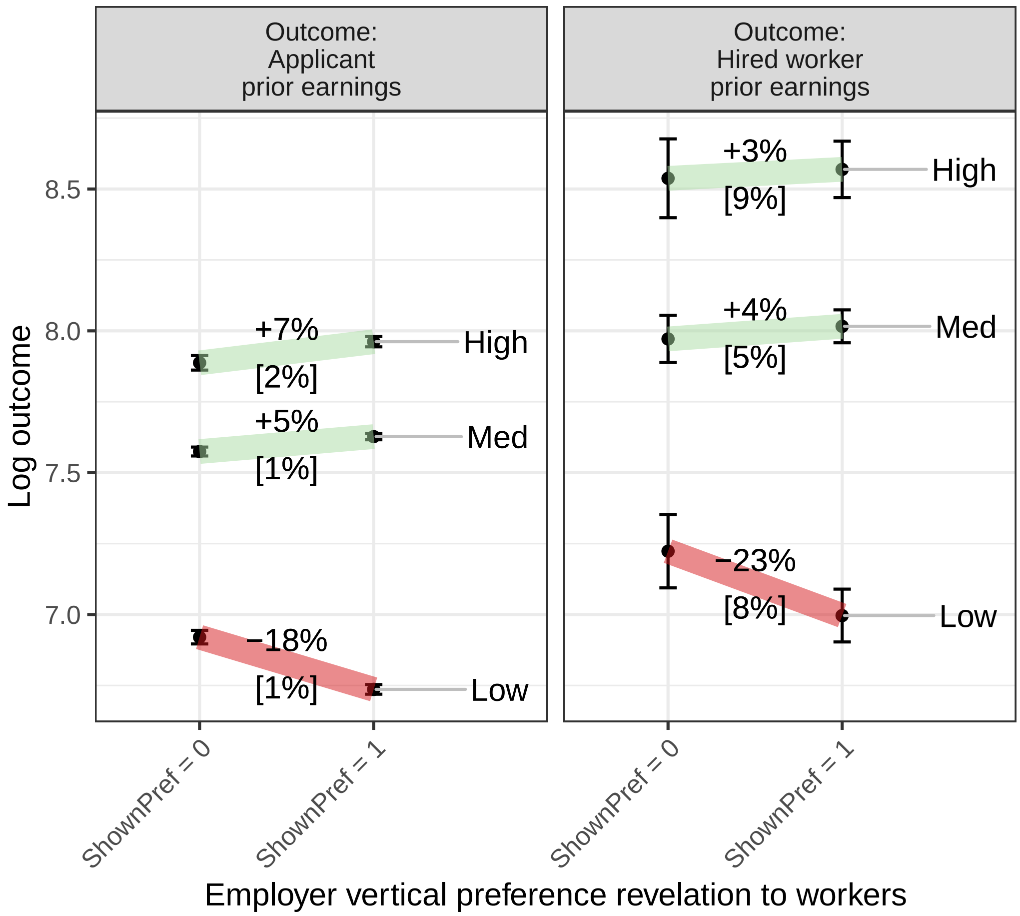

Notes: This figure plots coefficients from estimates of (16). The outcome in both panels is the applicant's cumulative prior hourly earnings at the time of application. The sample in the left panel is all applicants in the ambiguous arm of the experiment, whereas in the right panel, the sample is hired workers in the ambiguous arm. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with the lowest rates."; (2) Intermediate: "I am looking for a mix of experience and value."; (3) Expert: "I am willing to pay higher rates for the most experienced [workers]." We refer to these tiers as "low," "medium," and "high," respectively. Employers in the ambiguous arm were told that the platform *might* reveal their preferences to workers; whether workers were shown employer vertical preferences was randomly determined *ex post*. The error bars indicate the 95% confidence interval for the conditional mean.

In order to investigate these predictions, we can examine sorting based on deciles of the distribution of applicant experience (which is our proxy for ability), by job opening. Note that we are not proposing a quantile regression, but rather comparing the mean of the quantile. As an example of how this measure is constructed, consider three job openings for high message firms where the message was not revealed. Each opening attracts applicants with different abilities (measured by the past earnings of each of the applicants). If the median ability of each opening was $200, $300, $500, then the average median ability is $333. Similarly, we consider the openings of high message firms where the message was shown, and observe median earnings of $250, $300, $550 at each of these. These have a mean of $366. We then estimate treatment effects as simple means comparisons, in this case comparing $333 to $366. We do this in Figure 5, which not only considers the median, but also compares other quantiles. It plots the effects of message revelation on the ability of applicants at different quantiles, where ability is measured by log past earnings (in the top panel) and log hours-worked on the platform (in the bottom panel).

For the earnings measure, we can see that only above the 25th percentile, revelation indeed had worker sorting effects: past-experience was higher in the high tier, about the same in the medium tier, and lower in the low tier. The magnitudes are similar to those estimated from the application-level regressions presented in Figure 4. For the hours-worked measure, we see a roughly similar pattern: separation at higher quantiles, reflecting the changes we expect theoretically when differences in valuations are the main drivers of heterogeneity. At low quantiles for both measures, there is not evidence revelation affected composition, suggesting a limited role of cost differences.

We have studied the application and bidding behavior of workers, but the final hiring decision lies with employers. In the model, changes in the application pool composition are passed through into hires. However, in the experiment, it is possible that employer selection could "undo" any changes in the applicant pool composition.

To test whether revelation affects the characteristics of hired workers, we estimate (16), but with the sample restricted to hired workers. We return to Figure 4, which showed how the composition of the applicant pool changed with revelation, we now examine the right panel. In this panel, labeled "Outcome: Hired worker prior earnings" the outcome is still the job-seeker's log cumulative earnings at the time of application, but the sample is restricted to hired applicants. As before, standard errors are clustered at the level of the job opening. Observations are weighted by the inverse of the number of workers who were hired for that opening, so as to count all job openings equally (though the vast majority of employers hired only one). Note that we are not considering whether revelation affected the probability a match was formed—we will return to this question later, but to preview results, there is no strong evidence of a change in match formation probability.

Figure 5: Effects of showing employer vertical preferences on applicant pool composition with respect to experience, by tier quantile

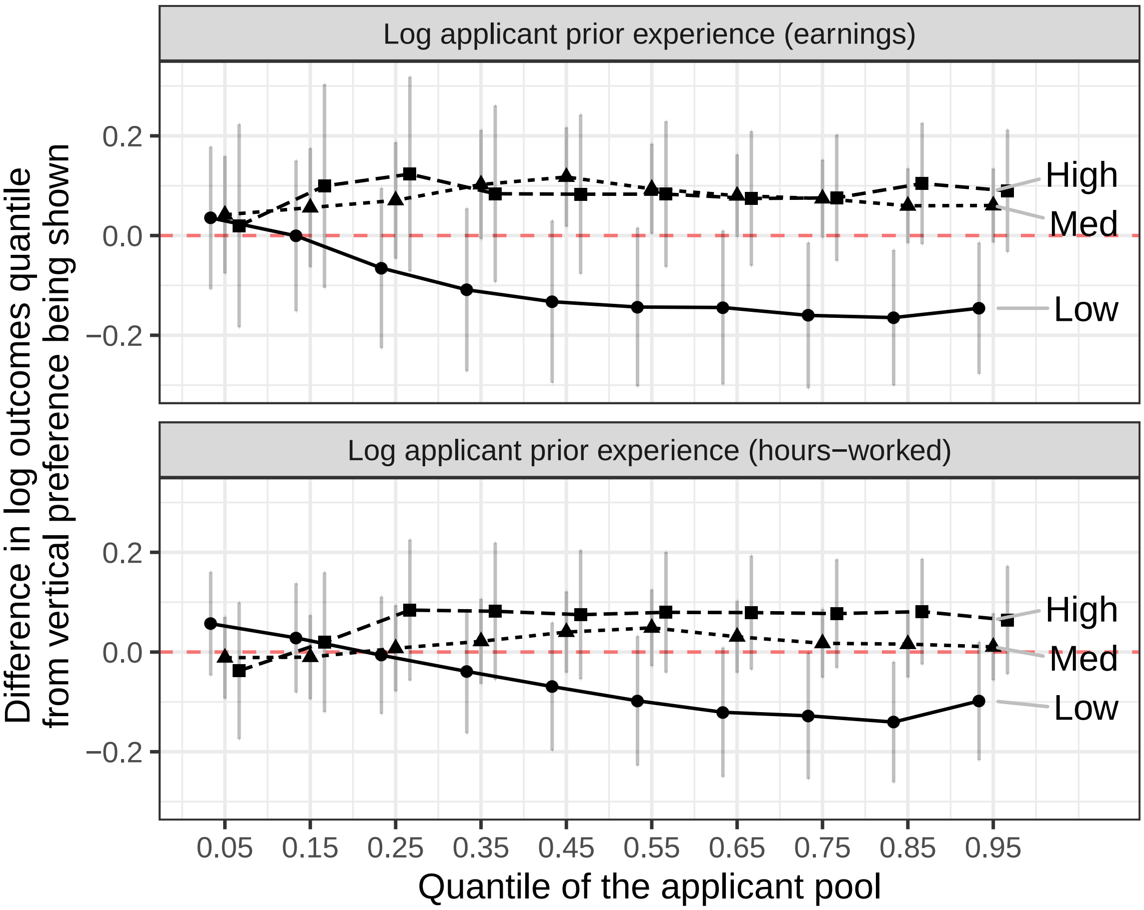

Notes: This figure shows the effects of employer vertical preference revelation on the composition of applicant pools in the ambiguous arm of the experiment for worker experience at the time of their application. The top panel is for cumulative earnings, while the bottom panel is for hours-worked. Workers with no experience are excluded. Each point is the mean effect of revelation of the employer's vertical preference on some applicant attribute at that quantile of the applicant pool. The error bars indicate the 95% confidence interval. Employers in the ambiguous arm were told that the platform *might* reveal their preferences to workers; whether workers were shown employer vertical preferences was randomly determined *ex post*.

In the right panel of Figure 4, we can see that although there is the same separation between the tiers when preferences are not shown, hired workers are—compared to applicants—systematically more experienced. For example, in the high tier, the prior cumulative earnings of applicants is about $\exp(8) \approx \$3,000$. In contrast, for hired workers, prior experience is closer to $\exp(8.6) \approx \$5,400$.

Employers that had their message revealed hired workers that were more like the kinds of workers they stated they were interested in. In the low tier, we can see that signal revelation caused hired workers to have -22.7% lower cumulative prior hourly earnings. In the other tiers, the effects of revelation are positive and broadly similar in magnitude to what was observed for the change in the applicant pool composition but the point estimates are quite imprecise, due to the much smaller samples when restricted to hires.

4.3 Number of job applications per opening, by message

Part 2 of Proposition 1 established theoretically that low type firms experience a decrease in applicant counts if the main difference between firms is their valuations. The opposite effect on low firm types arises if the main differences is differences in the minimum ability (i.e., the cost) needed to produce positive output. Opposite effects are arise for high type firms. To test these predictions we estimate

$$\begin{aligned}\log A_j = \beta_0 &+ \beta_1 \textsc{ShownPref}_j + \beta_2 \textsc{MedTier}_j + \beta_3 \textsc{HighTier}_j + \\ &\beta_4 \left(\textsc{MedTier}_j \times \textsc{ShownPref}_j\right) + \\ &\beta_5 \left(\textsc{HighTier}_j \times \textsc{ShownPref}_j\right) + \epsilon, \end{aligned} \tag{17}$$

where $\textsc{MedTier}_j$ and $\textsc{HighTier}_j$ are indicators for the medium and high tier employers, respectively. The low tier is the omitted category.

In the left panel of Figure 6, the sample is all job openings in the ambiguous arm, whereas in the right panel, the sample is all job openings in the explicit arm. For both arms, the samples are restricted to only those job openings receiving at least one applicant. This restriction removes about 1% of job openings. There is no evidence that the fraction of openings differs by ShownPref status.[15] Within each panel, the error bar (to the far right, above the label "Pooled") shows the group means (i.e., $\hat{\beta}_0$ versus $\hat{\beta}_0 + \hat{\beta}_1$) from (18). Note that the x-axis shows the employer tier selection, ordered from low to high; the colored lines correspond to the ShownPref value.

[15] Job openings sometimes receive no applicants because the employer removes the job post shortly after posting. As this could be affected in principle by the experimental group assignment, we make no attempt to drop these openings from our sample, with the exception of removing them for this specific purpose.

Figure 6: Effect of employer vertical preference revelation on the size of the applicant pool

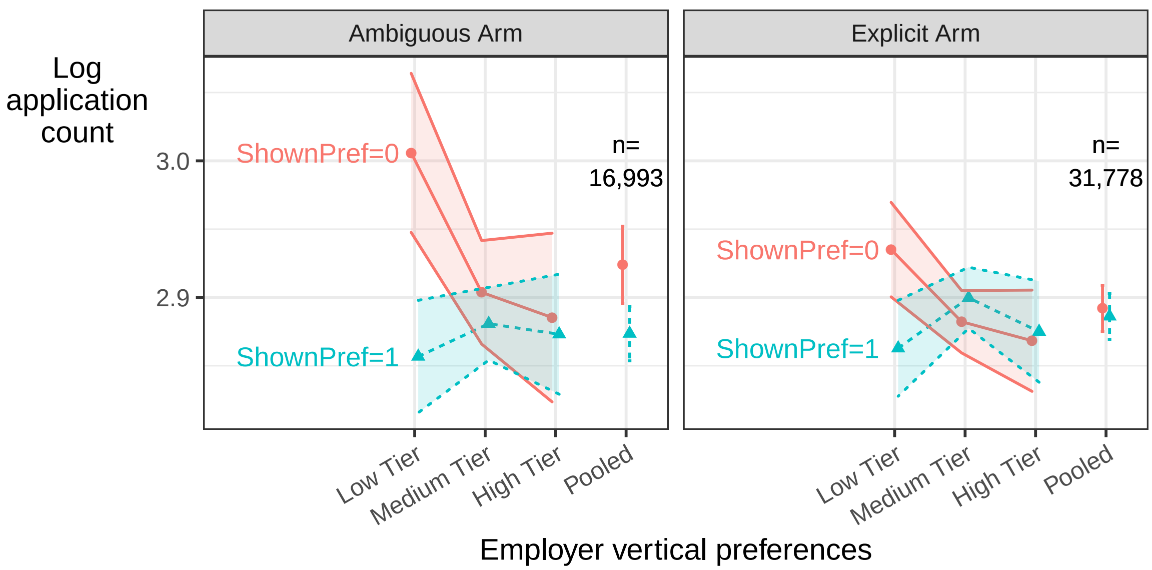

Notes: This figure reports regression results where the outcome is the log number of applications received by that opening. The right panel uses job openings from the explicit arm, whereas the left panel uses openings from the ambiguous arm. The samples are restricted to job openings receiving at least one application. In each panel, the far-right error bars indicate the overall treatment effect, not conditioning by the employer vertical preference tier. The rest of the point estimates in a panel are for the respective tiers. Standard errors are calculated for the conditional means and a 95% CI is shown. Standard errors are robust to heteroscedasticity.

We can see from Figure 6 that in both arms, applicant pool reductions are concentrated in the low tier, with message revelation having little discernible effect in the other tiers. For low types this seems broadly consistent with effects driven by heterogeneous valuations. We do not see opposite effects for high type firms, as we would expect if heterogeneity in valuation were the only difference between firms.

To assess the actual effects of message revelation, we can compare the total number of applicants by treatment assignment without conditioning on the tier. To do this, we can regress the log number of applications per job on the whether the message was revealed:

$$\log A_j = \beta_0 + \beta_1 \textsc{ShownPref}_j + \epsilon, \tag{18}$$

where A_j is the number of applications received by opening j.

We plot $\hat{\beta}_0$ versus $\hat{\beta}_0 + \hat{\beta}_1$ in Figure 6, with 95% CIs, for each arm. In the explicit arm, the point estimates imply that revelation leads to an overall decline of -1.9% in the size of the applicant pool. The ambiguous arm shows larger effects, with an overall decline of -5%. This is also the smaller sample, and so the larger effect could reflect sampling variation. However, the most likely interpretation of the data is that revelation of the message had a modest negative effect on applications sent overall.

4.4 Wages

We now explore how message revelation affected the wage bidding of applicants. Theoretically, part 3 of Proposition 1 establishes that message revelation increases wage bids relative to what workers would usually bid without message revelation at high type firms, and the opposite at low type firms.

Empirically we first estimate (16) but with the log wage bid as the outcome. In the first panel from the left of Figure 7—labeled "Outcome: Applicant wage bid"—we see the same pattern of separation in wage bids that we observed with applicant prior experience, even with $\textsc{ShownPref} = 0$: high type firms get higher bids and low type firms get lower bids. And as before, revelation of the tier intensified the effect: revelation caused wage bids to be 10% higher in the high tier, 4% higher in the medium tier and -13% lower in the low tier.

The empirical analysis of wages so far confounds two effects, though: Revealing messages changes the types of workers that a firm attracts, and the bids that these workers demand. Recall Figure 4): high tier employers who have their message revealed attract and hire more experienced applicants. Since more experienced workers tend to ask for higher pay, this alone could explain the wage effect. But composition does not have to be the only explanation: workers could also directly condition their bid on perceived employer willingness-to-pay. This is

Figure 7: Comparison of applicant mean log wage bids and profile rates by employer vertical preference tier and revelation of the signal

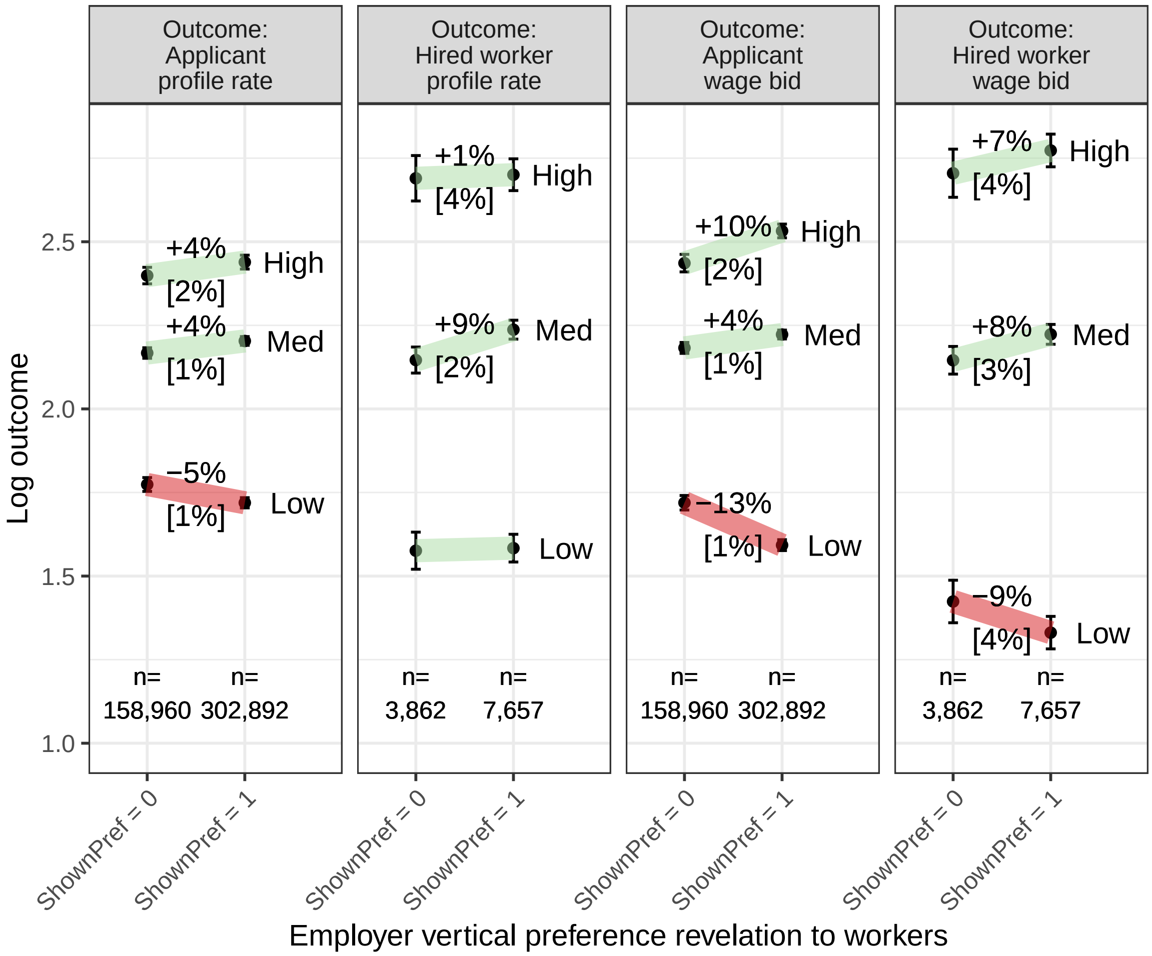

Notes: This figure plots predictions from estimates of (16), using the wage bid and profile rate as the outcomes. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with the lowest rates."; (2) Intermediate: "I am looking for a mix of experience and value."; (3) Expert: "I am willing to pay higher rates for the most experienced [workers]." We refer to these tiers as "low," "medium," and "high," respectively. If ShownPref = 1, would-be applicants could observe the employer's vertical preference before applying, otherwise they could not, ShownPref = 0. The sample is restricted to the ambiguous arm of the experiment. Employers in the ambiguous arm were told that the platform *might* reveal their preferences to workers; whether workers were shown employer vertical preferences was randomly determined *ex post*. The error bars indicate the 95% confidence interval.

the prediction of Proposition 1, bullet 3.

One way to disentangle the two effects on wage bidding—composition and conditional bidding—is to look at changes in the applicant *profile rate* (i.e., the rate declared on their profile) and compare it to changes in the wage bid. The profile rate is not likely to be conditioned on the job opening, whereas the wage bid can be conditioned on the specific features of the job opening, including the employer's tier choice, if available. The profile rate is set by the worker at his or her desired level, but it tends to closely follow a worker's typical hourly wage bid. This correlation is due in part to employers consider the profile rate when recruiting, and so workers have an incentive to keep it "honest."

Using the log profile rate as the outcome in the second panel from the right of Figure 7—labeled "Outcome: Applicant profile rate"—we see the same sorting pattern and revelation effect as we have for all outcomes. However, the message revelation effects are much smaller for the profile rate than they were for the wage bid: revelation raised the profile rates of applicants to high tier openings by 4%, raised them by 4% to medium tier openings, and lowered them by -5% for low tier openings. Note that these low and high tier revelation effects for the wage bid are about twice as large in magnitude compared to the profile rates. Finding smaller effects for profile rates than for wage bids is suggestive that workers are marking up or marking down their wage bids directly in response to the tier choice.[16]

We can directly test for wage bid conditioning by exploiting the fact that workers on the platform apply to multiple job openings. To do this, we estimate the application-level regression

$$\begin{aligned} \log w_{ij} = {} & \alpha_i + \beta_1 \textsc{ShownPref}_j + \beta_2 \textsc{MedTier}_j + \beta_3 \textsc{HighTier}_j + \\ & \beta_4 \left(\textsc{MedTier}_j \times \textsc{ShownPref}_j \right) + \\ & \beta_5 \left(\textsc{HighTier}_j \times \textsc{ShownPref}_j \right) + \epsilon. \end{aligned} \tag{19}$$

where α_i is a worker-specific fixed effect. This "within" estimator allows us to compare the decision-making of workers that applied to job openings with the same tier, but that differed in SHOWNPREF, as well as jobs that differed in their tier.

It is easier to appreciate the interaction of the tier and SHOWNPREF by plotting the mean predicted values from the estimate of (19) when the outcome is the log wage bid, which we do in Figure 8, in the left panel. The sample consists of all applications to job openings in the ambiguous arm of the experiment. We can see that even when workers cannot observe the tier choice, $\textsc{ShownPref} = 0$, they still "pick up" some of the employer's vertical preference, bidding

[16] As a direct measure of wage bid conditioning, we can use as an outcome the "markup" in the wage bid, or the difference between the wage bid and profile rate, divided by the profile rate. See Appendix E.1 for an analysis showing higher markups in response to revelation of a high-type signal and lower markups in response to revelation of a low-type signal.

Figure 8: Worker wage bid, profile rate and experience at time of application, by employer vertical preference and revelation status in the ambiguous arm

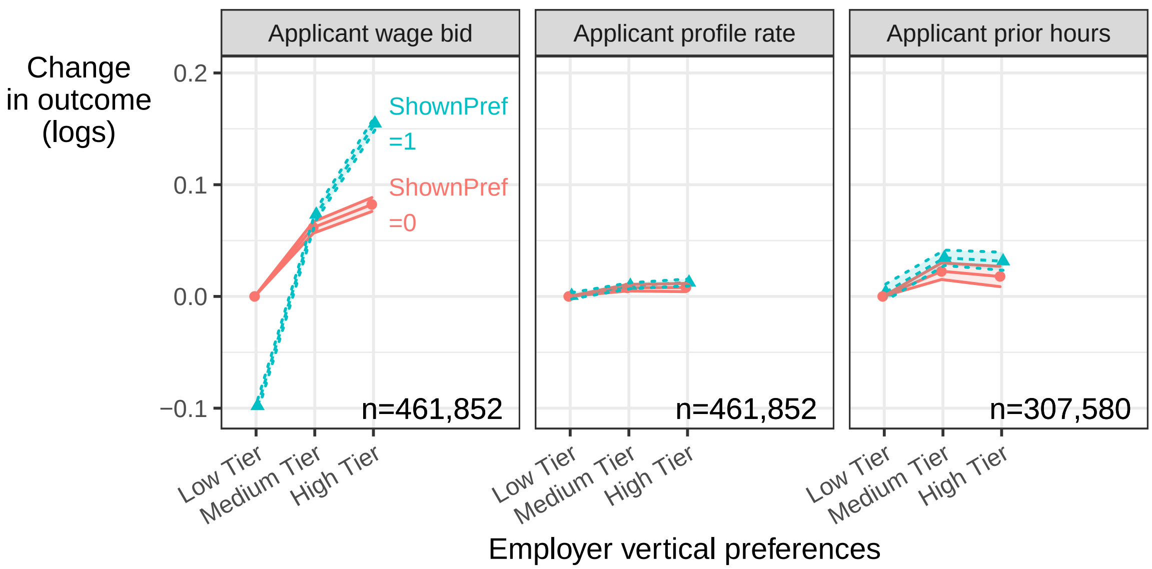

Notes: This figure reports estimates of (19). The sample consists of all applications sent to job openings in the ambiguous arm of the experiment. In each regression, a worker specific fixed effect is included. Standard errors are clustered at the level of the individual worker. The dependent variables are the worker's hourly wage bid, profile rate at time of application and past hours-worked at time of application. Standard errors are calculated for each of these conditional means and a 95% CI is shown. Standard errors are robust to heteroscedasticity.

more when facing a higher tier employer. The coefficient on MedTier implies workers increase their wage bids by 6.2%, and the coefficient on HighTier implies a 8.2% increase in the wage bid. Note again these results are from a regression with a worker-specific fixed effect, and so these changes in bids are not due to changes in composition.

When the message is revealed, workers adjust their wage bids much more strongly. They bid -9.8% less when ShownPref = 1 when they know it is a low tier opening; if the worker learns it is a high tier job opening, they bid an additional 7.3% more, on top of the 8.2% increase noted above.

If our within-worker approach removes worker composition effects, neither the tier nor the revelation of the tier should matter much for outcomes that are quasi-fixed attributes of the applicant. In the middle panel of Figure 8 the outcome is the applicant's log profile rate. In the rightmost panel the outcome is the worker's cumulative hours-worked on the platform at the time of application, if any (note the smaller sample). For both of these quasi-fixed worker attributes, experience and profile rates are slightly increasing in the vertical preference tier, but do not seem to depend strongly on ShownPref. This slight increase over tiers reflects that over the 5 month course of the experiment, workers gain experience and shift their applications to more demanding job openings "organically" and increase their profile rates. However, the effect sizes are only 1/10th of the size of the effects on the wage bid. This implies our within-worker approach effectively nets out composition effects.

The pattern of wage bidding results are consistent with workers directly conditioning on employer willingness to pay. However, there are alternatives explanations. For example, perhaps workers perceive the message as indicating the worker's likely costs. However, we view these alternative explanations as implausible. We show in Appendix E.2, there is no evidence that high tier employers were harsher reviewers when giving feedback even when preferences were not revealed. This suggests that high tier employers did not have costlier expectations for workers that would require a compensating differential. It is also unlikely that some firms are viewed as more attractive. The relatively impersonal nature of these online interactions, along with their short-duration and lack of brand-name firms all make it unlikely workers have strong non-monetary preference over firms *à la* Sorkin (2018). More generally, it is hard to square the sorting and bidding effects with compensating differential argument. For example, workers should submit *lower* bids to employers they preferred to work with, implying that low tier employers are the most desirable, and yet this was precisely the tier that applicants avoided applying to (recall Figure 6).

For the effects of revelation on the wage bids and profile rates of hired workers, we return to Figure 7. From the left, the second and fourth panels have samples restricted to hired workers. Hired worker profile rates were higher in the medium tier and high tier and about the same in the low tier, though again the effects are fairly imprecisely estimated. For the wage bid, we see that revelation raised the wage bids of hired worker in the medium and high tiers, and lowered the wage bid in the low tier by -9%.

4.5 Match outcomes and efficiency

We now examine whether revelation of the employer's tier affected the quantity and characteristics of matches formed. In the theory, the introduction of truthful cheap talk messages improves market efficiency (see Proposition 1, Part 4). Efficiency gains are important *per se*, but they are also important for a matching platform itself, as it can eventually profit from improvements, either by increasing fees or by attracting more business.

A challenge with assessing match quality effects is that we only observe match characteristics, such as hours-worked, if a match is formed. As such, we are inherently selecting samples that could be influenced by treatment assignment. This selection could matter, biasing "downstream" measures.

Selection forces us to be cautious in interpretation, but as we will see, there is no evidence that revelation affected the *quantity* of matches formed. Furthermore, there is no strong evidence that the kinds of job openings that filled differed by ShownPref with respect to pre-treatment attributes. In Appendix E.4, we show that job openings where a match was made had good balance on pre-treatment characteristics by treatment status, consistent with idiosyncratic factors affecting which openings were actually filled.

An additional inferential issue is that slightly less than half of all job openings are filled, and so we have less power than for outcomes that we always observe. To increase statistical power, we pool both the explicit and ambiguous arms. Furthermore, we include not only the first job opening, but all subsequent openings by that employer during the experimental period, adjusting for the hierarchical data that results. This gives us a total sample size of 220,510 jobs openings, of which 73,866 were filled.

Although our preferred estimates for match outcomes are made with the full sample, in Appendix E.6 we report estimates for all the different possible sample combinations (e.g., explicit arm, first openings; ambiguous arm, first openings, all arms, all openings, and so on). The point estimates differ with the sample, but the same general pattern of results is the same as reported when using all arms and all openings.

To measure match outcomes, our regression specification is

$$y_j = \beta_0^k + \beta_1^k \textsc{ShownPref}_j + \epsilon \big| k = \textsc{Tier}_j \tag{20}$$

where $\textsc{Tier}_j$ is the associated tier for the opening. We estimate separate regressions for each tier.

We also estimate the regression with all job openings pooled together, which we label "Pooled." To account for the nested structure of the data, we cluster all standard errors at the level of the employer.

The introduction of the cheap talk signaling opportunity changed many things about the match—the identity of the hired worker, the wage bid, and even the competitive environment. There are various "objective" measures of match quality we could observe, but perhaps one of the more straightforward measures is those that simply asked both sides how they felt. We split our analysis of outcomes into objective and subjective measures.

In looking at objective match outcomes, with our larger and different sample, we recapitulate some of the results from earlier—namely the number of applications and the hired worker wage. But we also add new outcomes, such as whether the job opening was filled. Using the full data, we report estimates of the coefficient on ShownPref from (20) in Figure 9, using as outcomes: (1) the log number of applications, (2) whether any worker was hired and then, selecting only filled openings, (3) the log wage of the hired worker, (4) the log hours-worked of the hired worker, and (5) the log total wage bill. Note that (3) and (5) are based on the actual mean wage over the contract, not the worker's original bid, so it can include raises.

For the filled openings, the sample is all job openings for which hired workers worked at least 15 minutes at a wage greater than 25 cents per hour.[17] If multiple workers were hired for a job opening, we average outcomes. For each estimate, we report the number of observations ("$n = \ldots$") and for the pooled regression, the number of distinct employers ("$g = \ldots$").

In the top panel of Figure 9, we see a reduction in applicant pool sizes from revelation in both the high and low tiers (recapitulating Figure 6, but with a larger sample). In the low tier, the reduction is about -3.3% and in the high tier about -1.6%. There is also a reduction in the medium tier, but it is quite small. As in Figure 6, applicant pool reductions seem to be concentrated in the low tier. However, these effects are smaller than those in the Figure 6.

Despite a reduction in the number of applications, there is no evidence of fewer matches formed, which we can see in the second panel from the top in Figure 9. The point estimates are positive and small, but the associated confidence intervals comfortably include zero.

Although the number of matches did not discernibly change, there are several pieces of evidence that the matches themselves changed. In the third panel from the top of Figure 9, we can see that revelation in the high tier increased hired worker wages by 4.6%, while revelation in the low tier decreases wages by -3.9%, with little effect on the medium tier. The net overall effect, indicated by "Pooled," is slightly negative. However, this does not necessarily imply workers

[17] We make this restriction on hours-worked and wages because a small number of employers (against the platform's wishes) create very low wage contracts to simply use the hours-tracking feature but not process payments through the platform.

were made worse-off, as we saw that substantially less experienced workers hired in the low tier (recall Figure 5). In Appendix E.5, we show that on a per-application basis and with worker specific fixed effects, workers had higher application success probabilities and higher expected values (wage bids times success probability), when applying to ShownPref $= 1$ job openings.

In the bottom two panels, we can see that revelation led to more hours-worked and a larger wage bill. Pooled across tiers, revelation increased hours-worked by 4.6%, with increases of 2.9% in the high tier and 5% in the low tier. Revelation increased the wage bill by 2.7%. Analogous to worker-employer tenure being a measure of match quality in conventional markets, these increases in quantities are suggestive of better matches being formed with revelation.

If buyers are less satisfied, they might leave worse feedback for the worker or the platform. The effects of revelation on these feedback measures is reported in Figure 10. All feedback outcomes are transformed into z-scores, and so point estimates are interpretable as fractions of a standard deviation. The top panel is the employer's feedback to the hired worker, the middle panel is the worker's feedback to the hiring employer, and the bottom panel is the feedback of the employer to the platform (framed as a probability of recommending the platform to someone else).

For the worker-on-employer and employer-on-worker feedback, parties are prompted to give feedback after the conclusion of a contract but are not obligated to, hence the sample of contracts for feedback is smaller than the number of contracts. For the platform feedback, employers are randomly sampled and asked for feedback about 1/3 of the time, explaining why this sample is considerably smaller.

From Figure 10, we can see that there is little change in the feedback to the worker. For the feedback to the employer, there is some evidence of better feedback to high tier employers and lower feedback to low tier employers who had their preference revealed. This would be consistent with worker feedback increasing in the hourly wage received, perhaps due to feeling grateful to the employer for the higher wage (Akerlof, 1982).[18] Despite somewhat lower feedback to workers, the platform itself got slightly higher marks from employers—effects were higher in all tiers, with an overall effect that is about 0.025 standard deviations, though the estimates are not very precise.

4.6 Separating equilibrium in the long-run

A limitation of our experimental design is that it does not directly shed light on the repeated choices of firms over time, nor on the full market equilibrium when all firms are treated. Our

[18] See Luca and Reshef (2021) for an interesting paper on how plausibly exogenous changes in the price of a good affect subsequent buyer ratings.

Figure 9: Effects of revealing employer vertical preferences on job opening outcomes

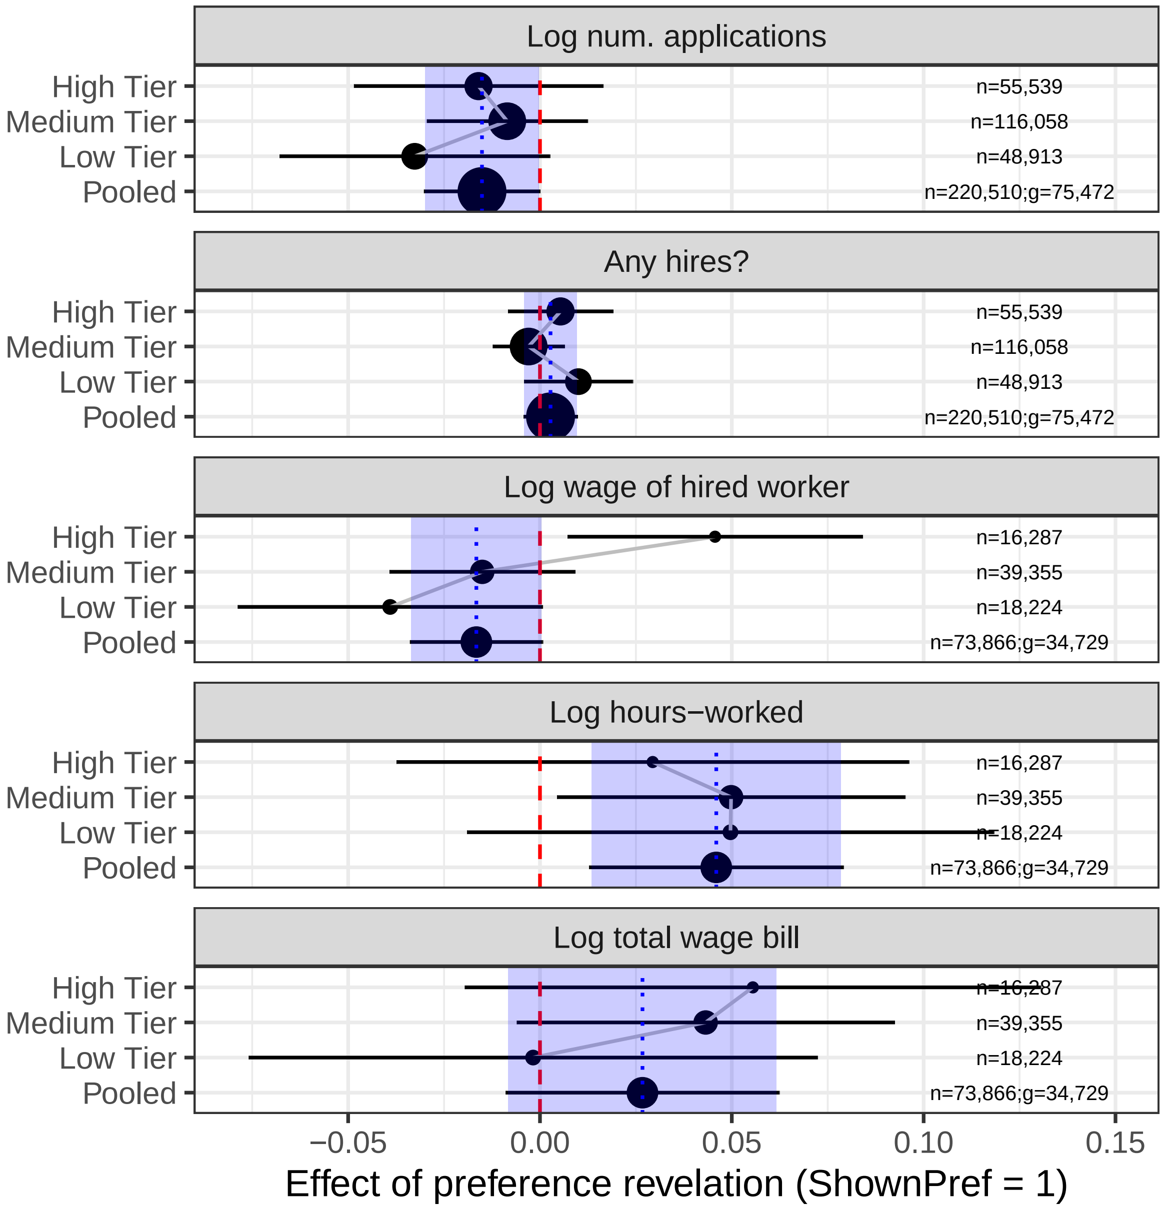

Notes: This figure shows the effects of revealing employer preferences, ShownPref = 1, on a number of outcomes. The sample consists of all job openings from both the ambiguous and explicit arms. Each point estimate is surrounded by at a 95% CI.

Figure 10: Effects of revealing employer vertical preferences on job opening feedback scores (z-scores)

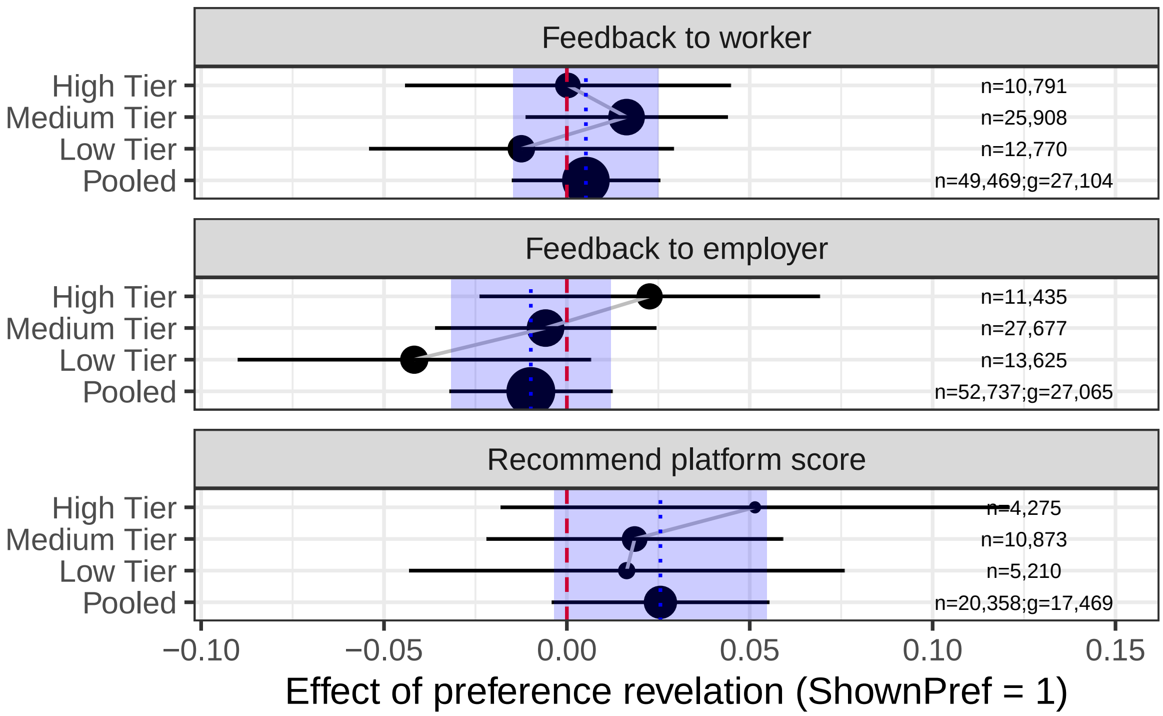

Notes: This figure shows the effects of revealing employer vertical preferences on various feedback measures. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with the lowest rates."; (2) Intermediate: "I am looking for a mix of experience and value."; (3) Expert: "I am willing to pay higher rates for the most experienced [workers]." We refer to these tiers as "low," "medium," and "high," respectively. The sample consists of job openings from both the ambiguous and explicit arms. Each point estimate is surrounded by at a 95% CI. Point sizes are scaled by the sample size.

variation comes from a setting where some firms get treated and others do not. If the block-recursive nature of our theory is valid, the insights from the within-market setting should carry over to a treatment of the full market, as the fraction of treated individuals did not matter. A particular implication in Part 5 of Proposition 1 is that the incentives to send particular messages should not be affected by the amount of firms that are treated in the market.

At the conclusion of the experiment, message revelation was rolled out to the entire market: all employers received an experience identical to the ShownPref $= 1$ cell in the explicit arm, meaning that all employers now knew their preferences would be revealed (and they were revealed). There are two empirical approaches that allow us to investigate whether a separating equilibrium persisted: we can (1) look at the fraction of job openings selecting the various tiers in the post period, and (2) look at trends within employer in the tier choice.

To investigate the evolution of choices into the post-experimental period, Figure 11 shows the fraction of employers choosing the various tiers over time, with the end of the experiment indicated. The figure shows that the share of high-tier choices remains remarkable constant over time. It does not seem to be affected by the fact that only parts of the market could show their messages during the experiment compared to the roll-out over the whole market after the experiment, in line with Part 5 of Proposition 1. There is a bit more volatility in the shares of medium and low-tier messages, with perhaps some evidence of an immediate post roll-out increase in low tier selections, but this does not persist and the long-run pattern seems to be one of relatively stable shares for each of the tiers.

We also looked at firms that posted jobs repeatedly during the experimental period. Recall that in the previous sections we only used the initial job posted by each of these firms. Our theory presumed that firms understand all features of the environment right away. It is conceivable that firms only slowly understand the strategic aspects of their environment, and even though they are truthful in their initial message they become more strategic over time, which could lead to pooling on particularly beneficial messages over time. In Appendix E.7 we document that reporting is very stable over time: there is a slight trend away from low and medium-tier choices towards high-tier choices as firms repeatedly post jobs, but even after the tenth job posting firms are only 2%pt more likely to choose the high tear. Importantly, this is not affected by whether the preference was shown or hidden, which is another indication that firms tell the truth even when the message is strategically relevant (in line with the evidence of Section 4.1).

So overall, message frequencies stay remarkably constant even as individual employers post jobs repeatedly, but also for the whole market over time and across a large roll-out to all firms.

Figure 11: Employer vertical preference signal choice over time, both before and after the experiment

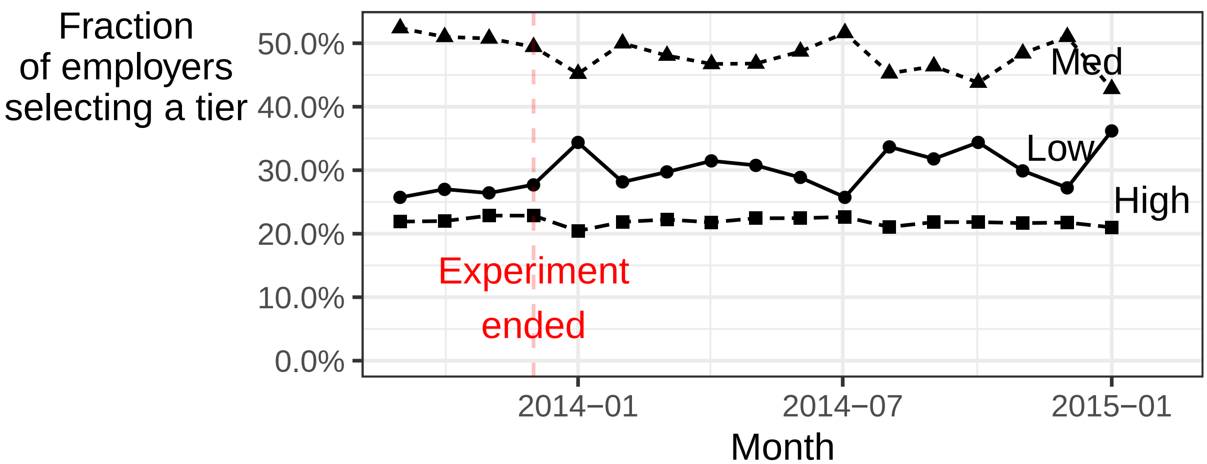

Notes: This figure shows the fraction of job openings each month selecting each of the three possible tiers. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with the lowest rates."; (2) Intermediate: "I am looking for a mix of experience and value."; (3) Expert: "I am willing to pay higher rates for the most experienced [workers]." We refer to these tiers as "low," "medium," and "high," respectively. The vertical red line indicates when the experiment ended and all employers were asked for their preferences. After the experiment, these preferences were always shown to applicants, and employers knew upfront that their signal choices would be revealed.

5 Conclusion

Platform-engineered cheap-talk opportunities can move designed markets to more desirable equilibria, both in theory and in practice. In our setting, match efficiency was improved and the quantity transacted in the market increased via a platform intervention that had essentially zero marginal cost. Given the platform's pricing structure of applying an *ad valorem* charge, the market intervention raised platform revenue by nearly 3% if the experimental estimates generalized. Despite this positive result, there are several open questions, such as whether the method could be applied to other preference dimensions and other market settings.

A feature of the cheap-talk opportunity described here is that workers were able to apply cross-tier. It would be straight-forward to design a version of the cheap-talk opportunity in which workers would have to choose a tier and only apply within a tier for some period of time. The tier selection could also be made centrally by the platform, using prior experience or feedback to create cut-off scores rather than allowing workers to self-select, which would turn the message game away from a cheap-talk game into a signalling game. This could lead to more sorting and more "refined" pools, but at the cost of greater intervention by the platform and the greater chance of leaving jobs under- or over-filled if supply is not managed.

In addition to determining which workers are allowed in which tier, another possible direction could be for the platform to define what different tiers "mean," such as by labeling them with experience requirements. Theory provides little guidance on the actual lables that are used in cheap-talk, but in reality the ability to easily identify with particular labels might alter the ability of the market to coordinate on different cheap-talk messages. Labels with experience requirements might get more informed separation, though it also increases the burden on the platform in deciding what are reasonable tier labels.

Although our context is an online labor market, the matching process in this market mirrors that found in conventional markets. The cheap-talk opportunity in this paper is with respect to vertical preference, but there are other potential pieces of information that might be conveyed by a cheap-talk mechanism. For example, if employers could choose to describe their project as "urgent" could we get a similar sorting equilibrium? Employers could also signal information about their management "style" (e.g., closely managed or hands-off), their degree of confidence in what works needs to be done, the degree of contract completeness, and so on. Essentially any feature of the economic relationship for which buyers and sellers have heterogeneous preferences or attributes and have imperfectly aligned incentives is a potential candidate for a cheap talk intervention.

References

Acemoglu, Daron and Robert Shimer, "Holdups and Efficiency with Search Frictions," *International Economic Review*, 1999, *40* (4), 827–849.

Ahn, Hie Joo, Bart Hobijn, and Ayşegül Şahin, "The Dual US Labor Market Uncovered," Technical Report, National Bureau of Economic Research 2023.

Akerlof, George A, "Labor contracts as partial gift exchange," *The Quarterly Journal of Economics*, 1982, *97* (4), 543–569.

Albrecht, James, Pieter A. Gautier, and Susan Vroman, "Directed search in the housing market," *Review of Economic Dynamics*, 2016, *19*, 218–231.

Bagger, Jesper and Rasmus Lentz, "An empirical model of wage dispersion with sorting," *The Review of Economic Studies*, 2019, *86* (1), 153–190.

Bajari, Patrick and Ali Hortaçsu, "The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions," *The RAND Journal of Economics*, 2003, *34* (2), 329–355.

Barach, Moshe A. and John Horton, "How do employers use compensation history? Evidence from a field experiment," *Journal of Labor Economics*, 2021, *39* (1).

Belot, Michèle, Philipp Kircher, and Paul Muller, "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice," *The Review of Economic Studies*, 2018, *86* (4), 1411–1447.

Bhole, Monica, Andrey Fradkin, and John Horton, "Information About Vacancy Competition Redirects Job Search," *Working paper*, 2021.

Blake, Thomas and Dominic Coey, "Why marketplace experimentation is harder than it seems: The role of test-control interference," in "Proceedings of the fifteenth ACM conference on Economics and computation" ACM 2014, pp. 567–582.

Burdett, Kenneth, Shouyong Shi, and Randall Wright, "Pricing and Matching with Frictions," *Journal of Political Economy*, 2001, *109* (5), 1060–1085.

Cai, Xiaoming, Pieter Gautier, and Ronald P Wolthoff, "Search, screening and sorting," 2021.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler, "Labor supply of New York City cabdrivers: One day at a time," *The Quarterly Journal of Economics*, 1997, *112* (2), 407–441.

Coles, Peter, Alexey Kushnir, and Muriel Niederle, "Preference signaling in matching markets," *American Economic Journal: Microeconomics*, 2013, *5* (2), 99–134.

Crawford, Vincent P and Joel Sobel, "Strategic information transmission," *Econometrica*, 1982, *50* (6), 1431–1451.

Davis, Steven J. and Brenda Samaniego de la Parra, "Application flows," *Working paper*, 2021.

Diamond, Rebecca and Enrico Moretti, "The geography of consumption inequality," *Working paper*, 2018.

Eeckhout, Jan and Philipp Kircher, "Sorting and decentralized price competition," *Econometrica*, 2010, *78* (2), 539–574.

_ and _ , "Identifying sorting—in theory," *The Review of Economic Studies*, 2011, *78* (3), 872–906.

Filippas, Apostolos, John Horton, and Joseph Golden, "Reputation inflation," *Working paper*, 2018.

Gautier, Pieter A and Coen N Teulings, "How large are search frictions?," *Journal of the European Economic Association*, 2006, *4* (6), 1193–1225.

Gee, Laura K, "The more you know: Information effects on job application rates in a large field experiment," *Management Science*, 2019, *65* (5), 2077–2094.

Gilchrist, Duncan S, Michael Luca, and Deepak Malhotra, "When 3 + 1 > 4: Gift structure and reciprocity in the field," *Management Science*, 2016, *62* (9), 2639–2650.

Gneezy, Uri and John A List, "Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments," *Econometrica*, 2006, *74* (5), 1365–1384.

Gregory, Victoria, Guido Menzio, and David G Wiczer, "The alpha beta gamma of the labor market," Technical Report, National Bureau of Economic Research 2021.

Hall, Robert E and Marianna Kudlyak, "The unemployed with jobs and without jobs," *Labour Economics*, 2022, *79*, 102244.

Harrison, Glenn W. and John A. List, "Field Experiments," *Journal of Economic Literature*, December 2004, *42* (4), 1009–1055.

Hoppe, Heidrun C, Benny Moldovanu, and Aner Sela, "The theory of assortative matching based on costly signals," *The Review of Economic Studies*, 2009, *76* (1), 253–281.

Horton, John J., "Online labor markets," *Internet and Network Economics: 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings*, 2010.

Horton, John J, "The effects of algorithmic labor market recommendations: Evidence from a field experiment," *Journal of Labor Economics*, 2017, *35* (2), 345–385.

_ , "Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution," *Management Science*, 2019.

_ **and Shoshana Vasserman**, "Job-Seekers Send Too Many Applications," *Working paper*, 2020.

Jin, Ginger Zhe and Phillip Leslie, "The effect of information on product quality: Evidence from restaurant hygiene grade cards," *The Quarterly Journal of Economics*, 2003, *118* (2), pp. 409–451.

Kim, Kyungmin, "Endogenous market segmentation for lemons," *The RAND Journal of Economics*, 2012, *43* (3), 562–576.

_ **and Philipp Kircher**, "Efficient Competition Through Cheap Talk: The Case of Competing Auctions," *Econometrica*, 2015, *83* (5), 1849–1875.

Kuhn, Peter and Hani Mansour, "Is Internet job search still ineffective?," *The Economic Journal*, 2014, *124* (581), 1213–1233.

_ **and Mikal Skuterud**, "Internet job search and unemployment durations," *The American Economic Review*, 2004, *94* (1), 218–232.

Kushnir, Alexey, "Harmful signaling in matching markets," *Games and Economic Behavior*, 2013, *80*, 209–218.

Lee, Soohyung and Muriel Niederle, "Propose with a rose? Signaling in internet dating markets," *Experimental Economics*, 2015, *18* (4), 731–755.

Lewis, Gregory, "Asymmetric Information, adverse selection and online disclosure: The case of eBay motors," *The American Economic Review*, 2011, *101* (4), 1535–1546.

Lowe, Airielle, "US Freelancing Surges to Record as 60 Million Americans Go Solo," 2022.

Luca, Michael, "Reviews, reputation, and revenue: The case of Yelp. com," 2016.

_ **and Oren Reshef**, "The Effect of Price on Firm Reputation," *Management Science*, 2021.

Marinescu, Ioana and Ronald Wolthoff, "Opening the black box of the matching function: The power of words," *Working paper*, 2016.

Mas-Colell, Andreu, "On a theorem of Schmeidler," *Journal of Mathematical Economics*, 1984, *13* (3), 201–206.

Menzio, Guido, "A theory of partially directed search," *Journal of political Economy*, 2007, *115* (5), 748–769.

Moen, Espen R, "Competitive search equilibrium," *Journal of political Economy*, 1997, *105* (2), 385–411.

Pallais, Amanda, "Inefficient hiring in entry-level labor markets," *The American Economic Review*, 2014, *104* (11), 3565–3599.

Peters, Michael, "Ex Ante Price Offers in Matching Games Non-Steady States," *Econometrica*, 1991, *59* (5), 1425.

_ , "A Competitive Distribution of Auctions," *The Review of Economic Studies*, 1997, *64* (1), 97.

_ , "Noncontractible heterogeneity in directed search," *Econometrica*, 2010, *78* (4), 1173–1200.

Roth, Alvin E, "Marketplaces, markets, and market design," *The American Economic Review*, 2018, *108* (7), 1609–1658.

Shi, Shouyong, "Frictional Assignment. I. Efficiency," *Journal of Economic Theory*, 2001, *98* (2), 232–260.

_ , "A Directed Search Model of Inequality with Heterogeneous Skills and Skill-Biased Technology," *The Review of Economic Studies*, 2002, *69* (2), 467–491.

Shimer, Robert, "The Assignment of Workers to Jobs in an Economy with Coordination Frictions," *Journal of Political Economy*, 2005, *113* (5), 996–1025.

_ **and Lones Smith**, "Assortative matching and search," *Econometrica*, 2000, *68* (2), 343–369.

Sorkin, Isaac, "Ranking firms using revealed preference," *The Quarterly Journal of Economics*, 2018, *133* (3), 1331–1393.

Stanton, Christopher T and Catherine Thomas, "Landing the first job: The value of intermediaries in online hiring," *The Review of Economic Studies*, 2015, *83* (2), 810–854.

Tadelis, Steven and Florian Zettelmeyer, "Information disclosure as a matching mechanism: Theory and evidence from a field experiment," *The American Economic Review*, February 2015, *105* (2), 886–905.

Teulings, Coen N and Pieter A Gautier, "The right man for the job," *The Review of Economic Studies*, 2004, *71* (2), 553–580.

Varian, Hal R, "Computer mediated transactions," *The American Economic Review*, 2010, *100* (2), 1–10.

Wright, Randall, Philipp Kircher, Benoît Julien, and Veronica Guerrieri, "Directed Search and Competitive Search Equilibrium: A Guided Tour," *Journal of Economic Literature*, 2021, *59* (1), 90–148.

Appendix to

"Sorting through Cheap Talk: Theory and Evidence from a Labor Market" by John Horton, Ramesh Johari and Philipp Kircher

(for Online Publication)

A Derivations and Proofs for Subsections 2.1 and 2.3

Payoff for workers in the presence of mass points: Consider a worker of type a who considers a job with observation (m, s_1) and intends to bid b_L (resp., b_H) after observing low (resp., high) second signal. The generalization of this worker's expected payoff (3) in the presence of mass points is:

$$V(m, s_1, a, b_L, b_H) \tag{21}$$
$$= \sum_{s_2, T} \Psi(s_2, T|m, s_1) b_{s_2} \left(e^{-\Gamma_+(\pi(a, b_{s_2}|T)|m, s_1, s_2, T)}\right) [\pi(a, b_{s_2})|T) \geq 0] M(\triangle(\pi(a, b|T)|m, s_1, s_2, T)).$$

The first three factors inside the summation are unchanged with respect to (3). The forth factor comprises an indicator function that takes the value of one if the argument is true and zero otherwise, capturing the fact that the worker only gets hires and generates non-zero expected payoff if his bid generates non-negative profits (otherwise the worker would not be hired). In the presence of mass points the worker not only has to worry about other workers who generate strictly higher profits for the firm, but also about workers who generate exactly the same profits. If such workers are present, then one of them gets selected at random. The final factor $M(\triangle)$ captures the probability that this worker is hired rather than any worker generating exactly the same profit. It is strictly below 1 if there is a mass of competitors ($\triangle > 0$) and is exactly 1 otherwise, and is constructed as follows.
The probability $M(\triangle)$ consists of several elements: probability $e^{-\triangle}$ that no worker with equal profit is present, the probability $\triangle e^{-\triangle}\frac{1}{2}$ that exactly one other worker with equal profit is present, but the other worker is not chosen, the probability $\frac{\triangle^2 e^{-\triangle}}{2}\frac{1}{3}$ that two other workers with equal profit is present, but the other workers are not chosen, and so forth. The probability of being hired (conditional on no strictly more profitable bidders) is then the infinite sum $M(\triangle) = \sum_{k=0}^{\infty} \frac{\triangle^k e^{-\triangle}}{k!} \frac{1}{k+1} = \frac{1-e^{-\triangle}}{\triangle}$. See Burdett et al. (2001) and Wright et al. (2021) for further references.

Proof of Lemma: Assume to the contrary that $\triangle(\pi|m, s_1, s_2, T) > 0$ at a point where $\Psi(s_2, T|m, s_1) > 0$ and $\pi \geq 0$. Then there must exist (a, b_L, b_H) in the support of $\Lambda(\cdot|m, s_1)$ with $\pi(a, b_{s_2}|T) = \pi$. In fact, there must exist such combination with $b_{s_2} > 0:$ Clearly $b_{s_2} \geq 0$ as otherwise the worker would make a loss when hired and would improve his utility by raising the wage demand at least to zero. It is not possible to have a mass of workers offering the same π and all offering $b_{s_2} = 0$, as in this case only a single worker type solves (1), and a single worker type does not create a mass point.

Now that we know that $b_{s_2} > 0$, it is easy to see that this is not optimal for the worker to demand this wage. If the worker chooses this demand, the fact that $\triangle(\pi) > 0$ implies that $M(\triangle(\pi)) < 1$, where for ease of readability we suppressed that this all conditions on (m, s_1, s_2, T). Alternatively, we can find for any $\varepsilon > 0$ a bid $b'_{s_2} \in (b_{s_2} - \varepsilon, b_{s_2})$ slightly below the original one with associated profit $\pi' > \pi$ such that $\triangle(\pi') = 0$. This is due to the fact that only a countable number of points can have a positive mass. This implies $M(\triangle(\pi')) = 1$, and the hiring probability captured by the third term in (3) jumps up discretely relative to the original wage demand. The second term in (3) weakly increases because $\pi' > \pi$, and the first and fifth term do not change. So the only negative comes from the forth term that captures the wage conditional on getting hired. Since b'_{s_2} can be made arbitrarily close to b_{s_2}, the loss in the wage can be made arbitrarily small and, thus, less than the gain from the discrete jump in the hiring probability. $\square$

Subsection 2.3 entailed the following claim that we prove here in the appendix: Consider firms that can send message $m \in \{L, H\}$ in a truthful equilibrium. At a given message m the worker abilities that enter in equilibrium form a non-empty interval $[\underline{a}_m, \bar{a}_m]$, and no single worker type enters with positive mass.

Proof: We will first show that there is a positive mass of workers that enter and approach message m. Assume not, and consider worker type a with $v_T(a - c_T) - u(a) > 0$. This worker could enter, approach firms with message m, and bid $b = v_m(a - c_m)$. Since he would be hired for sure as no better workers are present at any firm with such message, he obtains $V(m, a, b) > u(a)$, contradicting the third equilibrium condition.

Next, we rule out that only a single type a with mass $\delta > 0$ approaches firms with messages m. Assume this were true. Since we have proven that the distribution of profits that is offered has no mass point, they have to offer different profits in equilibrium. Let π_l be the infimum and π_h the supremum of such profits. We know by free entry that $u(a) = e^{-\Gamma(\pi_l|m)}\left[v_m(a - c_m) - \pi_l\right] = e^{-\Gamma(\pi_h|m)}\left[v_m(a - c_m) - \pi_h\right] =: U(a|m)$, where the last term denotes the utility conditional on entering and approaching firms with message m. Let $U_{-}(a|m)$ be the left derivative and $U_{-}(a|m)$ be the right derivative of $U(a|m)$ with respect to a. We can apply the envelop theorem to obtain:

$$
\begin{aligned}
U'_{-}(a|m) &= e^{-\Gamma(\pi_l|m)} v_m \\
U'_{+}(a|m) &= e^{-\Gamma(\pi_h|m)} v_m
\end{aligned}
$$

We also know that $u'(a) \le U'_{-}(a|m)$, since types a' below a have $u(a') \ge U(a'|m)$ as required by the third equilibrium condition and in order to reach $u(a) = U(a|m)$ the utility after entering has to rise at least as fast as the utility of the outside option. But since $U'_{-}(a|m) < U'_{+}(a|m)$, we obtain $U(a''|m) > u(a'')$ for types a'' that are larger but close to a. This yields a contradiction to the third equilibrium condition.

Finally we will show that the set of abilities that approach message m is an interval. Assume not. Then there exist a_l and $a_h > a_l$ that approach firms with message m but no firms inbetween (i.e., a_l and a_h are in the support of $\Lambda(a|m)$ but no $a \in (a_l, a_h)$ is in the support). Let π_h be the (supremum of) profits offered by type a_l, and π_h the (infimum of) profits offered by a_h to firms with message m. Since profits are rising in ability, we have $\pi_l \le \pi_h$. Free entry implies $u(a_l) = e^{-\Gamma(\pi_l|m)}\left[v_m(a_l - c_m) - \pi_l\right]$ and $u(a_h) = e^{-\Gamma(\pi_h|m)}\left[v_m(a_h - c_m) - \pi_h\right]$. Free entry also means for types $a \in (a_l, a_h)$ that their utility from entering and approaching message m needs to be weakly lower than the outside option: $U(a'|m) \leq u(a)$. Together with the previous equalities this requires $U'(a_l|m) = e^{-\Gamma(\pi_l|m)}v_m \leq u'(a_l)$ and $U'(a_h|m) = e^{-\Gamma(\pi_h|m)}v_m \geq u'(a_h)$. Since the outside option is strictly concave, this implies $e^{-\Gamma(\pi_h|m)} > e^{-\Gamma(\pi_l|m)}$. But since the distribution of profits has no mass points and there are no bids between π_h and π_l, it holds that $\Gamma(\pi_h|m) = \Gamma(\pi_l|m)$, yielding a contradiction. $\square$

Proof of Proposition 3: Consider firms that cannot send a cheap-talk message. Given that we assumed precise second signals, we can treat s_2 and the true type T interchangeably, and for easy of exposition we will only use the type T here. Since we only have one other signal left, we can simply write s instead of s_1. As long as we consistently refer to the same signal, we might also omit the signal from the notation. Now consider workers who see signal s and believe that the fraction of low (high) firms at this signal is $\Psi(L|s)$ ($\Psi(H|s)$) where we suppressed the message because it is absent ($m = \emptyset$). Note that these fractions are exogeneously given; for example, $\Psi(L|L) = \psi_1\delta_L/(\psi_1\delta_L + (1-\psi_1)\delta_H)$ and does not depend on the equilibrium interaction.

Similar to (5) and (6) we can now write the utility of a worker with ability a as

$$W(a|s) = \max_{\pi_L,\pi_H} \sum_T \Psi(T|s)e^{-\Gamma(\pi_T|T,s)}\left[v_T(a - c_T) - \pi_T\right]\mathbf{1}[a \geq c_T], \tag{22}$$

where π_L and π_H are the profits that the worker offers to the low and high type firm respectively, and $\Gamma(\pi_T|T,s)$ denotes the queue of workers who offers higher profits to these firms at this message. Workers can only submit profit bids larger than zero if their type exceeds the cost threshold, and therefore $\mathbf{1}[\cdot]$ is an indicator that takes the value of one if the ability is higher than the work requirement, and zero otherwise.

Due to the same supermodularity between a and π_T, higher ability workers leave higher profits. Again there cannot be any mass points nor holes in the distribution of offers $\Gamma(\pi_T|T,s)$. We focus for now on workers with ability larger then requirement c_H, and will later show that this is the relevant range. This allows us to drop the indicator function in (22). For such workers, utility increases in ability according to

$$W'(a|s) = e^{-\Lambda(a|s)}\hat{v},$$

where $\Lambda(a|s)$ is the queue of workers at message s with type a or higher and $\hat{v}$ is the average valuation for their work

$$\widehat{v} = \sum_T \Psi(T|s)v_T.$$

Clearly the average valuation is increasing in the belief that the firm is of high type. We can now apply the same steps as in the Section 2.3, just with this replaced term for the value. As in (7) the increase in value has to equal the increase in outside option to make free entry feasible, so

$$e^{-\Lambda(a|s)}\widehat{v} = u'(a),$$

and the second derivative yields in analogy to (8)

$$\lambda(a|s)e^{-\Lambda(a|s)}\widehat{v} = u''(a)$$

so that in analogy to we obtain

$$\lambda(a|s) = u''(a)/u'(a).$$

So tightness-density is completely governed by the outside option, exactly as in the previous section. So in terms of types that approach a given initial signal, the allocation remains random where the types between $s = L$ and $s = H$ overlap, and it is identical to the types that queue in the fully revealing equilibrium we already studied. Wage bids are also relatively straightforward. At a given first signal s and a second signal that reveals type T, the wages are given by a differential equation identical to (12)

$$b_T'(a|T, s) = v_T B'(a), \tag{23}$$

where $B(a)$ was the "normalized" wage of a firm with valuation one and cost zero as defined in the main text.

The lowest worker type $\underline{a}_s$ that approaches message s can again only win if no other worker is present, no matter whether the second signal indicates that it is a high or a low type. So at either firm type this worker extracts all output conditional on matching, and in this case obtains

$$\begin{aligned} & \sum_T \Psi(T|s)v_T(\underline{a}_s - c_T) \\ = \quad & \widehat{v}[\underline{a}_s - \widehat{c}] \text{ with } \widehat{c} := \sum_T \Psi(T|s)[v_T c_T] / \sum_T \Psi(T|s)v_T. \end{aligned}$$

where the second line simply re-orders the terms. The term $\widehat{c}$ equals c_L when workers think they face a low type for sure (i.e., $\Psi(L|s) = 1$), $\widehat{c}$ is monotonically increasing in the probability that the firm is a high type, and it equals c_H when workers think they face a high type for sure. For the same reasons as in Section 2.3 the lowest type is solely determined by the costs via implicit function

$$\underline{a}_s = \widehat{c} + u(\underline{a}_s)/u'(\underline{a}_s). \tag{24}$$

The lowest type $\underline{a}_s$ is increasing in $\widehat{c}$, and therefore increasing in the probability of facing a high type firm.

This fully characterizes the interaction at jobs that do not have cheap-talk available to them, except that we dropped the indicator function after equation (22). We first note that $\underline{a}_s$ is always above c_H, because even $\underline{a}_{T=L}$ is above because of our restriction on the difference between c_L and c_H.[19] So the worker with lowest ability according to our derivations so far would bid for jobs both at high and low firms in a way that leaves non-negative profits for the firms, and this worker would just recover his/her outside option. We also know from our previous analysis that no worker with ability below $\underline{a}_s$ can recover their outside option if this worker makes offers to both firms that leave non-negative profits to these firms. Workers with abilities below c_H would choose not to be hired by high firm types. They would never be able to recover their outside option, though, either, as we now show in the final step:

Type $\underline{a}_{T=L}$ has

$$
\begin{aligned}
& e^{-\Lambda(\underline{a}_s|s)} \sum_T \Psi(T|s) v_T(\underline{a}_{T=L} - c_T) = U(\underline{a}_s) \\
\Rightarrow \quad & e^{-\Lambda(\underline{a}_s|s)} \Psi(L|s) v_L(\underline{a}_{T=L} - c_L) < U(\underline{a}_{T=L}) \\
\Rightarrow \quad & e^{-\Lambda(\underline{a}_s|s)} \Psi(L|s) v_L < \underline{v}_{T=L}(\underline{a}_{T=L}),
\end{aligned} \tag{25}
$$

where $\underline{v}_{T=L}(a) = U(a)/(a - c_l)$ as defined in Section 2.3. Any type $a < \underline{a}_{T=L}$ who only bids for the low job obtains $e^{-\Lambda(\underline{a}_s|s)} \Psi(L|s) v_L(\underline{a} - c_L)$ but has outside option $U(a) = \underline{v}_{T=L}(a)(a - c_L)$. For this to justify entry we have to have $e^{-\Lambda(\underline{a}_s|s)} \Psi(L|s) \geq \underline{v}_{T=L}(a)$, but this is contradicted by (25) and the fact that $\underline{v}_{T=L}(\underline{a}_{T=L})$ is the global minimum of $\underline{v}_{T=L}(\underline{a})$.□

The proof of Corollary 1 entailed the claim that holding the lowest worker type $\underline{a}_m$ at truthful message m fixed but changing v_m does not alter the expected utility at the lowest type, and we prove this here: The expected utility at $\underline{a}_m$ is according to (11) given by $e^{-\Lambda(\underline{a}_m|m)} v_m(\underline{a}_m - c_m)$. We will see that the equilibrium utility at the lowest type does not change, and therefore the lowest type remains unchanged when the valuation is altered. To do this, we will take comparative statics with respect to v_m, holding $\underline{a}_m$ fixed, but allowing $\Lambda(\underline{a}_m|m)$ to change. The derivative is given by

$$
\begin{aligned}
& e^{-\Lambda(\underline{a}_m|m)}(\underline{a}_m - c_m) + e^{-\Lambda(\underline{a}_m|m)} \left[-\frac{d\Lambda(\underline{a}_m|m)}{dv_m} \right] v_m(\underline{a}_m - c_m) \\
= \quad & e^{-\Lambda(\underline{a}_m|m)}(\underline{a}_m - c_m) \left[1 - v_m \lambda(\bar{a}_m|m) \frac{d\bar{a}_m}{dv_m} \right] \\
= \quad & e^{-\Lambda(\underline{a}_m|m)}(\underline{a}_m - c_m) \left[1 - v_m \lambda(\bar{a}_m|m) u''(\bar{a}_m) \right] \\
= \quad & e^{-\Lambda(\underline{a}_m|m)}(\underline{a}_m - c_m) \left[1 - \frac{v_m}{u'(\bar{a}_m)} \right] \\
= \quad & 0
\end{aligned}
$$

The second line is obtained simply by rewrites the change in the queue of workers, the third line by implicitly differentiating (10), the forth line uses (9) to substitute out $\lambda(.)$, and the last line again uses (10).□

Proof of Proposition 1: Part 1 and 3 follow directly from Lemma 1 and 2. For Part 2, note

[19]Equation (24) applied to $\widehat{c} = c_L$ implies that we are looking for $\underline{a}_{T=L} = c_L + u(\underline{a}_{T=L})/u'(\underline{a}_{T=L})$. We assumed that $c_H - c_L < u(c_H)/u'(c_H)$ or equivalently that $c_H < c_L + u(c_H)/u'(c_H)$. Therefore, $\underline{a}_{T=L} > c_H$. Since the cutoff at any other $\widehat{c}$ is higher, it will also be above c_H.

that effects on the number of applications are ambiguous since low firm types who send a truthful message attract additional low ability workers but loose high ability workers relative to their no-message counterparts. How this affects applications depends on the number of workers that they add at low abilities and the number of workers they loose at high abilities. This in turn depends on two things: how many additional ability types are added or lost, and the queue of workers at each ability level (conditional on applying at all) which by (9) equals $u''(a)/u'(a)$. For $v_H \approx v_L$ the size of the high-ability segment that is lost shrinks to zero, and truthful cheap-talk attracts more workers to low-message firms. The opposite is the case for high-message firms. In the case of $v_H > v_L$ but $c_H \approx c_L$, the size of the low ability segment that is lost or added shrinks to zero, and truthful messages at low type firms induces a sizeable loss of high-ability workers without an offsetting gain at low-ability workers. For high type firm truthful cheap talk induces a gain in high ability workers with an offset of low ability workers.

For Part 4, note that efficiency per firm with a message is given by

$$\max_{\lambda^P_{m=L}(.),\lambda^P_{m=H}(.)} \delta_L S(v_L, c_L|\lambda^P_{s=L}(.)) + \delta_H S(v_H, c_H|\lambda^P_{m=H}(.)) \tag{26}$$

where

$$\begin{aligned} S(v,c|\lambda(.)) &= \int \left(v(a-c)\right) d\left(1-e^{-\int_a^\infty \lambda(a')da'}\right) - \int \lambda(a)u(a)da \\ &\int v(a-c)\lambda(a)e^{-\int_a^\infty \lambda(a')da'} da - \int \lambda(a)u(a)da \end{aligned}$$

represents the output of the firm minus the forgone outside option of all the workers who queue at this firm. Let

$$\mathcal{S}(v,c) = \max_{\lambda(.)} S(v,c|\lambda(.)). \tag{27}$$

The optimality conditions for this problem coincide nearly everywhere with the equilibrium conditions (7) for a given firm type, and it shares the same boundary conditions. This establishes that the equilibrium decentralizes the planner's solution. Evidently the value of program (26) is given by $\delta_L \mathcal{S}(v_L, c_L) + \delta_H \mathcal{S}(v_H, c_H)$.

For firms who cannot send a message the maximal efficiency per firm is attainable through queue of

$$\begin{aligned} \max_{\lambda^P_{s=L}(.),\lambda^P_{s=H}(.)} \quad & \delta_L \left[\psi_1 S(v_L, c_L|\lambda^P_{s=L}(.)) + (1-\psi_1) S(v_L, c_L|\lambda^P_{s=H}(.))\right] \\ & + \delta_H \left[\psi_1 S(v_H, c_H|\lambda^P_{s=H}(.)) + (1-\psi_1) S(v_H, c_H|\lambda^P_{s=L}(.))\right], \end{aligned} \tag{28}$$

where $S(.)$ still denotes the output minus the forgone outside option of workers, but which queue a firm obtains now depends on which signal it sends. The optimality conditions for this problem again coincide with conditions for the decentralzied equilibrium, and thus the decentralized equilibrium obtains the highest surplus given the information frictions.

Nevertheless, evidently value of maximization problem (26) is higher than that of maximization problem (28) since the queue length in the former can be chosen to maximize the surplus for each firm while in the latter it has to partially cater to both firm types. Formally, note that both $S(v_L, c_L|\lambda^*_{s=L}(.))$ and $S(v_L, c_L|\lambda^*_{s=H}(.))$ are bounded above by $\mathcal{S}(v_L, c_L)$, and $S(v_H, c_H|\lambda^P_{s=H}(.))$ and $S(v_H, c_H|\lambda^P_{s=L}(.))$ are bounded above by $\mathcal{S}(v_L, c_L)$. Therefore the value of problem (28) is bounded from above by $\delta_L \mathcal{S}(v_L, c_L) + \delta_H \mathcal{S}(v_H, c_H)$, which coincides with the value of problem (26).

In fact, one can prove that value of problem (28) is strictly below that of (26): the maximizer $\lambda^{*P}(a|v, c)$ of problem (27) is unique for given (v, c) and strictly different for different values of (v, c) as it inhits the same properties as the equilibrium queue length under truthful cheap talk. Equal value of the two problems above would require that that $\lambda^{*P}(.|v_L, c_L) = \lambda^P_{s=L}(.) = \lambda^P_{s=H}(.)$ and $\lambda^{*P}(a|v_H, c_H) = \lambda^P_{s=L}(.) = \lambda^P_{s=H}(.)$, which yields a contradiction to $\lambda^{*P}(.|v_L, c_L) \neq \lambda^{*P}(a|v_H, c_H)$.□

B Short discussion of model features and robustness

This appendix section briefly discusses the two extensions of Subsection 2.4 and the analysis of partially revealing equilibria from Subsection 2.5.

B.1 The role of the second signal

Under truthful cheap talk the second signal plays no role. In the absense of cheap talk the second signal ensures that workers can target their bidding to more information after having read the details of the job description. Taking the limiting case where this signal is arbitrarily precise simplifies the mathematical representation of the problem substantially as the bid can be expressed as if the worker knows the type of the firm, and from an ex-ante point of view (before the second signal is revealed) the bidding costs are just the sum of the costs at the different firm types. It turns out that this ensures that workers continue to approach firms at random at intermediate ability levels.

Consider instead a setting where workers remain unsure about the firm type at the point of bidding, and assign probability $\Psi(T)$ to it being type T. It is easy to show that their bids would follow differential equation[20]

$$b(a) / \left(\sum\nolimits_T \lambda(a) \Psi(T) / (v_T - b'(a)) \right) = 1.$$

This converges to the differential equation in (12) when $\Psi(T) \longrightarrow 1$, as do the associated endpoint conditions. Note that this ensures also that equilibria converge to the limit equilibrium, since free entry and associated block-recursivity imply that worker entry and behavior at these firms can be studied without effect on the other firms (i.e., on those with other first signals or with truthful messages), as each attracts their "own" set of workers.

[20]See our previous working paper Horton, Johari and Kircher (2021) that derives this and solves the related model, but at the expense of assuming that workers do not receive any first signal. This entails the very strong implication that workers select among firms at random in the absence of messages. The formulation in this paper allows clear and tractable insights that allows for selection based on some information even without signals.

B.2 More general production functions

Our production function is a special case of increasing and concave production function $\phi(a, T)$ that describes the output of a worker of type a at a firm of type T. Assume that this function is increasing in its arguments. Assume that it is concave in a. Similar steps as in the main body yield under fully revealing messages in analogy to (9) that

$$\lambda_m(a) = u''(a)/u'(a) - \phi_{aa}(a, m)/\phi_a(a, m), \tag{29}$$

with upper and lower end-point conditions $\phi_a(\bar{a}_m, m) = u'(\bar{a}_m)$ and $e^{-\Lambda(\underline{a}_m|m)}\phi_T(\underline{a}_m, m) = u(\underline{a}_m)$. The upper endpoint continues to be strictly higher for higher firm types if ϕ is strictly supermodular in a and T, which was assured in our baseline specification. The queue length $\lambda_T(a)$ was independent of T in our baseline specification, which made the setup particularly tractable as it implied that workers choose among firms "as if" at random. From the earlier equation it is apparent that this only continues to apply more generally if the semi-elasticity of ϕ_a with respect to worker ability is independent of T, as is the case for example for productions $\phi(a, T) = \ln(T(a - d_1)) + d_2$ or $\phi(a, T) = f(T)g(a - d_3) + h(T)$, where d_j are positive constants and $f(.)$, $g(.)$ and $h()$ are non-negative and increasing functions (and, to preserve concavity in ability, $g(.)$ needs to be concave). We assume this for the rest of this section (things become substantially more involved when the semi-elasticity of ϕ_a depends on the firm type, which exceeds the bounds of this paper.)

It remains to be clarifies when the lower bound on ability $\underline{a}_T$ is higher under for high than for low firm types. To study this, it is easiest to think of L and H as real numbers and to define our production function for all $T' \in [L, H]$ in a way that is twice continuously differentiable, preserves the assumptions we had placed on $\phi(a, T)$ for all T, and preserves the true functions $\phi(a, L)$ and $\phi(a, T)$ at the boundaries. For our example, this is for example achieved by setting $\phi(a, T) = \left(v_L + (v_H - v_L)\frac{T-L}{H-L}\right)\left(a - c_L - (c_H - c_L)\frac{T-L}{H-L}\right)$ so that valuations and costs rise linearly. The advantage of this extension is that we can use calculus to study the question at hand. We show that the lower bound weakly increases in T whenever it is possible to such a function that is weakly increasing, as is the case in the example above.

Proposition 5. *Consider more general production functions $\phi(a, T)$ with a semi-elasticity of ϕ_a in a that is independent of firm type. Strict supermodularity (submodularity) of ϕ ensures that the highest ability $\bar{a}_T$ is increasing (decreasing) in T. Strict log-supermodularity (log-submodularity) of ϕ ensures that the lowest ability type $\underline{a}_T$ is increasing (decreasing) in T.*

Proof of Proposition 5: Consider truthful cheap talk. Recall that the the queue length is given by $\lambda(a|T) = u''(a)/u'(a) - \phi_{aa}(a, T)/\phi_a(a, T)$, with end-point conditions $\phi_a(\bar{a}_T, T) = u'(\bar{a}_T)$ and $e^{-\Lambda(\underline{a}_T|T)}\phi(\underline{a}_T, T) = u(\underline{a})$.

Applying the implicit function theorem to the first endpoint condition and noting that $\phi_a a(a, T) - U''(a) < 0$ establishes the comparative statics on $\bar{a}_T$.

Applying the implicit function theorem to the latter and noticing that the denominator is negative (otherwise the types above a_T would not have positive surplus and would not enter) implies that the sign of $-\partial\underline{a}_T/\partial T$ equals the sign of

$$
\begin{aligned}
& e^{-\Lambda(\underline{a}_T|T)}\phi_T(\underline{a}_T,T)+e^{-\Lambda(\underline{a}_T|T)}\left[-\frac{d\Lambda(\underline{a}_T|T)}{dT}\right]\phi(\underline{a}_T,T) \\
= \;& e^{-\Lambda(\underline{a}_T|T)}\phi(\underline{a}_T,T)\left[\frac{\phi_T(\underline{a}_T,T)}{\phi(\underline{a}_T,T)}-\lambda(\bar{a}_T|T)\frac{d\bar{a}_T}{dT}\right] \\
= \;& e^{-\Lambda(\underline{a}_T|T)}\phi(\underline{a}_T,T)\left[\frac{\phi_T(\underline{a}_T,T)}{\phi(\underline{a}_T,T)}-\lambda(\bar{a}_T|T)\left[-\frac{\phi_{aT}(\bar{a}_T,T)}{\phi_{aa}(\bar{a}_T,T)-u''(\bar{a}_T)}\right]\right] \\
= \;& e^{-\Lambda(\underline{a}_T|T)}\phi(\underline{a}_T,T)\left[\frac{\phi_T(\underline{a}_T,T)}{\phi(\underline{a}_T,T)}-\frac{u''(\bar{a}_T)-\phi_{aa}(\bar{a}_T,T)}{\phi_a(\bar{a}_T,T)}\left[\frac{\phi_{aT}(\bar{a}_T,T)}{u''(\bar{a}_T)-\phi_{aa}(\bar{a}_T,T)}\right]\right] \\
= \;& e^{-\Lambda(\underline{a}_T|T)}\phi(\underline{a}_T,T)\left[\frac{\phi_T(\underline{a}_T,T)}{\phi(\underline{a}_T,T)}-\frac{\phi_{aT}(\bar{a}_T,T)}{\phi_a(\bar{a}_T,T)}\right]
\end{aligned}
$$

A sufficient condition for this to be weakly (strictly) negative is that ϕ is weakly (strictly) log-supermodular. In that case

$$
\frac{\phi_T(\underline{a}_T,T)}{\phi(\underline{a}_T,T)}-\frac{\phi_{aT}(\bar{a}_T,T)}{\phi_a(\bar{a}_T,T)}\leq\frac{\phi_T(\bar{a}_T,T)}{\phi(\bar{a}_T,T)}-\frac{\phi_{aT}(\bar{a}_T,T)}{\phi_a(\bar{a}_T,T)}\leq 0,
$$

where both the first and the second inequality are immediate implications of strict log-supermodularity. The opposite holds for strict log-submodularity.□

C Information Revelation

C.1 Incentive Compatibility

Proof of Proposition 4:

Consider first the limiting case where the second signal has very high precision. Note first that the probability that a firm at message m has no applicant with qualification above a is $e^{-\Lambda(a|m)}$. That means that the density that the best applicant is of type a is $\lambda(a|m)e^{-\Lambda(a|m)}$. So the expected profit at message m for firm type t is at least

$$
\begin{aligned}
\Pi(m|t) \;\geq\; & \int_{\underline{a}_m}^{\bar{a}_m}\max\{v_t(a-c_t)-b_m(a),0\}\lambda(a|m)e^{-\Lambda(a|m)}da. \\
= \;& \int_{\underline{a}_m}^{\bar{a}_m}\max\left\{\frac{v_t(a-c_t)-b_m(a)}{v_m},0\right\}u''(a)da. \qquad (30) \\
= \;& \int_{\underline{a}_m}^{\bar{a}_m}\max\left\{\frac{v_t(a-c_t)-\frac{v_m}{v_L}b_L(a)}{v_m},0\right\}u''(a)da. \\
= \;& \int_{\underline{a}_m}^{\bar{a}_m}\max\left\{\frac{v_t}{v_m}(a-c_t)-\frac{b_L(a)}{v_L},0\right\}u''(a)da,
\end{aligned}
$$

where it is suppressed that $\bar{a}_m$ is a function of v_m and $\underline{a}_m$ is a function of c_t. The first line holds exactly with equality if workers with higher ability offer more profits. In this case the highest a is always selected, and leaves profits $v_t(a-c_t)-b_m(a)$ to the firm whenever this is positive. The second line simply substitutes (8), and the third line uses $b_H(a)=\frac{v_H}{v_L}b_L(a)$ which follows directly from Corollary 1.

On the equilibrium path we have proven that higher types indeed offer more profits, so the first line holds with equality when $m=t=L$ or $m=t=H$. It also holds with equality when high types contemplate to send a low message ($t=H$ but $m=L$): since $v_L(a-c_L)-b_L(a)$ is increasing in a in equilibrium, also $v_L(a-c_L)-b_L(a)$ is increasing for all $a\geq c_H$, and worker types that have a lower a are always rejected as taken into account in the max operator. This makes it easy to analyze these three cases. It may not hold in the case of low firm types sending a high message ($t=L$ and $m=H$) which makes its analysis more challenging.

Truth-telling requires that $\Pi(H|H)\geq\Pi(L|H)$ and $\Pi(L|L)\geq\Pi(H|L)$.

We will first consider the truth-telling condition for high types $(T=H)$*:*

$$\Pi(H|H) \geq \Pi(L|H) \tag{31}$$

$$\Leftrightarrow \int_{\underline{a}_H}^{\bar{a}_H}\left[a-c_t-\frac{b_L(a)}{v_L}\right]u''(a)da \geq \int_{\underline{a}_L}^{\bar{a}_L}\max\left\{\frac{v_H}{v_L}(a-c_t)-\frac{b_L(a)}{v_L},0\right\}u''(a)da \tag{32}$$

For $0=c_L=c_H$, it holds that $\underline{a}_H=\underline{a}_L=0$. Fix v_L, and vary v_H. The derivatives are

$$\frac{d\Pi(H|H)}{dv_H} = \left[\bar{a}_H-\frac{b_L(\bar{a}_H)}{v_L}\right]u''(\bar{a}_H)\frac{d\bar{a}_H}{dv_H}=\bar{a}_H-\frac{b_L(\bar{a}_H)}{v_L}, \tag{33}$$

$$\frac{d\Pi(L|H)}{dv_H} = \int_{\underline{a}_L}^{\bar{a}_L}\frac{a}{v_L}u''(a)da, \tag{34}$$

where the second equality in the first line follows from implicitly differentiating (10). The first line goes to infinity since $b_L(a)$ is bounded away from v_La for all $a>\underline{a}_H+\varepsilon$ for some ε, which follows from (12) and the assumption that $au''(a)/u'(a)$ is bounded away from zero. The second line is simply a constant. So truthtelling is incentive compatible for high types when v_H is sufficiently large.

When $v_H\approx v_L$ and $c_H=0$, truth-telling requires

$$\bar{a}_L-\frac{b_L(\bar{a}_L)}{v_L}\geq\int_{\underline{a}_L}^{\bar{a}_L}\frac{a}{v_L}u''(a)da, \tag{35}$$

which follows since $\bar{a}_H\approx\bar{a}_L$.

Now consider truth-telling for low types $(T=L)$:

We will first consider truthtelling when valuations of high and low type firms are roughly similar, and will show that incentive compatibility generically fails when v_H approaches v_L and $c_H=0$. In this case, it has to hold that

$$\int_{\underline{a}_L}^{\bar{a}_L}\left[a-\frac{b_L(a)}{v_L}\right]u''(a)da\geq\int_{\underline{a}_H}^{\bar{a}_H}\max\left\{\frac{v_L}{v_H}a-\frac{b_L(a)}{v_L},0\right\}u''(a)da.$$

Taking the derivative with respect to v_H and evaluating it around $v_H \approx v_L$ (and $\bar{a}_H \approx \bar{a}_L$ as well as $\underline{a}_H = \underline{a}_L = 0$) yields

$$0 \geq -\int_{\underline{a}_H}^{\bar{a}_H} \frac{a}{v_L} u''(a) da + \left[\bar{a}_H - \frac{b_L(\bar{a}_H)}{v_L}\right]. \tag{36}$$

Inequality (36) is the opposite of inequality (35), and generically in the space of permissable $u(.)$ and v_L they do not both hold simultaneously.

Next, we consider truth-telling of low types when v_H is large (for fixed v_L) and will show that it is incentive compatible. We will first build intuition under the case where firms are willing to hire all workers, so that $0 = c_L = c_H$. It will then be easy to expand to the general case afterwards. For this part of the proof it is convenient to denote by $\check{a}$ the ability type with $u'(v) = 1$ or ability type $\bar{a}_L$, whichever one is larger. Let $\check{v} \geq v_L$ be such that $\bar{a}_H = \check{a}$ if $v_H = \check{v}$. Note that if the low firm type could only match with workers of abilities $a \in [0.\check{a}]$, it is easy to show that deviating and sending a high message is not profitable:

- For abilities $a \in [\underline{a}_H, \bar{a}_L]$ the queue does not change but the wage becomes $b_H(a) = b_L(a)v_H/v_L$, so the firm pays more for these workers but does not attract more of them.

- For abilities $a \in [0, \underline{a}_H]$ workers cease to approach the firm if it sends the high signal.

- For abilities $a \in [\bar{a}_L, \check{a}]$, let $\bar{g} = \sup_{a\in[\bar{a}_L,\check{a}]} g(a) > 0$ so that the wage $b_H(a) \geq \bar{g}v_H/v_L$, which exceeds $v_L\check{a}$ for v_H sufficiently large. So the low type firm would make a loss hiring these workers, and would refuse to do so. So their presence does not constitute a profitable deviation.

Now consider also abilities beyond $\check{a}$. In general, by raising v_H from $\check{v}$ to higher values, from the perspective of a low-type firm that contemplates to send message H this has two effects: For all ability types $a \in [0, \bar{a}_H]$ the wage $b_H(a) = v_H g(a)$ increases when v_H is marginally increased, which is a negative outcome from the firm's perspective. But one also raises the upper end of workers willing to match. We will show that even for the second effect, the additional wage will outweigh the benefit in matching. We know that $u'(\bar{a}_H) = v_H$. So $d\bar{a}_H/dv_H = 1/u''(\bar{a}_H)$. So the wage at the higest type increase by

$$\begin{aligned}
db_H(\bar{a}_H)/dv_H &= d(v_H g\,(\bar{a}_H))/dv_H \\
&= g\,(\bar{a}_H) + v_H g'(\bar{a}_H) d\bar{a}_H/dv_H \\
&= g\,(\bar{a}_H) + v_H g'(\bar{a}_H) d\bar{a}_H/dv_H \\
&= g\,(\bar{a}_H) + v_H[1 - g(\bar{a}_H)u''(\bar{a}_H)/u'(\bar{a}_H)]/u''(\bar{a}_H) \\
&= g\,(\bar{a}_H) + v_H/u''(\bar{a}_H) - g(\bar{a}_H)/u'(\bar{a}_H)
\end{aligned}$$

For the low firm type who considers sending message H, a higher value v_H implies the following change in profits when matching with the highest type at message $H : v_L\bar{a}_H - b_H(\bar{a}_H)$. So this changes in v_H according to

$$\begin{aligned}
& v_L d\bar{a}_H/dv_H - db_H(\bar{a}_H)/dv_H \\
= \;& v_L/u''(\bar{a}_H) - g\,(\bar{a}_H) - v_H/u''(\bar{a}_H) + g(\bar{a}_H)/u'(\bar{a}_H)
\end{aligned}$$

which is definitely negative as $\bar{a}_H \geq \check{a}$, which implies $u'(\bar{a}_H) \geq 1$. This establishes that for high enough v_H low types strictly prefer to send message $m = L$ to message $m = H$ in a truthtelling equilibrium. Since high types also prefer to tell the truth, this establishes that truth-telling is an equilibrium outcome when $c_L = c_H = 0$.

We finally have to establish that truth-telling is incentive compatible for low firm types when v_H is sufficiently large even in the more general setting where $0 = c_L < c_H$: We have already established that low firm types would not pretend to be high firm types if costs are zero. In the candidate equilibrium with truth-telling, note that strictly positive costs at high types do not change the workers who approach low types or their bids. On the other hand, at high types the set of worker types shrinks, those who do approach high messages do so with exactly the same queue $\lambda(a|H)$ by (9), and their bids increase. So if low types do not want to deviate and pretend to be high types under $c_H = 0$, they will not want to do so when $c_H > 0$. For high types we had already established incentive compatibility in the general case, which concludes the proof for the case when the second signal has limit precision of $\psi_2 = 1$.

Finally, consider the case where there remains some imprecision in the second signal ($\psi_2 < 1$), but precision is sufficiently high. In Section B.1 we established the under the assumption of truth-telling equilibria converge to the limit equilibrium when the precision ψ_2 of the second signal increases to 1. So payoffs from truth-telling and lying converge to the limit arguments. Since we have shown above that for sufficient differences in valuations both firm types strictly prefer truth-telling in a candidate equilibrium with full type revelation, this continues to hold when the signal is sufficiently precise. Furthermore, the proof above establishes that for equal costs and valuations that are very close only a vanishingly small set of parameter values induces truth-telling, while for all others at least one type has strict incentives to pretend to be the other type. For a given set of parameters where there are strict incentives to "lie" (i.e., strict incentives to send the message that does not correspond to the own type) when the second signal has limit precision, a sufficiently high second signals will again also induce incentives to lie. $\square$

C.2 Incentive Compatibility for more general production functions

Proposition 4 prove that in a candidate equilibrium with truth-telling each type has incentives to reveal the truth when valuations between high and low type firms are sufficiently far apart. This was done for the specific production function $v_T(a - c_t)$. This also holds for the somewhat broader class of production functions studied in Section B.2. To see this, consider again a production function $\phi(a, T)$ on R^2_+ that is increasing in both arguments, convex in a, and where the semi-elasticity of $\phi_a(a, T)$ with respect to a is constant in T. Again assume that for each $T > T$ at least some ability has positive surplus, and assume that the lowest firm type for which this arises remains bounded (i.e., for each $T > 0$ there exists $a \leq \bar{a}$ such that $\phi(a, T) > u(a)$). Fix some low type $L > T$.

We will show that there is incentive compatibility if the valuation of the high type is sufficiently large. In particular, assume that $\phi_a(a, T)$ at each a increases in a and is unbounded.

We then obtain the following insight: *Under the assumptions in this subsection, if high type $T = H$ is sufficiently large, then in a candidate equilibrium where each firm type sends a message equal to its type no firm has incentives to deviate to the other message.*

Establishing this follows the same logic as the proof for the specific production function of the main body of the paper. We highlight here the key steps. First, note that the same derivations that established the differential equation for the bidding function (12) in the main body of the paper reveal that the bidding function in equilibrium at message $m \in \{L, H\}$ is given by

$$b'_m(a) = \phi_a(a, m) - b_m(a) \left[u''(a)/u'(a) - C(a)\right] \tag{37}$$

where $C(a) := \phi_{aa}(a, m)/\phi_a a, m$ is the semi-elasticity of ϕ_a in a and is independent of T by our earlier assumptions. Since the bid-equation at the high message becomes increasingly steep with higher H, it is easy to show that the set A of worker abilities that would approach high type firms and make bid $b_H(a)$ below the output at the low firm (i.e., $b_H(a) < \phi(a, L)$) has measure that converges to zero and the highest such ability worker remains bounded.[21] So for a low firm type the profit after a deviation to a high message shrinks to zero, and therefore becomes lower than their profit on the equilibrium path.

For high type firms, the benefit of sending a high message is that of trading with an ever-expanding set of trading partners (as H increases the upper bounded increases to infinity and the set of active abilities has to increase to infinity as well as otherwise the lower end-point condition could not be satisfied). Trading with ever-larger set of types become eventually more profitable than deviating and trading with the fixed set $[a_L, a_L]$. This establishes incentive compatibility for both types when H is sufficiently large.

We omit the generic failure of truth-telling when types are very close together as this is a standard result that appears in much of the literature.

C.3 Discussion of the role of imprecise second signals:

When we establish incentive compatibility in Proposition 4 we use the limit formulation with perfect second signals, but establish that incentives for truth-telling are strict when valuations are far enough apart, and because of convergence this carries over to settings with imprecise second signals as long as the second signal is sufficiently precise.

While our results carry over when signals are sufficiently precise, it is worth discussing what happens if there is no second signal and workers simply know the true firm type at that point. This does change incentive compatibility for the following reason: our current setup studies truthful cheap talk, and if an L-firm deviated and sends message $m = H$ then workers believe that this is an H-firm (as only these send such messages in equilibrium) and approach and bid at such firms as if it is an H-firm. The deviant firm would attract better workers but these would demand higher wages. If workers were to learn the actual firm type through the second message, low type firm that mis-represents its type could attract more applicants, but these would learn the firm's type and then bid low. One can show that truth-telling in this setting would not be sustained. So our formulation is best viewed as an approximation to the setting where signals are not fully revealing. It also highlights the value of our empirical setting: it is not obvious that allowing for cheap-talk will increase market information, as this would fail under different informational assumptions, and off course under different equilibria such as the well-known "babbling" equilibrium where no information is transmitted.

[21] Recall for this that the queue length of any given worker type is constant in T by 29.

C.4 Partially revealing equilibria

Truthful cheap talk is not the only equilibrium in cheap-talk games, even if such equilibria exists. It is well-known that there always exists a babbling equilibrium where receivers disregard the messages because senders have choosen them at random. Whether parameters admit a truthful equilibrium or not, there also often exist partially revealing equilibrium where some additional information is transmitted, but not all types are perfectly revealed. Here we briefly point out that if a cheap-talk intervention displays any effect, it must be linked to some information transmission. We show here that all equilibria where cheap-talk reveals some additional information inherit the same broad empirical implications: Part 1,2 and 4 of Proposition 1 and the Corollaries 1 and 2 that lead up to it are based around comparative statics. Less information essentially acts to average the valuation and the costs. The base model does it for the extreme case of starting with perfect information revelation, but the comparative statics hold similarly when messages are only partially informative.[22]

D Balance

To assess the effectiveness of randomization, in Table 2 we report the mean values for various pre-randomization attributes of employers (the top panel) and their job openings (the bottom panel), for both the ambiguous and explicit arms pooled, by whether preferences were shown. We can see there is excellent balance on pre-treatment characteristics, both for employers and job openings. Balance is unsurprising, as the platform has used the software for randomization many times in previous experiments.

E Additional empirical results

E.1 Effects on wage bidding

In the bottom panel of Figure 13 (which also shows the quantile means for the wage bid and profile rate), we can see that markups were higher in the high tier and lower in the low tier following signal revelation. There is some evidence that revelation has little effect on markups for all tiers around the 80th percentile. But outside of this range, we can see clear effects on markups in the expected direction. The effect of revelation on the markup shows us that compositional changes do not explain all of the change in wage bids.

E.2 Do wage bids reflect compensating differentials?

A high type employer might also be a more demanding employer, expecting greater effort from their hires. Anticipating these great expectations, workers might bid more, as they know their

[22] Parts 1, 3 and 4 of Proposition Proposition 1 rely on comparative statics and carry over to partially informative equilibria. In contrast, Part 2 of Proposition 1 uses a limit argument for the the size of the high valuation, holding the amount of information that is conveyed constant. The approach to the proof extends if there is a lower bound on the information that a partially informative equilibrium conveys about firms, but the technique of the proof does not extend if the amount of information goes to zero as the valuation of the high type firm becomes large.

Figure 12: Comparison of outcomes using applicant-level regression in the ambiguous arm with extensive pre-treatment job level controls

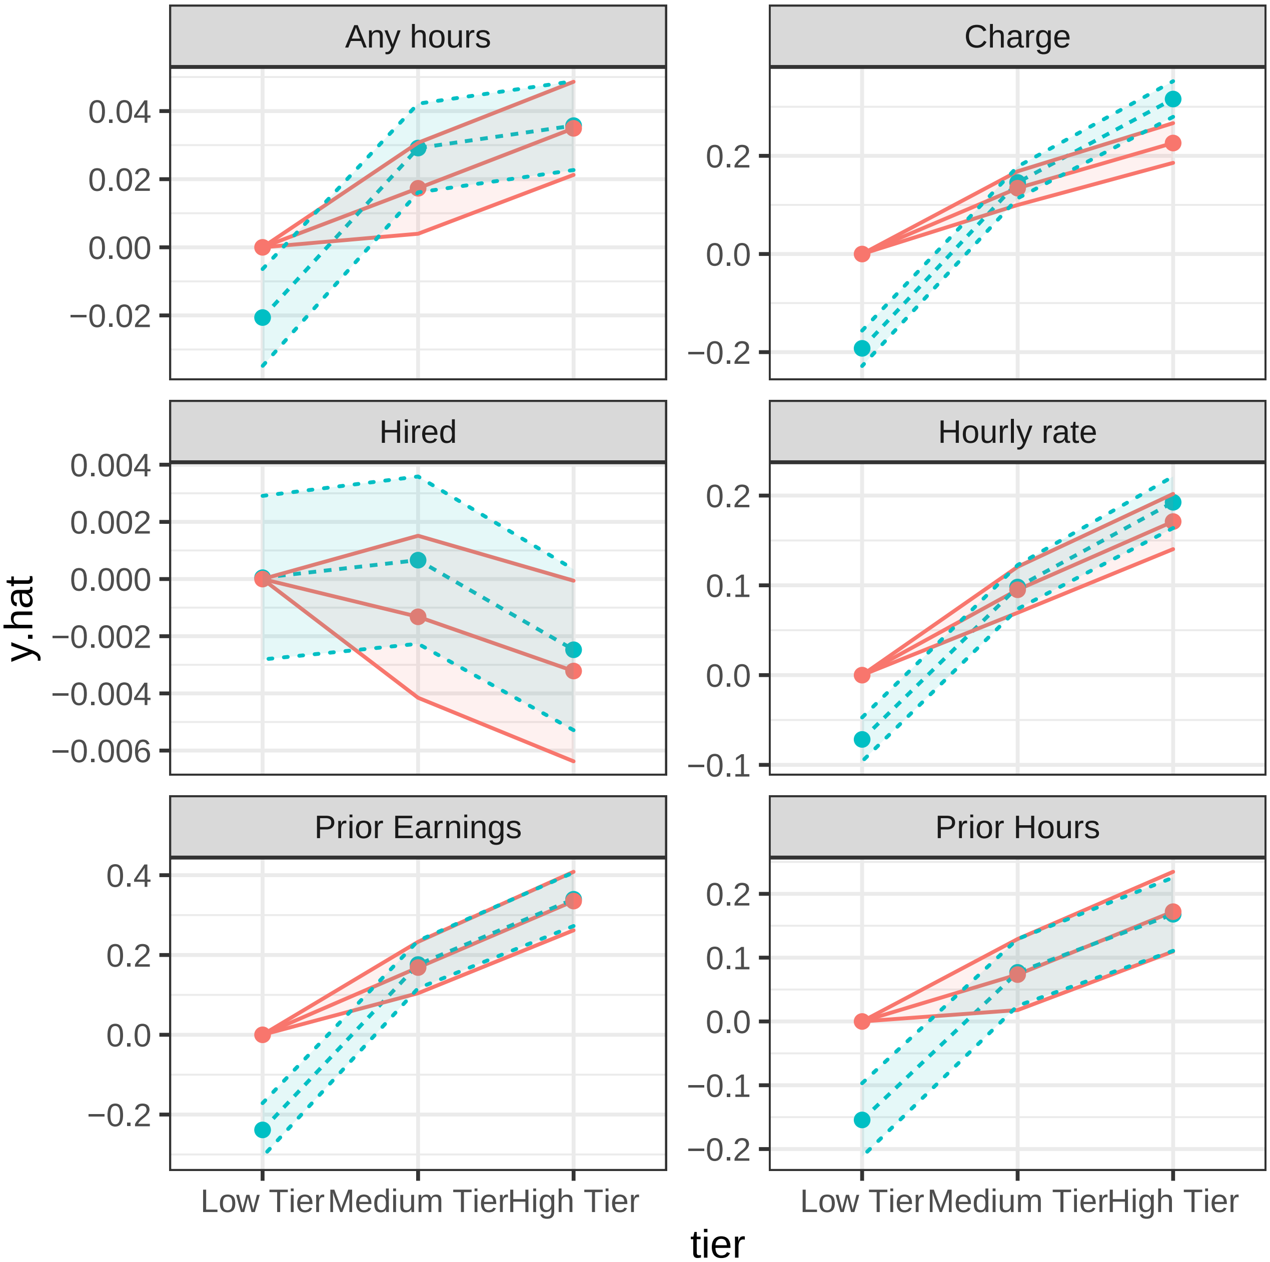

Notes: This figure plots coefficients from estimates of (16). The outcome in both panels is the applicant's cumulative prior hourly earnings at the time of application. The sample in the left panel is all applicants in the ambiguous arm of the experiment, whereas in the right panel, the sample is hired workers in the ambiguous arm. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with the lowest rates."; (2) Intermediate: "I am looking for a mix of experience and value."; (3) Expert: "I am willing to pay higher rates for the most experienced [workers]." We refer to these tiers as "low," "medium," and "high," respectively. Employers in the ambiguous arm were told that the platform *might* reveal their preferences to workers; whether workers were shown employer vertical preferences was randomly determined *ex post*. The error bars indicate the 95% confidence interval for the conditional mean.

Figure 13: Effects of showing employer vertical preferences, ShownPref = 1, on applicant pool composition with respect to wage bidding in the ambiguous arm

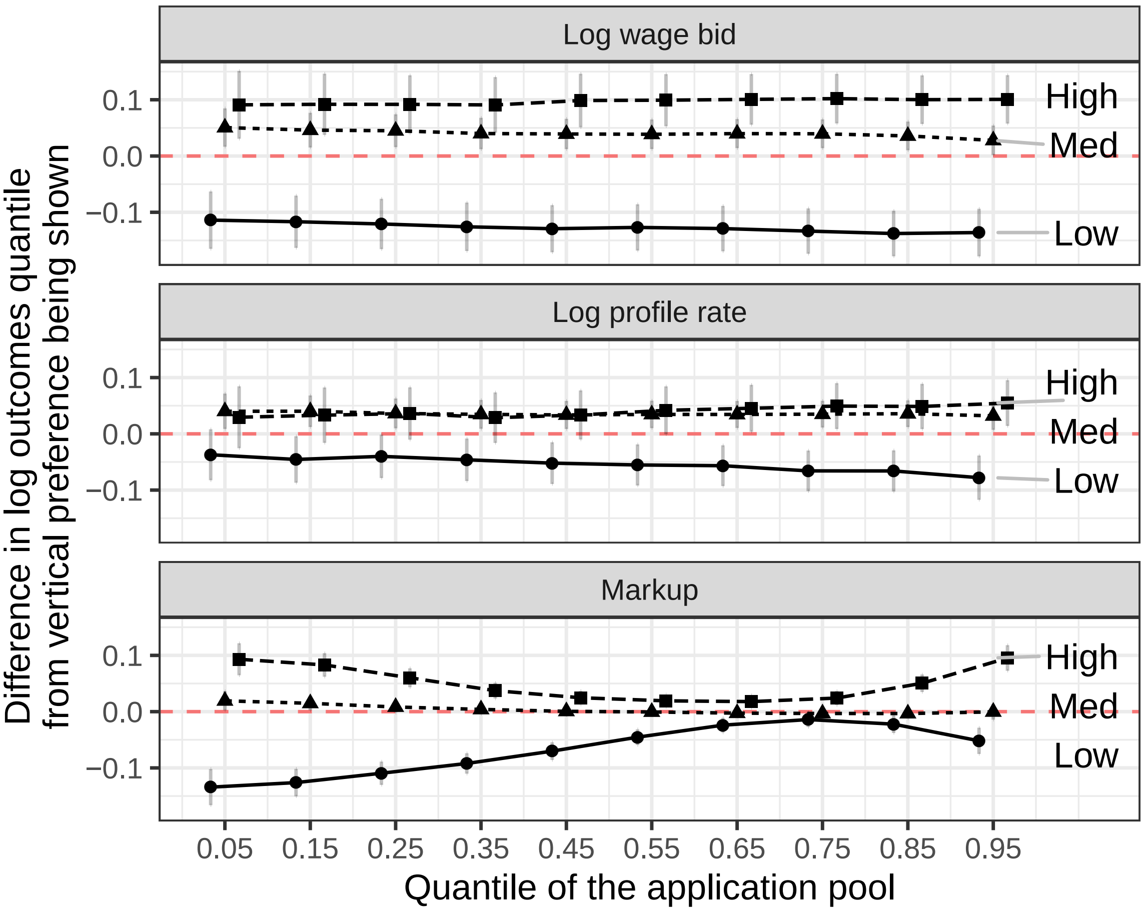

Notes: The figure shows the effects of vertical preference revelation on the composition of applicant pools with respect to wage bidding and profile rates. The sample is the ambiguous arm of the experiment. Each point is the mean effect of revelation on some applicant attribute at that quantile of the pool. For example, in the top facet, the effect of signal revelation for a high tier employer on the median applicant's wage bid is about 10 log points, of 10%. The error bars indicate the 95% confidence interval for the conditional mean.

Table 2: Employer and job opening characteristics by whether tier choice was shown, with job openings pooled from both arms of the experiment

	ShownPref=0	ShownPref=1	Δ	% Change
Employer attributes				
Prior job openings	4.29 (0.12)	4.18 (0.08)	-0.10 (0.13)	-2.44
Prior spend (log) by employers	7.12 (0.03)	7.11 (0.02)	-0.01 (0.03)	-0.11
Num prior workers	4.38 (0.15)	4.32 (0.09)	-0.06 (0.16)	-1.45
Job opening attributes				
Prefered experiance in hours	30.63 (0.80)	31.67 (0.75)	1.04 (1.10)	3.40
Estimated job duration in weeks	15.35 (0.14)	15.40 (0.12)	0.04 (0.18)	0.27
Job description length (characters)	553.29 (3.94)	556.64 (3.45)	3.35 (5.23)	0.61

Notes: This table reports means for a number of pre-randomization characteristics for the employer and job opening by ShownPref status. The data are pooled to include employers from both the ambiguous and explicit arms. Standard errors are reported next to the estimate, in parentheses. The far right column also reports the percentage change in the ShownPref = 1 group, relative to the mean in the control group. Significance indicators: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

costs will higher, either from greater effort or perhaps the greater probability of receiving bad feedback. As such, part of the higher wage bid observed in the high tier could reflect this anticipated greater, costly effort. Although we have no direct test of this hypothesis, several pieces of evidence make this compensating differential explanation relatively improbable relative to the straightforward perceived willingness to pay argument.

First, the pattern of results in Figure 3 is suggestive that employers selecting a high tier are not looking for harder work that would require more effort, but rather "smarter" work. A high tier selection is commonplace in highly skilled categories such as web and software development, whereas in categories like a support—which is largely data entry—the most common selection is low tier. Second, there is little empirical evidence for the notion that vertical preferences reflect higher employer expectations that might manifest in bad feedback if not met.

Among employers selecting a high tier in the ambiguous arm but not having their preferences revealed, there is no evidence that high tier employers are harsher evaluators. In Column (1) of Table 3, the outcome is the z-score of feedback (on a 1 to 5 point scale). Controls are included for the job category. The key independent variable are indicators for the employers (un-revealed) vertical preference—the sample is restricted to the ShownPref = 0 cell in the ambiguous arm. There is no evidence of systematically better or worse feedback scores by tier.

In Column (2), we report the same regression, but use the z-score of the employer's net promoter score (NPS) for the platform. Employers are randomly sampled to give a score, so the sample is smaller. Again, there is no evidence of a tier-related difference. In Columns (3) and (4), we still use the NPS measure but expand the sample. There is no overall effect of revelation on NPS, though there is some evidence of improved scores for employers that had medium- and high-vertical preferences revealed.

In addition to lack of empirical evidence that workers should "fear" high tier employers

Table 3: Measures of employer satisfaction by whether the firm's vertical preferences were revealed

	Dependent variable:			
	FB to worker (z)	Promotor score (z)		
	(1)	(2)	(3)	(4)
MedTier	0.031	−0.040	0.019	−0.039
	(0.027)	(0.043)	(0.025)	(0.041)
HighTier	0.0004	−0.072	0.019	−0.066
	(0.034)	(0.053)	(0.030)	(0.051)
ShownPref			0.028	−0.046
			(0.021)	(0.042)
MedTier x ShownPref				0.087*
				(0.050)
HighTier x ShownPref				0.128**
				(0.061)
Observations	8,441	3,482	10,432	10,432
R^2	0.018	0.027	0.017	0.017

Notes: This table reports regressions where the outcome variable is some measure of employer satisfaction afte the conclusion of a contract. The outcome in Column (1) is the feedback to the hired worker, normalized to a z-score (it is actually given on a 1 to 5 star scale). The outcome in the remaining columns is the normalized promotion score for the platform. Employers are not always asked for a promotor score at the conclusion of a contract, so it offers a smaller sample than the feedback sample. Significance indicators: †:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

because of increased expectations, there is little evidence that employers would justifiably think that paying higher wages would have anything but a selection effect: Gilchrist et al. (2016) shows via a field experiment in an online labor market that higher wages do not lead to greater measurable productivity. This is consistent with the relatively poor empirical support for persistent gift-exchange effects in labor settings (Gneezy and List, 2006).

E.3 Should workers consider changed applicant pool size when bidding?

As we saw, signal revelation had some effect on applicant pool size, particularly in the low tier. A natural question is whether these different pool sizes influenced wage bids. If workers thought they faced less competition, all else equal, they have an incentive to bid up. In settings where it can be examined, endogenous entry has proven empirically important (Bajari and Hortaçsu, 2003). However, in contrast to common value auctions, there is presumably a much greater role of idiosyncratic worker-specific surplus in the case of hiring, muting the effects. Whether this consideration is important in practice is an empirical question—the competition effects might be sufficiently small that the worker does not have to consider them from a worker's perspective. To test whether anticipated pool size "matters," we can test what workers do naturally, in the sense that we could consider how they adjust their bidding behavior on a job-to-job basis. Ideally we would estimate a regression of the form

$$\log w_{ij} = \alpha_i + \beta_1 \log A_j + \epsilon \tag{38}$$

where w is the individual wage bid of worker i to job opening j, α_i is an individual worker fixed effect and A_j is the number of applications opening j will receive, which is determined at random. Of course, in practice, A_j is very likely to be correlated with other factors that could affect the wage bid, such as how attractive or unattractive the job opening is to workers or how quickly a job opening is filled. However, there are factors that affect how many applications a job opening is likely to receive that is plausibly exogenous with respect to other opening characteristics, and so an instrumental variables approach is feasible.

To start, we ignore the endogeneity of A_j and simply estimate (38), reporting the results in Column (1) of Table 4. This regression uses the full set of applications to job openings in the ambiguous arm of the experiment. We can see that a larger applicant pool is associated with a lower wage bid—a worker bids about 0.36% less when facing a 10% larger applicant pool.

To account of the endogeneity in A_j, we construct an instrument. We use the mean log applicant pool size of other job openings in that same category, posted on that same day.[23] We include day-specific fixed effects in the second stage. The identifying assumption is that there is day-to-day variation in the number of jobs posted and the number of workers active that changes the number of applicants per job for exogenous reasons. In Column (2), we report the first stage of the IV estimate. We can see that is a powerful instrument, with a conditional F-statistic of 24156.93.

In Column (3) report the 2SLS estimate. We can see that the larger the pool, the lower the wage bid, with an effect size of -12.7%. As expected, when the applicant pool is larger for

[23] This is conceptually similar to the instrument used by Camerer et al. (1997).

Table 4: Effects of applicant pool size on individual wage bidding behavior

	Dependent variable:		
	Wage Bid	Log Apps	Wage Bid
	(1)	(2)	(3)
Log num apps	-0.036^{***}		
	(0.001)		
IV		0.780^{***}	
		(0.004)	
Log num apps (instrumented)			-0.127^{***}
			(0.003)
Worker FE	Y	Y	
Observations	583,492	583,303	583,303
R^2	0.919	0.555	0.915

Notes: This table reports regressions that explore the relationship between applicant pool size and individual wage bidding. In Column (1), the OLS estimate of log wage bids on log pool size is reported, with a worker-specific fixed effect. In Column (2), the first stage of an IV regression regression is reported, where the IV is the mean log number of applications received by job openings posted the same day, and in the same work category, as the "focal" job opening (but not including that opening). In Column (3), the second stage of the IV regression is reported. The sample consists of all applications to exeriment job openings that received at least two applications. Significance indicators: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

Table 5: Employer and job opening characteristics for filled job openings, by whether tier choice was shown, with job openings pooled from both arms of the experiment

	ShownPref=0	ShownPref=1	Δ	% Change
Employer attributes				
Prior job openings	5.50 (0.16)	5.77 (0.15)	0.26 (0.22)	4.80
Prior spend (log) by employers	7.22 (0.04)	7.21 (0.03)	-0.01 (0.05)	-0.19
Num prior workers	5.61 (0.17)	6.01 (0.17)	0.40 (0.24)	7.09
Job opening attributes				
Prefered experiance in hours	34.75 (1.39)	35.38 (1.28)	0.63 (1.90)	1.82
Estimated job duration in weeks	13.36 (0.22)	13.73 (0.20)	0.37 (0.29)	2.75
Job description length (characters)	572.52 (6.45)	563.74 (5.43)	-8.78 (8.37)	-1.53

Notes: This table reports means for a number of pre-randomization characteristics for the employer and job opening by ShownPref status. The data are pooled to include employers from both the ambiguous and explicit arms. Standard errors are reported next to the estimate, in parentheses. The far right column also reports the percentage change in the ShownPref = 1 group, relative to the mean in the control group. Significance indicators: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

plausibly exogenous reasons, a worker bids less. Despite being negative, the point estimate from Column (3) implies that that the equilibrium adjustment would be minuscule: for the low tier, where pool size dropped about 5%, workers would bid up by a bit more than 1/2 of 1%. The implication of these point estimates is that the change in bidding to perceived pool size—while in the expected theoretical direction—is relatively unimportant.

E.4 Selection on observables for filled openings

Table 5 compares the pre-randomization attributes of filled job openings, by ShownPref. The sample consists of all job openings pooled over the ambiguous and explicit arms of the experiment. Three is perhaps some slight evidence that more experienced employers were more likely to fill their job openings when their preferences were revealed, though the differences are not conventionally statistically significant.

E.5 Worker welfare

The overall effect of the signaling equilibrium on workers is challenging to estimate. For one, workers applied to both kinds of job openings, so it is not the case that we have treated and control workers whose outcomes we can compare. We can see, however, measure with applications had a higher expected value, on average, when they were sent to those employers whose preferences were shown. In Table 6, we report application level regressions in which the independent variable is the treatment assignment of the job opening. In Column (1), the outcome is an indicator for whether the worker was hired. In Column (2), the outcome is the indicator for whether the worker was hired times their wage bid, in levels.

Table 6: The effect of revelation on win probability and expected wage

	Dependent variable:	
	Hired	Expected wage
	(1)	(2)
ShownPref	0.001	0.006
	(0.001)	(0.009)
Worker FE	Y	Y
Observations	461,852	461,852
R^2	0.305	0.444

Notes: The unit of analysis is the individual job application. Significance indicators: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $* * *$:$p \leq 0.001$.

From Column (1), we see evidence of an increase in per-application win rate, which is consistent with the overall decline in the quantity of applications and no reduction in the probability a match was formed. This coefficient on the ShownPref indicator implies a 3.1% increase relative to the mean application success probability. In Column (2), the point estimate is positive, though fairly imprecise. At the mean value, this point estimate corresponds to a 2.2% increase. The average effect on workers was to increase in application success probability, leave the expected wage per-application the same or perhaps slightly higher.

E.6 Match outcome result robustness to sample definition

Figure 14 reports results for a number of outcomes using different sample definitions.

E.7 Tier choice over time

As employers can and do post multiple job openings during the experiment, we can observe if their tier choices change over time. Note that we only use the first observation for our experimental analysis. Table 7 reports estimates where the outcome is an indicator for a particular tier choice, and the key explanatory variable is the ordering of the opening, or OrderRank. The regressions show no change in probability of selecting low tier over time. However, there is some movement away from the medium tier, into the high tier. In Column (4), the order rank is interacted with the treatment assignment—there is no evidence that treatment assigned affected the choice over time.

This is obviously a short-run view, but it does show that there is no evidence that employers are experimenting with truthful revelation but then returning back to a "pooled" state after a bad experience. If anything, there appears to be less pooling over time.

Figure 14: Effects of revealing employer vertical preferences on job opening outcomes

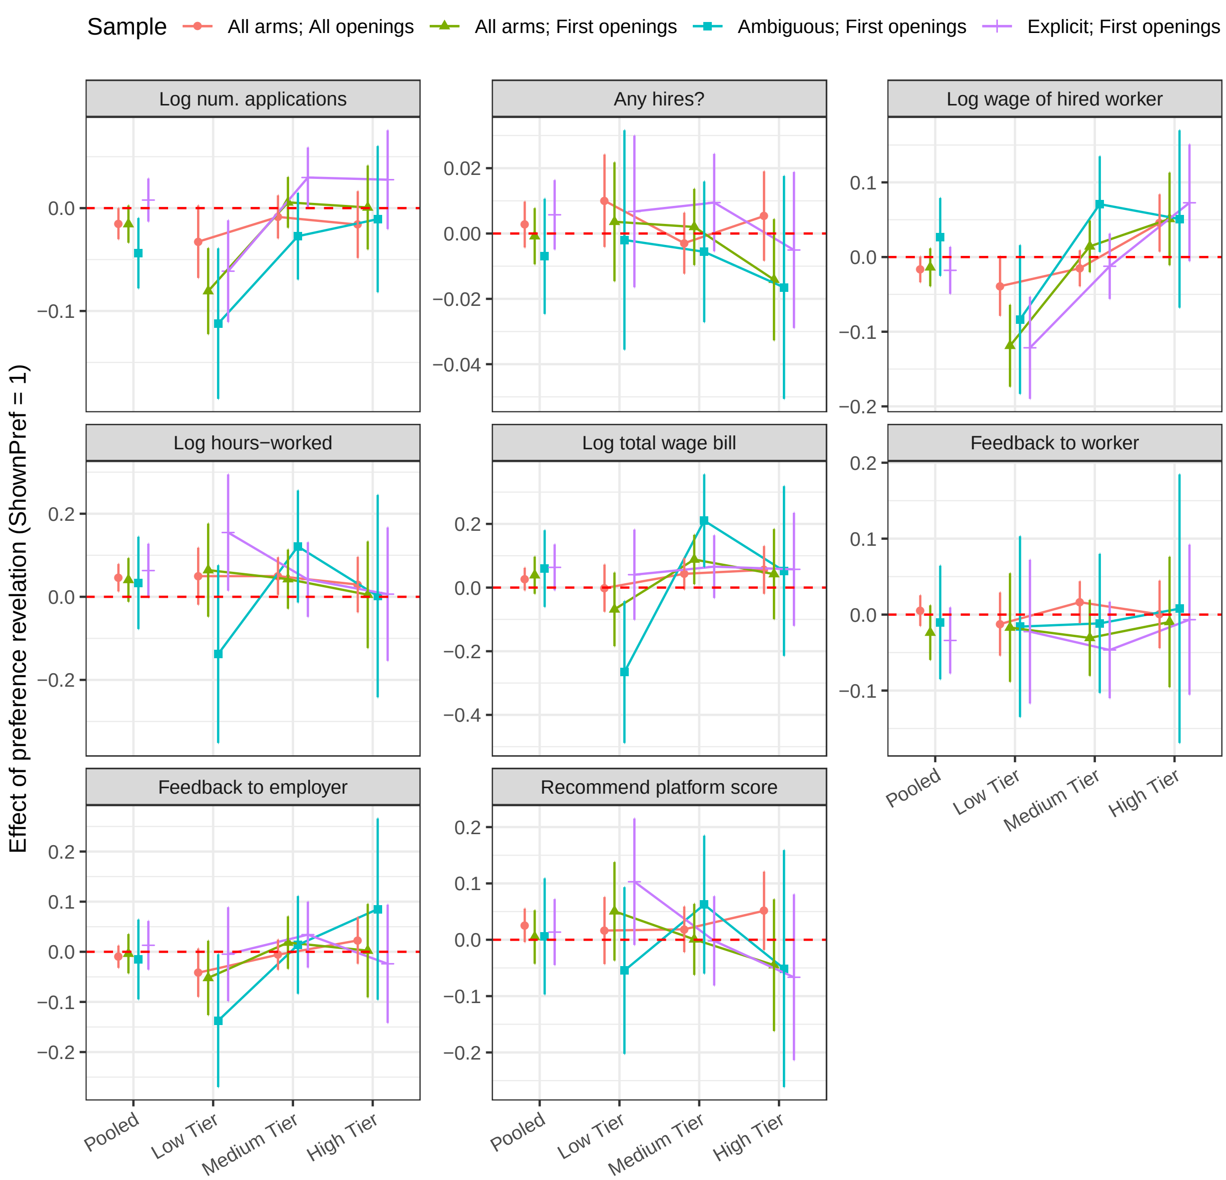

Notes: This figure shows the effects of revealing employer preferences, ShownPref = 1, on a number of outcomes, for several different samples. Each point estimate is surrounded by at a 95% CI.

Table 7: Employer vertical preference signal over time, by treatment assignemnt

	Dependent variable:			
	LowTier	MedTier	HighTier	
	(1)	(2)	(3)	(4)
OpeningRank	-0.001^{*}	-0.002^{***}	0.002^{***}	0.002^{***}
	(0.0003)	(0.0004)	(0.0003)	(0.0003)
ShownPref				-0.030
				(0.026)
ShownPref x OpeningRank				0.0004
				(0.001)
Observations	228,702	228,702	228,702	228,702
R^2	0.727	0.647	0.669	0.669

Notes: This table reports regressions where the dependent variable is an indicator for an employer's vertical preference selection and the independent variables are the chronological rank of the opening (ascending order) for that particular employer, OpeningRank, and its interactions with ShownPref. If ShownPref = 1, would-be applicants could observe the employer's vertical preference before applying, otherwise they could not, ShownPref = 0. The sample is restricted to employers assigned to the explicit arm that posted more than 1 but fewer than 10 openings. Employers in the two-cell explicit arm were told *ex ante* that the platform would reveal or would not reveal their vertical preferences to workers. In each regression, an employer-specific fixed-effect is included. Standard errors are clustered at the employer level. Significance indicators: †:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.