Cheap Talk Messages for Market Design: Theory and Evidence from a Labor Market with Directed Search*

John J. Horton — MIT & NBER

Ramesh Johari — Stanford University

Philipp Kircher — Cornell University

April 15, 2022

Abstract

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 which affects the number of applications, types of applications, and wage bids across firms. 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 to improve market outcomes.

*Author contact information available at http://www.john-joseph-horton.com/. Thanks to 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. Thanks to participants at the Marketplace Innovation Workshop at Stanford, the ACM-EC Conference, and the NBER Summer Institute for IT & Digitization. Thanks to Ada Yerkes Horton for help in preparing the manuscript. Slides: http://john-joseph-horton.com/papers/hot_towel.slides.html. Code (eventually): http://john-joseph-horton.com/papers/hot_towel.code.html.

1 Introduction

In this paper we study the facilitation of cheap talk in the sense of Crawford and Sobel (1982) as a tool to improve 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 empirically the usefulness of such cheap talk. 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 wages paid adjust to the messages that are posted. 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 exhibit a preference for workers with more experience (Pallais, 2014; Stanton and Thomas, 2015) and are therefore 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 to facilitate 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 apply to job openings and include a wage demand, and firms 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 for higher value jobs to gain 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 resolves worker uncertainty about precisely what kind of firm they face. We characterize when an informative cheap talk equilibrium is possible, showing that a necessary condition for an informative separating equilibrium is sufficiently heterogeneous preferences among employers.

We contrast the informative separating equilibrium to an equilibrium where cheap talk messages are not possible, but workers still obtain some information about what kind of firm they face once they decide to apply. When cheap talk signaling is possible, relatively high ability workers direct search only to high-type firms, leaving low-type firms without these higher ability applicants and thus fewer applicants overall. All workers adjust their wage bids: high-type firms "pay" for attracting high-ability workers as most workers raise their wage bids, while low-type firms are compensated for the lower application counts by workers lowering their wage bids. Note that the change in bids are conditional on worker type and are not just driven by ability sorting.

Our theory constructs a setting where cheap talk improves matching, but the theory also shows cheap talk is not always sustainable—there is always a "babbling" equilibrium where cheap talk messages convey no information. Whether an actual marketplace could benefit from introducing signaling therefore remains an empirical question. We approach this empirical question with our experiment.

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 towards 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 that experimental revelation of the cheap talk message induced substantial additional sorting of job-seekers. In aggregate, job-seekers avoided firms revealed to have a low willingness to pay, especially those with higher measured experience. Job-seekers bid up against high-type firms and bid down against low-type firms. This sorting and bidding strongly affected who was matched to whom, and at what price, without changing the quantity 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.

In terms of match outcomes, the revelation of employer preferences increased 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. In terms of subjective evaluations, there is marginally significant evidence that employers rated the platform more positively post-transaction.

One contribution of this paper is that we believe it is the first to explicitly explore the use of cheap talk messages to improve the search process in matching markets, underpinning theoretical insights with experimental evidence. Although several papers have considered the effect of cheap talk in matching markets, these all have focused on 1:1 matching scenarios in which participants are privately signaling their interest in a 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 who invested in 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 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 seeker.[1] 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.[2] As we can observe applications in our empirical setting, we can observe how much of 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 has 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 (2005b), 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. Although our empirical context is an online labor market, the

[1]E.g., Monster, Indeed, SimplyHired, LinkedIn, Craigslist, Facebook Jobs.

[2]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.

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 results the results of the experiment on the matching process and outcomes. We discuss future directions for research and conclude in Section 5.

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

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 larger number of employment options online; they have to select a small number (in the model, just one) to apply for work and include a wage demand. Employers 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 work ability, but is not codified and hard to search. Cheap talk is explicitly introduced by asking employers to choose one from 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 are collected in our key proposition (Proposition 4). 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.[3]

[3] 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 if 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: The economy consists of a mass of firms that try to hire, and a mass of workers looking for employment. There are two types of firms. Mass δ_L of firms has type $t = L$ and mass δ_H of firms has type $t = H$. This type is their private information. Workers do not observe the firm type, but see only a signal about the firm's type which reflects their prior knowledge about the firm and their inference from the specific features of the job description. The signal $s \in \{L, H\}$ of each firm matches its true type with probability $\psi \in (\underline{\psi}, 1)$. We will consider settings with relatively informative signals, i.e., with $\underline{\psi}$ large.

The measure of workers is normalized to unity. This is without loss of generality because only the ratio $\lambda \equiv 1/(\delta_L + \delta_H)$ of workers to firms matters. Workers differ in ability a drawn from distribution F with support on $[\underline{a}, \bar{a}] \subset \mathbb{R}_{++}$, with continuous strictly positive density $f(a)$. Ability is known and publicly observable.

Payoffs: Each firm can hire at most one worker, each worker can accept at most one job. A firm with valuation v that hires a worker of ability a produces va; and if it pays the worker b its profit π is

$$\pi = va - b \tag{1}$$

while the workers utility is b. Firms differ in their valuation: type t firms have valuation v_t. Throughout, we will consider only strictly positive valuations that are higher for high types ($v_H > v_L$) but that where it remains optimal for some workers to still visit low types (which turns out to imply the restrictions $e^{-1/\delta_H} v_H < v_L$). The utility of an unmatched firm or worker are normalized to zero.

Timing: First, if cheap talk is possible each firm first posts a message m, where without loss of generality we consider $m \in \{L, H\}$. When we study settings in which firms do not have the possibility to send a message we write $m = \emptyset$ and might omit the message indicator. Each firm also generates its signal $s \in \{L, H\}$.[4] Let $\iota = (m, s)$ denote the information about each firm. Workers then observe this information, and each worker chooses one firm to approach. Since there is no other information, workers can only choose whether to approach a high or low message firm with a high or a low signal, but then select among that group of firms at random.

Then the worker submits a wage bid $b_\iota(a)$ to its chosen firm. The bid conditions on his own type and the observable information about the firm. Each firm observes all the bids it receives, chooses the worker who delivers the highest payoff π, and pays this worker his wage bid. It can

[4] We will consider properties and existence of a fully revealing cheap-talk equilibrium, and for such analysis it does not matter whether the signal gets generated before or after the firm posted its message. For equilibrium that are not fully revealing this difference would matter.

reject all workers if it wants.

For the equilibrium analysis, let $\mu_m(v)$ be the equilibrium fraction of firms of valuation v that post message m. Let γ_ι denote the equilibrium fraction of workers that approach firms with information ι. For some of the analysis it will be useful to define

$$\lambda_{m,s} = \frac{\gamma_{m,s}}{\delta_m \mu_m(v_m)\Psi_{s,m} + \delta_{-m}\mu_m(v_{-m})\Psi_{s,-m}}$$

as the ratio of workers to firms with observable information $\iota = (m, s)$, where $-m = L$ if $m = H$ and vice versa, and where $\Psi_{s,m} = \psi$ if $s = m$ and $\Psi_{s,m} = 1 - \psi$ otherwise. Since workers apply to firms at random conditional on the information $\iota = (m, s)$, it is a well-known result from the directed search literature that the number of applicants at a firm with that information is Poisson distributed with parameter λ_ι.[5] Among workers that approach firms with this information, let $F_\iota(a)$ denote the fraction of workers that have type weakly below a. Denote its density by $f_\iota(a)$ whenever it exists. Within each message, workers lack further information and choose a firm at random (the usual anonymity assumption in directed search) and submit their bid according to their equilibrium bidding distribution $b_{\iota(}a)$. The equilibrium objects are $(\mu_m(v), \gamma_\iota, F_\iota(a), b_\iota(a))$ for all $\iota \in \{L, H\}^2$.

The economy where cheap talk is disabled is exactly analogous, except that firms can send no messages. That is, the setup without cheap talk is identical to the babbling equilibrium of the game with cheap talk where all firms send a message at random.

2.2 Analysis of truthful cheap talk in the market game

Assume cheap talk is enabled and firms truthfully reveal their types: $\mu_L(v_L) = \mu_H(v_H) = 1$. We will check at the very end when truthful revelation is incentive compatible. Truthful revelation implies that workers are certain about the type of the firm when they see the message, and the signal does not provide any additional information value. We therefore suppress the subscript s for the signal here, and denote by $\bar{\gamma}_m = \gamma_{m,s=L} + \gamma_{m,s=H}$ the visit strategy to firms of the same message independent of signal.[6] Given this, workers know the type of firms they face. Let

$$\pi_m(a) = v_m a - b_m(a) \text{ for } m \in \{L, H\}$$

denote the utility that the firm obtains when it hires a worker of type a in equilibrium.

[5] This is usually established by taking a market with a finite number of workers and firms with a worker-firm ratio of λ_m, and then by increasing the market size to infinity while keeping that ratio constant. See, e.g., Burdett et al. (2001), Peters (1997) or Wright et al. (2021).

[6] In principle the signal could serve as a coordination device in this market, but it is easy to show that in this economy the market tightness neither the $\lambda_{m,s}$ nor the bids $bm, s(a)$ vary in s in equilibrium.

Denote by D_m the equilibrium distribution of profits π_m at message m. That is, $D_m(\pi) = \int_{a:v_m a - b_m \leq \pi} dF_m$. In analogy to standard results in the auction literature, D_m cannot have a mass point: if there was a mass point then those workers would tie and have a strictly positive probability of losing the job to someone who bids exactly the same, while a tiny reduction in the wage demand would have negligible costs to the worker but he would always obtain the job in such circumstances.[7]

Now consider a worker of type a who leaves profit π to the firm with message m. He wins if no other worker at this firm leaves a higher profit. There are in total $\lambda_m(1 - D_m(\pi))$ such workers at message m, and divided by the mass of firms δ_m this yields the queue length of more profitable workers. Bidders are Poisson distributed across auctions, where the Poisson parameter is the queue length. So the chance that no better worker arrives at the firm is $e^{-\lambda_m(1-D_m(\pi))}$. This is the winning probability for this worker. To generate π, he had to bid $b = v_m a - \pi$, and therefore his total expected utility is

$$e^{-\lambda_m(1-D_m(\pi))}(v_m a - \pi). \tag{2}$$

Note that this is supermodular in (a, π). Therefore, if $\pi_m(a)$ is the optimal profit that worker type a wants to leave, then it is optimal for lower worker types to leave less and for higher worker types to leave more. Since D_m has no mass points, $\pi_m(a)$ is strictly increasing among those types that actually approach message m. We can therefore invert π_m on the support of F_m to find the worker type who left this amount of profit, and call the inverse α_m. The distribution of profits equals the distribution of types that generates these profits: $D_m(\pi) = F_m(\alpha_m(\pi))$. We can then rewrite the equilibrium utility (2) directly in terms of the bids as

$$e^{-\lambda_m(1-F_m(\alpha_m(v_m a - b)))}b. \tag{3}$$

In a truthful equilibrium workers choose where to go and what to bid to maximize (3), taking equilibrium objects $\{\bar{\gamma}_L, \bar{\gamma}_H, F_L(a), F_H(a), b_L(a), b_H(a)\}$ as well as truth-telling ($\mu_m(v_m) = 1$) as given.

We will now establish in sequence some key properties of the equilibrium: A) Indifference of the lowest worker type between both messages. B) Indifference of all low worker types between messages, implying that all types up to some cutoff $\hat{a}$ are indifferent between both messages. C) Random visits of low worker types, meaning that the queue length of types below $\hat{a}$ at both

[7] This argument does not apply to workers who bid zero as they cannot cut their wage lower. But bidding zero never happens in equilibrium as the worker obtains nothing, while he could obtain a strictly positive utility from demanding, e.g., $v_m a/2$ which is clearly accepted if the worker is the only bidder, which happens with strictly positive probability since bidders are Poisson distributed.

messages is as if these workers choose firms at random. D) High worker types like high messages, so types above $\hat{a}$ only visit those firms. We will show these properties in turn.

Indifference of the lowest worker type: If any workers visit the low message at all, let $\underline{a}_m$ be the lowest type that visits message m (i.e., the infimum of the support of F_m). Since in equilibrium the lowest type only wins if no higher type is present, this types demands all surplus $b = v_m \underline{a}_m$ and wins with probability $e^{-\lambda_m}$. He could have visited the firms of the other message and demanded all surplus there, but chooses not to, implying:

$$e^{-\lambda_L} v_L \underline{a}_L \geq e^{-\lambda_H} v_H \underline{a}_L,$$
$$e^{-\lambda_L} v_L \underline{a}_H \leq e^{-\lambda_H} v_H \underline{a}_H,$$

which readily establishes

$$e^{-\lambda_L} v_L = e^{-\lambda_H} v_H. \tag{4}$$

Clearly the lowest type on the population will approach one of the two messages, so $\underline{a} = \underline{a}_m$ for some message m, and from (4) he is indifferent between both messages.

Indifference of all low worker types: Now assume some worker type is indifferent between both messages. We will show that all types below will be indifferent as well, and proceed by proof by contradiction. If not, there would have to be types a' and $a'' > a'$ that are indifferent between both messages but all types inbetween strictly prefer message m over message m'. Therefore, the types inbetween will only approach message m. The marginal utility at m improves according to the envelope theorem applied to (3) by

$$U'_m(a) = e^{-\lambda_m(1-F_m(a))} v_m \tag{5}$$

The marginal utility at m' for $a \in [a', a'']$ improves according to $U'_{m'}(a) = e^{-\lambda_{m'}(1-F_{m'}(a'))} v_{m'}$ since no types above a' joins there and so the competition for the job is determined by $F_{m'}(a')$). Recall that a' and a'' are indifferent between messages, so intermediate types only prefer m' if their utility there improves more around a', and they return to indifference only if their utility there falls around a'':

$$e^{-\lambda'_m(1-F_{m'}(a'))} v_{m'} \geq e^{-\lambda_m(1-F_m(a'))} v_m,$$
$$e^{-\lambda_{m'}(1-F_{m'}(a'))} v_{m'} \leq e^{-\lambda_m(1-F_m(a''))} v_m,$$

which yields a contradiction because $F_m(a') < F_m(a'')$ as workers in $[a', a'']$ do approach message m.

Random visits by all low worker types: If worker types a below $\hat{a}$ choose among firms at random this implies that

$$\lambda_L F_m(a) = \lambda_H F_h(a), \tag{6}$$

i.e., the queue of workers with types weakly below a is equalized at all firms. Any equilibrium needs to have that feature. We know that the lowest worker type is indifferent between messages. For other low types to be indifferent, the marginal utility increase at message m according to (5) has to equal the corresponding expression at m'. Applying (4) and cancelling terms yields (6).

High worker types prefer the high message: Worker types above $\hat{a}$ all strictly prefer the same message m. This message has to be the high one $m = H$. Otherwise, there would be indifference at $\hat{a}$ and the utility rise $e^{-\lambda_m(1-F_L(a))}v_L$ at the low message would need to strictly exceed the marginal gain $e^{-\lambda_m(1-F_H(\hat{a}))}v_H$ at the high message for all $a \in (\hat{a}, \hat{a}+\epsilon)$ for some strictly positive ϵ. Again substituting (6) leads to a contradiction as $\lambda_L F_L(a)) > \lambda_H F_H(a))$ when types up to $\hat{a}$ choose firms at random and above they only choose firms with message L.

With these insights on where workers bid, we can now characterize the queue length at each message, and we can differentiate the bidding function (3) along the equilibrium path to characterize the bids: The bid distributions follow a similar path both at low and high message firms, only scaled up by the productivity differences. Since bidding is revenue-equivalent to a second price auction, workers choose to bid where their social contribution to surplus is highest conditional on other worker's strategies, which is exactly the condition for constrained efficiency in this market derived in Shimer (2005a).[8] These insights are summarized in more detail in the following proposition:

Proposition 1. *Assume firms reveal their type truthfully so that message $m = H$ is sent only by v_H firms and $m = L$ only by v_L firms. The remaining equilibrium conditions imply a unique market outcome, in which all workers with types below some cutoff $\hat{a} \in [\underline{a}, \bar{a})$ choose among all firms at random, while workers with type above $\hat{a}$ only approach firms with the high message. That means that the expected number of applicants with types below $a \in [\underline{a}, \hat{a})$ has the form*

$$\lambda_L F_L(a) = \lambda_H F_H(a) = \lambda F(a),$$

where the queue length at the firms is

$$\lambda_L = \frac{1 - \delta_H \ln(v_H/v_L)}{\delta_H + \delta_L}, \;\; \lambda_H = \frac{1 + \delta_L \ln(v_H/v_L)}{\delta_H + \delta_L} \tag{7}$$

[8] Constrained efficiency means that the planner cannot increase output if he can only change the strategies of the agents, but cannot avoid the matching frictions that arise when agents apply randomly for a given message as embedded in the Poisson matching process.

and the number of workers that queue for each message is $\gamma_m = \delta_m \lambda_m$. *The interval* $[\underline{a}, \hat{a})$ *is non-empty whenever* $1 - \delta_H \ln(v_H/v_L) > 0$. *For* $a \in [\hat{a}, \bar{a}]$ *it holds that* $F_L(a) = 1$ *and* $F_H(a) = 1 - (1 - F(a))/\lambda_H$. *The bid distribution is given by*

$$b_m(a) = v_m g_m(a) \tag{8}$$

independent of signal s, *where* $g_m(a)$ *is uniquely characterized by differential equation*

$$\lambda_m f_m(a) g_m(a) = 1 - g'_m(a).$$

with boundary condition $g_m(\underline{a}) = \underline{a}$. *In equilibrium the market outcome is constraint efficient. Bids are increasing at* a *if the type distribution is sufficiently dispersed, where* $f(a) \leq \delta_H/a$ *is a sufficient condition. For* $a \leq \hat{a}$ *it holds that* $g_L(a) = g_H(a) =: g(a)$, *i.e., workers extract an equal fraction of the surplus at either firm type.*

See Appendix A for proof of Proposition 1.

We are now equipped to consider the truth-telling behavior of firms. We focus on the parameter range where some workers choose message L after truthtelling (i.e., $[\underline{a}, \hat{a})$ non-empty) as our empirical estimates indicate that this is the relevant case. We obtain that truthtelling is an equilibrium when v_H is substantially larger than v_L, while it is generically not when v_H becomes arbitrarily close to v_L. Note that v_H large and workers visiting both messages (conditional on truthtelling) can only happen when δ_H is small, so conditions on δ_H are part of the proposition.

Proposition 2. *Fix* v_L, δ_L *and* $F(a)$. *Truthtelling is incentive compatible when* δ_H *is sufficiently small, and conditional on this,* v_H *is large. Truthtelling is generically not incentive compatible when* v_H *becomes arbitrarily close to* v_L.

A direct consequence of Proposition 2 is that a platform designer cannot use truthful cheap talk to separate valuations that are close, but is able to do so for far apart valuations. While it goes beyond this exposition, one can construct examples with a third type such that each type reports truthfully, again if each type is far enough from the other. What is not possible, by the same argument as in Proposition 2, is to have types that are very close together all reporting truthfully. As we will see in the empirical setting, the experiment only allows for far-apart expressions, and we find indications that messages are truthful.[9]

[9]With many firm types in $[v_L, v_H]$ full revelation is not possible by arguments as in Proposition 2. This implies only a bounded number of informative messages, as is common in the cheap talk literature since Crawford and Sobel (1982). A full analysis goes beyond this paper, but one can easily add a zero measure of such firms to the truthful revelation equilibrium above. Since these do not affect worker behavior, it is easy to show that higher

We now analyze the setting where cheap talk is disabled, i.e., where no information is transmitted in the first stage. Finally, we compare the predictions from both settings.

2.3 Analysis of the market without cheap talk

Now consider a world without the possibility of sending messages, so we can omit messages in the notation. The only information about a firm is its signal $s \in \{L, H\}$ that reveals the true type with probability ψ. Using Bayes' rule workers assign probabilities $\Psi_{t=0H,s=H}$ and $\Psi_{t=H,s=L}$ to the high type firms according to

$$\Psi_{H,s=H} = \frac{\psi\delta_H}{\psi\delta_H + (1-\psi)\delta_L}, \Psi_{H,s=L} = \frac{(1-\psi)\delta_H}{(1-\psi)\delta_H + \psi\delta_L}, \tag{9}$$

and low value firms have complementary probabilities $\Psi_{L,s=H} = 1 - \Psi_{H,s=H}$ and $\Psi_{L,s=L} = 1 - \Psi_{H,s=L}$. To compare this environment with the one of truthful cheap-talk, we proceed within the parameter restrictions that allow truthful cheap-talk in the previous section: low δ_H and high v_H. Given these, we assume a sufficiently precise but not perfect signal.

Approaching a firm with signal s is similar to approaching a firm with a truthful message in the previous section, only that now there is some uncertainty whether the firm type is truly equal to the signal. Interestingly the equilibrium structure remains very similar: low type workers still approach firms at random, while high type workers only approach high type firms. At the high signals the bids exceed the valuation of low types firms, who reject such offers, but differential equation that characterizes wage offers is identical to the setting where workers know the firm type for sure. At low messages the wage bids are targeted mainly to low firm types, but high types would also accept those, and the bids partially reflects these higher valuations. To summarize:

Proposition 3. *Fix v_L, δ_L and $F(\cdot)$, and consider δ_H sufficiently small, and conditional on this, v_H sufficiently large. For a high-precision signal it holds that*

Given all other parameters, there exists $\underline{v}$ and $\underline{\psi}(v_H)$ such for $v_H \geq \underline{v}$ and $\psi \geq \underline{\psi}(v_H)$ the following uniquely characterizes the equilibrium in the absence of cheap talk (or if firms babble with $\mu_L(v_L) = \mu_H(v_H)$): The queue length at all vacancies is $\lambda = 1/(\delta_L+\delta_H)$, the type distribution

types send (weakly) higher messages to attract more and better workers. These firms do not strictly speaking reveal their type, but when separately asked to choose between low-ability-low-wage-workers or higher-ability-higher-wage-workers again higher type firms choose weakly higher ability, akin to the results in the empirical setup. Empirical findings would be very different in uninformative (babbling) equilibria: more productive firms continue to prefer higher-ability-higher-wage workers, but here send messages at random.

is $F(a)$, and the equilibrium bidding strategy is determined by differential equations

$$\frac{\lambda f(a)}{v_H - b'_{s=H}(a)} b_{s=H}(a) = 1 \text{ with } b_{s=H}(\underline{a}) = v_H \underline{a}, \tag{10}$$

$$\left[\sum_{t \in \{L,H\}} \frac{\lambda f(a) \Psi_{t,s=L}}{v_t - b'_{s=L}(a)} \right] b_{s=L}(a) = 1 \text{ with } b_{s=L}(\underline{a}) = v_L \underline{a}. \tag{11}$$

See Appendix A for proof of Proposition 3.

What makes this proposition attractive is the ease at which it can be compared to the previous setting with cheap talk: the differential equation after the high signal (10) coincides with the differential equation at high messages in the truth-telling equilibrium in (8) for $m = H$ on $[\underline{a}, \hat{a})$. Here the signal is not perfect and the worker is not sure at a high signal whether the firm is truly high productivity, but he still submits a bid that is so high that only high firms accept and so the results coincide. That does not mean that truly high firms always get such bids, as they sometimes send low signals. After a low signal, workers now bid in a way that is acceptable both to the low-productivity firms which are very likely as well as to the high-productivity firms which are unlikely. So conditional on bid b, the differential equation (11) generates a derivative $b'(a)$ that lies between the derivatives given in the truth-telling equilibrium at the low and high message characterized by (8). This similarity allows us to analyze relatively transparently in the next section what differences we might expect between a setting with cheap talk and without.

2.4 Comparing a market with truthful cheap talk with a market without cheap talk

In the actual labor market experiment, we observe messages for firms where the message is shown to job-seekers (cheap talk), as well as in settings where they are not shown (no cheap talk, at least on this dimension). As we elicit messages even in the second setting, we can condition on them in our analysis. Under truth-telling this equates to conditioning on firm types. The following comparison of our stylized cheap talk environment with our stylized no-cheap talk environment similarly conditions on firm type.

Consider a market where parameters are as in Proposition 3, and truth-telling is an equilibrium as according to Proposition 2. That is, δ_H sufficiently small, v_H sufficiently large, and ψ sufficiently informative. We compare the environment with truthful cheap talk according to Proposition 1, which we sometimes simply call the truth-telling equilibrium, with the environment without cheap talk in Proposition 3.

In the truth-telling equilibrium worker types below cutoff ability $\hat{a}$ approach all firms at random. This is also true without cheap talk. With truthful cheap talk only the types above $\hat{a}$ are selective and only approach high type firms, while without cheap talk these types also visit firms at random. How pronounced the selection is depends on the level of $\hat{a}$. Since firms always hire the best worker that approaches them, improvements (deterioration) in the quality of applicants leads to improvements (deterioration) in the quality of hired workers.

Conditional on the visit of a worker of type a, we can compare the bids that firms receive. The main difference in bidding behavior is that workers effectively have more information about the type of firm they are facing when they receive also the messages rather than only signals. Consider first a high type firm. Under babbling, it receives bid $b_{s=H}(a)$ with probability ψ from the workers and $b_{s=L}(a)$ with complementary probability $1-\psi$. Under truth-telling it receives bid $b_{m=H}(a)$ for sure. For worker types below the cutoff $\hat{a}$ we know that the distribution of these types is unchanged, but for the bids we have $b_{m=H}(a) = b_{s=H}(a) > b_{s=L}(a)$. The first equality follow trivially from the equivalence of the differential equations. The inequality follows because at the lowest type $b_{s=L}(\underline{a}) = v_L\underline{a} < v_H\underline{a} = b_{m=H}(\underline{a})$, and if at any higher type $a > \hat{a}$ the bids should approach one another in the sense that $b_{s=L}(\underline{a}) \approx b_{m=H}(\underline{a})$, comparison of the differential equations readily reveals that $b'_{s=L}(\underline{a}) < b'_{m=H}(\underline{a})$, so that equality of these bids can never be achieved. Since the firm faces a mix between $b_{s=H}(a)$ and $b_{s=L}(a)$ under babbling, the expected bids per applicant are lower.

For high types firms that face a worker of type $a > \hat{a}$ we know that they are more likely to face such a type under truth-telling than babbling, as truth-telling implies that all high-ability workers will go only to high-value firms. The expected wage bid of a worker with ability slightly above $\hat{a}$ is higher under truthful cheap talk then without it. This arises because in the absence of cheap talk the bid is a weighted average of $b_{s=H}(a)$ and $b_{s=L}(a)$. We have seen in the previous paragraph that $b_{m=H}(\hat{a}) = b_{s=H}(\hat{a}) > b_{s=L}(\hat{a})$, so at the cutoff the average bid under babbling is strictly lower. By continuity this also holds for types that are higher. It is possible, though, that for very high worker types this no longer holds. The reason is that under cheap talk all high types compete for jobs at high value firms, and this increase in competition can decrease the wage bid.

For low type firms, we know that the expected number of applicants with types below any $a \in (\underline{a}, \hat{a})$ is identical under truthful cheap talk and without cheap talk. But any such type bids more aggressively in the absence of cheap talk under either signal: $b_{m=L}(a) < b_{s=L}(a) < b_{s=H}(a)$. The second inequality follows as in the previous paragraph. The first inequality follows because bids are equal at the lowest type, and again the bids increase faster under a low cheap talk message $m = L$ than under a low signal $s = L$ (i.e., $b'_{m=L}(a) > b'_{s=L}(a)$) whenever the bids are roughly similar (i.e., $b_{m=L}(a) \approx b_{s=L}(a)$).

Finally, the cheap talk equilibrium with truthful type revelation is constrained efficient, which means that a planner that knows the types of the agents cannot improve output by changing workers' strategies. Since such a planner could choose the assignment of workers as in the absence of cheap talk but chooses not to, the cheap talk equilibrium improves market outcomes.

The following proposition summarizes these findings:

Proposition 4. *Consider a market with parameters such that Proposition 3 applies, and truth-telling is an equilibrium in the cheap talk market according to Proposition 2. The market with truthful cheap talk compares to the market without cheap talk as follows:*

1. *Ability sorting: truthful cheap talk induces more assortative matching. For high value firms the average quality of the applicant increases under truthful cheap talk and therefore also the average quality of a hire, while for low value firms the average quality of applicants and hires decreases. This is driven through changes at the top: workers with high ability ($a > \hat{a}$) choose high type firms more often under truthful cheap talk, while workers with low ability ($a < \hat{a}$) choose high types equally likely in either setting.*

2. *Number of applications: The number of applications at low value firms decreases, while the number of applications at high value firms increases under truthful cheap talk.*

3. *Wage bidding: Any given worker type bids less at low value firms under truthful cheap talk. Any given worker type a bids more at high value firms under truthful cheap talk, except possibly those with the highest ability levels (i.e., bids increase for any ability $a < a'$ for some cutoff $a' > \hat{a}$, where $a' = \bar{a}$ is possible).*

4. *Market outcome: efficiency is higher under truthful cheap talk.*

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 sections return to the experimental environment, study whether truthfulness seems to characterize the market interaction, and investigate empirical analogues to the predictions we have just derived.

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 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.

Figure 1: Vertical preference signaling opportunity presented to employers when posting their job openings

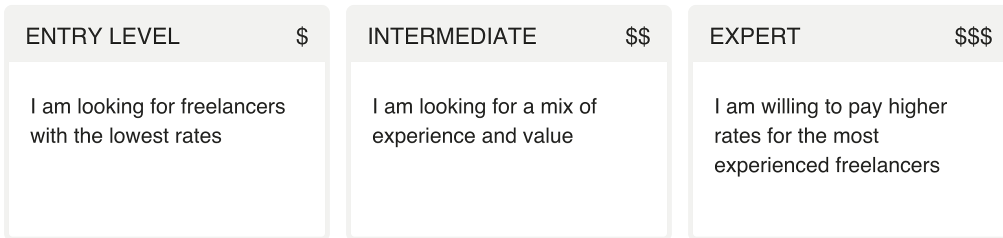

Notes: This figure shows the vertical preference signaling interface and language presented to employers when they posted a job.

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 1. 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)).

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" 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.[10] 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

[10] 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.

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 1 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 allocated to the experiment. These employers collectively posted 220,510 job openings.[11] 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, 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

[11] 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.

vertical preference in the interface shown to workers. The labeling was prominently displayed to make it salient to applying workers.[12] Randomization was effective—see Appendix B.

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, 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

Proposition 2 showed that truthful cheap talk can be sustained in equilibrium in some circumstances, but that this is not always possible. 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

[12] 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.

of the market. If they are not telling the truth, their distribution should be tilted relative to the previous group.

Figure 2 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 2 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, 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 2 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.

Figure 2: 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

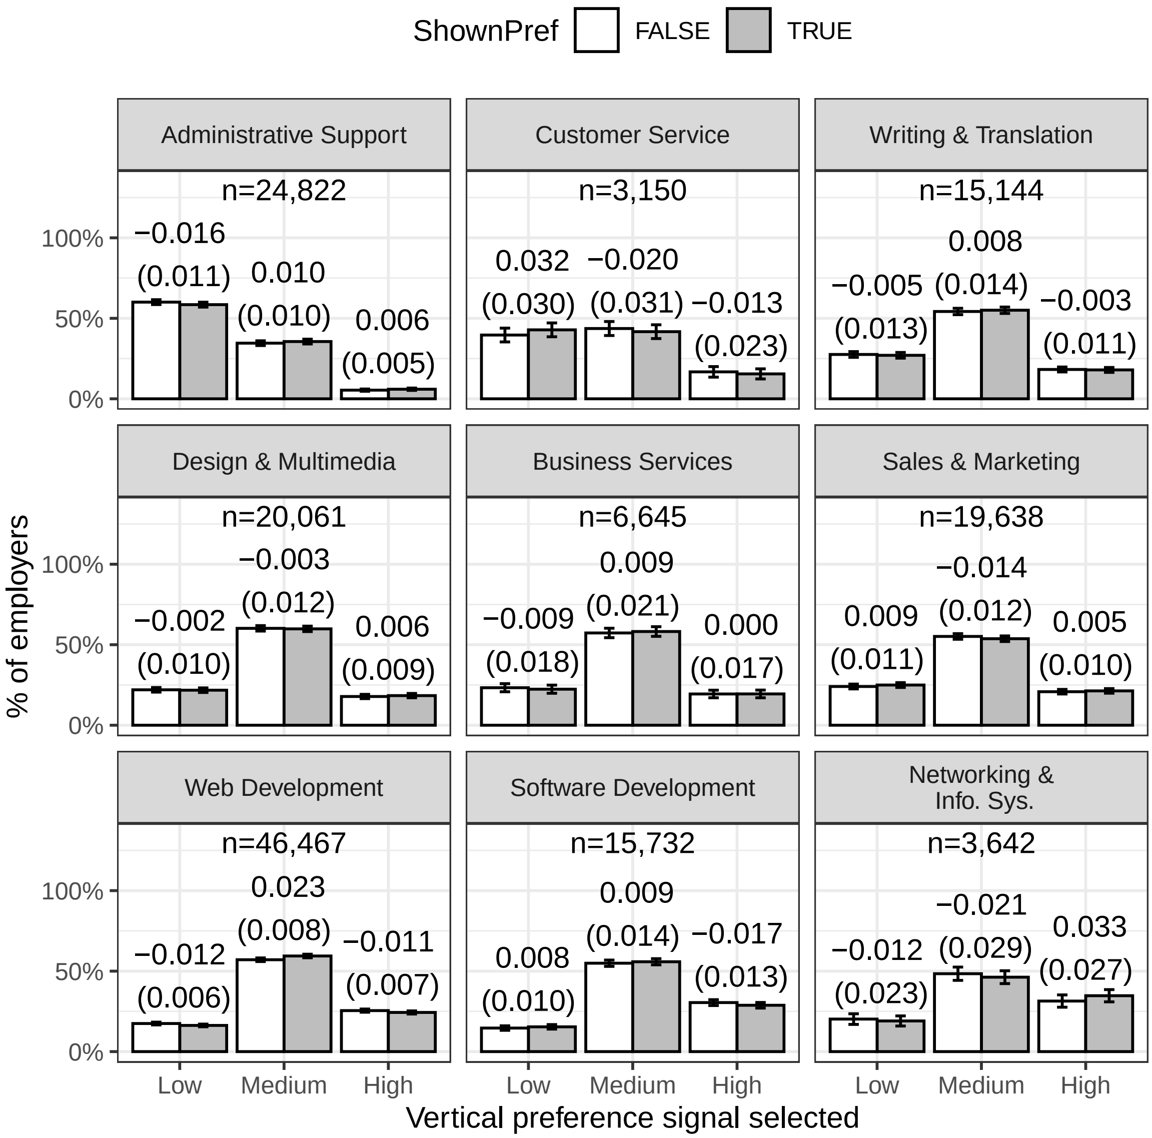

Notes: This figure shows the fraction of employers in the explicit arm of the experiment selecting the various vertical preference tiers, by ShownPref, and by the category of work of the associated job opening. When posting a job opening, employers had to select from one of three "tiers" to describe 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.

4.2 Sorting

The first implication of Proposition 4 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. 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

$$\log y_{ij} = \sum_k \beta_k^s \cdot \textsc{Tier}_{kj} + \epsilon_j \Big| s = \textsc{ShownPref}_j, \tag{12}$$

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 (12) are plotted in Figure 3, illustrating the sorting of job-seekers with and without access to the message. The left panel of Figure 3 plots both sets of $\hat{\beta}_k^s$ coefficients from (12) 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.[13] 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.

[13]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 $\textsc{ShownPref}$.

Figure 3: Comparison of applying and hired worker experience by employer vertical preference tier and message revelation in the ambiguous arm

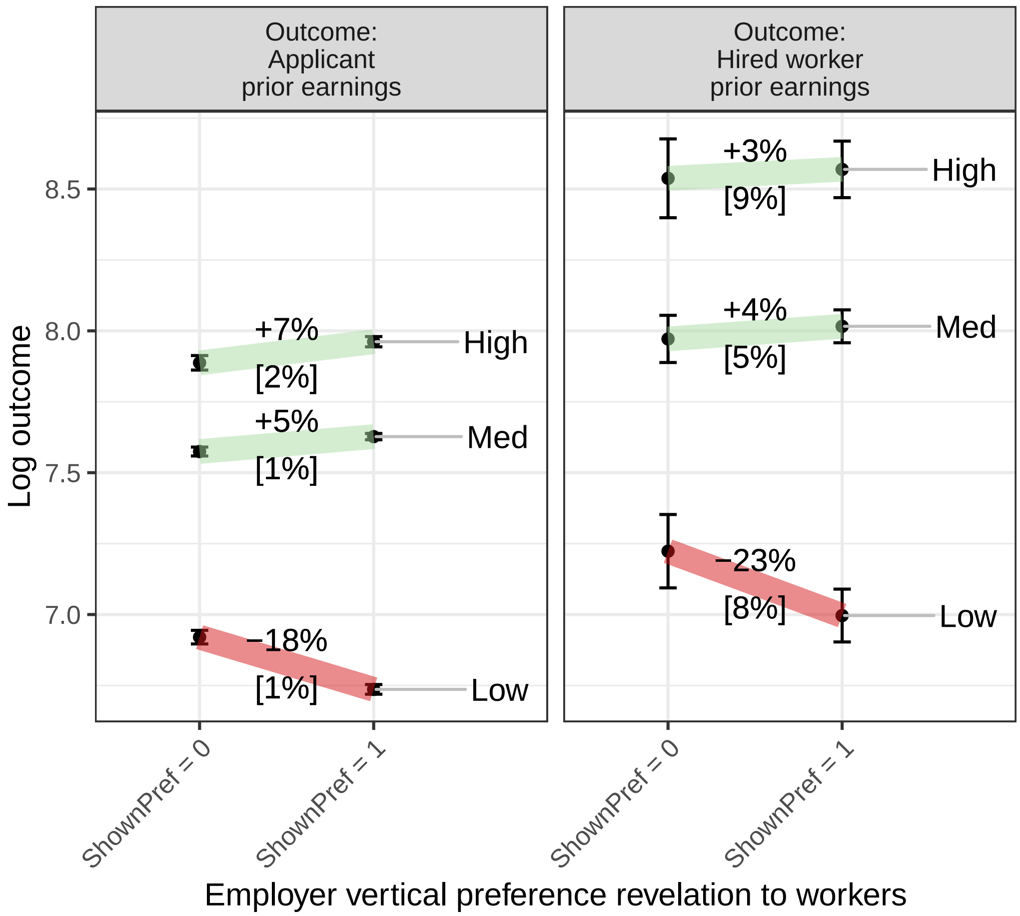

Notes: This figure plots coefficients from estimates of (12). The outcome in both panels is the applicant's cumulative prior hourly earnings at the time of application. The sample in the left panel is all applicants in the ambiguous arm of the experiment, whereas in the right panel, the sample is hired 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.

If we look at the $\textsc{ShownPref} = 1$ coefficients, $\hat{\beta}^1_k$, we can see that revelation increases sorting in the expected directions according to Proposition 4: 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%.

Figure 3 shows the average effects of sorting on applicant composition, but Proposition 4 predicts that applicant pool differences should be driven by relatively high ability workers. In the model, it is relatively high ability workers that sort more strongly towards high type firms, whereas low type applicants show no reaction to the message with respect to where they apply. That is, low type employers who have their message revealed lose relatively able applicants, but high type employers still get relatively low ability applicants even when their message is revealed.

In order to investigate this prediction, we can examine sorting based on quantiles of the distribution of applicant experience, by job opening. Assuming lower ability workers populate the lower quantiles of the applicant pool, then Proposition 4 predicts that revelation should have little effect. Note that we are not proposing a quantile regression, but rather comparing the mean quantile. As an example of how this measure is constructed, if we had three job openings, and the 10th percentile of past workers earnings for each was \$100, \$200, \$0, then the average 10th percentile would be \$300. Using this method allows us to create opening-level measures, which we can average by group and estimate treatment effects as simple means comparisons. We do this in Figure 4, which plots the effects of revelation on the experience of applicants, as 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 3. For the hours-worked measure, we see a roughly similar pattern: separation at higher quantiles but not at lower quantiles. At low quantiles for both measures, there is not evidence revelation affected composition: low-type workers seem to have ignored the message with respect to what job openings they approached. This is a surprising confirmation of the theoretical prediction of the model that it is high-types that do the sorting.

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.

Figure 4: Effects of showing employer vertical preferences on applicant pool composition with respect to experience, by tier quantile

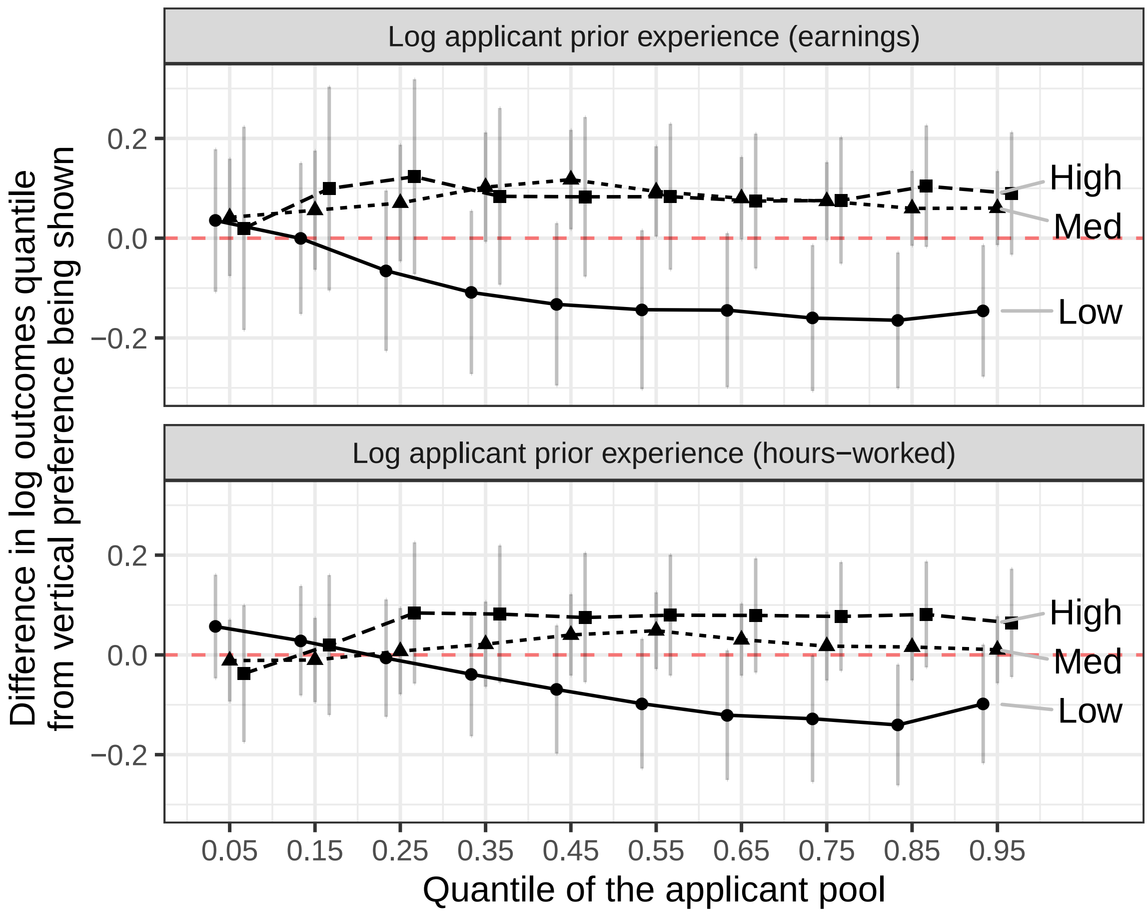

Notes: This figure shows the effects of employer vertical preference revelation on the composition of applicant pools in the ambiguous arm of the experiment for worker experience at the time of their application. The top panel is for cumulative earnings, while the bottom panel is for hours-worked. Workers 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*.

To test whether revelation affects the characteristics of hired workers, we estimate (12), but with the sample restricted to hired workers. We return to Figure 3, 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.

In the right panel of Figure 3, 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

Proposition 4 predicts that low type firms experience a decrease in applicant counts, while high type firms experience an increase. To test these predictions we estimate

$$\begin{aligned}\log A_j = {} & \beta_0 + \beta_1 \textsc{ShownPref}_j + \beta_2 \textsc{MedTier}_j + \beta_3 \textsc{HighTier}_j + \\ & \beta_4 \left(\textsc{MedTier}_j \times \textsc{ShownPref}_j\right) + \\ & \beta_5 \left(\textsc{HighTier}_j \times \textsc{ShownPref}_j\right) + \epsilon, \end{aligned} \tag{13}$$

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 5, 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

Figure 5: Effect of employer vertical preference revelation on the size of the applicant pool

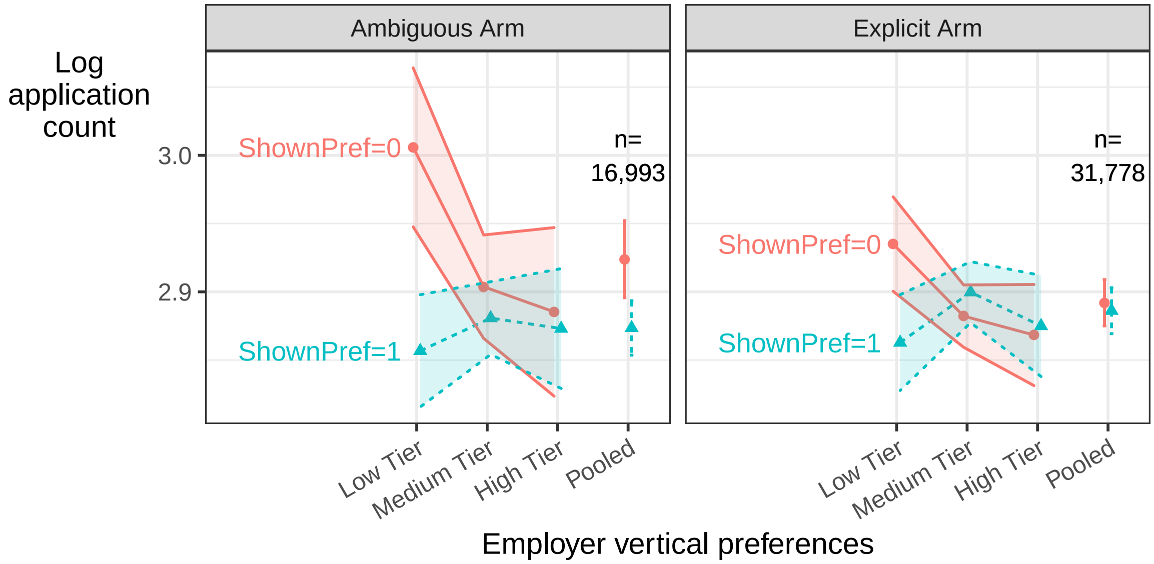

Notes: This figure reports regression results where the outcome is the log number of applications received by that opening. The right panel uses job openings from the explicit arm, whereas the left panel uses openings from the ambiguous arm. The samples are restricted to job openings receiving at least 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.

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 dropped differs by ShownPref status.[14] 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 (14). Note that the x-axis shows the employer tier selection, ordered from low to high; the colored lines correspond to the ShownPref value.

We can see from Figure 5 that in both arms, applicant pool reductions are concentrated in the low tier, with message revelation having little discernible effect in the other tiers. Note that this result is only partially consistent with Proposition 4, bullet point 2: we get a reduction in application counts for the low-type firms, but no discernible increase applications to the high-types.

In the model, there is no entry or exit: every worker sends a single application. In the actual marketplace, we have both entry and exit on both sides of the market. Job-seekers also decide

[14] 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.

how many applications to send, and so we have endogenous application quantities. To the extent message revelation changed the returns to sending an application, we might expect changes in the number of applications sent. However, it is theoretically ambiguous whether job-seekers should send more or fewer applications following a increase in win probability—on the one hand, it takes fewer applications to get a job, but on the other hand, each application is not more valuable in expectation (Shimer, 2004).

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:

$$\log A_j = \beta_0 + \beta_1 \textsc{ShownPref}_j + \epsilon, \tag{14}$$

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 5, 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. Proposition 4 predicts an increases in wage bids relative to what workers would usually bid without message revelation at high type firms, and the opposite at low type firms. To test this prediction, we again estimate (12) but with the log wage bid as the outcome. In the first panel from the left of Figure 6—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.

Observing a change in wage bids by revelation is expected given the compositional changes caused by message revelation (recall Figure 3): high tier employers who have their message revealed should receive higher wage bids because the workers who apply to those employers are more experienced. 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 the prediction of Proposition 4, bullet 3.

Figure 6: Comparison of applicant mean log wage bids and profile rates by employer vertical preference tier and revelation of the signal

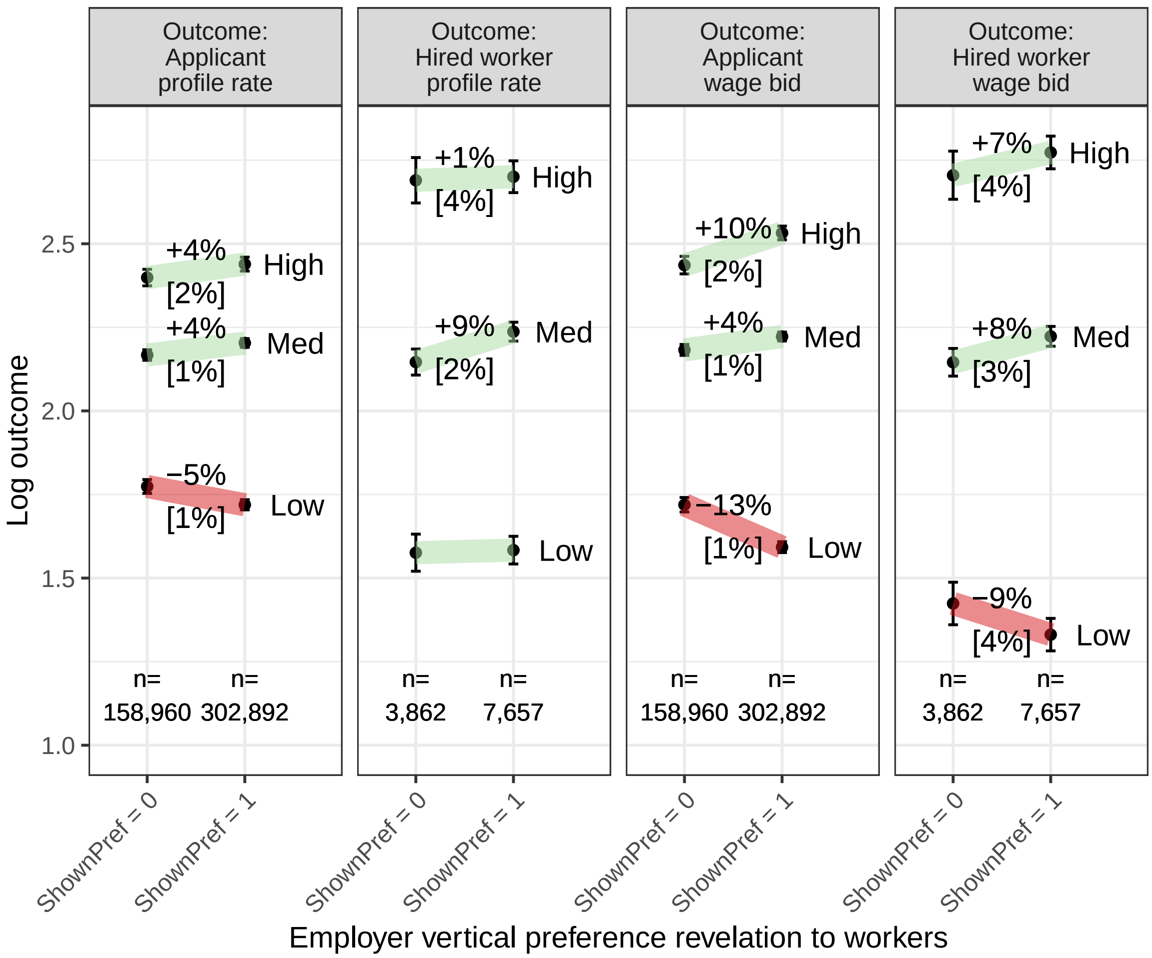

Notes: This figure plots predictions from estimates of (12), using the wage bid and profile rate as the outcomes. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] 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.

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 6—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.[15]

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

$$\begin{aligned}\log w_{ij} = \alpha_i &+ \beta_1 \textsc{ShownPref}_j + \beta_2 \textsc{MedTier}_j + \beta_3 \textsc{HighTier}_j + \\ &\beta_4 \left(\textsc{MedTier}_j \times \textsc{ShownPref}_j\right) + \\ &\beta_5 \left(\textsc{HighTier}_j \times \textsc{ShownPref}_j\right) + \epsilon. \qquad (15)\end{aligned}$$

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 (15) when the outcome is the log wage bid, which we do in Figure 7, 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, ShownPref = 0, they still "pick up" some of the employer's vertical preference, bidding more when facing a higher tier employer. The coefficient on MedTier implies workers increase

[15] 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 C.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 7: Worker wage bid, profile rate and experience at time of application, by employer vertical preference and revelation status in the ambiguous arm

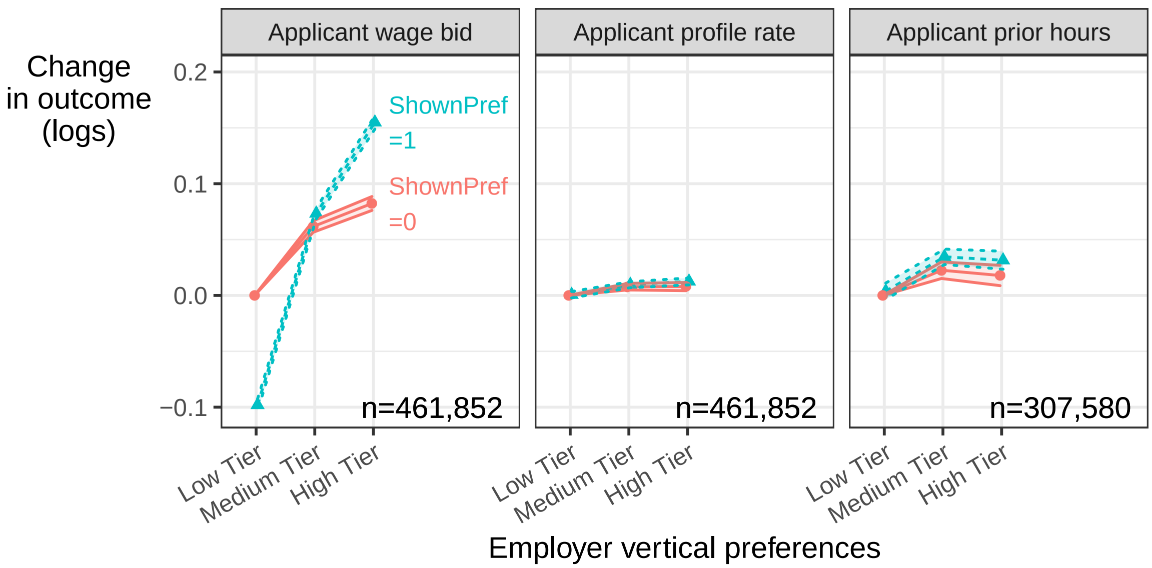

Notes: This figure reports estimates of (15). The sample consists of all applications sent to job openings in the ambiguous arm of the experiment. In each regression, a worker specific fixed effect is included. Standard errors are clustered at the level of the individual worker. The dependent variables 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.

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 7 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 C.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 5).

For the effects of revelation on the wage bids and profile rates of hired workers, we return to Figure 6. 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 4, bullet 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 C.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 C.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

$$y_j = \beta_0^k + \beta_1^k \textsc{ShownPref}_j + \epsilon \big| k = \textsc{Tier}_j \tag{16}$$

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 (16) in Figure 8, 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.[16] 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 8, we see a reduction in applicant pool sizes from revelation in both the high and low tiers (recapitulating Figure 5, 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 5, applicant pool reductions seem to be concentrated in the low tier. However, these effects are smaller than those in the Figure 5.

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 8. 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 8, 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 were made worse-off, as we saw that substantially less experienced workers hired in the low tier

[16] 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.

(recall Figure 4). In Appendix C.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 9. 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 9, 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).[17] 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.

5 Conclusion

A limitation of our experimental design is that it does not directly shed light on the full market equilibrium. At the conclusion of the experiment, all employers received an experience identical to the ShownPref = 1 cell in the explicit arm, meaning that all employers now knew their

[17] 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 8: Effects of revealing employer vertical preferences on job opening outcomes

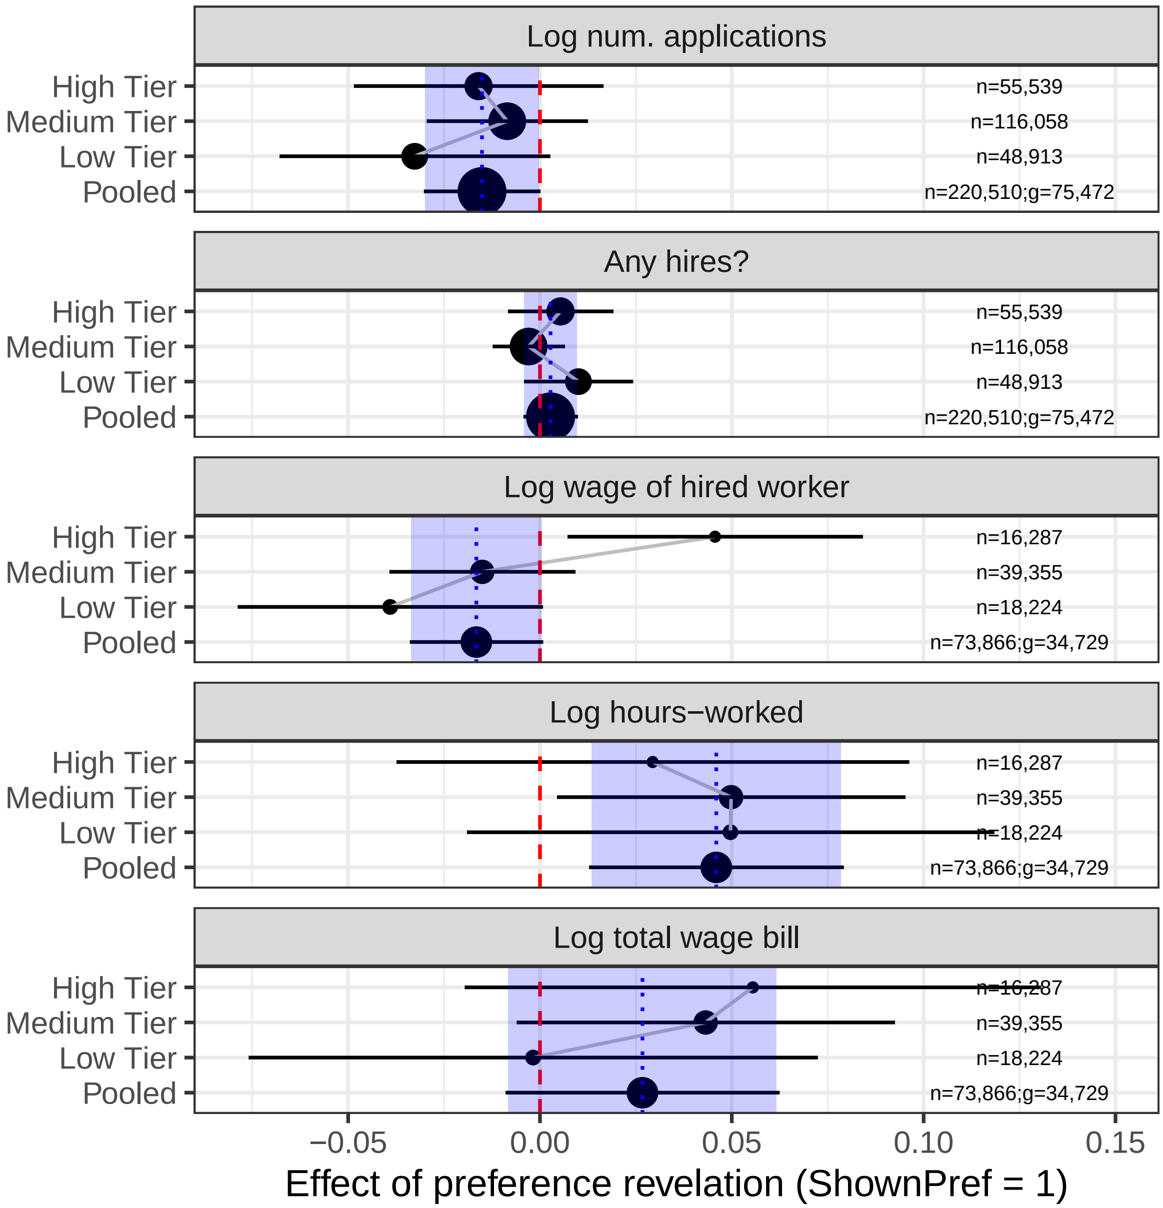

Notes: This figure shows the effects of revealing employer preferences, ShownPref $= 1$, on a number of outcomes. The sample consists of all job openings from both the ambiguous and explicit arms. Each point estimate is surrounded by at a 95% CI.

Figure 9: Effects of revealing employer vertical preferences on job opening feedback scores (z-scores)

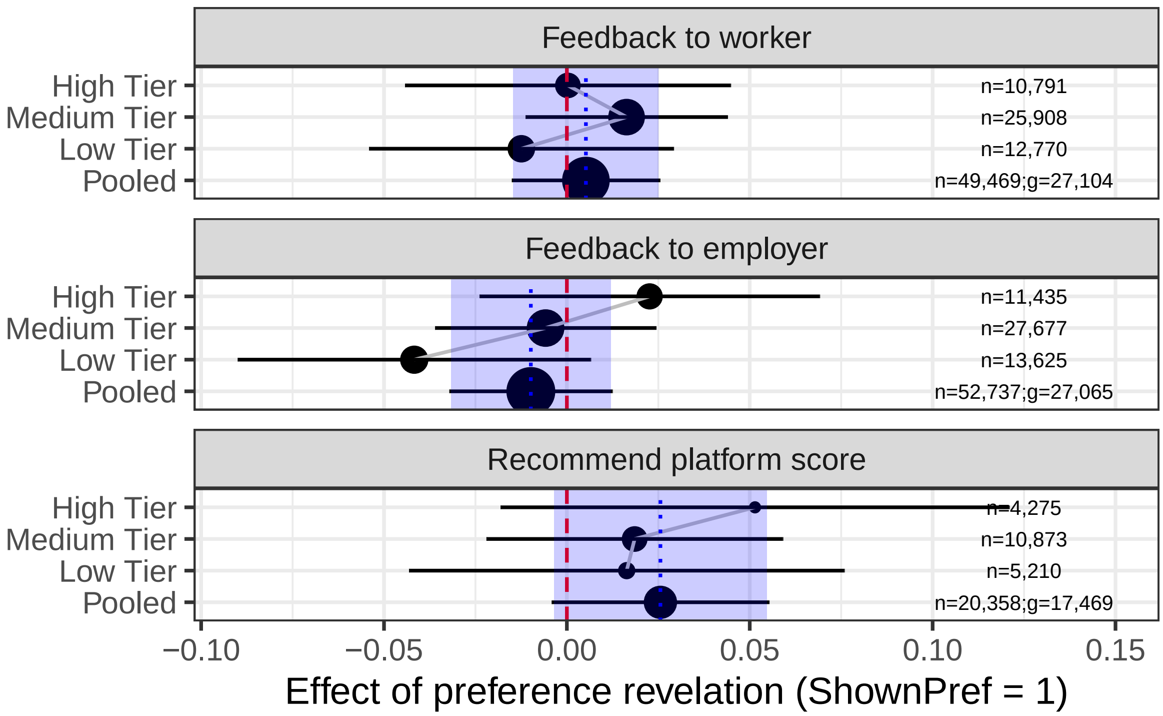

Notes: This figure shows the effects of revealing employer vertical preferences on various feedback measures. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with 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.

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 trends within employer in the tier choice and (2) look at the fraction of job openings selecting the various tiers in the post period.

During the experiment, among employers that posted multiple job openings, we can look for trends in their choice. If we saw employers pooling on a tier—the medium tier, which is the most common tier and seems like the most natural place for employers to "pool"—the long run viability of the separating equilibrium would be endangered. In Appendix C.7, we show that if anything, the trend is towards employers being more likely to select the high tier. Of course, we could have an uninformative high tier pooling equilibrium, but if employers were "moving up" because they were receiving bad applicants in the low tier, we would expect more medium tier employers as a first step, but this is not the case.

Another measure of whether the separating equilibrium persisted comes the period after the experiment ended. Figure 10 shows the fraction of employers choosing the various tiers over time, with the end of the experiment indicated. There is 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.

Platform-engineered signaling 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 a more separated equilibrium would be desirable and whether the method could be applied to other preference dimensions and other market settings.

A feature of the signaling opportunity described here is that workers were able to apply cross-tier. It would be straight-forward to design a version of the signaling opportunity in which workers would have to choose a tier and only apply within a tier for some period of time. This might address the theoretical and empirical problem that relatively low ability workers do not sort. 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. 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. This might get more informed separation, though it also increases the

Figure 10: Employer vertical preference signal choice over time, both before and after the experiment

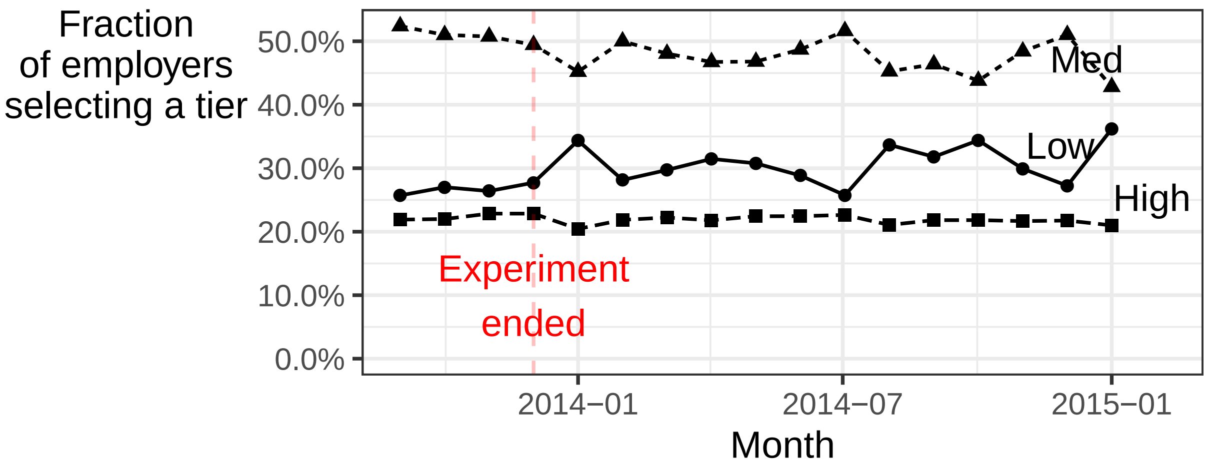

Notes: This figure shows the fraction of job openings each month selecting each of the three possible tiers. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with 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.

burden on the platform in deciding what are reasonable tier labels. These wage standards could be scaled by the category to try to induce equal shares selecting each tier.

Although our context is an online labor market, the matching process in this market mirrors that found in conventional markets. The signaling opportunity in this paper is with respect to vertical preference, but there are other potential pieces of information that might be conveyed by a signaling 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.

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.

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

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.

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.

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.

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, "Search intensity," Technical Report 2004.

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

_ , "The cyclical behavior of equilibrium unemployment and vacancies," *American economic review*, 2005, *95* (1), 25–49.

_ **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.

A Proofs

The following pages proof results about the equilibrium after truth-telling, and then studies the incentives of firms to tell the truth, and finally considers bidding when there is no message or (the message is uninformative).

A.1 Proposition 1

Proof. To establish Proposition 1, consider first the *equilibrium bidding behavior*. Expression 3 gives the utility of a worker with type a who bids b, taking the equilibrium distribution $b_m(a)$ of other workers as given. Clearly, in equilibrium the optimal choice of this worker needs to be $b = b_m(a)$. Note that in this case $\alpha_m(va - b) = a$. Taking first order conditions of the expected utility with respect to b evaluated at $b = b_m(a)$ gives

$$\text{FOC:} \ - e^{-\lambda_m(1-F_m(a))}\lambda_m f_m(a)\alpha_m'(\pi_m(a))b_m(a) + e^{-\lambda_m(1-F_m(a))} = 0$$

or, noting that $\alpha_m'(\pi) = \frac{1}{v_m - b_m'(a)}$, this gives

$$\lambda_m f_m(a) b_m(a) = v_m - b_m'(a). \tag{17}$$

This immediately means that $\pi_m(a)$ is indeed increasing in equilibrium, as the right hand side represents $\pi_m'(a)$. Note that the lowest type $\underline{a}$ only wins if no other worker is present. This type therefore asks for the full surplus:

$$b_m(\underline{a}) = v_m \underline{a}. \tag{18}$$

The differential equation (17) together with endpoint condition (18) uniquely determines $b_m(a)$. It can easily be verified that the solution takes the form

$$b_m(a) = v_m g_m(a), \tag{19}$$

where

$$\begin{aligned} \lambda_m f_m(a) g_m(a) &= 1 - g_m'(a), \\ g_m(\underline{a}) &= \underline{a}. \end{aligned} \tag{20}$$

The approach just taken is equivalent to choosing how much profit to leave to firms in (2) via its first order condition. Supermodularity then readily establishes optimality of the solution to the first order condition.[18] Clearly better workers will get higher expected utility in equilibrium, but in a search model this can be delivered through better matching probabilities rather than higher wage bids. Bids are increasing if we can ensure that $g_m'(a)$ is positive. Clearly workers cannot extract more than their type as firms would refuse to match at negative profit, we have $g_m(a) \leq a$, which yields a sufficient condition for bids to increase in types of

$$\lambda_m(a) f_m(a) \leq 1/a. \tag{21}$$

[18]This follows from standard arguments: strict supermodularity implies that it is optimal for a to choose a profit weakly lower than any optimal profit $a' > a$ and weakly smaller than any optimal profit of $a'' < a$.

Next, consider *where workers search for jobs.* One remaining difficulty is that the queue length λ_m and the type density $f_m(a)$ at each of the messages is endogeneous. It is a choice of workers where to attempt to get a job. We first proceed along the following conjecture.

Conjecture: Low types in $[\underline{a}, \hat{a})$ *mix between both announcements, while high types in* $[\hat{a}, \bar{a}]$ *only go to message* H (to be verified).

As we will see, the first interval may be empty. We prove constructively that such an equilibrium exists, but explicitly deriving the equilibrium bid distribution and the queue at each type of firms. Under the assumption that this is the unique equilibrium, we show that it is efficient. Finally, we prove that this is the unique equilibrium, which is slightly more involved.

Given the conjecture, low types have to mix between both firm types. Thewe worker types therefore have to be indifferent. We will exploit the indifference condition to find the endogeneous distribution types at each announcement. The indifference condition is:

$$e^{-\lambda_L(1-F_L(a))}b_L(a) = e^{-\lambda_H(1-F_H(a))}b_H(a). \tag{22}$$

For the lowest type, recall his bid in (18) to obtain indifference condition

$$\begin{aligned} & e^{-\lambda_L}v_L\underline{a} = e^{-\lambda_H}v_H\underline{a} \\ \Leftrightarrow\quad & e^{-(1-\gamma_H)/\delta_L}v_L = e^{-\gamma_H/\delta_H}v_H \end{aligned} \tag{23}$$

which can be achieved with $\gamma_H \in [0,1]$ if and only if $1 \geq \delta_H \ln(v_H/v_L)$. Otherwise all worker types only queue at the H message. We can solve the previous equation to get

$$\begin{aligned} \gamma_L &= \frac{\delta_L(1-\delta_H\ln(v_H/v_L))}{\delta_H+\delta_L}, \\ \gamma_H &= \frac{\delta_H(1+\delta_L\ln(v_H/v_L))}{\delta_H+\delta_L}. \end{aligned} \tag{24}$$

which are exclusively determined by exogeneous parameters.

But the indifference condition (22) has to hold at all types in $[\underline{a}, \hat{a})$. So we can differentiate (22) with respect to a and obtain in this range

$$\begin{aligned} & e^{-\lambda_L(1-F_L(a))}\lambda_L f_L(a)b_L(a) + e^{-\lambda_L(1-F_L(a))}b'_L(a) \\ = \quad & e^{-\lambda_H(1-F_H(a))}\lambda_H f_H(a)b_H(a) + e^{-\lambda_H(1-F_H(a))}b'_H(a). \end{aligned}$$

Recall that by the first order condition for optimal bidding (17) we have $\lambda_m f_m(a)b_m(a) = v_m - b'_m(a)$, so the previous inequality reduces to

$$e^{-\lambda_L(1-F_L(a))}v_L = e^{-\lambda_H(1-F_H(a))}v_H. \tag{25}$$

Note that this trivially holds for the lowest type, as seen in (23). Dividing each side in (25) by the same side in (23) and taking logs, we obtain

$$\lambda_L F_L(a) = \lambda_H F_H(a) \tag{26}$$

on $[\underline{a}, \hat{a})$. That means that at both messages, for any type a in that set, the ratio of worse workers to firms is equalized across messages. Recall that the number of types across the messages has to add up to the overall number of types in the population:

$$\gamma_L F_L(a) + \gamma_H F_H(a) = F(a). \tag{27}$$

Recalling that $\lambda_m = \gamma_m/\delta_m$, and that $\lambda = 1/(\delta_L + \delta_H)$, we can solve the system (26) and (27) to obtain

$$\lambda_L F_L(a) = \lambda_H F_H(a) = \lambda F(a) \tag{28}$$

for $a \in [\underline{a}, \hat{a}]$. This immediately implies that $\lambda_L f_L(a) = \lambda_H f_H(a) = \lambda f(a)$ and therefore $g_L(a) = g_H(a) =: g(a)$ in (21) for $a \in [\underline{a}, \hat{a}]$. That is, the bidding distribution at message H is simply v_H/v_L higher than that at message L, at least in the area where agents visit both messages.

Equality (28) further implies

$$\begin{aligned} F_L(a) &= \frac{\delta_L}{\gamma_L(\delta_H + \delta_L)} F(a) = \frac{1}{1 - \delta_H \ln(v_H/v_L)} F(a), \\ F_H(a) &= \frac{\delta_H}{\gamma_H(\delta_H + \delta_L)} F(a) = \frac{1}{1 + \delta_L \ln(v_H/v_L)} F(a). \end{aligned} \tag{29}$$

on $[\underline{a}, \hat{a})$, where the second equality in each line follows from (24). For $a > \hat{a}$, $F_L(a) = 1$ and $F_H(a) = 1 - (1 - F(a))/\gamma_H$.[19] The boundary type $\hat{a}$ is simply determined by $F_L(\hat{a}) = 1$, or equivalently $F(\hat{a}) = 1 - \delta_H \ln(v_H/v_L)$. This concludes the construction of the conjectured equilibrium.

Note that these results simplify (21) to $f(a) \leq (\delta_L + \delta_H)/a$ for $a \in [\underline{a}, \hat{a}]$ and $f(a) \leq \delta_H/a$ otherwise. So $f(a) \leq \delta_H/a$ is a sufficient condition for bids to increase in ability.

To understand the efficiency properties of the market, note that workers bid in a first price auction, which is revenue equivalent to bidding in a second price auction. This means that the payoff of worker a participating in market m is equal to the probability that no higher-ability worker arrives multiplied by: the match output ($v_m a$) minus the expected match-out of the next-best worker, where the next-best worker output includes zero's in case no other worker is present.[20] This coincides exactly with the additional surplus that a worker brings: he only adds surplus if there is no higher bidder, and even then the added surplus is only his output minus what is created by the next best worker. Workers join the the message where this is highest. Shimer (2005a) Proposition 1 proves that this, together with feasibility of queue lengths which are implied by the construction of the equilibrium, is necessary and sufficient for constrained efficiency of the market. That proof is for finite number of types, but since it holds for any finite approximation of our type distribution it also holds in the limit, and since our equilibrium is the unique equilibrium in the limit, the efficiency result extends.

[19] The expression for $F_H(a)$ at $a > \hat{a}$ arises because the mass of all higher types $1 - F(a)$ goes to the high message, and so has to equal the probability of going to that message times the conditional probability of having a type above a, which equals $\gamma_H(1 - F_H(a))$.

[20] (Formally, this is equal to $e^{-\lambda_m(1-F_m(a))}$ multiplied by: $(v_m a)$ minus $\sum_{n>1} P(n|\lambda_m F_m(a)) \int_{\tilde{a}=\underline{a}}^{a} v_m \tilde{a} \ d\left(\frac{F_m(\tilde{a})}{F_m(a)}\right)^n$, which is given by the Poisson probability $P(n|\lambda_m F_m(a))$ of n such workers when the Poisson parameter is $\lambda_m F_m(a)$ times the output of the highest such worker which is given by the $n'th$ order statistic.

Uniqueness of the equilibrium outcome after truthful messages is the remaining point we still have to prove. To this end, it suffices to show that the conjecture above is true (in the confines of this conjecture we constructed the unique equilibrium). To show that the conjecture has to hold in any equilibrium, we will prove that if one worker type is indifferent between high and low firms, then any type below also has to be indifferent (which requires these firms to mix between both announcements). This also means that there is highest type that is indifferent. We will show that types above approach high firms only. This proves that all equilibria conform to our conjecture, noting that the conjecture allowed the indifference region to be empty and that it is trivial to rule out equilibria where all workers go only to low types (as deviating and bidding $v_H a$ at a high type firms would be profitable).

It is easiest to work with the formulation where workers bid by leaving profit to firms, as in (2). Let $U_m(a)$ be the highest utility that worker type a can obtain in equilibrium when approaching firms with type and message m, and let $\pi_m(a)$ be the optimal profit he would leave to the firm. By (2) we can write

$$U_m(a) = P_m(\pi_m(a))(v_m a - \pi_m(a)). \tag{30}$$

where $P_m(\pi) = e^{-\lambda_m(1-D_m(\pi))}$. In connection with (2) we proved that $\pi_m(a)$ is weakly increasing. Moreover, on the support of F_m it is strictly increasing because D_m has no mass points, and therefore $P_m(\pi_m(a))$ is strictly increasing in a in this range. If one considers type a' that is not approaching message m in equilibrium (i.e., a' not in the support of F_m) then $\pi_m(a') = \pi_m(a'')$ is locally constant, where $a'' = \max\{a \in suppF_m | a \leq a'\}$ is the maximum type in the support of F_m that is smaller than a. The reason is the following: if a'' is the maximal type in the support of F_m, then bidding $\pi_m(a'')$ is sufficient to win for sure and bidding more is not necessary. Otherwise there exists $a''' = \min\{a \in suppF_m | a \geq a'\}$, where since D_m has no holes we have $\pi_m(a'') = \pi_m(a''')$. Since π_m is increasing, this implies $\pi_m(a') = \pi_m(a'')$. Since $\pi_m(a')$ is locally constant in worker type, this also means that the winning probability $P_m(\pi_m(a'))$ is locally constant.

Now consider any type $\hat{a}$ that is indifferent between both messages, so that $U_L(\hat{a}) = U_H(\hat{a})$. We first show that then also all types $\hat{a}' < \hat{a}$ are indifferent. We prove this by contradiction. Assume there exists $\hat{a}' < \hat{a}$ such that $U_m(\hat{a}) > U_{-m}(\hat{a})$ for $m \in \{L, H\}$ and $-m \in \{L, H\}/\{m\}$. Note that by the envelope theorem

$$U_m'(a) = P_m(\pi_m(a))v_m.$$

Let $\hat{a}''$ be the lowest type above $\hat{a}'$ that is indifferent. Given that $\hat{a}'$ strictly prefers m, to achieve the indifference at $\hat{a}''$, we need

$$P_m(\pi_m(\hat{a}''))v_m < P_{-m}(\pi_{-m}(\hat{a}''))v_{-m}.$$

But note that this implies

$$P_m(\pi_m(a))v_m < P_{-m}(\pi_{-m}(a))v_{-m} \tag{31}$$

for all $a \in (\hat{a}', \hat{a}'')$. The reason is that types in $(\hat{a}', \hat{a}'')$ by assumption strictly prefer m, and therefore are in the support of F_m but not in the support of F_{-m}. As proven earlier, therefore $P_m(\pi_m(a))$ declines as one goes to lower worker types, while $P_m(\pi_{-m}(a))$ is constant. But note that (33) means that $U'_m(a) < U'_{-m}(a)$ for $a \in (\hat{a}', \hat{a}'')$. Since we have $U_m(\hat{a}'') = U_{-m}(\hat{a}'')$, this means that $U_m(a) < U_{-m}(a)$ for $a \in (\hat{a}', \hat{a}'')$, which contradicts our initial assumption that these types prefer message m. This establishes that if one type is indifferent between the messages, then all lower types will also be indifferent.

This immediately establishes that there exists a highest worker type that is indifferent between both messages (as are all types below it). For notational convenience call this type $\hat{a}$. Types above this either all strictly prefer message L or all strictly prefer message H. It cannot be that there are strictly higher types some of which prefer L and some H, as by continuity of $U_m(a)$ this would imply that some strictly higher type would be indifferent, contradicting that $\hat{a}$ is the highest type that is indifferent.

We will now rule out that types above $\hat{a}$ strictly prefer L. We prove this by contradiction. Assume types above $\hat{a}$ do strictly prefer L. Since $\hat{a}$ is indifferent, this requires $U'_L(\hat{a}'') > U_H(\hat{a}'')$, or equivalently

$$P_L(\pi_L(\hat{a}''))v_L > P_H(\pi_H(\hat{a}''))v_H.$$

But since higher types strictly prefer L and, therefore, do not approach firms with message H, we have $P_H(\pi_H(\hat{a}'')) = 1$, so this inequality cannot hold. This contradiction establishes that all types above $\hat{a}$ strictly prefer H, which establishes that any equilibrium satisfies our conjecture, under which we found a unique equilibrium. □

A.2 Proof of Proposition 2

Proof. Note first that the probability that a firm at message m has no applicant with qualification above a is $e^{-\lambda_m(1-F_m(a))}$. That means that the density that the best applicant is of type a is $\lambda_m f_m(a)e^{-\lambda_m(1-F_m(a))}$. So the expected profit at message m for firm type v_m is

$$\begin{aligned}\Pi_m &= \int (v_m a - b_m(a))\lambda_m f_m(a)e^{-\lambda_m(1-F_m(a))}da. \\ &= v_m \int (a - g(a))x(a)e^{-\lambda_m(1-F_m(a))}da.\end{aligned}$$

It simply integrates the profits over the highest type worker that arrives. Consider first the deviation condition for v_H firms. We know that $v_L a - b_L(a)$ is positive and increasing, so we also know that $v_H a - b_L(a)$ is also positive and increasing. So if a v_H firm chooses message L, it still wants to hire the highest worker type. Truthtelling now becomes

$$\begin{aligned} v_H \int (a - g(a))x(a)e^{-\lambda_H(1-F_H(a))}da &\geq \int_{a\leq\hat{a}} (v_H a - v_L g(a))x(a)e^{-\lambda_L(1-F_L(a))}da \\ \iff v_H \int (a - g(a))x(a)e^{-\lambda_H(1-F_H(a))}da &\geq v_H \int_{a\leq\hat{a}} (a - \frac{v_L}{v_H}g(a))x(a)e^{-\lambda_H(1-F_H(a))}\frac{v_H}{v_L}da, \\ \iff \int (a - g(a))x(a)e^{-\lambda_H(1-F_H(a))}da &\geq \int_{a\leq\hat{a}} (\frac{v_H}{v_L}a - g(a))x(a)e^{-\lambda_H(1-F_H(a))}da, \end{aligned} \tag{32}$$

where the second line used the workers indifference condition (25). For high firm types the effects of sending an L message are threefold: they loose the best worker types who no longer apply, any given type is asking for less money, and finally the distribution of bids is slightly shifted. So in the last line this can be reduced to two effects: they get less of the good applicants, but get more out of each of the applicants that do apply.

Fixing v_L, δ_L and δ_H, let $\tilde{v}_H$ be such that $\delta_H \ln(\tilde{v}_H/v_L) < 1$. As v_H approaches $\tilde{v}_H$ from below, clearly truthtelling condition (32) is satisfied as the right hand side goes to zero since $\hat{a}$ goes to $\underline{a}$.

Now consider the opposite alternative where $v_H \approx v_L$. In particular, take derivatives of both side of (32) with respect to v_H, and evaluate it at $v_H = v_L$ (and therefore $\hat{a} = \bar{a}$). For truthtelling we need that the right hand side grows weakly less than the left hand side:

$$0 \geq \frac{\partial \hat{a}}{\partial v_H}(\bar{a} - g(\bar{a}))x(\bar{a})e^{-\lambda_H} + \int_{a \leq \hat{a}} \frac{1}{v_L} a x(a) e^{-\lambda_H(1-F_H(a))} da.$$

Recall that $F(\hat{a}) = 1 - \delta_H \ln(v_H/v_L)$, so that $\frac{\partial \hat{a}}{\partial v_H} = -\frac{-\delta_H \frac{v_L}{v_H}\frac{1}{v_L}}{-f(\hat{a})}$. So we can write the inequality as

$$0 \geq -\delta_H(\bar{a} - g(\bar{a}))x(\bar{a}) + f(\bar{a}) \int a x(a) e^{\lambda_H F_H(a))} da.$$

Clearly for $f(\bar{a})$ small enough this is satisfied, and clearly one can choose $f(\bar{a})$ so that this is violated.

Now consider the low types. Studying deviations for low types is much harder, as after a deviation to message H it is no longer obvious that the profit $v_L a - b_H(a)$ is increasing in worker type. If it is not increasing, then they would not necessarily choose the highest worker type who applies to them. We will therefore look at two extreme cases that mirror the previous analysis of high types: the case where v_H is large, and the case where $v_H \approx v_L$. Consider first the part where v_H is large. We can establish that low types do not want to deviate if $b_H(a) \geq v_L a$ for all skill levels. In this case the low types can never benefit from deviating. Note that this is clearly the case at $a = \underline{a}$, as $b_H(\underline{a}) = v_H \underline{a}$. At higher levels $v_L a - b_H(a)$ can fall, as long as it always stays positive. We are done if we can show that $b'_H(a) - v_L \geq 0$ at any a where $b_H(a) - v_L a = 0$. This reduces to the requirement that

$$\begin{aligned} v_H - \lambda_H f_H(a) v_L a - v_L &\geq 0 \\ \Leftrightarrow v_H - v_L(\lambda_H f_H(a) a + 1) \geq & \quad 0 \end{aligned} \tag{33}$$

Intuitively this should hold when v_H becomes large. But to make this formal we have to take into account that $\lambda_H f_H(a)$ is endogeneous and can vary with v_H. Moreover, v_H is only possible when δ_H is sufficiently small, as otherwise we are in the uninteresting case where workers do not visit the low signal. So, fix v_L, δ_L and $F(a)$. Then fix δ_H sufficiently small. Finally let v_H be sufficiently large, but still workers mix between both, i.e., below but close to $\tilde{v}_H(\delta_H)$ which is defined by $1 = \delta_H \ln(\tilde{v}_H(\delta_H)/v_L)$. It is useful to note that $\delta_H \tilde{v}_H(\delta_H)$ goes to infinity as δ_H becomes small.

On $[\hat{a}, \bar{a}]$ where types only visit message H, we have $\lambda_H f_H(a) = f(a)/\delta_H$. Insertion into (33) readily establishes that the inequality holds when δ_H small and v_H close to $\tilde{v}_H(\delta_H)$, precisely since $\delta_H \tilde{v}_H(\delta_H)$ goes to infinity.

On $[\underline{a}, \hat{a})$ in the area of mixing

$$\begin{aligned}
\alpha_L f_L(\alpha) + \alpha_H f_H(a) &= f(a) \\
\Leftrightarrow \delta_L \lambda_L f_L(\alpha) + \delta_H \lambda_H f_H(a) &= f(a) \\
\Leftrightarrow \delta_L \lambda_H f_H(\alpha) + \delta_H \lambda_H f_H(a) &= f(a) \\
\Leftrightarrow \lambda_H f_H(a) &= \lambda f(a) = \lambda_L f_L(a).
\end{aligned}$$

where $\lambda = 1/(\delta_H + \delta_L)$. Clearly $\lambda_H f_H(a)$ remains bounded even for vanishing δ_H, and so (33) holds when δ_H small and v_H large. So we have established that all wage bids at the high message are above the valuation of low types, and deviation to the high message is strictly unprofitable.

Consider now the case where v_H is close to v_L. In this part the surplus continues to increase, and we can write the truthtelling condition in analogue to the previous one

$$\begin{aligned}
v_L \int_{a \leq \hat{a}} (a - g(a)) x(a) e^{-\lambda_L(1-F_L(a))} da &\geq \int (v_L a - v_H g(a)) x(a) e^{-\lambda_H(1-F_H(a))} da \\
\Longleftrightarrow v_L \int_{a \leq \hat{a}} (a - g(a)) x(a) e^{-\lambda_L(1-F_L(a))} da &\geq v_L \int (a - \frac{v_H}{v_L} g(a)) x(a) e^{-\lambda_L(1-F_L(a))} \frac{v_L}{v_H} da, \\
\Longleftrightarrow \int_{a \leq \hat{a}} (a - g(a)) x(a) e^{-\lambda_L(1-F_L(a))} da &\geq \int (\frac{v_L}{v_H} a - g(a)) x(a) e^{-\lambda_L(1-F_L(a))} da, \qquad (34)
\end{aligned}$$

Taking derivatives gives

$$\begin{aligned}
\frac{\partial \hat{a}}{\partial v_H} (\bar{a} - g(\bar{a})) x(\bar{a}) e^{-\lambda_L} &\geq -\int \frac{1}{v_L} a x(a) e^{-\lambda_L(1-F_L(a))} da \\
\Leftrightarrow f(\bar{a}) \int a x(a) e^{\lambda_L F_L(a)} da - \delta_H(\bar{a} - g(\bar{a})) x(\bar{a}) &\geq 0.
\end{aligned}$$

Note that $\lambda_L F_L(a) = \lambda_H F_H(a)$, so that whenever truthtelling for the L type is ensured, it is violated for the H type. Therefore, for small differences in valuation truthtelling cannot be achieved. □

A.3 Proof of Proposition 3

Proof. Consider the setting without messages, or a babbling equilibrium in the game with messages. In either case messages are uninformative, and we drop the reference to them. Workers therefore approach firms at random, but condition their bid on the signal s that is common to the firm they face.

Consider first workers facing a firm with high signal $s = H$. With sufficient signal precision workers are nearly sure to face a high value firm. Therefore their bids are mostly targeted to them, and one can prove that a supermodularity condition in analog to that associated with (2) applies when ψ is sufficiently large. This assures that for high value firms $\pi_{H,s=H}(a) = v_H a - b_{s=H}(a)$ is increasing in a. In analogy to the previous proofs, let $\alpha_{H,s=H}$ be the inverse of π_H. It will become clear that firms only want to cater high types. Assume that all workers submit bids only acceptable to the high value firms. Then the equilibrium payoff of workers a bidding b and facing a firm with signal $s = H$ is

$$\left(\Psi_{H,s=H} e^{-\lambda(1-F(\alpha_{H,s=H}(va-b)))}\right) b,$$

since $e^{-\lambda(1-F(\alpha_{H,s=H}(va-b)))}$ is the probability that no other bidder is present that offers a better value to high type firms, with probability $\Psi_{H,s=H}$ the firm is actually of high type and accepts (see definition (9)) in which case the return is b. The first order condition evaluated at $b = b_{s=H}$ is (10). From the first order condition it is again clear that π_m is increasing, as mentioned. This means that the lowest type can only win if no other bidder is present, in which case he can demand the whole surplus if high types, given the boundary condition in (10). It remains to be verified that low firms would not accept any of the bids, i.e., $b_{s=H}(a) > v_L a$ for all $a \in [\underline{a}, \bar{a}]$. The bidding function here is identical to the bidding function at the high message in the separating equilibrium since $\lambda_H f_H(a) = \lambda f(a)$, and so the proof that bids outpace the valuation of low firms in (33) applies also here. Finally, we have to check that firms never want to attract low types: since bids are bounded away by some $\triangle$ from the valuation of low types implies that this would substantially lower the payoff at least by $\Psi_{H,s=H} e^{-\lambda} \triangle$ while the gains are vanishing as ψ goes to one.

Now consider workers facing a firm with low signal $s = L$. When signal precision is high workers are nearly fully convinced that they face a low type firm, and it is easy to show that it is strictly profitable to make bids that low types will accept. And since high types value the service even more all firm types will accept. One can also again prove a single crossing condition in analogue to (2) on the profit for low types when the signal is sufficiently precise, so $\pi_{L,s=H}(a) = v_L a - b_{s=L}(a)$ is again strictly increasing. Then $\pi_{H,s=L} = v_H a - b_{s=L}(a)$ is also strictly increasing. Let $\alpha_{L,s=L}$ and $\alpha_{H,s=H}$ be the inverse of them, respectively.

The lowest worker type $\underline{a}$ can only win if no other worker is present. So he will extract the surplus from low firms $b_{s=L}(\underline{a}) = v_L \underline{a}$, and obtain utility $e^{-\lambda} v_L \underline{a}$. Clearly this exceeds the value increasing his bid further and extracting all value from high types $(1 - \Psi_{H,s=L}) e^{-\lambda} v_H \underline{a}$ as long as signal precision is high. A similar argument holds for all other types. In general the expected utility of a worker of type a who bids b is given by

$$\left(\Psi_{L,s=L} e^{-\lambda(1-F(\alpha_L(v_L a-b)))} + \Psi_{H,s=L} e^{-\lambda(1-F(\alpha_H(v_H a-b)))}\right) b.$$

That is, he wins if he meets a low type firm (probability $\Psi_{L,s=L}$) and if none of the other workers are ranked higher by this firm, and with complementary probability the same for a high type firm. The first order condition evaluated at $b = b_{s=L}(a)$ is

$$\begin{aligned} \sum_{t \in \{L,H\}} \Psi_{t,s=L} \left(-e^{-\lambda(1-F(a))} \lambda f(a) \alpha_t'(\pi_t(a)) b(a) + e^{-\lambda(1-F(a))}\right) &= 0 \\ \Leftrightarrow \sum_{t \in \{L,H\}} \Psi_{t,s=L} \lambda f(a) \frac{1}{v_t - b'(a)} b(a) &= 1, \end{aligned}$$

where we omitted the subscripts of the bidding function for notational brevity. $\square$

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: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

B 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.

C Additional empirical results

C.1 Effects on wage bidding

In the bottom panel of Figure 12 (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.

C.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

Figure 11: Comparison of outcomes using applicant-level regression in the ambiguous arm with extensive pre-treatment job level controls

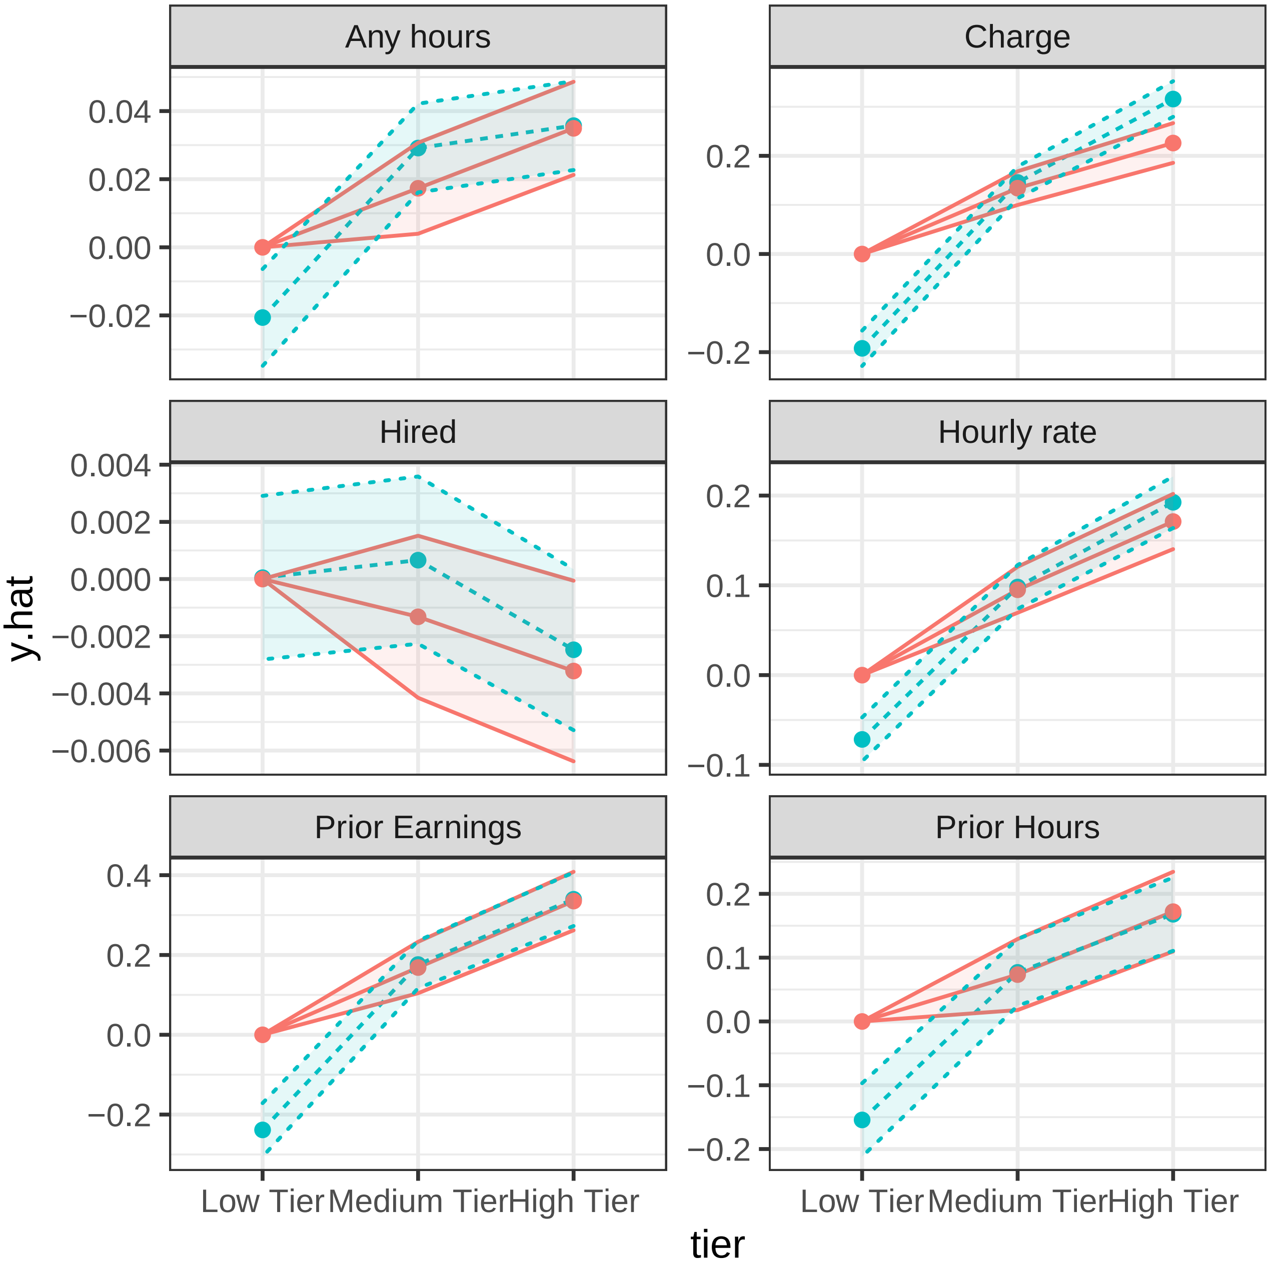

Notes: This figure plots coefficients from estimates of (12). The outcome in both panels is the applicant's cumulative prior hourly earnings at the time of application. The sample in the left panel is all applicants in the ambiguous arm of the experiment, whereas in the right panel, the sample is hired 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 12: Effects of showing employer vertical preferences, ShownPref = 1, on applicant pool composition with respect to wage bidding in the ambiguous arm

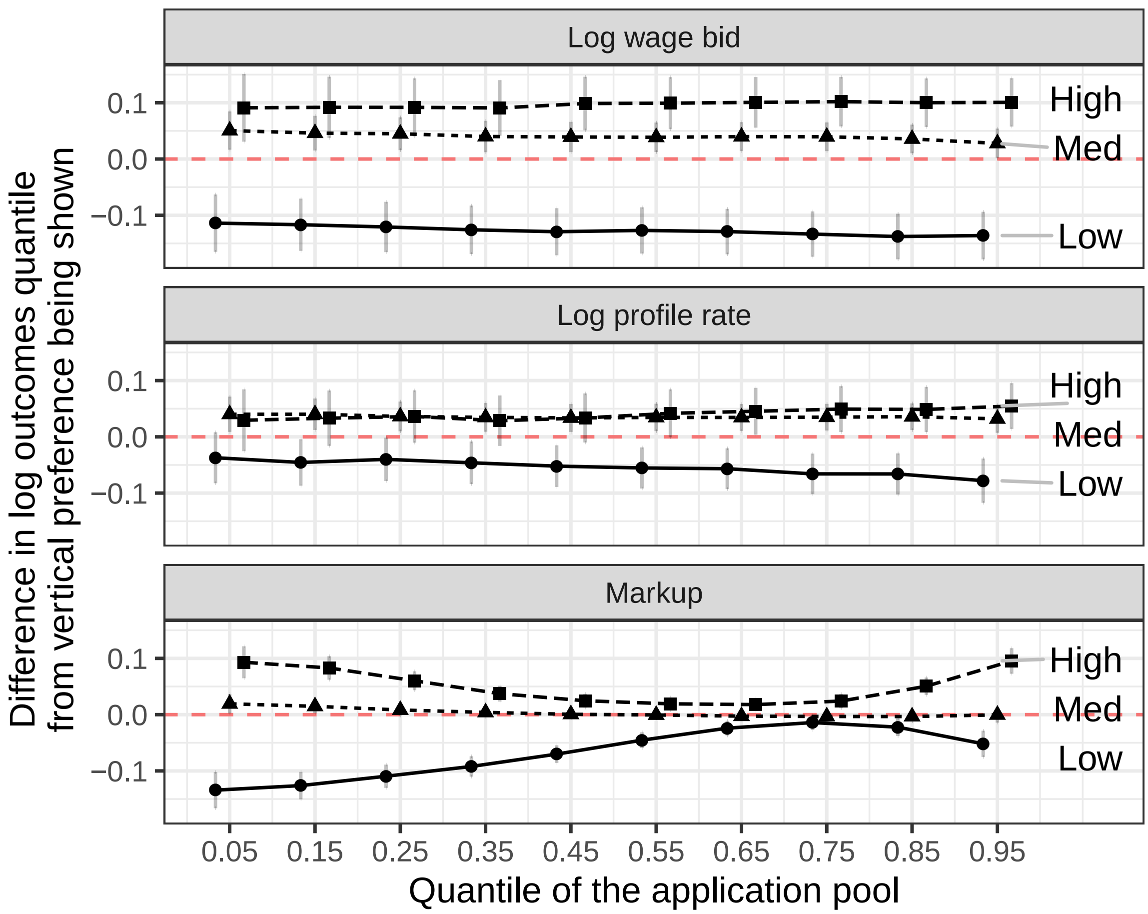

Notes: The figure shows the effects of vertical preference revelation on the composition of applicant pools with respect to wage bidding and profile rates. The sample is the ambiguous arm of the experiment. Each point is the mean effect of revelation on some applicant attribute at that quantile of the 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.

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 2 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 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).

C.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

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: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

job-to-job basis. Ideally we would estimate a regression of the form

$$\log w_{ij} = \alpha_i + \beta_1 \log A_j + \epsilon \tag{35}$$

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 (35), 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.

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: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

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.[21] We

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

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: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

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 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.

C.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.

C.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.

C.6 Match outcome result robustness to sample definition

Figure 13 reports results for a number of outcomes using different sample definitions.

C.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.

Figure 13: Effects of revealing employer vertical preferences on job opening outcomes

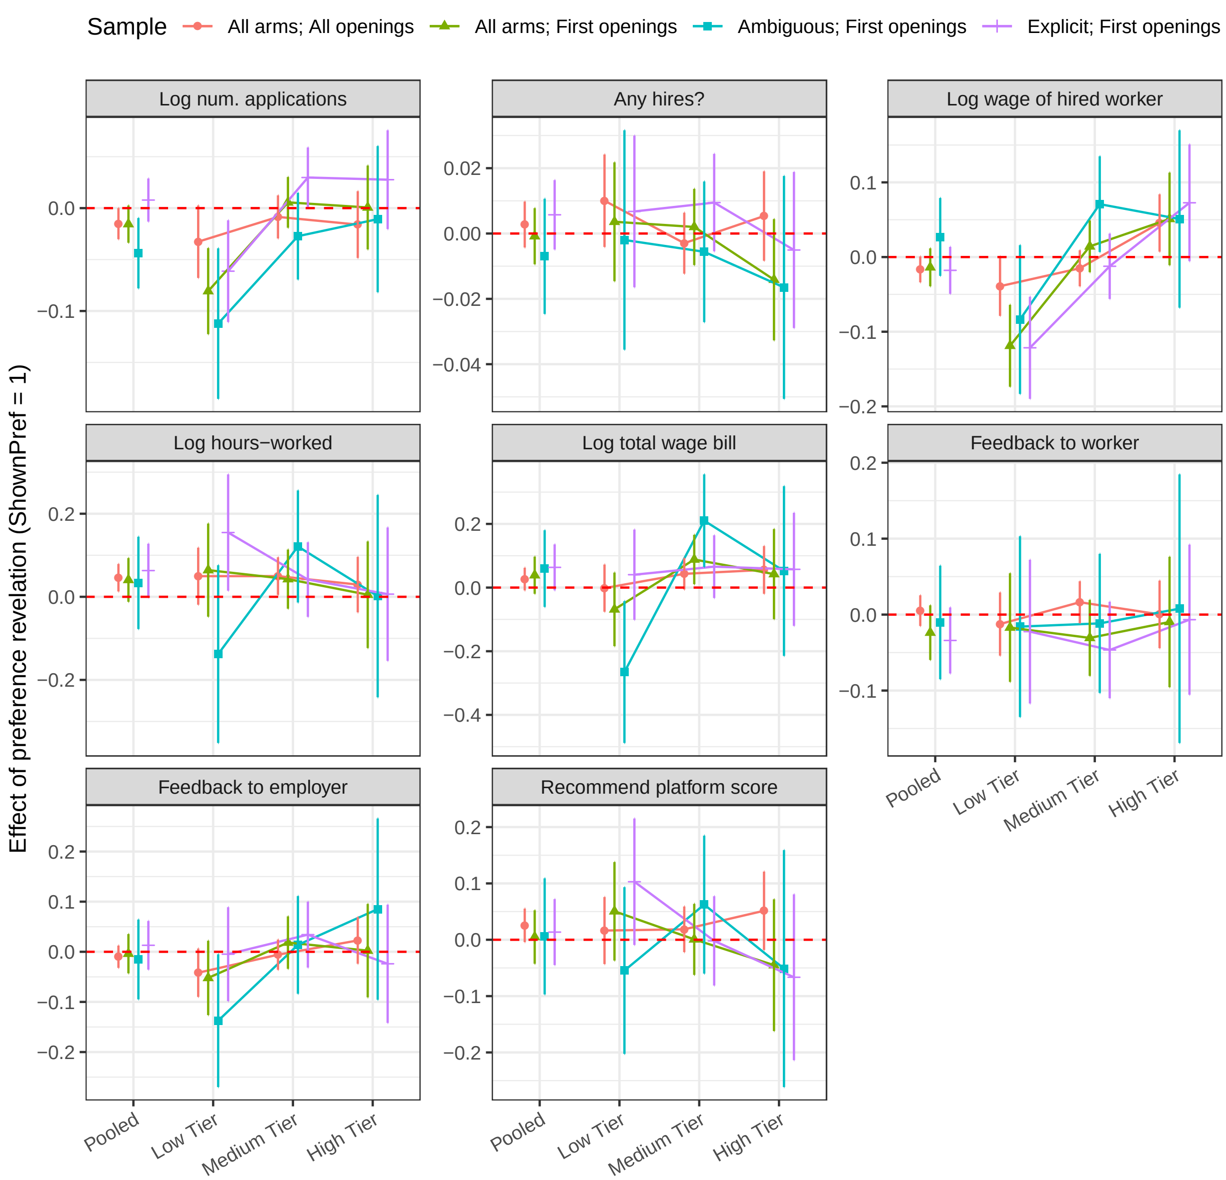

Notes: This figure shows the effects of revealing employer preferences, ShownPref = 1, on a number of outcomes, for several different samples. Each point estimate is surrounded by at a 95% CI.

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.

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: $\dagger$:$p < 0.10$, $*$:$p < 0.05$, $**$:$p < 0.01$, $***$:$p \leq 0.001$.

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.