Buyer Signaling Improves Matching: Evidence from a Field Experiment*

John J. Horton
NYU Stern

Ramesh Johari
Stanford University

October 26, 2018

Abstract

In a large online market, buyers were given the opportunity to signal their relative preferences over price and quality—first experimentally, then later as the default experience in the market. The possibility of signaling caused substantial sorting by sellers to buyers of the right "type." However, sellers clearly tailored their bids to the type of buyer they faced, bidding up against buyers with a high revealed willingness to pay. Despite this markup, a separating equilibrium was sustained over time, post-experiment, suggesting buyers found revelation incentive compatible. We find evidence that informative signaling improved both matching efficiency and match quality.

Keywords: signaling, matching, market design, experimentation

1 Introduction

In many matching markets buyers and sellers have incomplete information about the attributes and preferences of the other side of the market. This

*Author contact information available at http://www.john-joseph-horton.com/. Thanks to Andrey Fradkin, Mike Luca, Meng Liu, Steve Tadelis, 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.

causes both sides to engage in costly information acquisition (Stigler, 1961). To the extent residual information gaps remain, participants are likely to make sub-optimal matches. In response to this shortcoming, market designers might consider introducing signaling opportunities that allow one side of the market to reveal information to the other side. Although economists have long-recognized the potential for signaling to overcome market informational problems (Spence, 1973; Crawford and Sobel, 1982), it is far from clear that (1) the introduction of such an opportunity would lead to a separating equilibrium, and (2) even if it were obtained, whether such an equilibrium would be better than the status quo.

In this paper, we investigate the introduction of a new signaling opportunity in a large online labor market. In our empirical setting, sellers are freelance workers that differ in their experience and degree of expertise. All employers presumably value these qualities, but they differ in their willingness to pay for them—some employers want to hire the very best workers, even at a very high wage, while other employers require only "good enough" work and are very price-sensitive. Heterogeneity in employer "vertical" preferences surely exists in conventional settings as well, and are important to would-be applicants, as they are informative about which job openings to apply to, and at what terms.

Historically, job-seekers in our marketplace had to infer employer vertical preferences from employer and job characteristics. As in the conventional labor market, there was no formal, standardized way for employers to communicate their vertical preferences. As a platform-initiated change to the market, employers were given the opportunity to explicitly state their vertical preferences to would-be applicants, potentially revealing their relative willingness to pay for "better" workers.[1]

The signaling opportunity was simple: when posting a job opening, employers selected one of three "tiers" to describe the kinds of applicants they

[1] Our use of the term "employer" simply reflects the typical usage of these terms in the literature and is not a comment on the nature of the legal relationship created on the platform.

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. When the signaling opportunity was introduced market-wide (which occurred after an experimental period), the tier choice was revealed publicly to all job-seeking workers.

An employer that believes his or her tier choice will be revealed must consider the two effects the signal will have: the effect on (1) the pool of applicants he or she receives and thus who can be hired, and (2) the ultimate wage that is paid to the hired worker. As workers can condition their wage bids on the preferences signaled by employers, truthful revelation might not be in the employer's best interest: an employer signaling "high" may get more experienced applicants, enabling a better match, but at the cost of higher wage bids by those applicants; conversely, an employer signaling "low" might benefit from lower wage bids, but at the cost of a lower quality applicant pool. In short, employers have to weigh the bargaining and sorting effects relative to their own preferences when deciding what to signal.

Our empirical analysis focuses on two questions. Does a separating equilibrium emerge following the introduction of the signaling opportunity? And if such a separating equilibrium emerges, what are the consequences for welfare? Although our context is an online labor market, the same questions would apply in other markets where this signaling intervention could be implemented. A similar signaling internvetion could be introduced in online job boards, and professional social networks, which in turn cover a large fraction of all job openings (Azar et al., 2017; Marinescu and Wolthoff, 2016). CareerBuilder, LinkedIn, Indeed, Monster.com and other sites where employers upload job openings could introduce similar centrally-managed signaling mechanisms.[2]

For the first question, we show that a separating equilibrium is obtain-

[2]TechCrunch, "Google to launch a jobs search engine in the U.S." Posted May 17, 2017 https://techcrunch.com/2017/05/17/google-to-launch-a-jobs-search-engine-in-the-u-s/ TechCrunch, "Facebook threatens LinkedIn with job opening features" Posted November 7, 2016, https://techcrunch.com/2016/11/07/jobbook/.

able: employers have substantial heterogeneity in their vertical preferences; they reveal these preferences through use of the signal; and applicants sort in response to employers' signals. However, workers are also strategic and bid up against employers with a revealed high willingness to pay and bid down when facing employers with a revealed low willingness pay. These combined bidding and sorting effects by tier created a dilemma for employers regardless of type, but as we show, they evidently found truthful revelation of their preferences incentive compatible. For the second question, we provide evidence that a separating equilibrium is more efficient: fewer job applications are needed to form just as many matches, and the matches that are formed are of higher quality, on average.

Our data comes primarily from an experimental period that preceded market-wide roll-out of the signaling opportunity. During this experimental period, employers were assigned to different treatment groups upon posting a job opening. As we later discuss in detail, these treatment groups differed on two dimensions: (1) whether the employer's tier choice was revealed to the would-be applicant; and (2) the *ex ante* knowledge employers had about this revelation. The purpose of each experimental group will be apparent as we discuss various research questions they were designed to answer.

We organize our research findings as follows. First, we discuss the heterogeneity of employer preferences, and their signaling choices. Next, we show that applicants respond to the employers' revealed signals both by sorting and by conditioning their wage bids. Finally, we study match outcomes and the welfare implications of a separating equilibrium.

Employer preferences and signaling behavior. We begin our analysis using the "explicit arm" of the experiment, in which employers were told, *ex ante*, whether or not their tier choice (i.e., low, medium, or high) would be shown to applicants. For those employers who knew their signal would not be revealed, there is no strategic aspect in deciding which signal to send. We regard their tier choices as the "ground truth," in the sense that their choices were not influenced by an anticipated worker response.

Among employers that did not have their tier choice revealed, we find that there is substantial variation in vertical preferences. Some of the variation in employer preferences can be explained by the nature of the work to be performed. For example, in the "Administrative Support" category, 59% employers chose the low tier, compared to just 17% choosing the low tier in "Web Development"; in "Networking & Information Systems" the fraction of employers choosing the high tier was 33%, whereas only 6% of employers in "Administrative Support" chose the high tier.

Despite the importance of the job category in explaining the variation in vertical preferences, there is still substantial within-category variation, with no tier signal being selected by even close to 100% of employers within a category of work. This residual within-category variation is important for our purposes, as it implies that a signal *could be* informative to workers.

The existence of variation in preferences among employers does not guarantee they are willing to reveal it, particularly if they know workers can condition upon that information. Employers in the explicit arm who knew their tier choice would be revealed might have chosen a different tier had they counterfactually been told their choice would be kept private. For example, a high type employer might choose to signal "medium" if they expect signaling "high" will lead to large markups in the wage bids they face. Or a low-type employer might also choose to signal "medium" if they expect a "low" signal will lead to a very poor pool of applicants. It was precisely this possibility of strategic misreporting that motivated the creation of the explicit arm.[3]

Despite the possibility of employer strategic misreporting, we find no evidence for this hypothesis. Indeed, we find that the distributions of tier choices by employers in a category were essentially identical, regardless of whether their choice was going to be revealed to workers. It is possible that employers did not anticipate the changes revelation would bring about—changes we will discuss. But another possibility is that they did anticipate the changes and

[3] The design here is similar to that used in Bursztyn et al. (2017), who manipulate the expectation among survey-takers that their survey responses will be shared with peers or kept private, with the experimental outcome being the survey responses.

preferred them, consistent with truthful revelation being incentive compatible. [4]

Applicant sorting based on employer preference revelation and a separating "equilibrium". Next, we investigate how the revelation of an employer's vertical preferences to would-be applicants affected the applicant pool. Our interest is in whether workers sorted to employers of the right "type," with highly experienced workers applying to high tier employers and inexperienced workers applying to low tier employers. This better sorting could reduce the total number of applications needed, which could improve matching efficiency. We are also interested in how revelation changed the collection of wage bids employers faced, both from the compositional change in the applicant pool, as well as from workers' conditioning their bids on the employers' willingness to pay.

For this analysis, we use the "ambiguous arm" of the experiment, in which employers were told that the platform *might* reveal their preferences to workers, with actual revelation determined at random, *ex post*. Because employers did not know whether their signal would be revealed, we can treat their signal choice as exogenous, allowing us to see how revelation of the signal affected the composition of the applicant pool and applicant behavior, by tier. Note that (as discussed above) we found no evidence that employers conditioned their tier choice on whether their signal will be revealed to workers; however, the platform did not know this would be the case when it designed the experiment, which is why it created the ambiguous arm.

We find that workers strongly respond to the signals sent by employers. Among high tier employers, those who had their preference revealed received applicants that were 7.4% more experienced at the time of application (as measured by total prior earnings on the platform). The equivalent revelation effect for low tier employers was a -18.4% reduction in experience.

Applicant pools were not only more sorted—they were also smaller on

[4] We note that in principle, there could be hidden sorting that leaves the fraction of employers choosing different tiers unchanged, though this seems relatively improbable and inconsistent with other pieces of evidence we present.

average. Overall, applicant counts per opening were reduced by -5%. The effects were concentrated among low tier employers, who experienced pool sizes that shrunk by more than 10%. By reducing application costs, signal revelation could improve matching efficiency, so long as the quantity and quality of matches does not decline; we explore this hypothesis further in our discussion of match outcomes below.

In addition to sorting in response to the revealed signal, workers also conditioned their wage bids on the signal they observed. For high tier employers, revelation of the signal increased the average wage bids they received by 9.7%. For low tier employers, revelation lowered wage bids by -12.7%.

These changes were not solely due to the compositional changes induced by the signal revelation; exploiting the fact that workers send multiple applications, we use a within-worker analysis to show that workers directly condition their wage bids on the revealed signal. In short, the same worker bids more when they know they are facing an employer with a higher willingness to pay. We also show that this conditioning in wage bids was not caused by applicants perceiving job amenities/disamenities that might correlate with the tier.

Once an applicant pool is formed, an employer has to decide which workers to hire, if any. We find that there was no detectable effect of signal revelation on hiring: the quantity of matches formed remained unchanged. But revelation did change the characteristics of matches formed, with the applicant pool changes wrought by revelation being largely "passed through" and reflected in the attributes of the hired workers. The pass-through was particularly pronounced among low tier employers, where signal revelation caused hired workers to have -22.7% lower cumulative past hourly earnings. Revelation in the other tiers shifted hired worker composition in the expected direction (given the change in applicant pools)—among high tier employers, revelation caused hired workers to have 3.1% higher earnings, though the points estimates are not precise.

Changes in the applicant pool wage bids caused by revelation were reflected in the wage bids of hired workers. For low tier employers, those that had their preference revealed hired workers that bid -9% less, whereas for high tier employers they hired workers that had bid 7% more. This pass-through shows that employers did not fully "undo" the compositional effects on the applicant pool through selection.

Taken together, our results provide strong evidence that the introduction of the signal created a separating "equilibrium" in the marketplace. Employers truthfully revealed their relative preference for price and expertise, and workers sorted themselves accordingly. Despite strategic wage adjustments by workers in response to employers' preference revelation, our analysis reveals that these sorting effects were maintained in the matches formed.

Match outcomes. The central goal of the platform in creating the signaling opportunity was to create "better" matches and improve matching efficiency. In terms of matching efficiency, we first show revelation of the signal did not lower the total number of matches formed. Given the decline in the total application counts, the non-decline in matches is *prima facie* evidence of a superior equilibrium. However, it also matters whether the quality of matches formed changed. Defining match quality is challenging, as there are several dimensions potentially of interest;accordingly, we describe analysis of a number of proxies that could be suggestive of better or worse matches.

One characterization of a better match is one where the buyer (happily) buys more of what the seller is selling. In our setting, if we observe an employer buy more hours from a hired worker, it suggests they valued those hours more. Of course, the nature of labor presents some challenges to this interpretation; workers vary in their productivity, and so an employer might pay twice as much per hour for a twice-as-productive worker but buy half as many hours, keeping the wage bill unchanged. As such, the total amount paid in wages (hours-worked times the hourly rate) within a contract is also a useful quantity measure. Using the wage bill as a proxy for match quality allows us to sidestep the complication created by the fact that hours-worked are not recorded in efficiency units.

Examining our two quantity proxies for match quality, we find that revelation increased hours-worked by 4.6% and the total wage bill by about 2.7%.

The by-tier point estimates are close to the overall "pooled" estimate not conditioned by tier. To the extent that our quantity measures are proxies for match quality, it seems that the separating equilibrium obtained by revelation of employers' signals is preferred.

In addition to quantity measures of match quality, we also observe the feedback that each party gave at the end of the contract. These measures of feedback suggest employers were as satisfied—and potentially more satisfied—if their signal was revealed, undercutting a concern that increased hours-worked or wage bills reflected padded hours. However, worker satisfaction seems to depend somewhat on the tier, with scores increasing in the tier, consistent with a gift exchange conception of the employment relationship (Akerlof, 1982).

At the conclusion of the experimental period, the signaling opportunity was made available to all employers and all changes were revealed. Although we lack experimental evidence from this period, the fractions of employers choosing the various tiers stayed more or less constant after the market-wide implementation, suggesting the separating equilibrium persisted.

Contribution. The effects of introducing the signaling opportunity are consistent with an improved equilibrium: employers received more applicants of the right "type," somewhat fewer applications had to be sent to make just as many matches, and employers of all kinds seemed to form better matches. A particularly noteworthy feature of the new equilibrium is the increased hiring of less experienced workers, who tend to find entry into these kinds of markets challenging (Pallais, 2014; Stanton and Thomas, 2015).

The dilemma we posited—that employers faced when deciding what tier to choose—was largely borne out in the data: high tier employer must ask whether the better applicant pool is worth the higher price; and low tier employers must ask whether lower wage bids are worth the lower quality applicant pool. Our answer is that for both kinds of employers, revelation was apparently worth it. They revealed their preferences (at least as measured by what they would state to the platform) during the experiment—and continued to reveal their preference over time, after the experiment was over and all choices were revealed. As we observe employers over time as they post additional jobs, we can also test for within-employer trends in tier selection that would be consistent with regret of learning. We find little evidence of a strong trend, and to the extent there is one, it is away from the medium tier, again consistent with employers finding revelation attractive.

Although our context is a labor market, our results are potentially relevant to all designed markets that have some degree of differentiation and decentralized matching. The results are particularly relevant to large markets where participants are unlikely to have fully-formed preferences over all potential trading partners. In such markets, search and information acquisition are unavoidable, but as we show, the platform can help. Many markets fit this characterization, and it may explain the growing market design interest in understanding how preferences are learned and communicated (Acemoglu et al., 2017; Ashlagi et al., 2017). The "design pattern" illustrated in this paper can be applied to many kinds of marketplaces, particularly those with sufficient heterogeneity in both preferences and quality to make sorting "worth it" and the computer-mediation that makes market interventions relatively cheap (Varian, 2010).

In addition to the market design relevance of our results, the paper also has implications for our understanding of labor markets more generally. This is the first paper to show experimentally that worker perceptions of employer willingness to pay affects sorting and bargaining. On the issue of bargaining, there is much indirect evidence for the importance of rent-sharing (Blanchflower et al., 1996; Card et al., 2015). We show its importance in a setting where perceived willingness to pay is exogenously manipulated and many competing explanations are "off the table" because of the online, spot market-like context.

The typical perspective on rent-sharing is that it is something that firms would like avoid (Goldschmidt and Schmieder, 2017; Weil, 2014). However, in our setting, there is heterogeneity in employer preferences, with some employers/firms sacrificing bargaining position to get a better applicant pool, with other employers/firms sacrificing applicant pool size and quality to improve their bargaining position. This endogenous rent-sharing is, as far as we know, a novel finding.

The rest of the paper is organized as follows: In Section 2, we explain the empirical context and related work on signaling in matching markets. In Section 3, we explain the experimental design. In Section 4, we show that employers have vertical preferences and are willing to signal them. How workers sort and bid in response to these revealed signals, and how the composition of filled job openings changes are presented in Section 5. The effects of revelation on match formation and match outcomes are reported in Section 6. We discuss future directions for research and conclude in Section 7.

2 Empirical context

The setting for our analysis is a large online labor market. In these online labor 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 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 signaling 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, 2017).

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 arguably be thought of as the party choosing hours-worked.

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

2.1 Sources and uses of information

The matching process in our setting is largely about the acquisition and evaluation of information. Although any self-reported information on the platform is cheap talk, employers and workers have little strategic incentive to misreport much of it, as incentives are aligned (Farrell and Rabin, 1996). For example, an employer looking for a web developer has no incentive to claim that he or she wants to hire an accountant.

Even if there are some pieces of self-reported information that could self-servingly influence a match if misstated—such as a worker over-claiming experience with some technology—the shadow of bad feedback following bad performance—and the potential verifiability of some claims during the pre-hire screening process used by the employer—probably curtails gross misrepresentation. 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.

For information the platform verifies, such as past hours-worked and wages, the platform serves as a kind of labor market intermediary (Autor, ed, 2009). Like other labor market intermediaries, many online platforms focus their product development and improvement efforts on injecting new and better sources of information into the market.

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, 2017) and capacity to take on more work (Horton, Forthcoming). 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.

Seller improvement is not the only effect of information disclosures—in some cases, they can be used to thicken markets by coordinating buyers and sellers, particularly with respect to buyer vertical preferences. This coordinating function is both possible and useful because in many markets, buyers differ in their willingness to pay for the vertical attribute, and sellers/goods are heterogeneous. For example, there are buyers specifically looking for distressed assets, damaged cars, and in the case of labor markets, entry level workers. In these markets, revealing information about the quality of the good can beneficially lead buyers to sellers of the right "type," as in Tadelis and Zettelmeyer (2015).

2.2 Signaling as a source of information

Information does not have to be generated solely by the platform; platforms can also facilitate the revelation of information by creating signaling opportunities. Several recent papers have considered the effect of signaling in matching markets, though all have focused on scenarios in which participants are privately signaling their interest in a particular counter-party (Lee and Niederle, 2015; Coles et al., 2013). For example, the economics job market signaling mechanism allows job-seekers to reveal some idiosyncratic preference for a school—by conventional wisdom, a school that would not otherwise anticipate the job-seeker's interest. By making such signals scarce, the platform makes them informative, though as Kushnir (2013) shows, the total number of matches can be reduced by this kind of preference signaling.

The vertical preference signaling we consider differs fundamentally from buyer-specific "pin point" preference signaling exemplified by the economics job market signaling mechanism. First, the party signaling is the buyer and his or her stated preferences are relevant to many sellers. Second, given that a signal is informative about a buyer's willingness to pay, there is scope for strategic misreporting, wheres for pin point preferences, the market function is mostly coordination.

2.3 Measuring seller vertical differentiation

A separating equilibrium depends on the existence of vertical differentiation in sellers. In the competitive labor market model, "verticality" is nicely summarized by a worker's market wage, as it equals his or her marginal product. For workers that have worked recently in the platform market, this past market wage is available. However, any particular wage could reflect idiosyncratic job amenities/disamenities and furthermore, the farther back in time the wage was observed, the more likely that it does not reflect current worker productivity or market conditions. As an alternative measure of worker productivity that suffers from fewer of these defects, we use a worker's profile rate. The profile rate is an hourly wage that is displayed publicly on each worker's profile page.

The profile rate is set by the worker at his or her desired level, but it tends to closely follow a worker's 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." However, workers can and do tailor their bids to the job opening they are applying to. We also use a worker's experience on the platform at the time of application as an outcome to measure sorting. The preference in online labor markets for workers with more on-platform experience is well-established (Pallais, 2014; Stanton and Thomas, 2015).

3 Experimental design and internal validity

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 experiment was run by the platform from 2013-07-18 to 2013-12-05. A

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

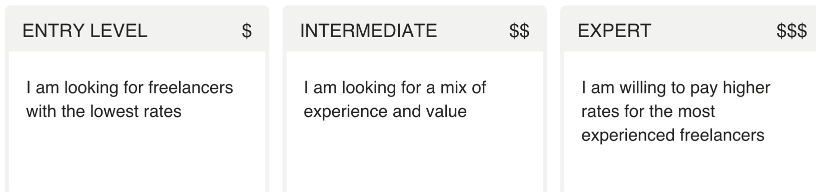

Notes: This figure shows the vertical preference signaling interface and language presented to employers when they posted a job.

total of 50,877 employers were allocated to the experiment. These employers collectively posted 220,510 job openings, though we typically only use the first opening posted by an employer in our analysis.[5] 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. All the data we use in this paper was obtained directly from the platform.

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

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

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

	Allocation	Signal 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 $= 0$.

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

was labeled with the employer's vertical preference in the interface shown to workers. The labeling was prominently displayed to make it salient to applying workers.[6]

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.

[6] 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.

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 get 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 Existence and revelation of vertical preferences

We first examine whether there is variation in employer vertical preferences. Using only data from the explicit arm of the experiment, 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 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.

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

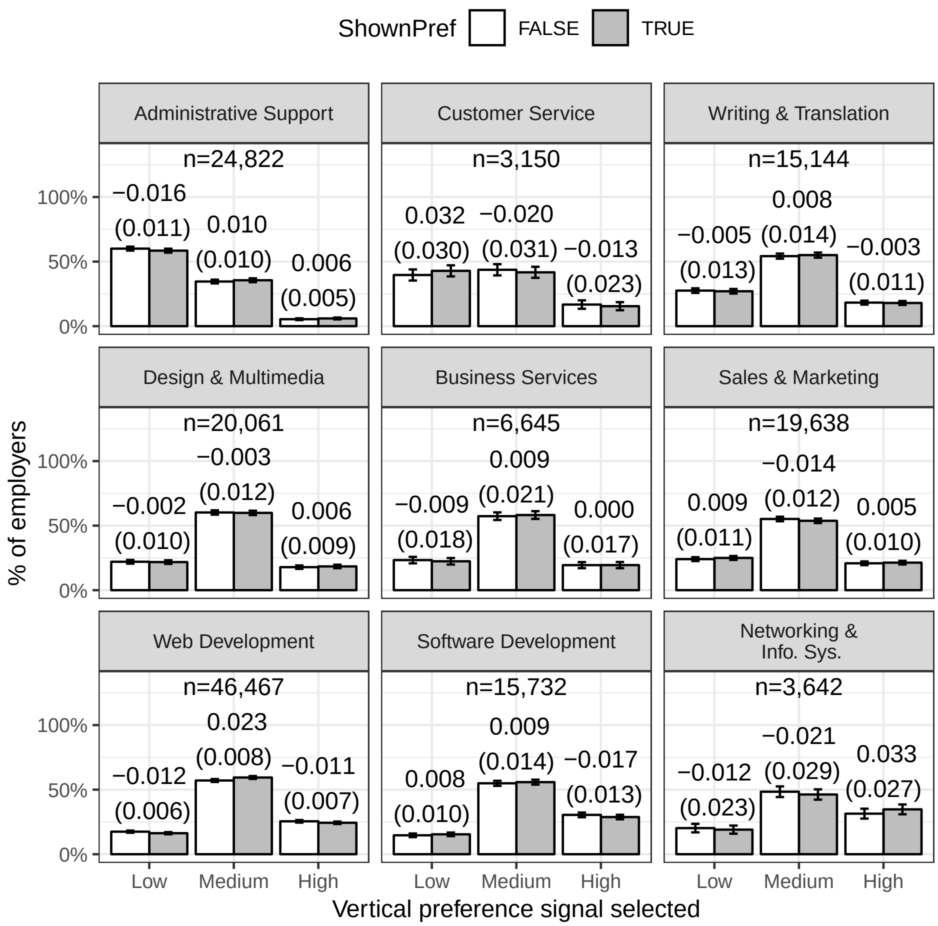

Notes: This figure shows the fraction of employers in the explicit arm of the experiment selecting the various vertical preference tiers, by ShownPref, and by the category of work of the associated job opening. When posting a job opening, employers had to select from one of three "tiers" to describe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

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. Vertical preferences clearly vary between categories, but 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 by knowing the category of work.

When the experiment was designed, it was expected that employers would condition their tier choice on the nature of their project—and also on whether their choice would be shown to would-be applicants. In particular, we expected to observe more "pooling" in the medium tier when ShownPref = 1. The design intent of the explicit arm was to test this "endogenous tier" hypothesis.

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. However, 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.

5 Worker sorting in response to the signal

We now turn to the question of how tier revelation affected applicant pool composition. 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. 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

$$\log y_{ij} = \sum_k \beta_k^s \cdot \textsc{SignalTier}_{kj} + \epsilon_j \big| s = \textsc{ShownPref}_j, \tag{1}$$

where y_{ij} is some outcome of interest for worker i applying to job opening j, $\textsc{SignalTier}_{kj}$ 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 models, 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, regardless of the number of applications received. We cluster standard errors by the job opening.

The left panel of Figure 3 plots both sets of $\hat{\beta}_k^s$ coefficients from Equation 1 where the outcome is the applicant's total prior earnings at the time of application. 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.[7] Workers with no experience at the time of application are dropped from the sample.

We can see from the pattern of $\hat{\beta}_k^0$ that even when preferences are not revealed, there is already substantial sorting. High tier employers get more experienced applicants and low tier employers get less experienced applicants, with medium tier employers getting applicants in the middle.

Despite the clear evidence of sorting—with workers presumably respond-

[7]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 revelation of the signal using applicant-level regression in the ambiguous arm

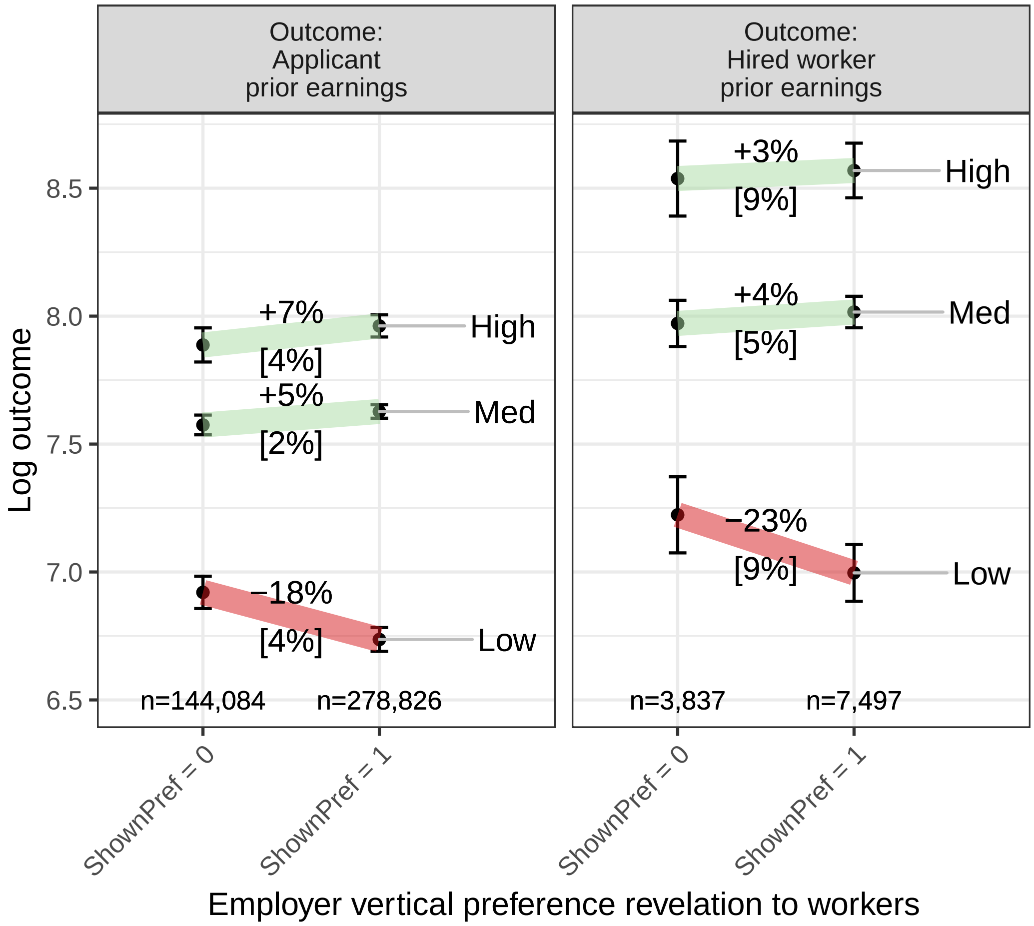

Notes: This figure plots coefficients from estimates of Equation 1. The outcome in both panels is the applicant's cumulative prior hourly earnings at the time of application. The sample in the left panel is all applicants in the ambiguous arm of the experiment, whereas in the right panel, the sample 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.

ing to the category and other observable attributes correlated with vertical preferences—if we look at the ShownPref = 1 coefficients, $\hat{\beta}_k^1$, we can see that revelation increases sorting in the expected directions. 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 prior earnings by -18.4%.

In addition to the application level approach, we can also detect changes in composition by examining quantiles of the distribution of applicant experience, by job opening. 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 above the 25th percentile, revelation had the desired 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 more or less the same pattern of sorting, with more evidence of separation between medium and high tiers. This greater separation in hours-worked could reflect workers interpreting the signal as specifically referring to experience as measured by hours-worked.

5.1 Number of job applications per opening

By making some job openings more or less attractive to workers, signal revelation could change the number of applicants received. The size of the applicant pool (in logs) is the outcome in the regression results reported in Figure 5.

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

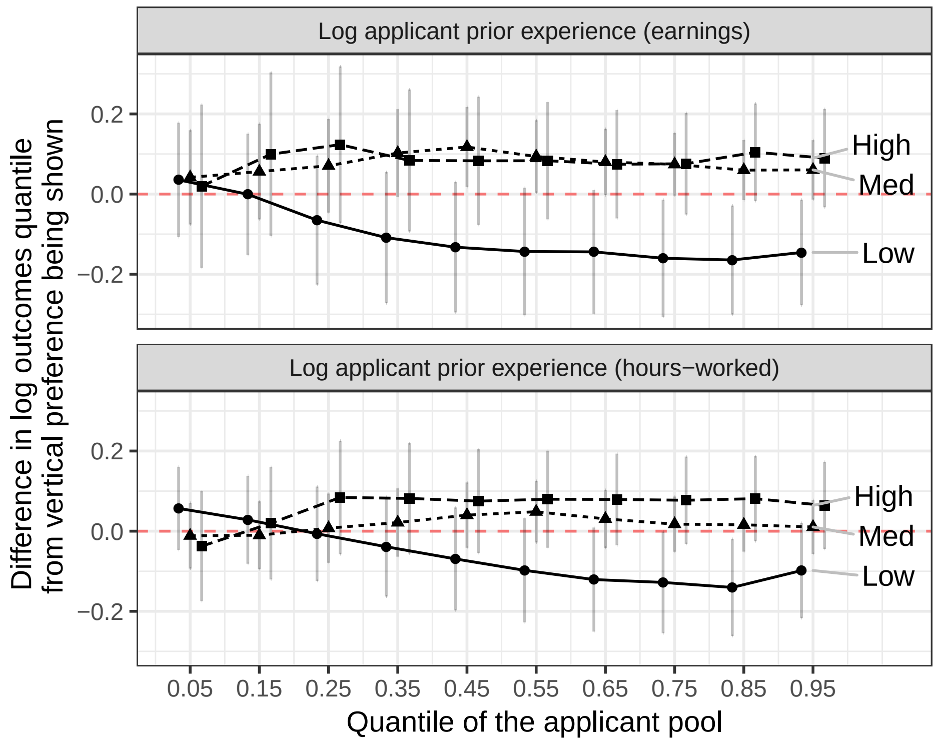

Notes: This figure shows the effects of employer vertical preference revelation on the composition of applicant pools in the ambiguous arm of the experiment for worker experience at the time of their application. The top panel is for cumulative earnings, while the bottom panel is for hours-worked. Workers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

Our baseline regression is

$$\log A_j = \beta_0 + \beta_1 \textsc{ShownPref}_j + \epsilon, \tag{2}$$

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

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 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.[8]

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 Equation 2. 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%.

We would expect reductions in the applicant pool to differ by tier. To test for heterogeneous effects, we estimate

$$\begin{aligned}\log A_j = \beta_0 &+ \beta_1 \textsc{ShownPref}_j + \beta_2 \textsc{MedTier}_j + \beta_3 \textsc{HighTier}_j + \\ &\beta_4 \left(\textsc{MedTier}_j \times \textsc{ShownPref}_j\right) + \\ &\beta_5 \left(\textsc{HighTier}_j \times \textsc{ShownPref}_j\right) + \epsilon, \end{aligned} \tag{3}$$

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.

We can see from Figure 5 that in both arms, applicant pool reductions are concentrated in the low tier for both samples, with revelation having little discernible effect in the other tiers. It seems that the low tier was uniquely off-putting to would-be applicants.

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

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

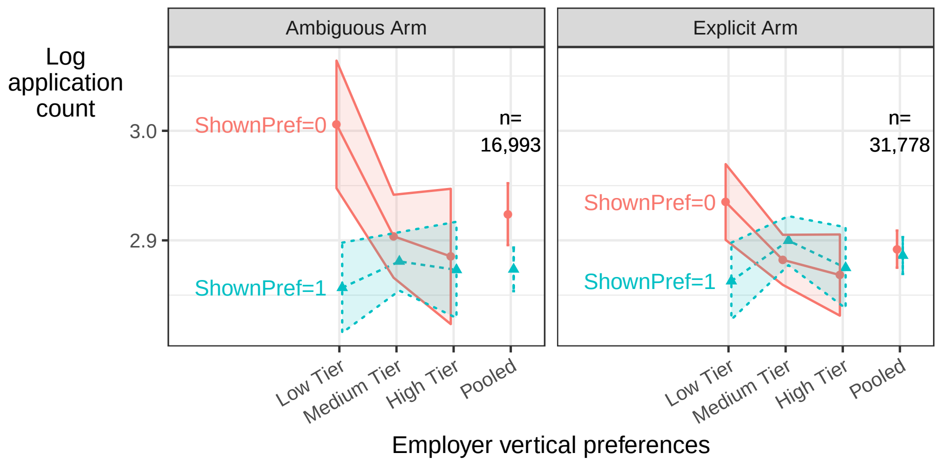

Notes: This figure reports regression results where the outcome is the log number of applications received by that opening. The right panel uses job openings from the explicit arm, whereas the left panel uses openings from the ambiguous arm. The samples are restricted to job openings receiving at least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

5.2 Worker wage bidding in response to the signal

We now turn to the question of how revelation of the signal affected the wage bidding of applicants. We again estimate Equation 1 but the outcome is the log wage bid. In the second panel from the right of Figure 6, we see the same pattern of separation in wage bids that we observed with experience, even with ShownPref $= 0$. And as before, revelation of the tier intensified the sorting. 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 unsurprising given the compositional changes caused by revelation. High tier employers who have their signal revealed should receive higher wage bids because the workers who apply to those employers are more experienced (recall Figure 3). But composition does not have to be the only explanation: workers could also directly condition their bid on perceived employer willingness-to-pay.

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 (unless the worker changed it specifically for that opening), whereas the wage bid can be conditioned on the specific features of the job opening, including the employer's tier choice, if available.

Using the log profile rate as the outcome in the leftmost panel of Figure 6, we see the same sorting pattern and revelation effect as we have for all outcomes. However, the 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. As a direct measure of a bargaining effect, we can use as

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

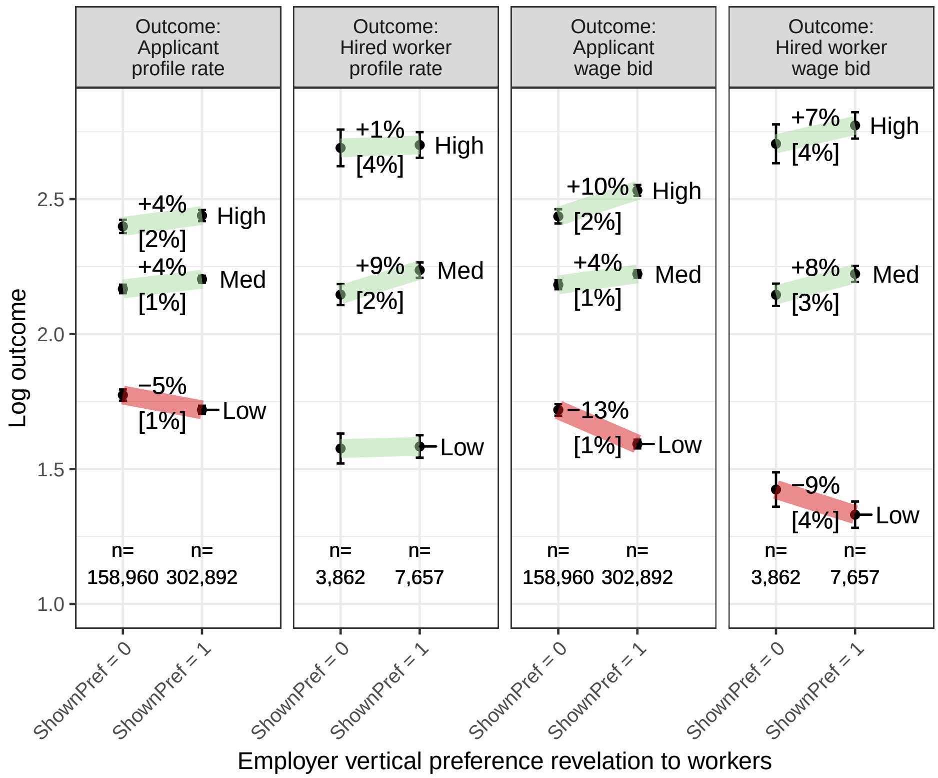

Notes: This figure plots predictions from estimates of Equation 1, using the wage bid and profile rate as the outcomes. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers]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

Figure 7: Effects of showing employer vertical preferences, $\textsc{ShownPref} = 1$, on applicant pool composition with respect to wage bidding in the ambiguous arm

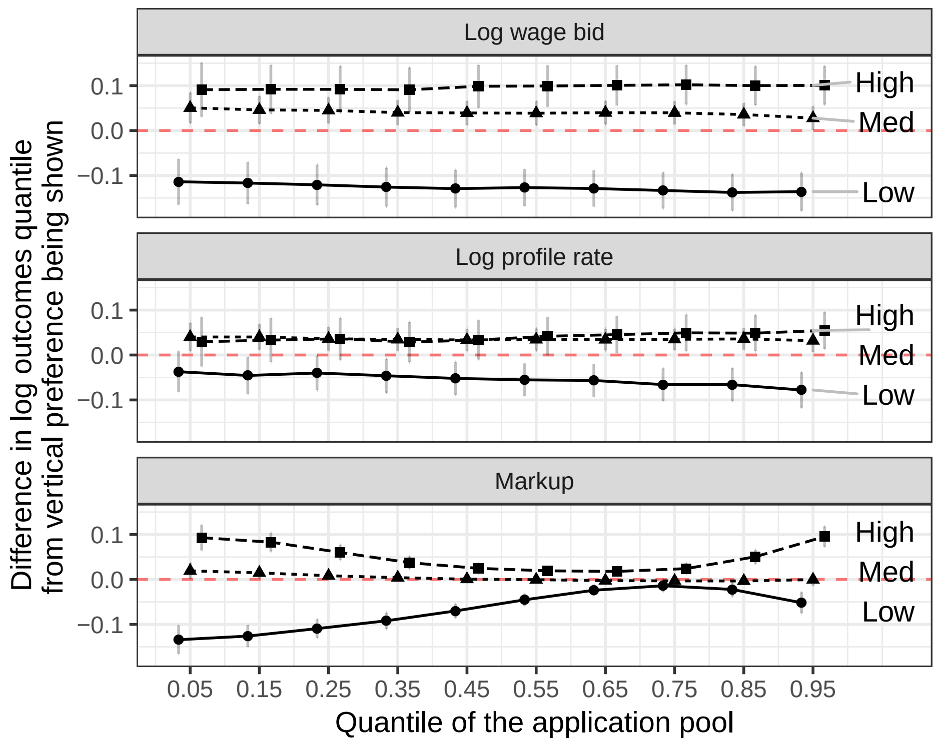

Notes: The figure shows the effects of vertical preference revelation on the composition of applicant pools with respect to wage bidding and profile rates. The sample is the ambiguous arm of the experiment. Each point is the mean effect of revelation on some applicant attribute at that quantile of the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

an outcome the "markup" in the wage bid, or the difference between the wage bid and profile rate, divided by the profile rate. Using this outcome, in the bottom panel of Figure 7 (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.

5.3 Within-worker approach to detecting conditioning

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

$$\begin{aligned}\log w_{ij} = \alpha_i &+ \beta_1 \textsc{ShownPref}_j + \beta_2 \textsc{MedTier}_j + \beta_3 \textsc{HighTier}_j + \\ &\beta_4 \left(\textsc{MedTier}_j \times \textsc{ShownPref}_j\right) + \\ &\beta_5 \left(\textsc{HighTier}_j \times \textsc{ShownPref}_j\right) + \epsilon. \end{aligned} \tag{4}$$

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.

In Figure 8, the left panel shows the mean predicted values from the estimate of Equation 4 when the outcome is the log wage bid. 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 their wage bids by 6.2%, and the coefficient on HighTier implies a 8.2% increase in the wage bid.

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

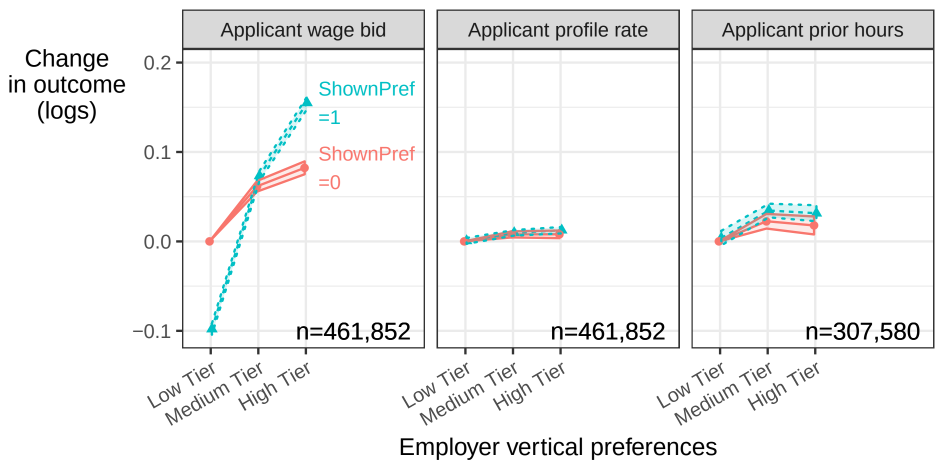

Notes: This figure reports estimates of Equation 4. The sample consists of all applications sent to job openings in the ambiguous arm of the experiment. In each regression, a worker specific fixed effect is included. Standard errors are clustered at the level of the individual worker. The dependent variables 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.

When the tier is revealed, workers adjust their wage bids much more strongly. They bid -9.8% less when ShownPref = 0 when they know it is a low tier opening; if the worker learns it is a high tier job opening, they bid an additional 7.3% more, on top of the 8.2% increase noted above.

If our within-worker approach removes worker composition effects, neither the tier nor the revelation of the tier should matter (much) for outcomes that are quasi-fixed attributes of the applicant. In the middle panel of Figure 8 the outcome is the applicant's 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).

Regardless of ShownPref, for both of these quasi-fixed worker attributes, experience and profile rates are slightly increasing in the vertical preference tier. This increase reflects that over the 5 month course of the experiment, workers gain experience and shift to 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 more or less "works" at netting out composition effects.

5.4 Alternative explanations for bargaining effects

The effects on wage bidding are consistent with workers directly conditioning on employer willingness to pay, but there are alternative explanations. For example, a worker might infer something about an employer's expectations—and hence the worker's costs—if they are hired. Despite the plausibility of this "compensating differentials" argument, there are several pieces of evidence against this explanation.

First, 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). Furthermore, as we show in Appendix A.1, there is no evidence that high tier employers were harsher reviewers when giving feedback even when preferences were not revealed. Finally, the kinds of job openings where employers were more likely to claim high tier preferences (e.g., software development, networking and information systems—recall Figure 2) were not those where we would expect "harder" work so much as smarter work.[9]

Another non-amenities channel by which the treatment could alter wage bids is through its effect on applicant pool size. For example, applicants to the low tier might anticipate they will be facing less competition and bid up, offsetting some of the effect of the low perceived willingness to pay. We explore this possibility in Appendix A.2, finding that any wage bid adjustment would be exceedingly modest, even if workers could fully anticipate changes in pool sizes.

5.5 Effects of revelation on the characteristics of hired applicants

To see whether revelation affects the characteristics of hired workers, we estimate Equation 1, but with the sample restricted to hired workers. 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.

First, we look at the experience of hired workers. Returning to Figure 3, which showed how the composition of the applicant pool changed with revelation, we now examine the right panel, where the outcome is still the worker's cumulative earnings at the time of application, but the sample is restricted to hired applicants. 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 $\exp(8.6) \approx \$5,400$.

The effects of revelation on hired worker experience mirror the effects of

[9] 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).

revelation on applicant experience. 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.

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 the sample 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 it in the low tier by -9%.

5.6 Discussion of signaling, sorting, and bidding results

Our results demonstrate that a separating equilibrium is obtainable, with employers credibly signaling their relative vertical preferences. Further, workers sort based on experience in a manner consistent with the revealed preferences of employers. Finally, our analysis shows these sorting effects are sustained through the formation of matches.

Why would employers choose to credibly reveal their preferences? To understand the trade-offs they face, consider a simple model with two types of employers: H ("high") type and L ("low") type. Workers are endowed with a "quality" θ. If an employer of type t hires a worker of quality θ at wage w, their payoff is:

$$v_t\theta - w, \tag{5}$$

where $v_H > v_L$ represents the relative preference of quality vs. price for each type of employer. We assume employers maximize expected payoff (where expectations are over all uncertainty related to matching and wages).

To obtain a separating equilibrium, high (respectively, low) type employers must prefer to declare themselves high (respectively, low) type. Let q_t denote the expected quality obtained by an employer when they declare type $t = H, L$, and let w_t be the resulting wage. Then these equilibrium conditions are that:

$$v_H q_H - w_H > v_H q_L - w_L; \quad v_L q_L - w_L > v_L q_H - w_H. \tag{6}$$

If these pool effects carry over to the attributes of hired workers, then in the marketplace, $q_H > q_L$ and $w_H > w_L$: in other words, declaring that one is a high type employer could lead to a higher quality worker as a match, but at the cost of a higher wage. Similarly, declaring that one is a low type employer leads to a lower wage, but at the cost of a lower quality worker. In a separating equilibrium, these effects work out favorably for each type of employer, and thus preference revelation is sustained.

6 Results on match outcomes

We now examine whether revelation of the employer's tier affected the quantity and characteristics of matches formed. Of course, 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 A.3, 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 A.5 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.

6.1 Objective match outcomes

Our regression specification for match outcomes is

$$y_j = \beta_0^k + \beta_1^k \textsc{ShownPref}_j + \epsilon \Big| k = \textsc{SignalTier}_j \quad (7)$$

where $\textsc{SignalTier}_j$ 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.

Using the full data, we report estimates of the coefficient on ShownPref from Equation 7 in Figure 9, using as outcomes: (1) the log number of applications, (2) whether any worker was hired and then, selecting only filled openings, (3) the log wage of the hired worker, (4) the log hours-worked of the hired worker, and (5) the log total wage bill. Note that (3) and (5) are based on the actual mean wage over the contract, not the worker's original bid.

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.[10] 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

[10] 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.

pooled regression, the number of distinct employers ("$g = \ldots$").

In the top panel of Figure 9, we see a reduction in applicant pool sizes from revelation in both the high and low tiers (recapitulating Figure 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, suggesting that our previous estimate might be an overestimate due to sampling variation.

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

Although the number of matches did not change, there are several pieces of evidence that the matches themselves changed. In the third panel from the top of Figure 9, we can see that revelation in the high tier increased hired worker wages by 4.6%, while revelation in the low tier decreases wages by -3.9%, with little effect on the medium tier. The net overall effect, indicated by "Pooled," is slightly negative. However, this does not necessarily imply workers were made worse-off, as we saw that substantially less experienced workers hired in the low tier (recall Figure 4). In Appendix A.4, 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%. These increases in quantities are suggestive of better matches being formed.

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

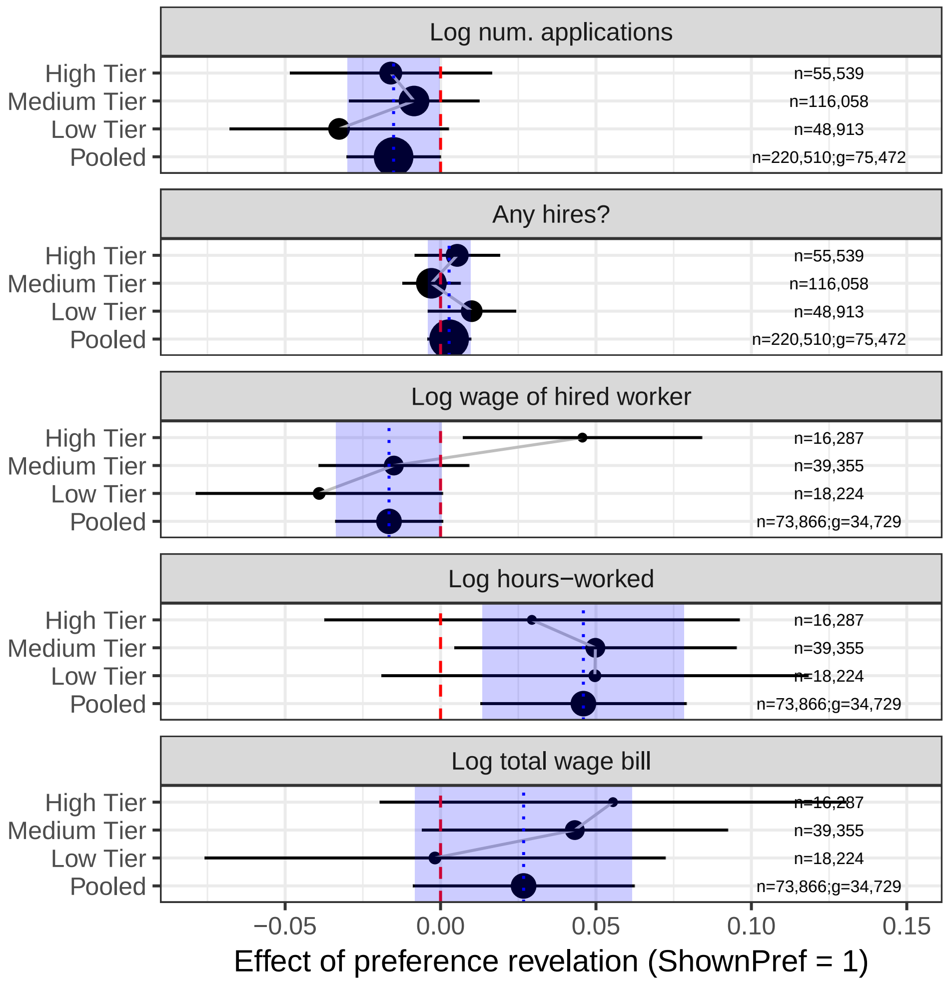

Notes: This figure shows the effects of revealing employer preferences, ShownPref = 1, on a number of outcomes. The sample consists of all job openings from both the ambiguous and explicit arms. Each point estimate is surrounded by at a 95% CI.

6.2 Subjective match outcomes

As we discussed, a problem with the hours-worked and wage bill measures of match quality is that both could be inflated by hours-padding. As such, we might see buyers buying more hours or paying more, but being less satisfied. For these reasons, we also use *ex post* subjective feedback ratings as an outcome.

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

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

From Figure 10, we can see that there is little change in the feedback to the worker. For the feedback to the employer, there is some evidence of better feedback to high tier employers and lower feedback to low tier employers who had their preference revealed. This would be consistent with worker feedback increasing in the hourly wage received, perhaps due to feeling grateful to the employer for the higher wage (Akerlof, 1982). 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.

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

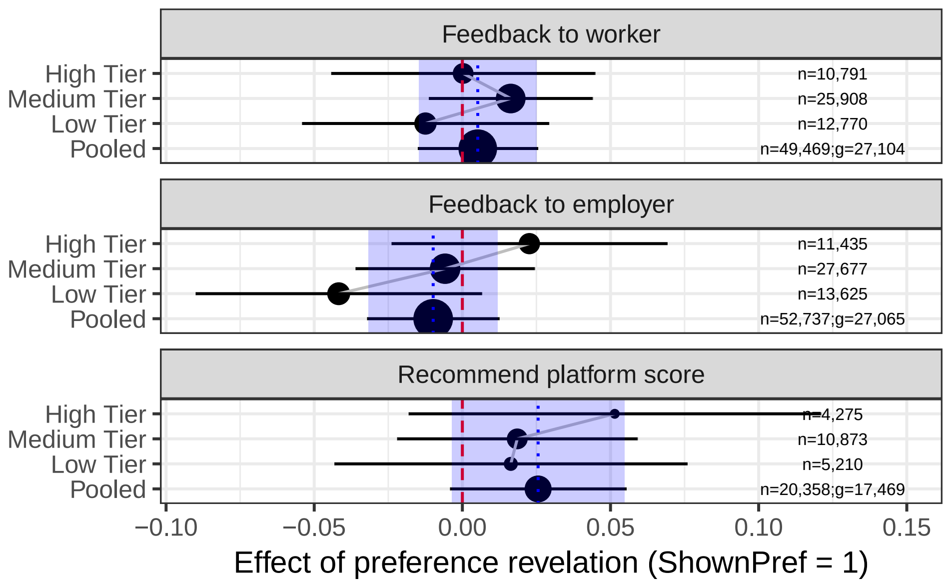

Notes: This figure shows the effects of revealing employer vertical preferences on various feedback measures. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with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

6.3 Discussion of match outcome results

Our results suggest that the separating equilibrium likely improved match quality, at least as measured by our quantity outcomes. This raises the more general question: under what conditions can we expect an overall increase in welfare in the separating equilibrium? To obtain some insight into this question, we return to the simplified framework introduced in Section 5.6.

If we presume that workers' payoffs are only the wages they earn, then since wages are just transfers, overall market welfare in the equilibrium is simply the average value of $v_t\theta$ over all employers (where t is the employer type and θ is the quality of the hired worker). Suppose there is a unit mass of employers in the market, with a fraction μ being high type, and $1-\mu$ being low type. It follows that aggregate market welfare in the separating equilibrium is:

$$\mu v_H q_H + (1-\mu)v_L q_L. \tag{8}$$

On the other hand, suppose in the pooling equilibrium that average hired worker quality is q; we make the simplifying assumption that this is independent of the employer type, since no signaling takes place in the pooling equilibrium.[11] Thus welfare in the pooling equilibrium is:

$$(\mu v_H + (1-\mu)v_L)q. \tag{9}$$

Welfare in a separating equilibrium dominates welfare in a pooling equilibrium when

$$\alpha_H q_H + \alpha_L q_L > q, \tag{10}$$

where:

$$\alpha_H = \frac{\mu v_H}{\mu v_H + (1-\mu)v_L}, \quad \alpha_L = \frac{(1-\mu)v_L}{\mu v_H + (1-\mu)v_L}.$$

We can interpret α_H (resp., α_L) as the *value-weighted* fraction of high (resp., low) types in the marketplace. As one simple case where Equation 10 holds,

[11]Note this may not be completely accurate, since wage negotiation is endogenous, and employers' type may be partially revealed by wage negotiations.

suppose that:

$$\mu q_H + (1-\mu) q_L \geq q.$$

In other words, this implies that the average quality of hired workers in the separating equilibrium is at least as high as the average quality of hired workers in the pooling equilibrium. For example, in supply-constrained markets, where most workers are hired, we would expect that this condition approximately holds. As long as $q_H > q_L$, then since $v_H > v_L$ it is straightforward to verify that Equation 10 holds.

The preceding discussion provides some insight into how to interpret the suggested welfare increase in our setting. Essentially, the welfare gains to high type employers are sufficiently high to offset the welfare losses of low type employers, either because there is a sufficient differentiation in quality ($q_H \gg q_L$), or there is a sufficient mass of high type employers ($\alpha_H \gg \alpha_L$).

6.4 Separating equilibrium in the long-run

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 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 A.6, 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

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

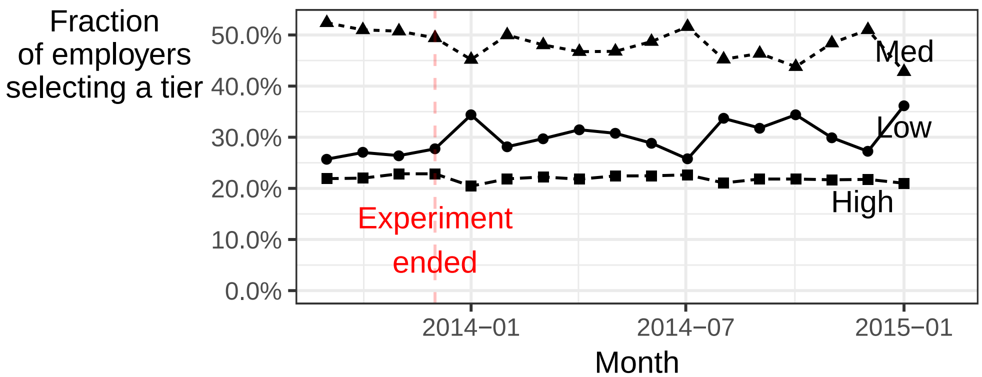

Notes: This figure shows the fraction of job openings each month selecting each of the three possible tiers. When posting a job opening, employers had to select from one of three "tiers" to describe the kinds of applicants they were most interested in: (1) Entry level: "I am looking for [workers] with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

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 11 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.

7 Conclusion

Platform-engineered signaling opportunities can move designed markets to more desirable equilibria. 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%. 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. 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 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 signaling "treatment."

References

Acemoglu, Daron, Mohamed Mostagir, and Asuman Ozdaglar, "Flash organizations," *Working paper*, 2017.

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

Ashlagi, Itai, Mark Braverman, Yash Kanoria, and Peng Shi, "Communication requirements and informative signaling in matching markets," *Working paper*, 2017.

Autor, David, ed., *Studies of labor market intermediation*, University of Chicago Press, 2009.

Azar, José, Ioana Marinescu, and Marshall I Steinbaum, "Labor market concentration," *Working paper*, 2017.

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," *Working paper*, 2017.

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.

Blanchflower, David G, Andrew J Oswald, and Peter Sanfey, "Wages, profits, and rent-sharing," *The Quarterly Journal of Economics*, 1996, *111* (1), 227–251.

Bursztyn, Leonardo, Thomas Fujiwara, and Amanda Pallais, "'Acting wife': Marriage market incentives and labor market investments," *American Economic Review*, 2017, *107* (11), 3288–3319.

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.

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women," *The Quarterly Journal of Economics*, 2015, *131* (2), 633–686.

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.

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

Farrell, Joseph and Matthew Rabin, "Cheap talk," *The Journal of Economic Perspectives*, 1996, *10* (3), 103–118.

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

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.

Goldschmidt, Deborah and Johannes F Schmieder, "The rise of domestic outsourcing and the evolution of the German wage structure," *The Quarterly Journal of Economics*, 2017, *132* (3), 1165–1217.

Horton, John J, "Online labor markets," *Internet and Network Economics: 6th International Workshop*, 2010.

_ , "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*, Forthcoming.

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.

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.

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

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

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

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

Spence, Michael, "Job market signaling," *The Quarterly Journal of Economics*, 1973, *87* (3), 355–374.

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.

Stigler, George J, "The economics of information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

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.

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

Weil, David, *The fissured workplace*, Harvard University Press, 2014.

A Appendix

A.1 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 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.

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

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).

A.2 Should workers consider changed applicant pool size when bidding?

As we saw, signal revelation had some effect on applicant pool size, particularly in the low tier. A natural question is whether these different pool sizes influenced wage bids. If workers thought they faced less competition, all else equal, they have an incentive to bid up. In settings where it can be examined, endogenous entry has proven empirically important (Bajari and Hortaçsu, 2003). However, in contrast to common value auctions, there is presumably a much greater role of idiosyncratic worker-specific surplus in the case of hiring, muting the effects.

Whether this consideration is important in practice is an empirical question—the competition effects might be sufficiently small that the worker does not have to consider them from a worker's perspective. To test whether anticipated pool size "matters," we can test what workers do naturally, in the sense that we could consider how they adjust their bidding behavior on a job-to-job basis. Ideally we would estimate a regression of the form

$$\log w_{ij} = \alpha_i + \beta_1 \log A_j + \epsilon \tag{11}$$

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 Equation 11, 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^{***}	
		(0.005)	
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.[12] We include day-specific fixed effects in the second stage. The identifying assumption is that there is day-to-day variation in the number of jobs posted and the number of workers active that changes the number of applicants per job for exogenous reasons. In Column (2), we report the first stage of the IV estimate. We can see that is a powerful instrument,

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

with a conditional F-statistic of 18584.4.

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.

A.3 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.

A.4 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.

From Column (1), we see evidence of an increase in per-application win

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

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

	Dependent variable:	
	Hired	Expected wage
	(1)	(2)
ShownPref	0.001***	0.003
	(0.0001)	(0.002)
Worker FE	Y	Y
Observations	7,052,680	7,052,680
R^2	0.099	0.257

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

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 3.3% 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 1.2% 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.

A.5 Match outcome result robustness to sample definition

Figure 12 reports results for a number of outcomes using different sample definitions.

A.6 Tier choice over time

As employers can and do post multiple job openings during the experiment, we can observe if their tier choices change over time. Note that we only use the first observation for our experimental analysis. Table 7 reports estimates where the outcome is an indicator for a particular tier choice, and the key explanatory variable is the ordering of the opening, or OrderRank. The regressions show no change in probability of selecting low tier over time.

However, there is some movement away from the medium tier, into the high tier. In Column (4), the order rank is interacted with the treatment assignment—there is no evidence that treatment assigned affected the choice over time.

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

Figure 12: Effects of revealing employer vertical preferences on job opening outcomes

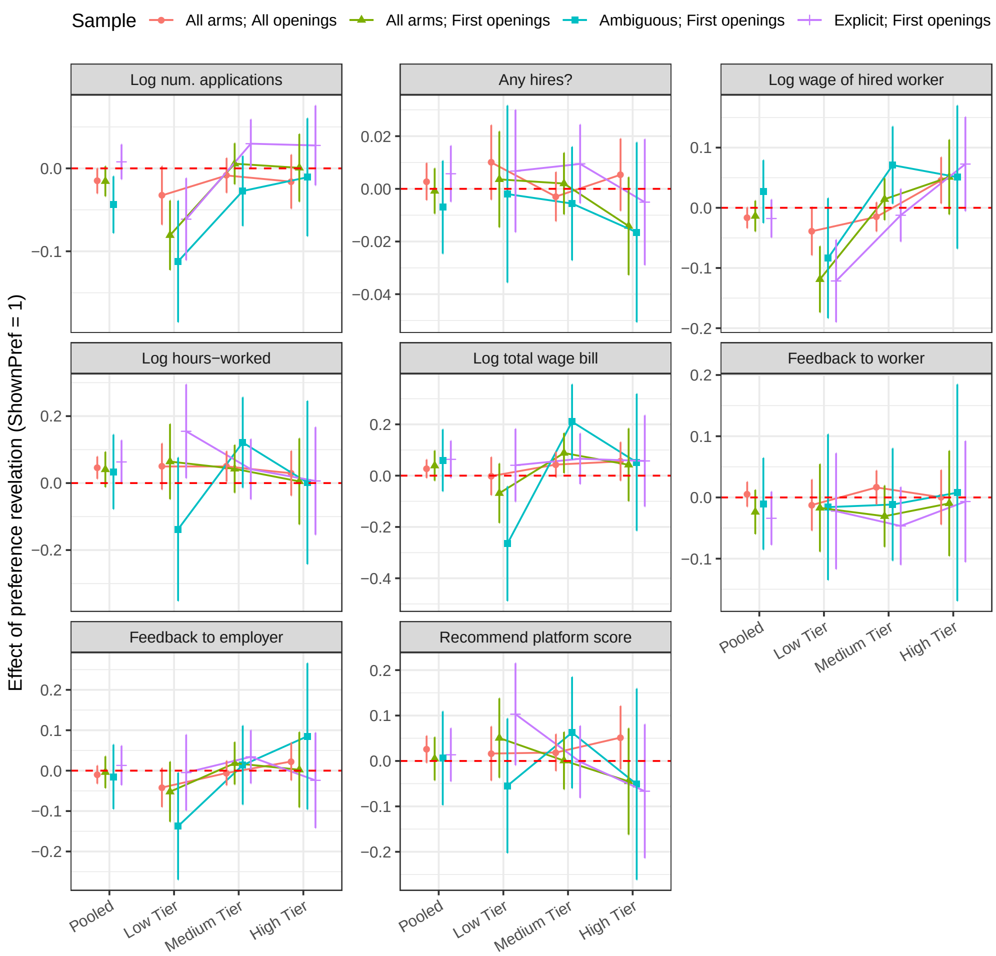

Notes: This figure shows the effects of revealing employer preferences, ShownPref = 1, on a number of outcomes, for several different samples. Each point estimate is surrounded by at a 95% CI.

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

	Dependent variable:			
	LowTier	MedTier	HighTier	
	(1)	(2)	(3)	(4)
OpeningRank	-0.001	-0.002^{***}	0.002^{***}	0.002^{***}
	(0.0004)	(0.001)	(0.0004)	(0.0005)
ShownPref				-0.030
				(0.035)
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: $\dagger{:}p < 0.10$, $*{:}p < 0.05$, $**{:}p < 0.01$, $*{*}*{:}p \leq 0.001$.