

Costly Capacity Signaling Increases Matching Efficiency: Evidence from a Field Experiment*

Apostolos Filippas John J. Horton Prasanna Parasurama Diego Urraca

September 20, 2024

Abstract

Buyers pursuing unavailable sellers is a common source of inefficiency in matching markets. We report the results of a field experiment in a large online labor market where workers could pay to signal higher capacity via a badge that simply said “available now.” All workers could rent this signal, but only randomly treated employers could see it. We find that workers who rented this signal were positively selected, treated employers sought them out more, and matching efficiency increased. We show empirically that mere statements about worker capacity had become uninformative in this marketplace, and we develop a dynamic matching model that explains why costly signaling is necessary to facilitate this coordination. Two years after the experiment, we show that workers renting this signal continue receiving substantially higher employer attention.

JEL Codes: D82, D83, J01, L8, M3

*We gratefully acknowledge generous financial support from the NET Institute (www.netinst.org). Thanks to seminar participants at NYU Stern and Microsoft Research New England. Thanks to Susan Athey, Kyle Bagwell, Michael Ostrovsky, Adam Ozimek, Brad Shapiro, and Robert Somogyi for their helpful comments and suggestions. “Diego Urraca” is a pen name for an employee of the marketplace where this experiment was conducted, and is used to preserve the anonymity of the platform. This work was not subject to prior review and reflects only the opinions of the authors. The latest version of this paper can be found at <https://www.apostolos-filippas.com>.

1 Introduction

In many matching markets, sellers signal that they are “on the market” in order to attract buyer attention and facilitate trade. Examples of such signals include availability disclosures on dating apps, “accepting new patients” banners on doctors’ websites, and “open to work” badges on LinkedIn. One problem with these signals is that sellers often have insufficient incentive to reveal their true status: buyer attention and offers are valuable *per se* to sellers even if not pursued.¹ This creates a market inefficiency because pursuing unavailable sellers is often wasteful (Horton, 2019; Fradkin, 2023).

The goal of this paper is to describe and evaluate a market design approach to remedy this problem. Our approach is to introduce a costly signal that sellers can rent when they want to indicate that they have higher capacity. If successfully implemented, costly signaling can make search less “random” and more “directed,” thereby increasing market efficiency (Mortensen and Pissarides, 1994; Wright, Kircher, Julien and Guerrieri, 2021). In practice, however, this mechanism will increase market efficiency only if it creates a separating equilibrium where higher-capacity sellers rent the signal and lower-capacity sellers do not. If lower-capacity sellers rent the signal, buyers will soon learn to simply ignore it; and even if only higher-capacity sellers rent the signal, buyers need to view the signal as a positive sign of capacity and not as a sign of desperation.²

We evaluate the effectiveness of costly capacity signaling empirically, through a field experiment in a large online marketplace for services. For the first time, workers (sellers) were given the opportunity to pay to rent a costly signal of capacity.³ The signal took the form of a “badge” with the text “available now,” and appeared next to a worker in the platform’s search results. The employers (buyers) could see a notice that workers had to pay to rent this badge. Importantly, renting the badge did not give workers greater visibility to employers: it did not change search rankings nor the size of the displays in the search results (Edelman, Ostrovsky and Schwarz, 2007; Athey and Ellison, 2011; Decarolis and Rovigatti, 2021). During the experiment, all workers could rent this badge but only randomly treated employers could see the badge. This random exposure to the badge information allows us to estimate the causal effects of costly capacity signaling on employer and worker outcomes.

Our key results are summarized as follows. First, we find positive effects of costly capacity signaling on employer outcomes. Treated employers (i) sought out badge-renting workers more by inviting them to apply to their projects, (ii) sent out more invites to workers in total,

¹For example, job offers are valuable even when not taken, as they can be used to negotiate higher pay in one’s current job.

²For example, there is an ongoing debate on whether LinkedIn’s “open to work” badge is beneficial to job-seekers (see <https://www.businessinsider.com/linkedin-open-to-work-banner-badge-job-pros-cons-2024-9>.)

³We use the terms “workers” (sellers) and “employers” (buyers) following the online labor market literature, not as a comment on the legal status of the relationship between the parties involved in these transactions.

(iii) received more positive responses to their invites, and (iv) were more likely to form contracts with workers. Critically, this increase was not at the expense of workers who did not rent the badge—there was a *net* increase in contracts of about 2%. We also find positive effects on worker outcomes: renting the badge increased the likelihood of a worker receiving an invite and forming a contract with an employer.

We next describe our evidence in more detail. Because we observe which workers chose to rent the badge, we can compare badge renters to non badge renters. We find strong evidence that badge renters had higher capacities to take on more work. Counterintuitively, workers renting the badge were busier on average—they had more active contracts and were already receiving more employer invites. Despite their greater busyness, badge renters were much more likely to respond positively to an employer’s invite by submitting a proposal to the project listing. As such, a naive algorithmic approach that directed more employer attention to less busy workers would likely exacerbate the problem of employers pursuing unavailable workers.⁴ Consistent with this virtuous selection into badge renting, treated employers who could see the badge were 4.59% more likely to contact badge renters, and sent on average 9.71% more invites to badge renters. This led to a net increase in matches: treated employers formed 2.63% more contracts compared to the control group.

Non-badge renters were not crowded out by badge renters, at least in aggregate. The reason is that treated employers sent more recruiting invites in total, presumably because the added information made more workers appear worthy enough to send an invite to. This lack of crowd-out matters because signaling that solely redirects employer attention from one worker to another might have little welfare import.

An important question is whether signaling has to be costly to reveal capacity information. The answer lies in the incentives faced by workers. From the worker’s perspective, invites are options, and options are valuable even if not pursued. When the cost of receiving an invite is low, workers have little incentive to reveal their capacity truthfully. Before the experiment, workers could indicate they had high capacity at no cost. This “free” signaling had existed for over a decade in this marketplace (Horton, 2019). We show that, before the experiment, nearly all workers stated that they had high capacities to take on more work and rarely changed their capacity status, thereby making the “free” signal uninformative.

One potential concern with our results is that the benefits of costly capacity signaling may dissipate in the long run. For example, if employers only seek out badge renters because badges are novel or eye-catching, we would expect that the positive effects of costly capacity signaling would diminish greatly over time, especially if badge renters are adversely selected. We provide evidence against this view by examining the effects of costly capacity signaling over time. Using a within-worker analysis, we find that two years after the platform-wide roll-out

⁴It was empirically true in our case that “if you want something done, ask a busy person.”

of costly signaling, badge renters continued to receive about 50% more employer invites. This evidence suggests that costly signaling changes the economics of the platform and its effects are not due to a change in the badge renters' salience.

To rationalize our results and provide a welfare analysis, we develop a dynamic model of a matching market. In this model, “busy” workers are more likely to have higher costs than “available” workers, but employers search for workers without knowing these realized costs. In the equilibrium of a market without costly signaling, welfare is reduced because employers attempt to transact with low-capacity workers. We show that introducing costly signaling creates a separating equilibrium where all workers with capacities above a threshold choose to signal, and all other workers do not. As a result, employers direct a higher degree of their recruiting efforts towards workers with higher capacities, and matching efficiency increases. In short, costly signaling increases efficiency in our model by solving a coordination failure.

Our model also rationalizes why a separating equilibrium can exist in our setting. A busy worker may still choose to signal if the cost of signaling is less than the marginal value of employer invites and offers, leading to a pooling equilibrium. In the model, signaling only affects the likelihood of receiving an invite, but not the likelihood of contract formation conditional on an invite. This means that even if busy workers signal to increase their chances of receiving an invite, they will be less likely to form a contract once they receive the invite due to higher costs. This prevents busy workers from signaling in the first place, leading to a separating equilibrium where only available workers signal. Our empirical results are consistent with this assumption: badge renting increases the likelihood of receiving an invite, but has no significant effect on the likelihood of contract formation conditional on an invite.

The main contribution of this paper is to show that costly signaling can remedy a common market inefficiency of buyers pursuing unavailable sellers. Costly signaling coordinates employers and workers and results in higher matching formation. Our context is one particular digital market, but our results are informative about the role of costly signaling in markets more generally. First, the economic problem of buyer uncertainty about seller capacity or suitability is ubiquitous. Insofar as not all sellers are equally interested in more sales at a moment in time, costly signaling can play a role in directing search ([Bagwell and Ramey, 1994](#)). Second, the economic problem of strategically missing information that we describe is quite general. The information revealed by costly signaling was missing for economic reasons—not technical reasons that digitization alone can solve.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes our study's empirical context and experimental design. We examine the effects of the treatment for employers in Section 4, and for workers in Section 5. Section 6 shows the long-run effects of costly signaling. Section 7 presents a simple model of costly signaling that can rationalize our results. We conclude in Section 8 with thoughts on future research directions.

2 Related work

This paper contributes to the market design literature on improving matching efficiency in two-sided platforms. Prior studies have examined the effects of information disclosure (Tadelis and Zettelmeyer, 2015; Bhole, Fradkin and Horton, 2021; Horton, 2019; Lewis, 2011; Huang, Burtch, He and Hong, 2022), the platform’s recommendation algorithm (Horton, 2017; Fradkin, 2023; Jung, Lim, Lee and Kim, 2022), signaling and advertising (Sahni and Nair, 2020; Yang, Sahni, Nair and Xiong, 2023), and reputation systems (Filippas, Horton and Golden, 2022; Fradkin and Holtz, 2023) on matching in online marketplaces.

This work is most closely related to the experimental studies that take a market design approach to address congestion in markets where capacity constraints are important (e.g., in online labor markets, dating markets, and short-term rental markets). Huang et al. (2022) show that disclosing demand information (i.e., popularity) increases matching efficiency in an online dating platform. Jung et al. (2022) find that increasing the choice capacity of female users in a dating platform increases matching outcomes. Bhole et al. (2021) show that disclosing the number of applications for a job increases overall applications by redirecting job seekers to apply to jobs with fewer applications. Horton (2019) shows the self-disclosure of worker capacity increases the number of matches in an online labor market—but as we show in this paper, this effect is short-lived. This paper shows that costly signaling of capacity can increase matching efficiency, and is sustained in the long run, by directing employer search towards workers with higher capacities.

Costly signaling works in our setting because the act of signaling, in and of itself, reveals credible information about a worker’s capacity to employers (Spence, 1973; Nelson, 1974; Kihlstrom and Riordan, 1984; Milgrom and Roberts, 1986). This is because high capacity workers have more to gain from signaling than low capacity workers. This prevents low capacity workers from mimicking high capacity workers, creating a separating equilibrium. We formalize this in a dynamic model of a matching market. In the classical models of signaling (Spence, 1973; Nelson, 1974; Kihlstrom and Riordan, 1984), sellers differ in some vertical attribute (e.g., quality), which is presumed to be exogenous. However, in our setting, workers (sellers) differ in their capacity, and this capacity is endogenous. In our model, we do not assume that workers differ in capacity but rather micro-found it with a matching process that affects worker capacity.

3 Empirical context and experimental design

Our study is conducted in a large online labor market (Horton, 2010; Agrawal, Horton, Lacetera and Lyons, 2015; Horton, Kerr and Stanton, 2017). In online labor markets, employers form contracts with workers to complete projects remotely. These projects include computer pro-

programming, graphic design, data entry, research, and writing. Each market differs in scope and focus, but platforms commonly provide ancillary services, including maintaining project listings, hosting employer and worker profile pages, arbitrating disputes, certifying worker skills, and maintaining feedback systems (Filippas, Horton and Zeckhauser, 2020). They are broadly similar to a host of other online marketplaces that have arisen in recent years (Einav, Farronato and Levin, 2016).

Several features of conventional labor markets also exist in our context. Employers and workers are free to enter and exit the market anytime. Employers post project descriptions, and workers apply to projects. Employers and workers can negotiate over wages and form contracts. More generally, employers and workers face substantial search frictions (Horton, 2017, 2019), barriers to entry (Pallais, 2013; Stanton and Thomas, 2016), and information asymmetries during the matching process (Benson, Sojourner and Umyarov, 2019; Filippas et al., 2022).

3.1 Search and matching in the market

The matching process can be initiated by either the workers or by the employers. Workers can initiate the matching process by searching for and applying to projects. To do so, workers can view an algorithmically determined ranking of all available projects of interest, access project descriptions and employer profiles, and apply to projects by sending a proposal and placing a wage bid. Applications use up “coins,” an in-platform currency sold through the platform and costing \$0.15 each (Filippas, Horton and Zeckhauser, 2023).

Employers may also initiate the matching process by inviting workers to apply to their project. Upon posting a project, employers can view rankings of workers who are determined algorithmically to be good matches for their projects. Employers can further explore worker profiles and, if interested, invite workers to submit proposals for their project. We call this employer invitation to apply an “invite.” For each project post, employers may send a fixed number of free invites and can purchase the right to send additional invites. A worker application following an employer invite uses up no coins. As such, employer invites are valuable to workers because they both reduce application costs and because they signal employer intent that might lead to a paid project.

Employer invites are common on the platform: the majority of the employers send at least one invite after they post a project. Because invites are limited and searching for workers is costly, employers prefer to send invites to workers who are likely to accept them. As such, we would expect employers to try to infer worker capacity to take on new projects.

3.2 The near uselessness of costless capacity signaling

In its earlier days, the platform had introduced a signaling mechanism that allowed workers to self-report their “availability”—their capacity to take on new projects—on their profiles. The availability signaling feature allowed workers to put one of three messages on their profiles about their availability: (1) “More than 30 hrs/week,” (2) “Less than 30 hrs/week,” and (3) “As Needed - Open to Offers.” Workers could change their availability at any point in time. [Horton \(2019\)](#) showed that this feature was effective: workers signaling high capacity received more employer invites, rejected fewer invites, and were more likely to form a contract.


Although the self-reported worker availability feature appeared promising initially, its effectiveness deteriorated. [Figure 1a](#) shows the distribution of self-reported availability and their respective invite acceptance rates using a cross-section spanning half a year before the commencement of the experiment. Workers overwhelmingly reported that they had high capacities: 88.6% reported they were available 30+ hrs/week, 4.2% reported less than 30 hrs/week, and 7.3% reported they were available as needed (see [Figure 1a](#) panel (i)). This is in stark contrast with [Horton \(2019\)](#), who showed that when this feature was introduced, 45% of workers reported 30+ hrs/week, 33% reported less than 30 hrs/week, and 22% reported they were available as needed. Despite, the vast majority of workers reporting high capacity, their acceptance rate (37.9%) was not much higher than those who reported lower capacities—37.8% for those who reported “as needed” and 25.2% for those who reported “less than 30 hrs/week”.

Workers also did not change their availability status often. [Figure 1b](#) plots the percentage of workers who changed their self-reported availability status each month, using data spanning more than two years before the experiment. The percentage of workers changing their availability status decreased over time from about 4.5% to about 0.6%.


Taken together, the evidence above suggests that costless capacity signaling was ineffective in the long run: (i) workers seldom changed their self-reported availability and likely overstated their capacity to take on new projects; (ii) employers were aware of this misreporting and responded by “mixing” their invites to workers with lower self-reported availability statuses as well. In other words, simply allowing workers to signal higher capacity “for free” did not work in the long run.

Figure 1: Statistics on costless capacity signaling

(a) Self-reported availability and invite acceptance rates



(b) Workers changing their self-reported availability over time



Notes: This figure reports statistics on the platform’s costless capacity signaling feature. Figure 1a plots summary statistics for the workers’ self-reported availability. Panel (i) plots the distribution of workers’ self-reported availability. Panel (ii) plots the invite acceptance rates for each availability choice. Figure 1b plots the percentage of active workers who changed their self-report availability at least once each month. The period covers about 2.5 years before the commencement of the experiment. The sample comprises active workers, defined as those workers who applied for at least one project during the current or the previous month.

3.3 Experiment

The platform introduced a mechanism for costly signaling of worker capacity via an experiment. During the experiment, workers in select technical categories became eligible to engage in costly signaling by renting a “badge.” The badge displayed the phrase “Available Now” on the worker’s profile and in all search tiles where the worker appeared. Importantly, the badge had no other effect. The price of renting the badge was fixed to 2 coins per week throughout the experiment. Workers were notified of this opportunity upon logging into the platform. A total of 243,126 workers were engaged in the experiment, that is, they became eligible to rent the badge. There was high badge uptake—by the end of our data, 39.8% of the active workers were renting the badge (see Appendix A.3).

Employers were randomized into a treatment and a control group upon posting a project in the same select technical categories. The only difference between the two groups was that treated employers could see the worker badges, but control employers could not (see Appendix A.1 for more details on the employer search interface).

The experiment began on July 26, 2021 and ended on October 01, 2021. A total of 84,425 employers were engaged in the experiment. Of that total, 42,474 employers (50.31%) were allocated to the treatment group, and 41,951 (49.69%) to the control group. The experimental groups were well-balanced across several pre-experimental observables (see Appendix A.2). Table 1 reports summary statistics of employer-, project-, and worker-level outcomes during the experimental period.

Recall that badge renting did not affect the workers' placements in the employers' search rankings. These rankings were determined algorithmically and did not take badge renting into account throughout the experiment. As such, it is possible that employers could have little exposure to badge renters, depending on the interplay of search rankings and badge renting decisions. However, this was not the case: badge renters made up about 50% of the workers that employers were exposed to (see Appendix A.3).

Table 1: Summary statistics during the experimental period

	Mean	Median	SD
<i>Employer-level</i>			
Projects posted	1.48	1.00	1.48
<i>Project-level</i>			
Invites sent	3.17	1.00	8.12
Invites sent to badge renters	1.44	0.00	3.31
Proposals received following invite	1.38	0.00	2.82
Overall proposals received	14.90	9.00	17.92
Contracts formed	0.30	0.00	0.55
<i>Worker-level</i>			
Invites received	5.61	1.00	15.52
Proposals sent following invite	2.83	1.00	10.38
Overall proposals sent	26.60	6.00	93.59
Contracts formed	0.53	0.00	1.49

Notes: This table reports summary statistics of employer-level, project-level, and worker-level outcomes. The employer sample includes all employers who posted at least one project during the experimental period. The project sample includes all projects posted during the experimental period. The worker sample includes all workers who applied to at least one project during the experimental period.

4 The effects of costly capacity signaling on employers

We begin our analysis by examining the effects of costly capacity signaling on employer outcomes. We consider 5 main employer outcomes: (a) invites sent to workers, (b) invites sent to badge renters, (c) proposals received following invite, (d) overall proposals received, and (e) contracts formed. For each outcome, we examine the treatment effect on the extensive margin (whether there was any change in the outcome) and intensive margin (how much the outcome changed). To get the extensive margin, we apply the indicator variable transformation on each outcome.

Because some employers may have posted many projects during the experiment, we analyze the experiment using two different samples: (1) “All projects” uses the entire sample of projects posted during the experiment, and (2) “First project” restricts the sample to the first project each employer posted during the experiment. Our preferred specification is “All projects” because it captures the most behavior; comparing it to the “First project” specification allows us to examine how the effects of the treatment varied over time.

For each sample, we regress each outcome on indicators for the treatment, i.e.,

$$y_{jp} = \beta_0 + \beta_1 \text{TRT}_j + \epsilon,$$

where y_{jp} is the outcome of interest for project p posted by employer j , TRT_j indicates whether employer j was assigned to the treatment group, and ϵ is an error term. For the “All projects” specification, we cluster standard errors at the employer-level.

Figure 2 reports the estimated effects as percentage changes over the control group outcome, by plotting the least squares estimate $\hat{\beta}_1/\hat{\beta}_0$ for each of the specifications. We also replicate the analysis with Poisson regression. All regression tables can be found in Appendix A.4.

Treated employers were more likely to send at least one invite: compared to a baseline of 53.09% for the control group, the increase was 1.33 percentage points (2.51%)⁵. Similarly, treated employers sent 0.24 (7.74%) more invites per post. Notably, treated employers recruited badge renters more. Employers who were able to see badges were 1.86 percentage points (4.59%) more likely to send an invite to a badge renter, and sent on average 0.13 more invites (9.71%) to badge renters.

Treated employers received more positive responses to their invites. Treated employers were 1.33 percentage points (3.04%) more likely to receive a worker proposal following an invite, and received 0.07 (5.03%) more such proposals.

Costly capacity signaling increased matching efficiency. Treated employers were 0.83 percentage points (3.09%) more likely to form at least one contract, and formed 0.01 (2.63%) more

⁵Note that here—and throughout the paper—for differences in levels where the outcome is discussed as a fraction, we label level differences as “percentage points.” For percentage changes with respect to the outcome of the treatment group, we use the “%” symbol.

contracts.

It is worth noting that almost all estimates in the “First project” sample are smaller than in the “All project” sample. This suggests that the experiment’s effects are unlikely to be transitory, as they seem to increase in magnitude over time.

4.1 Treated employers did not substitute away from non badge renters

Treated employers directed their recruiting efforts towards badge renters. Next, we examine whether this was at the expense of non badge renters. To do so, we compare the treatment effects on invites sent to badge renters and non badge renters, and report the results in Table 2.

Interestingly, we do not see any net decline in invites going to non badge renters. This is surprising, as we might expect that employers who send few invites may shift their attention to badge renters, substituting away from non badge renters. To build confidence in this result, we also compare the distributions of employer invites by treatment status and worker badge renting status (see Appendix A.5). We find that the treatment increased invites to badge renters; however, there is no discernible shift in the distribution for non badge renters.

Table 2: Treatment effect estimates – invites to badge renters vs. non badge renters

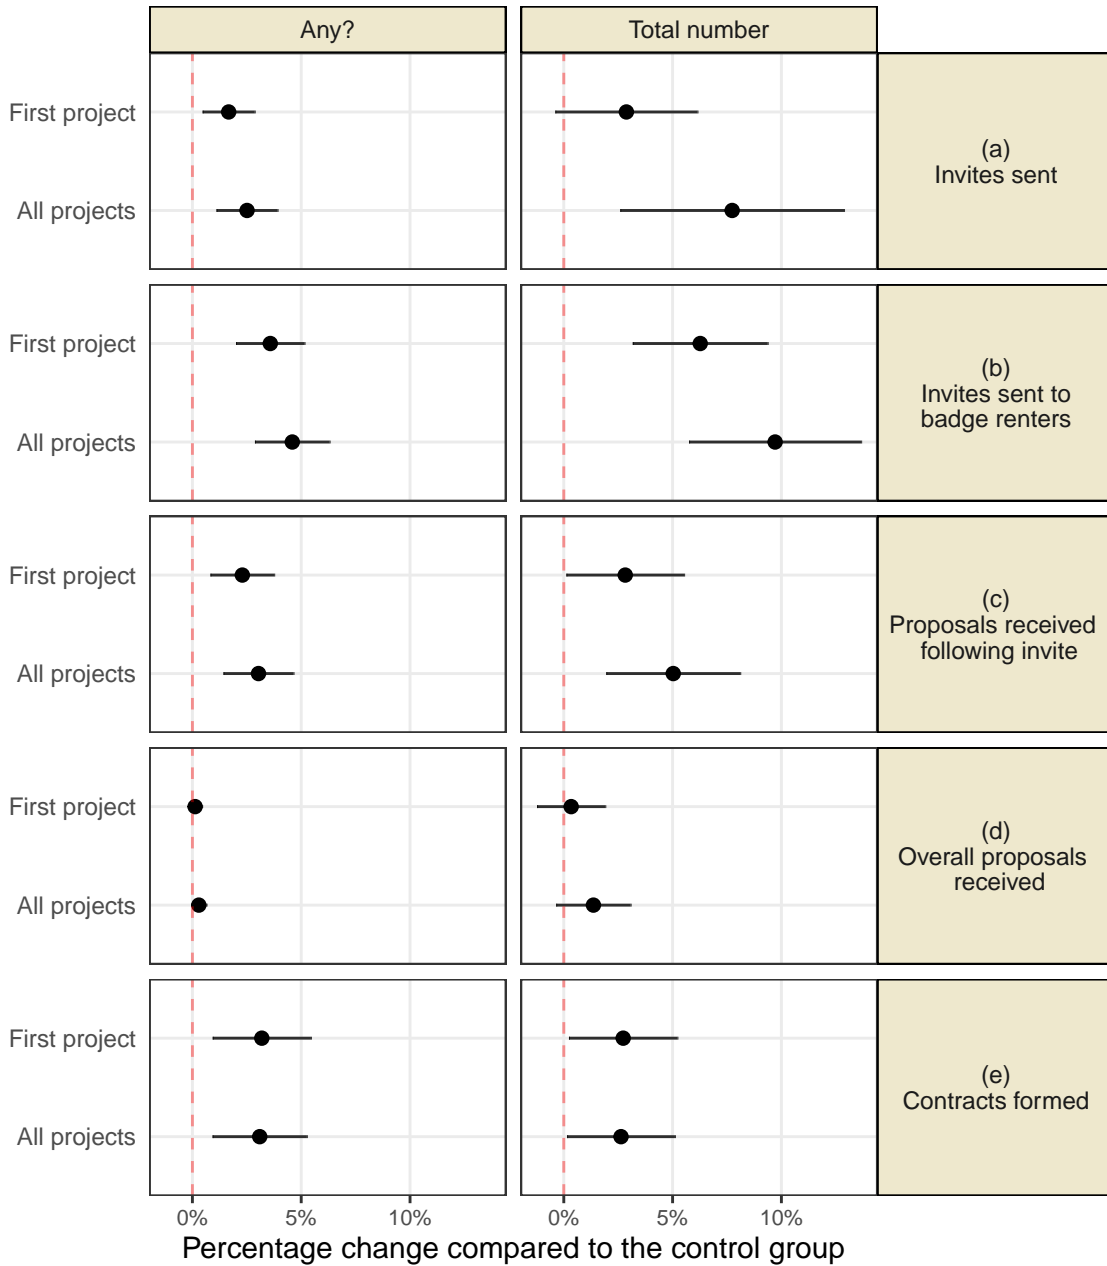
Dependent Variables: Sample:	Invites sent to badge renters		Invites sent to non badge renters	
	All	First	All	First
	(1)	(2)	(3)	(4)
Constant	1.370*** (0.015)	1.397*** (0.015)	1.682*** (0.024)	1.655*** (0.023)
TRT	0.133*** (0.028)	0.088*** (0.022)	0.103‡ (0.056)	0.000 (0.032)
<i>Fit statistics</i>				
Observations	126,550	84,425	126,550	84,425

Clustered (Employer) standard-errors in parentheses

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports regression estimates where the dependent variables are invites sent, and the independent variable is a treatment indicator. Estimates are computed for two different samples: (i) “All projects” uses the entire sample of projects and clusters standard errors at the employer-level, (ii) “First project” only uses each employer’s first project post during the experiment. Columns (1-2) report invites sent to badge renting workers, and Columns (3-4) report invites sent to non badge renting workers.

Figure 2: Treatment effect estimates for employer outcomes



Notes: This figure plots estimates of the treatment effects on employer outcomes, using cross-sectional data. Each panel reports point estimates as the percentage change in the treatment group over the control group, along with a 95% confidence interval. Panels on the left examine the extensive margin effect, with the dependent variable being the indicator variable transformation of each outcome. Panels on the right examine the intensive margin effect, with the dependent variable being the “raw” outcome. Estimates are computed for two different samples: (i) “All projects” uses the entire sample of projects and clusters standard errors at the employer-level, (ii) “First project” only uses each employer’s first project post during the experiment. The regression tables can be found in Appendix A.4.

4.2 Novelty does not explain the effects of the treatment

One concern might be that employer invites to badge renters were due to a novelty effect: badge renter profiles may have stood out to employers because the badge was new to the platform and eye-catching. We provide evidence suggesting this was not the case.

First, in the aggregate, workers did not turn off the badge over time (see Appendix A.3), suggesting that the badge continued to be effective in attracting invites. Second, employers continued to send more invites to badge renters not just in the first project they posted, but in subsequent projects as well (see Figure 2). Third, we show in Section 6 that badge renting was still effective two years after the experiment—by which time the feature had been rolled out platform-wide. Together, this evidence suggests that the effects of costly capacity signaling were not due to novelty.

5 The effects of costly capacity signaling on workers

5.1 Worker selection into costly capacity signaling

We begin our analysis on worker outcomes by first examining the characteristics of workers who select into badge renting. Table 3 compares badge renters and non badge renters based on their pre-experimental and experimental outcomes.⁶ For the purposes of this exercise, we define “badge renters” to be the workers who rented the badge for at least two full days during the experimental period. To focus on workers who were active, we restrict our sample to workers who sent at least one proposal during the experimental period.

Workers who chose to rent the badge look quite different from workers who did not. Before the start of the experiment, badge renters were already applying for more projects, receiving more invites, accepting invites at a higher rate, and forming more contracts. During the experiment, the differences between the two groups remain directionally consistent. Badge renters continued to send more proposals, receive more invites, accept invites at a higher rate, and form more contracts.

These data show that badge renters were positively selected in their capacity to take on new projects. As we will show in Section 7, this positive selection is in line with a signaling equilibrium.

⁶We provide an alternative approach to modeling the workers’ selection into badge renting in Appendix A.7. There, we quantify the relative importance of various factors that predict whether a worker will choose to rent the badge, by reporting the results of a logistic regression where the outcome is an indicator variable for selection into badge renting, and the independent variables are pre-experiment worker attributes and outcomes.

Table 3: A comparison of the characteristics of badge renters and non badge renters

	Non badge renters (mean)	Badge renters (mean)	Difference (percentage)	p-value
BEFORE THE EXPERIMENT (April 20, 2021 - July 20, 2021)				
stated availability	39.37	39.42	0.15%	0.025
invites received	2.95	4.51	34.68%	0
acceptance rate	0.44	0.53	17.5%	0
proposals sent	12.33	19.49	36.74%	0
contracts formed	0.3	0.52	42.4%	0
DURING THE EXPERIMENT (July 26, 2021 - October 01, 2021)				
invites received	4.25	7.23	41.15%	0
acceptance rate	0.47	0.55	13.83%	0
proposals sent	18.96	35.76	46.98%	0
contracts formed	0.35	0.75	53.45%	0
<i>Observation counts</i>	46,845	39,088		

Notes: This table reports profile information and project application statistics before and during the experimental period for badge renters and non badge renters. The sample comprises workers who applied for at least one project during the experimental period. “Badge renters” are defined as workers who rent the badge for at least 48 hours during the same period. For each outcome, we report the mean value for each group, the percentage difference for badge renters, and the p-value of a two-sided test of equal means between the two groups. For each worker, we report the number of invites received, and, conditional on receiving at least one invite, the acceptance rate, the number of proposals sent, and the number of contracts formed.

5.2 Impression-level effects of costly capacity signaling on worker outcomes

Although randomization took place at employer-level, the design of the experiment allows us to estimate the causal effects of badge renting on worker outcomes. To do so, we turn to impression-level data and examine the effects of badge renting on worker outcomes conditional on receiving an impression. An impression occurs when an employer sees a worker on her search interface. We construct a sample where each observation is a worker impression in an employer’s search interface during the experimental period. Our estimation strategy is to regress each worker outcome of interest on (i) the treatment indicator of the employer who posted the project and viewed the impression, (ii) the badge renting status of the worker, and (iii) the interaction of the two terms. Our specification is:

$$y_{ip} = \beta_0 + \beta_1 \text{TRT}_p + \beta_2 \text{BADGE}_{ip} + \beta_3 (\text{TRT}_p \times \text{BADGE}_{ip}) + \epsilon_{ip},$$

where i is a worker, p is a project post, y_{ip} is an outcome for the employer-worker interaction for project p , TRT_p indicates whether employer who posted project p was assigned to the treatment group, BADGE_{ip} indicates whether the worker was renting the badge during that

interaction, and ϵ_{ip} is an error term.

Badge renting is randomly visible to employers due to the randomized assignment, and hence the estimates for the coefficients β_1 and β_3 can be interpreted causally.

It is worth elaborating on the interpretation of these coefficients. β_1 measures the effect of badge visibility on non badge renters in the treated employers' search interfaces. We expect this estimate to be equal to zero unless the treatment had crowd-out or spillover effects. A negative coefficient would indicate a crowd-out effect, e.g., employers who can see the badge substitute away from non badge renters. A positive coefficient would indicate a spillover effect, e.g., employers who can see the badge become more active in their recruiting, which in turn affects non badge renters positively. β_2 measures the difference in outcomes between badge renters and non badge renters in the control employers' search interfaces. β_3 measures the effect of badge visibility on *badge renters* in the treated employers' search interfaces. This is the key coefficient of interest, as it measures the causal effect of badge renting on worker outcomes.

We report the estimated effects in Table 4. The estimate $\hat{\beta}_1$ on TRT is close to zero for all dependent variables, suggesting that the treatment did not have a strong spillover or crowd-out effect. This is consistent with what we learned from our employer-focused analysis showing no reduction in invites sent to non badge renters.

For inquiries, $\hat{\beta}_2$ is negative, meaning that badge renters were less likely to receive an invite compared to non badge renters who appeared in the control employers' search interfaces.⁷ Conditional on receiving an invite, these badge renters were more likely to respond positively to the invite, and they were overall more likely to form a contract.

The estimate $\hat{\beta}_3$ is positive for all outcomes, meaning that badge renting had a positive effect for workers. Renting the badge increased the probability of receiving an invite after getting an impression by 0.39 percentage points (4.83%), and the probability of responding positively to the invite 0.44 percentage points (8.63%). Although the effect is somewhat imprecise, badge renting also increased the probability of forming a contract by 0.02 percentage points (7.72%).

5.3 Total effects of costly capacity signaling on worker outcomes

One shortcoming with the impression-level analysis of Section 5.2 is that workers differ in how many impressions they receive, yet the badge had a fixed cost independent of impressions received. Therefore, in deciding whether or not to rent the badge, what matters for workers is

⁷Note that this seems counter to the results presented in Section 5.1, where we showed that badge renters received more invites overall. This discrepancy is due to the fact that this impression-level analysis conditions on receiving an impression, which is a different sample than the one used in Section 5.1. Badge renters received more invites overall, but conditional on receiving an impression, they were less likely to receive an invite compared to non badge renters.

Table 4: Effects of costly capacity signaling on worker outcomes at the impression level

Dependent Variables:	Worker received invite (1)	Worker applied after invite (2)	Contract formed (3)
CONTROL	0.0802*** (0.0012)	0.0316*** (0.0005)	0.0028*** (0.0001)
TRT	0.0031‡ (0.0018)	0.0000 (0.0006)	-0.0001 (0.0001)
BADGE	-0.0102*** (0.0011)	0.0089*** (0.0006)	0.0005*** (0.0001)
TRT × BADGE	0.0039** (0.0013)	0.0035*** (0.0005)	0.0002‡ (0.0001)
<i>Fit statistics</i>			
Observations	3,427,112	3,427,112	3,427,112

Clustered (Project post & Worker) standard-errors in parentheses

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports OLS regressions where the independent variables are each employers’s treatment status, an indicator for whether a worker viewed in the search was renting the badge, and an interaction term between the two. Observations are on the project post employer impression level. The dependent variables are indicators of whether (1) the the worker received an invite, (2) the worker sent a proposal after receiving the invite, and (3) whether the employer and the worker formed a contract.

not the per-impression effect of badge renting but the net effect over some period of time.

To examine whether badge renting “made sense” workers, we switch our focus from the worker-impression level to the worker-period level. We construct a two-period panel: before and after badge renting was introduced to the market, which we use for a difference-in-differences analysis of the effect of badge renting on worker outcomes.

Table 5 reports summary statistics for the panel. Note that rather than simply having a badge on/off indicator as an independent variable, we now have badge renting days as our key dependent variable. This is useful as badge renters can decide how many days to rent the badge.

To estimate the overall effects of badge renting on worker outcomes, we use the following specification:

$$y_{it} = \alpha_i + \beta_0 \text{EXPPERIOD}_t + \beta_1 (\text{BADGEDAYS}_i \times \text{EXPPERIOD}_t) + \epsilon_i,$$

where y_{it} is the outcome of interest for worker i in period t , α_i is a worker-specific fixed effect, EXPPERIOD is an indicator for the post-period, and BADGEDAYS measures the number of days the worker rents the badge in the post-period. Table 6 reports the results.

Table 5: Summary statistics for worker difference-in-differences panel

	Mean	Median	SD
Badge renter	0.45	0.00	0.50
Contracts formed	0.47	0.00	1.32
Days of badge renting	19.99	0.00	26.34
Invites received	4.63	1.00	13.07

Table 6: Effects of costly capacity signaling for workers at the worker level

Dependent Variables:	Employer inquiries received (1)	log(wage bid) (2)	Contracts formed (3)	Contracts formed (4)
EXPERIOD	1.154*** (0.030)	0.038*** (0.002)	0.050*** (0.005)	-0.036*** (0.006)
EXPERIOD \times BADGEDAYS	0.040*** (0.002)	0.000 (0.000)	0.004*** (0.000)	0.003*** (0.000)
log(proposals sent)				0.456*** (0.006)
<i>Fixed-effects</i>				
Worker	✓	✓	✓	✓
<i>Fit statistics</i>				
Observations	171,866	119,052	171,866	151,024

Clustered (Worker) standard-errors in parentheses

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, †: 0.1*

Notes: This table reports regressions where the dependent variables are an indicator for the experimental period, and the same variable interacted with the number of days each worker rented the badge during the experimental period. The dependent variables are: (1) the number of invites the worker received from employers, (2) log wage bid, (3-4) the number of contracts the worker formed.

Our primary coefficient of interest is, $\text{EXPERIOD} \times \text{BADGEDAYS}$. In Column (1), we can see that each day of badge renting led to approximately 0.04 more invites. The coefficient on EXPERIOD captures this mean of the difference between the two periods. The experimental period lasted 66 days, and the pre-period lasted 91 days, but our sample is made up of workers who applied for at least one project during the experimental period. Because looking for new projects is episodic, we expect the pre-period to be characterized by less activity compared to the experimental period.

In Column (2), the outcome is the average log wage bid. There is no evidence that badge renting workers increased or lowered their wage bids while renting the badge.

In Column (3), the outcome is the number of contracts formed. We find that badge renters form around 0.004 more contracts per day compared to non badge renters. However, the estimate in Column (3) may be too high, as workers might have also increased their project-finding efforts by applying to more projects. This channel could be strongly linked to the decision to rent the badge. Although our worker fixed effect helps control for this, workers could still choose to rent the badge and increase their search intensity simultaneously.

To deal with this issue, in Column (4), we include a control for the number of proposals sent during the period to proxy for search intensity. Although we are controlling for a downstream outcome, this outcome is likely still informative about whether renting the badge was really driving increased business for workers. Including this control reduces the treatment effect by about 1/4, but the estimate remains positive and significant.

In terms of the returns badge renting, our results imply that if an employer invite is worth at least a \$1 and a contract at least \$10, then renting the badge has a positive return on investment at the 2 coins per week price (see Section 3 for more details on coins).

6 The long-run efficiency of costly capacity signaling

So far, we have focused on the time period following the introduction of costly capacity signaling to the market. One concern is that its effectiveness may dissipate in the long run. We next present evidence that costly capacity signaling sustained its efficiency in the long run.

We collected additional data on employer invites and worker badge renting two years after the experiment took place. To stay close to the design of the original experiment, we restrict our sample to workers who received at least one invite during July 2021, and we collect data for these workers for the month of August 2023. It is worth noting that during this time period (i) workers from all platform categories were eligible to rent badges, and (ii) workers' badge renting was still not taken into account when determining worker rankings.

Our empirical approach relies on utilizing within-worker variation in badge renting to estimate its effect on the number of invites received by the workers. Specifically, we estimate the following regression:

$$y_{it} = \beta_0 + \beta_1 \text{BADGE}_{it} + f_i + \tau_t + \epsilon_{it},$$

where y_{it} is the number of invites received by worker i on day t , BADGE_{it} is a dummy variable indicating whether worker i rented the badge on day t , f_i is a fixed effect for worker i , and τ_t is a fixed effect for day t . The coefficient β_1 captures the effect of badge renting on the number of invites. The main identifying assumption is that a worker who engages in badge renting doesn't do so when her profile becomes more attractive or visible to the employers. We cluster standard errors at the individual worker level.

Table 7 reports the results. The first column reports the results of an OLS regression

without fixed effects, and the second column reports the results with fixed effects. The latter estimates suggest that one day of badge renting increases the number of employer invites by about 0.036 per day. This is more than a 50% increase relative to the average number of invites received by workers on days when they did not rent the badge. The point estimate is remarkably close to the experimental estimate (see Table 6). Two years after its introduction, badge renting appears to continue to lead to more employer invites for the workers renting the badge. We take this result as evidence that the virtuous equilibrium we identified in Section 5.1 is sustained in the long run.

Table 7: Effects of costly capacity signaling on workers two years after the experiment

Dependent Variable:	Employer inquiries received	
	(1)	(2)
CONSTANT	0.069*** (0.001)	
BADGE	0.096*** (0.003)	0.036*** (0.001)
<i>Fixed-effects</i>		
Worker		✓
Day		✓
<i>Fit statistics</i>		
Observations	2,118,075	2,118,075

Clustered (Worker) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, †: 0.1*

Notes: This table reports regressions where the dependent variable is the number of employer invites worker i received on day t , and the independent variable is an indicator for badge renting activity by i on day t . The sample consists of all workers who received at least one employer invite in July 2021, and we use August 2023 for the analysis.

7 How costly capacity signaling can increase efficiency

A key result of our experiment is that costly capacity signaling shifted employer attention to more available workers. These workers were more likely to transact, leading to an overall increase in matches formed. In this section, we develop a model that helps us to examine the conditions under which costly capacity signaling can increase market efficiency.

7.1 Model

Consider a labor market in which a unit mass of employers and a unit mass of workers interact to form contracts. Each worker is characterized by her state $s \in \{a, b\}$, where a means the worker is “available” and b means she is “busy.” The worker’s state is private information and affects the worker’s output cost, C_s .

Contract formation between employers and workers is a 2-stage process. First, employers and workers need to meet. Upon a meeting, a contract is formed if the worker can deliver a service that the employer values at $v > 0$, and if the worker’s cost is $C_s \leq v$.

Workers and employers potentially meet in several markets. If the worker cannot signal her availability, only one market exists. But if a successful signaling technology is available, there could be separate markets for busy and available workers. If a market has x employers and y workers, then $m(x, y)$ “meetings” occur. We make the following standard assumptions about the matching function m :

Assumption 1. The matching function $m(x, y)$ is continuously differentiable, quasiconcave, increasing, homogenous of degree 1 (constant returns to scale⁸), with $m(x, 0) = m(0, y) = 0$ and $m(1, 1) \leq 1$.

When a worker in state s and an employer meet, the worker draws a cost $C_s \geq 0$ for completing the project from a distribution with cdf F_s . We assume that a busy worker is more likely to have a higher cost of completing the project than an available worker:

Assumption 2 (Stochastic dominance). For any cost $c \in \mathbb{R}_+$, $F_a(c) = \Pr(C_a \leq c) \geq \Pr(C_b \leq c) = F_b(c)$, where the inequality is strict for $c = v$.

If the value to the employer exceeds the cost to the worker, $C_s \leq v$, then a contract is formed with surplus $v - C_s$. A fraction $\alpha \in (0, 1)$ of the surplus goes to the worker, and the rest to the employer.

Contract formations determine how workers transition between the two states. Time is discrete. A worker who is busy in period t and forms a contract in period t continues being busy in period $t + 1$; otherwise, the worker becomes available in period $t + 1$. A worker available

⁸Constant returns to scale means that $m(\lambda x, \lambda y) = \lambda m(x, y)$ for all $\lambda \geq 0$.

at t who forms a contract in that period becomes busy in $t + 1$; otherwise, the worker remains available in $t + 1$. A new unit mass of myopic employers enters the market at every period.

7.2 Payoffs and the law of motion

Payoffs. An employer and a worker in state s who meet each other form a contract if $C_s \leq v$. Let p_s be the probability of contract formation conditional on a meeting. Clearly, $p_s = F_s(v)$, with $p_a > p_b$. Let w_s be the expected value of the surplus conditional on the meeting, equal to $w_s = \mathbb{E}[\max\{v - C_s, 0\}]$. We will see later that all objects in our model depend on the distributions F_s only through the values of p_s and w_s . For example, the expected payoff to the worker of type s from a meeting is αw_s , and the expected payoff to the employer is $(1 - \alpha)w_s$. The following lemma demonstrates that we can treat the values (p_a, p_b, w_a, w_b) as model fundamentals, with the only restrictions on them being that $p_a > p_b$ and $w_a/w_b \in [1, +\infty)$.

Lemma 1. *For any ratio $w_a/w_b \in [1, +\infty)$ and probabilities p_a, p_b satisfying $p_a > p_b$, there exist a pair of distributions F_a and F_b such that $F_a(v) = p_a$, $F_b(v) = p_b$, and F_b dominates F_a in the sense of Assumption 2.*

Proof. This proof and all other omitted proofs can be found in Appendix B. □

The law of motion. Let $A_t \geq 0$ and $B_t \geq 0$ be the measures of all available and busy workers in period t , with $A_t + B_t = 1$. Suppose that all employers and workers meet in one market. Then the number of contracts M_t in period t is given by $M_t = m(1, 1)B_t p_b + m(1, 1)A_t p_a$. This is because, with random meeting, we expect the fraction of meetings with type s to be proportional to their mass. As the number of contracts is equal to the number of workers forming a contract, and all workers who formed a contract become busy in the next period, the law of motion for the mass of busy workers becomes:

$$B_{t+1} = M_t = m(1, 1)B_t p_b + m(1, 1)A_t p_a. \quad (1)$$

Now suppose there are two distinct markets for available workers and busy workers. In that case, we also need to distinguish between the employers shopping in the market for available workers, with a mass of R_t^a , and those employers shopping for the busy workers, with a mass of R_t^b . The total number of contracts formed is then $M_t = m(R_t^a, A_t)p_a + m(R_t^b, B_t)p_b$, and the law of motion for the number of busy workers is

$$B_{t+1} = m(R_t^a, A_t)p_a + m(R_t^b, B_t)p_b. \quad (2)$$

7.3 Pooling (“no costly signaling”) equilibrium

We study the existence of an equilibrium in which workers cannot credibly signal their availability. We call this the pooling equilibrium. We restrict attention to the economy’s steady state, in which the number of busy and available workers remains fixed over time. With only one market, the agents do not make any choices, making the characterization of the pooling equilibrium straightforward.

Definition 1. A *stationary pooling equilibrium* is a collection $(R^a, R^b, A, B) \in \mathbb{R}_+^4$ such that (1) $A + B = 1$, (2) $R^b = B$, and (3) $B = m(R^a, A)p_a + m(R^b, B)p_b$.

The definition of the stationary pooling equilibrium is written as if there are two separate markets, with A workers and R^a employers shopping in the market for available workers, and B workers and R^b employers shopping in the market for busy workers. This choice will become convenient when we study the costly signaling equilibrium in which there will be two separate markets. In deriving the law of motion (1) for the economy with one market, we noted that the total number of meetings $m(1, 1)$ would be split between available and busy workers with weights A and B , respectively. Constant returns to scale implies $Am(1, 1) = m(A, A)$ and $Bm(1, 1) = m(B, B)$, meaning that we can think of the matching process with one market as if it is actually taking place in two separate markets, where the number of employers and the workers are equal to each other. We now establish the existence and uniqueness of the stationary pooling equilibrium.

Proposition 1. *There exists a unique stationary pooling equilibrium.*

7.4 Costly signaling equilibrium

We now study the existence of a “costly signaling” equilibrium, in which the two types of workers are able to credibly signal their availability through costly signaling. In such an equilibrium, employers and workers must decide which market they want to transact in. While workers are free to transact in any market, we will construct an equilibrium in which they choose the market that matches their state.

If a worker of type s enters the available market, the probability that she meets an employer is $m(R^a, A)/A$ —this is the total number of meetings divided by the number of workers in that market. The expected surplus conditional on a meeting is w_s , which is determined by the worker’s type and not by the market she is in. The worker receives a fraction α of the surplus. Finally, the worker must pay π to signal her availability. Therefore, the payoff from transacting in the “available” market is $U_s(a) = \alpha w_s m(R^a, A)/A - \pi$, and the payoff from transacting in the “busy” market is $U_s(b) = \alpha w_s m(R^b, B)/B$. Note that, as long as $m(R^b, B)$ is positive, workers are guaranteed a positive payoff from participating in the busy market, meaning that the “participation constraint” is slack.

For both available and busy workers to reveal their types, the following two incentive compatibility conditions should hold.

$$\alpha w_a m(R^a, A)/A - \pi \geq \alpha w_a m(R^b, B)/B \quad (3)$$

$$\alpha w_b m(R^b, B)/B \geq \alpha w_b m(R^a, A)/A - \pi \quad (4)$$

These constraints can be equivalently rewritten as two constraints on the price that can support the costly signaling equilibrium—i.e., the price of signaling must be cheap enough for the available workers to want to signal their type but costly enough for busy workers to not want to signal they are available.

$$\alpha w_a \left(\frac{m(R^a, A)}{A} - \frac{m(R^b, B)}{B} \right) \geq \pi \geq \alpha w_b \left(\frac{m(R^a, A)}{A} - \frac{m(R^b, B)}{B} \right). \quad (5)$$

If the allocation (R^a, R^b, A, B) admits prices satisfying (5), then that allocation is incentive-compatible for the workers. Notice that the non-emptiness of the set of signaling prices (5) is equivalent to the condition

$$\frac{m(R^a, A)}{A} \geq \frac{m(R^b, B)}{B} \quad (6)$$

We summarize this in Proposition 2:

Proposition 2. *For a costly signaling equilibrium to exist, a necessary condition is that workers have to enjoy higher meeting rates when they choose to signal.*

We are now ready to define a costly signaling equilibrium and prove its existence.

Definition 2. A *stationary costly signaling equilibrium* is a collection $(R^a, R^b, A, B, \pi) \in \mathbb{R}_+^5$ such that (i) $A + B = 1$, (ii) $m(R^a, A)w_a/R^a = m(R^b, B)w_b/R^b$, (iii) $B = m(R^a, A)p_a + m(R^b, B)p_b$, and (iv) π satisfies condition (5).

In the definition above, condition (ii) is the indifference condition for the employers. Since all employers are identical, they should be indifferent between the two markets in equilibrium. Interestingly, since $w_a \geq w_b$, this implies that employers must have a *lower* meeting rate in the “available” market, but they are compensated by a higher matching rate, conditional upon a meeting.

Proposition 3. *A stationary costly signaling equilibrium exists and in such equilibrium $B \geq R^b$, where the inequality is strict if $w_a > w_b$.*

7.5 Welfare properties

We have characterized the two types of equilibria that a market can have—one where workers can pay to signal their availability and one where they cannot. Next, we study the welfare

properties of the two equilibria.

Social welfare at a point (R^b, B) that lies on the frontier of stationary allocations $B = p_a m(1 - R^b, 1 - B) + p_b m(R^b, B)$ is

$$W(R^b, B) = w_a m(1 - R^b, 1 - B) + w_b m(R^b, B) \quad (7)$$

Definition 3. A stationary allocation is *constrained-efficient* if it maximizes Equation 7 subject to the stationarity constraint.

Characterizing the welfare-maximizing allocation and comparing it to the two types of equilibria that we study is challenging at the level of generality that we have maintained so far. Meetings in the available market produce higher surplus w_a but are, in a sense, more expensive because they are more likely to lead to a contract ($p_a > p_b$) and exhaust the endogenous “budget” of busy workers B the economy can sustain in a stationary environment. Our model places limited restrictions on (p_a, p_b, w_a, w_b) (see Lemma 1) and the form of the meeting function m , and, as a result, we cannot obtain strong welfare results. Proposition 4 below presents a welfare claim that holds under our general conditions, while Proposition 5 demonstrates that stronger results require additional assumptions on the model. We then show how additional structure placed on the model can strengthen the welfare results.

Proposition 4. *At the point of the stationary allocations frontier that defines the pooling equilibrium ($B(R^b) = R^b$), the derivatives of total contracts B and aggregate welfare $W(R^b, B(R^b))$ with respect to the number of employers in the busy market R^b are negative.*

Proposition 4 shows that as we *move* along the stationary frontier from the pooling equilibrium toward a costly signaling equilibrium, both the welfare and the total contracts have to increase. Intuitively, the pooling equilibrium is characterized by the “proportionality” condition $B = R^b$ and the stationarity condition $B = p_a m(1 - R^b, 1 - B) + p_b m(R^b, B)$, both of which ignore the difference in surplus w_a vs. w_b that different types of meetings produce. In the costly signaling equilibrium, employers take into account the difference in expected surpluses, and they “move” away from the pooling equilibrium in the “right” direction. However, the costly signaling equilibrium, once reached, is not necessarily efficient. In fact, we can produce examples where the costly signaling differs from the welfare-maximizing allocation.

Proposition 5. *The welfare in the costly signaling equilibrium is not necessarily higher than that in the pooling equilibrium.*

To make the model tractable, we impose two additional assumptions. The first assumption puts additional structure on the relationship between the contract formation probabilities (p_a, p_b) and the expected surplus from a meeting (w_a, w_b) .

Definition 4. A “no price dispersion” market is one in which $\frac{p_a}{p_b} = \frac{w_a}{w_b}$, or that expected pay-off to a worker conditional upon contract formation is the same regardless of the worker type.

Recall that $\mathbb{E}[\max\{v - C_s, 0\}] = \mathbb{E}[v - C_s | v - C_s \geq 0]p_s$. The type of the worker—or the level of her costs C_s —affects the social surplus through both the probability with which she forms a contract *and* the size of the pie conditional on the contract. In a “no price dispersion” market, the value $\mathbb{E}[v - C_s | v - C_s \geq 0]$ does not depend on the type of the worker. This is true in any class of models where the distribution of C_s conditional on being below v is the same for both types; one simple example is C_s taking two values (“high” and “low”) where only the “low” one leads to match. This assumption aligns the weights that the welfare function in Equation 7 puts on meetings in different markets with the “prices” those meetings have in the stationary allocation equation, leading to the following result.

Proposition 6. Under “no price dispersion,” maximizing welfare (Equation 7) is synonymous with maximizing the number of contracts B , or the equilibrium number of busy workers.

Proof. Social welfare is maximized by solving

$$\begin{aligned} \max_{R^b, B} w_a m(1 - R^b, 1 - B) + w_b m(R^b, B) \\ \text{s.t. } B = p_a m(1 - R^b, 1 - B) + p_b m(R^b, B). \end{aligned}$$

Multiplying the objective function by a constant p_b/w_b and using the fact that $\frac{p_a}{p_b} = \frac{w_a}{w_b}$, we can transform the objective function into the expression for B . \square

Proposition 6 makes the characterization of the socially-efficient outcome simple. Intuitively, this is because when realized surplus is the same for all matches, the social planner simply wants to maximize the number of people who are working. The following restriction on the meeting function m , in turn, simplifies the characterization of the costly signaling equilibrium.

Assumption 3. The elasticity of the meeting function with respect to the number of employers in a market is constant: $\frac{\partial m(x, y)}{\partial x} \frac{x}{m(x, y)} \equiv \eta$.

Assumption 3 effectively restricts the meeting function to be Cobb-Douglas and does not hold in the case of the more general CES matching function, as the elasticity depends on the particular values of x and y .⁹ In Proposition 7, we combine Definition 4 and Assumption 3 to show that the costly signaling equilibrium is constrained-efficient. It is an equilibrium that

⁹Empirically, there is work estimating matching function parameters with Cobb-Douglas and the more general CES matching function (see Bernstein, Richter and Throckmorton (2022) for an overview), along with the non-parametric approaches Lange and Papageorgiou (2020). Every example we are aware of is from conventional labor markets. The literature is unsettled on how consequential the constant elasticity assumption is in practice.

maximizes the total number of busy workers, B , which is also the number of contracts formed each period.

Proposition 7. *With no price dispersion and constant matching efficiency, the costly signaling equilibrium is constrained-efficient.*

Proof. Proposition 6 shows that the social planner would like to maximize the total number of busy workers B . To show this B -maximizing equilibrium is the same as the costly signaling equilibrium, consider the first order condition for the B -maximization problem:

$$\begin{aligned} \frac{d}{dR^b} [p_a m(1-R^b, 1-B) + p_b m(R^b, B)] &= 0 \\ -\eta p_a \frac{m(1-R^b, 1-B)}{1-R^b} + \eta p_b \frac{m(R^b, B)}{R^b} &= 0 \\ w_a \frac{m(1-R^b, 1-B)}{1-R^b} &= w_b \frac{m(R^b, B)}{R^b}, \end{aligned}$$

which is the employer indifference condition for the stationary costly signaling equilibrium from Definition 2. \square

Figure 3 illustrates the situation. The y-axis is the number of “busy” workers, B , and the x-axis is the share of employers recruiting in the “busy” sub-market, R^b . The heavy dark line depicts the set of stationary allocations $B = m(R^a, A)p_a + m(R^b, B)p_b$. These are combinations of employers and workers in the “busy” market such that these shares would remain unchanged.

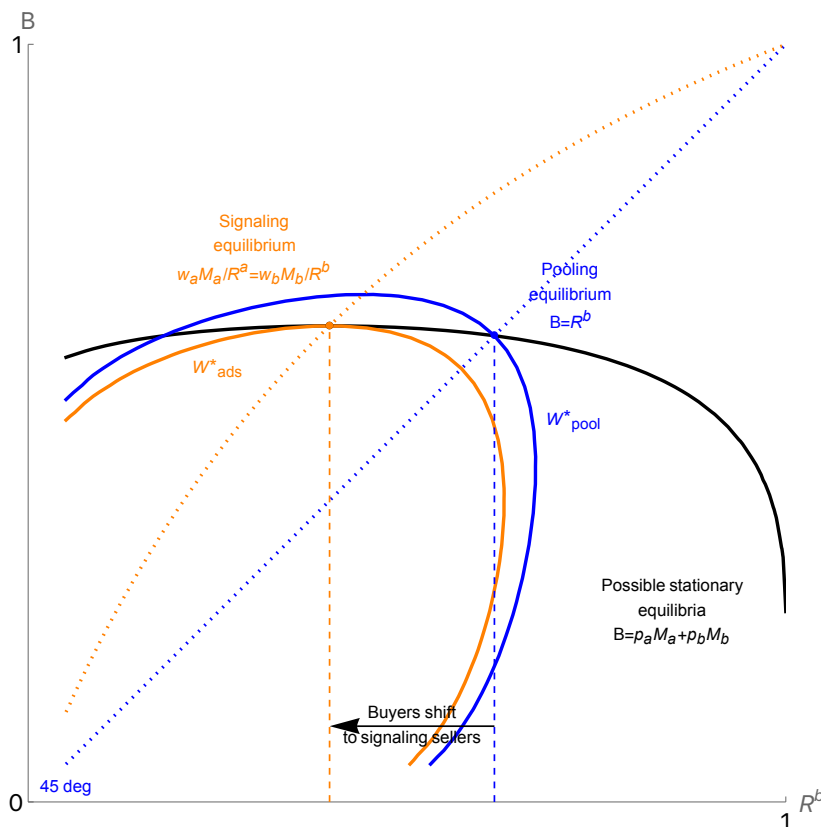
The 45-degree line is consistent with the pooling condition $R^b = B$ when meetings are random. The curved dashed line indicates the employer indifference condition $m(R^a, A)w_a/R^a = m(R^b, B)w_b/R^b$, which is met when pursuing busy and available workers offer the same expected pay-off for employers. Where these two curves intersect the curve of stationary equilibria, we have the costly signaling and pooling equilibria.

7.6 The model predictions versus the data

In this section, we connect some of our theoretical results to our empirical findings. First, Proposition 2 shows that workers who signal should enjoy higher “meeting” rates for the separating costly signaling equilibrium to exist. A reasonable proxy for a “meeting” in the model is an employer invite in the empirical context, as it may or may not lead to a contract formation. Consistent with our model, badge-renting workers enjoy more employer invites (see Table 3).

Second, in the model, signaling only increases the meeting rate, but not the probability of contract formation conditional on a meeting. This is also borne out in the data. As we’ve already seen, renting the badge increases the worker’s likelihood of receiving an invite (i.e., meeting) and the likelihood of forming a contract (see Table 8). However, once we condition

Figure 3: Illustration of the busy worker equilibria



Notes: The heavy dark line indicates stationary equilibria: $B = m(R^a, A)p_a + m(R^b, B)p_b$. The employer indifference condition between the busy and available worker markets is indicated by the curved line from the origin, $m(R^a, A)w_a/R^a = m(R^b, B)w_b/R^b$. Where it intersects the stationary equilibria frontier is the costly signaling equilibrium. The random matching condition is a 45-degree line from the origin, $R^b = B$, and where it intersects the stationary equilibria frontier is the pooling (no costly signaling) equilibrium. The iso-welfare curves for these two equilibria are indicated by W^*_{pool} and W^*_{ads} . Note that the costly signaling equilibria iso-welfare curve is tangent to the frontier. This is the social welfare-maximizing equilibrium.

on receiving an invite (column (3) in Table 8), there is no significant effect of badge renting on the likelihood of contract formation. This is consistent with the model, as the meeting rate is higher when signaling, but the probability of contract formation conditional on a meeting is unaffected by the signaling decision.

Third, our setting allows for a test of Proposition 4. As one starts “moving away” from the equilibrium without costly signaling to the one with it, the total number of employers searching in the “busy” market should decrease, and the number of matches should increase. Both findings are strongly supported by the empirical findings of Section 4.

Lastly, we showed that the costly signaling equilibrium is constrained-efficient under the assumptions of constant matching efficiency and no price dispersion. Although we cannot test for the first assumption, our data is consistent with the second. We find no evidence that

Table 8: Effects of costly capacity signaling on worker outcomes at the impression level

Dependent Variables:	Worker received invite		Contract formed	
	(1)	(2)	(3)	(3)
CONTROL	0.0802*** (0.0012)	0.0028*** (0.0001)	-0.0001 (0.0001)	-0.0001 (0.0001)
TRT	0.0031‡ (0.0018)	-0.0001 (0.0001)	-0.0002‡ (0.0001)	-0.0002‡ (0.0001)
BADGE	-0.0102*** (0.0011)	0.0005*** (0.0001)	0.0008*** (0.0001)	0.0008*** (0.0001)
TRT × BADGE	0.0039** (0.0013)	0.0002‡ (0.0001)	0.0001 (0.0001)	0.0001 (0.0001)
Worker received invite				0.0363*** (0.0006)
<i>Fit statistics</i>				
Observations	3,427,112	3,427,112	3,427,112	3,427,112

Clustered (Project post & Worker) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports OLS regressions where the independent variables are each employers’s treatment status, an indicator for whether a worker viewed in the search was renting the badge, and an interaction term between the two. Observations are on the project post employer impression level.

workers change their prices depending on whether or not they are engaging in costly signaling (see Table 6).

8 Conclusion

This paper studies the experimental introduction of costly capacity signaling in a large online labor market. On the worker-side, we find that badge-renting increases the number of invites from employers and the number of contracts formed. On the employer-side exposure to availability badges led to more invites to workers in total, more positive responses to their invites, and more contracts formed. Critically, the increase in transaction probability was not at the expense of workers who did not rent the badge—the overall effect was market expanding and there was a *net* increase in contracts of about 2%.

These results show costly capacity signaling can help to overcome a market failure by serving as a signal that helps to coordinate employers and workers. To serve this function, signaling had to be costly, as costless signaling had become uninformative. In our context, signaling was about capacity, but it is easy to imagine this could vary based on the context. The common economic problem is employer uncertainty about seller suitability—which could

be price, capacity, quality, or some other vertical attribute.

Other economic institutions have evolved to solve the capacity problem, although we are aware of no cases where the signal cost is determined centrally. In many of these scenarios, the cost of signaling a willingness to “trade” is more of a hassle cost or a technical cost, and the solution is imperfect. In our setting, the platform can pick a price for the signal that maximizes its informational content. We do not explore this optimal quantity problem, but it poses an interesting market design problem for future work.

References

- Agrawal, Ajay, John J Horton, Nicola Lacetera, and Elizabeth Lyons**, “Digitization and the contract labor market: A research agenda,” in “Economic analysis of the digital economy,” University of Chicago Press, 2015, pp. 219–250.
- Athey, Susan and Glenn Ellison**, “Position auctions with consumer search,” *The Quarterly Journal of Economics*, 08 2011, 126 (3), 1213–1270.
- Bagwell, Kyle and Garey Ramey**, “Coordination economies, advertising, and search behavior in retail markets,” *The American Economic Review*, 1994, pp. 498–517.
- Benson, Alan, Aaron Sojourner, and Akhmed Umyarov**, “Can reputation discipline the gig economy? Experimental evidence from an online labor market,” *Management Science*, 2019.
- Bernstein, Joshua, Alexander W. Richter, and Nathaniel A. Throckmorton**, “The matching function and nonlinear business cycles,” Technical Report 2201, FRB of Dallas Working Paper February 2022.
- Bhole, Monica, Andrey Fradkin, and John Horton**, “Information about vacancy competition redirects job search,” *SocArXiv*, April 2021, (p82fk).
- Decarolis, Francesco and Gabriele Rovigatti**, “From mad men to maths men: Concentration and buyer power in online advertising,” *American Economic Review*, October 2021, 111 (10), 3299–3327.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz**, “Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords,” *American economic review*, 2007, 97 (1), 242–259.
- Einav, Liran, Chiara Farronato, and Jonathan Levin**, “Peer-to-peer markets,” *Annual Review of Economics*, 2016, 8, 615–635.
- Filippas, Apostolos, John Horton, and Richard Zeckhauser**, “The surprisingly low cost of free goods: Stimulating consumption with in-kind transfers,” *Working paper*, 2023.
- , **John J Horton, and Joseph Golden**, “Reputation inflation,” *Marketing Science*, 2022.
- , – , and **Richard J Zeckhauser**, “Owning, using, and renting: Some simple economics of the “sharing economy”,” *Management Science*, 2020, 66 (9), 4152–4172.
- Fradkin, Andrey**, “Search, matching, and the role of digital marketplace design in enabling trade: Evidence from Airbnb,” 2023. Working Paper.

- **and David Holtz**, “Do incentives to review help the market? Evidence from a field experiment on Airbnb,” *Marketing Science*, September 2023, 42 (5), 853–865.
- Horton, John J**, “Online labor markets,” *Internet and Network Economics*, 2010, pp. 515–522.
- , “The effects of algorithmic labor market recommendations: Evidence from a field experiment,” *Journal of Labor Economics*, 2017, 35 (2), 345–385.
- , “Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution,” *Management Science*, 2019, 65 (8), 3518–3540.
- , **William R Kerr**, and **Christopher Stanton**, “Digital labor markets and global talent flows,” Technical Report, National Bureau of Economic Research 2017.
- Huang, Ni, Gordon Burtch, Yumei He, and Yili Hong**, “Managing congestion in a matching market via demand information disclosure,” *Information Systems Research*, December 2022, 33 (4), 1196–1220.
- Jung, Jaehwuen, Hyungsoo Lim, Dongwon Lee, and Chul Kim**, “The secret to finding a match: A field experiment on choice capacity design in an online dating platform,” *Info. Sys. Research*, December 2022, 33 (4), 1248–1263.
- Kihlstrom, Richard E. and Michael H. Riordan**, “Advertising as a signal,” *Journal of Political Economy*, June 1984, 92 (3), 427–450.
- Lange, Fabian and Theodore Papageorgiou**, “Beyond Cobb-Douglas: Flexibly estimating matching functions with unobserved matching efficiency,” Technical Report 26972, National Bureau of Economic Research April 2020.
- Lewis, Gregory**, “Asymmetric information, adverse selection and online disclosure: The case of eBay Motors,” *American Economic Review*, June 2011, 101 (4), 1535–1546.
- Milgrom, Paul and John Roberts**, “Price and advertising signals of product quality,” *Journal of Political Economy*, August 1986, 94 (4), 796–821.
- Mortensen, Dale T and Christopher A Pissarides**, “Job creation and job destruction in the theory of unemployment,” *The review of economic studies*, 1994, 61 (3), 397–415.
- Nelson, Phillip**, “Advertising as information,” *Journal of Political Economy*, July 1974, 82 (4), 729–754.
- Pallais, Amanda**, “Inefficient hiring in entry-level labor markets,” *American Economic Review*, 2013.


- Prada-Sarmiento, Juan David**, “A note on concavity, homogeneity and non-Increasing returns to scale,” *Working paper*, 2010.
- Sahni, Navdeep S and Harikesh S Nair**, “Does advertising serve as a signal? Evidence from a field experiment in mobile search,” *The Review of Economic Studies*, May 2020, 87 (3), 1529–1564.
- Spence, Michael**, “Job market signaling,” *The Quarterly Journal of Economics*, August 1973, 87 (3), 355.
- Stanton, Christopher T and Catherine Thomas**, “Landing the first job: The value of intermediaries in online hiring,” *The Review of Economic Studies*, 2016, 83 (2), 810–854.
- Tadelis, Steven and Florian Zettelmeyer**, “Information disclosure as a matching mechanism: Theory and evidence from a field experiment,” *American Economic Review*, February 2015, 105 (2), 886–905.
- Wright, Randall, Philipp Kircher, Benoît Julien, and Veronica Guerrieri**, “Directed search and competitive search equilibrium: A guided tour,” *Journal of Economic Literature*, 2021, 59 (1), 90–148.
- Yang, Joonhyuk, Navdeep S. Sahni, Harikesh S. Nair, and Xi Xiong**, “Advertising as information for ranking E-Commerce search listings,” *Marketing Science*, July 2023.

A Additional experimental details and results

A.1 Employer view of the badge renting information

Figure 4 shows an example of an employer’s view of workers during the experiment. It depicts the case where the worker on the top has chosen to rent the badge, and the worker in the bottom has chosen not to. Upon hovering over the badge, the employer could see the text: “This worker is promoting that they’re open to more work.” Badges were visible only while the employer was searching for workers to send invites. Crucially, the ability to view the badge was the sole difference between treated and control employers.

Figure 4: An example of an employer interface during the experiment



A.2 Internal validity


The experimental groups were well-balanced across several pre-experimental observables. To assess whether the randomized assignment was performed correctly, we test for systematic differences in observable pre-treatment outcomes between employers assigned to the control and the treatment groups. Table 9 reports two-sided t-tests for various employer project-specific outcomes. Figure 5 plots the number of employers allocated to the experimental cells over time.

Table 9: Balance tests for treated employers.

	Control mean \bar{X}_{CTL}	Treatment mean \bar{X}_T	p-value
<i>Project-specific outcomes</i>			
number of project posts	2.36	2.31	0.269
number of invites sent to workers	4.04	4.71	0.337
number of worker applications received	20.91	21.04	0.694
number of contract offers extended	0.67	0.67	0.858
<hr/>			
<i>Observation counts</i>	41,951	42,474	0.072

Notes: This table reports averages and p-values of two-sided t-tests for various pre-treatment outcomes, for employers assigned to the control and treatment group. The reported outcomes are (i) the number of projects posted, (ii) the number of invites sent to workers per post, (iii) the number of worker applications received per post, (iv) the number of contract offers extended per post.

Figure 5: Employers allocated to the control and treatment groups over time

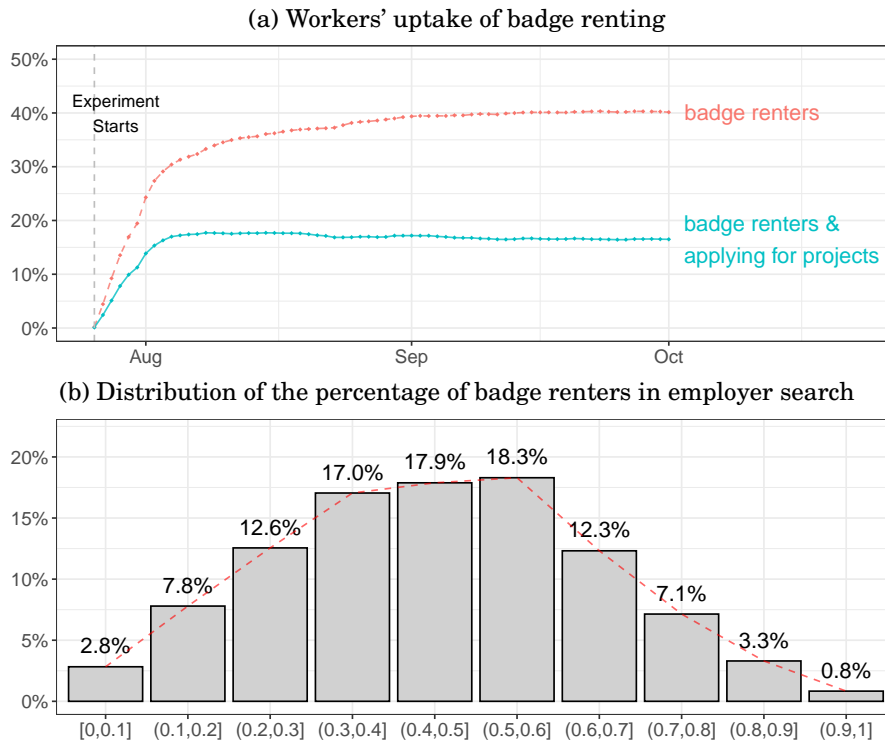


Notes: This figure plots the number of employers allocated to the control and treatment groups each day of the allocation period. The allocation period began on July 26, 2021 and ended on September 27, 2021.

A.3 Badge uptake and exposure

Figure 6 provides more details on the workers’ uptake of badge renting over time, and the employers’ exposure to badge renting over time. Figure 6a displays the number of workers renting the badge each day during the experiment. We also examine dynamic measures of worker capacity and compare them to badge uptake. Specifically, we look at whether workers who rented the badge were actively applying for projects organically. To do so, we define a worker as actively applying for projects at a given time if they had applied for at least one project in the previous seven days. We normalize this number by the total number of workers who applied to at least one project during the experimental period. Our findings indicate that many badge renters also actively applied to projects without being recruited via an invite.

Figure 6: Details on badge uptake and exposure



Notes: Panel (a) plots the workers’ badge renting uptake over time. It plots the number of workers renting the badge, and both renting the badge and actively applying for projects, for each day during the experimental period. We define a worker as “active” on a given day if she applied for projects within the previous seven-day window. Panel (b) depicts the distribution of the percentage of badge renters in employer search, using data for the last week of the experiment. Panel (b) uses data only for treated employers.

Badge renters were displayed prominently in employer search, with about 49.6% of all impressions and 52.6% of first-page impressions coming from badge renters. Nevertheless, each employer’s experience may have differed, even if workers renting the badge were commonplace. Figure 6b shows the distribution of the percentage of badge renters seen by employers

using data from the last week of the experiment. We can see that employers' exposure to badge renters was widespread, and that seeing no badge renters was a rare event for employers.

A.4 Treatment effect estimates of employer outcomes

Table 10: Treatment effect estimates of employer outcomes, OLS regression

(a) Outcomes (intensive margin)										
Dependent Variables:	Invites sent		Invites sent to badge renters		Proposals received following invite		Overall proposals received		Contracts formed	
Project Sample:	All	First	All	First	All	First	All	First	All	First
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Constant	3.052*** (0.037)	3.052*** (0.035)	1.370*** (0.015)	1.397*** (0.015)	1.350*** (0.013)	1.410*** (0.013)	14.795*** (0.093)	15.503*** (0.087)	0.292*** (0.003)	0.288*** (0.003)
TRT	0.236** (0.080)	0.088‡ (0.051)	0.133*** (0.028)	0.088*** (0.022)	0.068** (0.021)	0.040* (0.019)	0.201 (0.129)	0.053 (0.124)	0.008* (0.004)	0.008* (0.004)
<i>Fit statistics</i>										
Observations	126,550	84,425	126,550	84,425	126,550	84,425	126,550	84,425	126,550	84,425
<i>Clustered (Employer) standard-errors in parentheses</i>										
<i>Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1</i>										
(b) Outcomes (extensive margin)										
Dependent Variables:	Invites sent (any)		Invites sent to badge renters (any)		Proposals received following invite (any)		Overall proposals received (any)		Contracts formed (any)	
Project Sample:	All	First	All	First	All	First	All	First	All	First
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Constant	0.531*** (0.003)	0.559*** (0.002)	0.404*** (0.003)	0.425*** (0.002)	0.439*** (0.003)	0.459*** (0.002)	0.937*** (0.001)	0.945*** (0.001)	0.267*** (0.002)	0.265*** (0.002)
TRT	0.013*** (0.004)	0.009** (0.003)	0.019*** (0.004)	0.015*** (0.003)	0.013*** (0.004)	0.011** (0.003)	0.003 (0.002)	0.001 (0.002)	0.008** (0.003)	0.008** (0.003)
<i>Fit statistics</i>										
Observations	126,550	84,425	126,550	84,425	126,550	84,425	126,550	84,425	126,550	84,425
<i>Clustered (Employer) standard-errors in parentheses</i>										
<i>Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1</i>										

Notes: This table reports OLS regressions where the dependent variables are employer outcomes and the independent variable is a treatment indicator. Estimates are computed for two different samples: (i) “All projects” uses the entire sample of projects and clusters standard errors at the employer, (ii) “First project” only uses each employer’s first project post during the experiment. Panel (a) reports the raw outcomes, and Panel(b) reports the indicator variable transformation of each outcome.

Table 11: Treatment effect estimates of employer outcomes, Poisson regression

Dependent Variables: Project Sample:	Invites sent		Invites sent to badge renters		Proposals received following invite		Overall proposals received		Contracts formed	
	All	First	All	First	All	First	All	First	All	First
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Constant	1.116*** (0.012)	1.116*** (0.012)	0.315*** (0.011)	0.334*** (0.011)	0.300*** (0.010)	0.344*** (0.009)	2.694*** (0.006)	2.741*** (0.006)	-1.232*** (0.009)	-1.246*** (0.009)
TRT	0.075** (0.025)	0.028‡ (0.016)	0.093*** (0.019)	0.061*** (0.015)	0.049** (0.015)	0.028* (0.014)	0.014 (0.009)	0.003 (0.008)	0.026* (0.012)	0.027* (0.013)
<i>Fit statistics</i>										
Observations	126,550	84,425	126,550	84,425	126,550	84,425	126,550	84,425	126,550	84,425

Clustered (Employer) standard-errors in parentheses


*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports Poisson regressions where the dependent variables are employer outcomes and the independent variable is a treatment indicator. Estimates are computed for two different samples: (i) “All projects” uses the entire sample of projects and clusters standard errors at the employer, (ii) “First project” only uses each employer’s first project post during the experiment.

A.5 Model-free evidence on the shift towards badge renting workers

Table 2 showed evidence that treated employers sent more invites to badge renting workers but no clear evidence of a shift away from non badge renting workers. This is surprising, as we might expect that employers who send few invites may shift their attention to badge renters, at the expense of non badge renters. To build confidence in this result, we plot the distribution of employer invites by treatment status and worker badge renting status. The plot shows the empirical cumulative distribution functions for invites sent to badge renting workers (left panel) and non badge renting workers (right panel). We can see clear evidence that the treatment worked to increase invites to badge renting workers. However, there is no discernible shift for non badge renting workers.

Figure 7: Employer invite quantiles by worker badge renting and employer treatment status



Notes: This figure shows the empirical cumulative distribution functions of the invites sent by employers to badge renting and non badge renting workers. In each panel, we split the data by the employer’s treatment status.

A.6 Effects of costly capacity signaling by search position

In this section, we compare employer invites by search position, treatment status, and worker badge renting status. We use worker impressions presented to employers during the experimental period. An impression occurs when an employer sees a worker on her search interface.

Table 12 reports summary statistics on the impressions data, pooling all observations. There were about 3.4M impressions in total. A total of 75,622 unique employers saw at least one impression. A total of 136,174 unique workers received at least one impression.

The number of worker tiles an employer sees during a search session depends on how extensively the employer searches. Some searches go quite deep, but most are fairly shallow,

Table 12: Summary statistics for impressions ($n = 3,427,112$)

	Min	Mean	Median	SD	Max
Position in search (1 = top)	1	72.2	25	162	2.89e+03
Worker renting badge	0	0.492	0	0.5	1
Employer invite	0	0.0778	0	0.268	1
Worker accepts employer invite	0	0.0368	0	0.188	1
Contract formed	0	0.00304	0	0.055	1

Notes: This table reports summary statistics for impressions of workers presented to employers. The reported outcomes are (i) the impression position in the employer search, (ii) whether the worker was renting the badge when the impression took place, (iii) whether the employer sent an invite after the impression, (iv) whether the worker responded to the invite by applying for the project, and (v) whether a contract was formed.

as indicated by the maximum “Position in search” value in Table 12. On average, employers see badge renters in about half of their impressions. Of all the impressions made, employers make invites in about 8%, and about 4% of those invites receive a positive response from the worker, i.e., the response rate is about 50%. We can see that the conditional probability that the employer forms a contract with a worker who has accepted the invite is a bit less than 10%.


We are interested in how the worker’s position, badge renting status, and the employer’s treatment status affect the probability that the employer sends an invite. We begin by exploring this graphically in Figure 8. The x-axis is the worker’s position in search (1 = top of page), and the y-axis is the fraction of those workers that received an invite. As expected, employer invite rates are strongly declining in search position. Although these are organic listings, the consumer search pattern of starting at the top and working down is evident. However, recall that badge renters were not given additional prominence in our experiment.

Figure 8a slices our data by the workers’ badge renting status. In the left facet, we plot the invite rates of treatment and control employers when they encounter non badge renting workers. Treatment and control employers respond similarly for every position in the left facet. While it may seem mechanical, as the workers in this comparison are all non badge renters, the lack of difference is not necessarily assured: if the treatment prompted employers to switch from non badge renters to badge renters generally, we would expect to see a drop in invite rates for treated employers in this panel, since both types of workers appear mixed together in search results. The absence of this decline provides evidence of little crowd-out in the aggregate. In the right facet, we see the badge renting substantially increased the probability that employers sent invites to workers renting the badge. The signal conveyed by badge renting appears to be as valuable as a higher position in search.

Figure 8b slices our data by the employers’ treatment status. We see that treated employers have a significantly higher invite rate than control employers. Importantly, comparing the two

Figure 8: Worker search position and employer invites

(a) Receiving invites by worker badge renting status



(b) Sending invites by employer treatment status



Notes: This figure plots estimates of the probability of a worker impression resulting in an employer invite. The x-axis is the impression position in the employer's search page (1 = top of page), and the y-axis is the probability that an impression led to an employer invite. The left panel restricts the sample to non badge renting workers and the right panel to badge renting workers. Estimates for control employers who could not see the workers' badge renting status are depicted by red circles. Black triangles depict estimates for treated employers who could see the workers' badge renting status. We report 95% confidence interval for each point estimate.

facets suggest that the ability to view badges does not decrease employer invites to non badge renters, but rather it increases employer invites to workers who rent the badge.

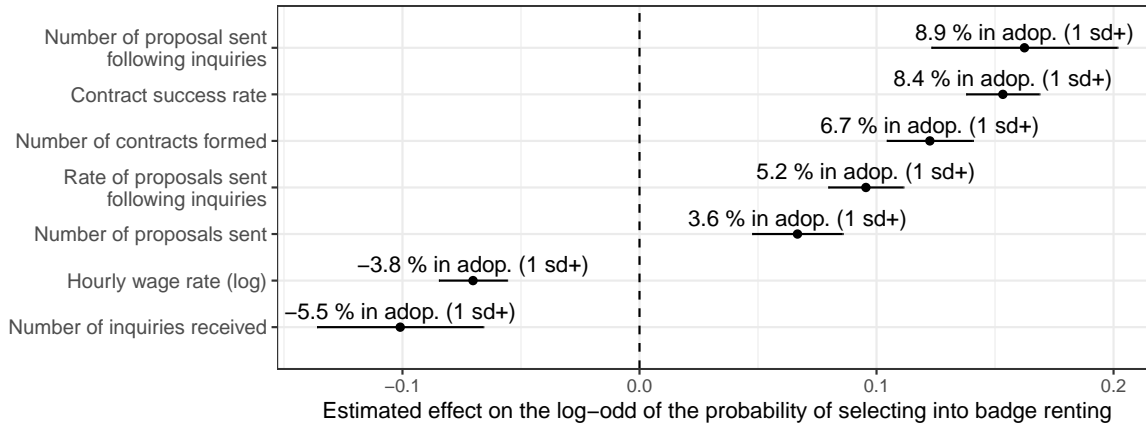
A.7 Statistically modeling worker selection into badge renting

We quantify the relative importance of various factors that predict whether a worker will choose to rent the badge. To do this, we report the results of a logistic regression where the outcome is an indicator variable for selection into badge renting, and the independent variables are pre-experiment worker attributes and outcomes. We standardize the independent variables (mean 0 and 1 standard deviation) and use them as predictors in the logistic regres-

sion. Standardization allows for comparing the relative importance of the predictors in the model.

Figure 9 reports those coefficients, ordered from largest to smallest. The coefficients are the effects on log odds; above each effect, we report the implied percentage change in badge renting probability from a one standard deviation increase in that measure. The baseline advertising adoption level is 45% of workers.

Figure 9: Relative importance of factors predicting worker selection into badge renting



Notes: This figure reports estimates of the effects of pre-experiment worker attributes on the probability that they select into badge renting. The estimates are obtained through a logistic regression where the outcome is a binary indicator for the worker selecting into badge renting, and the independent variables are standardized attributes. We report a 95% confidence interval around each estimate and the implied percentage change in the probability of selecting into badge renting from one standard deviation increase in the corresponding attribute.

The positive predictors of badge renting are being highly active: many accepted employer invites (“Number of bids placed following employer invites”), more contracts (“Number of contracts formed”), a higher acceptance rate of employer invites (“Rate of bids placed following employer invites”), and the total number of bids (including when not solicited by employers). We can also see that workers with higher reputation scores—as indicated by their successful contract completion rate (“Contract success rate”)—are more likely to select into badge renting. In contrast, those with higher wage asks are less likely to select into badge renting. Regarding magnitudes, the coefficient on the contract success score is 0.15, which implies that a worker with a 1 SD higher score has a 8.4% higher probability of renting the badge than the baseline, assuming log-odds are linear in the predictors.

The two factors that predict being less likely to rent the badge are (a) a higher hourly rate and (2) a greater number of employers invites already received. With all of the caveats needed for this cross-sectional analysis, a simple interpretation is that employers renting the badgewere interested in more work, as evinced by a relatively lower number of worker invites and a lower wage ask.

B Proofs of our theoretical results

B.1 Proof of Lemma 1

Proof. Consider $G(x; \beta) = p_b(x/v)^\beta$ when $x < v$. Let $F_b(x)$ coincide with $G(x; \beta)$ for all $x < v$ and set $F_b(v) = p_b$. Let $F_a(x)$ coincide with $G(x; 1)$ for all $x < v$ and set $F_a(v) = p_a$. The functions can be arbitrarily extended for $x > v$ and in a manner where $F_a(x) \geq F_b(x)$ for any $x > v$.

If $\beta = 1$, then $F_a(x)$ and $F_b(x)$ coincide for any $x < v$. It is easy to see that, in this case, $w_a = w_b$, implying $w_a/w_b = 1$. In the limit as $\beta \rightarrow +\infty$, the distribution $F_b(x)$ converges to a point mass at $x = v$ for $x \leq v$, which corresponds to $w_b = 0$. Since expected surplus is a continuous function of β , any value of $w_a/w_b \in [1, +\infty)$ can be achieved with the right choice of the parameter.

If continuity is imposed on both F_a and F_b , it becomes impossible for the ratio of w_a/w_b to be exactly equal to 1. However, one can construct a similar sequence of cdfs that get the ratio arbitrarily close to unity, showing that the lower bound of 1 is not a degeneracy that can be improved with additional assumptions on the distributions. \square

B.2 Proof of Proposition 1

The stationary pooling equilibrium could be obtained in closed form. Using conditions (1) and (2), condition (3) can be written as

$$B^* = m(1 - B^*, 1 - B^*)p_a + m(B^*, B^*)p_b$$

By constant returns to scale, the above becomes $B^* = m(1, 1)p_a(1 - B^*) + m(1, 1)p_bB^*$, which implies

$$B_{pool} = \frac{m(1, 1)p_a}{1 + m(1, 1)(p_a - p_b)}, \quad A_{pool} = \frac{1 - m(1, 1)p_b}{1 + m(1, 1)(p_a - p_b)} \quad (\text{A1})$$

B.3 Proof of Proposition 3

Our proof proceeds by demonstrating that there always exists a collection (R^a, R^b, A, B) that satisfies conditions (1)-(3) in the definition of a stationary costly signaling equilibrium (Definition 2). We then show that those conditions guarantee the existence of the price π that supports the separation of the types. Since $R^a = 1 - R^b$ and $A = 1 - B$, we can work with the pair (R^b, B) .

Step 1. In this step, we demonstrate that the stationarity condition (3) can be rewritten as a concave function $B(R^b)$ with $B(0) > 0$ and $B(1) < 1$.

To show this, fix $R^b \in [0, 1)$, and consider the function $f(B) = m(1 - R^b, 1 - B)p_a + m(R^b, B)p_b - B$. We have $f(0) = m(1 - R^b, 1)p_a > 0$ and $f(1) = m(R^b, 1)p_b - 1 < m(1, 1)p_b - 1 \leq p_b - 1 < 0$. Since

$f(B)$ is continuous, a solution $B(R^b)$ to $f(B) = 0$ exists for any value of $R^b \in [0, 1)$. Note, in particular, that for $R^b = 0$ we have $f(0) > 0$ and hence $B(0) > 0$.

This solution $B(R^b)$ is unique because $f(B)$ is a concave function. If $B_1 < B_2$ are two distinct solutions, then there exists $\lambda \in (0, 1)$ such that $B_1 = \lambda \times 0 + (1 - \lambda) \times B_2$. By concavity of f , we get

$$0 = f(B_1) = f(\lambda \times 0 + (1 - \lambda) \times B_2) \geq \lambda f(0) + (1 - \lambda)f(B_2) = \lambda f(0) > 0,$$

which is a contradiction.

The concavity of $f(B)$ follows from two observations. First, under the assumptions that we make, $m(x, y)$ is a concave function (Prada-Sarmiento, 2010). Second, it is simple to show that if $m(x, y)$ is concave, then so is $m(1 - x, 1 - y)$. Finally, a sum of concave functions is concave.

One remark is in order about the solution of the equation $f(B) = 0$ when $R^b = 1$. In that case, we are solving $m(1, B)p_b - B = 0$. Under our assumptions, $B = 0$ is a solution, and any solution is less than 1. In the uniqueness argument above, we assumed that $f(0) > 0$, which doesn't hold when $R^b = 1$. This development can generate one additional solution (more than one would still be ruled out by concavity unless $m(1, B)p_b - B$ is 0 everywhere). However, the stationary equilibrium with $(R^b = 1, B = 0)$ is of little economic interest because, in such an equilibrium, all the employers shop in the market for busy workers while all the workers are available. If another solution exists, then we use it instead, but if not—our results are not affected by this corner case.

Step 2. In this step, we demonstrate that the indifference condition (2) can be expressed as an increasing function $B(R^b)$ that satisfies $B(0) = 0$, $B(1) = 1$, and $(w_a > w_b) \implies (B(R^b) > R^b$ for $R^b \in (0, 1))$.

The collection of all points (R^b, B) which make the employers indifferent between the two markets is given by

$$w_b \frac{m(R^b, B)}{R^b} = w_a \frac{m(1 - R^b, 1 - B)}{1 - R^b} \quad (\text{A2})$$

Using constant returns to scale, this can be written as

$$w_b m\left(1, \frac{B}{R^b}\right) = w_a m\left(1, \frac{1 - B}{1 - R^b}\right) \quad (\text{A3})$$

For any $R^b \in (0, 1)$, there is a unique value of B that satisfies this equation. To see that, consider $f(B) = w_a m\left(1, \frac{1 - B}{1 - R^b}\right) - w_b m\left(1, \frac{B}{R^b}\right)$. We have $f(R^b) = w_a - w_b \geq 0$. If $w_a = w_b$, it is easy to see that $B = R^b$ is the only solution. In the interesting case of $w_a > w_b$, note that $f(1) = -w_b m(1, 1/R^b) < 0$. The solution exists by continuity of f and it is unique by the monotonicity of $m(1, x)$. Note that, unless $w_a = w_b$, the solution satisfies $B > R^b$ for $R^b \in (0, 1)$. It is trivial to

see that the solution $B(R^b)$ is an increasing function of R^b .

We now show that $\lim_{x \rightarrow 1} B(x) = 1$. First of all, note that $B(R^b)$ is bounded above by 1: for any $R^b \in (0, 1)$, $B = 1$ is too 'large' to satisfy equation (A3). Since $B(R^b)$ is an increasing function, then by the monotone convergence theorem and the continuity of $B(R^b)$, a limit at $R^b \rightarrow 1$ exists. Now suppose that $\lim_{x \rightarrow 1} B(x) < 1$. Then

$$\lim_{x \rightarrow 1} w_a m \left(1, \frac{1-B}{1-x} \right) \geq w_a m(1, 1),$$

since $m(1, \infty)$ is either infinity, in case $m(1, x)$ is bounded, or at least exceeds $m(1, 1)$. Then, by taking the limits of both sides in equation (A3), we get

$$w_b m \left(1, \lim_{x \rightarrow 1} B(x) \right) = w_a m(1, \infty) \geq w_a m(1, 1).$$

This is a contradiction, as, generally, $w_b < w_a$ and $m(1, \lim_{x \rightarrow 1} B(x)) < m(1, 1)$. Instead, if $\lim_{x \rightarrow 1} B(x) = 1$, then both the numerator and the denominator of $(1-B)/(1-x)$ are going to 0 as $R^b \rightarrow 1$. That allows the equation to be satisfied, provided that the convergence to 0 happens at a particular rate:

$$w_b m(1, 1) = w_a m \left(1, \lim_{x \rightarrow 1} \frac{1-B(x)}{1-x} \right) \tag{A4}$$

$\lim_{x \rightarrow 1} \frac{1-B(x)}{1-x} = L$ where L satisfies

$$m(1, L) = \frac{w_b}{w_a} m(1, 1) \tag{A5}$$

The limit of $B(x)$ as $x \rightarrow 0$ depends on the functional form of $m(x, y)$. If $m(x, y)$ is unbounded, then $\lim_{x \rightarrow 0} B(x) = 0$, as we will show momentarily. However, if the function is bounded above, then it is no longer the case. One example is $m(x, y) = \min(x, y)$, where $B(x)$ takes the form $B(x) = 1 - (w_b/w_a)(1-x)$. However, if $m(1, 1/x) \rightarrow \infty$ as $x \rightarrow 0$, then $B(x)$ has to approach 0 for equation (A3) to hold.

Step 3. There is a pair (R^b, B) that satisfies both the indifference condition (2) and the stationarity condition (3).

Let $B = f_1(R^b)$ be the function that describes the indifference condition and $B = f_2(R^b)$ be the function that describes the stationarity condition. We established that $f_1(0) = 0$, while $f_2(0) > 0$. Similarly, we showed that $f_1(1) = 1$ and $f_2(1) < 1$. Since both functions are continuous, there exists a value R^b where the two cross.

Step 4. If (R^b, B) is a point described in Step 3, then there exists a price of costly signaling π that fulfills the separation condition (6).

We established that all the points that satisfy the indifference condition are such that $B \geq R^b$. That implies $A = 1 - B \leq 1 - R^b = R^a$ and, hence, $R^a/A \geq R^b/B$. The separation condition (6) is satisfied when

$$\frac{m(R^a, A)}{A} = m(R^a/A, 1) \geq m(R^b/B, 1) = \frac{m(R^b, B)}{B},$$

which holds since $m(x, 1)$ is an increasing function.

B.4 Proof of Proposition 4

Let $B(R^b)$ be the frontier of all stationary allocations. Total welfare at any point R^b can then be written as

$$W(R^b, B(R^b)) = w_b m(R^b, B(R^b)) + w_a m(1 - R^b, 1 - B(R^b)) \quad (\text{A6})$$

We take the derivative of the welfare with respect to R^b and evaluate it at the point where $R^b = B$, which characterizes the pooling equilibrium. That derivative is negative, as we now show. The fact that the welfare-maximizing level R^b lies to the left of the point $R^b = B$ then follows from the concavity of $W(R^b)$, which we also demonstrate.

The derivative of interest has three components:

$$\frac{d}{dR^b} W(R^b, B(R^b)) = \left. \frac{\partial W(R^b, B)}{\partial R^b} \right|_{R^b=B} + \left. \frac{\partial W(R^b, B)}{\partial B} \right|_{R^b=B} \times \left. \frac{dB}{dR^b} \right|_{R^b=B} \quad (\text{A7})$$

Let us evaluate them all.

$$\frac{\partial W(R^b, B)}{\partial R^b} = w_b \frac{\partial m(R^b, B)}{\partial R^b} + w_a \frac{\partial m(1 - R^b, 1 - B)}{\partial R^b} \quad (\text{A8})$$

$$\left. \frac{\partial W(R^b, B)}{\partial R^b} \right|_{R^b=B} = \frac{\partial m}{\partial R^b}(1, 1)(w_b - w_a) < 0 \quad (\text{A9})$$

Here we used the fact that partial derivatives of a function that is homogenous of degree 1 are homogenous of degree 0, i.e., $\frac{\partial m}{\partial R^b}(1, 1) = \frac{\partial m}{\partial R^b}(x, x)$ for any x .

An identical exercise shows that

$$\left. \frac{\partial W(R^b, B)}{\partial B} \right|_{R^b=B} = \frac{\partial m}{\partial B}(1, 1)(w_b - w_a) < 0 \quad (\text{A10})$$

The last part is of special interest because this derivative tells us the change in total matches as we move away from the pooling equilibrium towards a costly signaling equilib-

rium where R^b is lower. Since $B(R^b)$ is defined implicitly as the solution of $B = m(1 - R^b, 1 - B)p_a + m(R^b, B)p_b$, we use the inverse function theorem to obtain $B'(R^b) = \frac{dB}{dR^b}$.

$$B'(R^b) = p_b m_1(R^b, B) + p_b m_2(R^b, B)B'(R^b) - p_a m_1(1 - R^b, 1 - B) - p_a m_2(1 - R^b, 1 - B)B'(R^b)$$

Plugging in $B = R^b$ and using the homogeneity of m (and its partial derivatives m_1 and m_2) we get

$$B'(R^b)|_{R^b=B} = -\frac{(p_a - p_b)\frac{\partial m}{\partial R^b}(1, 1)}{1 + (p_a - p_b)\frac{\partial m}{\partial B}(1, 1)} < 0 \quad (\text{A11})$$

The equation above proves that introducing costly signaling should locally increase the number of matches.

The signs of our three derivatives are not enough to determine the sign of the overall expression A7. We evaluate that expression to find

$$\frac{dW}{dR^b}\Big|_{R^b=B} = -\frac{\frac{\partial m}{\partial R^b}(1, 1)(w_a - w_b)}{1 + \frac{\partial m}{\partial B}(1, 1)(p_a - p_b)} < 0 \quad (\text{A12})$$

B.5 Proof of Proposition 5

It is proof by example.

Let $m(x, y) = \min\{x, y\}$, $p_a = 1/2$, $p_b = 2/5$, $w_b = 1$, $w_a = 5/4$. Then the pooling equilibrium has $R^b = B = 5/11$ and total welfare is $25/22$. The costly signaling equilibrium has $R^b = 1/4$, $B = 2/5$, and $W = 1$. Note that, while $m(x, y) = \min\{x, y\}$ doesn't fully satisfy the assumptions we placed on the meeting function, its CES approximation does. We find that $m(x, y) = ((1 - \alpha)x^\rho + \alpha y^\rho)^{1/\rho}$ with $\alpha = -1/2$ and $\rho = -10$ gets very close to the Leontief example and also produces higher welfare in the pooling equilibrium.

Interestingly, if we set $\rho = 0$ and get a Cobb-Douglas meeting function in the example above, the pooling equilibrium and welfare do not change. However, the costly signaling equilibrium and the welfare with that allocation do change to $R^b \approx 0.35$, $B \approx 0.457$, and $W \approx 1.1425 > 25/22 = 1.136$. Ultimately, there appears to be a connection between the elasticity of substitution between employers and workers in the matching function and the efficiency of the costly signaling equilibrium.

Figure 10 provides an example economy where welfare is identical in the pooling and the costly signaling equilibria. The point where welfare is maximized is where the frontier of stationary allocation is tangent to iso-welfare curves. As the diagram plot shows, the welfare-maximizing allocation lies somewhere between pooling and costly signaling equilibria. For this example, although moving from the pooling equilibrium to the costly signaling equilibrium is welfare-improving, the actual costly signaling equilibria offers the same welfare.

Figure 10: Welfare indifference curves and the conditions defining the costly signaling and the pooling equilibria.

