How Do Intermediaries Shape Labor Market Efficiency? An Equilibrium Model and Experimental Evidence *

Daniel Haanwinckel† Navdeep S. Sahni‡ Caio Waisman§

September 2025

PRELIMINARY AND INCOMPLETE

Abstract

This paper studies how intermediaries affect labor market efficiency through their monetization and match-steering strategies. We conduct a large-scale experiment in partnership with an online job platform and show that steering employer attention toward premium-tier workers—those who pay for preferential treatment in the matching mechanism—reduces overall hiring. To interpret these findings, we develop a frictional labor market model with imperfect competition among profit-maximizing platforms. Platforms choose both their fee structures and the extent to which they distort the matching process in favor of premium-tier workers. Worker differentiation into tiers arises endogenously, as workers experiment with platforms until finding a good match. We calibrate the model to match both observational and experimental moments and use it to simulate counterfactual policies. We find that banning premium membership tiers consistently improves welfare, despite their potential to reduce congestion. Our results highlight how monetization strategies that distort matching can reduce efficiency, even when intermediaries compete and provide valuable search infrastructure.

*We would like to thank participants at the Saieh Family Fellows Conference and at the Bank of Portugal research seminar for their comments on an earlier version of this paper.

†Department of Economics, UCLA. E-mail: haanwinckel@econ.ucla.edu

‡Graduate School of Business, Stanford University. E-mail: navdeep.sahni@stanford.edu.

§Kellogg School of Management, Northwestern University. E-mail: caio.waisman@kellogg.northwestern.edu.

1 Introduction

Workers and firms utilize intermediaries to facilitate matching in labor markets. Examples include job posting boards, university career fairs, union hiring halls, labor recruiters who manage seasonal migrant work in developing countries, and hiring agencies. This paper investigates how this type of intermediation affects labor market efficiency. More specifically, we study the consequences of three features that are present in many labor markets: (i) intermediaries compete with each other and may have some market power over workers and firms; (ii) they have some ability to steer matches toward specific individuals; and (iii) they have discretion over their monetization strategies, including the possibility of charging market participants for preferential treatment in the matching process.

These features may have significant but theoretically ambiguous effects on welfare. On one hand, participants willing to pay for preferential treatment may be positively selected: when matched with a counterparty, they are more likely to form a productive job relationship. In this case, steering matches toward those individuals can be welfare-improving. On the other hand, intermediaries may intentionally reduce match rates for non-paying participants to extract more rents from those with high willingness to pay, akin to a monopolist introducing an inferior version of a product to implement second-degree price discrimination. In that case, average match rates may fall, leading to lower welfare.

This paper's central contribution is to develop and implement a methodology to quantify these welfare effects. In our quantitative exercises, we study the market for temporary, remote work ("freelancers") intermediated by online job platforms. The main revenue source for most platforms is a proportional fee on workers' earnings. But many also feature "premium memberships" that promise higher match rates for workers or firms who pay for them. These memberships are typically implemented by restricting non-paying participants' access to resources with zero marginal cost for the platform, such as information, messaging, and job applications. This raises the possibility that premium memberships may have a negative effect on welfare arising from artificially generated scarcity in the search technology. On the other hand, platforms often host large contingents of workers with very low hiring probabilities (Kässi, Lehdonvirta and Stephany, 2021). In that case, giving those workers less weight in the matching mechanism may increase efficiency by reducing congestion externalities.

We partnered with a large platform to conduct an experiment quantifying how increasing the

advantage given to premium workers affects internal platform outcomes. The experiment separated employers searching for workers into treatment and control groups. Treated employers saw advertisements for workers whenever they ran search queries on the platform. By design, the advertised workers came from the premium membership tier. Using click-level data, we show that the experiment successfully diverted employers' attention from basic- to premium-tier workers. We then demonstrate that treated employers were significantly less likely to hire basic-tier workers than control employers. This effect is not offset by higher hiring rates of premium-tier workers; overall hiring rates decline.

To interpret these findings and infer market-level equilibrium outcomes, we develop a model that embeds intermediaries into the classic search and matching framework in the tradition of Diamond, Mortensen, and Pissarides (1984; 1994). In our model, workers and firms can only match if they participate in the same platform. Platforms possess market power for two reasons. First, workers and firms incur idiosyncratic costs when joining platforms, making their participation responsive to the expected value delivered. Second, a worker's suitability for a given platform is only revealed upon a successful match. If it takes too long to find a job, workers may exit and turn to another platform. This feature introduces endogenous differences in willingness to pay for platform services, as "experienced" workers face higher switching costs than those in the "experimental stage."

We use this framework to characterize when it is optimal for platforms to raise revenues through "match fees" (fixed percentages of labor earnings), premium membership fees, or both. We focus on premium memberships for workers to match the experimental design. We propose a microfoundation for the platform's matching mechanism, closely related to our experimental setting, based on the idea that the platform can direct employers' limited attention toward premium workers. We demonstrate that if the platform implements the premium tier, it is optimal to design it so that every experienced worker chooses that tier, while the basic tier is populated by workers in the experimental stage. Whether the platform wants to implement the premium tier depends on parameters that govern the elasticity of new users to the value promised by the platform, the share of experienced workers on the platform, and the elasticity of separations for experienced users.

Next, we augment the model with two extensions that incorporate realistic features of the market we study. The first is workers' imperfect information about job-finding rates. Specifically, while workers know the average job-finding rate in the market, they only learn the rate on a specific platform after using it for some time. The second is the existence of "phan-

toms"—inactive accounts that reduce match efficiency by increasing congestion (Cheron and Decreuse, 2016). While both extensions increase the likelihood that premium memberships are offered in equilibrium, they have opposite implications for the question of whether premium memberships are welfare-improving.

After calibrating the model using observational and experimental data, we use it for a series of counterfactual simulations. First, we study the choices of a welfare-maximizing planner who can control platform creation and design, but is still subject to search frictions. We find that the optimal number of platforms can be higher or lower than the competitive outcome, depending on a parameter that governs the strength of the "thin markets" effect within platforms. However, in all of these scenarios, the planner would opt to eliminate the premium membership tier entirely, financing the platforms solely through match fees.

Next, we consider a counterfactual policy that would prevent platforms from using premium membership fees. In all cases we analyze, this policy leads to welfare gains of around five percent (measured in terms of aggregate output net of all search costs, including platform operation costs). We use additional exercises to shed light on how the experiment informs these welfare results. We show that what matters is not the overall magnitude of the treatment effects on employment, but rather the ratio of the effects on non-premium workers to the overall employment effects. This ratio helps identify the fraction of basic-membership accounts that could form productive matches.

Note: Future versions of this paper will investigate the role of competition between platforms on outcomes, as well as the determinants of the Beveridge curve in this model.

Related Literature

We embed a model of competition between intermediaries in two-sided markets (Rochet and Tirole, 2003) into the canonical Diamond–Mortensen–Pissarides model of frictional labor labor markets (Pissarides, 2000). Within the industrial organization literature, we are closest to models of competition between intermediaries who make choices on pricing strategies (Caillaud and Jullien, 2003) or who can steer matches toward specific market participants (Hagiu and Jullien, 2014). Our work contributes to this literature in several ways. First, it incorporates simultaneous choices of monetization and steering. Second, we introduces endogenous differentiation between market participants based on their previous interactions with the platform. Third, we show that our model can replicate detailed features of the data

and serve as the basis for quantitative policy analysis.[1]

On the labor market side of the theoretical literature on intermediaries, our findings on platforms' market power and optimal entry echo the theoretical insights of Farboodi et al. (2025) and Farboodi, Jarosch and Shimer (2023), although those papers analyze different environments. Shimer and Wu (2023) also examine platform competition, but focus on private information and sorting rather than platform choices on monetization and steering.

A growing empirical literature uses online labour markets as field laboratories. Randomised experiments have examined algorithmic recommendations (Horton, 2017), information access (Barach and Horton, 2021), minimum-wage floors (Horton, 2025), sponsored advertising (Filippas et al., 2025*b*), and capacity signalling (Horton, Barach and Golden, 2020; Filippas et al., 2025*a*). We contribute to this literature in two ways. First, our experiment shows that advertising workers is effective at redirecting employer attention, yet the impact on overall hiring is negative. Second, we provide a theoretical framework that can be used to extrapolate the general-equilibrium implications of interventions studied in these experiments, as well as measure welfare effects which may not always align with employment.

Our paper is also related to work on how the internet and online platforms shape labour-market efficiency and distribution (e.g. Kuhn and Skuterud, 2004; Autor, 2008; Stevenson, 2008; Bhuller et al., 2023). Our focus on premium tiers and matching distortions connects most closely to studies of displacement effects, whereby limited employer attention in frictional markets can generate inefficiencies (Crépon et al., 2013; Pallais, 2014; Cheron and Decreuse, 2016). We explicitly model the "matching with phantoms" problem highlighted by Cheron and Decreuse (2016) and relate our findings to Pallais (2014) in the discussion of counterfactuals.

The paper is organized as follows. The following section describes the empirical setting and the experiment. The third section presents the model. The fourth section reports the quantitative analysis based on the calibrated model. The fifth section concludes and suggests directions for future research.

[1] Our focus on the equilibrium consequences of competition between intermediaries differentiate our work from a series of interesting papers that study related design choices in models with a single platform, such as Hagiu and Jullien (2011), Hagiu and Wright (2015), de Cornière (2016), Dinerstein et al. (2018), de Cornière and Taylor (2019), Marx and Schummer (2021), Teh, Wang and Watanabe (2024), Horton, Johari and Kircher (2024), and Lee and Musolff (2025). Atkeson, Eisfeldt and Weill (2015) study optimal entry of intermediaries in over-the-counter derivatives markets.

2 Online Job Platforms and Experimental Evidence

2.1 Description of the Market

Since their emergence at the turn of the century, online labor markets have experienced rapid growth over recent decades. By 2009, workers in this market earned around $700 million annually (Horton, 2010) and this figure reached $1 billion by the end of 2012 (Agrawal et al., 2015). More broadly, the so-called gig economy currently involves tens of millions of workers and generates hundreds of billions of dollars in gross volume of transactions (Statista Research Department, 2022).

In short, online labor markets are digital platforms through which individual employers can hire the services of independent workers for specific tasks. These platforms are not passive marketplaces as they operate the matching process between workers and employers and facilitate transactions between them. Typically, employers can post and describe vacancies and search for workers, whereas workers can search and apply for jobs. More information about agents on both sides of the market is also made publicly available.

Platforms compete against one another, and there are several options available to workers and employers. Whereas some platforms specialize in specific types of service, others host workers of varied backgrounds. Nevertheless, the monetization tools of all these platforms are roughly similar, as summarized in Table 1.

Most often, platforms charge a fee on successful matches between workers and employers. However, they also commonly offer premium memberships through which market participants can access additional and obtain more favorable services, such as enhanced visibility. Participation fees are very rarely charged.

2.2 The Platform: Searching, Matching, and Hiring

We partnered with a large online platform to run our experiment, which we describe in more detail in the next section. This platform does not specialize in specific types of tasks; it hosts workers from several areas, including marketing, IT, and translation.

In this platform, an employer creates listings with the job description and a range of prices they are willing to pay. Employers can also search for workers and invite specific ones to apply for their job, and the platform provides employers with algorithmic recommendations based on machine-learning models. Employers can observe workers' profiles, portfolios,

Table 1: Monetization Features of Some Online Labor Platforms

Platform	Fees on Matches	Premium for Workers	Premium for Employers	Participation Fee
99designs	Yes	No	Yes	No
Care.com	No	Yes	Yes	Yes
Catalant	Yes	No	Yes	No
DesignCrowd	Yes	No	Yes	No
Fiverr	Yes	Yes	Yes	No
Freelancer	Yes	Yes	Yes	No
Guru	Yes	Yes	Yes	No
Handy	Yes	No	No	No
Hubstaff Talent[†]	No	No	No	No
Outsourcely	No	Yes	Yes	Yes
PeoplePerHour	Yes	Yes	No	No
Rover	Yes	Yes	No	No
TaskRabbit	Yes	No	No	No
Thumbtack	Yes[††]	No	No	No
Toptal	Yes	No	No	No
Truelancer	Yes	Yes	No	No
Upwork	Yes	Yes	Yes	No
WeGoLook	Yes	No	No	No
Wonolo	Yes	No	No	No
Workana	Yes	Yes	Yes	No

Notes: Information from each platform was obtained by browsing their websites as of May 2025. [†] Hubstaff Talent does not charge fees because they function as an advertising tool for Hubstaff's time-tracking and workforce management paid products. [††] On Thumbtack, workers pay fees on "leads," that is, when employers try to contact them after seeing their profile on the website.

and reviews, and they have the ability to interview workers. In addition, the platform offers a premium option for employers that offers enhanced support in exchange for additional fees.

Workers create a profile and search for jobs in the platform. They can submit proposals for specific jobs, which may require the payment of an internal currency amount that has to be purchased. This currency can also be used to increase visibility and show availability. Following the conclusion of jobs and reviews, workers can be awarded specific badges that indicate their high quality. Furthermore, the platform also offers a premium tier for a monthly payment, which gives workers additional internal currency and other benefits.

Workers and employers need to come to terms for a match to be successful. The platform allows for both fixed-price contracts and hourly contracts, whose respective rates must be agreed upon by the two parties. The platform charges a fee on earnings in addition to potential additional payments from elective premium options.[2]

[2]The platform can effectively retain fees from matches because one of the services it provides is facilitating the payment for labor services. Firms and workers benefit from using the platform for settlement because they

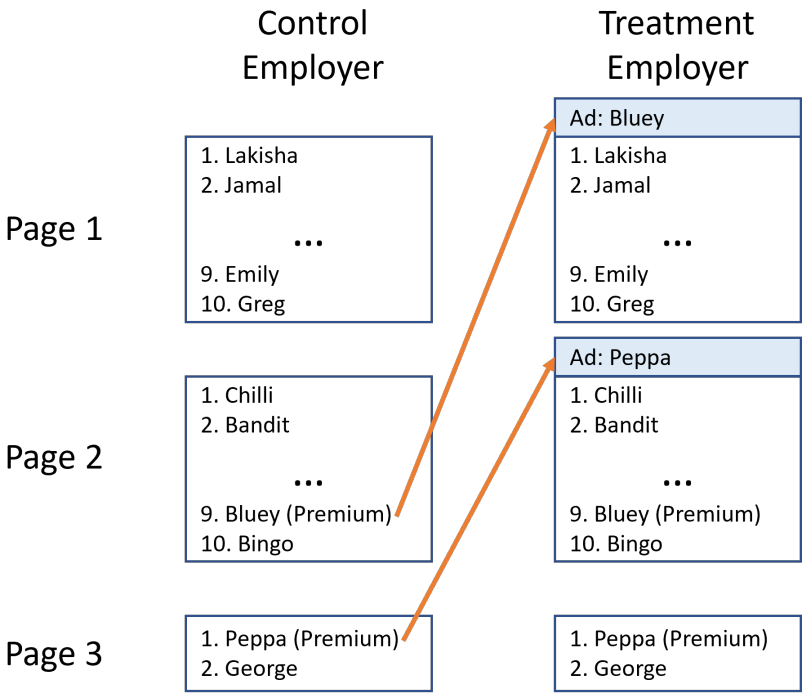


Figure 1: Premium workers featured as "ads" in search results

2.3 Experiment

We leverage data from an experiment that the platform ran between May and July of 2019. In this experiment, the platform introduced ads for the first time. More specifically, the platform boosted workers that belonged to the premium-tier to the top of the page, where they were displayed as an ad, based on the order in which they showed up as an organic result. Figure 1 illustrates how this boosting was implemented.[3]

The experiment randomized employers into treatment and control groups equiprobably.[4] Employers who belonged to the control group saw the platform as it always had been, without advertisements. In turn, treated employers were eligible to be exposed to advertised workers. The specific workers features as ads were those who would appear at some point in the regular search results, and by design, they are overwhelmingly sampled from the pool of workers who are currently enrolled in the premium membership tier.

Table 2 reports experimental results for outcomes measured at the employer level. The first

offer dispute resolution services. Additionally, the platform offers monitoring tools for workers engaged in online-only work.

[3]See DeVaro and Gürtler (2018) for a qualitative discussion of advertising in labor markets.

[4]Table A1 shows balance checks.

Table 2: Experimental results, employer-level regressions

	All workers	Sometimes premium	Never premium
*Panel A: dependent variable is **any clicks on these workers***			
Treatment	-0.00439***	0.0155***	-0.0123***
	(0.00147)	(0.00129)	(0.00134)
Constant	0.493***	0.255***	0.306***
	(0.00104)	(0.000906)	(0.000957)
*Panel B: dependent variable is **number of clicks on these workers***			
Treatment	-0.179**	0.0388***	-0.0795***
	(0.0769)	(0.00967)	(0.0181)
Constant	4.153***	0.703***	1.097***
	(0.0582)	(0.00717)	(0.0141)
*Panel C: dependent variable is **any hire***			
Treatment	-0.00280**	-0.000534	-0.00334***
	(0.00141)	(0.00109)	(0.00116)
Constant	0.366***	0.165***	0.196***
	(0.00100)	(0.000772)	(0.000825)
*Panel D: dependent variable is **number of hired workers***			
Treatment	-0.0155	-0.00299	-0.00782*
	(0.0112)	(0.00233)	(0.00445)
Constant	0.967***	0.252***	0.357***
	(0.00797)	(0.00169)	(0.00311)
N	464986	464986	464986

Data sources: Internal data from the platform. Note: the "Constant" rows in Panels B and D, Columns (2) and (3) do not add up to the value in Column (1) because there are workers in the platform for whom premium-membership status is not observed, such that they are included in the first column but do not appear in the others.

two panels shows the employer's interaction with search results in the platform, in the form of clicks. The dependent variable in Panel A is whether the employer clicked in the profile of a particular worker, while Panel B's dependent variable is the total number of clicks. While the first column shows clicking behavior for all workers, the second and third columns condition on the employer clicking on particular subsets of workers: those that sometimes have a premium membership during the experimental period, and those who never have a premium membership in the same period.

The results show that ads seem to successfully divert employer attention during the search process toward premium-membership workers. Compared to control employers, treated ones are more likely to click premium workers and less likely to click basic-membership ones. The overall effect is negative.

Next, we turn to hiring outcomes in Panels C and D. Treated employers are less likely to hire any workers during the experimental period—an effect of around 0.8%. The analysis by worker group shows that this decline is concentrated on workers who never have a premium membership. The point estimate suggests small, negative effect on the likelihood of "sometimes premium" workers being hired, but it is not possible to reject an *increase* of 0.9% at standard confidence levels.

To shed further light on the results, Table 3 reports results of alternative regression specifications. We employ a design where the unit of observation is an employer-worker pair. The workers paired to each employer are those who appear in the first page of search results ever seen by the employer in the experimental period.[5] The reason why restrict to that pool of workers is because it is not a selected sample. That is because the decision to look at more pages of search results is potentially endogenous to treatment itself. In addition, only workers for whom premium membership can be verified in the day of the search are included in the regression. The regressions include worker fixed effects, a dummy for whether the firm j is treated, a dummy for whether worker i has a premium membership on the day of the search t, and the interaction between those two variables.

The results in the first two columns support the interpretation that the experiment diverted

[5]We only include workers who appear as an *organic listing* in that page; that is, we exclude those who only appear as an ad. That is because including those appearing only as ads would create a potential imbalance in match quality, given that those appearing as organic listings of the first page are those that the platform's algorithm believes are the best matches for that particular employer and search query. The ones appearing as ads may be workers who would appear as organic listings in further pages, meaning that they could have lower match quality.

employers' attention but did not lead to a net increase in hiring of premium workers. The third column shows regressions where the dependent variable is an indicator for whether the worker appears as an add in that first page of search results, in addition to appearing as a regular listing. It confirms the experimental design: premium membership users are about 20 times more likely to be advertised than basic-tier workers.[6]

2.4 Discussion

In this section, we described the market for online temporary job platforms, focusing on how they obtain revenues from workers and firms. Most platforms offer premium memberships of some kind. To induce workers or firms to pay for those memberships, platforms must introduce a difference in match rates between premium and non-premium platform users. This wedge is often implemented by exploiting the fact that agents have limited attention, and that it can thus be directed towards premium users. This redirection of attention can be done by preventing basic users from contacting counterparties or by advertising premium users.

The experiment provides strong evidence that such approaches can be highly effective in introducing differences in hiring rates. Treated firms are significantly more likely to click on premium workers during the search process, and less likely to click on basic workers. As a result, fewer basic workers are hired, whereas effects on premium memberships are close to zero.

The fact that overall hiring is lower for treated firms may seem to suggest a reduction in match efficiency coming from premium memberships, but there are several reasons why that interpretation may be misleading. First, control users may be indirectly affected by the experiment because both groups face the same pool of workers. Thus, the difference in treatment and control groups may not reflect a thought experiment where a platform implements ads. Second, it is possible that a *marginal* increase in the distortion toward premium users is negative, but that a lower level is efficient. Third, because the experiment only comprises one platform and lasts for a limited period, it provides limited information on the equilibrium implications of premium memberships.

[6]The small number of basic workers showing up as ads could reflect a small degree of measurement error in the premium membership variable (which is defined at the daily level in our design, and thus could contain errors in the day of a change in membership status). It could also reflect very specific search queries where none of the organic listings is a worker with a premium membership.

Table 3: Experimental results, pairwise regressions (first page of search results)

	Hired	Clicked	Advertised
Treatment_j	-0.000448***	-0.00315***	0.00257***
	(0.000149)	(0.000575)	(0.000302)
$\text{Treatment}_j \times \text{Premium}_{i,t}$	0.000502*	0.00238**	0.0479***
	(0.000267)	(0.000948)	(0.00274)
$\text{Premium}_{i,t}$	-0.000191	-0.000639	
	(0.000554)	(0.00185)	
Constant	0.00175***	0.0203***	
	(0.000232)	(0.000784)	
N		415566	
r2	0.143	0.120	0.0367
Unique employers		223187	
Unique workers		22040	
Worker fixed effects	✓	✓	

Data sources: Internal data from the platform.

These limitations constitute the key motivation for the model developed in the next section. As we will show, the model allows us to consider the broader impacts of premium memberships in the market. Still, the data provided by the platform is critical in assessing the quantitative implications of the model's mechanisms, as we show towards the end of our analysis.

3 Model

In this section, we develop the continuous-time frictional model of job search with profit-maximizing intermediaries, which we refer to as platforms. Although the model is dynamic, we will focus on steady-state equilibria. In addition, throughout the paper, we restrict attention to symmetric equilibria where platforms make identical choices.

Our focus is to study the cases where platforms manipulate the matching mechanism to favor some workers relative to others, aligning with the experiment described in the previous subsection. While the same kind of bias can exist towards specific firms, we abstract from that dimension to keep the model as simple as possible.

3.1 Workers, Firms, Platforms, and Frictions

The economy is populated by N_w workers and a large number of potential employers, who may pay a fixed flow cost c to sustain an open vacancy. We use "firms" or "employers" interchangeably. Workers and firms are ex-ante identical, infinitely-lived, and risk-neutral. A worker-firm pair is only capable of producing output if their idiosyncratic productive mismatch d, a $Uniform[0,1]$ random variable, is at most $1/D$. Mismatch is not directly observable; firms and workers can only verify if they can form a productive match through a process that takes time. Thus, if the exogenous parameter D is large, very few matches will be formed if firms and workers only meet at random.

The role of platforms in the model is providing a more efficient means for workers and firms to match with each other. Specifically, platforms possess a technology that allows them to find, among the workers who participate in the platform, the best match for a given vacancy. Details on this matching technology will be discussed below.

Risk-neutral entrepreneurs must pay a flow cost F to operate a productive platform. Throughout the paper, we will index platforms by i, denote by I the endogenous number of operating platforms in the economy, and assume that workers and firms can only belong to at most one platform at a time.[7]

Platforms are horizontally differentiated in two ways. First, firms and workers have idiosyncratic one-time costs of joining a platform for the first time, and workers in the platform are also subject to idiosyncratic, negative shocks to utility that can be avoided by leaving. Those two types of idiosyncratic shocks provide a tractable means to introduce imperfect competition between platforms. That's because they make inflows and outflows of participants responsive to the expected value delivered by the platform, but with finite elasticities.

The second form of differentiation is that a platform may not be effective for all workers. Specifically, when a worker joins a platform, an unobservable random variable x is drawn from a Bernoulli distribution with mean $\bar{x}$. If $x = 0$, the worker is "unsuitable" to the platform, meaning that the worker cannot form productive relationships with firms that participate in that platform.[8] Workers learn that $x = 1$ only if they are successfully matched to a firm

[7] In other words, our model does not allow for *multi-homing*. In the context of online labor platforms this is not an insensible assumption; see, for example, Horton (2025). In Section 3.6.1, we incorporate an extension to the model that can at least partly account for the possibility of multi-homing in the data. See that section for further details.

[8] This feature can represent worker-platform idiosyncratic issues that are hard to observe before using the platform, such as whether the worker finds the platform's interface intuitive and can effectively operate it.

while using the platform. This model component creates endogenous ex-post differentiation between workers in the platform: some are "inexperienced" workers experimenting with the platform to discover whether they are suitable or not, while others already know that $x = 1$ and, hence, have a higher willingness to pay for staying in the platform.

We assume that platforms can only raise revenues in two ways. The first is through a match fee f_m that is extracted from successful job matches. The second is a "premium membership" fee f_p paid by workers who choose that tier, if the platform chooses to offer it. We do not allow platforms to charge participation fees. This assumption is central to our results. It is consistent with the fact that very few platforms charge those kinds of fees, as shown in Table 1. In Appendix A.1, we discuss how zero participation fees could emerge endogenously in augmented versions of the model with additional mechanisms.

3.2 Timing, State Transitions, Value Functions, and Wage-Setting

Employers can be in one of three states. The first is a newly created vacancy. In this stage, the firm has to search for a *platform* and join it. The second stage is an open vacancy within the platform. At this point, the firm searches for workers, but may also receive an exogenous vacancy destruction shock. The third and final stage is reached after matching with a worker, at which point the firm has a filled job that is subject to a job destruction shock. These states are illustrated as the orange and green boxes in Figure 2.

In the first state, firms receive the opportunity of being matched to a randomly-drawn platform i at Poisson rate λ_0. When this happens, they draw the cost of joining the platform ε from an Exponential distribution with mean $1/\rho_\varepsilon$. The corresponding flow value is given by

$$rV_0^e = -c + \lambda_0 \int_0^\infty \left[\frac{1}{I} \sum_i \max \left\{ V_i^e - \varepsilon - V_0^e, 0 \right\} \right] \rho_\varepsilon e^{-\rho_\varepsilon \varepsilon} d\varepsilon,$$

where r is the discount rate.

The value of being an employer in the second state—an open vacancy at platform i—is given by

$$rV_i^e = -c + q_i \left[J_i^e - V_i^e \right] + s_v \left[V_0^e - V_i^e \right],$$

where q_i is the platform's vacancy-filling rate and s_v is an exogenous rate of separations from

Suitability could also reflect unmodeled sorting of employers with different skill requirements, or of different geographies, to different platforms, so long as that kind of sorting is not easily observable ex ante.

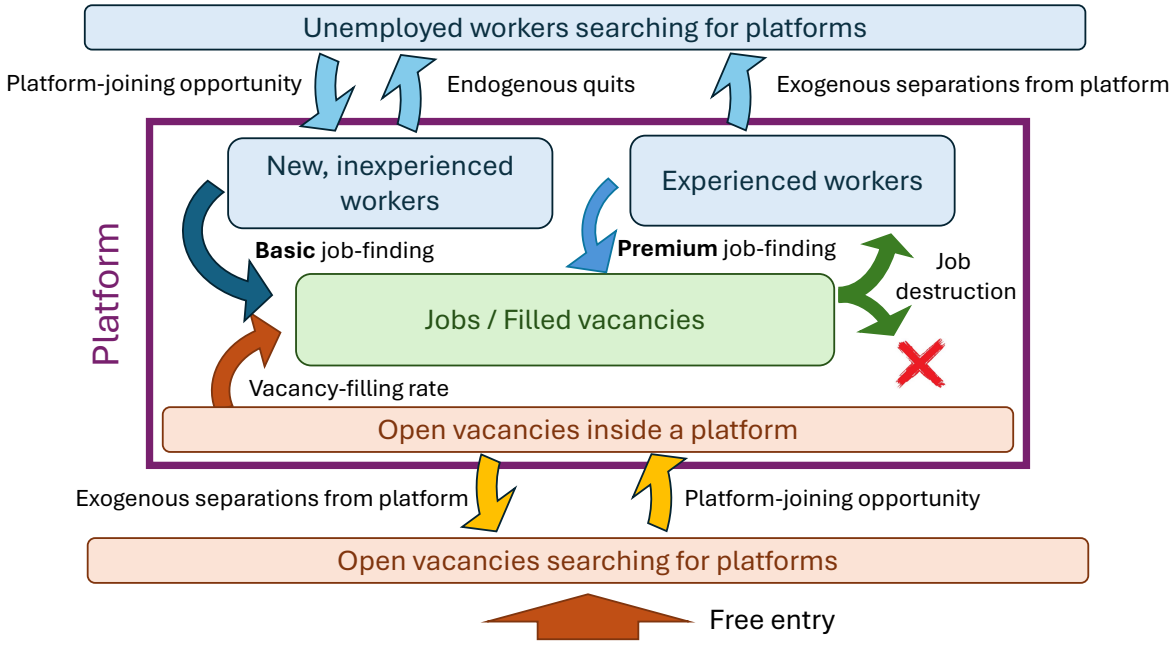


Figure 2: The states in which workers and firms may be at any moment in the model. The figure also shows how the ways through which workers and firms transition between states.

the platform while in the open vacancy stage.

Finally, the value of a filled vacancy at platform i, J_i^e, is defined by

$$rJ_i^e = y - w_i + \delta\left[0 - J_i^e\right]$$

where y is the raw match output, w_i is the bargained wage (to be defined below), and δ is the job destruction rate. When the job is destroyed, the firm goes back to the pool of potential entrants, with present value normalized to zero.

Workers may be in one of *five* possible states. Like firms, the beginning of the match process is the search for a platform. Upon joining one, they remain in the second state—an *inexperienced* worker—until they are matched to an employer or, alternatively, if they decide to leave the platform. Employed workers (the third state) leave that state following job-destruction shocks. At that point, they remain in the platform and can choose between one of the two last states: an experienced unemployed worker with either a basic or a premium membership.

We start with the equation that defines the value of the platform-search state,

$$rV_0^w = \lambda_0 \int_0^\infty \left[\frac{1}{I} \sum_i \max \left\{ V_i^{iw}(0) - V_0^w - \varepsilon, 0 \right\} \right] \rho_\varepsilon e^{-\rho_\varepsilon \varepsilon} d\varepsilon,$$

where $V_i^{iw}(t)$ corresponds to the *perceived* value of being an "inexperienced worker" who has been in the platform for t units of time without having received a match. That value function is defined by

$$rV_i^{iw}(t) = \Pr\left(x = 1 | t, \bar{h}_b\right) \bar{h}_b \left[J_i^w - V_i^{iw}(t) \right] + \frac{dV_i^{iw}(t)}{dt}.$$

where $\tilde{J}_i^w$ is the perceived value of being an employed worker in platform i (to be defined below).

This expression defines a *perceived* value, rather than an actual value, for two reasons. First, the worker does not observe their idiosyncratic suitability within that platform, $x \in \{0,1\}$. Thus, the value function reflects their posterior probability regarding their suitability given how long they have stayed in the platform. For workers just joining the platform, that perceived probability is $\Pr\left(x = 1 | t = 0, h_{b,i}\right) = \bar{x}$. As time passes, that perceived probability falls, approaching zero as $t \to \infty$. When $V_i^{iw}(t)$ falls so much that it becomes equal to the outside value V_0^w, it becomes optimal to leave the platform and return to the platform-search stage. See Appendix A.3 for more details.

Second, we assume that inexperienced workers cannot observe the platform-specific job-finding rate for basic-membership workers, $h_{i,b}$. Instead, workers infer that the platform's job-finding rate is equal to the average rate across all platforms, denoted by $\bar{h}_b$. In a symmetric equilibrium, worker's beliefs regarding job-finding rates are correct. However, this assumption has implications for the platform's optimal choices, as we discuss below. Appendix A.4 contains an extended discussion regarding this assumption. There, we show that it is more consistent with our data than the alternative assumption that inexperienced workers are fully informed about job-finding rates. We also discuss the robustness of our main quantitative result as we allow a fraction of inexperienced workers to be well-informed about match rates.

Note that, in defining $V_i^{iw}(t)$, we assumed that inexperienced workers choose the basic membership in the platform, even if the premium membership tier is offered. We will show later that this assumption is correct in equilibrium.

The value function for employed workers is given by

$$rJ_i^w = (1 - f_{m,i})w + \delta\left[\max\left\{V_i^b, V_i^p\right\} - J_i^w\right]$$

where $f_{m,i}$ is the platform-specific percentage fee collected on payments from the firm to the worker (referred to as "match fees"). V_i^b and V_i^p denote the values of being an experienced, but unemployed worker in the platform with either a basic or a premium membership, respectively.

Before presenting the equations defining these last two value functions, we must introduce the process of separations from the platform. At an exogenous rate s_U, the worker receives a "separation shock" ν drawn from an Exponential distribution with mean $1/\rho_\nu$. The worker has then the choice of incurring an instantaneous utility loss of ν or return to the platform-search stage. The resulting asset values then satisfy:

$$rV_i^b = h_{b,i}\left[J_i^w - V_i^b\right] + s_U \int_0^\infty \max\left\{-\nu, V_0^w - V_i^b\right\} \rho_\nu e^{-\rho_\nu \nu} d\nu$$
$$rV_i^p = -f_{p,i} + h_{p,i}\left[J_i^w - V_i^p\right] + s_U \int_0^\infty \max\left\{-\nu, V_0^w - V_i^p\right\} \rho_\nu e^{-\rho_\nu \nu} d\nu,$$

where $f_{p,i}$ is the premium membership fee and $h_{p,i}$ is the premium job-finding rate.

We end this section by describing wage setting in the platform. As in the standard DMP model, wages are set by Nash bargaining. The threat points for both worker and firm correspond to returning to the search stage within the platform. Denoting the worker's bargaining power as β, that means the bargained wage at platform i is the implicit solution to the equation

$$(1-\beta)\left[J_i^w(w) - \max\left\{V_i^b, V_i^p\right\}\right] = \beta\left[J_i^e(w) - V_i^e\right]. \tag{1}$$

Note that wages do not depend on whether it's the worker's first match in the platform or not. That is because upon the first match, the worker learns that $x = 1$, and thus the threat point is to return to the platform as an experienced worker.[9]

[9]When we defined the value for inexperienced workers $V_i^{iw}(t)$ above, we used a *perceived* value for the value of a job $\tilde{J}_i^w$, rather than the real J_i^w. That is because, as discussed above, inexperienced workers do not perfectly observe job-finding rates—which also means they cannot perfectly predict wages that result from Nash bargaining. The expression for $\tilde{J}_i^w$ is obtained after substituting $\bar{h}_b$ for $h_{b,i}$, and $\bar{h}_p$ for $h_{p,i}$, in the expressions for unemployed workers in the platform, and then using the results instead of the real values in the expression for J_i^w and in the Nash bargaining expression. $\tilde{J}_i^w$ is still platform-specific because inexperienced workers can perfectly observe fees $f_{m,i}$ and $f_{p,i}$.

3.3 Platform's Matching Technology

In this Section, we present the equations that define the platform's matching technology. We use lowercase n to denote numbers of workers and firms in each platform. Specifically, $n_{V,i}$ is the number of firms with open vacancies, $n_{J,i}$ is the number of filled job positions, $n_{b,i}$ counts unemployed workers with a basic membership, and $n_{p,i}$ counts unemployed workers with a premium membership. The total number of firms in the platform is thus $n_{V,i} + n_{J,i}$, while the total number of workers is $n_{b,i} + n_{p,i} + n_{J,i}$.

The basic logic of the matching technology is as follows. At an exogenous rate λ, the platform is able to identify the best match for a randomly-drawn vacancy in the platform. That is, it identifies the worker with the lowest mismatch d, though it does not directly observe d itself. The platform then presents that worker to the firm. If the worker is suitable to the platform (that is, $x = 1$) and the mismatch is low enough ($d < 1/D$), then a productive match will be formed.

We also allow the platform to restrict attention to premium-membership workers when it has the chance to identify the best match for a given vacancy. If it does so, then it identifies the lowest-d worker within the premium tier, ignoring workers outside of that tier. The frequency with which the platform engages in this kind of distortion is a choice, denoted by $\gamma \in [0,1]$.[10]

Given this structure, the vacancy-filling rate can be approximated as

$$q_i = \frac{\lambda}{n_v}\left\{\gamma_i\left[1 - e^{-n_{p,i}/D}\right] + (1-\gamma_i)\left[1 - e^{-(n_{b,j}+n_{p,i})/D}\right]\phi_i\right\}. \tag{2}$$

$$\text{where } \phi_i = \frac{n_{p,i} + n_{b,i}\bar{x}}{n_{p,i} + n_{b,i}}.$$

This formula is derived in Appendix A.2, which provides additional discussion on the platform's matching technology and how it relates to our experimental setting.

Four points about this expression are noteworthy. First, it is increasing in the number of workers in the platform. With few workers, match probabilities become small due to a "thin

[10] We view the "best-match" arrivals as a combination of technical design choices by the platform—such as the quality of the matching algorithm and website design—and the effort and attention of users. There are different mechanisms that the platform can use to implement $\gamma > 0$. One is to deny access to platform features that have zero marginal cost, such as arbitrary constraints on messaging features or access to information. Another is advertising, whereby the platform deviates employer attention to "featured" premium workers that may not necessarily be the most suitable workers to that employer given the platform's matching algorithm. The latter interpretation corresponds to the experimental part of the paper.

markets" effect. That is, even when the platform is able to identify the most suitable worker, that worker may still not be a good enough match.

Second, it may be either increasing or decreasing in γ. On one hand, restricting attention to premium workers means that there would be situations where a suitable basic-membership worker would be a good enough match, but none of the premium-membership ones are good enough. On the other hand, that kind of restriction may be efficient if many platform recommendations coming from the unrestricted pool do not leading to jobs because few of those workers are suitable (that is, if ϕ_i is low). Thus, our formulation of the matching process does not imply that the existence of premium tiers is necessarily detrimental to market efficiency. In Section 3.6 below, we extend the model in a way that provides further mechanisms through which premium memberships could increase welfare by allowing workers to credibly reveal private information.

The third remark is that, conditional on $\lambda > 0$, the vacancy-filling rate is strictly increasing in the share of the platform's workers who belong to the premium tier. That is because having more workers in that tier makes the conditioning to that pool less restrictive. This observation has an important implication for the profit-maximizing problem of the platform. As we will show below, conditional on implementing a premium membership, the platform will want to induce all experienced workers to participate in it, in order to make it less distortionary.

The fourth and final remark is that, while both the λ and D affect overall match rates, they have independent effects on the way the matching technology works. Specifically, when D is high relative to the number of workers in the economy, platforms suffer severely from a "thin markets" problem. When we solve the problem of a welfare-maximizing central planner in this economy in Section 3.7, we will show that a higher D is associated with larger optimal platform sizes.

We end this discussion by noting the corresponding job-filling rates for suitable workers in the platform, depending on whether they belong to the basic or premium tiers:

$$h_{b,i} = \frac{\lambda(1-\gamma_i)\left(1-e^{-(n_{b,i}+n_{p,i})/D}\right)}{n_{b,i}+n_{p,i}} \tag{3}$$

$$h_{p,i} = h_{b,j} + \frac{\lambda\gamma_i\left(1-e^{-n_{p,i}/D}\right)}{n_{p,i}}. \tag{4}$$

3.4 Firm and Worker Choices, Flows, and Steady-State Quantities

We now start to characterize a market equilibrium in this model. We focus on symmetric equilibria in which all platforms make identical choices. When referring to platform-related variables, we will use $-i$ for the single value chosen by all other platforms where needed, and we will omit subscripts where the distinction is not necessary. In this section, we describe optimal behavior of workers and firms. The following section uses these results to formulate and solve the profit-maximizing problem of the platforms.

Optimal vacancy creation by firms requires a zero-profit condition,

$$V_0^e = 0. \tag{5}$$

Firms and workers in the platform-search stage who receive the opportunity to join a platform will do so as long as the joining cost, ε, is low enough. That condition is $\varepsilon < V_i^e - V_0^e = V_i^e$ for firms and $\varepsilon < V_i^{iw}(0) - V_0^w$ for workers. Denoting by N_0^e and N_0^w the number of workers and firms in the platform-search stage, respectively, we can then calculate the inflow of workers into any particular platform as follows:

$$\begin{aligned} \text{Inflow of firms into platform } i: &\quad \frac{N_0^e}{I}\lambda_0\Big[1-\exp\big(-\rho_\varepsilon V_i^e\big)\Big] \\ \text{Inflow of workers into platform } i: &\quad \frac{N_0^w}{I}\lambda_0\Big[1-\exp\big(-\rho_\varepsilon (V_i^{iw}(0)-V_0^w)\big)\Big]. \end{aligned}$$

It is easy to show that the elasticities of inflows into a platform with respect to the platform value are strictly positive, strictly decreasing, diverge to positive infinity as $V_i^e \to V_0^e$ or $V_i^{iw}(0) \to V_0^w$, and converge to zero as V_i^e or $V_i^{iw}(0)$ grow large. The parameters λ_0 and ρ_ε regulate the degree of competitiveness between platforms.

Firms with open vacancies have no action; they passively wait for the arrival of a match or a platform separation shock.

Inexperienced workers stay in the platform for a certain amount of time, but leave if it takes too long to be matched to a firm. In Appendix A.3, we show that the optimal stopping time is

$$T_i \;=\; \max\left\{0,\; \frac{1}{\bar{h}_b}\left[\ln\Big(\frac{\bar{x}}{1-\bar{x}}\Big)+\ln\left(1+\frac{\bar{h}_b}{r}\frac{J_i^w-V_0^w}{V_0^w}\right)\right]\right\}. \tag{6}$$

The argument for why these workers choose the basic membership tier will be made in the

following subsection, when we discuss optimal platform choices.

Matched workers and firms have no choices. They remain matched until receiving a job destruction shock. Experienced unemployed workers have the option of choosing between premium or basic memberships. They pick the one that maximizes utility.

We now proceed to characterize the size of the equilibrium pools of workers and firms in each possible state. To do so, we need to make assumptions about the membership options of different workers. Specifically, we assume that inexperienced workers choose basic memberships while experienced workers choose premium memberships. We will show later that these choices are consistent with an equilibrium. If the platform does not offer a premium membership, one should interpret the premium membership and the basic membership as being identical, such that the assumption above can be maintained even in this case.

The equilibrium number of firms in the platform-search stage is pinned down by the firm free entry condition. For workers, that number is instead determined by the equation

$$n_p s_U I + \frac{I n_{iw}\left[(1-\bar{x}) + \bar{x}\exp\left(-h_b T\right)\right]}{T} = N_0^w \lambda_0 \Big[1 - \exp\big(-\rho_\varepsilon (V^{iw}(0) - V_0^w)\big)\Big]. \quad (7)$$

The left-hand side represents the inflow rate into the platform-search stage. The first term corresponds to experienced (premium-membership) workers who receive separation shocks, while the second accounts for inexperienced workers who leave after the optimal stopping time.[11] The right-hand side represents outflows of workers successfully joining platforms.

The number of open vacancies within a specific platform i is pinned down by

$$\frac{N_0^e}{I}\lambda_0 \Big[1 - \exp\big(-\rho_\varepsilon (V_i^e - V_0^e)\big)\Big] = (q_i + s_V) n_{V,i}, \quad (8)$$

where, again, the left-hand side represents inflows into that pool and the right-hand side represents outflows.

[11] In that second term, the bracketed term in the numerator corresponds to the probability that an inexperienced worker will not be matched before reaching the optimal stopping time. Dividing by the stopping time T converts that probability into a normalized outflow rate, which is then multiplied by the aggregate number of inexperienced workers in all platforms.

The steady-state number of workers in the experimentation stage in platform i satisfies

$$\frac{N_0^w}{I}\lambda_0\Big[1-\exp\big(-\rho_\varepsilon(V_i^{iw}(0)-V_0^w)\big)\Big](1-\theta)=n_{iw,i}\left[\frac{(1-\bar{x})+\bar{x}\exp\big(-h_{b,i}T_i\big)}{T_i}+\bar{x}h_{b,i}\right]. \tag{9}$$

Based on put previous assumption regarding membership choices, the equation above is also sufficient to pin down the number of basic membership unemployed workers in the platform:

$$n_{b,i}=n_{iw,i}. \tag{10}$$

The pools of employed and premium membership workers are pinned down by

$$\begin{aligned} n_{iw,i}\bar{x}h_{b,i}+n_{p,i}h_{p,i}&=\delta n_{J,i}\\ \delta n_{J,i}+\frac{N_0^w}{I}\lambda_0\Big[1-\exp\big(-\rho_\varepsilon(V_i^{iw}(0)-V_0^w)\big)\Big]\theta&=\big[h_{p,i}+s_U\exp\big(-\rho_\nu(V_i^p-V_0^w)\big)\big]n_{p,i}, \end{aligned}$$

which can be rewritten as

$$n_{p,i}=\frac{\bar{x}h_b}{s_U\exp\big(-\rho_\nu(V_i^p-V_0^w)\big)}n_{iw,i} \tag{11}$$

$$n_{J,i}=\frac{h_p+s_U\exp\big(-\rho_\nu(V_i^p-V_0^w)\big)}{\delta}n_{p,i}. \tag{12}$$

3.5 Optimal Platform Choices and Equilibrium

The equations in the last section are sufficient to characterize all endogenous variables of the model given a guess of symmetric platform choices (f_m,f_p,γ), along with the number of operating platforms I. To finish the characterization of a decentralized equilibrium, we assume that platform choices maximize steady-state flow profits, and that the number of platforms is pinned down by a free entry condition.

To make this problem more tractable, we make a "monopolistic competition" assumption: platforms take the outside values V_0^w and V_0^e, along with the number of platforms, as given. Then, we can write the platforms profit-maximizing problem as

$$\Pi(V_0^w,V_0^e,I)=\max_{f_{m,i},f_{p,i},\gamma_i} f_m w(\cdot)n_J(\cdot)+f_p n_p(\cdot)-F, \tag{13}$$

where quantities $n_J(\cdot), n_p(\cdot)$ and wages $w(\cdot)$, are functions of both (V_0^w, V_0^e, I) and platform choices.

To shed light on this problem, we reformulate it as a two-step process,

$$\Pi(V_0^w, V_0^e, I) = \max_{f_{m,i}, \gamma_i} \left[\max_{f_{p,i}} f_m w_i(\cdot) n_J(\cdot) + f_p n_p(\cdot) - F \right].$$

First, the platform chooses f_m and the distortion γ, and then picks f_p in a second stage.

To analyze the "inner" problem, it is necessary to understand the implications of the choice of matching distortion. If $\gamma = 0$, the platform does not offer any advantage to premium workers, and workers would not pay positive fees $f_p > 0$ for it. In that case, we assume without loss that the inner solution is $f_p = 0$, such that we can maintain our previous assumptions regarding experienced workers choosing the premium memberships (which are, in this case, identical to the basic membership).

If $\gamma > 0$, then the platform may extract fees from premium-membership workers in exchange for higher match rates. To do so, the utility associated with the premium membership tier must be at least as high as that of the basic membership as perceived by experienced workers, that is, $V^p \geq V^b$. But we show in Appendix A.5 that, if this inequality is strict, then the platform can strictly increase profits with a marginal increase in the premium fee f_p, accompanied by correctly-sized reduction in match fees f_m. Thus, the profit-maximizing platform equates the utility of the basic and premium memberships for experienced workers.

Equalizing $V^p = V^b$ yields

$$\begin{aligned} f_{p,i} =& (h_{p,i} - h_{b,i})(J_i^w - V_i^b) \\ =& \frac{\lambda \gamma_i \left(1 - e^{-n_{p,i}/D}\right)}{n_{p,i}} (J_i^w - V_i^b) \end{aligned}$$

That is, the premium fee must be proportional to the increase in the match rate provided by the premium membership. Multiplying both sides by $n_{p,j}$ yields an expression for total revenues raised from premium membership fees

$$f_{p,i} n_{p,i} = \lambda \gamma_i \left(1 - e^{-n_{p,i}/D}\right) (J_i^w - V_i^b).$$

Now consider the second stage part of the platform's problem: choosing $f_{p,i}$ conditional on

$(f_{m,i}, \gamma_i)$. Note that $(J_i^w - V_i^b)$ does not depend directly on $f_{p,i}$ when the utilities of both tiers are equalized. This is because if utilities by tier are the same, the value of being unemployed or employed in the platform can be calculated by assuming that workers always choose the basic tier. Similarly, the inflow of workers and firms into the platform is also invariant to $f_{p,i}$ conditional on $(f_{m,j}, \gamma_j)$ and on the utilities being equalized. Thus, the expression above implies that revenues from premium memberships are strictly increasing in $n_{p,i}$, even accounting for the fact that the firm has to reduce f_p to maintain utilities equalized between tiers.

The intuition for this result is that a larger pool of premium workers makes the matching mechanism more efficient conditional on $\gamma > 0$. Hence, if the firm wants to distort the matching process to extract rents from premium workers, the most efficient way to do so is to charge fees that are low enough to induce all experienced workers to join that tier, making the preference awarded to them less distortionary.

However, inexperienced workers will not see any benefit from joining the premium membership. This is because, for them, the expected benefit in increased match rates is multiplied by their perceived probability of being a good match to the platform. Thus, if experienced workers are indifferent between the two tiers, inexperienced workers will surely prefer the basic membership.

Why would the platform not set fees low enough so that inexperienced workers become indifferent, and thus join the premium membership as well—making it even less distortionary? If everyone is in the premium tier, then γ is irrelevant. The model becomes one in which the platform charges a participation fee for all workers. However, we ruled out this possibility early on by disallowing participation fees (see Appendix A.1).

Using these results and plugging in some expressions from the previous section, the problem of the platform can be rewritten as

$$\Pi = \max_{f_{m,i},\gamma_i} \underbrace{n_{iw}\Big(f_{m,i},\, h_b(\cdot)\Big)}_{\text{Acquisition}} \underbrace{\bar{x}\, h_b(\cdot)}_{\text{Conversion}} \Bigg[\underbrace{\frac{h_p(\cdot)+s_U}{s_U\,\delta} f_{m,i}}_{\substack{\text{Match}\\ \text{fees}}} + \underbrace{\frac{1}{s_U}\left[h_p(\cdot)-h_b(\cdot)\right] \frac{(y-f_{m,i})(r+s_U)\beta \;-\; r\,s_U\,V_0^w}{(r+\delta)\,(r+s_U) \;+\; r\,h_b(\cdot)}}_{\substack{\text{Premium}\\ \text{membership fees}}} \Bigg] \;-\; \underbrace{F}_{\substack{\text{Operational}\\ \text{costs}}}, \quad (14)$$

where the $(\cdot)$ notation is used to emphasize which components are functions of platform choices.

This expression highlights the margins through which $f_{m,i}$ and γ_i affect profits. In the absence of a premium tier, the trade-off associated with the choice of f_m is fairly standard. On the one hand, higher $f_{m,i}$ increases revenues from match fees, the first term inside square brackets. On the other hand, those fees reduce the expected value of being in the platform, which reduces the "acquisition" component. Note that, as f_m increases, the reduction in the number of employers increases the job-finding rates, $h_b(\cdot)$ and $h_p(\cdot)$, because there are less congestion effects between workers. That effect dampens the first negative effect: even though the inflow of inexperienced workers is reduced, the "Conversion" component increases, as do employment rates within the platform.[12]

Now consider the choice of γ_i. Its main effect is to reduce $h_b(\cdot)$, which has a direct positive effect on the "premium fees" component of total profits. That is for two reasons. First, it increases the job-finding differential, $h_p(\cdot) - h_b(\cdot)$. Second, it increases the value of being employed in the platform relative to being unemployed, which corresponds to the fractional term in the "premium fees" component.

The reduction in h_b decreases n_{iw} mostly due to an effect similar to that of increasing f_m: it reduces the value of being employed in the platform, J^{iw} (see Equation (**??**) in Appendix B). However, unlike the previous case, where h_b partly offsets the decline in n_{iw}, here h_b falls as well. Thus, increases in γ_i can cause significant drops in the number of employed workers in the platform.

There is no guarantee that this problem has an internal solution in which the platform uses both revenue sources. Given the negative impacts on inflows and employment noted above, it may well be the case that the platform may choose $\gamma_i = 0$. Alternatively, it is also possible to have corner solutions where $f_m = 0$. To see why, note that there is some degree of substitutability between revenue sources. When $f_m > 0$, an increase in γ will reduce revenues from match fees by reducing employment in the platform. And when $\gamma > 0$, an increase in f_m decreases premium membership fees by compressing the gap between the values of employment and unemployment. Thus, in a scenario where the platform derives significant revenues from membership fees, a marginal increase in f_m, starting from zero, may reduce

[12] If the platform is simply choosing f_m with $\gamma = 0$, the solution is to set the elasticity of employment in the platform with respect to f_m equal to -1. Because of the dampening effects mentioned above, the elasticity of inflows n_{iw} at the optimal point will be somewhat lower than -1.

revenues. We view the possibility of these corner solutions an appealing feature of our modeling approach, given the fact that real-world platforms differ in their revenue strategies (as shown in Table 1).

When is the platform more likely to implement premium membership fees? The profit expression shows that a key determinant is the s_U parameter. Specifically, when s_U is low relative to job-finding rates, the platform will likely have a large pool of unemployed workers. This yields the key reason for introducing premium memberships: it provides a means to extract rents from those workers, who are less elastic than new entrants. In addition, the premium membership tier for workers is more likely to be useful for firms when β is large, that is, when workers capture a larger share of the match surplus. Again, that is related to the fact that premium membership fees are designed to extract rents specifically from workers, and there is more room for such extraction when the surpluses they accrue are larger.

3.6 Extensions

Now, we discuss three extensions that incorporate realistic features of online job platforms and help increase the accuracy of our quantitative exercises in the next section. The first two increase the likelihood that premium memberships are an effective revenue-generating strategy, but have different welfare implications. The third extension constrains the platform's market power on experienced workers, and thus limits the extent to which their rents can be extracted via premium memberships.

3.6.1 Phantoms

We now extend the model to allow for phantom (i.e., inactive) worker accounts left behind in platforms after a worker switches. Specifically, when a worker leaves a platform (which happens at the exogenous rate s_U), their account remains in the platform as a "phantom". Phantoms disappear from the platform at an exogenous rate $s_\dagger$.[13] This parameter should be interpreted as the speed with which the platform can reliably identify phantom accounts and exclude them from the matching algorithm. Crucially, we will assume that phantom accounts always choose the basic membership, as workers who leave the platform would

[13] One could generalize this model such that only a fraction ϕ of the workers leaving the platform generate a phantom. For practical purposes, such a model would be isomorphic to the one we use once the phantom separation rate is adjusted to $s'_\dagger = s_\dagger/\phi$.

have no incentives to keep paying for a premium membership.[14]

The presence of phantoms among basic users shifts the platform's trade-offs on the optimality of introducing a premium membership tier. Specifically, it provides another rationale for such a tier: it provides a means for active workers to convey their higher availability through a costly signal. To the extent that the matching procedure downweights basic-membership workers, that can be efficient if that pool contains many phantoms. However, this could also have distributional implications, as inexperienced workers may face welfare losses compared to experienced ones.

The mechanics of the model are very similar, with a few changes to the equations that determine the steady-state pools. Let $n_{ph,i}$ denote the number of phantoms in platform i. The equation that determines the number of basic-membership users in the platform becomes

$$n_{b,i} = n_{iw,i}\bar{x} + n_{\dagger,i}. \tag{10.a}$$

The pool of phantom workers is pinned down by

$$n_{p,i}s_U = n_{\dagger,i}s_\dagger, \tag{15}$$

which equates inflows into to outflow from the phantom pool. The main effect of phantoms is to reduce the job-finding rate $h_{b,i}$ for workers in the platform.

Phantom accounts can also interpreted as a device to incorporate the consequences of multihoming in our model. Suppose that inexperienced workers can split their attention across several platforms at the same time. When they do that, their job-finding rate in each platform is reduced in proportion to the number of platforms they participate in, reflecting diluted search efforts. This alternative model would imply that the number of inexperienced worker

[14] Under the literal logic of the model, the platform would be able to perfectly identify phantoms by membership switching, and thus eliminate such workers from the matching pool whenever such a switch is identified. We do not allow platforms to use this strategy because, in real online job platforms, there are unmodeled considerations that could make such a strategy sub-optimal. One example are workers who exit the premium membership temporarily due to cash flow constraints, but who would still accept jobs (and that could later return to the premium membership if the match occurred). Another consideration is that workers who switch memberships may be marginal in the sense that they are in the process of switching to another platform. The cost of retaining those workers could be lower than the cost of acquiring new ones, so that removing those workers from the matching pool would not be optimal; if anything, the platform should be increasing their match rates instead. The "phantom" problem we model here should be viewed as reflecting the platform's limitation in identifying which of those workers are truly phantoms and which are marginal workers for whom targeted promotion would be effective.

accounts is larger than the overall number of experienced workers, and that job-finding rates for inexperienced workers fall the larger that gap is. Those are precisely the mechanical implications of reducing $s_\dagger$ in the model with phantoms.

3.6.2 Private information regarding suitability

...

3.7 Welfare and the Planner's Problem

We define aggregate welfare flow as total output net of platform operation costs and joining costs for workers and firms

$$\begin{aligned} rW = (n_J y - F)I \\ - \frac{\lambda_0 N_0^w}{\rho_\varepsilon}\left[1 - \exp\left(-\rho_\varepsilon(V^{iw}(0) - V_0^w)\right)\left(1 + \rho_\varepsilon(V^{iw}(0) - V_0^w)\right)\right] \\ - \frac{\lambda_0 N_0^e}{\rho_\varepsilon}\left[1 - \exp\left(-\rho_\varepsilon(V^e - V_0^e)\right)\left(1 + \rho_\varepsilon(V^e - V_0^e)\right)\right]. \end{aligned} \tag{16}$$

We are also interested in how welfare is split across workers and firms (note that the platform's share of aggregate welfare is zero due to the free entry condition). The share of welfare going to employers is given by

$$WS^e = \frac{N_0^e V_0^e + In_V V^e + In_J J^e}{W}. \tag{17}$$

The corresponding shares for experienced workers is

$$WS^{ew} = \frac{In_J J^w + In_p V_p^w}{W}. \tag{18}$$

Finally, the share accruing to inexperienced workers—either at platforms or currently seeking a platform—is given by $WS^{iw} = 1 - WS^e - WS^{ew}$.

As a benchmark for studying inefficiencies introduced by imperfect competition between platforms and pricing constraints, we introduce the concept of a planner who can control platform creation and operation, but cannot implement lump-sum transfers, control firm and worker behavior directly, nor overcome informational frictions by fiat. We also assume the planner may not implement participation fees in the platforms, a modeling choice that we

discuss in Appendix A.1.

The planner chooses the number of platforms in the economy I and their design choices (f_m, f_p, γ) with the goal of maximizing welfare. As with the individual platform's problem, the planner's problem has two degrees of freedom. Given a choice of (I, f_m, γ), premium fees f_p must be such that experienced, unemployed workers are indifferent between the basic and premium memberships. And given (I, γ), match fees f_m must be such that the planner's budget is balanced, that is, platforms make zero profits in the aggregate.

4 Quantitative Exercises

4.1 The Experiment Through the Lens of the Model

We use data from the platform—including the experimental data—to calibrate the model. That requires a clear conceptual mapping between the model and the experiment.

We begin with the assumption that the platform in our data is one of many operating platforms in the symmetric market equilibrium. The experiment is modeled as an exogenous increase $\Delta\gamma > 0$ in the distortion parameter, holding fees f_m and f_p fixed; in other words, the platform deviates from its profit-maximizing choices. The increase in γ applies only to firms randomly assigned to the treatment group.

To analyze the experiment's implications for the platform's internal outcomes, we extend the model to track treated and untreated vacancies separately. Although the inflow of firms into each pool is equalized by design (each firm is assigned to treatment with probability 0.5 at its first search during the experimental period), the *stocks* of vacancies diverge because the experiment generates different vacancy-filling rates: control firms exit the platform slightly faster than treated firms.

The worker pools, by contrast, are *not* split between treatment and control groups. Hence our modeling allows for the indirect effect that treated firms can have on control firms through their shared interactions with workers.

4.2 Calibration

To evaluate the welfare implications of premium memberships, we calibrate the model with three sources of information: (i) parameter values drawn from the literature, (ii) direct cali-

Table 4: Externally calibrated parameters

Parameter	Value	Rationale
r	0.004	Yearly discount of 5%
δ	2.00	Av. emp. spell of half a month (see Appendix)
N_w	10000	Normalization
y	100.0	Normalization
β	0.50	Symmetric bargaining
λ_0	150.00	Five arrivals per day while platform-searching
D	7000.00	Other values considered in robustness exercises.

Table 5: Internally calibrated parameters

Parameter	Value	Moment	Target	Sim.
F	910.2	Number of platforms	10.00	10.00
λ	3445.825	Average number of jobs found by workers over two months	0.124	0.124
s_V	1.448	Share of platform vacancies eventually filled	0.293	0.293
s_U	2.282	Share of premium-membership workers in platform	0.115	0.115
ρ_ε	0.196	Match fees as fraction of gross wage	0.150	0.150
ρ_v	0.283	Premium fees as fraction of platform revenue	0.100	0.100
c	11.856	Ratio of workers to firms in platform	3.549	3.549
$\bar{x}$	0.423	Optimal wait time (months)	1.000	1.000
$\Delta\gamma$	0.0199	Treat. effect on prob. hiring a basic membership worker	-0.0176	-0.0176
$s_\dagger$	0.2194	Treat. effect on prob. hiring any worker	-0.0090	-0.0090

Data sources: Internal data from the platform.

bration based on institutional knowledge of the platform and its market, and (iii) a moment-matching procedure. The unit of time in the model is a month.

The first subset of parameters appears in Table 4, along with brief justifications for each choice. We model a labor market of 10,000 workers, which should be viewed as a segmented, observable category of workers in the overall market (for example, "C# .NET programmers"). We choose a high value for the arrival rate of platform offers in the platform-search stage, so that most time spent in search occurs inside platforms. We calibrate the job destruction δ to two, implying an expected job length of half a month. We choose this parameter so that, when all other targets are met, the monthly price of a premium memberships relative to the hourly wage is approximately the same as the true values for the platform around the time of the experiment. See Appendix **??** for details.

The remaining ten parameters, listed in Table 5, are chosen by simulating the model and matching a corresponding number of targets. Most of the targets correspond to observable patterns in the data, including the experimental effects on employment of any workers and, specifically, basic-membership workers. Two of them—match fees as a fraction of gross wage and premium fees as a share of platform revenue—are obtained directly from the reports from the platform itself. The equilibrium number of entering platforms is based on the relatively large number of platforms that we observe operating in the United States (see Table 1), while also noting that some of them cater to specific workers or firms (such that the number of platforms that could be used by any given worker is less than the overall number of platforms). Finally, we set the optimal waiting time for inexperienced workers to one month as a rough guess; future versions of this paper will include robustness exercises regarding this choice.[15]

Recall from the presentation of the matching technology that the parameter D regulates the "thin markets" effect in the platforms. Our data does not provide a clear way of identifying that parameter. For that reason, we use three different values in our counterfactual simulations. The baseline value of 7,000 is chosen such that the planner's optimal choice of platform creation is close to the market outcome of 10 platforms in the estimated model. In some counterfactual exercises, we consider alternative calibrations of the model where D is either 4,000 or 10,000.

4.3 The Planner's Solution and Premium-Membership Prohibition

Table 6 presents a series of model outcomes for three different scenarios. The first column corresponds to the baseline estimated equilibrium. The second column measures the same outcomes under the planner's optimal solution. The last column reports outcomes for a market equilibrium where platforms are prevented from using premium memberships, that is, they can only raise revenues through match fees.

Panel A shows effects on welfare and employment. Both increase by around five percent, both in the planner solution and under the zero-γ policy. The policy achieves a similar level

[15] We numerically verify that the solution of this square system of equations seems to have a single solution, by trying a large number of starting points in the estimation process and checking that all solutions are identical. We also numerically verify that the equilibrium of the model seems to be unique for the estimated parameters, and that endogenous variables vary smoothly with exogenous parameter changes (suggesting that the equilibrium is stable). While we do not have a formal proof of uniqueness, we have verified a large number of parameter combinations and have never found a case of multiplicity of equilibria.

Table 6: Planner Solution and Premium Membership Prohibition Policy

Outcome	Mkt. Equilibrium	Planner	Market, no premium
Panel A: Aggregate metrics			
Welfare, normalized (base = 100)	100.0	105.3	105.2
Employment rate	0.099	0.105	0.105
Search costs as share of output	0.604	0.605	0.605
Panel B: Market outcomes			
Number of platforms	10.0	10.4	10.3
Match fees	0.150	0.163	0.162
Match distortion γ	0.050	0.000	0.001
Premium membership fees	0.53	0.00	0.01

of welfare as the planner's solution for two reasons. First, and most importantly, the parameters inferred from the data imply that the efficiency losses from distorting the matching mechanism outweigh the benefits of reducing congestion, such that the planner has no use for premium memberships. The second is that the planner's optimal level of firm creation is very close to the baseline level of platform entry, as discussed above regarding the choice of D. We will see below that, once we consider alternative values for D, the planner's solution will feature higher welfare than the prohibition policy scenario.

Our finding that premium memberships reduce welfare becomes even more informative when we connect it to the results in Pallais (2014). That paper shows that, once we account for the possibility of inefficient hiring of inexperienced workers, interventions that encourage employers to hire such workers can be welfare-improving. Premium memberships push in the opposite direction: they raise experienced workers' hiring rates while lowering hiring rates for inexperienced workers. We therefore conjecture that incorporating the channel highlighted by Pallais (2014) would strengthen our main conclusion.

4.4 Thin Markets and Optimal Platform Creation

As previously discussed, a high level of D implies a more severe thin-markets effect. For any finite D, creating more platforms is productive because the rate of matches that the platform can create is fixed, given by the λ parameter. But it is easy to show that in the limit where D and λ go to infinity but their ratio remains constant, the rate of match formation in the platform is proportional to that ratio times the number of workers in the platform. That is, match rates become independent of platform size. In that limit, a planner would thus prefer to create a single platform to avoid paying platform operating costs. More generally, the

Table 7: Thin Markets and Optimal Platform Creation

Outcome	$D = 4000.0$	$D = 7000.0$	$D = 10000.0$
Panel A: Planner solution			
Welfare gain (%)	7.0	5.3	5.8
Employment rate change (%)	10.4	5.5	4.6
Number of platforms	14.6	10.4	8.4
Match fees	0.212	0.163	0.135
Panel B: Premium membership prohibition policy			
Welfare gain (%)	5.0	5.2	5.3
Employment rate change (%)	4.3	5.4	5.8
Number of platforms	10.2	10.3	10.4
Match fees	0.161	0.162	0.162

higher the D parameter, the smaller the optimal number of platforms.

Table 7 reports counterfactual simulations for alternative calibrations of the model using different values for the matching function parameter D. The baseline value $D = 7,000$, used in the previously-discussed table, corresponds to the middle column. The other columns show results for a lower and a higher value for this parameter. Note that in every column, all parameters are estimated to match the same set of moments in the data. In particular, in all columns, the number of operating platforms in the descentralized market equilibrium is 10.

As expected, the planner's choice on the number of platforms is decreasing in D. The welfare gains of moving to the planner's solution are the lowest for the baseline value of $D = 7000$. That's because, for the other values, the planner can improve on the market allocation by not only eliminating premium memberships but also correcting the entry problems. But notably, the welfare gains from the prohibition policy are remarkably robust to the choice of D.

4.5 How Does the Experiment Inform the Welfare Results?

Our final result discusses how the information provided by the experiment is used by the model to infer welfare impacts. Recall that the estimation procedure uses two outcomes from the experiment: the treatment effect on the hiring of any type of worker (a reduction of around 0.9%) and the treatment effect on the hiring of workers with a "basic" membership (a reduction of around 1.8%). Table 8 shows alternative scenarios where we change the value of these moments, re-estimate the model, and re-calculate the counterfactual scenario where we eliminate premium memberships. The first column shows the baseline estimation. In the

Table 8: How the Experiment Informs Counterfactual Results

Outcome	Baseline estimation	Both outcomes 20% lower	Any hiring 20% lower	Basic hiring 20% lower
Panel A: Predicted gains from prohition of premium memberships				
Welfare gain (%)	5.24	5.25	4.50	6.25
Panel B: Estimated parameters				
$\bar{x}$	0.423	0.422	0.477	0.367
$s_\dagger$	0.219	0.222	0.146	0.537
$\Delta\gamma$	0.020	0.018	0.020	0.018
Panel C: Composition of basic accounts in the platform				
Suitable inexperienced workers	0.23	0.23	0.19	0.28
Unsuitable inexperienced workers	0.31	0.31	0.21	0.49
Phantoms	0.46	0.46	0.60	0.22

second column, we reduce the magnitude of both coefficients by 20%. The third column reduces only the coefficient on any hiring, while the fourth reduces the coefficient on basic-membership hiring.

The results in Panel A show that changing both coefficients makes no difference for the predicted welfare gains from the policy. On the other hand, changing only one of the coefficients leads to sizable changes in those gains. Panels B and C clarifies why that is the case. When both coefficients change, the model's inference is simply that the experiment corresponded to a smaller change in the matching technology, that is, $\Delta\gamma$ is 20% smaller. On the other hand, a change in the ratio of the effects on all workers relative to basic workers is informative about the composition of the basic-membership tier. Specifically, when the effects on any hiring are small relative to basic-membership hiring, that means increasing the advantage to premium workers does not seem to impact overall hiring in the platform by much. Through the lens of the model, that means the basic-membership tier is more heavily populated by phantoms, relative to suitable—but inexperienced—workers. In that case, the premium membership is not as distortionary, and thus the gains of eliminating it are smaller.

5 Conclusion

In this paper, we develop a search-and-matching model to study labor markets in which intermediaries charge users for preferential treatment in the matching process. Our theory clarifies the trade-offs in intermediaries' decisions to offer premium memberships and high-

lights the resulting welfare and distributional consequences.

We calibrate the model with rich internal data and a large-scale field experiment from an online labor platform. Counterfactual simulations show that premium memberships have a significant negative effect on market performance. These losses arise because platforms degrade the match efficiency of basic users to induce them to upgrade. The model also implies that the decentralized number of operating platforms may be higher or lower than optimal. While our experiment does not allow us to identify the optimal level of platform entry, our analysis suggests that, unless that number is very different from the observed market outcome, the potential gains from correcting entry are probably smaller than the gains from eliminating premium memberships.

Our modeling framework can be extended to study other dimensions of platform design. For example, Horton (2017) provides experimental evidence that well-calibrated recommendation algorithms can raise match efficiency. A version of our model that includes endogenous investment in match quality could shed light on the equilibrium effects of such investments and their interaction with imperfect competition among platforms. More broadly, we believe that our model can be used to study other forms of labor market intermediation where workers interact repeatedly with intermediaries, and those intermediaries are allowed to charge fees in return for preferential treatment in the matching process. Potential examples include staffing agencies, international labor recruiters, union halls with dues-based priority, or creative talent agencies.

References

Agrawal, Ajay, John J Horton, Nicola Lacetera, and Elizabeth Lyons. 2015. "Digitization and the contract labor market: A research agenda." In *Economic Analysis of the Digital Economy*. , ed. Avi Goldfarb, Shane M Greenstein and Catherine E Tucker, 219–250. University of Chicago Press.

Atkeson, Andrew G., Andrea L. Eisfeldt, and Pierre-Olivier Weill. 2015. "Entry and Exit in OTC Derivatives Markets." *Econometrica*, 83(6): 2231–2292.

Autor, David H. 2008. "The Economics of Labor Market Intermediation: An Analytic Framework." Working Paper 14348.

Barach, Moshe A., and John J. Horton. 2021. "How Do Employers Use Compensation History? Evidence from a Field Experiment." *Journal of Labor Economics*, 39(1): 193–218.

Bhuller, Manudeep, Dominico Ferraro, Andreas Ravndal Kostøl, and Trond Christian Vigtel. 2023. "The Internet, Search Frictions and Aggregate Unemployment." Unpublished manuscript.

Caillaud, Bernard, and Bruno Jullien. 2003. "Chicken and egg: competition among intermediation service providers." *RAND Journal of Economics*, 34(2): 309–328.

Cheron, Arnaud, and Bruno Decreuse. 2016. "Matching with Phantoms." *The Review of Economic Studies*, 84(3): 1041–1070.

Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora. 2013. "Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment." *Quarterly Journal of Economics*, 128(2): 531–580.

Cumming, Douglas, Lars Hornuf, Moein Karami, and Denis Schweizer. 2023. "Disentangling Crowdfunding from Fraudfunding." *Journal of Business Ethics*, 182(4): 1103–1128.

de Cornière, Alexandre. 2016. "Search Advertising." *American Economic Journal: Microeconomics*, 8(3): 156–88.

de Cornière, Alexandre, and Greg Taylor. 2019. "A model of biased intermediation." *The RAND Journal of Economics*, 50(4): 854–882.

DeVaro, Jed, and Oliver Gürtler. 2018. "Advertising and Labor Market Matching: A Tour through the Times." *Journal of Labor Economics*, 36(1): 253–307.

Diamond, Peter A. 1984. *A search-equilibrium approach to the micro foundations of macroeconomics.* Vol. 1982, MIT press.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2018. "Consumer Price Search and Platform Design in Internet Commerce." *American Economic Review*, 108(7): 1820–1859.

Dyck, Alexander, Adair Morse, and Luigi Zingales. 2010. "Who Blows the Whistle on Corporate Fraud?" *The Journal of Finance*, 65(6): 2213–2253.

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer. 2023. "The Emergence of Market Structure." *Review of Economic Studies*, 90(1): 261–292.

Farboodi, Maryam, Gregor Jarosch, Guido Menzio, and Ursula Wiriadinata. 2025. "Intermediation as Rent Extraction." Forthcoming, Journal of Economic Theory.

Filippas, Apostolos, John J. Horton, Prasanna Parasurama, and Diego Urraca. 2025*a*. "Costly capacity signaling increases matching efficiency: Evidence from a field experiment." Unpublished manuscript.

Filippas, Apostolos, John J. Horton, Prasanna Parasurama, and Diego Urraca. 2025*b*. "Sponsored advertising in labor markets: Evidence from a field experiment." Unpublished manuscript.

Hagiu, Andrei, and Bruno Jullien. 2011. "Why do intermediaries divert search?" *The RAND Journal of Economics*, 42(2): 337–362.

Hagiu, Andrei, and Bruno Jullien. 2014. "Search diversion and platform competition." *International Journal of Industrial Organization*, 33: 48–60.

Hagiu, Andrei, and Julian Wright. 2015. "Marketplace or Reseller?" *Management Science*, 61(1): 184–203.

Horton, John J. 2010. "Online labor markets." 515–522. Berlin, Germany:Springer.

Horton, John J. 2017. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment." *Journal of Labor Economics*, 35(2): 345–385.

Horton, John J. 2025. "Price floors and employer preferences: Evidence from a minimum wage experiment." *American Economic Review*, 115(1): 117–146.

Horton, John J, Moshe A Barach, and Jacob Golden. 2020. "Buyer Uncertainty About Seller Capacity: Causes, Consequences, and a Partial Solution." *Management Science*, 66(11): 5114–5129.

Horton, John J., Ramesh Johari, and Philipp Kircher. 2024. "Sorting through Cheap Talk: Theory and Evidence from a Labor Market." LIDAM Discussion Paper CORE.

Kuhn, Peter, and Mikal Skuterud. 2004. "The Internet and Job Search." *American Economic Review*, 94(1): 218–232.

Kässi, Otto, Vili Lehdonvirta, and Fabian Stephany. 2021. "How many online workers are there in the world? A data-driven assessment." *Open research Europe*, 1: 53.

Lee, Kwok Hao, and Leon Musolff. 2025. "Two-sided markets shaped by platform-guided search." Unpublished manuscript.

Marx, Paul, and Joachim Schummer. 2021. "Revenue from Matching Platforms." *Theoretical Economics*, 16(3): 883–925.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *The Review of Economic Studies*, 61(3): 397–415.

Pallais, Amanda. 2014. "Inefficient Hiring in Entry-Level Labor Markets." *American Economic Review*, 104(11): 3565–99.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory.* . 2nd ed., MIT Press.

Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association*, 1(4): 990–1029.

Shimer, Robert, and Liangjie Wu. 2023. "Assortative Matching with Private Information." Working paper, University of Chicago.

Statista Research Department. 2022. "Gig economy in the U.S." *https://www.statista.com/topics/4891/gig-economy-in-the-us/*, Accessed: 2025-05-09.

Stevenson, Betsey. 2008. "The Internet and job search." National Bureau of Economic Research.

Teh, Christopher, Chengsi Wang, and Makoto Watanabe. 2024. "Strategic limitation of market accessibility: Search platform design and welfare." *Journal of Economic Theory*, 216: 105798.