Wage Setting Protocols and Labor Market Conditions: Theory and Evidence*

Stanislav Rabinovich[†] Brenda Samaniego de la Parra[‡]
Ronald Wolthoff[§]

December 1, 2023

Abstract

We theoretically and empirically examine how firms' choices of wage-setting protocols respond to labor market conditions. We develop a simple model in which workers may be able to send multiple job applications and firms choose between posting wages and Nash bargaining. Posting a wage allows the firm to commit to lower wages than would be negotiated ex post, but eliminates the ability to respond to a competing offer, should the worker have one. We show that higher productivity lowers both the application-vacancy ratio and the fraction of firms posting a wage. On the other hand, an increase in the number of applications per worker raises the application-vacancy ratio while lowering the fraction of firms posting a wage. As a result, the equilibrium fraction of firms posting a wage may be positively or negatively correlated with the application-vacancy ratio, depending on the source of shocks. The model also implies that an increase in the number of applications per worker may lead to a decrease in the number of posting firms rather than a change in the wages posted by those firms. Empirically, we demonstrate that the model's predictions are confirmed in a novel dataset from an online job board.

JEL Classification: J31, J64, C78, D21

Keywords: search frictions, applications, vacancies, wages, wage posting, bargaining

---

*We thank Yuri Bykov, Rachel Ceccarelli, Courtney Chamberlain, Jennifer Milan and Elizabeth Schillo for extensive consultations regarding the Dice.com data, platform, and business model. Samaniego de la Parra received compensation from DHI Group, Inc. for developing the DHI Database. Wolthoff gratefully acknowledges financial support from the Social Sciences and Humanities Research Council.

[†]University of North Carolina, Chapel Hill

[‡]University of California, Santa Cruz

[§]University of Toronto

1 Introduction

How wages are determined in frictional labor markets, and how wage determination depends on economic conditions, are classic yet unsettled economic questions. According to one prominent paradigm, wages are posted by firms, and workers select the best offer when faced with multiple offers (e.g. as in Burdett and Judd (1983) or Burdett and Mortensen (1998)). According to another classic paradigm, wages are not committed to, but instead determined by bilateral bargaining between workers and their employers (e.g. as in Pissarides (2000)), and possibly renegotiated in case of receiving a competing job offer (e.g. as in Cahuc et al. (2006)). These strands of the literature have operated in parallel and, with a few exceptions discussed below, have largely been separate.

This parallel but separate treatment of distinct wage-setting protocols is restrictive. As shown by Hall and Krueger (2012), job postings with and without explicit wages coexist in the data. Moreover, there may well be important interactions between firms who post wages and firms who don't: e.g. the wages firms choose to post may be affected by how many firms opt to not commit to wages, and the decision of whether to post a wage may depend on the wages posted by other firms. The coexistence of wage-setting modes raises important substantive questions: what determines the choice of wage-setting protocol? How does it vary with labor market conditions? And, does endogenous choice of wage-setting protocol amplify/dampen the effects of economic shocks or policies? To a large extent, the big obstacle to making progress on these dimensions has been the availability of data. We contribute to this effort both theoretically, by providing and analyzing a simple model with endogenous wage-setting protocols, and empirically, by confirming the model's predictions in a novel dataset with information on job applications as well as wage-posting choices.

Theory. The theory is a simple static model that introduces the possibility of wage bargaining into the classic framework of Burdett and Judd (1983). Firms post vacancies, at a cost, and decide whether or not to commit to a wage. Workers send either one or two applications each, and may therefore obtain multiple offers. The novelty is that not all firms may have chosen to post a wage. If all the job offers obtained by a worker come from firms who posted a wage, the worker simply chooses the best offer, as in Burdett and Judd (1983). If a firm has not posted a wage, and it is the worker's only match, the worker and firm engage in Nash bargaining, as in the classic Pissarides (2000) model. If a worker has two offers, one from a firm who posted a wage and the other from a firm who did not, the latter firm bargains with the worker, with the worker's outside option being accepting the wage of the former firm. Finally, if the worker has met two firms and neither firm has posted a wage, the two firms engage in Bertrand competition for the worker. Equilibrium is determined such that firms' posting vs. bargaining choices are made optimally: in particular, if some firms post and others bargain, each firm is indifferent between bargaining and posting.

The model incorporates the classic tradeoff for a firm who decides to post a wage: a higher wage results in lower profits conditional on hiring, but leads to a higher probability of hiring, more so if there is a larger chance of the worker having a second offer. The novelty is the tradeoff involved when deciding whether to post a wage at all. By posting a wage, the firm is able to commit to a lower wage than would be negotiated ex post. On the other hand, by posting a wage, the firm eliminates the ex post option to respond to a competing offer, should the worker have one. This logic suggests that the relative benefits of posting a wage are higher when there is a lower chance that the worker has a competing offer. Because firms' entry – and hence the number of offers – is determined endogenously, the probability that the worker has a competing offer is (in contrast to Burdett and Judd (1983)) dependent on labor market conditions, e.g. productivity. This is the case even though the number of applications per worker is taken as fixed; instead, it is the probability that an application is successful that varies endogenously.

The model yields results consistent with the above intuition. An increase in productivity raises firms entry and thereby lowers the ratio of applications per vacancy, which is the relevant measure of labor market slack in our model. This raises the probability that an individual application is successful, and therefore the probability that an applicant at a firm has a competing offer. Consistent with the above intuition, firms then have a weaker incentive to post a wage. This leads to the first main result: an increase in productivity lowers the application-vacancy ratio and lowers the fraction of firms who post wages. Our second result has to do with the effect of increasing the fraction of workers who send two rather than one application. This increase in average applications per worker likewise increases the chance that an applicant at a given firm as a competing offer, and therefore also lowers the fraction of firms who post wages.

The second set of results has to do with the effect on the distribution of posted wages. In the classic Burdett and Judd (1983) framework, either an increase in productivity or an increase in the fraction of workers sending two applications would raise posted wages, in the stochastic dominance sense. An increase in productivity raises equilibrium posted wages by raising the firm's benefit from hiring a worker. An increase in the fraction of workers sending two applications raises equilibrium posted wages by increasing the change that a firm's applicant has a competing offer. Consistent with these results, our model implies that an increase in productivity raises equilibrium posted wages. However, an increase in the fraction of workers sending two applications has no effect on posted wages, as long as firms are indifferent between bargaining and posting. Instead, in response to an increase in expected applications per worker, it is the fraction of firms posting wages that adjusts, rather than the distribution of wages posted by those firms.

Empirical analysis. Our model makes predictions regarding how the fraction of firms advertising an explicit wage, and the average wage they post, varies with labor market conditions, such as productivity and the number of applications per worker. In our empirical analysis, we test the model's predictions using proprietary data from `www.Dice.com` (henceforth, the DHI data), an online job board.[1] The DHI data links the universe of vacancy postings advertised on Dice.com between January 1, 2012 and December 31, 2017 to all of the job seekers that applied to them. On the job posting side, the data includes detailed information about the job posting, including the position's job title, location, tax-terms, work-schedule, and – crucially for our analysis – whether the posting offers an explicit wage (or wage range) and, if so, the amount offered. On the applicant side, the data includes the candidate's job title, their location, and the exact time-stamp and job for each of their applications.

In addition to providing information about the incidence of wage-posting and wage amounts, the data lets us construct two key measures of labor market conditions: the application-vacancy ratio and the number of applications per applicant, which both have model counterparts. Because productivity is not observed directly, we devise an alternative strategy to test the above theoretical predictions. The theory implies that, to the extent that productivity is driving the variation in the data, it would induce a *positive* correlation between the application-vacancy ratio and the fraction of firms posting a wage. It would also induce a negative correlation between the application-vacancy ratio and the average posted wage. On the other hand, if the variation in the data is driven by variation in applications per worker, the model implies that such variation would induce a *negative* correlation between the application-vacancy ratio and the fraction of firms posting a wage, and *no* correlation between the application-vacancy ratio and the average posted wage.

This set of predictions can be taken to the data. The first implication is that the model-implied correlation between the application-vacancy ratio and the incidence of wage posting is, in general, ambiguous, but is more likely to be positive once a measure of labor market competition (i.e. applications per worker) is controlled for. The second implication is that this measure of labor market competition itself is likely to be positively associated with the incidence of wage posting. Turning to the posted wages, our results imply that the application-vacancy ratio should be negatively associated with posted wages if applications-

[1] See Davis and Samaniego de la Parra (2019) for a detailed description of the data.

per-worker are controlled for, but this correlation should weaken if applications-per-worker are not controlled for. We show that these model predictions are consistent with the data.

Implications. Our theory implies that the incidence of wage posting is endogenous and therefore not invariant to labor market conditions. Our empirical analysis confirms this by showing that the incidence of wage posting covaries systematically with labor market slack. This result underscores the importance of modeling the wage-setting choice of a firm, as this wage-setting choice can compound the effects of economic shocks on labor market outcomes. The model with endogenous choice of wage-setting protocol also makes qualitatively different predictions from a model where a particular wage-setting protocol is assumed exogenously. Consider an increase in average applications per worker, which through the lens of our model, can be thought of as a measure of competition between employers. The standard Burdett and Judd (1983) framework implies that such an increase would induce employers to post higher wages and thereby affect workers' realized wages. In our model, an increase in applications per worker certainly affects labor market outcomes, but it does so not through a change in posted wages, but a change in the number of firms who post.

2 Relationship to the Literature

Our paper contributes to the literature that theoretically and empirically examines the coexistence, interaction, and consequences of various wage-setting protocols. On the theoretical side, a number of different frameworks contain meaningful tradeoffs between wage posting and wage bargaining. In Michelacci and Suarez (2006), this tradeoff arises because workers are heterogeneous in their productivity, which is unverifiable ex ante. Posting rather than bargaining allows the firm to attract a desired number of workers but eliminates the option to adjust the wage based on productivity ex post. In Cheremukhin and Restrepo-Echavarria (2020), workers and firms are heterogeneous and face search costs that prevent them from seamlessly locating the right partner. Posting a wage allows the firm to attract the right workers through self-selection but lowers the ex post profits compared to bargaining. In our model, workers and firms are homogeneous, and the firm's tradeoff is between being able to commit to a low wage and being able to respond to a competing offer. The key role assigned to the prevalence of competing offers makes our model more similar to Doniger (2023) and Flinn and Mullins (2021), who both consider environments with on-the-job search nesting Burdett and Mortensen (1998) and Cahuc et al. (2006). In Doniger (2023), renegotiation in response to competing offers is costly. In Flinn and Mullins (2021), the tradeoff is similar to ours, though their model does not have free entry and the specific comparative statics results of our model are novel. The key difference is where the outside offers come from. Whereas on-the-job search generates competing offers in Flinn and Mullins (2021) similarly to Burdett and Mortensen (1998) and Cahuc et al. (2006), in our framework they come from receiving multiple offers as in Burdett and Judd (1983). We also endogenize the probability of receiving multiple offers, thereby allowing the choice of wage-setting to respond to labor market conditions, leading to both a key feature of our mechanism and key testable predictions. Relative to all this prior work, we are able to test the model predictions empirically thanks to a novel dataset in which the number of applications per worker – an object with a clear model counterpart in the Burdett and Judd (1983) framework – can be measured.

On the empirical front, Hall and Krueger (2012) and Brenzel et al. (2014) document the co-existence of wage bargaining and wage posting. Banfi and Villena-Roldan (2018) and Marinescu and Wolthoff (2020) examine whether higher posted wages attract more applicants. We contribute to this literature by showing that the incidence of wage posting covaries systematically with measures of labor market slack. Our paper is also relevant for the literature on labor market competition and its effect on wages. Azar et al. (2020) and Azar et al. (2022) argue that labor market power, as measured by employer concentration, is relevant for wages. We view our work as complementary and conceptually related, though distinct: in our analysis, it is the number of applications per worker that serves as a measure of labor market competition; furthermore, we argue that an increase in labor market competition can manifest itself not only (or not necessarily) through the magnitude of posted wages but through the fraction of firms who post wages at all.

3 Theory

3.1 Environment

We consider a static model with an exogenous measure u of workers and a larger measure of firms. Each firm can post one vacancy at cost κ. Workers apply to firms randomly and get matched according to the process described below. A matched worker and firm produce y, where $y > \kappa$, whereas workers and firms that remain unmatched produce zero.

Matching. An exogenous fraction $\alpha \in [0, 1]$ of workers sends out 2 applications each. The remaining fraction $1 - \alpha$ of workers sends out one application each. The total number of applications is therefore

$$s = u(1 - \alpha + 2\alpha) = u(1 + \alpha).$$

There is a meeting technology that matches applications to vacancies bilaterally. If v is the measure of vacancies and s is the measure of applications, the total number of meetings is given by $M(v,s)$, where M is concavely increasing in both arguments, exhibits constant returns to scale, and satisfies $0 < M(v,s) \leq \min\{v,s\}$ as well as $dM(0,s)/dv = dM(v,0)/ds = 1$. We define the vacancy-unemployment ratio $\theta = v/u$, and the application-vacancy ratio $\lambda = s/v = (1+\alpha)/\theta$. The probability that a vacancy receives an application is then

$$\frac{M(v,s)}{v} = M\left(1, \frac{s}{v}\right) \equiv q(\lambda).$$

while the probability that an application gets matched to a vacancy is given by

$$\frac{M(v,s)}{s} = M\left(\frac{v}{s}, 1\right) \equiv p(\lambda) = q(\lambda)/\lambda.$$

Our assumptions regarding the meeting function M imply that q is increasing in λ from $q(0) = 0$ to $q(\lambda) \to 1$ when $\lambda \to \infty$. In contrast, p is decreasing in λ from $p(0) = 1$ to $p(\lambda) \to 0$ for $\lambda \to \infty$.

Competition. We next derive the probability, from a firm's point of view, that the worker who has a job offer from that firm also has a competing job offer. Conditional on receiving an application, the probability that the worker sending the application has sent a second application equals $\frac{2\alpha}{1+\alpha}$, and the probability that the latter application is successful is $p(\lambda)$. Therefore, the probability that the worker has another job offer is equal to

$$\psi \equiv \frac{2\alpha}{1+\alpha} p(\lambda). \tag{1}$$

This will be a key object from the firm's standpoint, as it determines the probability that the firm is competing with another firm for the applicant. This probability is increasing in the application-vacancy ratio λ and will therefore be endogenously determined.

Wage setting protocols. Each firm decides whether to post a wage (as in Burdett and Judd, 1983) or to bargain (as in Cahuc et al., 2006). This decision is made before meeting an applicant and hence *before* the firm knows whether its applicant has a competing offer. The wage in the match is then determined as follows. If the worker only has one job offer, and that firm posted a wage w, the worker either accepts or rejects. If the worker has only one job offer from a firm that chose to bargain, the worker receives βy, and the firm receives $(1-\beta)y$, where $\beta \in (0,1]$ is worker's bargaining weight.[2] Next, consider a worker with two

[2] We restrict attention to positive β as it is easy to show that all firms will bargain when $\beta = 0$.

job offers. If both firms have posted a wage, the worker picks the firm with the higher wage and becomes employed at that wage, with the other firm remaining vacant; if both firms posted the same wage, one is chosen at random. If both firms have chosen to bargain, the two firms engage in Bertrand competition over the worker, which drives the wage to y (and the worker randomly picks one of the firms). Finally, if one of the firms has chosen to bargain and the other has posted a wage of w, then the wage with the first firm is decided by Nash bargaining, with worker bargaining weight β and the outside option being the wage w at the second firm. The outcome is that the worker is employed by the bargaining firm at the wage of $\beta y + (1 - \beta)w$.

Discussion of modeling choices. Our static model is deliberately kept very simple and designed to transparently illustrate the role of the novel element, namely the endogenous choice of bargaining vs. posting.[3] Apart from the endogenous wage-setting protocol, the model also departs from the original Burdett and Judd (1983) framework by assuming free entry. This assumption implies that the probability of a worker facing a competing offer is endogenous: this probability depends not only on the probability that the worker sent a second application (which is assumed exogenous here) but also on the chance that the second application is successful.[4] This model feature is generates an endogenous link between economic conditions (e.g. more firms entering due to higher productivity) and the competitiveness of the labor market, which in turn affects the incentives to post a wage.

3.2 Equilibrium

An individual firm chooses whether or not to post a vacancy, and – conditional on posting a vacancy – whether or not to post a wage. Let λ be the application-vacancy ratio, let $\phi \in [0, 1]$ be the endogenous fraction of firms that choose to post a wage, and let $F(w)$ be the probability distribution of posted wages among them. When making its individual decisions, a firm takes these objects as given.

Consider a firm posting a wage w. Conditional on having an applicant, there are two ways for the firm to obtain a payoff $y - w$. First, with probability $1 - \psi$, the applicant does not have another offer. Second, with probability ψ, the applicant has another offer. In that case, the firm will succeed to hire the worker if the competing firm posts a wage that is lower than w (where ties in the posted wage are broken randomly). Hence, the expected payoff of

[3]The model can be extended in a number of ways, e.g. allowing for an arbitrary stochastic number of applications, without changing the main results and intuition.

[4]For a similar feature, see Kaplan and Menzio (2016).

a firm posting w is

$$\Pi(w) = -\kappa + q(\lambda)\left[1 - \psi + \psi\phi \int \left(\mathbf{1}_{w'<w} + \frac{1}{2}\mathbf{1}_{w'=w}\right) dF(w')\right](y - w). \tag{2}$$

Naturally, the firm will choose the wage that maximizes $\Pi(w)$, and thus obtain a payoff

$$\Pi = \max_w \Pi(w). \tag{3}$$

Next, consider a firm that decides to bargain. Conditional on having an applicant, there are again a few possibilities. With probability $1-\psi$, the applicant has no other offer and the firm obtains a payoff $(1-\beta)y$. With probability $\psi(1-\phi)$, the applicant has a competing offer from a bargaining firm. In this case, Bertrand competition implies that the worker extracts the entire surplus, leaving a zero payoff for either firm. Finally, with probability $\psi\phi$, the applicant has a competing offer from a firm that has posted some wage $w \sim F(w)$. In this case, the bargaining firm hires the worker and obtains a payoff $(1-\beta)(y-w)$. Hence, the expected payoff for a firm that decides to bargain equals

$$\Pi_b = -\kappa + q(\lambda)\left[(1-\psi)(1-\beta)y + \psi\phi(1-\beta)\int (y-w)\,dF(w)\right]. \tag{4}$$

We can then define an equilibrium for this economy as follows:

Definition 1. *An equilibrium consists of a number λ, a distribution F, a number $\phi \in [0,1]$, and numbers Π_b and Π such that*

1. Π_b*, and* Π *satisfy* (4) *and* (3) *with* $\Pi(w)$ *given by* (2)*;*

2. $\Pi(w) = \Pi$ *for every* w *in the support of* F*;*

3. ϕ *satisfies*

$$\phi \begin{cases} = 0 & \text{if } \Pi_b > \Pi, \\ \in (0,1) & \text{if } \Pi_b = \Pi, \\ = 1 & \text{if } \Pi_b < \Pi. \end{cases} \tag{5}$$

4. *Free entry holds:* $\max\{\Pi_b, \Pi\} = 0$.

To characterize the equilibrium, we first consider the equilibrium wage-setting choices of firms taking as given the application-vacancy ratio λ. We then use the free-entry condition to determine the equilibrium value of λ.

We first rule out the case where all the firms bargain. Intuitively, if a firm's applicant has a competing offer from a firm that has chosen to bargain, the former firm gets zero profits. Thus, if all firms bargain, then an individual firm only obtains a positive payoff if it has an applicant with no other offers. However, the firm's chances of meeting and hiring such an applicant are exactly the same if it posts a wage of zero instead, which is a profitable deviation. The following lemma formalizes this result.

Lemma 1. *Given $\beta > 0$, there does not exist an equilibrium in which all the firms bargain.*

Proof. See Appendix A.1. □

Given that some firms post wages, the next step in our equilibrium derivation is to characterize the distribution $F(w)$. This part of the analysis closely follows the standard logic of Burdett and Judd (1983), adjusted for the fact that some firms might bargain in equilibrium. Specifically, a firm must be indifferent between any two wages in the support of the wage distribution. This uniquely pins down the distribution.

Lemma 2. *In any equilibrium, the distribution F of wages among the firms that post wages has connected support $[0, \bar{w}]$ and no mass points, and satisfies*

$$F(w) = \frac{1}{\phi}\frac{1-\psi}{\psi}\frac{w}{y-w} \tag{6}$$

and

$$\bar{w} = \frac{\phi\psi}{1-(1-\phi)\psi}y, \tag{7}$$

where $\psi = \frac{2\alpha}{1+\alpha}p(\lambda)$.

Proof. See Appendix A.2. □

Intuitively, suppose that a firm slightly increases its posted wage, leading to lower profits conditional on hiring but, possibly, a greater likelihood of hiring. The strength of the latter depends on the probability ψ that the firm's applicant has a competing offer. This is the standard result of Burdett and Judd (1983), captured by the presence of ψ in (6) and (7). The novelty is that the profit increase from increasing one's posted wage also depends on the likelihood that the competing offer is from a firm that posted a wage, which is captured by the presence of ϕ in the expressions (6) and (7). If fewer firms post wages, then an individual wage-posting firm is less likely to attain the worker whether it raises its wage or not.

Having characterized the distribution of posted wages, we are then in a position to consider the fraction of firms that choose to post a wage rather than bargain.

Lemma 3. *In equilibrium, all firms post, i.e. $\phi = 1$, if and only if*

$$\psi \leq 1 - \exp\left(-\frac{\beta}{1-\beta}\right). \tag{8}$$

Otherwise, the equilibrium fraction of firms that post equals

$$\phi = \frac{(1-\psi)}{\psi}\left[\exp\left(\frac{\beta}{1-\beta}\right) - 1\right] < 1. \tag{9}$$

Proof. See Appendix A.3. □

A firm's decision whether to post a wage depends on the worker's bargaining power β and the (endogenous) probability ψ that a firm is competing with another firm for its applicant. The benefit from posting a wage is the ability to commit to a lower wage than would be bargained ex post; this benefit is stronger if the worker bargaining power is high. On the other hand, the cost of posting a wage consists of the inability to respond to a competing offer should the worker has one; this cost is stronger if the likelihood of a competing offer is high. These two forces determine the firm's incentive to post rather than bargain. Furthermore, from the perspective of an individual firm, the aforementioned cost of posting a wage is strongest if the competing offer is likely to be from a firm who *also posts*: otherwise, the ability to respond to the competing offer would be irrelevant. In equilibrium, it follows that one of two cases occur. The equilibrium may be such that a firm strictly prefers to post even when every other firm posts, so that $\phi = 1$; this occurs if worker bargaining power is sufficiently high. Alternatively, the equilibrium may feature a positive fraction of bargaining firms, such that a firm is just indifferent between bargaining and posting.

The final step in the equilibrium derivation is to characterize the application-vacancy ratio. The following lemma does this and establishes uniqueness of the equilibrium.

Lemma 4. *A unique equilibrium exists. The equilibrium application-vacancy ratio solves*

$$-\kappa + q(\lambda)\left[1 - \frac{2\alpha}{1+\alpha}p(\lambda)\right]y = 0. \tag{10}$$

Proof. See Appendix A.4. □

The intuition is straightforward. By Lemma 1, any equilibrium features some wage posting. Lemma 2 then implies that the equilibrium payoff of any firm entering the market must equal the payoff of a firm posting the wage of zero. Since such a firm only hires if the applicant has no competing offer, we immediately obtain λ from (10). The equilibrium λ, in turn, pins down the equilibrium ϕ via Lemma 3 and the wage distribution via Lemma 2.

Substituting $\psi = \frac{2\alpha}{1+\alpha}p(\lambda)$ in Lemma 3, we then obtain a complete equilibrium characterization. If the λ solving (10) is such that $\frac{2\alpha}{1+\alpha}p(\lambda) \leq 1 - \exp\left(-\frac{\beta}{1-\beta}\right)$, the equilibrium has $\phi = 1$, and the wage distribution is given by

$$F(w) = \frac{1-\psi}{\psi}\frac{w}{y-w}, \tag{11}$$

where $\psi = \frac{2\alpha}{1+\alpha}p(\lambda)$. The worker bargaining power is so high that all the firms post wages in equilibrium; the economy behaves like a standard Burdett and Judd (1983) economy, except for the addition of free entry. On the contrary, if the λ solving (10) is such that $\frac{2\alpha}{1+\alpha}p(\lambda) > 1 - \exp\left(-\frac{\beta}{1-\beta}\right)$, the equilibrium has

$$\phi = \frac{1-\psi}{\psi}\left[\exp\left(\frac{\beta}{1-\beta}\right) - 1\right] < 1, \tag{12}$$

and the wage distribution is now given by

$$F(w) = \left[\exp\left(\frac{\beta}{1-\beta}\right) - 1\right]^{-1}\frac{w}{y-w}, \tag{13}$$

where again $\psi = \frac{2\alpha}{1+\alpha}p(\lambda)$. In this case, a positive measure of firms decide not to post a wage. A firm is indifferent between posting and bargaining, and, conditional on posting, indifferent among all the wages in the support of the wage distribution.

3.3 Comparative Statics

Our focus is on generating model-based comparative statics predictions regarding the behavior of the application-vacancy ratio λ, the incidence of wage posting ϕ and the distribution of posted wage F, which can then be confronted with data. Since productivity is a likely driver of labor market conditions, we first examine comparative statics with respect to y. We then contrast them with the effect of α, which can be thought of as measuring the degree of competition in the labor market.

Lemma 5. *An increase in productivity y*

1. *decreases the application-vacancy ratio λ.*

2. *decreases the fraction ϕ of firms that post a wage, as long as $\phi < 1$.*

3. *increases the distribution of posted wages in a first-order stochastic dominance way.*

Proof. See Appendix A.5. □

A higher value of productivity encourages entry by firms, lowering the equilibrium application-vacancy ratio λ. With fewer applications per vacancy, the probability $p(\lambda)$ that an application results in a job offer then increases, as evident from (10). For a firm making a job offer to an applicant, it then becomes more likely that the applicant has a competing offer. This make it less attractive to post a wage, all else equal. As a result, the equilibrium fraction ϕ of firms that post a wage decreases, as long as the economy is already in the interior region where $\phi < 1$, as evident from (12). Finally, the increase in productivity has a direct effect on the wage distribution via either (11) or (13).

We next examine the effect of an increase in α, the fraction of workers sending two applications.

Lemma 6. *An increase in the fraction α of workers sending two applications*

1. *increases the application-vacancy ratio λ.*
2. *decreases the fraction ϕ of firms that post a wage, as long as $\phi < 1$.*
3. *increases the distribution of posted wages in a first-order stochastic dominance way if $\phi = 1$, but has no effect on wages if $\phi < 1$.*

Proof. See Appendix A.6. □

An increase in α directly increases the probability ψ that an applicant has a competing offer. Although the effect is mitigated by a reduction in entry, Lemma 6 establishes that the direct effect dominates; competition for workers rises, making wage posting less attractive and thereby reducing ϕ. Of particular interest is the effect of α on the posted wage distribution. If all the firms are posting wages ($\phi = 1$), the marginal effect of an increase in α is to raise wages, in the first-order stochastic dominance sense. This well-known and intuitive result from Burdett and Judd (1983) is evident from (11): an increase in competition for workers raises wages. The interior ($\phi < 1$) case, illustrated by absence of α in the expression in (13), is the less intuitive result, which runs quite contrary to the standard Burdett and Judd (1983) logic. All else equal, an increase in α raises a firm's marginal benefit from raising its wage as it essentially makes its labor supply more elastic. However, at the same time, an increase in α lowers the fraction of firms who post wages, thereby lowering an individual firm's marginal benefit from raising its wage. These two effects exactly cancel each other out. Note that this does not mean that labor market competition has no effect on workers' wage prospects: it will affect realized wages both directly, by enabling workers to sample more offers on average, and indirectly, by inducing more firms to bargain. The novelty is that the change in firm's behavior manifests through a lower probability of posting wages, not through the wages that they post conditional on doing so.

3.4 Predictions

Taking stock of the above comparative statics, we formulate model-based predictions to be taken to data. We focus on the interior case ($\phi < 1$) as this will be the empirically relevant case in the data. As explained below, our data allows for measurement of the application-vacancy ratio in a local labor market and, less trivially, also permits the construction for a proxy for labor market competition, α. It also contains data on the fraction of firms posting wages, and on the wages they post. We do not have direct information on productivity, y. Instead, the main predictions that we wish to test will have to do with the model-implied correlations between the application-vacancy ratio and the fraction of firms posting a wage, as well as between the application-vacancy ratio and the average posted wage.

The key idea is the following. Suppose that the variation in these endogenous objects in the data is driven by some combination of variation in α and variation in y. To the extent that the variation in the data is driven by variation in y, Lemma 5 implies that such variation would induce a *positive* correlation between the application-vacancy ratio and the fraction of firms posting a wage. It would also induce a negative correlation between the application-vacancy ratio and the average posted wage. On the other hand, if the variation in the data is driven by variation in α, Lemma 6 implies that such variation would induce a *negative* correlation between the application-vacancy ratio and the fraction of firms posting a wage, and (since $\phi < 1$) *no* correlation between the application-vacancy ratio and the average posted wage.

This set of predictions can be taken to the data. The first implication is that the model-implied correlation between the application-vacancy ratio and the incidence of wage posting is, in general, ambiguous, but is more likely to be positive once a measure of labor market competition (i.e. applications per worker) is controlled for. The second implication is that this measure of labor market competition itself is likely to be positively associated with the incidence of wage posting. Turning to the magnitude of the posted wages, our results imply that the application-vacancy ratio should be negatively associated with posted wages if the proxy for labor market competition α is controlled for, but this correlation should weaken if this proxy is not controlled for.

4 Empirical Evidence

In this section, we test the model's predictions regarding the correlation between applications per vacancy, incidence of wage posting, and the magnitude of posted wages. Confronting these predictions with empirical evidence requires data on the fraction of firms posting an explicit wage, as well as the magnitude of the wages posted by those firms. It also requires measuring the empirical counterparts of λ, the average applications per vacancy, and α, which is proxied for by the average applications per job seeker. First, we describe the data used in these analyses. Then, we explain how we construct empirical analogs to the objects λ and α. Next, we develop our empirical specification and present the estimation results.

4.1 Data Description

We use proprietary data from www.Dice.com (henceforth, the DHI data),[5] an online job board that focuses on technology sectors, software development, other computer-related occupations, engineering, financial services, business and management consulting, and other careers that require technical skills.[6] The DHI data links the universe of vacancy postings advertised on Dice.com between January 1, 2012 and December 31, 2017 to all of the job seekers that applied to them. It includes detailed information about the job posting, including the position's job title, location, tax-terms, work-schedule, and, crucially for our analysis, whether the posting offers an explicit wage (or wage range) and, if so, the amount offered.[7] We can also observe the exact time-stamp when the job posting was first visible to potential applicants, the number of seconds it was visible to applicants on each day, and the number of views and applications it received on each day that it was active on the platform. On the applicant side, we can observe the candidate's job title, their location, and the exact time-stamp and job for each of their applications.[8,9]

We focus on jobs posted by "Direct Hire" employers. These refer to companies that hire workers to add to their payroll or directly employ as contractors, as opposed to jobs posted by recruitment or staffing firms, who post jobs to find candidates for other firms to consider hiring or leasing. These employers represent 82 percent of all job posting accounts and a third of posted vacancies on the Dice.com platform. This sample restriction allows us to focus on jobs in which the employer and job poster coincide with the organization that determines the wage setting protocol. We further restrict the sample to jobs in the United States (which represent 99% of all jobs). The Dice.com platform allows employers to post,

[5]DHI Group Inc owns and operates the Dice.com platform.

[6]Dice.com has operated in the US since 1996, first, offering services to recruiting and staffing firms and later (1999) expanding to direct hire employers.

[7]The "tax-terms" variable includes information about the contract type (i.e., 1099 or W2) and the hiring modality (i.e., corporation-to-corporation or corporation-to-individual). Work schedule specifies whether the job is part-time, full-time, remote, and at times whether there are travel requirements.

[8]The applicant's job title is self-reported and may refer to the current, most recent, or desired job title. Based on conversations with DHI, our understanding is that many of the applicants that are active on the platform search on-the-job, so the job title likely reflects current positions. However, we cannot distinguish between those who are and are not employed.

[9]See Davis and Samaniego de la Parra (2019) for detailed variable description and summary statistics.

remove, and re-post the job several times. We focus on cases where the employer post the job and then permanently removes it from the platform within 31 days after first posting. Three-fourths of vacancies on Dice.com exhibit this posting pattern.[10] Henceforth, we refer to these postings as "standard postings."

For each posting, employers can choose to include information on the job's compensation. This variable is a free-text field in which employers can include specific values or ranges of wages/salaries, information about benefits, bonuses, etc. The field can also be left blank or include generic terms like "Negotiable," "Competitive," or "Market-based." There is also variation in terms of the frequency of the reported compensation, that is, a job posting may state that it pays \$50 per hour or \$96,000 per year.[11]

We construct three variables using the information available on compensation. The first variable is an indicator ($wage\_posted_{i,j,t}$) equal to 1 if the job posting i posted by employer j with first-active date on calendar month t includes a specific pay level or range, and zero if it does not. For the second variable, conditional on $wage\_posted_{i,j,t} = 1$, we calculate the natural logarithm of the wage or the lower end of a range if the posting mentions a range ($\ln(w_{i,j,t})$).[12,13] The third variable is categorical and describes the pay rate frequency. While there are six possible values to this variable (daily, hourly, weekly, biweekly, monthly, and annual), job postings with per hour and per year wages account for 58 percent and 39 percent of all jobs that offer an explicit wage.[14] We trim the bottom and top 5 percent of wages within each pay frequency. For our baseline analyses, we focus on the set of job postings that do not post a wage or that post an hourly wage. Our results are robust to instead focusing on jobs with annual posted wages.

We define a labor market using job function-location categories. We distinguish between 52 locations, consisting of the 49 states in the continental United States, D.C., and Puerto Rico. Job functions refer to occupational categories such as "Programmer, or "Consultant." We group postings into 56 different job functions categories.[15] This categorization covers

[10]Focusing on jobs with this "standard" posting pattern allows us to likely identify job postings that advertise a single position. In contrast, jobs that are posted and remain active for long periods of time likely reflect recurring hiring needs and multiple job openings. See Davis and Samaniego de la Parra (2023) for additional information on standard postings and the distribution of posting duration.

[11]Under a full-time, 8-hour day, 5-day week schedule with 48 workweeks per year, these hypothetical jobs have equivalent pay.

[12]We express wages in 2016 dollars using the GDP implicit price deflator.

[13]Results are robust to using the midpoint of the range. However, since using the midpoint increases measurement error due to data cleaning challenges, we use the smallest value of the range for the baseline results.

[14]5% of jobs offer an explicit wage but do not specify the pay rate frequency. For these jobs, we impute the frequency using the distribution of wages across the six pay rate frequency categories.

[15]A job posting can mention more than one job function in its description. We classify postings based on the first job function mentioned in the job title.

93% of all standard job postings in the Dice.com platform. We further restrict our sample to job markets with at least 20 distinct employers and 20 distinct job postings in each month of the sample period. Our baseline sample includes 2,846,814 job postings, posted by 4,100 employers across 217 job markets. [16] Table 1 shows additional descriptive statistics.

Table 1: Descriptive Statistics (Means and Std. Dev.)

Panel A: Across Job Postings (N = 2,846,814)					
Applications	Views	Posting duration	% Jobs w/ posted wages	Median Hourly Wage	
10.6	156.1	9.4	51.8	10.1	
(28.1)	(191.0)	(8.1)	(538.7)	(30.1)	
Panel B: Across Labor Markets (Job Function X Location; N= 217)					
Job postings	Firms	Applications	Job Seekers	Competing offers ($\bar{\alpha}$)	Applications/Vacancies (λ)
237.4	116.2	2260.8	1169.1	17.5	8.6
(148)	(83)	(1023.5)	(714.5)	(13.1)	(6.1)
Panel C: Across Job Seekers' Search Spells (N = 1,977,718)					
Applications	N Spells	Spell duration			
9.2	1.9	9.5			
(57.3)	(1.8)	(21.1)			

Source: Own calculations using the DHI Vacancy and Application Flow Database (2012-2016).

Notes: Standard deviations reported in parenthesis. Wages are expresed in 2016 USD using the GDP implicit price deflator.

4.2 Application-Vacancy Ratio and Labor Market Competition

In the data, we can directly observe the application-vacancy ratio for each labor market in each time period. Let A_i be the total number of applications received by posting i. Let

[16] Importantly, the set of job seekers include only those individuals applying to the jobs in our selected sample. However, we consider all the applications submitted by these individuals, including those to jobs that we exclude from the sample. This allows us to correctly capture employers' competition for job seekers, across all job postings regardless of whether the competing jobs meet our sample criteria.

$\mathbf{I}[i \in l \times f, t]$ be an indicator function equal to 1 if job posting i belongs to job market $l \times f$ and was first active on the platform in month t. The application-vacancy ratio, $\lambda_{l\times f,t}$ is given by:

$$\lambda_{l\times f,t} = \frac{\sum_{i=1}^{I} A_i \times \mathbf{I}[i \in l \times f, t]}{\sum_{i=1}^{I} \mathbf{I}[i \in l \times f, t]} \tag{14}$$

A firm's decision whether to post a wage or bargain with a worker upon matching depends on whether the employer believes the worker has a competing offer. As shown in equation 1, the probability of a competing offer depends on the application-vacancy ratio and on the probability that the worker submitted other applications (α).[17] We next describe how we construct an empirical analog to α using a measure of the expected number of competing applications submitted by each job posting's applicants.

In the data, we can directly observe the applications submitted by each job seeker.[18] To calculate the number of competing applications for each job posting, we first construct search spell as the set of applications that are likely simultaneously considered for each job seeker. We define a search spell as the set of all applications submitted by a job seeker such that the length of time between applications is less than 30 days.[19] The average applicant has 2 search spells on the platform with a mean duration of 8 days. On an average search spell, job seekers submit 9 applications. Let $A_{-i,k}$ be the number of applications submitted to job postings, different from posting i but within the same search spell, by job seeker k (who has also applied to i). Then, the mean number of competing applications submitted by the set of job seekers that applied to job posting i is given by:

$$\alpha_i = \frac{\sum_{k=1}^{A_i} A_{-i,k}}{A_i} \tag{15}$$

Employers do not directly observe the number of competing applications submitted by the job seekers that apply to their job postings before determining whether to offer an explicit wage. Instead, the firm considers the number of competing applications submitted by the

[17]In the theoretical model, we assume α is an exogenous parameter. Arguably, the optimal number of applications submitted by optimizing job seekers may vary. Importantly, in a related paper Samaniego de la Parra (2023) finds that, while labor market tightness affects the type of job postings that job seekers target, the number of applications submitted is not significantly affected.

[18]Some job seekers are flagged by the platform or self-identify as "3rd-party applicants." These applications are difficult to track to a single job seeker and we therefore exclude them when calculating the number of competing applications submitted by the average job seeker.

[19]The results are robust to using 15 or 45 days as the maximum time between applications when defining search spells.

mean job seeker applying to postings in the same labor market.

$$\alpha_{l\times f,t} = \frac{\sum_{i=1}^{I} \alpha_i \times \mathbf{I}[i \in l \times f, t]}{\sum_{i=1}^{I} \mathbf{I}[i \in l \times f, t]} \tag{16}$$

4.3 Empirical Analysis

We consider two outcome variables: first, each firm's decision to post an explicit wage for a given job posting, and second, the chosen wage level, conditional on choosing to post a wage. As discussed in section 3.4, the model implies an ambiguous correlation between the application-vacancy ratio (λ) and the incidence of wage posting (ϕ). The reason for the ambiguity is that shocks to the number of applications submitted by job seekers (α) or to productivity (y) imply different signs for the correlations between λ and ϕ. If neither of these shocks are not controlled for in an empirical specification, then a regression of wage posting incidence on the application-vacancy ratio includes the combined effect of both sources of variation. The model provides guidance on how we should expect this coefficient to change as we control for some of the potential sources of variation.

We consider a set of empirical specifications that seek to control for different potential sources of variation. Specifically, we estimate regressions of the form

$$wage\_posting_{i,j,t} = \tilde{\beta}_1 \ln\left(\lambda_{l\times f(i),t}\right) + \gamma_{l\times t} + \eta_{f\times t} + \xi_{j,t} + \epsilon_{i,j,t} \tag{17}$$

$$wage\_posting_{i,j,t} = \beta_1 \ln\left(\lambda_{l\times f(i),t}\right) + \beta_2 \ln\left(\bar{\alpha}_{l\times f(i),t}\right) + \gamma_{l\times t} + \eta_{f\times t} + \xi_{j,t} + \epsilon_{i,j,t} \tag{18}$$

We include a high-dimension set of fixed effects in order to control for idiosyncratic time-variation within firms, within job functions, and within locations. The two specifications (17) and (18) differ in whether we control for the expected competition for job seekers (α) or allow it to vary. In the latter case, the interpretation of the coefficient of λ includes the effect caused by variation in α. We additionally include progressively fewer time fixed effects. The logic behind these alternative specifications is to allow the coefficient on λ to encompass the effect from a wider range of productivity shocks.

Table 2 presents the results. Consistent with the model's predictions, column (2) shows that the coefficient on λ is positive when controlling for α, and the coefficient on α itself is negative. The same result holds across different fixed-effects specifications. An additional observation is that, for any given choice of fixed effects, the coefficient on λ is larger when α is not controlled for than when it is (e.g. the coefficient on λ is smaller in column (1) than in column (2), and a similar observation holds when comparing (3) and (4), or (5) and (6)). This can likewise be rationalized by the reasoning above: failure to control for α means that the regression coefficient on λ captures the variation induced by α, which leads to the opposite sign compared to productivity shocks.

Table 2: Labor Market Conditions and the Probability of Wage Posting

	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(\lambda)$	0.005***	0.006***	0.006***	0.007***	0.003***	0.005***
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
$\ln(\alpha)$		-0.002***		-0.002***		-0.003***
		(0.001)		(0.001)		(0.001)
Observations	2,825,656	2,825,656	2,846,807	2,846,807	2,846,807	2,846,807
R-squared	0.378	0.378	0.260	0.260	0.260	0.260
Firm × Month FE	Y	Y	Firm only	Firm only	Firm only	Firm only
State × Month FE	Y	Y	Y	Y	Y	Y
JF × Month FE	Y	Y	Y	Y	JF only	JF only
absorbed FE	133613	133613	7152	7152	5854	5854

Source: Own calculations using the DHI Vacancy and Application Flow Database (2012-2016).

Notes: Standard errors are clustered at the labor market × time level where a labor market is defined by job function and location. *, **, and *** indicate significance at the 0.1, 0.05, and 0.01 levels, respectively.

We next look at the effect on offered wages. The theory abstracts from idiosyncratic firm heterogeneity that is likely relevant to affect both firms' incentives to post a wage and the distribution of wages. In the data, we of course only observe wages for those jobs in which the firm chooses to post a wage. As a result, the estimates from a fixed effect model like that shown in equation (19) would be subject to sample selection bias:

$$\ln\left(wage_{i,j,t} | wage\_posting_{i,j,t}\right) = \delta X_{l\times f(i),t} + \gamma_{l\times t} + \eta_{f\times t} + \xi_{j,t} + \epsilon_{i,j,t} \tag{19}$$

We instead take advantage of the firm-time, job function-time, and location-time fixed effects estimated in equations (17) and (18), and estimate

$$\ln\left(wage_{i,j,t}|wage\_posting_{i,j,t}\right) = \tilde{\delta_1}\ln\left(\lambda_{l\times f(i),t}\right) + \delta_3\widehat{\gamma_{l\times t}} + \delta_4\widehat{\eta_{f\times t}} + \delta_5\widehat{\xi_{j,t}} + \tilde{\epsilon}_{i,j,t} \quad (20)$$

$$\begin{aligned}\ln\left(wage_{i,j,t}|wage\_posting_{i,j,t}\right) =& \delta_1\ln\left(\lambda_{l\times f(i),t}\right) + \delta_2\ln\left(\bar{\alpha}_{l\times f(i),t}\right)\\ &+\delta_3\widehat{\gamma_{l\times t}} + \delta_4\widehat{\eta_{f\times t}} + \delta_5\widehat{\xi_{j,t}} + \tilde{\epsilon}_{i,j,t} \quad (21)\end{aligned}$$

Doing so essentially controls for idiosyncratic factors that are either firm-time-specific, location-time-specific, or job-time-specific and affect the probability of posting a wage.

Table 3: Labor Market Conditions and Posted Hourly Wages

	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(\lambda)$	0.003	-0.055***	-0.023***	-0.084***	0.029***	-0.020***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
$\ln(\alpha)$		0.090***		0.100***		0.075***
		(0.003)		(0.003)		(0.003)
N	255,268	255,268	259,178	259,178	259,178	259,178
R-squared	0.118	0.131	0.088	0.106	0.154	0.159
$\widehat{\text{Firm} \times \text{Month}}$ FE	Y	Y	Firm only	Firm only	Firm only	Firm only
$\widehat{\text{State} \times \text{Month}}$ FE	Y	Y	Y	Y	Y	Y
$\widehat{\text{JF} \times \text{Month}}$ FE	Y	Y	Y	Y	JF only	JF only

Source: Own calculations using the DHI Vacancy and Application Flow Database (2012-2016).

Notes: Sample includes postings that explicitly mention a wage per hour rate. For postings with a numerical posted wage but without a rate frequency, we input the frequency based on the distribution of wages with explicitly payment frequency. We trim that bottom and top 5% of the hourly wages (after imputing frequencies). We use the fixed effects estimated in equations (17) and (18) as controls. Bootstrapped standard errors with 1,000 replications are clustered at the labor market × time level where a labor market is defined by job function and location. *, **, and *** indicate significance at the 0.1, 0.05, and 0.01 levels, respectively.

Table 3 presents the results of estimating (20) and (21). Consistent with productivity being the driving source of variation, there is a negative association between λ and the average level of posted wages. Also consistent with the model's implications, the coefficient on α is positive, controlling for λ. The theory predicts that an increase in α, all else equal, should increase λ and leave wages unaffected. The interpretation of the coefficient is therefore that it is measuring the effect of an increase in α with a countervailing increase in productivity so as to keep λ fixed, which would increase posted wages. In summary, the empirical evidence in both Table 2 and Table 3 is consistent with the predictions of the theory when productivity is the source of residual variation.

5 Concluding Remarks

We have extended a simple search model to incorporate a choice between wage posting and wage bargaining and shown theoretically that the model yields clean predictions about how this choice of wage-setting protocol responds to labor market conditions. Despite its simplicity, our theory yields predictions regarding co-movement of key outcome variables that are consistent with empirical evidence from a rich novel dataset. We believe that the theoretical and empirical analysis here provides new insights on wage setting. In particular, our theoretical results imply that the choice of wage-setting protocol can act as an additional mechanism through which shocks transmit to labor market outcomes.

The contribution here has been of a qualitative nature. A full quantitative analysis of a model with endogenous wage-setting protocols, which would likely require a dynamic model, is a natural step forward in this line of research.

References

Azar, J., Marinescu, I., and Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, 57(S):S167–S199.

Azar, J., Marinescu, I., Steinbaum, M., and Taska, B. (2020). Concentration in us labor markets: Evidence from online vacancy data. *Labour Economics*, 66:101886.

Banfi, S. and Villena-Roldan, B. (2018). Do high-wage jobs attract more applicants? directed search evidence from the online labor market. *Journal of Labor Economics*, forthcoming.

Brenzel, H., Gartner, H., and Schnabel, C. (2014). Wage bargaining or wage posting? evidence from the employers' side. *Labour Economics*, 29:41–48.

Burdett, K. and Judd, K. L. (1983). Equilibrium price dispersion. *Econometrica*, 51:955–969.

Burdett, K. and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2):257–273.

Cahuc, P., Postel-Vinay, F., and Robin, J.-M. (2006). Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, 74(2):323–364.

Cheremukhin, A. and Restrepo-Echavarria, P. (2020). Wage setting under targeted search. *FRB St. Louis Working Paper*, (2020-41).

Davis, S. and Samaniego de la Parra, B. (2019). The DHI Vacancy and Application Flow Database: Record Layouts, Variable Descriptions, and Summary Statistics. Technical report, DHI Group Inc.

Davis, S. and Samaniego de la Parra, B. (2023). Application flows. unpublished manuscript.

Doniger, C. L. (2023). Wage dispersion with heterogeneous wage contracts. *Review of Economic Dynamics*.

Flinn, C. and Mullins, J. (2021). Firms' choices of wage-setting protocols. Technical report, Discussion paper, New York University.

Hall, R. E. and Krueger, A. B. (2012). Evidence on the incidence of wage posting, wage bargaining, and on-the-job search. *American Economic Journal: Macroeconomics*, 4(4):56–67.

Kaplan, G. and Menzio, G. (2016). Shopping externalities and self-fulfilling unemployment fluctuations. *Journal of Political Economy*, 124(3):771–825.

Marinescu, I. and Wolthoff, R. P. (2020). Opening the black box of the matching function: the power of words. *Journal of Labor Economics*, 38(2):535–568.

Michelacci, C. and Suarez, J. (2006). Incomplete wage posting. *Journal of Political Economy*, 114(6):1098–1123.

Pissarides, C. A. (2000). *Equilibrium Unemployment Theory*. Oxford University Press.

Samaniego de la Parra, B. (2023). Labor market tightness, recruitment and search behavior. unpublished manuscript.

Appendix A Proofs

A.1 Proof of Lemma 1

Suppose all firms choose to bargain, i.e. $\phi = 0$. In that case, firms' payoff equals Π_b, which reduces to

$$-\kappa + q(\lambda)(1-\psi)(1-\beta)y. \tag{22}$$

If a firm deviates to posting a wage of zero, it will obtain a payoff $\Pi(0)$, which equals

$$-\kappa + q\left(\lambda\right)\left(1-\psi\right)y. \tag{23}$$

Clearly, (23) is strictly larger than (22) given that $\beta > 0$. $\square$

A.2 Proof of Lemma 2

The proof that F has no mass points and connected support on $[0, \bar{w}]$ closely resembles Burdett and Judd (1983), We therefore keep the discussion of these results brief and refer to their paper for additional detail. Subsequently, we derive the expressions for F and $\bar{w}$ as given in (6) and (7), respectively.

As a first result, note that no firm will post a wage w equal to y in equilibrium, since that would lead to a payoff of $-\kappa < 0$ with certainty, which is worse than staying out of the market. Suppose then that F has a mass point at some wage $w_0 \in [0, y)$. The probability that a firm posting w_0 competes for an applicant with another firm posting the same wage is $q(\lambda)\psi\phi\mathbb{P}[w = w_0]$ which is strictly positive. In that case, the firm hires the worker with probability $1/2$. A profitable deviation then exists: increasing the wage by $\varepsilon \to 0$ leads to a certain hire in this scenario, increasing the payoff by $\frac{1}{2}q(\lambda)\psi\phi\mathbb{P}[w = w_0](y - w_0) > 0$, while only marginally changing payoffs in other scenarios. Hence, mass points in F cannot be part of an equilibrium.

Next, suppose that F has a gap in its support, denoted by (w_1, w_2). A firm offering w_2 then has an incentive to deviate to a wage between w_1 and w_2 which would not affect the firm's probability of hiring, but would increase the payoff conditional on hiring. Note that this argument also rules out a lower bound of the support of F that is strictly larger than zero. Hence, F must have connected support on $[0, \bar{w}]$ for some $\bar{w} \in (0, y)$.

The above results imply that $\Pi(w)$ reduces to

$$\Pi(w) = -\kappa + q(\lambda)\left[1 - \psi + \psi\phi F(w)\right](y - w). \tag{24}$$

To derive F, note that firms must be indifferent between all wages in its support. That is, $\Pi(w) = \Pi(0)$ for any $w \in [0, \bar{w}]$, where $\Pi(0) = -\kappa + q(\lambda)(1-\psi)y$. Solving for F yields equation (6). The upper bound of the support of F, as given in (7), then readily follows from the fact that $F(\bar{w}) = 1$. $\square$

A.3 Proof of Lemma 3

To derive the equilibrium value of ϕ, note that Lemma 2 implies that

$$\psi\phi \int (y-w)\, dF(w) = (1-\psi)y \int_0^{\bar{w}} \frac{1}{y-w}\, dw$$
$$= -(1-\psi)y \log\left(\frac{1-\psi}{1-(1-\phi)\psi}\right).$$

Substituting this into (4) implies that the payoff for a bargaining firm equals

$$\Pi_b = -\kappa + q(\lambda)(1-\psi)(1-\beta)y\left[1-\log\left(\frac{1-\psi}{1-(1-\phi)\psi}\right)\right]. \tag{25}$$

This expression is strictly increasing in ϕ, i.e. the more firms post, the more attractive it becomes for an individual firm to bargain. It is easy to see that $\Pi_b < \Pi(0)$ when $\phi = 0$. An interior solution for ϕ then requires that $\Pi_b > \Pi(0)$ when $\phi = 1$, which is the case if and only if

$$\beta < \frac{-\log(1-\psi)}{1-\log(1-\psi)} \in [0,1]. \tag{26}$$

In that case, the equilibrium value of ϕ follows from $\Pi_b = \Pi(0)$, which is equivalent to equation (9). $\square$

A.4 Proof of Lemma 4

By Lemma 3, any equilibrium features some wage posting. Lemma 2 then implies that the equilibrium payoff of any firm entering the market must equal

$$\Pi(0) = -\kappa + q(\lambda)\left[1-\frac{2\alpha}{1+\alpha}p(\lambda)\right]y, \tag{27}$$

which must equal zero in equilibrium. Taking the derivative of (27) with respect to λ yields

$$\frac{d\Pi(0)}{d\lambda} = q'(\lambda)\left[1-\frac{2\alpha}{1+\alpha}p(\lambda)\right]y - q(\lambda)\frac{2\alpha}{1+\alpha}p'(\lambda)\,y. \tag{28}$$

Since $q'(\lambda) > 0$ and $p'(\lambda) < 0$, this derivative is strictly positive, i.e. $\Pi(0)$ is strictly increasing in λ. Note further that $\Pi(0) = -\kappa < 0$ when $\lambda = 0$ and $\Pi(0) \to y - k > 0$ when $\lambda \to \infty$. Hence, there exists a unique value of λ such that the free entry condition is satisfied. $\square$

A.5 Proof of Lemma 5

By Lemma 4, the free entry condition can be written as $\Pi(0) = 0$ where $\Pi(0)$ satisfies (27). Taking the derivative with respect to y yields

$$\frac{d\Pi(0)}{dy} = \frac{d\Pi(0)}{d\lambda}\frac{d\lambda}{dy} + q(\lambda)\left(1 - \frac{2\alpha}{1+\alpha}p(\lambda)\right) = 0.$$

As discussed around equation (28), the derivative of $\Pi(0)$ with respect to λ is strictly positive, which implies that $d\lambda/dy < 0$. The lower value of λ increases $\psi = \frac{2\alpha}{1+\alpha}p(\lambda)$. This, in turn, decreases ϕ as long as $\phi < 1$, as follows from (9).

To consider the effect of the productivity shock on the distribution of posted wages, note that when ϕ is interior, we can use (9) to rewrite (6) as

$$F(w) = \frac{1}{\exp\left(\frac{\beta}{1-\beta}\right) - 1}\frac{w}{y-w}. \tag{29}$$

Hence, while there is a direct effect of y on $F(w)$, there is no indirect effect that operates through ϕ or ψ. The direct effect of y on $F(w)$ is negative, since

$$\frac{dF(w)}{dy} = -\left[\exp\left(\frac{\beta}{1-\beta}\right) - 1\right]^{-1}\frac{w}{(y-w)^2}. \tag{30}$$

That is, when ϕ is interior, wages are increasing in a first-order stochastic dominance way.

When $\phi = 1$, the wage distribution is given by (11). Hence,

$$\frac{dF(w)}{dy} = \frac{\partial F(w)}{\partial \psi}\frac{d\psi}{d\lambda}\frac{d\lambda}{dy} + \frac{\partial F(w)}{\partial y}. \tag{31}$$

Since $\frac{\partial F(w)}{\partial \psi} < 0$, $\frac{d\psi}{d\lambda}\frac{d\lambda}{dy} > 0$ and $\frac{\partial F(w)}{\partial y} < 0$, it follows that wages are again increasing in a first-order stochastic dominance way.

A.6 Proof of Lemma 6

By Lemma 4, the free entry condition can be written as $\Pi(0) = 0$ where $\Pi(0)$ satisfies (27). Taking the derivative with respect to α yields

$$\frac{d\Pi(0)}{d\alpha} = \frac{d\Pi(0)}{d\lambda}\frac{d\lambda}{d\alpha} - q(\lambda)\,y\frac{\partial \psi}{\partial \alpha} = 0.$$

As discussed around equation (28), the derivative of $\Pi(0)$ with respect to λ is strictly positive. Further, $\partial\psi/\partial\alpha > 0$. Hence, $d\lambda/d\alpha > 0$. Next, consider the total effect of α on ψ, including the effect that operates through the application-vacancy ratio. We get

$$\begin{aligned}\frac{d\psi}{d\alpha} &= \frac{\partial\psi}{\partial\alpha} + \frac{d\psi}{d\lambda}\frac{d\lambda}{d\alpha} \\ &= \frac{\partial\psi}{\partial\alpha}\left[\frac{d\Pi\left(0\right)}{d\lambda} + \frac{d\psi}{d\lambda}q\left(\lambda\right)y\right]\left(\frac{d\Pi\left(0\right)}{d\lambda}\right)^{-1}.\end{aligned}$$

This expression has the same sign as

$$\frac{d\Pi\left(0\right)}{d\lambda} + \frac{d\psi}{d\lambda}q\left(\lambda\right)y = q'(\lambda)(1-\psi)y > 0.$$

That is, the probability ψ that an applicant has a competing offer is increasing in α. This, in turn, decreases ϕ as long as $\phi < 1$, as follows from (9).

For the wage distribution, we again distinguish between $\phi = 1$ and $\phi < 1$. When $\phi < 1$, the wage distribution is given by (29), which is clearly independent of α. In contrast, when $\phi = 1$, the wage distribution is given by (11), which depends on α through ψ. In particular, the increase in ψ that results from the increase in α increases wages in a first-order stochastic dominance way.