

# How Do Firms Search for Workers Over Time?\*

Zoë Cullen<sup>†</sup>

Mitchell Hoffman<sup>‡</sup>

Felix Koenig<sup>§</sup>

February 2023

## Abstract

How do firms adjust wages and non-wage policies over time to fill their jobs? Using administrative data and two RCTs on a large staffing platform, we show that even when jobs are unfilled days or hours before the start of work, most firms do not raise wages. Rather, firms are much more likely to make non-wage adjustments such as lowering skill requirements. When wage adjustment occurs, it is usually in the positive direction and occurs close to when work begins. Firms' reluctance to raise wages does not appear to be due to inattention or time constraints. Firms reject opportunities to use different forms of automatic wage adjustments presented by the platform in the first RCT. Instead, firms are uncertain about the labor supply elasticity and concerned about spillovers from raising wages, both across workers and over time. A second worker-side RCT that randomly adjusts wages confirms wage increases are effective at increasing completed work, allowing us to quantify the production ramifications of firm search frictions. Our results help rationalize why many vacancies go unfilled.

*JEL Classifications:* M51, M52

*Keywords:* Hiring, wages, skills

**EXTREMELY PRELIMINARY. PRIMARY RCT YET TO BE CONDUCTED.**

---

\*Julia Gilman provided outstanding research assistance. We thank Charlie Brown for useful comments. Financial support from SSHRC is gratefully acknowledged.

<sup>†</sup>Harvard Business School and NBER

<sup>‡</sup>U. Toronto Rotman and NBER and CEPR

<sup>§</sup>Carnegie Mellon

The hiring process, in which workers and firms match, is a central economic topic. A large literature studies how workers search for jobs, often focusing on the worker’s problem over time, e.g., how do workers search as they near the date of expiration for their unemployment benefits.<sup>1</sup> However, firms also face time constraints. For example, a theater needs someone to sell tickets and snacks for a movie showing at a certain time, or a university needs someone to teach a class in the winter semester. How do firms adjust to fill their jobs over time, both in terms of wage and non-wage factors? Unfortunately, relative to what is known on worker search over time, little is known about how firms search for workers over time, consistent with a general dearth of knowledge in personnel economics on firm hiring strategies (Oyer & Schaefer, 2011; Bloom & Van Reenen, 2011).

The question of how firms search for workers over time is of particular importance during tight labor markets, such as the one occurring during much of the COVID-19 pandemic and its aftermath. In 2022, the ratio of vacancies to jobseekers reached unprecedented levels, and this was especially true in low-skill labor markets. Of course, the issue of labor shortages is hardly specific to the pandemic. In many occupations, from truckers and construction tradespeople, to nurses, physicians, and teachers, it is frequently argued that there are substantial labor shortages.

This paper sheds light on how firms adjust wage and non-wage factors using administrative data and two randomized controlled trials (RCTs) on a leading national staffing platform (hereafter, the “Platform”). On the Platform, firms hire workers for a large variety low-skill jobs, usually, jobs lasting one day. The setting is well suited for our research questions. First, we observe rich, high-frequency data on how firms adjust features and requirements of the jobs. For each job, we see all the adjustments which are made, including hourly wage adjustments, hours adjustments, skill requirements, and restrictions based on worker productivity. We observe the exact time at which worker is set to begin, so we can observe how adjustments to wage and non-wage features of jobs occur relative to the start of work. We are not aware of other data with these features. Second, the Platform is broadly representative of many low-skill jobs, for which labor shortages seem especially acute. In addition, since the Platform is close to a spot market, we can rule out many alternative explanations for results such as the existence of long-term contracts, insurance against income fluctuations, or wage negotiations. Third, collaborating with the Platform, we have

---

<sup>1</sup>A recent paper containing a rich review of the literature is DellaVigna *et al.* (2017).

the opportunity to analyze RCTs, which present firms with different opportunities to fill their jobs.

Our analysis is grounded in a simple conceptual framework. In a very simple model of a monopsonistic firm, we show that firms increase wages when labor markets are tighter or as it gets closer to the start of the job. However, the degree to which firms prefer to adjust wage factors depends on potential spillovers from increasing wages.

Using the universe of tens of thousands of workers and businesses on the Platform during 2017-2019, we document that firms appear hesitant to adjust wages. As time runs closer to the start of a job, firms do raise wages, but only by very limited amounts, namely, about \$0.05 per hour on average. This occurs across a wide range of jobs and many jobs go unfilled as a result.

As we document, firms' limited tendency to increase wages as time fills out appears part of a broader tendency of firms not to adjust wages. For example, wages also adjust little to (a) tightening labor markets, (b) filling positions at unpopular times of day, and (c) local wage differences. We view this general lack of wage adjustment as broadly consistent with a recent paper by Hazell *et al.* (2022), who show that large firms perform relatively little adjustment of wages to the conditions of a local labor market.

The spot market nature of the Platform means that a number of leading explanations based on long-term work arrangements cannot explain our findings. Another possible explanation is that firms adjust their wages at lower time frequencies. However, we find only limited evidence for that in our data. One possible class of explanations is that firms are inattentive or adjustment costs in responding rapidly to unfilled positions. A second possibility is that firms face important frictions, potentially psychological ones, that limit their ability to quickly adjust wages. To better test between these explanations, we turn to an RCT.

In the primary RCT described in Section 3, firms make binary decisions about hypothetical firm policies related to dynamic wage adjustment and other policies. However, the Platform is deeply interested in actually implementing these policies. Working with the Platform's legal and product teams, we collaborated on policies that were of strong interest to them, and that could potentially be implemented in the future. Thus, firm responses were incentivized based on the idea that policies could actually be implemented in the future. In the RCT, we find strong evidence that firms shy away from adjusting wages due to frictions, both across workers and over time. We find little evidence that firms do not adjust wages simply because they do not know about opportunities

to adjust wages or that it is too costly in terms of time to do so.

A final explanation for why firms perform very limited wage adjustment is that wage increases have limited benefits in our setting in terms of filling jobs.

Our paper relates to several literatures. First, it relates to work in labor and personnel economics on hiring. Models of job search focus on the process by which workers search for jobs. Much less is known about how firms search for workers, and especially about the detailed nature of the search, and how they trade off wage versus non-wage factors. Using highly granular data, our findings substantially expand our understanding of the process by which firms search and the frictions that they face.

Second, it relates to work in labor economics on wage rigidity. It has long been known that wages exhibit downward wage rigidity, with firms being hesitant to cut wages in bad times. Our results speak instead to the issue of upward wage rigidity and the frictions firms face. Given our finding that firms are concerned about intertemporal spillovers (i.e., the problem that workers will feel frustrated if wage yesterday was higher), our paper suggests that downward and upward wage rigidity are closely related. Firms may be hesitant to raise wages now if there are issues with lowering them later.

Third, it relates to work in industrial organization on dynamic pricing. Growing work in various industries from airlines to ticket pricing shows that firms engage in dynamic pricing, and that there are often significant profit returns from doing so. For example, many firms drop prices in the hours very close to a concert or flight. While we document that wages raise, the degree of rigidity we observe is much lower than for consumer markets.

Fourth, it relates to work on behavioral firms. A small but growing body of research suggests that not only do workers sometimes act in “behavioral” factors, but sometimes so do firms. In our case, our findings broadly support that firms are not necessarily making mistakes, but are optimizing in a constrained environment, where adjusting wages imposes a variety of costs on workers and the broader organization.

# 1 The Platform and the Data

## 2 Observational Analysis of Data on the Platform

**Firm reactions to hiring difficulties.** The above results show that firms raise wages only limited amounts prior to the start of jobs. While this seems to suggest that firms are hesitant to adjust wages to fill jobs, other explanations are also possible. One possibility is that firms do adjust wages to fill jobs, but face challenges in performing wage adjustments at daily frequencies, e.g., perhaps they adjust wages on a weekly timeframe. A second possibility is that firms wish to adjust wages, but do not know the going the wage.

To help shed light on such possibilities, we examine how firms adjust their wages and other factors when they fail to fill their jobs. For each establishment, we locate the first event in our data when they experience a zero fill rate. We estimate regressions of the form:

$$(1) \quad y_{it} = \sum_k \beta^k D_{it}^k + X_{it} \delta + \epsilon_{it}$$

where  $y_{it}$  is an outcome of firm  $i$  in week  $t$  such as firm wages;  $D_{it}^k = 1(t - \tau = k)$  are event time dummies, with  $\tau$  being the time at which the establishment experiences a hiring shock;  $X_{it}$  are controls, and  $\epsilon_{it}$  is an error. We restrict to the five weeks before and after the first negative hiring event, focusing on the first event as it is likely to contain the most information and to avoid complications with multiple shocks per establishment.

Figure 7 shows that firms do increase wages somewhat after a zero hiring shock, but the magnitude of adjustment is relatively small, with an average post-treatment increase of roughly 20 cents. Panel (a) of shows using the full sample of shocks, showing that firms appear to gradually raise their hourly wage after a zero hiring shock, increasing by roughly \$0.10 per hour in the week after the zero fill event, to roughly \$0.40 more per hour roughly five weeks after the event, and an average post-treatment effect of \$0.21 more per hour. A challenge with panel (a) is that the sample is highly unbalanced with many firms not hiring in week. Thus, panel (b) remakes the results requiring that an establishment post work in all five weeks after a zero fill event. The result is similar, having a raise of \$0.35 more per hour five weeks out, and an average post-treatment effect of \$0.20 more per hour.

One concern with these results is whether the event is endogenous, e.g., firm experience hiring shocks when they lower wages. We address this by instrumenting the presence of a zero hiring shock with whether the fill rate of other plants in the same three-digit zip code and week. As shown in the Appendix, we find ABC.

### **3 Primary RCT: Firm Demand for Dynamic Pricing**


### **4 Secondary RCT: The Impact of Wage Increases on Job Filling**

### **5 Conclusion**

## References

- BLOOM, NICHOLAS, & VAN REENEN, JOHN. 2011. Human Resource Management and Productivity. *Handbook of Labor Economics*, **1**, 1697–1767.
- DELLAVIGNA, STEFANO, LINDNER, ATTILA, REIZER, BALAZS, & SCHMIEDER, JOHANNES F. 2017. Reference-Dependent Job Search: Evidence from Hungary. *Quarterly Journal of Economics*, **132**(4), 1969–2018.
- HAZELL, JONATHON, PATTERSON, CHRISTINA, SARSONS, HEATHER, & TASKA, BLEDI. 2022. *National Wage Setting*. NBER Working Paper.
- OYER, PAUL, & SCHAEFER, SCOTT. 2011. Personnel Economics: Hiring and Incentives. *Handbook of Labor Economics*.


Figure 1 Hourly Wages Leading Up to Job Start



Notes: Point estimates with 95% confidence intervals in brackets, using job fixed effects and robust standard errors. The regression only includes observations from jobs that exist at least 5 days prior to job start (i.e., day 0). On each day, a job is only included in the sample if it is unfilled, defined as fewer slots filled than requested. Wages winsorized at .5 and 99.5 percentiles.




Figure 2 Total Wages Leading Up to Job Start



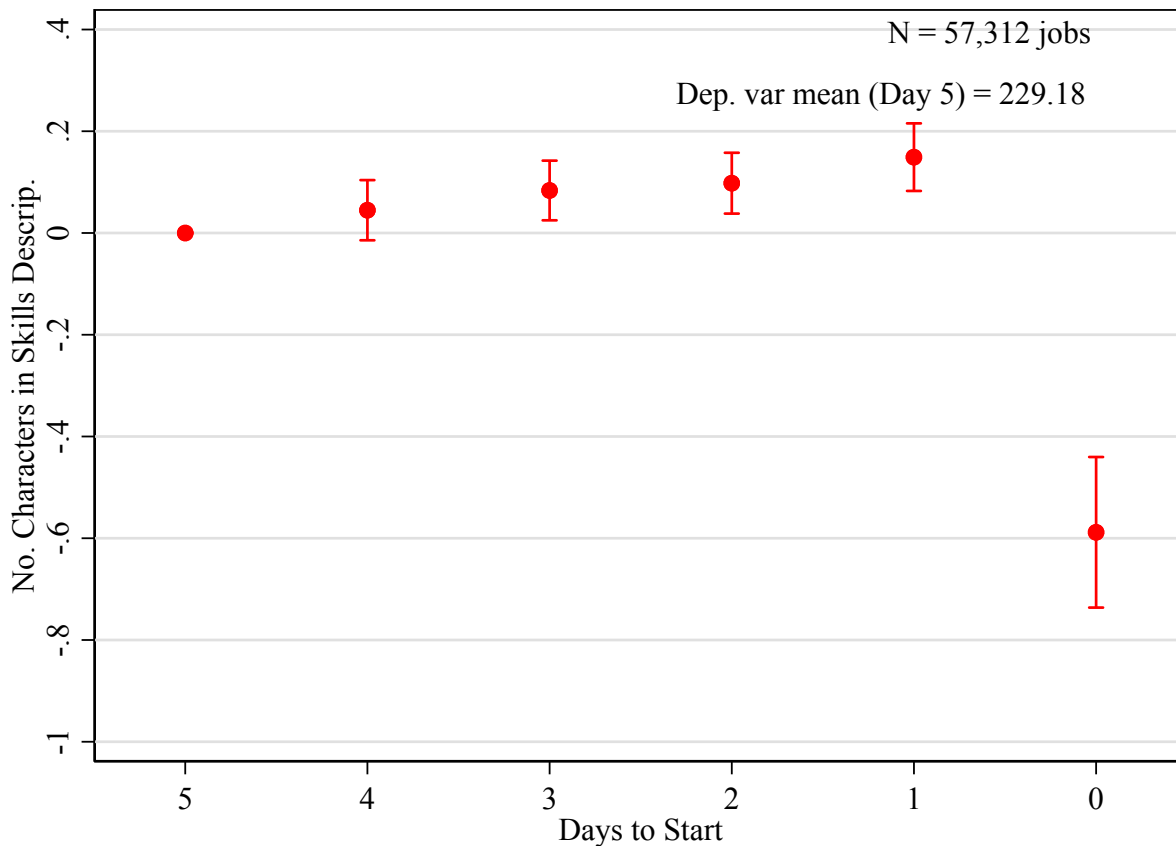
Notes: Point estimates with 95% confidence intervals in brackets, using job fixed effects and robust standard errors. The regression only includes observations from jobs that exist at least 5 days prior to job start (i.e., day 0). On each day, a job is only included in the sample if it is unfilled, defined as fewer slots filled than requested. Wages winsorized at .5 and 99.5 percentiles.

Figure 3 Hours Leading Up to Job Start




Notes: Point estimates with 95% confidence intervals in brackets, using job fixed effects and robust standard errors. The regression only includes observations from jobs that exist at least 5 days prior to job start (i.e., day 0). On each day, a job is only included in the sample if it is unfilled, defined as fewer slots filled than requested. The outcome, hours, measures the length in time of the job posting. 54% of outcome values are imputing using either the total wage compared to the hourly wage or the start and end time of the job posting.

Figure 4 Number of Characters in Skills Descriptions Leading Up to Job Start




Notes: Point estimates with 95% confidence intervals in brackets, using job fixed effects and robust standard errors. The regression only includes observations from jobs that exist at least 5 days prior to job start (i.e., day 0). On each day, a job is only included in the sample if it is unfilled, defined as fewer slots filled than requested. The outcome measures the number of characters in the description of skills needed for the job inputted by the firm.

Figure 5 Number of Items in Skills Descriptions



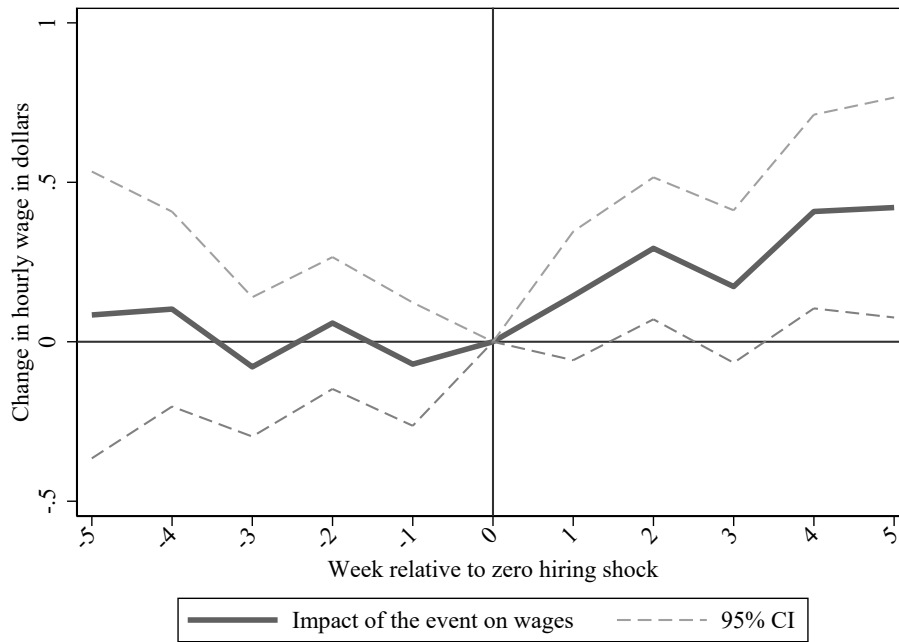
Notes: Point estimates with 95% confidence intervals in brackets, using job fixed effects and robust standard errors. Using jobs that exist at least 5 days prior to job start. The sample only includes observations from jobs that exist at least 5 days prior to job start. On each day, a job is only included in the sample if it is unfilled (i.e. fewer slots filled than requested). The outcome measures the number of listed items (“bullet points”) in the description of skills needed for the job inputted by the firm.

Figure 6 Badges: Conditional on an existing badge on the previous day



Notes: Point estimates with 95% confidence intervals in brackets, using job fixed effects and robust standard errors. Using jobs that exist at least 5 days prior to job start. The sample only includes observations from jobs that exist at least 5 days prior to job start. On each day, a job is only included in the sample if it is unfilled (i.e. fewer slots filled than requested). The outcome measures the total number of 5 available badges a job posting requires that workers have in order to accept the job. Moreover, an observation is only here included if, on the previous day, the job requires at least one badge, such that the employer may either add or remove a badge on the current day.

Figure 7 Firms' Adjustment to their Hourly Wage After Experiencing a Zero Fill Hiring Shock



(a) Use All Shocks

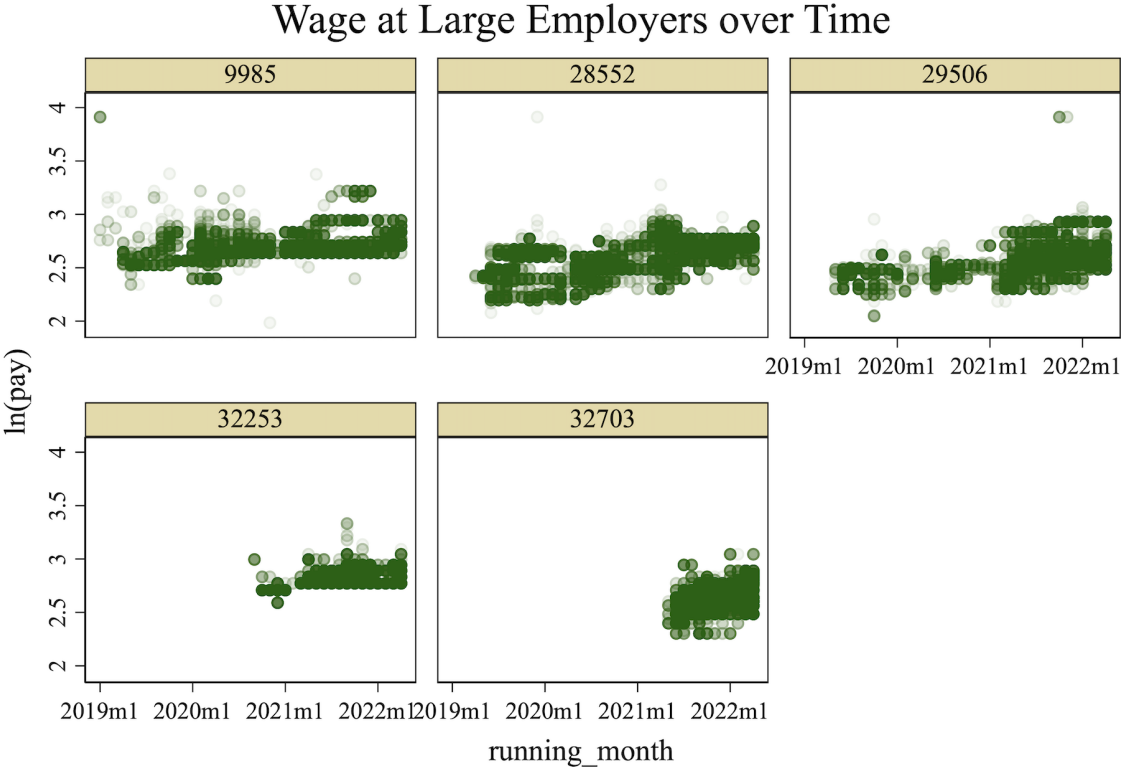


(b) Restrict to Shocks Where Firm Hires in All Five Weeks After the Shock

Notes: Standard errors clustered by establishment. Week zero refers to the week of a zero hiring shock. We take the first zero hiring shock for each establishment, and we exclude establishments with no zero hiring shocks. The sample for the event study is restricted to run from five weeks before a zero hiring event until five weeks after a zero hiring week. All regressions control for year-month dummies.

Web Appendix, “How Do Firms Search for Workers Over Time?”, by Cullen, Hoffman, and Koenig

Figure A1 Wage at Large Employers over Time



Graphs by CLIENT\_ID