THE PAYOFFS OF HIGHER PAY:

ELASTICITIES OF PRODUCTIVITY AND LABOR SUPPLY WITH RESPECT TO WAGES

Natalia Emanuel · Emma Harrington[1]

(Job Market Paper)

This version: October 30, 2020

Latest Version: Click here

Abstract

Firm wage-setting trades off the potential benefits of higher wages — including increased productivity, decreased turnover, and enhanced recruitment — against their direct costs. We estimate productivity and labor supply elasticities with respect to wages among warehouse and call-center workers in a Fortune 500 retailer. To identify these elasticities, we use rigidities in the firm's pay setting policies that create heterogeneity relative to a changing outside option, as well as discrete jumps when the firm recalibrates pay. We find evidence of labor market frictions that can give firms wage-setting power: we estimate moderately large, but finite, turnover elasticities (-3 to -4) and recruitment elasticities (3 to 4.5). In addition, we find productivity responses to higher pay in excess of \$1. By comparing warehouse workers' responses to higher wages both across workers and within the same worker, we find that over half of the turnover reductions and productivity increases arise from behavioral responses as opposed to compositional differences. Our results suggest historical pay increases are consistent with optimizing behavior. However, these aggregate patterns mask heterogeneity. For example, women's productivity responds more to wages than men's, while women's turnover is less responsive than men's, which can lead to occupational wage differences.

[1]Contact: Harvard University, 1805 Cambridge Street, Cambridge, MA 02138, emanuel@g.harvard.edu · eharrington@g.harvard.edu. We thank Claudia Goldin, Lawrence Katz, Nathan Hendren, Edward Glaeser, Jeffrey Liebman, Amanda Pallais, Lawrence Summers, and participants of the Public Finance and Labor Economics Workshop at Harvard for helpful comments. We appreciate input from Isaiah Andrews, Zoe Cullen, David Cutler, Gabriel Chodorow-Reich, Jerry Green, Jeffrey Miron, Matthew Rabin, Andrei Shleifer, and Elie Tamer. We are grateful to our colleagues Jenna Anders and Augustin Bergeron, as well as to Alyssa Bilinski, Valentin Bolotnyy, Justin Bloesch, Harris Eppsteiner, Benny Goldman, Omeed Maghzian, Dev Patel, Jonathan Roth, Gregor Schubert, Ben Sprung-Keyser, and Anna Stansbury. This project would not have been possible without the curiosity and commitment to research of our colleagues at the firms who shared data: Dave and Tommy, Lauren and Trevor. We are grateful for financial support from the National Science Foundation [Emanuel] and the Lab for Economic Applications and Policy. The findings and conclusions expressed are solely those of the authors and do not reflect the opinions or policy of the organizations that supported this work.

Firm wage-setting decisions must balance the benefits to the firm of higher pay — lower turnover, higher worker effort, and enhanced recruitment — against the direct costs of higher wages. Recent high-profile cases of large employers of low-wage workers (such as Costco) voluntarily implementing large wage increases and/or company minimum wages suggest that firms are aware of the potential gains from paying workers above their outside options.[2] The trade-offs firms face of higher net productivity against greater direct compensation costs from increasing wages are formalized in efficiency wage models (Yellen, 1984; Katz, 1986) and are a key component of company personnel policies. However, the causal impact of higher wages relative to worker outside options on firm performance has been difficult to assess given the endogeneity of wages to difficult-to-measure outside options. In this paper, we provide new evidence on what firms gain from higher pay. We estimate elasticities of turnover, productivity, and recruitment in the context of warehouse workers and customer service workers at a major online retailer, using sharp, discrete changes in wages or outside options. Our approach permits us to calculate the return to the firm of paying higher wages, inclusive of productivity effects.

In particular, we leverage rigidities in the firm's pay-setting policies to estimate these elasticities using three complementary empirical strategies. Nationally, the firm has sticky wages, which leads to exogenous variation in the value of wages relative to workers' local outside options. By comparing changes in the relative wage of workers in various cities over the course of the year to changes in the turnover in those cities, we are able to estimate the effect changes in relative wages on workers' behavior. Second, when the firm gets "unstuck" and recalibrates its wages, it changes its pay discretely, leading to plausibly exogenous variation in wages. We leverage this large, abrupt jump in pay to look at the difference in turnover and productivity before and after. Finally, to estimate elasticities of recruitment, we use the fact that the firm sets wages nationally for its remote customer service workers, regardless of location, which again creates heterogeneity in relative wages.

We estimate a turnover elasticity between 3 and 4 and a recruitment elastic between 3 and 4.5. While large relative to other estimates of labor supply elasticities, these elasticities are definitively finite, suggesting an upward sloping labor supply curve, consistent with firms having some monopsony power even in these labor markets that feature competitive pressure. While firms could use monopsony power to lower wages, the response of productivity to higher pay that we estimate is substantial (elasticity of 1.1-1.2), pointing toward a force that would push wages upward. Since these effects could arise from workers' behavioral responses to higher pay or to selection, we use data from a staffing agency to estimate how much of the decrease in turnover arises within the same worker when facing different wages for comparable work. We find that 80% of the turnover effects and suggestively 50% of firm satisfaction arises from

[2] Of course, public relations and goodwill gains could also play a role in the decision-making of large and visible employers who are concerned with potential regulatory actions and consumer responses.

workers' individual responses. Finally, we estimate gender-specific responses to higher pay to understand what our model suggests about the gender pay gap. We find that while women have lower labor supply elasticities than do men, women have much larger productivity responses, so that higher wages for women would be reasonable.

Our paper makes five contributions. First, we document the effect of higher pay on productivity for warehouse workers and customer service representatives, using objective productivity metrics of calls answered and boxes moved. Further, we estimate that the increase in productivity caused by raising wages fully pays for itself. This builds on the important literature on efficiency wages.[3] Our findings echo the analysis of Ford Motor Company where high wages reduced turnover rates and elicited greater effort from workers Raff and Summers (1987) and Cappelli and Chauvin (1991), who find that higher relative pay in a multi-plant firm reduced disciplinary infractions.[4] Moreover, we find that the gross productivity returns to higher pay are larger than one both when the firm voluntarily raises pay and when the firm keeps wages constant.

Second, we estimate turnover elasticities in two thick labor markets—warehousing and customer service—both of which are characterized by many workers, many firms employing workers to do a very similar job, and substantial churn of workers across different firms. Warehouses are often located close to cargo hubs, where many logistics firms all draw on the same pool of workers. The retailer's call-centers are also located in markets with many openings for customer service workers. Nonetheless, we find turnover elasticities between 3 and 4.5, suggesting that workers' labor supply to the firm is not perfectly elastic even in labor markets that would likely contain substantial competitive pressure. This finding contributes to mounting empirical evidence of wage-setting power of firms in a wide variety of contexts, from nurses (e.g., Sullivan, 1989; Staiger et al., 2010) and civil servants (Dal Bó et al., 2013) to online workers doing narrowly defined tasks (Dube et al., 2020) and school teachers(Ransom and Sims, 2010). Several papers further use linked employer-employee data to draw a connection between firm labor supply and workers' earnings(e.g., Bassier et al., 2020; Webber, 2015).

Third, we separately estimate recruitment elasticities and turnover elasticities. These elasticities are critical in the new monopsony literature, which argues that upward sloping labor supply curves give firms wage-setting power (Manning, 2003; Dube et al., 2020), and where the two elasticities are assumed to be roughly equivalent. This assumption is reasonable if workers joining one firm tend to be leaving another firm, for example. The labor supply elas-

[3] Under the efficiency wage hypothesis, employers may pay a premium above the market to give the worker an incentive to try to keep their job, to lower recruiting and turnover costs or to increase morale and effort (Shapiro and Stiglitz, 1984). If all employers raise wages, they employ fewer workers; thus, there is more labor supplied to the market than is demanded by firms so some workers who want jobs are unable to find them (Dickens et al., 1986).

[4] Krueger and Summers (1988), Dickens et al. (1986) and Orszag and Zoega (1996) explore whether intra-industry pay differentials can be attributed to efficiency wages.

ticity facing a firm is the combination of the job-to-job recruitment elasticity and the turnover elasticity (Manning, 2003). Since recruitment elasticities to an individual firm are difficult to estimate,[5] they are often assumed to be the same as the turnover elasticity. We find that relative pay is a significant determinant of workers' decisions but recruitment elasticities are definitively finite, suggesting limited information, heterogeneous preferences, or other frictions also shape recruitment. Further, we find that recruitment and separation elasticities are similar in magnitude in the aggregate, confirming a key assumption made throughout the literature.

Fourth, we estimate the extent to which the turnover effects we measure are due to sorting of better workers to higher-paying firms rather than the direct behavioral responses of workers to higher pay. We use data from staffing agency that places many warehouse workers in temporary jobs to assess how much of association between reduced quits and higher pay persists when we focus on the same worker placed in multiple jobs with different pay. We document that 80 percent of the turnover effect arises from behavioral responses when we look at the same individual's responses to higher and lower wages, consistent with efficiency wages operating on the worker effort margin. The remaining selection effects lead to negative spillovers onto other firms: when a high-paying firm is filling positions, other firms that are hiring for the same type of worker at that time end up with higher turnover rates. This evidence microfounds the literature that documents cross-firm wage elasticities since our estimates explain the need to raise wages in rival firms.[6]

Finally, by estimating gender-specific pay elasticities, we can shed light on how much responses to higher pay may explain gender pay gaps. Our estimates reveal that women's turnover is less responsive to pay then that of men in customer service, which would be consistent with an 8-cent pay gap. Our findings on turnover are in line with those of Ransom and Sims (2010) in the context of grocery store clerks and Hirsch et al. (2010) in the context of German workers. But our findings differ from the recent work by Caldwell and Oehlsen (2018) who find a small labor supply gap among Uber drivers, suggesting that gig work or male-dominated occupations may feature different dynamics than other sectors in terms of gendered labor supply elasticities. We contribute to this literature in two ways: (1) we leverage quasi-random changes in relative pay *within* a job rather than variation in pay *across* jobs to more cleanly identify differences in turnover elasticities and (2) we estimate gender-specific productivity responses to higher pay. These productivity responses are important because they suggest that in this context women should be paid *more* than men, which further complicates the puzzle of existing gender pay

[5]Recruitment elasticity estimates that do exist typically do not consider elasticities to specific firms and do not hold the attributes of the work fully constant (Katz and Krueger, 1991; **?**), with the notable exception of Dal Bó et al. (2013), who successfully randomize wages and measure what happens to recruitment in the Mexican civil service.

[6]Specifically, Staiger et al. (2010) finds that when Veteran's Affairs hospitals increase their wages for nurses, nearby hospitals do as well. Derenoncourt et al. (2020) examine the effect on local wages at other firm of wage raises such as Walmart's increases in pay from $9 to $11 in 2015-2018. They find a cross-employer wage elasticity of 0.25.

gaps.

Our paper also illustrates how firms' wage rigidities can be used to estimate the effects of relative pay on turnover and performance. We introduce a novel instrument that leverages the fact that nationally sticky wages lead to greater real wage depreciation in places with faster aggregate wage growth. Given the ubiquity of nominal wage rigidity, these empirical designs could likely be applied in other settings to analyze how relative pay affects outcomes of interest.

In the process of analyzing the effects of higher pay, our work highlights wage-setting strategies among major firms that may deviate from the optimum in the short run. One pay setting strategy we leverage is sticky pay, which may be optimal if there are large adjustment costs. We also highlight uniform wage-setting, a tendency which parallels firms' uniform pricing strategies that leave money on the table given variation in local demand and competition (DellaVigna and Gentzkow, 2019). While not the focus of our paper, our findings suggest that uniform wage-setting policies are unlikely to be optimal long-run strategies, given variation in local labor supply and competitive pressure from other employers. That said, if the costs of tailoring pay to each geography are too large or the firm is concerned about inequities on remote-teams where each worker is located in a different area, uniform wage setting may still be worth the costs we estimate.

The rest of the paper is organized as follows. Section 1 presents a conceptual framework for structuring our empirical investigation and Section 2 introduces our datasets. Sections 3 - 5 document our findings on the elasticities of turnover, productivity, and recruitment with respect to pay. In Section 6 we conduct a cost-benefit calculation to estimate the returns to our firm of marginally higher pay. In Section 7, we explore the degree to which selection versus behavioral responses contribute to our results and, accordingly, whether higher pay at one firm has negative spillovers on other local firms. Section 8 explores heterogeneity in responsiveness to pay by gender and its implications for the gender pay gap. We conclude in Section 9.

1 Conceptual Framework

We lay out a simple conceptual framework to illustrate how the the three parameters that we estimate — the turnover elasticity, the recruitment elasticity, and the productivity response to pay — matter for wages in a partial-equilibrium setting where firms are able to optimally set pay. This framework adds productivity to the key Burdett-Mortensen-Manning model linking the labor supply elasticity to wages (Burdett and Mortensen, 1998; Manning, 2003), by allowing the output of the worker to depend directly on her wage, as in efficiency wage models (Yellen, 1984).

We consider a single firm in a market with several other firms, and a large number of homogeneous workers.

The firm sets wages to maximize profits: the marginal product of the worker p, less her wages w, times the number of workers N:

$$\max_{x}(p(w_t) - w_t)N_t$$

The number of workers at a firm in period t will reflect the number of new recruits, R, plus the share of workers in the previous period who do not leave the firm (turn over), $(1 - T)$, multiplied by the number of workers last period:

$$N_t = [1 - T(w_t)]N_{t-1} + R(w_t),$$

Note that we treat $p(w)$ and $R(w)$ as separable, in effect modeling $p(w)$ as a measure of average workforce productivity. As such, $p(w)$ reflects two components: both the composition of the workforce and the effort exerted by the workers employed. We attempt to decompose the two in Section 7.

In steady state, the number of recruits must balance the number of quits at the constant wage, giving us:

$$R(w) = T(w)N \implies N = \frac{R(w)}{T(w)}$$

We can solve for the firm's optimal steady-state wage:

$$\max_{w}(p(w) - w)\frac{R(w)}{T(w)}.$$

We then have:

$$(p'(w) - 1)\frac{R(w)}{T(w)} + (p(w) - w)\left(\frac{R'(w)}{T(w)} - \frac{R(w)}{T(w)^2}T'(w)\right) = 0$$

Dividing through by $\frac{R(w)}{T(w)}$ allows us to isolate the elasticities of recruitment and turnover with respect to wages, where $\epsilon_{R,w} = \frac{R'(w)w}{R(w)}$ and $\epsilon_{T,w} = \frac{T'(w)w}{T(w)}$, noting that $\epsilon_{T,w}$ will be a negative number since we are considering turnover — namely separations from the firm — which we expect to decrease with an increase in wages.

$$p'(w) - 1 + \frac{p(w) - w}{w}(\epsilon_{R,w} - \epsilon_{T,W}) = 0$$

$$\frac{p(w) - w}{w} = \frac{1 - p'(w)}{\epsilon_{R,w} - \epsilon_{T,W}}$$

This equation captures the markdown, the percentage below marginal product that the worker is paid. The markdown is decreasing in the responsiveness of productivity to the wage and the elasticities of labor supply to the firm.

We can further rearrange to arrive at an expression for the optimal wage:

$$w^* = \frac{p(w)}{1 + \frac{1-p'(w)}{\epsilon_{R,w} - \epsilon_{T,W}}}$$

This expression captures predictions made by both the new monopsony literature and the efficiency wage literature. It shows that as elasticities of turnover and recruitment grow in magnitude, so do wages. Intuitively, if workers are unwilling to come except at high wages, or are willing to leave at lower wages, wages must be driven upward. The expression also shows that if productivity is increasing in wages, then wages will also be larger.

2 Data

We use data from two large firms: a Fortune 500 retailer and a leading staffing agency. Though both organizations function throughout the United States and abroad; we focus on their U.S. operations.

Online Retailer Data. The first data source is a major online retailer, which employed 8,597 warehouse workers and 4,551 customer service representatives between 2018 and 2019.

We use the human resources records from the retailer, which detail each active worker's job title, level of employment (e.g., entry-level, associate, senior), pay rate, and location. For warehouse workers, we can observe the shifts they worked.

In addition to HR records, we have two datasets to measure on-the-job productivity of workers at the retailer. First, we observe productivity of each of the retailer's warehouses on each week.[7] The key productivity metrics are boxes moved per hour and boxes moved per moving hour, (total hours worked excluding hours spent eating lunch or attending team meetings).[8]

We look predominantly at a single warehouse that featured a large, plausibly exogenous pay jump, calling this the "treated" warehouse. Summary statistics for this warehouse and other comparison warehouses are shown in the first three columns of Table 1, Panel A. In the three months before the pay jump, fully 13.4 percent of workers in the treated warehouse left in a given month. The treated warehouse tended to move 4.92 boxes per hour before the pay jump.

Second, we are able to directly observe the productivity of each customer service representative on each day they handle calls. These metrics include the total number of calls each representative answered and the average customer satisfaction reviews that day. Customer service representatives handle incoming calls from customers, potentially inquiring about a delivery, a return, or damaged product. Since these representatives do not make outgoing calls or handle

[7] The retailer does not track productivity of individual warehouse workers.

[8] At the time of data extraction, the retailer did not track data on damages or petty theft. This suggests that while these metrics may be important in principle, they are not first-order concerns for the retailer.

incoming sales requests, the metrics we observe represent the key measures of productivity that the firm cares about.

We use two subsets of customer service representatives in our analyses. To investigate recruitment we use the subsample of 593 remote workers, who are drawn from all over the US (column five in Table 1, Panel A). These workers are paid, on average, \$14.35/hour, which is 10 cents below the average entry rate for customer service agents in their metropolitan statistical areas (MSAs). On average, about 6.6 percent leave per month. They handled about 25 calls per day, which entails addressing a call in less than 19 minutes.

We also use the sample of 3,061 workers whose wages are sticky, which creates heterogeneity in their relative pay as their outside option changes. These workers include on-site workers, who are paid more, as well as remote workers; in the sample, the average pay is \$16.02/hour, well above the \$13.52 entry pay in their MSAs. There is lower turnover in this sample: 4.3 percent of workers leave in a given month. Daily call volumes are comparable to the remote sample.

We supplement these administrative records with data from Economics Modeling Specialists, International (Emsi) to find measures of the local pay for customer service representatives. Emsi compiles data from government sources including the Bureau of Labor Statistics and the Census, online profiles and resumes, online job postings and compensation data. Many companies, including the retailer, use Emsi's granular occupation- and labor-market-specific data on wages and labor supply to guide their decisions. While our own checks of Emsi's data against Bureau of Labor Statistics records, as well as the Quarterly Census of Employment and Wages (QCEW) and American Community Survey (ACS) suggest it is highly accurate, the widespread use of Emsi by companies suggests that their metrics reflect the local outside option as understood by firms.

Staffing Agency Data. Our second source of data is the segment of a large staffing agency that provides temporary staffing for production and warehouse companies, which placed workers in over 222,000 warehouse jobs between 2016 and 2018. Data from the staffing agency includes all of the assignments a worker was placed in through the staffing agency. For each assignment, we observe the pay rate, the firm that hired them, the reason that the temporary assignment concluded (e.g., the work was over, the worker quit, they were fired for poor performance, etc.), as well as the rating given by the manager at the firm ("Excellent," "Good," "Fair," "Poor"). Appendix 11 includes additional information about how individuals are placed in jobs at the staffing agency.

On average, temporary warehouse jobs through the staffing agency last 3.4 months, with an hourly pay of \$11.74/hour (see Panel B of Table 1). Only 44 percent of these jobs are completed, with fully 31 percent of people quitting and 27 percent having a bad ending, which includes being fired for attendance or performance problems or receiving a "Poor" or "Fair" evaluation from the manager at the hiring firm.[9] Only 13 percent of workers receive an "Excellent" evaluation from the hiring firm's managers.

One client of the staffing agency, a shipping company, regularly hires many warehouse workers to load and move boxes. This shipper employs temporary help throughout the entire United States, always employing workers at the same wage — $17/hour — regardless of location. In the three years that our data cover, the shipper hires 5,701 workers for 6,664 positions. On average, these temporary assignments last only 30 days, with 83 percent of the assignments completed. Moreover, 20 percent of workers receive an "Excellent" review from their manager.

Because we observe the same worker in several jobs, the staffing agency data offer a valuable opportunity to decompose the effect of pay on retention into selection of better workers versus incentives within the same worker. Further because we see many firms hiring for the same jobs in this dataset, we can estimate the spillovers of one firm's wage setting on other firms' turnover.

3 Turnover Elasticity

According to both estimates provided by the retailer and analysis of both warehouse and call-center data, turnover is costly, even for workers in jobs that are relatively routine and do not require an advanced degree. Internal estimates from the retailer suggest that training a new warehouse worker costs $1849 or 5.5 percent of the average worker's annual income.[10] Our estimates of the cost associated with a new customer service worker amount to $2990.[11] Moreover, objective metrics of productivity decrease when firms face turnover: fewer boxes are moved in warehouses and new customer service representatives answer fewer calls.

In weeks when workers leave a warehouse, the productivity in the warehouse decreases by 8 percent (0.75 fewer boxes per moving-hour off of an average of 9.14). Diminished productivity lasts three weeks. On average, each warehouse loses 2.8 workers per week.

It takes a new customer service representative about 6 months to reach the call volume of the average customer service representative who is answering calls on the same day within the same time-zone. As illustrated in Figure B.1, new representatives — who have just finished their 3 weeks of formal training — answer nearly 3 fewer calls per day, the equivalent of work-

[9] Note that quits and bad endings are not mutually exclusive categories. One could quit and also receive a poor evaluation, for example.

[10] Internal estimates suggest that training costs $689, reduced productivity and the associated overtime cost $860, and other costs including advertising, background checks, employee badges cost $300. This estimate does not include costs of recruiting and interviewing new candidates.

[11] These estimates include the decreased productivity over the course of training and are calculated based on the lower observed productivity times the average price the firm pays per call ($4.60) as well as the use of trainers' time. It also includes $300 in costs for advertising, background checks, and employee badges.

ing one fewer hour per day for the firm.[12] This pattern persists when we consider a balanced panel of representatives who stay at the retailer for at least 6 months (in the dotted line), suggesting that selection alone is not driving the observed trajectory. Given the trajectory of learning, a higher rate of churn means that at any given time more workers will be new to the firm and have developed less skill in answering calls. This dynamic also suggests that retention of senior customer service representatives is more valuable than retention of junior ones because they will walk away with more human capital accumulated in the firm.

3.1 Turnover Elasticity Estimates

We explore whether higher absolute and relative pay reduces turnover. An ideal experiment would randomize wages, allowing one estimate the causal relationship between turnover and wages:

$$\text{Turnover}_{it} = \beta_0 + \beta_{\$}\$_{it} + \epsilon_{it} \tag{1}$$

where $\beta_{\$}$ is the coefficient of interest, capturing the effect of hourly wages, $\$_{it}$ on turnover. In the absence of this experiment, we rely on two natural experiments that arise from firms' wage setting practices. At the retailer we study, wages are often sticky. This leads to two types of variation that we exploit in our analyses. First, when pay is ultimately changed, it is often done in a large, discontinuous manner. This is the case in one of the retailer's warehouses. The ensuring large, abrupt change allows us to compare worker performance in the warehouse when pay is lower to when pay is higher.

Second, when pay remains constant, changes in prevailing wages in workers' local areas changes their outside options, and thus the relative values of their wage at the retailer. We leverage the fact that pay remains constant from 2018 onward for all customer service agents to estimate how turnover varies with the changes in relative wage in various different metropolitan statistical areas (MSAs).

Turnover in a Warehouse. We use a one-time pay-jump in a single warehouse to investigate the effect of higher pay on turnover. In late July 2019, average pay was \$16.20/hour in this warehouse. One week later, the firm had increased the average pay to \$17.39/hour and by the first of September, it was solidly at \$18.00/hour, an 11% increase in pay over the course of a month. At the same time, pay remained essentially flat at other warehouses owned by the retailer. Figure 1 depicts the pay bump at the treated warehouse along with relative pay constancy at other warehouses at the retailer.

The pay bump in question arose out of long-standing concerns about high turnover at this warehouse in particular according to the Field Director at the retailer. Indeed, in the quarter before the pay change, turnover at the treated location was nearly twice as high as in other

[12]Note that in the first six months, all service workers are given easy calls (e.g., change of address). More complicated calls are reserved for senior representatives.

warehouses. As Slichter (1919) observes, high turnover is often cause for raising pay. The Field Director presumed that turnover was higher at this warehouse than at other warehouses because (a) it is in a highly competitive local labor market where other firms' warehouses are located in very close proximity, and (b) the work can be especially grueling given that this warehouse handles larger parcels than other nearby retailer warehouses (e.g., refrigerators or sofas rather than tea towels or books). The Field Director further confirmed that the nature of the work did not change around the pay jump and that it did not coincide with consumer holidays that could have affected work intensity. Correspondingly, there is no significant change in the demographics of those working at the warehouse (see Appendix Table A.1). Thus the treated warehouse differs in important ways from other warehouses but the timing of the pay jump is near random, and we do not see contemporaneous changes in the warehouse.

Table 1 describes the treated warehouse as compared to other warehouses in the quarter before the pay jump. The bulk of warehouse workers are men in their mid-30s working full time. On average, they have been with the firm for 10 months. Of the people working during the quarter before the pay jump, fully 63 percent of those in the treated warehouse and 50% of those in the other warehouses will ever leave the firm.

We compare the turnover in the treated warehouse before and after the pay change in an interrupted time series design. We scale our results so that they reflect the change in turnover that would arise from a single dollar's change in hourly pay. We use a two stage least squares approach. Our second stage is

$$\text{Turnover}_{i,t} = \alpha_1 + \beta_{\$}\hat{\$}_{i,t} + \epsilon_{i,t} \tag{2}$$

and our first stage predicts wages based on being before or after the pay jump: $\$_{i,t} = \alpha_0 + \delta \mathbb{1}_{\text{Post}} + \nu_{i,t}$, where $\mathbb{1}_{\text{Post}}$ is an indicator for whether the observed day occurs after the pay change. $\beta_{\$}$ is our parameter of interest. Because our data includes daily observations of each worker, but the warehouse may be subject to shocks in any given week, we use two-way clustered standard errors, clustering at both the week and employee level.[13]

The bump in pay was just a rightward shift in the whole distribution of pay and thus did not affect relative pay within the warehouse or workers' dynamic incentives to strive for promotions.[14] Figure B.2 shows the distribution of wages in the week before the first pay change and the pay one month later. The standard deviation in pay beforehand is 1.18 and afterward is 1.21. Since this pay change occurred throughout the entry level workforce, but not the managerial workforce, one might worry that the pay differential between these two rungs was compressed. However, to date, no one has been promoted between those two levels of the warehouse, so we

[13]While we use an interrupted time series design that includes only the treated warehouse, we have also run this as a difference-in-differences approach, and include those results in Appendix Table **??**.

[14]In contrast, increases in the minimum wage often compress the wage distribution of firms with low-wage workers, potentially tempering the workers incentives to climb the ranks of the firm.

don't believe the constant managerial pay affected worker incentives.

Table 2, Panel A, which displays the results of estimating equation above, shows that increased pay decreases worker turnover. In the three months before the pay increase, out of every 100 workers in the warehouse, on average 13.4 would be leave per month – a monthly retention rate of 86.6 percent. Paying an additional $1/hour decreases turnover by 2.5 individuals – a decrease in attrition of 18.7 percent, and an increase in retention of 2.8 percent. Since our point estimate captures the effect of a $1/hour increase off of $16.20/hour, our point estimate reflects an elasticity of turnover of 3.03. The overall effect on turnover is driven by voluntary quits, which decrease by 21.3 percent — 2.2 fewer quits relative to the base of 10.66 quits out of every 100 workers. There is no effect on being fired for performance.

We present results from a bandwidth of 3 months on either side of the pay jump. Table A.2 show the effects across one-, two- and three-month windows. We do not extend the window beyond 3 months after the pay jump because subsequent months include the holiday shipping season, which has its own impacts on warehouses independent of the late-summer pay bump (e.g., local demand shocks).

We test our results with both a placebo and a permutation test. First, we perform the same analysis on all other in-state warehouses. Since two of the three of these warehouses are within a 13-minute drive of the treated warehouse, if there were a shock to the local labor market for warehouse workers that were driving the decreased turnover, one would expect to see it decrease turnover in these warehouses as well. However, as Table A.3 shows, there is no decrease in turnover in other in-state warehouses.[15]

Our second test looks within the treated warehouse, asking whether similarly large decreases in turnover have been seen at other time periods. To do this, we place the date of treatment at every other week in 2019 and estimate the effect size over a three-month bandwidth. We do not extend into 2018 because the holiday period is an unusual time that may be subject to other treatments (e.g., local demand shocks). We require that the entirety of our artificial treatment window not overlap with the true post-treatment window to avoid biasing the results. Figure B.3 shows the results, in which two of the thirty permutations (6.6 percent) are lower than our treatment. A one-sided test is most relevant here because our prior is that turnover should decrease in this context, not increase.

Turnover among Customer Service Representatives. Our second context looks at customer service representatives at the same retailer. We use the fact that the retailer has maintained sticky wages over time; the retailer has not adjusted its entry-level wages for remote or on-site representatives since at least 2018, when our administrative data begins. The stickiness of the

[15]For the placebo analysis, we do not scale by the size of the pay jump since the other warehouses do not feature a pay jump. We run the reduced-form regression.

retailer's pay contrasts sharply with the changing nature of representatives' outside options: local pay increases over this period among other firms, and does so more steeply in some MSAs than others. Where pay in local customer service jobs rose faster, the retailer's sticky pay depreciated more compared to the outside option. For example, in Tampa, FL, entry-level wages for customer service representatives rose considerably between 2018 and 2019, whereas in Sarasota, FL wages barely budged. We can consequently evaluate how the change in relative pay translates into a change in productivity among the representatives drawn from MSAs with faster and slower wage growth. This strategy allows us to difference away any fixed disparities in productivity across MSAs, while accounting for general trends within the retailer.

Particularly, we consider the first-difference specification, in which we relate the change in wages in an MSA from 2018 to 2019, $\Delta\$_{MSA,'18\rightarrow'19}$ to the change in turnover in the same location during the same time period:

$$\Delta\text{Turnover}_{MSA,'18\rightarrow'19} = \delta_0 + \delta_{\$}\Delta\$_{MSA,'18\rightarrow'19} + \zeta_{MSA}. \tag{3}$$

Since the retailer's pay is sticky between 2018 and 2019, the change in its relative pay is entirely driven by the growth (or stagnation) of the outside options in the MSA. To fully leverage the daily nature of our data and account for fluctuations in consumer demand within a given year, we focus on the analogous individual-level analysis, which allows us to include more granular controls, particularly date-by-time-zone fixed effects.[16] While this approach utilizes individual data, it does not limit the changes to within an individual — thus, these estimates will reflect the changing selection of the retailer's representatives as well as the changing incentives they face. This individual-level approach yields nearly identical point estimates as the collapsed analysis but smaller standard errors, since it absorbs daily fluctuations in call volume.

Our coefficient of interest is $\delta_{\$}$, which reflects the relationship between a \$1/hr change in relative pay and the parallel change in the MSA's turnover between 2018 and 2019. For $\delta_{\$}$ to be an unbiased estimate of the effect of relative pay, other MSA-level changes that would affect productivity must be orthogonal to changes in the relative wage. In particular, changes in the pool of customer service representatives must be orthogonal to changes in the wage. While this is still a strong assumption, the short time-frame of our analysis makes it a credible one: over the span of a single year, it seems more plausible that fluctuations in the demand for customer service representatives would drive changes in wages than would changes in the supply, in terms of either quantity or quality.

[16] We drop the 3.75% of customer service representatives (1.82% of days) with missing wage information. We further drop the 1.47% of representatives (1.55% of days) who are missing information on the local outside option for customer-service representatives — either because their address is missing or because too few customer service representatives work in the MSA for Emsi to construct an outside option. We exclude representatives in the 3 physical call-centers constructed in 2018 and 2019 — 21.6% of representatives (982 of 4551). We also exclude 1424 representatives (37.3%) hired in 2020, since our outside option information from Emsi is only available for 2018 and 2019. Finally, we exclude 170 representatives (7.1%) in MSAs with hires in only one of the two years.

As shown in Table 2 Panel B, in places where the retailer's pay lost more ground to the outside option, monthly turnover increased more precipitously. Each $1/hr loss in relative pay is associated with a 28 percent increase in monthly turnover — 1.2 percentage points off an average of 4.3 percent, reflecting an elasticity of 4.5 (see Appendix Table **??** for the MSA-level analysis).

We find that the reduction in turnover stems from both a reduction in quits — worker-initiated departures from the retailer, which are not due to family emergencies or geographic moves — and a reduction in fires for poor performance. The final two columns of Table 2 Panel B suggest that fires are especially sensitive to relative pay. The effect on fires is consistent with managers' expectations for workers not fully adjusting to diminishing relative pay: this may be especially likely in contexts where the nominal pay at the firm does not change and instead the firm's pay only changes in relative terms. This contrasts with the first case study where the retailer actively changed its own pay and we see no changes in fires, which is consistent with no change in the alignment between performance and expectations.

Turnover among Temporary Warehouse Workers. We turn to the shipper who hires warehouse workers at $17/hour regardless of where the job is located to investigate how relative pay affects turnover in temporary warehouse work.[17]

We consider the relationship between relative pay and the outcomes of interest:

$$\text{Turnover}_{ijt} = \beta(17 - \overline{\$}^{cz}_{ijt}) + \gamma D + \mu_t + \epsilon_{ijt}$$

where $\overline{\$}^{cz}_{ijt}$ reflects the average hourly pay rate for other warehouse jobs in the season and commuting zone in the staffing agency.D is a vector of expected duration variables up to a quartic in case workers are less likely to complete a longer job. We calculate the expected duration based on the duration that other such jobs at that firm tend to last, which reflects the information that recruiters would be able to give to potential hires. We include season fixed effects μ_t to address the fact that work and work availability may differ season by season in warehouses. We cluster our standard errors at the commuting-zone level in case commuting-zone shocks to the labor market affect workers' on the job performance.[18]

Where the $17/hour wage represents a greater premium over the local outside options, com-

[17]Notably, the shipper is not hiring the warehouse workers to test them out for a permanent position: of the thousands of individuals hired by this firm as warehouse workers, only 16 are offered a permanent position. As such, the possibility of individuals exerting more effort with an eye toward a permanent offer is effectively shut down.

[18]We construct the sample by limiting to season-commuting zone pairs that have more than 10 assignments from the shipper during peak seasons where the shipper hires in more than one commuting zone. We further eliminate the 267 assignments (3.2% of the 8,477 assignments) at the shipper that are hired at a different rate, since we believe these are different jobs. Of the 8,477 temporary assignments that the shipper secures through the staffing agency, 75% are retained in our sample. To construct the outside option, we include all other warehouse jobs begun in the same season and in the same commuting zone filled through the staffing agency. The comparison between the jobs at the shipper and the outside options can be seen in Panel B of Table 1.

pletion rates of the temporary job are higher. For each $1/hr increase in relative pay, workers are 2.0 percent more likely to complete the job off a base of 83 percent completion and 9.0 percent less likely to quit, off a base of 5.9 percent (see Panel C of Table 2).

4 Productivity Response to Higher Pay

The efficiency wage hypothesis suggests that higher pay may induce greater on-the-job performance due either due to fear of losing a well-paying job or the morale-boosting effects of higher pay. We use the productivity metrics used by the retailer to assess how boxes moved per hour and daily call volumes react to higher absolute and relative pay. We use the same empirical approaches as we did when estimating the effects on turnover (see Section 3): in the warehouse context, we exploit a large, quasi-random increase in wages; among customer service workers pay is sticky, creating heterogeneity in wages at the reltailer compared to the changing outside option.

We find that in both the retailer's warehouse and among their on-site customer service agents, productivity increases when pay, or relative pay, increases. In the warehouse, when pay increases the number of boxes moved per hour by 7 percent (0.325/4.92 boxes per hour), reflecting an elasticity of 1.2. Among customer service representatives, paying $1/hour more than the local outside option increases calls taken per day by 7 percent, reflecting an elasticity of 1.12.

Warehouse Productivity. Using the same pay jump used to estimate the effect of pay on turnover in the warehouse context, we estimate the effects of pay on the warehouse's productivity. Three metrics capture warehouse-level productivity: boxes moved per person-hour; boxes moved per *moving* hour, which removes from the denominator the time spent on non-moving activities like morning meetings or lunch; and the ratio of moving hours to total hours. We might expect the ratio of moving hours to total hours to decrease if the team works more seamlessly.

In the three months before the pay change, the treated warehouse moved an average of 4.9 boxes per hour, or 7.7 boxes per moving hour. The time-series of boxes moved per hour is shown in Figure 2.

As shown in Table 3, Panel A, in the three months following the pay jump at the warehouse, boxes moved per hour increased by 0.328 off a base of 4.92 boxes moved per hour, an increase in productivity of 7 percent. This corresponds to an elasticity of 1.2. Our metric of boxes moved per *moving* hour is 0.316, an increase of 4 percent. Finally, we find an increase of 0.018 in the ratio of moving to total hours, which corresponds to an increase of 8.6 minutes of moving per person per day. This increase in productivity could come from a number of sources: it could arise from attracting and retaining more productive workers, from workers exerting more effort, or from workers collaborating more seamlessly in light of reduced turnover. We discuss mechanisms more in Section 7.

To contextualize this increase in a more generalizable way we can translate the increase in productivity into dollar terms. In the quarter before the pay jump, the retailer paid \$3.29 per box moved. An increase of 0.336 boxes per hour thus represents an hourly savings of \$1.10 for the retailer (see Section 6 for greater detail).

We again test our results with both a placebo and a permutation test. The placebo test examines whether the same increase in productivity may be found at other warehouses. In this instance, we compare to the "twin" warehouses that handle the same size parcel as the treated warehouse.[19] Since warehouses handling the same type of product are most likely to be hit by similar demand shocks and have similar interpretation of their units of productivity, considering the twin warehouses is most suitable. As Table A.4 shows, there is no increase in productivity in the twin warehouses.[20]

Our second test, looks within the treated warehouse, asking whether similarly large increases in productivity have been seen at other times. As with turnover, we assign the other weeks in 2019 to be the week of treatment and run our normal analysis. Figure *[B.4]* shows the results.

Customer Service Productivity. We likewise explore whether higher relative pay is associated with greater number of calls handled by customer service representatives. As in Section 3, we use the retailer's sticky wages alongside changes in the local pay for customer service representatives as in Equation 3 to assess the value of an additional dollar in relative pay to reach these estimates. One advantage of these data is that they are extremely granular: the data track each person's daily calls.

We find that each \$1/hr increase in relative pay is associated with a 7.5% increase in call volume, 1.9 additional calls per day off of a based of 26 (see Table 3, Panel B). Intuitively, in MSAs where the retailer's sticky pay depreciated more substantially relative to the representatives' rising outside options, daily call volume fell between 2018 and 2019 compared to what would be expected. To contextualize this figure, the average customer service call costs the firm \$4.60, so an increase of 1.9 calls per day saves the firm \$8.74/day/worker.

Higher relative pay has limited but positive impacts on customer satisfaction, as shown in Table 3, Panel B, Column 2). This is reassuring to the extent that higher call volumes are not coming at the expense of less satisfactory customer experiences. However, the high rate of five-star evaluations and relatively little variation suggest that this metric of performance may not be particularly telling. By contrast, there is no statistically significant change in the share of absences that are unapproved by a manager in advance and thus difficult for the retailer to respond to.

[19]The treated warehouse and its two twin warehouses handle large parcels the size of refrigerators or sofas. Other warehouses handle parcels the size of toasters or tea towels.

[20]As with turnover, for the placebo analysis, we do not scale by the size of the pay jump since the other warehouses do not feature a pay jump. We run the reduced-form regression.

Notably, relative pay seems to have limited impact on hours worked, total absent hours, and overtime hours, as detailed in Table A.5. Thus, such effects do not complicate the interpretation of our key metrics. It is unsurprising that relative pay does not appreciably move the needle on hours worked because relative pay does not necessarily relate to the purchasing power of the earnings of a marginal hour, which is typically the key consideration in extensive-margin labor supply choices. While one could tell stories where relative pay would still affect representatives' scheduling decisions — e.g. because representatives were balancing multiple jobs or balancing job search against hours worked — it is less obvious that relative pay should impact intensive-margin choice of how much labor to supply at one's chosen firm than that it should impact the extensive-margin choice of where to work.

5 Recruitment Elasticity

Higher pay may be effective in recruiting more people and more talented people to a firm. An assumption often made in the new monopsony literature is that elasticities of recruitment and turnover are equal in magnitude (Manning, 2003). The notion that recruitment and turnover elasticities might be similarly sized is motivated by the idea that one firm losing a worker is balanced by another firm gaining a worker. In this model of job-to-job moves, recruitment and turnover elasticities are two sides of the same coin. This need not be true if workers also transition in and out of non-employment. For a specific occupation, it may also be violated if higher wages can more effectively retain workers in that occupation than recruit workers into the sector or vis-versa. Our estimates of the effect of higher pay on number of workers recruited allows us to assess this assumption.

To test the effect of higher pay on recruiting more people, we use the fact that both the retailer and employs individuals at the same wage, regardless of their location. This wage-setting strategy creates variation in the advertised wages relative to the local outside option. We test the effect of higher relative wages on the number of people recruited to the retailer, and the quality of the workers through the staffing agency. We find that when the retailer's advertised wages are $1/hour higher than the local outside option, they recruit 23 to 30 percent more employees in the MSA, reflecting a recruitment elasticity between 3.2 and 4.2. Likewise $1/hour higher wages are associated with a 5 percent increase in the likelihood of employing a worker rated as excellent by their manager.

Quantity of workers recruited. The online retailer hires entry-level remote customer service representatives at $14/hour throughout the United States, despite heterogeneity in the local pay for customer service representatives (which is shown in Figure **??**). If this market were perfectly competitive, the retailer would not hire anyone from MSAs with higher pay. Likewise, if the relative wage were the only determinant of recruitment, the retailer would only attract representatives from the lowest paying MSAs. If instead, limited information or heterogeneous preferences contribute to recruitment, we would expect representatives to be drawn from MSAs with a range of pay. We find that the retailer hires throughout the country and higher relative pay increases recruitment in MSAs throughout the country.

The uniformity of the retailer's wage creates heterogeneity in the retailer's pay, relative to the representatives' local outside options. For example, in Dallas, TX, the retailer's pay is far below the average entry-level rate for customer service; by contrast, in Lufkin, TX, a couple hours from Dallas, the retailer's pay exceeds many of the less lucrative alternatives.

In relative terms, representatives in Lufkin are paid more than representatives in Dallas for the exact same work. We use variation in relative pay to draw inferences about pay's impacts on the number of recruits and the turnover and productivity of those recruits once at the retailer. We define relative pay at the retailer to be the difference between its uniform \$14/hr rate and the entry-level pay for customer service in the MSA according to Emsi.[21]

We consider the the relationship between relative wage in the MSA and the number of people hired in the MSA, in the cross section:

$$\#\ \text{Hired}_{MSA} = \beta_0 + \beta_{\$}(\text{Entry Relative Wage})_{MSA} + \Gamma(\text{MSA Controls})_{MSA} + \epsilon_{MSA}. \tag{4}$$

$\beta_{\$}$ reveals the relationship between relative pay and recruitment, holding fixed the nature of the work. MSA controls include the number of customer service workers in the MSA. To test the robustness of these estimates, we also include in the MSA controls other features of the MSA, including whether the retailer has a warehouse in that MSA and the number of people employed by the retailer in that MSA.

For $\beta_{\$}$ to offer an unbiased assessment of the effect of the retailer raising its own wage, determinants of recruitment other than pay and the size of the available pool of customer service workers must be orthogonal to pay. This assumption is plausible since we are looking at the number of individuals recruited (not the quality thereof), which may reasonably depend only on the number of available workers and the pay relative to the outside option. Note that this assumption allows relative pay to affect the selection of workers drawn from the pool of available workers — indeed, this is an important component of the return of higher pay from the perspective of the retailer, which we investigate in the context of the shipper.

As shown in Table 4, Panel A, every additional dollar the retailer pays above the average, local entry-level rate is associated with between 0.17 and 0.22 more customer service recruits in the MSA off of an average of 0.73. This translates into an elasticity of recruitment with respect to the wage of between 3.2 and 4.2.[22] When customer service representatives are considering dif-

[21]We approximate the changing entry pay in the MSA according to the average of the 25th and 50th percentiles of the local customer-service wage distribution from EMSI.

[22]We estimate the recruitment elasticity of a specific firm rather than the job-to-job recruitment elasticity or job-

ferent options at the recruitment stage, their decision-making seems heavily swayed by relative pay. Nevertheless, the fact that this elasticity is finite suggests that informational limitations or working preferences also seem to weigh in decision-making.[23]

We return to the assumption that the elasticity of recruitment equals the elasticity of turnover (Manning, 2003). Our elasticity estimates for turnover (-4.48) and recruitment ($3.19 - 4.22$) in the customer service context are similar in magnitude. Thus our estimates provide a measure of confidence in the assumption used in many parts of this literature.

Quality of workers recruited: We likewise examin whether higher relative pay enhances the quality of workers placed in a given job. We do this using the fact the shipper in question hires warehouse workers through the staffing agency at $17/hour, regardless of the location of the warehouse. Overall, the $17 per hour far exceeds the average pay in the staffing agency for warehouse work of $11.74. However, there is variation in the going rate for temporary warehouse workers across the country – some areas pay $15 per hour, some $11 per hour (see Figure B.5). Again, this creates variation in the relative value of the shipper's wages.

Through the staffing agency, we can see workers' reviews from their on-site managers, which we use as a for the quality of the worker. Among all warehouse workers, only 13 percent of workers receive an "Excellent" review.

To evaluate the quality of the workers placed in a job, we construct a prediction of workers' evaluations based on their prior assignments, job evaluations, and job endings.[24] Only 8 percent of workers are predicted to earn an excellent evaluation, and another 8.9 percent are predicted to earn a poor evaluation. Another 62 percent are new workers, and thus do not have evaluations from which to predict their quality.

As shown in Table 4, Panel B, we find that an additional dollar in relative hourly pay means the shipper is 5 percent (0.4 percentage points off a base of 8 percent) more likely to have a worker to is predicted to be reviewed excellently and 8.5 percent (0.76 percentage points off a base of 8.9 percent) less likely to have a worker who is predicted to be reviewed poorly.

6 THE RETURN TO HIGHER PAY

When considering how to procure sufficient effective labor, firms must weigh the benefits of higher pay against the cost of paying more. To better inform this debate, we estimate the returns

from-non-employment elasticities that might reflect elasticities relevant at a market level. The elasticities captured here reflect those that are relevant to an individual firm.

[23]Our findings are consistent with reports from firms who have voluntarily raised wages. For example, Doug McMillon, the CEO of Wal-Mart, said after a wage hike in 2015 "[o]ur job applications are going up and we are seeing some relief in turnover" (Layne, 2015).

[24]We can, in principle, run the same exercise with evaluations for this job. However, if managers are aware of the outside option, this may change their baseline expectations.

to paying higher wages, using our estimates of turnover and productivity elasticities.[25] We find that in both the warehouse context, where estimates arise from a deliberate increase in pay, and the customer service setting, where estimates arise from keeping pay constant, productivity shifts are instrumental in offsetting the costs of higher wages.

Warehouse Workers. At the retailer's warehouse, a \$1/hour increase in pay yields a gross return of \$1.42 to 1.54 from reduced turnover costs and increased warehouse efficiency.

The gross returns from decreased turnover in the warehouse are \$0.28 to \$0.40. Internal estimate of the cost of training (\$689), overtime while new workers get up to speed (\$860), drug testing, badges and other overhead (\$300) suggest that the retailer pays at least \$1849 per new recruit. We find that an increase of \$1/hour means the warehouse has 2.5 fewer workers per hundred employees leave each month, yielding a savings of (2.5 fewer turnovers x \$1849) \$4623 per month. If the firm had to pay 100 workers who worked 21 eight-hour days in a month, \$1/hour more in order to affect this change, the cost to the firm would be (100 workers x 168 hours/month) \$16,800. Thus, their gross return on a \$1 investment would be \$0.28 (\$4623/\$16800). However, the data from our firm suggests that each worker was only working 116 hours per month, in which case the gross return would be \$0.40.

The gross returns of increased productivity in the warehouse are \$1.14. Based on hourly pay in the treated warehouse, in the quarter before the pay jump, the firm was spending \$3.29 dollars per box moved (\$16.20 in hourly wages / 4.92 boxes moved per person-hour). Since the higher pay increased the warehouse level productivity by 0.336 boxes per person-hour, the gross return on a \$1 investment is \$1.10.[26]

Customer Service Representatives. Among customer service representatives at this retailer, the gross return on a \$1/hour increase in the relative wage is \$1.25.

Among customer service representatives, we find moderately small decreases in turnover from increasing relative pay. Each additional \$1/hour is associated with a decrease in monthly turnover of 1.3 representatives out of 100. We estimate the cost of replacing a customer service representative to be \$2,100, consisting of \$1800 over the course of their 3-week training and \$300 in badges and other administrative costs. According to these estimates, increased retention would thus reflect a savings of \$2,730. To achieve these savings, 100 customer service representatives working 21 eight-hour days, would have to be paid an additional dollar (totaling \$16,800), implying a gross return of \$0.16 (\$2,730/\$16,800).

[25] A key drawback of these estimates is that they do not incorporate the recruitment elasticities since the retailer did not have estimates of the cost of recruiting a worker.

[26] These figures do not include any changes in petty theft (which we assume would go down with a better paid workforce), damages (which we could imagine might go up due to increased congestion, or down due to a more practiced team working in the warehouse), or a slower warehouse footprint expansion. Thus, on balance, we suspect this is an underestimate of the returns to boosted productivity.

The gross returns from an increase in productivity among customer service representatives is \$1.06. Each call costs the firm roughly \$4.60 (\$15.60 average wage rate · 18 minutes/call). A higher wage increases call volume by 1.90 calls per day, so the return on an \$8/day in higher wages is \$8.74 (\$4.60 x 1.90) – or \$1.09 on the \$1/hour investment.

This estimate may yet be an underestimate if firms are able to leverage higher pay's effect on recruitment. The magnitude of the recruitment elasticity suggests that higher pay might appreciably reduce the time it takes recruiters to find acceptable candidates.

Optimal Wages. One could look at these estimates and conclude that this Fortune 500 retailer did not set wages optimally. Two items are worth noting. First, the estimates in the warehouse reflect the fact that the firm recognized that the wages were suboptimally low and raised them accordingly. From this perspective, it is not surprising that our estimates would suggest that it was profitable for the firm to take the steps that it did.

Second, more curious, perhaps, is that our estimates in the customer service context, which arise from sticky wages, should also suggest that wages could be profitably raised. But this firm uses a sticky wage strategy, in which firms set prices intending to keep them fixed for some time — a strategy that could make sense if there are substantial adjustment costs. The sticky wage strategy acknowledges that before wages are adjusted, they may be suboptimal. So while we do highlight that wages need to be adjusted, this does not mean that the firm is not, therefore, rational.

We can also use our wage equation from Section 1 to estimate what pay should be in the customer service context according to our model.[27] If we assume that marginal productivity, $p(w)$ was captured to a first order by sticky wages, we find that pay would only have been marginally higher at \$15.72 rather than \$15.60. [28] Of course, this is an out-of-sample estimate, and thus makes several assumptions, including constant elasticities and productivity responses.

7 Mechanisms: Selection and Behavioral Responses

A key question underlying our results is the mechanisms by which changes in pay affect measured turnover and productivity. Does raising wages for existing workers incentivize higher performance, or is the key advantage of higher pay attracting and retaining more reliable or productive workers?

To answer this, we use staffing agency data that follow the same worker in multiple jobs to assess how much pay affects the behavior of the individual worker. We find that over half of

[27]We refrain from estimating optimal wages in the warehouse context since the recruitment elasticity can only be estimated in the customer service setting.

[28]$w^* = \frac{p(w)}{1+\frac{1-p'(w)}{\epsilon_{R,w}-\epsilon_{T,W}}} = \frac{15.60}{1+\frac{1-1.06}{3.19+4.48}} = \frac{15.60}{1-\frac{0.06}{7.67}} = 15.72$

the turnover reduction and productivity increase arises from behavioral responses of the same worker facing different wages.

Since some of the effects of higher pay arise from sorting better workers to higher-paying firms, we also estimate the negative spillovers on other firms that are hiring for the same position at the same time. In the staffing agency context, we do find that when a high-paying firm is filling positions, the people placed in other, similar, local jobs perform slightly less well.

7.1 Retaining Better Workers

Higher pay not only attracts quality workers, but retains them.

We leverage the stickiness of the retailer's wage for customer service representatives to assess retention elasticities for workers with different baseline productivity, using Equation 3. We hypothesize that the pay in local outside options is more important for more productive workers, who are better able to convert lucrative outside options into job offers that draw them out of the retailer. We test this hypothesis by investigating whether turnover rises more sharply for highly productive workers in those MSAs where the retailer's sticky wage loses more ground to the local alternatives.

We find that higher pay is particularly effective at retaining representatives who start in the top third of daily call volume in their first month, as shown in Table 6, Panel A. Each $1/hr of additional relative pay reduces turnover by 44% for initial top performers, implying a turnover elasticity of 6.6. By contrast, for the rest of the representatives, the same increase in relative pay decreases turnover by 17%, implying a turnover elasticity of 2.76. This suggests that *selective retention* may be an important driver of increased aggregate productivity.

7.2 Incentivizing Better Work

To assess whether higher pay incentivizes better work, we consider the heterogeneous effects of higher relative pay across workers with different baseline productivity. If less skilled workers at baseline are more at risk of termination or are less likely to be promoted than higher-productivity workers, one might expect the output of these workers to be more sensitive to the relative pay of the retailer. Indeed, in Table 6, Panel B, we find that call-volume effects are concentrated among representatives who are in the bottom two-thirds of call volumes in their first month after training at the retailer, as consistent with these representatives being more concerned about the possibility of termination or that they won't be promoted. Representatives in the top third of first-month daily call volumes have no appreciable change in their call volumes when their relative pay quasi-randomly changes.[29]

[29] While one might worry about mean-reversion when comparing those who start in the top of the cohort to those who start in the bottom, we would expect this to be symmetrical, and in these analyses we do not find symmetrical mean reversion.

7.3 Source of the effect

To understand what share of the effects come from the same worker facing different wages and adjusting their behavior accordingly, we leverage data from a staffing agency. While the dataset is distinct from the retailer data, the staffing agency places many workers in similar warehouse jobs, allowing us to consider the effects of pay on this occupation. Because we observe the same worker in multiple, comparable jobs with different pay, we can see what percent of the reduced-form relationship is present when the same worker faces different pay rates.

For this analysis, we focus on the sample of warehouse workers placed in temporary jobs by the staffing agency (N=222,904), since this offers a relatively homogeneous set of jobs. We begin by estimating the reduced-form relationship between pay and performance:

$$\text{Turnover}_{ij} = \beta_0 + \beta_{\$} \cdot \text{Pay}_j + \mu_{oc} + \mu_{dct} + u_{ij}. \tag{5}$$

where $\beta_{\$}$ captures the relationship of interest. To isolate the pay premium of the firm above the local market, we include occupation-by-commuting-zone fixed effects and industry-by-commuting-zone-by-month fixed effects that together soak up variation in the local labor market. Our estimates are thus identified off of variation in hourly pay across firms and workers in the same local labor market and industry.

To isolate the incentive effects of higher pay, we look at the relationship between completion and pay *within* individual workers who work multiple jobs at the staffing agency, using the regression:

$$\text{Turnover}_{ij} = \psi_{\$} \cdot \text{Pay}_j + \underbrace{\mu_i}_{\text{Worker FE}} + \mu_{oc} + \mu_{dct} + v_{ij}.$$

We estimate both of these specifications for the sample of workers with multiple jobs through the agency, since these workers are used to identify the within-worker effect of higher pay. Table 7 presents the results of this analysis. We find that an additional dollar of pay increases job completion by 2.6 percentage points, off a base of 40 percent completion. This is equivalent to an elasticity of 0.72. We estimate that 83 percent of that effect arises within the same worker. A similar share of the quits may be attributed to behavioral responses.

We can also use “Excellent” evaluations as a metric of whether the firm is satisfied. While not the same as on-the-job productivity, it is nevertheless a useful metric of worker performance. In this case, we find that 50% of the increase associated with higher pay arises within the same worker.

7.4 Spillovers to Local Firms

Since a portion of the boost in relative completion rates may be attributed to sorting more reliable workers to higher paying firms, a natural question arises: do other local firms suffer the consequences? If so, is zero-sum between firms?

We address how higher wages may have spillovers to other firms in the context of the shipper, since the shipper (a) hires through the staffing agency only during the holiday season when they need more workers, and (b) pays all their workers \$17/hour regardless of location (as discussed in Section 5). We explore two types of spillovers: direct poaching and selection. We first look at whether workers within the agency leave their assignments at the time when the shipper is hiring to see if the shipper causes other firms to have unexpected turnover. We then look at whether the shipper may simply attract better workers, such that firms hiring for the same type of position in the same time periods have lower job completion and performance.

We use a difference-in-differences approach, comparing the change in worker quality around the holiday season in commuting zones where the central firm is present to the change in worker quality in control commuting zones. Because where to locate is a considered decision for firms, the choice of a control group is key to the validity of our strategy. We leverage the fact that we see rival firms that perform almost exactly the same function in our data to construct a control group. Among these rival firms, the considerations about where to locate are likely fairly similar – a supposition borne out by the fact that a great number of the locales overlap. We use the places where the rival firms have located but the central firm has not as the control location. The sample is described in Table 1, Panel B, Column 3.[30] The commuting zones where the shipper locates tend to have slightly higher pay than locations where only rivals locate. Nevertheless, more workers quit and more have bad endings in areas where the shipper locates, which might be due to having jobs that are expected to last longer.

We use a simple difference-in-differences approach, where we fully interact the specification with year to ensure that we do not put negative weight on any of our comparisons, in keeping with the recent literature on two-way fixed effects models (e.g., Goodman-Bacon, 2018; Abraham and Sun, 2018; Imai and Kim, 2019; de Chaisemartin and D'Haultfoeuille, 2020).

$$\text{Job Completion}_{it} = \alpha_0 \mathbb{1}^{cz}_{ijt} + \alpha_1 \mathbb{1}^{season}_{ijt} + \alpha_2 y_{ijt} + \alpha_3 y_{ijt} \mathbb{1}^{cz}_{ijt} + \alpha_4 y_{ijt} \mathbb{1}^{season}_{ijt} + \beta(y_{ijt} \cdot \mathbb{1}^{cz}_{ijt} \cdot \mathbb{1}^{month}_{ijt}) + \epsilon_{ijt} \tag{6}$$

where $\mathbb{1}^{cz}_{ijt}$ is an indicator for each commuting zone, $\mathbb{1}^{season}_{ijt}$ is an indicator for each season, y_{itj} indicates the year, and the β coefficients are aggregated into our coefficient of interest using inverse-variance weighting. We cluster standard errors at the MSA level for the regression and calculate bootstrapped standard errors for the weighted coefficient that aggregates the yearly estimates.

We would expect that in areas and times where the central firm's going rate of \$17 per hour is greater than the average pay for a warehouse worker, the negative effects on other firms would be larger. As such, in our next specification, we fully interact Equation 6 with $\overline{\text{Pay}}_{ijt}$, the

[30] We define a treated commuting zone to have at least one month in which the shipper hires more than 45 individuals in that month and to have had at least 20 hires outside the shipper. We require that control months have at least 20 job placements outside the shipper. The entire sample is limited to warehousing jobs. We further restrict the sample to the three months just before the shipper seeks workers to account for potential seasonality.

average pay differential in the commuting zone - season pair. In this case, β thus captures how much an additional dollar of relative pay offered by the shipper impacts job outcomes at rival firms in each of our treated years.

To assess the parallel trends assumption, we plot in Figure B.6 the average pay rates for warehouse jobs in the treated and untreated commuting zones in orange and blue dots, respectively. The shaded areas show the months where the central firm is hiring more than 50 individuals. While the central firm tends to locate in commuting zones that tend to have slightly higher pay than the areas where their rivals alone locate, the trends in pay are fairly similar throughout the time period. Additionally, we test for pretrends analytically, by adding treatment-season fixed effects, as Pischke (2019) recommends, and we find no significance.

As seen in Table 8, when the shipper is hiring at all, quits at rival firms increase by 12.4 percentage points off a base of 28 percent. An additional dollar of pay over the outside option is associated with a 1.45 percentage point increase in quits. We also assess bad endings–namely when workers be terminated for performance or attendance, or to receive a "Poor" Evaluation. When the shipper is hiring, bad endings at rival firms increase by 8 percentage points, off a base of 24 percent.

If selection is at work, when the shipper is hiring rival firms may also hire lower quality workers. We assess this possibility in Table A.7. Workers hired into rival firms when the shipper is hiring are 0.98 percentage points less likely to be predicted excellent off a baseline of 9.7 percent in control commuting zones and months and also 2 percentage points less likely to be new workers, off a baseline of 40.7 percent.

If pay is so much better at the shipper, one could imagine workers at the agency leaving their existing gigs in order to take higher-paid positions.[31] To see whether the higher pay at the shipper leads workers to quit their existing jobs, we conduct a difference-in-differences regression, comparing the warehouse jobs that end commuting zones and months where the shipper is hiring to the job endings in other locations where the shipper's rivals locate. This is distinct from the analysis above, where we were comparing workers *placed* in jobs at the same time; here we examine those jobs that *end* in the months when the shipper is hiring.

We consider all warehouse jobs in commuting zones in which the shipper or its rivals are located. Each ongoing job in a given month has an observation for that job-month since the worker *could* choose to terminate in that month. Thus the interpretation of β from Regression 6 in this context is the change in the percent of ongoing jobs that are completed/quit in a month when the shipper is hiring at a pay differential of \$1/hour more than the outside option.

Table A.8 shows that there is not an uptick in staffing workers leaving their job or otherwise

[31] The staffing agency does *not* prohibit workers from moving between client firms. Indeed in our data, 13,949 assignments end because the worker switches to another job within the same industry.

having an unsatisfactory end in order to take jobs at the central firm, suggests that in the context of temporary jobs, there is relatively little poaching.

8 Gender Differences in Elasticities

Elasticities of retention, turnover, and productivity may differ by gender.[32] Different elasticies would imply that workers of different genders may command different wages, an idea that goes back as far as Robinson (1933). We investigate whether there is indeed heterogeneity in the genders' responsiveness to pay and what this implies about wage gaps.

In the context of customer service agents, we find that labor supply elasticities may explain a \$0.08 gap in pay, but that women's productivity is considerably more responsive to pay, suggesting that women should be paid *more* than men.

Of course, in the face of anti-discrimination laws, this exercise does not capture what firms can legally carry out in setting wages. However, even in the presence of anti-discrimination laws, gender differences in elasticities can help explain occupation-level gaps in pay for particularly gendered occupations. Moreover, the exercise may be useful even within an occupation. As long as an individual firm is satisfied by certain group of workers, they can adjust wages, knowing that they may end up with a workforce that is mostly composed of one gender. For instance, if men have greater labor supply elasticities, a firm that doesn't mind an all-female workforce can simply keep pay low. Thus at an occupation level, we can still see a wage gap emerge when firms stay within the bounds of the law.

Using the wage expression from Section 1, we can unpack the implications for the gender wage gap.

$$\frac{w_f}{w_m} = \frac{\frac{p_f(w)}{1+\frac{1-p_f'(w)}{\epsilon_{f,R,w}-\epsilon_{f,T,W}}}}{\frac{p_m(w)}{1+\frac{1-p_m'(w)}{\epsilon_{R,w}-\epsilon_{T,W}}}}$$

Assuming a constant production function across genders, we have:

$$\frac{w_f}{w_m} = \frac{\frac{1}{1+\frac{1-p_f'(w)}{\epsilon_{f,R,w}-\epsilon_{f,T,W}}}}{\frac{1}{1+\frac{1-p_m'(w)}{\epsilon_{R,w}-\epsilon_{T,W}}}} = \frac{1+\frac{1-p_m'(w)}{\epsilon_{R,w}-\epsilon_{T,W}}}{1+\frac{1-p_f'(w)}{\epsilon_{f,R,w}-\epsilon_{f,T,W}}}.$$

If we set $p'(w) = 0$ and assume that $\epsilon_{R,w} = \epsilon_{T,w}$—an assumption that we validate in the customer service context, we have the equation used by Ransom and Oaxaca (2010) to calculate gender disparities in grocery stories.

[32]We refer to differences by gender because the data received from the retailer delineates gender. We guess that this field then documents self-reported gender and that we lack information on workers' sex.

8.1 Heterogeneity in Elasticities

We explore the degree to which turnover and productivity elasticities differ by gender in the context of customer service representatives.

Elasticities of Turnover. When deciding whether to leave the retailer, men are more elastic than women in customer service.An additional dollar of pay reduces turnover by over 40 percent for male customer service representatives, implying a turnover elasticity of 6.6. Female customer service representatives' response to higher relative pay is economically small and statistically indistinguishable from zero. These findings are consistent with findings from Wiswall and Zafar (2018), among others, that women prefer job stability whereas men prefer jobs with higher earnings growth.

These results are in sharp contrast to the findings in Caldwell and Oehlsen (2018), who suggest there are minimal gender differences in daily labor supply responses among Uber-drivers. Our setting of customer-service representatives contrasts with Caldwell and Oehlsen (2018)'s setting in a few key ways. First, most customer service representatives have full-time jobs at the retailer; thus, our estimates reflect the frictions that keep workers at their primary employers, rather than the rigidities in their decision-making about gig work. Second, in contrast to the male-dominated setting of Uber-driving, our setting is one in which women make up the majority — suggesting that the women in this occupation may be less strongly selected than those in Uber-driving. Conversely, men in the customer service may be selected.Finally, our setting may have more workers of child-bearing age than Uber, where the gender differences appear most pronounced. Our results are consistent with Hirsch et al. (2010) who find in matched employer-employee data from Germany that women's labor supply elasticities to the firm are smaller than those of male workers.

We are underpowered to estimate elasticities of recruitment by gender since there are only 93 male remote customer service workers. As such, when we trace out implications for the gender pay gap, we use the assumption—common in the new monopsony literature—that $\epsilon_{R,w} = \epsilon_{T,w}$, which we found to be true in the aggregate (see Sections 3 and 5). Nevertheless, results can be found in Appendix Table A.6

The wage equation derived in Section 1 helps us trace out implications for the gender pay gap of our different elasticities. We can take the ratio of female to male wages understand how labor supply elasticities would act on the wage gap.

$$\frac{w_f}{w_m} = \frac{\frac{1}{1+\frac{1}{\epsilon^f_{R,w}-\epsilon^f_{s,w}}}}{\frac{1}{1+\frac{1}{\epsilon^m_{R,w}-\epsilon^m_{s,w}}}} = \frac{1+\frac{1}{\epsilon^m_{R,w}-\epsilon^m_{s,w}}}{1+\frac{1}{\epsilon^f_{R,w}-\epsilon^f_{s,w}}}$$

Substituting in our elasticities:

$$
\begin{aligned}
&= \frac{1 + \frac{1}{2*-6.6}}{1 + \frac{1}{2*-3}} \\
&= \frac{1.07}{1.16} \\
&= 0.92
\end{aligned}
$$

Thus looking simply at the difference in extensive labor supply elasticity, we would end up with a slight wage gap, with women earning 92 cents for each dollar men earn if the law did not prohibit gender-based wage discrimination.

Elasticities of Productivity. We explore whether female and male customer service representatives are differentially responsive to pay in their productivity. We find that when relative pay is higher, female workers' productivity increases by 9 percent (2.3 off of a base of 25.58 calls per day) whereas male workers' productivity increases up by just 5 percent (1.24 calls off of a base of 24.41 calls per day), statistically indistinguishably from zero with our sample size. For women, this reflects an elasticity of 1.41, while for male workers it represents an elasticity of 0.8.

To translate this into dollar differences, consider that the retailer pays $4.60 per call. So the male increase of 0.8 calls per day amounts to a savings of $3.68/day or $0.46/hour. In contrast, the female increase of 2.3 calls/day amounts to a savings of $10.67/cay or $1.33/hour. This would suggest that if productivity responses to pay were incorporated into wages, women should have *higher* wages than men.

9 CONCLUSION

In this paper we present evidence that warehouse workers and customer service representatives are responsive to wages, not only with regard to recruitment and turnover, but also with regard to their on-the-job productivity. We estimate recruitment elasticities in excess of 3, turnover elasticities between -3 and -4.5, as well as productivity elasticities in excess of one. The productivity response to higher pay yields a net positive return. We estimate that 80 percent of the improvement in turnover arises from workers' behavioral responses to higher pay.

This paper also estimates gender differences in these elasticities. We find that while women's labor supply is slightly less elastic than men's, women increase their productivity in response to higher pay more than do men. The gender difference in labor supply elasticity is important because it suggests that when the concentration of firms is used as a measure of monopsony power, we may underestimate firms' power to set female wages. The productivity response is particularly intriguing because it suggests that if wage discrimination were not illegal, women should be paid more than men in this context. It also suggests that understanding workers' productivity responses will be particularly important in gaining a more comprehensive understanding of gender pay gaps.

Our results do have limitations and leave a number of questions unanswered. While we provide suggestive evidence for the mechanisms underpinning our results, we cannot perfectly estimate the relative contributions of attracting better workers versus eliciting greater effort from the existing workforce. This is a significant shortcoming insofar as it means we are unable to make general equilibrium predictions, such as what might happen if wages were raised universally in a given geography. If the effect we document is coming entirely from greater effort, then all firms might see an increase in productivity when all workers' pay is raised. If higher pay generates greater productivity only because better workers gravitate toward higher-paying firms, then a global increase in pay will not induce greater productivity since no sorting would occur.

Moreover, if on-the-job productivity increases with pay because reduced turnover itself increases output, then the resultant question is whether turnover is a function of relative pay or absolutely higher pay. We find that turnover is responsive to both relative and absolute pay. A more thorough investigation into the mechanisms would be valuable.

Our results that explore the spillovers of high-paying firms on local rival firms can only document the effect in terms of worker turnover and firm satisfaction, not in terms of objective measures of productivity. We would love to know whether work is slower at rival firms when a local firm raises pay, and view this as an important avenue for future work.

While there is ample room for additional research, this paper contributes by (a) estimating turnover and recruitment elasticities with respect to wages among warehouse and customer service workers, (b) bringing objective productivity metrics to bear on the question of how pay affects workers, (c) providing suggestive evidence about the relative contributions of selection and workers' behavioral responses to pay and about the spillovers on other firms that can arise from selection, and (d) estimate gender-specific responses to pay.

References

Abraham, Sarah and Liyang Sun, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Available at SSRN 3158747*, 2018.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu, "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies," Technical Report, National Bureau of Economic Research 2020.

Bó, Ernesto Dal, Frederico Finan, and Martín A Rossi, "Strengthening state capabilities: The role of financial incentives in the call to public service," *The Quarterly Journal of Economics*, 2013, *128* (3), 1169–1218.

Burdett, Kenneth and Dale T Mortensen, "Wage differentials, employer size, and unemployment," *International Economic Review*, 1998, pp. 257–273.

Caldwell, Sydnee and Emily Oehlsen, "Monopsony and the Gender Wage Gap: Experimental Evidence from the Gig Economy," 2018.

Cappelli, Peter and Keith Chauvin, "An interplant test of the efficiency wage hypothesis," *The Quarterly Journal of Economics*, 1991, *106* (3), 769–787.

de Chaisemartin, Clement and Xavier D'Haultfoeuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2994.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform pricing in us retail chains," *The Quarterly Journal of Economics*, 2019, *134* (4), 2011–2084.

Derenoncourt, Ellora, Clemens Noelke, and David Weil, "Do social norms around pay influence the wage-setting behavior of firms?," 2020.

Dickens, William T, Lawrence F Katz, and Kevin Lang, "Are Efficiency Wages Efficient?," Technical Report, National Bureau of Economic Research 1986.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri, "Monopsony in online labor markets," *American Economic Review: Insights*, 2020, *2* (1), 33–46.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," Technical Report, National Bureau of Economic Research 2018.

Hirsch, Boris, Thorsten Schank, and Claus Schnabel, "Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer-employee data from Germany," *Journal of Labor Economics*, 2010, *28* (2), 291–330.

Imai, Kosuke and In Song Kim, "On the use of two-way fixed effects regression models for causal inference with panel data," *Unpublished paper: Harvard University*, 2019.

Katz, Lawrence F, “Efficiency wage theories: A partial evaluation,” *NBER Macroeconomics Annual*, 1986, *1*, 235–276.

— **and Alan B Krueger**, “Changes in the Structure of Wages in the Public and Private Sectors,” Technical Report, National Bureau of Economic Research 1991.

Krueger, Alan B and Lawrence H Summers, “Efficiency wages and the inter-industry wage structure,” *Econometrica: Journal of the Econometric Society*, 1988, pp. 259–293.

Layne, Nathan, “Wal-Mart says staff turnover down after hike in minimum pay,” *Reuters*, 2015.

Manning, Alan, *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press, 2003.

Orszag, J.Michael and Gylfi Zoega, “Wages ahead of demand,” *Economics Letters*, December 1996, *53* (3), 341–347.

Pischke, Jörn-Steffen, “Differences-in-Differences,” 09 2019. `http://econ.lse.ac.uk/staff/spischke/ec533/did.pdf`, accessed 2020-01-05.

Raff, Daniel M G and Lawrence H Summers, “Did Henry Ford Pay Efficiency Wages?,” *Journal of Labor Economics*, October 1987, *5* (4), 31.

Ransom, Michael R and David P Sims, “Estimating the firmâĂŹs labor supply curve in a âĂIJnew monopsonyâĂİ framework: Schoolteachers in Missouri,” *Journal of Labor Economics*, 2010, *28* (2), 331–355.

— **and Ronald L Oaxaca**, “New market power models and sex differences in pay,” *Journal of Labor Economics*, 2010, *28* (2), 267–289.

Robinson, Joan, *The economics of imperfect competition*, London: Macmillan, 1933.

Shapiro, Carl and Joseph E Stiglitz, “Equilibrium unemployment as a worker discipline device,” *The American Economic Review*, 1984, *74* (3), 433–444.

Slichter, Sumner H, “The management of labor,” *Journal of Political Economy*, 1919, *27* (10), 813–839.

Staiger, Douglas O, Joanne Spetz, and Ciaran S Phibbs, “Is there monopsony in the labor market? Evidence from a natural experiment,” *Journal of Labor Economics*, 2010, *28* (2), 211–236.

Sullivan, Daniel, “Monopsony power in the market for nurses,” *The Journal of Law and Economics*, 1989, *32* (2, Part 2), S135–S178.

Webber, Douglas A, “Firm market power and the earnings distribution,” *Labour Economics*, 2015, *35*, 123–134.

Wiswall, Matthew and Basit Zafar, "Preference for the workplace, investment in human capital, and gender," *The Quarterly Journal of Economics*, 2018, *133* (1), 457–507.

Yellen, Janet, "Efficiency wage models of unemployment," *American Economic Review*, 1984, *74* (1), 200 – 250.

10 Tables

TABLE 1: SUMMARY STATISTICS

Panel A: Warehouse and Customer Service Samples from an Online Retailer

	Warehouse			Customer Service		
	Treated	In-State	Twin	All CSR	All Remote	Sticky Pay
$/hour	16.20	15.66	16.24	15.83	14.35	16.02
Entry CSR $/hr in MSA				13.81	14.46	13.52
% Turnover/Month	13.40	9.30	5.15	4.40	6.59	4.31
% Quit/Month	10.66	7.50	4.05	3.57	5.55	3.50
% Fired/Month	2.02	1.24	0.70	0.62	0.83	0.58
% Turnover/Month in MSA				6.25	6.96	6.27
Days in Co	276.32	314.17	235.17	325.91	172.21	333.89
% Female	21.89	52.50	20.31	70.75	88.58	69.25
Age	36.09	37.55	33.59	33.82	36.48	33.65
Boxes/Hour	4.92	6.51	2.76			
Boxes/Moving Hour	7.69	10.48	5.16			
Moving/Total Hours	0.64	0.62	0.55			
Calls/Day				25.11	25.32	25.27
Calls/Hour				3.27	3.28	3.30
Absent Unapproved Hrs				0.43	0.43	0.43
# Employees	368	690	896	4,551	593	3,061
# Days	20,824	48,401	59,994	1,289,980	115,685	854,614

Panel B: Temporary Warehouse Positions from a Staffing Agency

	All Warehouses	High Roller	Local Warehouses
$/Hour	11.74	17.00	12.51
% Job Completed	44.15	83.57	41.72
% Quit	31.84	5.85	33.98
% Bad Ending	27.93	6.15	30.64
% Excellent Eval	13.15	20.87	6.45
Expected Duration	102.87	30.70	106.02
# Workers	140,664	5,701	32,009
# Assignments	222,904	6,664	45,454
# Commuting Zones	374	83	83
# Firms	3,950	1	1,448

Note: We use data from an online retailer's warehouse workers and customer service representatives (Panel A) as well as from a staffing agency's warehouse placements (Panel B). Statistics are aggregated from daily data in Panel A, meaning that workers who are present longer have greater weight than workers who are present for a short period. For the warehouse workers in Panel A, we limit to the three months before the pay change analyzed in the paper. In Panel B, statistics are aggregated from job-level data, so each job is weighted equally.

TABLE 2: HIGHER PAY'S EFFECTS ON TURNOVER

Panel A: Turnover Effects In the Retailer's Warehouse

	First Stage	Monthly Turnover	Quits	Fires
Post	1.755***			
	(0.079)			
$1/hour		−2.504**	−2.270**	0.061
		(1.255)	(0.957)	(0.563)
Elasticity		-3.03**	-3.45***	0.49
		(1.52)	(1.45)	(4.53)
Base Mean	16.2	13.4	10.66	2.02
Workers	514	514	514	514
Observations	50,478	50,478	50,478	50,478

Panel B: Turnover Effects Among Customer Service Representatives

	Monthly Turnover	Quits	Fires
Entry Relative $1/hr	−1.208**	−0.671	−0.206**
	(0.610)	(0.561)	(0.090)
Elasticity	-4.484**	-3.071	-19.148**
	(2.264)	(2.57)	(8.375)
Date Fixed Effects	✓	✓	✓
Mean $/hr	16.02	16.02	16.02
Dependent Mean	4.31	3.5	0.17
MSAs	42	42	42
Workers	3061	3061	3061

Panel C: Turnover Effects Among Temporary Warehouse Workers

	Job Completed	Quits	Bad Ending
Relative Hourly Pay	1.165*	−0.482***	−0.319*
	(0.678)	(0.168)	(0.183)
Elasticity	0.24	-1.4	-0.86
	(0.14)	(0.49)	(0.5)
Season Fixed Effects	✓	✓	✓
Controls	Days Quartic	Days Quartic	Days Quartic
Base Mean	83.4	5.9	6.3
Workers	5,763	5,763	5,763
Observations	6,398	6,398	6,398
R^2	0.127	0.116	0.021

Note: *p<0.1; **p<0.05; ***p<0.01

We calculate retention responsiveness to (relative) wages in three contexts: Panel A shows where the retailer quasi-randomly increased wages within a single warehouse. Standard errors are clustered at the worker level. Panel B shows the effect of relative wages on customer service representative retention, using a sticky pay design. Panel C shows temporary warehouse workers who are subject to national wage setting.

TABLE 3: HIGHER PAY'S EFFECTS ON PRODUCTIVITY

Panel A: Pay's Effects on Productivity in the Warehouse

	First Stage	Boxes/Hr	Boxes/Moving Hr	Moving/Total Hrs
Post	1.746***			
	(0.054)			
$1/hour		0.336***	0.309**	0.017**
		(0.088)	(0.138)	(0.008)
Productivity ϵ		1.1	0.65	0.43
		(0.29)	(0.29)	(0.2)
Pre Jump Mean	16.2	4.93	7.7	0.64
Observations	26	26	26	26

Panel B: Pay's Effects Among Customer Service Representatives

	Daily Call Volume	Satisfaction (out of 5)	% of Absences Unapproved
Entry Relative $1/hr	1.904**	0.012***	−1.642
	(0.916)	(0.003)	(4.437)
Elasticity	1.2**	0.038***	0.383
	(0.58)	(0.01)	(1.036)
FE: Date	✓	✓	✓
Mean $/hr	15.96	15.96	16.11
Dependent Mean	25.27	4.89	68.31
MSAs	41	41	41
Workers	2687	2687	2782

Note: *p<0.1; **p<0.05; ***p<0.01

In Panel A, we show the interrupted time series and difference-in-differences estimates of the effect of pay on retention and productivity in the warehouse. The first two columns of both panels reflects the first stage, showing that after the pay change, hourly pay increased by $1.76 within the warehouse off of a mean of $16.20, a 10.8% increase. The next columns report the two-stage least squares estimates of the effect of pay on three types of monthly turnover. The estimates reflect warehouse level data and a 3-month bandwidth on either side of the pay jump. Appendix Table A.2 shows robustness to different bandwidths. Standard errors are clustered at the individual level.

In Panel B, we show how customer service productivity responds to higher pay relative to the local rate for customer service representatives. Standard errors are clustered at the MSA-level.

TABLE 4: EFFECTS OF HIGHER PAY ON RECRUITMENT

Panel A: Recruitment Quantity of Customer Service Representatives

	# Customer Service Representatives Hired				
Entry \$14/hr - MSA Entry \$/hr	0.168^{*}	0.195^{**}	0.206^{**}	0.219^{**}	0.222^{**}
	(0.087)	(0.089)	(0.094)	(0.101)	(0.101)
Elasticity	3.19^{*}	3.71^{**}	3.92^{**}	4.18^{**}	4.22^{**}
	(1.66)	(1.69)	(1.79)	(1.92)	(1.92)
# MSA Customer Service Workers	Linear	Log	Quartic	Quartic	Quartic
Retailer Non-CSR Presence				✓	✓
Retailer n-CSR Counts					✓
Mean Recruits/MSA	0.73	0.73	0.73	0.73	0.73
# MSAs	920	920	920	920	920
R^2	0.232	0.289	0.297	0.300	0.388

Panel B: Recruitment Quality of Warehouse Workers

	Predicted Excellent	Predicted Poor	New Worker
Relative Hourly Pay	0.869^{***}	-0.269	-0.746^{*}
	(0.301)	(0.196)	(0.426)
Elasticity	1.14^{***}	0.93	0.32
	(0.4)	(0.68)	(0.19)
Season Fixed Effects	✓	✓	✓
Dependent Mean	12.91	4.92	39.14
Workers	5,763	5,763	5,763
Observations	6,398	6,398	6,398
R^2	0.065	0.001	0.036

Note: We consider the relationship between relative pay of the employer and the number of customer service representatives ever recruited and hired in the MSA. Each observation is an MSA, excluding MSAs with on-site call-centers which have different advertising. Relative pay is the gap between the retailer's \$14/hr rate and the typical rate for entry-level workers, which we approximate with the average of the 25th and 50th percentiles of the local wage distribution. In the first column, we control only for a linear effect of the number of local customer service representatives in the MSA, whom the retailer could potentially draw from. In the second column, we instead control for a log in employment. In the third column, we control for a quartic in local employment in customer service. In the fourth column, we add indicators for the retailer having a warehouse in the MSA and the retailer having a corporate or sales' office in the MSA. In the final column, we also include controls for counts of the number of warehouse and other non-customer-service workers in the retailer in the MSA.

TABLE 5: HETEROGENEITY IN ELASTICITIES BY GENDER IN CUSTOMER SERVICE

Panel A: Pay's Effects on Recruitment by Gender

	Female	Male	Δ
Effect on Recruitment	0.13**	0.01	0.12**
	(0.06)	(0.02)	(0.05)
Elasticity of Recruitment	3.27**	1.86	1.41**
	(1.58)	(3.43)	
Mean Recruited/MSA	0.55	0.1	
Mean Pay	14	14	0
# Workers	508	93	
# MSAs	96	40	

Panel B: Pay's Effects on Turnover by Gender

	Female	Male	Δ
Effect on Turnover	-0.91	-1.87**	0.96*
	(0.6)	(0.6)	(0.54)
Elasticity of Turnover	-3.47	-6.63**	-3.16*
	(2.3)	(2.13)	
Mean Turnover	4.17	4.57	-0.78***
			(0.23)
Mean Pay	15.94	16.19	-0.07
			(0.07)
# Workers	2097	901	
# MSAs	39	23	

Panel C: Pay's Effects on Productivity by Gender

	Female	Male	Δ
Effect on Calls	2.32**	1.24	1.08**
	(0.95)	(1.03)	(0.54)
Elasticity of Calls	1.41**	0.8	0.6**
	(0.58)	(0.66)	
Mean Calls	25.58	24.41	0.3
			(0.34)
Mean Pay	15.47	15.75	-0.07
			(0.05)
# Workers	1555	618	
# MSAs	33	20	

Note: *p<0.1; **p<0.05; ***p<0.01

We calculate responsiveness to the retailer's (relative) wages, separately male and female customer service representatives. Recruitment differences utilize the same national-wage setting strategy as described in Section 5; turnover and productivity elasticities use the sticky pay design described in Section 3. Differences in means among customer service representatives are calculated base on regressions that include date and MSA fixed effects and have standard errors clustered at the MSA-level. Productivity analyses limit to those hired in 2018 or later because representatives hired earlier are able to handle different types of calls, which changes their call volume. Sensitivity of recruitment elasticities to other specifications can be viewed in Appendix Table A.6.

TABLE 6: EFFECTS OF PAY BY INITIAL PRODUCTIVITY IN CUSTOMER SERVICE

Panel A: Turnover Effects by Initial Productivity

	Start in Top Third	Rest of Workforce	Δ
Effect on Turnover	-2.13**	-0.77**	-1.36
	(1.06)	(0.33)	(1.05)
Elasticity of Turnover	6.68**	2.76**	3.92
	(3.31)	(1.19)	
Mean Turnover	4.98	4.39	0.32
			(0.44)
Mean Pay	15.62	15.78	-0.27***
			(0.04)
# Workers	615	1207	
# MSAs	17	25	

Panel B: Productivity Effects by Initial Productivity

	Start in Top Third	Rest of Workforce	Δ
Effect on Turnover	1.1	2.7**	-1.59**
	(0.71)	(0.87)	(0.59)
Elasticity of Turnover	0.55	1.8**	-1.25**
	(0.35)	(0.58)	(-0.23)
Mean Turnover	31.43	23.66	6.26***
			(0.37)
Mean Pay	15.62	15.78	-0.27***
			(0.04)
# Workers	615	1207	
# MSAs	17	25	

Note: *p<0.1; **p<0.05; ***p<0.01

We leverage the stickiness of the retailer's wage to estimate the effect of higher relative pay on turnover and call volumes for workers with different baseline productivities. We assess baseline productivity according to representatives' daily call volumes in their first month of calls after formal training. We find higher relative pay has a more pronounced effect on the daily call volume of representatives in the bottom two-thirds of productivity, as consistent with these workers being more concerned about termination and thus facing greater incentives to increase call volumes in response to higher relative pay. Regressions include time-zone and date fixed effects. Standard errors are clustered the MSA level.

TABLE 7: TURNOVER EFFECTS WITHIN AND ACROSS WORKERS

	Job Completion		Quits		Excellent Eval.	
	Across	Within	Across	Within	Across	Within
$1/hr	2.601***	2.164***	−2.758***	−2.321***	1.192***	0.5971*
	(0.216)	(0.219)	(0.217)	(0.209)	(0.139)	(0.119)
% of Full Effect		83.2%		84.2%		50.1%
Duration	✓	✓	✓	✓	✓	✓
FEs	✓	✓	✓	✓	✓	✓
Mean $/hr	11.19	11.19	11.19	11.19	11.19	11.19
Dependent Mean	40.59	40.59	33.56	33.56	11.07	11.07
Workers	93175	93175	93175	93175	93175	93175

Note: *p<0.1; **p<0.05; ***p<0.01 To understand how much the measured effects of pay arise from workers' behavioral responses versus sorting of workers, we examine effect sizes when the same worker faces different wages for comparable work. We look within all the warehouse jobs that the staffing agency places people in, to see at the effects on turnover that arise from a single worker facing different pay rates, and the share that arise from . Regressions include only those workers who have completed multiple jobs at the staffing agency. Regressions include occupation by commuting zone fixed effects, industry by commuting zone by time fixed effects as well as controls for expected duration as a quartic. Standard errors are clustered at the worksite-firm level.

TABLE 8: SPILLOVERS FROM THE SHIPPER'S HIRING ON OTHER LOCAL FIRMS

	Diff-in-diff		Continuous	
	Quits	Bad Ending	Quits	Bad Ending
Effect on Turnover	12.4*	8.11*	1.45*	0.8*
	(1.72)	(1.7)	(0.38)	(0.37)
Elasticity of Turnover	0.73*	0.56*	0.55*	0.35*
	(0.1)	(0.12)	(0.14)	(0.16)
Dependent Mean	28.19	24.38	28.19	24.38
Mean Pay	10.63	10.63	10.63	10.63
# Workers	16448	16448	16448	16448
# CZs	51	51	51	51

Note: *p<0.1; **p<0.05; ***p<0.01
We perform both a difference-in-differences and a continuous difference-in-differences exercise to see how the shipper's hiring affects firms hiring in the same local labor market in the same month. When pay at the shipper is one dollar higher than the outside option, workers hired into other local firms are more likely to quit or otherwise have a bad ending (e.g., be terminated for poor performance or poor attendance). To ensure that we do not put negative weights on any of our observations, we fully interact the specification with the three years in which we see treatment. Standard errors for yearly coefficients were clustered at the commuting-zone level. We then aggregate the point estimates using inverse-variance weighting and report bootstrapped standard errors.

11 Figures

Figure 1: Pay Change in Treated and Untreated Warehouses

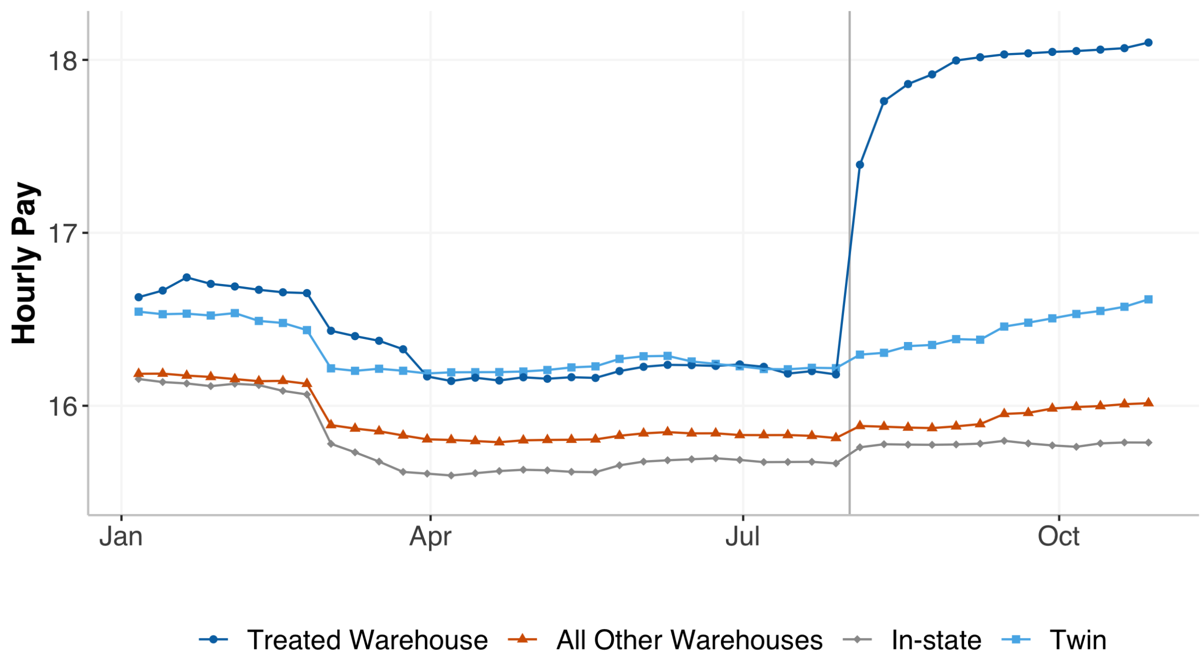

Note: We plot the average weekly pay within warehouses over the course of 2019. The grey line indicates August 2019. Average pay for all other warehouses are denoted in orange triangles, for warehouses in the same state as the treated warehouse in grey diamonds, and for "twin" warehouses that handle the same type of package in blue squares.

Figure 2: Impact of Warehouse Wage Increase on Turnover and Productivity (Number of Boxes Moved/Hour)

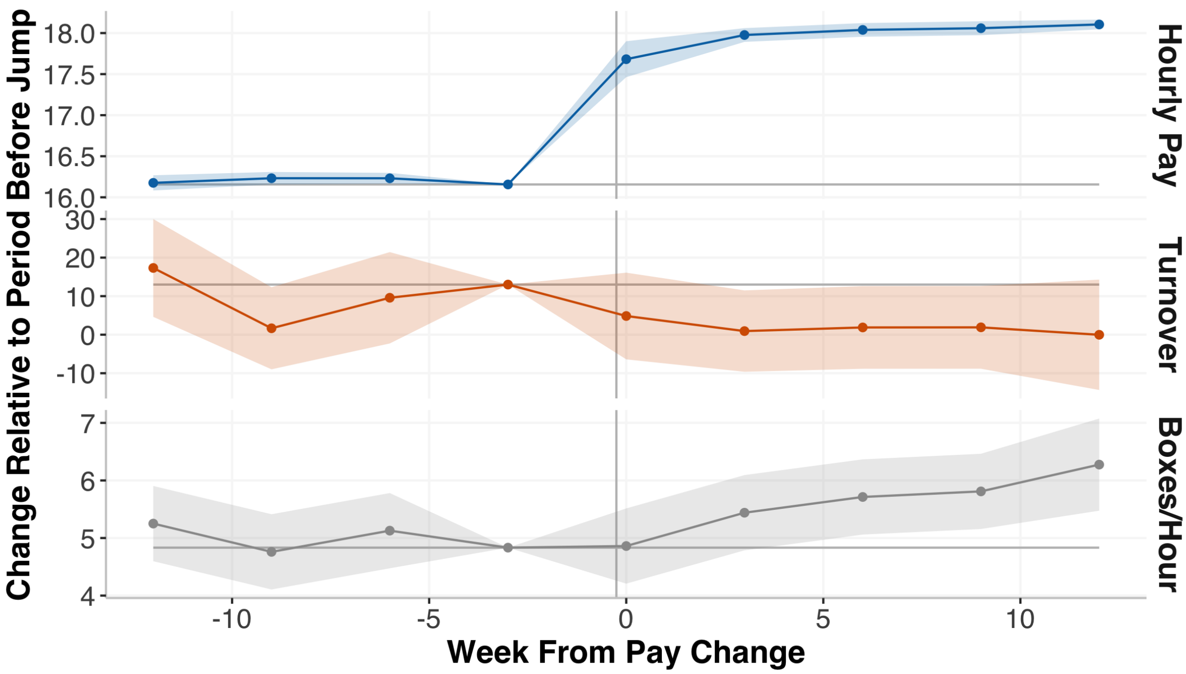

Note: This figure plots averages within three week bins around August 2019 in the treated warehouse. Panel A reports the change in average hourly pay among warehouse workers. Panel B reports the turnover rate in the warehouse. Panel C shows the average warehouse-level productivity (total boxes moved in a week/ total hour worked in a week). Standard errors are clustered at the employee- and week- levels. The shaded areas display 95 percent confidence intervals. The coefficients plotted are unscaled, in contrast to the estimates in the regressions displayed in Panel A of Tables 2 and 3, which are scaled by size of the pay jump so that point estimates reflect changes per $1/hour increase.

FIGURE 3: CHANGE IN RELATIVE PAY AMONG RETAILERS' CUSTOMER SERVICE AGENTS AND CHANGE IN TURNOVER AND PRODUCTIVITY

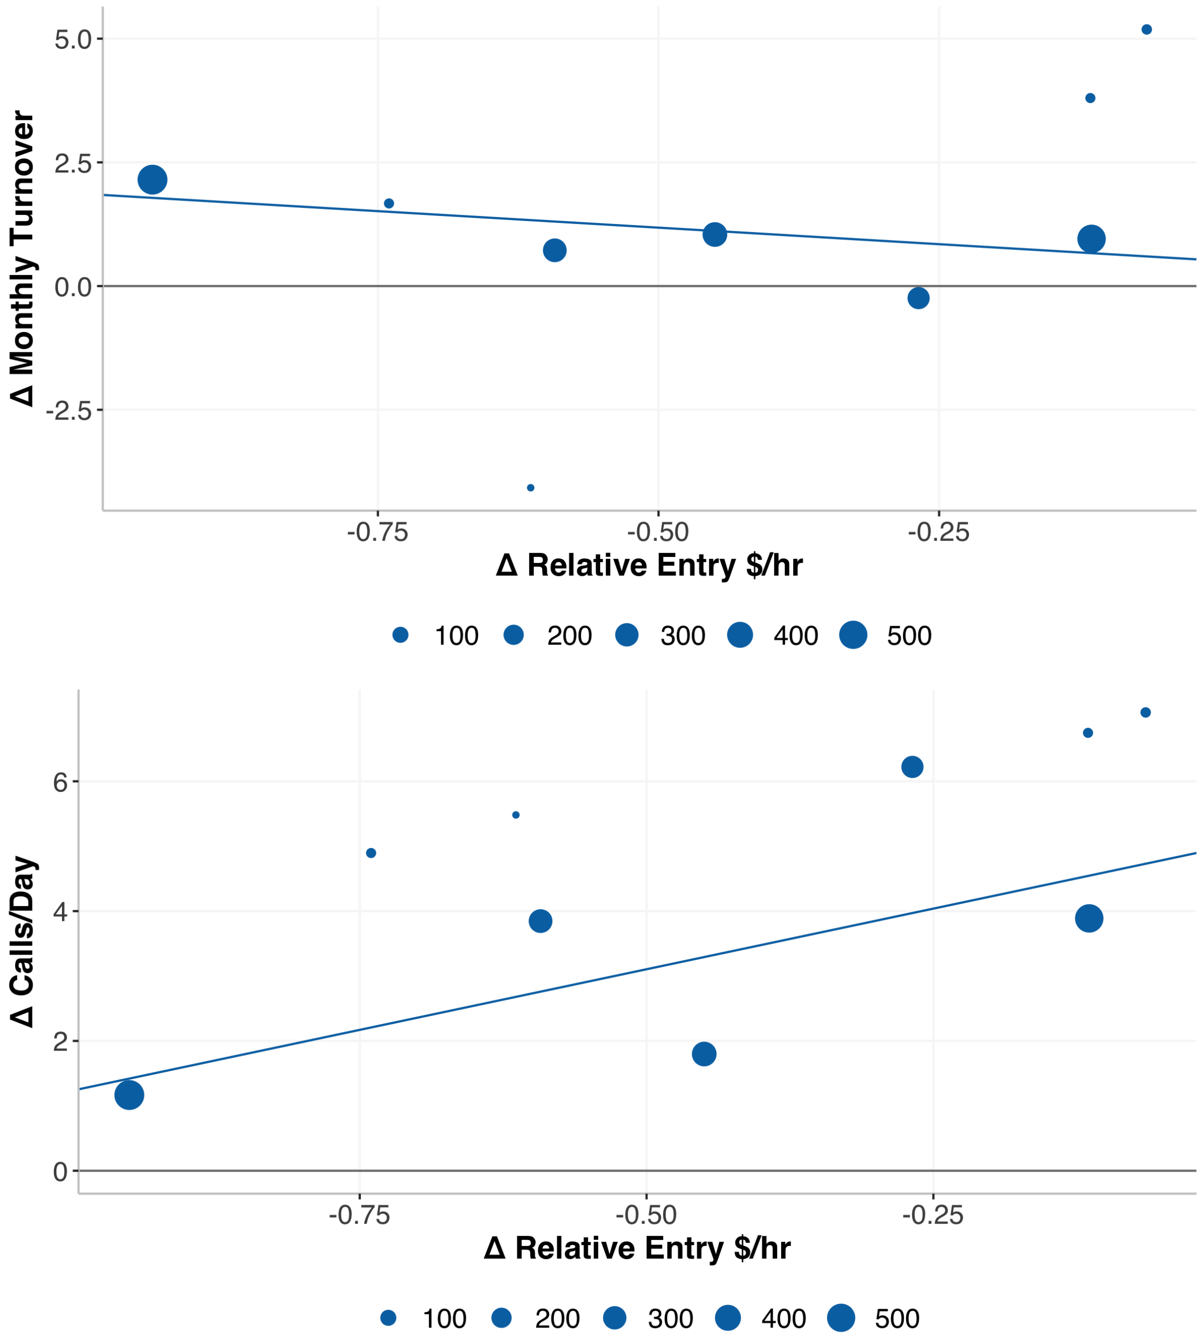

Note: To understand how the change in the gap between retailers' pay relative to the outside option affects turnover and productivity, we plot the change in relative pay in each MSA from 2018 to 2019 against the change in turnover/productivity in that MSA from 2018 to 2019. The regression form is displayed in Table 2, Panel B and Table 3, Panel B, respectively. We plot the MSAs with more than 15 customer service representatives.

Figure 4: Impacts of Relative Pay on Recruitment of Remote Customer Service Representatives

Panel A: Entry-Level Wages for Customer Service Representatives By MSA

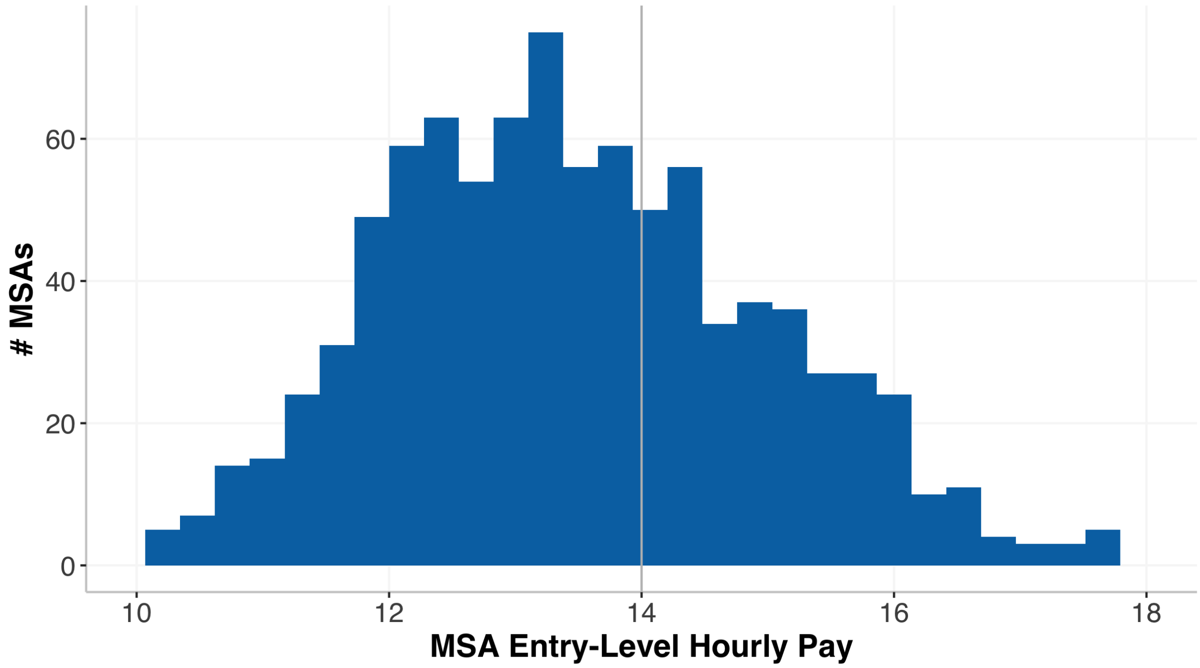

Panel B: Impact of Relative Pay on Recruits

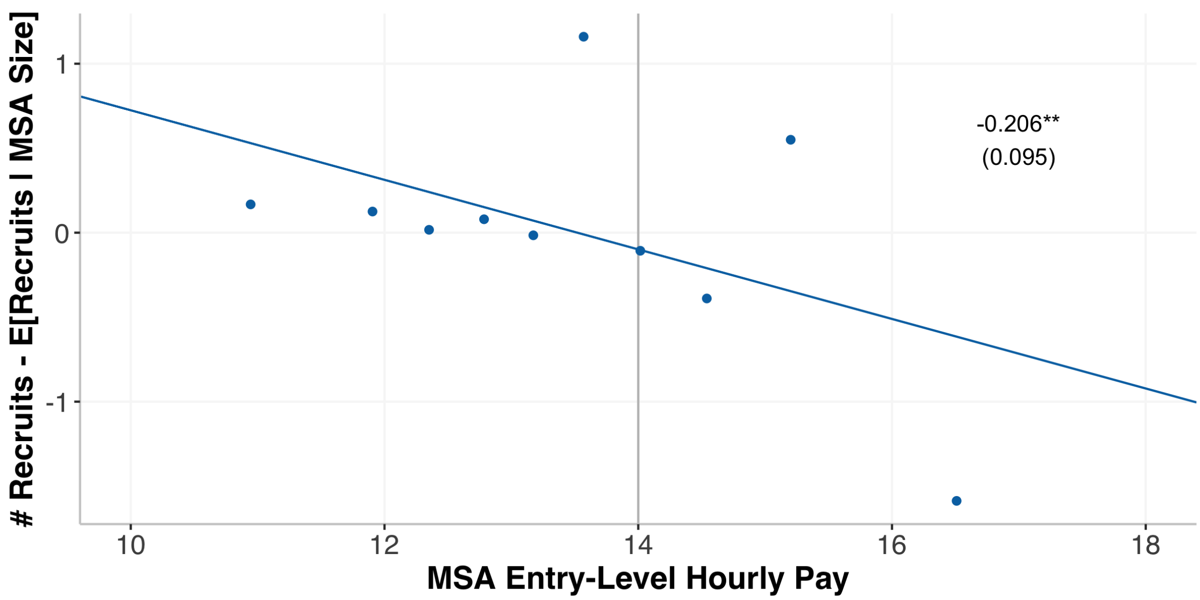

Note: Panel A plots the entry level hourly wages for customer service representatives by MSA, weighted by the number of customer service representatives in that MSA, across the United States in 2018. Panel B presents a binscatter of the number of recruits in a given MSA relative to what would be expected based only on the size of the MSA, as a function of the entry level pay for customer service representatives in the MSA. The grey, vertical line shows the retailer's offered wages; the blue line shows the regression line, controlling for the number of customer service workers in the MSA as a quartic. The standard error on the slope is shown in parentheses. The regression form is displayed in Table 4.

APPENDIX A: ADDITIONAL TABLES

TABLE A.1: DEMOGRAPHIC CHANGES IN WAREHOUSE AROUND PAY CHANGE

	Age	Female	Specialization	Share Hires
Post	0.051	−0.005	−0.075	−0.002
	(0.501)	(0.017)	(0.057)	(0.002)
Constant	36.090***	0.219***	0.792***	0.007***
	(0.733)	(0.026)	(0.061)	(0.001)
Workers	514	514	514	514
R^2	0.00000	0.00003	0.001	0.043
Note:			*p<0.1; **p<0.05; ***p<0.01	

Note: We consider the effects on turnover of a discrete pay change in a warehouse. The Field Director of the warehouse confirmed that this pay change was unexpected and did not accord with any change in work or work structure. Accordingly, we test for other changes in the warehouse around the time of the pay change. We compare the before and after period; standard errors are clustered at the week and employee levels.

TABLE A.2: SENSITIVITY OF TURNOVER EFFECTS TO TIME SPANS INCLUDED

	Monthly Turnover			Quits			Fires		
	1 Mo	2 Mo	3 Mo	1 Mo	2 Mo	3 Mo	1 Mo	2 Mo	3 Mo
$1/hour	−3.698**	−2.631**	−2.504**	−3.737***	−2.482**	−2.270**	0.280	0.188	0.061
	(1.571)	(1.307)	(1.255)	(1.132)	(1.003)	(0.957)	(1.318)	(0.687)	(0.563)
Observations	29,401	41,360	50,478	29,401	41,360	50,478	29,401	41,360	50,478
Note:									*p<0.1; **p<0.05; ***p<0.01

Note: We consider the effects on turnover of a discrete pay change in a warehouse. In the main body of the paper, we present results for a 3-month bandwidth on either side of the pay jump. Here we present robustness to other time spans. Here we present one-, two- and three month results. We do not extend the window beyond 3 months after the pay jump because we enter the holiday shipping season, which has it's own set of impacts on warehouse functioning. All regressions are scaled to show the effect of a single dollar; standard errors are clustered at the week and employee levels.

TABLE A.3: PLACEBO TEST OF TURNOVER IN IN-STATE WAREHOUSES

	First Stage	Monthly Turnover	Quits	Fires
Post	0.114	1.810	2.249*	−0.294
	(0.115)	(1.503)	(1.338)	(0.501)
Base Mean	15.66	9.3	7.5	1.24
Workers	1068	1068	1068	1068
Observations	99,178	99,178	99,178	99,178
Note:				*p<0.1; **p<0.05; ***p<0.01

Note: We perform a placebo test, exploring the change in turnover at other warehouses in the same state. Since 2 of the three other warehouses are within a 13-minute drive of the treated warehouse, if a shock to the local labor market for warehouse workers caused the effects in the treated warehouse, one would expect to see turnover decreases in these warehouses as well. While the regression in Table 2 scales by the size of the pay change, the unscaled coefficients are presented here since there is no significant first stage in our context. Standard errors are clustered at the employee and week-by-warehouse levels.

TABLE A.4: PLACEBO TEST OF PRODUCTIVITY IN TWIN WAREHOUSES

	First Stage	Boxes/Hr	Boxes/Moving Hr	Moving/Total Hrs
Post	0.207***	0.003	−0.333*	0.027**
	(0.031)	(0.142)	(0.169)	(0.013)
Pre Jump Mean	16.24	2.79	5.19	0.55
F	0	0	3.85	4.27
Observations	26	26	26	26
Note:				*p<0.1; **p<0.05; ***p<0.01

Note: We perform a placebo test, exploring the change in productivity at other warehouses that handle the same type of parcel. Since warehouses that handle the same type of parcel have similar units of productivity and since demand shocks are likely to hit them all similarly, we suspect we would see an increase in productivity in twin warehouses if it were driven by an uptick in consumer demand for large parcel goods. While the regression in Table 3 scales by the size of the pay change, the unscaled coefficients are presented here since there is no significant first stage in our context. Standard errors are clustered at the employee and week-by-warehouse levels.

TABLE A.5: CUSTOMER SERVICE DAILY LABOR SUPPLY EFFECTS

	Work Hrs	Absent Hrs	Absent Unapproved Hrs	Overtime Hrs
Entry Relative \$1/hr	−0.118	0.078	−0.036	0.065***
	(0.131)	(0.072)	(0.025)	(0.015)
Elasticity	0.26	1.52	1.33	4.11***
	(0.29)	(1.42)	(0.95)	(0.97)
FE: date-timezone	✓	✓	✓	✓
Mean \$/hr	16.02	16.02	16.02	16.02
Dependent Mean	7.14	0.82	0.43	0.25
MSAs	41	41	41	41
Workers	2871	2871	2871	2871

Note: Note: *p<0.1; **p<0.05; ***p<0.01 This table leverages stickiness of the retailer's pay to evaluate how the change in relative pay from 2018 to 2018 affect customer service representatives daily labor supply designs.

TABLE A.6: EFFECT OF RELATIVE PAY ON CUSTOMER SERVICE RECRUITMENT BY GENDER

	# Customer Service Representatives Hired				
Entry Relative \$/hr	0.013	0.010	0.008	0.008	0.015
	(0.025)	(0.009)	(0.011)	(0.011)	(0.010)
Female : Entry Relative \$/hr	0.115**	0.142**	0.156**	0.167**	0.157**
	(0.046)	(0.069)	(0.074)	(0.080)	(0.079)
Recruitment Elasticity for Men	1.86	1.41	1.09	1.17	2.14
	(3.43)	(1.26)	(1.52)	(1.57)	(1.35)
Recruitment Elasticity for Women	3.27	3.87	4.14	4.45	4.38
	(1.58)	(1.92)	(2.06)	(2.21)	(2.21)
Employment	Linear	Log	Quartic	Quartic	Quartic
Retailer Non-CSR Presence				✓	✓
Retailer n-CSR Counts					✓
F	117.55	29.19	59.53	45.41	37.22
Mean Female Recruits/MSA	0.55	0.55	0.55	0.55	0.55
Mean Male Recruits/MSA	0.1	0.1	0.1	0.1	0.1
# MSAs	920	920	920	920	920
R^2	0.204	0.172	0.264	0.272	0.280

Note: *p<0.1; **p<0.05; ***p<0.01

Note: We consider the relationship between relative pay of the employer and the number of customer service representatives ever recruited and hired in the MSA. Each observation is an MSA, excluding MSAs with on-site call-centers which have different advertising. Relative pay is the gap between the retailer's \$14/hr rate and the typical rate for entry-level workers, which we approximate with the average of the 25th and 50th percentiles of the local wage distribution. In the first column, we control only for a linear effect of the number of local customer service representatives in the MSA, whom the retailer could potentially draw from. In the second column, we instead control for a log in employment. In the third column, we control for a quartic in local employment in customer service. In the fourth column, we add indicators for the retailer having a warehouse in the MSA and the retailer having a corporate or sales' office in the MSA. In the final column, we also include controls for counts of the number of warehouse and other non-customer-service workers in the retailer in the MSA.

TABLE A.7: WORKER QUALITY SPILLOVERS FROM THE SHIPPER'S HIRING

	Predicted Excellent		Predicted Poor		New Worker	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated Month	2.009	10.050***	0.335	−2.043***	11.910**	4.343***
	(4.076)	(0.746)	(1.386)	(0.422)	(5.187)	(0.810)
Treated Commuting Zone	−3.011	−34.370***	2.656	4.081***	2.712	−29.490***
	(1.832)	(0.715)	(2.510)	(0.405)	(9.992)	(0.776)
Pay Diff. x Treated x 2016	−0.882	−2.855***	−0.014	0.045	−3.335**	−2.544***
	(1.075)	(0.287)	(0.428)	(0.163)	(1.311)	(0.312)
Pay Diff. x Treated x 2017	−1.304**	−3.423***	−0.951***	−0.504**	−4.484***	−1.070***
	(0.652)	(0.374)	(0.117)	(0.212)	(1.387)	(0.405)
Pay Diff. x Treated x 2018	−1.564	1.617***	−0.295	0.162	−4.831***	1.978***
	(1.147)	(0.515)	(0.393)	(0.292)	(1.299)	(0.559)
Year Fixed Effects	✓	✓	✓	✓	✓	✓
Mean	9.1	45.4	36.6	20.7	30.9	70.2
Workers	2,170	13,693	2,170	13,693	2,170	13,693
Observations	2,439	24,050	2,439	24,050	2,439	24,050
R^2	0.011	0.184	0.009	0.010	0.011	0.113

Note: Note: *p<0.1; **p<0.05; ***p<0.01 When pay at the shipper is one dollar higher than the outside option, workers at other local firms are less likely to be predicted high quality. Odd numbered columns display results from regressions on the Rival sample while even numbered columns display results from the state sample. All regressions include year fixed effects. Standard errors for Rival regressions are clustered at the firm level.

TABLE A.8: JOB ENDINGS IN MONTHS AND LOCATIONS WHERE THE SHIPPER HIRES

	Quits		Bad Ending	
	(1)	(2)	(3)	(4)
Treated Month	1.216	−8.398*	19.030*	−6.960
	(11.090)	(4.427)	(11.210)	(5.197)
Treated Location	−3.352	−4.272	21.950***	0.834
	(4.631)	(2.792)	(5.623)	(4.999)
Pay Diff x Treat X 2016	0.804	1.492	4.158*	1.322
	(2.044)	(1.295)	(2.136)	(1.519)
Pay Diff x Treat X 2017	0.962	0.157	5.314**	−1.077
	(2.615)	(0.646)	(2.316)	(0.684)
Pay Diff x Treat X 2018	−0.501	2.598**	4.307	3.478**
	(2.740)	(1.219)	(3.000)	(1.698)
Year Fixed Effects	✓	✓		
Mean	33	31.9	12.4	25.4
Workers	4,231	13,557	4,231	13,557
Jobs	5,147	13,557	5,147	13,557
R^2	0.010	0.007	0.011	0.010
Note:	*p<0.1; **p<0.05; ***p<0.01			

Note: *[The yearly estimates will be aggregated.]* To see whether the higher pay at the shipper leads workers to quit their existing jobs, we conduct a difference-in-differences regression, comparing the endings of warehouse moving jobs in commuting zones and months where the shipper is hiring to other locations where the shipper's rivals locate. We see no change in jobs completed nor increase in worker quits or bad endings as might be expected if workers were leaving for the higher paying job. Odd numbered columns display results from regressions on the Rival sample while even numbered columns display results from the state sample. All regressions include year fixed effects. Standard errors for Rival regressions are clustered at the firm level.

Appendix B: Additional Figures

Figure B.1: Productivity Trajectory for New Customer Service Representatives

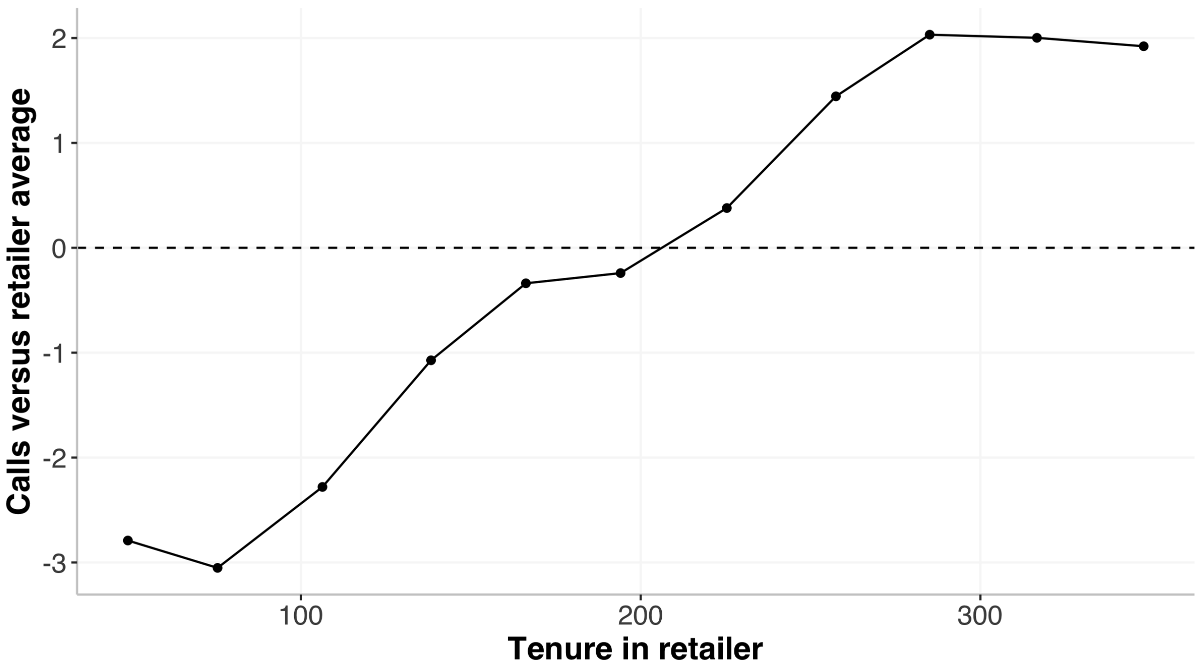

Note: This figure presents the daily calls taken by a worker relative to a typical customer service representative in the same time-zone on the same day as a function of their tenure. The x-axis plots representatives' days in the company after their training was completed. We plot the means of each month.

FIGURE B.2: DISTRIBUTION OF PAY BEFORE AND AFTER WAREHOUSE PAY JUMP

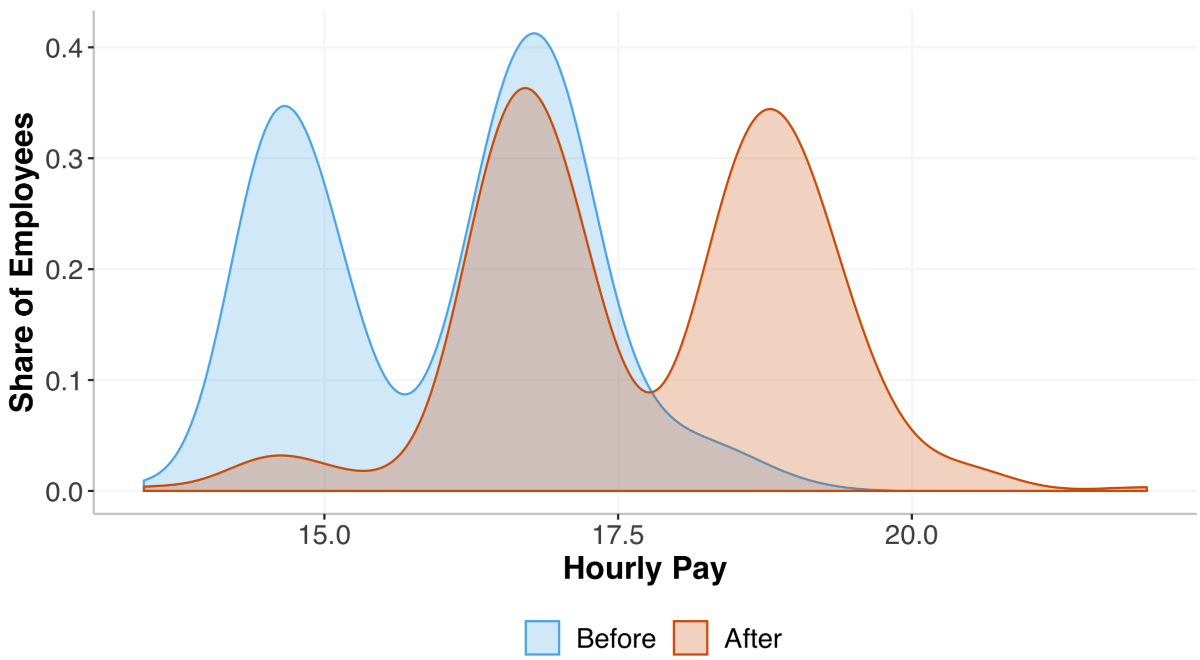

Note: This figure presents the distribution of pay among entry level warehoue workers within the treated warehouse one week before August 2019 and one month afterward. There are two sources of wage variation: a level shift if the worker works an unpleasant shift or is certified to work on specialized machinery, which generates a bimodal distribution, and wage variation based on when hired, which generates variation around these means. The distribution of pay before August had a standard deviation of 1.18; afterward it was 1.21.

FIGURE B.3: PERMUTATION TEST OF WAREHOUSE TURNOVER EFFECTS

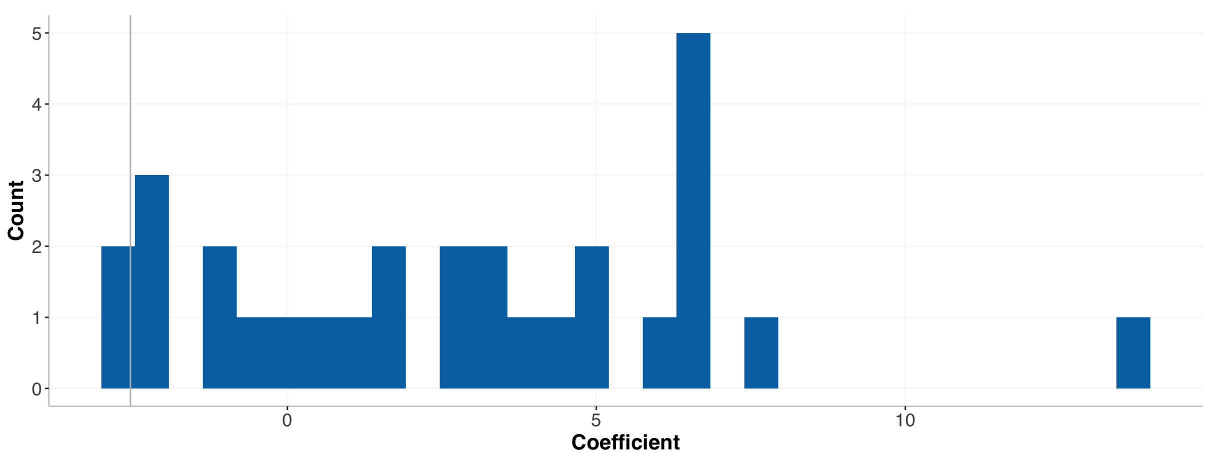

Note: We use a permutation test to explore whether similarly large decreases in turnover have been seen at other time periods in the treated warehouse. We place the date of treatment at every other week in 2019 and estimate the effect size over a three month bandwidth. We do not extend into 2018 because the holiday period is an unusual time that may be subject to other treatments. We require that the entirety of our artificial treatment window not overlap with the true post-treatment window so that we don't bias the results. For this analysis, we do not scale by the size of the pay jump since most periods do not feature a pay jump. Standard errors are clustered at the employee and week levels.

FIGURE B.4: PERMUTATION TEST OF WAREHOUSE PRODUCTIVITY EFFECTS

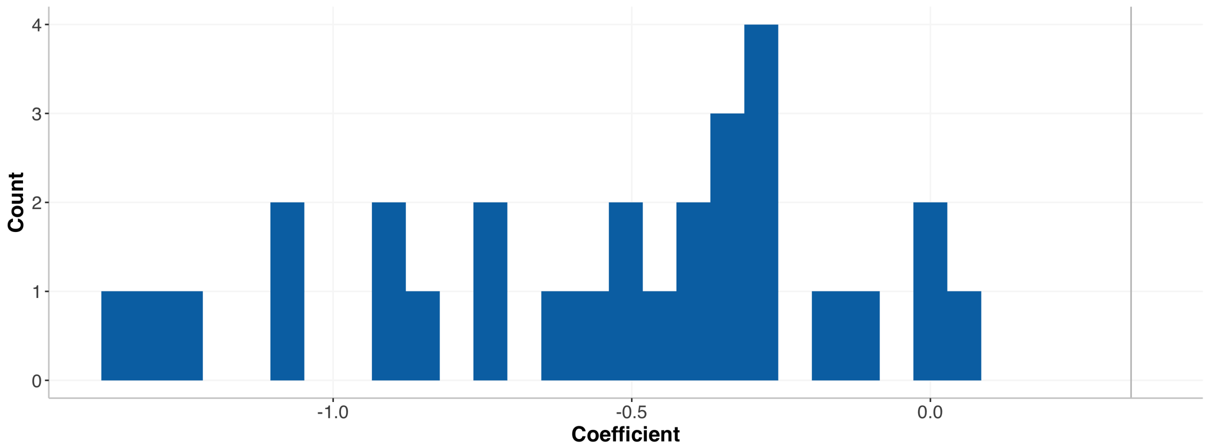

Note: We use a permutation test to explore whether similarly large decreases in turnover have been seen at other time periods in the treated warehouse. The grey line to the right shows the point estimate of the increase in boxes per hour from our main analysis. We place the date of treatment at every other week in 2019 and estimate the effect size over a three month bandwidth. We do not extend into 2018 because the holiday period is an unusual time that may be subject to other treatments rather than just the pay jump in question. We require that the entirety of our artificial treatment window not overlap with the true post-treatment window so that we don't bias the results. For this analysis, we do not scale by the size of the pay jump since most periods do not feature a pay jump. Standard errors are clustered at the employee and week levels.

FIGURE B.5: SHIPPER'S PAY RELATIVE TO THE OUTSIDE OPTION

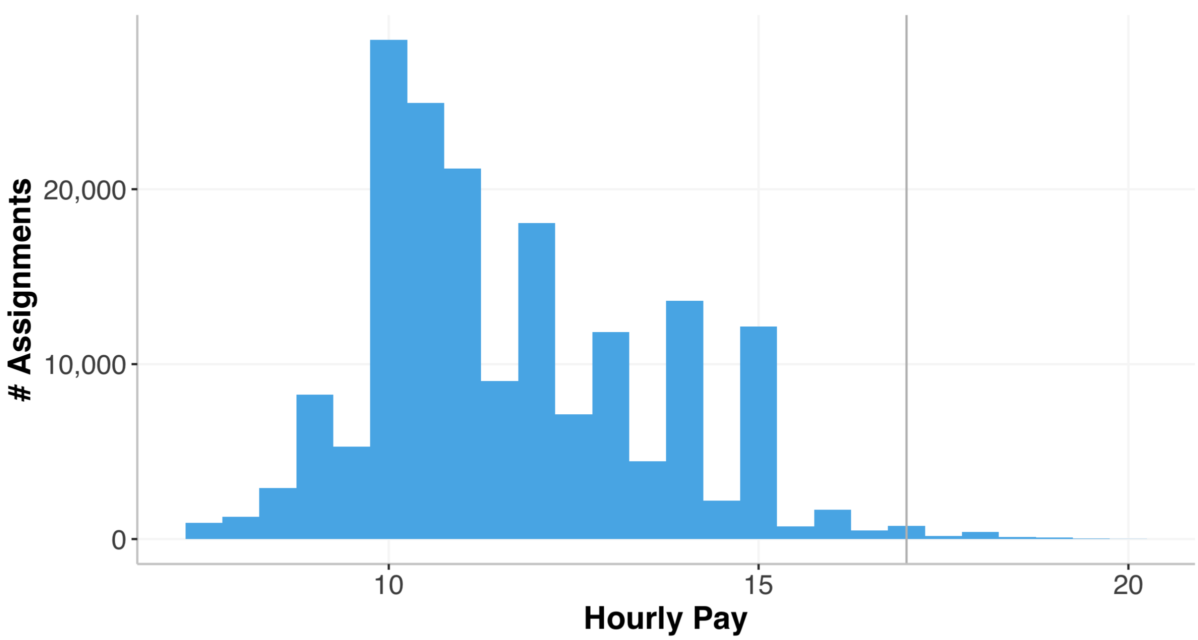

Note: The shipper pays all of it's workers the same high rate relative to the going rate in the geographic area. We leverage the variation in the local hourly rate for loader-movers to look at the effect of relative pay on performance. Above we plot the wages paid to loadter-movers in other firms in the same industry in the same commuting zones as the shipper locates. The grey line shows the wages paid at the shipper.

FIGURE B.6: AVERAGE PAY FOR WAREHOUSE JOBS IN TREATED AND UNTREATED COMMUTING ZONES

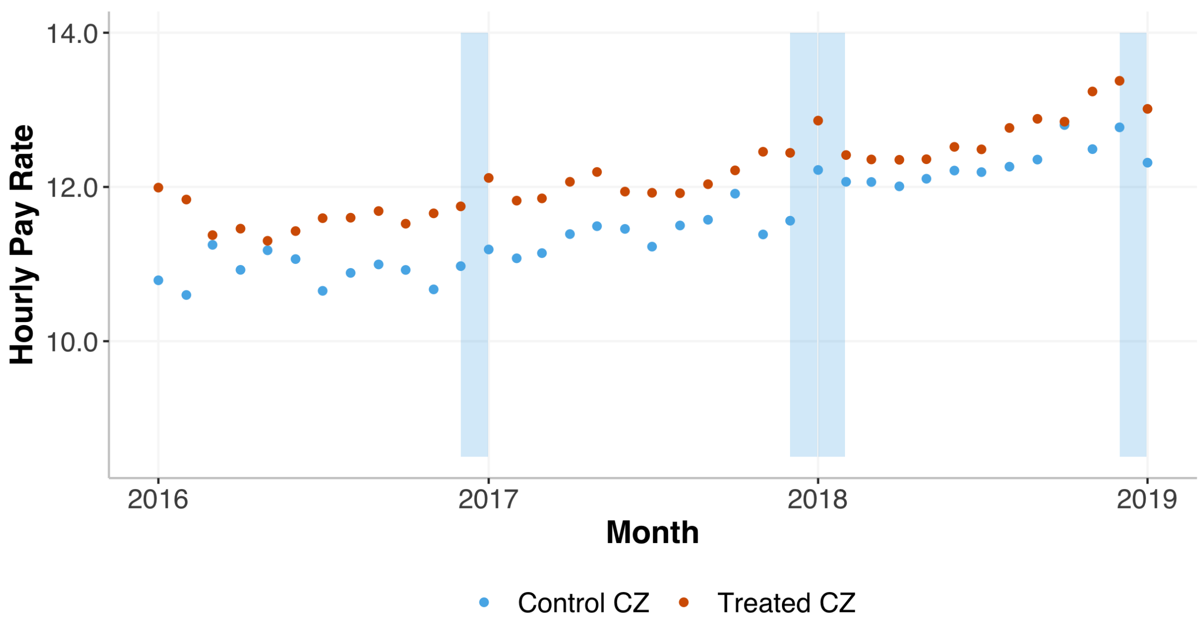

Note: We plot average pay for warehouse jobs by month that the job starts in the treated and untreated commuting zones in orange and blue, respectively. The shaded areas reflect the months when the central firm hires more than 50 workers. The central firm tends to locate in commuting zones that pay slightly more as seen by the fact that the orange dots are consistently above the blue ones. However, the trends appear to be fairly parallel throughout the time period. Also note that pay at the central firm is $17 per hour, which is considerably higher than the going for warehouse jobs in either treated or untreated areas.

Appendix C: Hiring in the Staffing Agency

For context, it may be helpful to review how hiring occurs through a staffing agency. When a firm hires through this staffing agency, they send to the staffing agency a description of the job their looking to fill and the pay rate. In select cases, the firm may ask the staffing agency for a particular worker with whom they have had a positive prior experience, but in most cases it is up to the recruiter to locate and present potential candidates. Some firms allow room for negotiation on staffer's wages, however, many refuse to negotiate on wages since they have set their advertised wages in relation to the wages of their full-time workers and they do not want to create strife.

The firms hiring through the Agency range from small, local companies to nationwide firms with hundreds of thousands of employees. While some firms seek tryout for long-term positions, many appear to be filling intrinsically temporary needs such as additional workers for holiday rush seasons. Indeed, only 7.5 percent of workers transitioned to a more permanent placement with the client firm during our sample period. Of the workers who took a job through the staffing agency, 64 percent did not return in our period for a second job. But for a notable minority of workers, the Agency provided continuing stints of work: 5.5 percent of workers take at least five jobs with the Agency and are employed for a total of 263 days on average.