THE PAYOFFS OF HIGHER PAY

ELASTICITIES OF LABOR SUPPLY AND PRODUCTIVITY WITH RESPECT TO WAGES

Natalia Emanuel · Emma Harrington[1]

Abstract

While paying less would seem to save firms' money, we ask whether paying more can payoff in higher productivity. We analyze this question in the context of a major online retailer that hires both warehouse and call-center workers. We estimate the returns to higher pay in settings where the retailer active chooses to raise pay and passively lets pay fall below the market. The nominal wage rigidity of the retailer creates two types of quasi-random variation in pay: (1) when wages get "unstuck," they often jump discretely; (2) when wages remain "stuck," their relative value depreciates more quickly in cities with faster wage growth. Using the resulting two designs, we find elasticities of separation between 3 and 4.5 and elasticities of output between 1.1 and 1.2 — in terms of boxes moved and calls taken. When the retailer actively chooses to raise pay and when it passively lets wages fall below the market, the returns to higher wages exceed one — in other words, higher wages pay for themselves.

Persistent gaps in pay exist across predominantly male and female jobs, especially those that attract women in their child-bearing years. We ask whether these gaps could be driven by differences in firms' incentive to raise pay. We find the returns to increasing pay differ by workers' gender and age: while effects on productivity are comparable, turnover effects are not. The turnover elasticity is 4.2 times higher for workers under 30 than those over 30 and 1.9 times higher for male workers than female workers in call-centers.

Finally, we use data from a staffing agency to show that 50-80 percent of the effect of pay on turnover arises from workers' behavioral responses, with the remainder attributable to sorting of better workers to higher-paying firms. We find this sorting generates negative spillovers on other firms in the local labor market: firms hiring at the same time as a high-paying firm experience higher turnover.

---

[1]Contact: Harvard University, 1805 Cambridge Street, Cambridge, MA 02138, emanuel@g.harvard.edu · eharrington@g.harvard.edu. We thank Isaiah Andrews, Claudia Goldin, Lawrence Katz, Nathan Hendren, Edward Glaeser, Jeff Liebman, Amanda Pallais, Lawrence Summers, and participants of the Public Finance and Labor Economics Workshop at Harvard for helpful comments. We appreciate input from Zoe Cullen, Jerry Green, Jeffrey Miron, Matthew Rabin, and Andrei Shleifer. We are grateful to our colleagues, Jenna Anders and Augustin Bergeron, as well as Alyssa Bilinski, Justin Bloesch, Dev Patel, Jon Roth, Gregor Schubert, 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 [Natalia] 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.

We seek to understand both the extensive labor supply elasticity — namely the elasticities of recruitment and turnover — as well as the intensive labor supply elasticity — the productivity achieved — in the context of routine workers. Assembling estimates of the range of responses helps shape an understanding of the effect of higher pay throughout various different points in the firm's interactions with a worker.

Our paper provides four new pieces of evidence. First, we offer estimates of recruitment elasticities as distinct from our estimates of separation elasticities. Manning (2003) shows that steady-state elasticity of the labor supply facing a firm is the combination of the job-to-job recruitment elasticity and the separation elasticity. Since the latter can be difficult to estimate, he argues that the steady-state elasticity can be approximated as as twice the separations elasticity. We estimate recruitment elasticities by leveraging a firm that offers the same wage for all remote workers across the USA, and comparing how the wage relative to the local outside option affects workers' attraction to the job. We do, indeed, find that in this context, the size of the recruitment and separation elasticities are similar in magnitude.

Second, we document productivity elasticities for warehouse workers and customer service representatives, using the objective metrics that the firm uses to evaluate these workers' output. This is especially important when calculating the cost effectiveness of higher pay because the returns associated with reduced turnover are dwarfed by the returns associated with increased productivity. While we find that for this retailer, raising wages for these workers may be profitable, this result is heavily reliant on the elasticities of productivity that we estimate. We leverage both a quasi-random increase in pay at a large online retailer's warehouse to see the effect on boxes moved in that warehouse. We also use the fact that the retailer kept wages for customer service representatives constant for several years while local outside options changed to assess the effect of relative pay on number of calls handled.

Third, we estimate how much of the reduced turnover is due to sorting better workers to higher paying firms and how much is due to the behavioral response of the worker. We use staffing agency data to assess how much of the improved job completion is associated with higher pay persists when we look within the same person placed in multiple jobs with different pay. We document that 20-50 percent of the effect arises due to selection of workers. We argue that this is consistent with efficiency wages. Accordingly, we also document negative spillovers: when a high-paying firm is filling positions, other firms that are hiring for the same type of worker see higher turnover rates.

Finally, we are able to estimate elasticities separately by gender, revealing that women are less responsive to pay then are men. This finding differs from recent work (Caldwell and Oehlsen, 2018), but highlights that differences in elasticities facing the firm may underpin some of the gender pay gap.

Warehouse workers and customer service representatives are particularly well-suited to studying the effect of pay for several reasons. First, these jobs produce objective metrics of performance that the firm cares about. Second, these are precisely the jobs that many managers consider routine and with low-training requirements for which finding a replacement worker may be easy. So understanding the impacts of turnover and how sensitive it may be to pay is worthwhile. Third, of jobs available to those without college degrees, these are relatively well-paying jobs and are comparatively abundant. Finally, the jobs studied tend to be in fairly thick markets. Large warehouses are often located within commuting distance of large cargo hubs and thus all logistics firms are drawing on the same pool of talent. The areas where customer service workers are hired by our retailer also feature relatively high demand for customer service workers. We find upward-sloping labor supply curves as well as gains in productivity from higher pay, which are particularly surprising in thick labor markets, where we might expect to find the possible benefits of higher pay have been realized due to intense competition.

This work contributes to existing scholarship on the relationship between pay and performance, labor supply elasticities, and labor market concentration.

Estimating the relationship between pay and worker output builds on an important literature of efficiency wages. Under the theory of efficiency wages, employers pay a premium above the market to give the worker an incentive to try to keep their job (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 (Katz, 1987). Raff and Summers (1987) use Ford Motor Company as a case study noting how above-market wages can slash turnover rates and elicit greater effort from workers. Krueger and Summers (1988) and Orszag and Zoega (1996) explore whether intra-industry pay differentials can be attributed to efficiency wages. Cappelli and Chauvin (1991) use centralized, uniform wage-setting in a multi-plant firm to document that a greater pay premium leads to lower rates of disciplinary infractions, which is consistent with our results. By randomizing the wages of public sector job adverts in Mexico, Dal Bó et al. (2013) capture the selection effect of higher wages. We contribute to this literature in three ways: first, we bring evidence about workers' on-the-job performance using the metrics that firms use to measure productivity. Second, we compare how the same worker responds to different pay rates. Third, we capture how higher wages at one firm may impact other firms' access to talent.

Our estimates contribute to an intimately related literature on labor supply elasticities and monopsony power. Manning (2003) has laid out an important theoretical framework, arguing that the elasticity of turnover and the elasticity of recruitment combine to form a measure of monopsony power. Several authors estimate the labor supply elasticities facing firms that employ workers doing seemingly interchangable jobs such as nurses and nurse aides (e.g., Sullivan, 1989; Staiger et al., 2010). Webber (2015) uses linked employer-employee data to draw a connection between between firm labor supply elasticities and workers' earnings. Dube et al.

(2020) estimate labor supply elasticities around 0.1 even in an online market place, suggesting that employers wield monopsony power even there. We likewise find upward sloping labor supply curves, even for warehouse workers (where the labor market is particularly thick), remote workers and temporary workers where one might conjecture that workers are highly mobile and thus labor supply elasticities are flat. We further contribute by estimating the elasticity of recruitment, which is rarely captured. We find a large positive elasticity of recruitment with respect to wages, which significantly increases the overall labor supply elasticity that Manning addresses. We also contribute to the literature exploring gender differences in labor supply. While Caldwell and Oehlsen (2018) finds no gender difference in labor supply among Uber drivers, we do find gender differences in our context, which may reflect differences in who elects into each occupation.

Additionally, a literature explores the spillovers when one firm raises wages. 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. This paper contributes to the examination of spillovers by documenting short-term effects on the quality of work done at rival firms, which helps explains rival firms' need to raise wages. Our evidence from a staffing agency of the job completion spillovers to other firms microfounds these cross-price elasticities since our estimates explain the need to raise wages in rival firms.

Finally, our work highlights pricing strategies that may or may not be perfectly optimal. This relates to work on X-inefficiency, since the firms may not always perfectly minimize their costs especially in the absence of competitive pressure (Leibenstein, 1966; Frantz, 2013). We study warehouses, which cluster at key geographic hubs creating intensive competitive pressure. Nevertheless, we find positive returns to higher wages in this context in case studies one and three. While not the focus of this paper, our empirical strategy in case studies two and three highlights that large firms engage in uniform wage-setting, a tendency parallel to retail chains using uniform pricing despite variation in local demand and competition(DellaVigna and Gentzkow, 2019).

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 the elasticities of recruitment, retention, and productivity with respect to pay. In Section 6 we conduct a cost-benefit estimate. We explore the degree to which selection versus incentives contribute to our results and accordingly whether higher pay has negative spillovers on local firms in Section 7. We conclude in Section 8.

1 CONCEPTUAL FRAMEWORK

Our conceptual framework aims to characterize the sufficient statistics for firms' wage-setting decisions. Intuitively, the firms' wage-setting decisions will be determined by the effect that higher pay has on retention and worker productivity.

A firm requires L units of effective labor. The amount of labor done at the firm depends on the number of workers, N, and the effective labor supplied by each worker, ψ. Together, we have $L = \psi N$. We allow effective labor to be a function of the wage paid by the firm, w, and the tightness of the local labor market, z. We assume ϕ is increasing in w and decreasing in z, and empirically evaluate this assumption. We allow the firm can choose the number of workers to hire independently of the wage it sets.

The firm needs to pay turnover costs T for each new worker, which may include recruitment and training costs. We assume that the wage has no impact on the number of workers a firm can recruit or the costs of recruiting them. The quit rate, $q(w, z)$ is a decreasing function of wages.[2] The turnover costs are given by $TNq(w, z)$.

Firms thus seek to minimize the cost of L units of effective units of labor:

$$\min_{w} wN + TNq(w, z) \text{ s.t. } \phi(w, z)N = L$$

If we substitute the constraint in for N, take the first order condition and solve we arrive at an equation for wages (Appendix 10 shows the entire derivation). We define the elasticity of quits such that when wages go up, quits go down, which entails flipping the sign of $\epsilon_{q,w}$. Thus we arrive at:

$$w = \frac{Tq(w, z)(\epsilon_{\phi,w} + \epsilon_{q,w})}{1 - \epsilon_{\phi,w}}$$

The optimal wage is increasing in the elasticity of effort with respect to the wage from both the denominator — which reflects the fact that fewer people need to be paid each period if each of them can do more work — and in the numerator — which reflects the fact that fewer people employed means fewer quits. We also see that the optimal wage is increasing in the elasticity of quits with respect to the wage since when this elasticity is greater, the firm will save more in turnover costs by increasing wages.

We empirically estimate the $\epsilon_{\phi,w}$, the elasticity of productivity with respect to wages; $\epsilon_{q,w}$, the elasticity of quits with respect to wages, $q_z(w, z)$ the relationship between outside options and quits, and bring internal estimates of T, the turnover cost.

[2]The quit rate derives from a poisson process. For any of the N workers, there is a λ probability of quitting at any given moment. When they quit, the firm must pay the turnover cost.

We explore how the outside option, z, impacts wages. We make the simplifying assumption that the outside option impacts wages only through turnover, as this might be the most obvious metric to many firms.

$$
\begin{aligned}
w &= \frac{Tq(w,z)(\epsilon_{\phi,w} + \epsilon_{q,w})}{1-\epsilon_{\phi,w}} \\
\frac{\partial w}{\partial z} &= \frac{q_z T(\epsilon_{\phi,w} + \epsilon_{q,w})}{1-\epsilon_{\phi,w}} + \frac{Tq\frac{q\cdot q_{wz} - q_w q_z}{q^2}}{1-\epsilon_{\phi,w}} \\
&= \frac{q_z T(\epsilon_{\phi,w} + \epsilon_{q,w})}{1-\epsilon_{\phi,w}} + \frac{Tq_{wz} - \frac{T}{q} q_w q_z}{1-\epsilon_{\phi,w}}
\end{aligned}
$$

The sign of the first term is positive since, $q_z > 0$, quits increase when the outside option increases as long as $0 < \epsilon_{\phi,w} < 1$.[3] The second term may also be considered positive since it is not unreasonable to think that $q_w z > 0$ if quits decrease as wages increase, rapidly at first and more slowly later, showing a slope that is strongly negative then more weakly negative. Finally, the third term is negative since $q_w < 0$ and $q_z > 0$. Thus overall we find that wages increase as the outside option increases.

2 Data

We use data on three case studies from two large firms. We first outline the data and then detail our empirical approaches.

Online Retailer Data. The first source of data is an Online Retailer, where we focus on data from their warehouse workers and their customer service representatives. These data begin in 2018 and go midway through 2020.

We study the 8,597 individuals who work in warehouses and the 4,799 customer service workers employed by the retailer. We have workforce data, which captures how many individuals were employed to work on a given day. This information is segmented based on workers' level of employment (e.g., entry-level), and location. The dataset also contains information about workers' job titles and pay rates. For warehouse workers, we can also see the shifts they worked and whether they have any skills (e.g. forklift certification) that warrant a pay premium over the base rate.

Warehouse productivity is measured in boxes moved per hour and boxed moved per moving hour – an hour in which workers are indeed moving boxes, rather than eating lunch or attending a team meeting. Our data consists of weekly measures of productivity for each of

[3]Given that we defined the elasticity of quits with respect to wages in a manner that flipped the sign, the sign of the first term may not be self-evident. To see that it is positive, recall that the term was $T(\epsilon_{\phi,w} - \frac{q_w w}{q})$ and that $q_w < 0$.

the retailer's warehouses. The retailer has distinct warehouses to handle large packages (e.g., a refrigerator or a couch) and smaller packages (e.g., a book or waterbottle). Boxes moved per hour will be higher at warehouses that handle smaller packages. Since the treated warehouse handles larger packages, it moves fewer boxes per hour than the average in the retailer.

Customer Service productivity metrics include information about the number of calls each representative handles each day, as well as the average customer satisfaction reviews left for them that day. Customer service representatives handle incoming calls from customers, potentially inquiring about a delivery, a return, or damaged product – these representatives do not make outgoing calls or handle incoming sales requests.

When a customer calls into the retailer, her call will be routed to a representative working at that time who has the skills that appear to be necessary to handle her concern. Calls that appear more challenging — like those about damaged products or refunds — are routed to representatives who have leveled up to more senior roles, while more routine calls — like those about delivery timing or the company's return policy — are routed to entry-level representatives. Representatives are fairly autonomous but are nonetheless organized into teams of about a dozen representatives all of a common skill-level, who share a single manager and are compared to one another for the purposes of quarterly performance bonuses.

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 use their information on granular occupation- and labor-market-specific data on wages and labor supply to guide their short-term decisions. In this regard, even if the data we use from Emsi is biased, it nevertheless reflects the local outside option as understood by firms.

Staffing Agency Data. The second dataset comes from a large staffing agency, capturing every person placed on an assignment in the US from 2016 through to the end of 2018.

The staffing agency places workers in over 222,000 warehouse jobs between 2016 and 2018. Because we see the same worker in several jobs, the staffing agencies offers a valuable opportunity to estimate how much of pay's effect on retention arises from attracting better workers versus from increasing incentives for the same worker. Further the nature of the staffing agency data allows us to ascertain whether firms were satisfied with their workers, lending us a measure of employer satisfaction. Finally, because we see many firms hiring for the same jobs, we can estimate the spillovers of one firm's wage setting on other firms' turnover. Appendix 10 includes additional information about how individuals are placed in jobs at the staffing agency.

3 Higher Pay Boosts Recruitment

Higher pay relative to the local outside option attracts more workers to the retailer. When the retailer's advertised wages are $1/hour higher than the local outside option, they recruit 23-30 percent more employees in the metropolitan statistical area (MSA), reflecting a recruitment elasticity between 3.2 and 4.2.

The online retailer hires entry level remote customer service representatives at $14/hour regardless of where the job-seeker resides. Job-seekers in low-paying places may more frequently find the $14/hr rate attractive relative to the local alternatives and thus apply for and accept jobs at the retailer at higher rates. Since the retailer does not care where its customer service representatives are located, this translates into more customer service representatives in MSAs where the retailer's uniform wage is more competitive.

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.

When considering the number of recruits in a metropolitan statistical area (MSA), we consider the MSA-level specification:

$$Y_m = \beta_0 + \beta_{\$,\text{Uniform}}(\text{Entry Relative Wage})_m + \epsilon_m. \tag{1}$$

$\beta_{\$,\text{Uniform}}$ reveals the relationship between relative pay and recruitment, holding fixed the nature of the work. For $\beta_{\$,\text{Uniform}}$ to offer an unbiased assessment of the effect of the retailer raising its own wage, relative pay must be orthogonal to the inherent attributes of an MSA's pool of available workers. 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. By contrast, this assumption does not allow the pool of workers in Lufkin to differ from pool of workers in Dallas and thus the outcomes of recruited workers to differ for reasons other than retailer's posted pay.

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 (see Table 2, Panel A). This translates into an elasticity of recruitment[4] with respect to the wage of between 3.2 and 4.2. When customer service representatives are considering different options at the recruitment stage, their decision-making seems heavily swayed by relative pay.

Our evidence suggests that women are more elastic at the recruitment stage than are men. While we have less power than one might wish when we split our sample by sex, we nevertheless find that in areas where the retailer's pay is higher than local entry-level pay, there are 28 percent more female recruits and 15 percent more male recruits (see Table 2, Panel B). Our estimates suggest that women have recruitment elasticities between 3.2 and 4.4 whereas men have recruitment elasticities around 1.1 and 2.1, though we are underpowered to distinguish men's elasticities from zero.

4 Higher Pay Increases Retention

Higher pay decreases turnover within firms. This is important because turnover is costly, even for workers in jobs that seem relatively routine and do not require an advanced degree. 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, see Figure 1. On average, each warehouse loses 2.8 workers per week.

It takes a new customer service representative about 6 months to reach the calls volume of the average customer service representative, who is answering calls on the same day within the same time-zone. As illustrated in Figure 2, new representatives — who have just finished their 3 weeks of formal training — answer nearly 3 fewer calls per day – the equivalent of working one fewer hour per day for the firm. 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 reps is more valuable than retention of junior ones because they will walk away with more human capital accumulated in the firm. Thus, our findings that pay primarily affects the turnover of senior representatives is particularly meaningful for the retailer's bottom line.

4.A Higher Pay Reduces Turnover

We show that higher absolute and relative pay reduces turnover in three contexts. In the first case, we look at the effects when the retailer raised pay at a single warehouse with quasi-

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

random timing. In the second case, we leverage the fact that the retailer kept it's pay constant for customer service representatives over a two year period, so we can see how the changing outside option affects the retailer's relative pay. This case study is useful because it allows us to look at heterogeneity by demographic group. In the third case, we use a firm that hires warehouse workers at a fixed price across the US; again, we can look at the advantages conferred by wages higher than the outside option. This example is particularly useful in that the staffing agency data allows us to decompose how much of the effect arises from selection and how much from incentivizing the same worker, which we explore in Section 7.

Retention in a Warehouse. When the retailer raised wages in one warehouse quasi-randomly, fewer workers left that warehouse.

In late July 2019, average pay was $16.20/hour. One week later, it had increased to $17.39/hour and by the first of September 2019, it was solidly at $18.00/hour, an 11% increase in pay over the course of a month. Figure 3 depicts the pay bump at the treated warehouse along with relative pay constancy at other samples of warehouses.

The timing of the pay change is quasi-random. The retailer's local Field Director shared that the pay bump arose out of long-standing concerns about high turnover at this warehouse in particular. Indeed, in the quarter before the pay change, turnover at the treated location was nearly twice as high as in other warehouses. 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 grueling given that this warehouse handles larger parcels than other warehouses for the retailer nearby. The Field Director further confirmed that the nature of the work did not change around the pay jump and it did not coincide with consumer holidays that may affect work intensity. Thus we consider that the treated warehouse differs in important ways from other warehouses but the timing of the pay jump is essentially random.

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.

The bump in pay occurred throughout the distribution of Level-1 workers at this warehouse and thus did not impact worker' dynamic incentives to strive for promotions.[5] Figure 4 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, reflecting a 2.4% in-

[5] 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.

crease in the spread of wages.

We use two empirical strategies to understand the changes resulting from this pay change: an interrupted time series and a difference-in-differences strategy. In the interrupted time series, we compare the turnover and productivity before and after the pay change. In the difference-in-differences approach, we compare the change in the treated warehouse to the change in other warehouses at the retailer. When considering retention, the difference-in-differences specification compares to those warehouses that are in the same state as the treated warehouse. This group of warehouses is situated in the same thick labor market: two of the three in-state warehouses are located within a 13-minute drive of the treated warehouse. These warehouses are located right by a major cargo airport and many large companies, including Amazon, have extensive warehouse real estate in close proximity. Thus, we believe they are a particularly suitable comparison group.

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 first stage is

$$\$_i = \alpha_0 + \delta \mathbb{1}_{\text{Post}} \cdot \mathbb{1}_{\text{Treated}} + \nu_i$$

and our second stage is

$$Y_i = \alpha_1 + \beta_\$ \hat{\$}_i + \beta_{\text{Post}} \mathbb{1}_{\text{Post}} + \beta_{\text{Treated}} \mathbb{1}_{\text{Treated}} + \epsilon_i$$

where Y_i reflects the potential outcomes (turnover and productivity), $\mathbb{1}_{\text{Post}}$ is an indicator for whether the individual day occurs after the pay change and $\mathbb{1}_{\text{Treated}}$ is an indicator for whether the individual works at the treated warehouse.[6] $\beta_\$$ is our parameter of interest.[7]

Because our data includes daily observations of each worker, we cluster our standard errors at the individual level to handle serial correlation. While there may be correlation within warehouses, since we only have one treated warehouse, we cannot cluster at the warehouse level. Where possible, we cluster at the state level, which will yield conservative estimates. For the other estimates, we thus turn to permutation test. We consider all the times between March 2018 and June 2019 that the pay increase, excluding the time after the pay jump so that we do not include post-treatment outcomes.

To test the parallel trends assumptions necessary for our difference in differences approach, we

[6] Clearly, in the interrupted time series specification, $\mathbb{1}_{\text{Treated}}$ is uniformly one so the term drops out and $\mathbb{1}_{\text{Post}}$ is subsumed in the intercept.

[7] Note that we run this as a classic difference-in-differences model since Imai and Kim (2019) point out that multi-period difference-in-differences estimators put negative weight on some of the observations. This is a suitable approach because we have no reason to believe that conditions at the other warehouses change substantially.

look at the significance of leads and lags in our model as suggested by Pischke (2019):

$$Y_i = \gamma_w + \lambda_t + \sum_{j=-4}^{3} \beta T_{wt+j} + \epsilon_{iwt}$$

where γ_w reflects fixed effects for each warehouse, λ_t reflects fixed effects for the month, and T_{wt} is an indicator for whether the observation is treated in time t. All of our leads are statistically indistinguishable from zero. One can also see this, visually, in Figure 5.

Kahn-Lang and Lang (2020) argue that in addition to testing parallel trends, any differences in levels must also be explained in the context of the data. As noted above, we would expect to find a difference in the levels of turnover because this is the underlying problem with the treated warehouse. We should also expect to see a difference in levels of productivity as measured in units moved per hour based on the type of item moved in the warehouses.

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 this measures an increase of pay off of $16.20/hour, our point estimate reflects an elasticity of turnover of 3.03 (see Table **??**, Panel A). Quits 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. When we compare to turnover rates at other in-state warehouse, the treatment effects are somewhat larger: a 28 percent reduction in turnover, reflecting an elasticity of 4.6, and a 35 percent reduction in quits, reflecting an elasticity of 5.77. The specifications noted here capture three months before and after the pay change. Appendix Table **??** show the effects across several bandwidths.

Retention 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 HR data begins. The stickiness of the retailer's pay contrasts sharply with the changing nature of representatives' outside options: local pay increases over this period, 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 a particular MSA. This strategy allows us to difference away any fixed disparities in productivity across MSAs, while accounting for the general trends within the retailer.

Particularly, we consider the first-difference specification:

$$\Delta Y_m = \alpha_0 + \alpha_{\$,\text{Sticky}} \Delta(\text{Entry Relative Wage from 2018 to 2019})_m + \zeta_m. \tag{2}$$

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:

$$Y_{imt} = \alpha_0 + \alpha_{\$,\text{Uniform}} \Delta(\text{Entry Relative Wage})_m \cdot \mathbb{1}[\text{Year} = 2019] + \mu_t + \mu_{MSA} + \zeta_{imt}, \tag{3}$$

with standard errors clustered at the MSA-level. This specification asks how the change in the retailer's relative pay between 2018 and 2019 relates to daily call volume and other outcomes of interest in 2019 after accounting for the fixed effect of the MSA and the general changes in the retailer over time. 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 accounts for the number of observations that contributed to our inferences about changes at an MSA-level. In Appendix Table **??**, we also present the MSA first-differences (specification 2), weighted by the geometric mean of representative-days across the two years to reflect the precision with which we estimate the differences in outcomes.

Our coefficient of interest is $\alpha_{\$,\text{Sticky}}$ — which reflects the relationship between a \$1/hr change in relative pay and the parallel change in the MSA's turnover or productivity. Extrapolating from $\alpha_{\$,\text{Sticky}}$ to the effect of the retailer raising its own pay requires weaker assumption than under the national wage-setting design. Particularly, the sticky-pay strategy allows fixed characteristics of the pool of customer service representatives to vary from place to place, since these discrepancies will be differenced away, when looking at changes rather than levels. This strategy still requires that changes in the pool of customer service representatives are 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 such a short horizon, it seems more plausible that fluctuations in the demand for customer service representatives would drive changes in wages than changes in the supply, in terms of either quantity or quality.

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.

We drop the 3.75% of reps (1.82% of days) with missing wage information. We further drop the 1.47% of reps (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 reps (982 of 4551). We also exclude 1424 reps (37.3%) hired in 2020, since our outside option information from Emsi is only available for 2018 and 2019. We finally exclude 170 reps (7.1%) in MSAs with hires in only one of the two years.

Changes in relative pay translated to changes in turnover. 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 — 1.2pp off an average of 4.3 percent, reflecting an elasticity of 4.5 (see Panel B of Table 3 for the individual-level analysis and 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 3 Panel B suggests that fires are especially sensitive to relative pay. The effect on fires is consistent with managers' expectations for workers not fully keeping up with changes in their 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 saw no statistical change in the alignment between performance and expectations.

Retention among Temporary Warehouse Workers. One firm that employs workers through the staffing agency hires several thousand warehouse loader-movers at exactly the same pay—$17 per hour—regardless of where the job is located. Overall, the $17 per hour far exceeds the average in the staffing agency of $12.85. However, there is variation in the going rate for temporary warehouse loader-movers across the country – some areas pay $15 per hour, some $11 per hour (see Figure 6). We leverage the variation in how much of a premium $17/hr lends the firm over the outside options to gain purchase on our core question of how higher (relative) pay affects performance.[8] We dub this firm High Roller for ease of exposition.

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

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

where Y_{ijt} represents the outcome of interest, and $\overline{\text{Pay}}^{cz}_{ijt}$ 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

[8]Notably, High Roller is not hiring the warehouse workers to test them out for a permanent position: of the 8,477 individuals hired by this firm as warehouse loader-movers, 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. This case study therefore captures the effect of higher wages alone, and *not* the effect of a possible permanent position.

job. 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.

We construct the sample by limiting to season-commuting zone pairs that have more than 10 assignments from High Roller during peak seasons where High Roller hires in more than one commuting zone. We further eliminate the 267 assignments (3.2% of the 8,477 assignments) at High Roller that are hired at a different rate, since we believe these are different jobs. Of the 8,477 temporary assignments that High Roller 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 High Roller and the outside options can be seen in Panel B of Table 1.

Where the $17/hour wage represents a greater premium over the local outside options, completion 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 3). This reflects an elasticity of quits of 0.4 with respect to a dollar extra in *relative* wages. Workers are also 10 percent less likely to have a bad ending (which includes being fired for performance or attendance problems or having a "Poor" evaluation with the staffing agency), off a base of 6 percent. Since much is standardized across branches of this firm, we interpret this to mean that workers' increase in performance is not outpaced by managers' expectations.

4.B Young and Male Workers are Especially Responsive to Pay

Some groups of workers are more responsive to pay than others. When deciding whether to leave the retailer, younger people are more elastic than older workers. Men are more elastic than women. Workers with good shifts are more elastic than those with unpleasant shifts.

In both the retailer's warehouse and the retailer's customer service team, we assess elasticities of retention with respect to wages. These figures inform which workers firms can expect to stick with them. However, some of them can also help shed light on the dynamics of the market more broadly: if women are less elastic than men, then such differences can help us understand the gender pay gap.

Difference by Sex. Female workers appear to be less elastic than male workers. Male customer service representatives have a turnover elasticity of 6.63 and warehouse workers have an elasticity of 2.45. An additional dollar of pay reduces turnover by over 40 percent for male customer service representatives and 15 percent for male warehouse workers. Female customer service representatives' response to higher relative pay is economically small and statistically indistinguishable from zero. We have too few female warehouse workers to discern a reliable elasticity. These findings are consistent with findings from Wiswall and Zafar (2018) that women prefer job stability whereas men prefer jobs with higher earnings growth.

These gender differences are accentuated among customer service representatives of child-bearing age (see Appendix Table **??**). While men of child-bearing age have retention elasticities of 0.51, women of the same age have a retention elasticity of 0.14, which is statistically indistinguishable from zero. Despite different elasticities, women and men of child-bearing age have similar pay and turnover rates, after accounting for MSA and time effects — thus, women and men leave the retailer at similar rates but appear to do so for different reasons.

Figure 8 illustrates retention elasticity of female and male customer service representatives over the life-cycle. Among young representatives — the minority of whom have child-care responsibilities (see Appendix Figure **??**) — female representatives are more responsive to pay than male representatives. However, as customer service representative enter child-bearing ages, their sensitivity to pay diverges: men become increasingly responsive to pay, while women become less responsive to pay, perhaps because they put increasing weight on other attributes of the job, such as regularly scheduled hours or the potential to work remotely.

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. Consistent with gendered selection into occupations, women in the warehouse appear to have higher elasticities (though we are underpowered to detect differences due to few women working in this context). 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.

Difference by Age. Employees under 30 are more elastic than workers 30 or older. Customer service workers under 30 show an elasticity of 7.3 and warehouse workers under 30 have an elasticity of 4.4. Workers 30 and over have significantly smaller elasticities, that are not distinguishable from zero.

Difference by Worker Type. We find that workers who might be more able to secure another job — customer service representatives who performed in the top third of representatives in the one month after training, and warehouse workers in a favorable shift — have greater responsiveness to pay. Top-performing customer service representatives show elasticities of 6.68 relative to 2.76 among lower-performing representatives. Warehouse workers with daytime shifts have elasticities of 4.08 relative to 0.76 among nighttime workers. The difference in elasticities is consistent with either the workers who take night shifts having fewer options or preferring their existing bundle of disamenities and the associated additional $0.50 per hour that poor hours confers.

5 Higher Pay Increases Productivity

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, we find that 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. After the introduction of higher pay, productivity in the treated warehouse increases relative to productivity beforehand, and also relative to the changes in productivity in other warehouses that handle similar types of parcels. The elasticity of boxes moved per hour is greater than one in both specifications.

In the difference-in-differences specification, we compare to the "twin" warehouses at the retailer that handle the same type of object in different locations. Whereas other warehouses at the retailer handle smaller objects, these warehouses also handle larger objects, just like the treated warehouse.

Three metrics capture warehouse productivity: boxes moved per 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, perhaps due to increased retention.

Our contacts at the retailer confirm "twin" warehouses handle similar types of parcels, though the baseline productivity between the treated and "twin" warehouses differ considerably (see Figure **??**). 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. In contrast, the "twin" warehouse moved about 3 boxes per hour or 5 boxes per moving hour. The time-series is given in Figure 7 shows the change in units moved per hour. There is a gentle slope upward in the weeks following the pay change, followed by a leveling out. There is no such change at the "twin" warehouses. Given the different baseline levels, we believe the event study may be the most accurate representation. Nevertheless for completeness, we present both findings.

We find an increase of 0.328 boxes moved per hour in the three months following the pay jump at the warehouse, an increase in productivity of 7 percent (see Table **??**, Panel B). This corresponds to an elasticity of 1.15. Our metric of boxes moved per *moving* hour is 0.316, an increase of 4 percent and an elasticity of 0.67. 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 day.

Customer Service Productivity. Higher pay increases the number of calls that customer representatives handle by 7 percent (without compromising customer satisfaction. It also decreases the share of unscheduled absent hours. While unscheduled absent hours are not among the retailer's formal productivity metrics, they are costly to the firm since it can be hard to fill unanticipated gaps in call-capacity.

Each $1/hr increase in relative pay is associated with a 7.2% increase in call volume, 1.84 additional calls per day off of a based of 26 (see Table **??**). As in Section 4, we use changes in the local pay for customer service representatives to assess the value of an additional dollar in relative pay to reach these estimates. 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 given the dynamics in the rest of the retailer.

Higher relative pay has limited impacts on customer satisfaction. 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, the final two columns suggest that the share of absences that are unapproved in advance and thus difficult for the retailer to respond to is a decreasing function of pay. Each $1/hr increase in relative pay is associated with a 12pp decrease in the share of absences that are unapproved (18 percent off the base of 68 percent).

Notably, relative pay seems to have limited impact on hours worked, total absent hours, and overtime hours — as detailed in Table **??**. 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 labor-supply choice of how much to work at one's chosen firm than that it should impact extensive marginal labor supply choice of where to work.

6 Higher Pay May be Cost Effective

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 to paying higher wages. To help clarify the cost calculus that a firm may consider, we perform cost-benefit calculations using the case studies from the retailer. In both the warehouse and customer service instance, productivity shifts alone pay for the costs of higher wages.

At the retailer's warehouse, each $1/hour yields a gross return $1.42-1.54 from reduced turnover costs and increased warehouse efficiency.

The gross returns from decreased turnover in the warehouse are $0.28-0.40. 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 Lauren & Co pay at least $1849 per new recruit.[9] 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 $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, their gross return on a $1 investment would be $0.28. 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/hour / 4.92 boxes moved per person-hour). Since the higher pay increased the warehouse level productivity by 0.348 boxes per person-hour, the gross return on a $1 investment is $1.14.[10]

Among customer service representatives at this retailer, the gross return on a $1/hour increase in the relative wage is $1.10.

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. We estimate the cost of replacing a customer service representative to be $2,990 – including $1800 in training[11], $300 in badges and other administrative costs. Again, this is an under-estimate because it does not include recruiter time — the magnitude of the recruit-

[9] This is likely an underestimate because the full duration of time workers need to get up to speed and the costs of recruiters' time are not included.

[10] 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. On balance, we suspect this is an underestimate of the returns to boosted productivity.

[11] This is 3-weeks of wages during the period of formal training. We assume that representatives are on-net neutral during this period since they take some real customer calls but they are also more intently supervised and advised by a trainer, whose time is also valuable.

ment elasticity suggests that higher pay might appreciably reduce the time it takes recruiters to find acceptable candidates. According to these conservative estimates, increased retention would thus reflect a savings of $3,887. To achieve these savings, 100 customer service representatives working 21 eight-hour days, would have to be paid an additional dollar, implying a gross return of 23 cents.

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 1.85 calls per day according to the sticky pay specification. The return on an $8 investment/day is $8.51 – or $1.06 on the $1/hour investment. This estimate may yet be an underestimate if firms are able to leverage higher pay's effect on recruitment. We found that higher pay attracted more workers, with an elasticity of over 3. If firms can distinguish between high-productivity workers, then securing greater interest from prospective employees is valuable. Of course, if firms are unable to ascertain workers' talent *a priori*, then greater recruitment is not helpful.

7 Mechanisms: Both Selection and Behavioral Responses Contribute

A key question is the mechanisms underlying changes in turnover and productivity. Does raising wages for existing workers incentivize higher performance, or is the key advantage of higher pay attracting and retaining talent? We find that 20-50 percent of the effects can be attributed to sorting better workers to higher paying jobs, with the rest attributable to improvements in the same workers' behavior.

We first use staffing agency data that shows the same worker in multiple jobs to assess how much pay affects the behavior of the individual worker, finding that as much as 50-80 percent of the effect arises from behavioral responses to higher pay. We then document that higher pay both attracts and retains better workers: when higher pay is offered at the retailer's warehouse, those who switch into the warehouse are positively selected. Similarly, those who secure higher paying jobs at the staffing agency are more likely to be predicted to be excellent. Among customer service representatives, higher relative pay appears to be more effective in preventing turnover among high quality and more senior customer service representatives. Likewise those who leave the warehouse after the pay jump tend to be newer workers.

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

7.A Higher Pay Incentivizes Better Work

While attraction and retention are important factors, they do not account for all of the effect of higher pay. Indeed, if the new workers at the retailer's warehouse accounted for the entirety of the productivity shift, they would have to work at 177 percent the level that workers had been before the pay bump. That is, they would have to move an additional 3.8 boxes per hour over the baseline 4.92 boxes per hour. [12]

We can also assess the incentive-effects of higher pay in the customer service context. We first 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 than more able workers, one might expect the output of these workers to be more sensitive to the relative pay of the retailer. Indeed, in Table 8, we find that call-volume effects are concentrated among representatives, who are in the bottom two-thirds of call volumes in their first month at the retailer, as consistent with these representatives being more concerned about the possibility of termination. 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.

Moreover, we can investigate within-representative effects of changes in the retailer's relative pay, by adding in individual fixed effects to specification 3 to arrive at:

$$Y_{imt} = \gamma_{\$,\text{Uniform}} \Delta(\text{Entry Relative Wage})_m \cdot \mathbb{1}[\text{Year} = 2019] + \mu_t + \mu_i + \zeta_{imt}. \quad (4)$$

Table 8 details the result of this approach. Unfortunately, this design substantially reduces the precision of our estimates, but our results nonetheless suggest that much of the productivity effects persist within a representative. Taken at face value, the within-representative call effect in column two explains 64% of the aggregate effect in column one. Similarly, 37% of the reliability-effect on the share of absences that are unapproved persists within worker. Finally, the customer satisfaction changes are larger within worker than across workers, although both changes are economically small and statistically insignificant.

We can employ a similar design in the context of the staffing agency, where we often see the same worker in multiple jobs, each with different pay.

We first perform the same exercise in the context of the High Roller, where we can leverage national wage-setting to identify variation in real pay for similar work done in different geographies. Results in the context of High Roller are noisy, but go in a similar direction, with about half of the effect on job completion arising from incentives (see Figure **??**). In this context, all of the reduction in worker quits can be attributed to the incentives of workers.

[12] Scaling each worker by the amount of time they worked at Lauren & Co., there are 241.8 effective retained workers during the post-period and 44.27 effective new workers during the post-period. If retained workers move 4.92 boxes per hour, new workers would have to move 8.72 boxes per hour for the whole team to reach 5.51 boxes per hour.

More broadly, the staffing agency affords us many more observations with which to draw inferences about the relative importance of selection and incentives, albeit in settings that are less clean than our case studies. These estimates suggest about half to two-thirds of the effect of higher pay on job completion and excellent evaluations persist within worker.

For this analysis, we focus on the sample of warehouse workers, since this offers a relatively homogeneous set of jobs that account for a third of agency placements (222,904 of 712,275) and concord with one of the two routine jobs in our case studies. We begin by estimating the reduced-form relationship between pay and performance:

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

where Y_{ij} may be the job completion rate or the rate of workers quitting and $\beta_{\$}$ 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.

We seek to leverage the staffing agency context to understand how higher pay changes performance *within* a given worker. This is instructive for a firm considering how changing pay will impact the productivity of existing workers or for a policy-maker considering the aggregate effects of higher pay on output across many firms. To isolate the incentive effects of higher pay, we can look at the relationship between completion and pay *within* individual workers who work multiple jobs at the staffing agency, using the regression:

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

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.[13] Table 9 presents the results of this analysis. In the first panel, each additional dollar of pay is associated with a 2pp higher rate of job completion off of a base of 55 percent in the sample of workers who take multiple jobs. We find a 1.95pp (or 11 percent) decrease in workers quitting. There is likewise an increase in excellent evaluations of 0.47pp (1.8 percent). Panel B turns to *within worker* effects. Within a given worker, each additional dollar of pay is associated with a 1.3pp higher rate of job completion off of a higher base of 55 percent in the sample of workers

[13] Appendix Table **??** reports the reduced-form relationship between pay and job completion and worker evaluations for all warehouse jobs and for inbound call center jobs, which are the most comparable jobs to those in our case studies. Appendix Table **??** presents this reduced-form relationship for all routine jobs in the staffing agency, including manufacturing and cleaning jobs as well as warehousing and call-center work.

who take multiple jobs. We find a 1.7pp (or 11 percent) decrease in workers quitting. There is likewise an increase in excellent evaluations of 0.34pp (1.8 percent).

Since the aggregate effect of higher pay combines the *within-worker* incentive effect and the cross-worker selection effect, we can use our incentive estimate to decompose the effect of higher pay into the incentive-effect and selection-effect. We find that 60% of the effect of higher pay on performance comes from the within-worker incentive effect while the remaining 40% comes from the cross-sectional selection effect. By contrast for quits, 88% of the effect comes from within-worker incentives, supporting the idea that quits are under the worker's control and a function of her decisions whereas other terminations may reflect the worker's abilities irregardless of her efforts. As consistent with this, only 26% of the effect on bad endings occurs within worker.

7.B Higher Pay Attracts Better Workers

At the staffing agency, the quality of workers place in a job responds to hourly pay. An additional dollar in hourly pay means a firm is 8.5 percent less likely to have a worker who is predicted to do poorly and about 5 percent more likely to have a worker to is predicted to do excellently (see Table **??**). To create this analysis, we construct a prediction of workers' evaluations based on their prior assignments, job evaluations, and job endings. 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.

When pay is increased at the retailer's warehouse, 14 people switch from the retailer's other local warehouses to the treated warehouse.[14] These transferees are positively selected: they have higher pay than their peers at the warehouse of origin and longer tenure than the existing workers at the treated warehouse.

Positive selection is consistent with our findings below that higher pay in one warehouse has negative impacts on local, rival warehouses.

7.C Higher Pay Retains 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. We hypothesize that the pay in local outside options is more important for more productive worker, 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 produc-

[14] These people are not considered terminations and therefore do not impact our estimates of pay's effects on turnover where we use the in-state warehouses as a control.

tive 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 (see Table 6). Each $1/hr of additional relative pay reduces turnover by 44% for initial top performers, implying a retention elasticity of 0.35. By contrast, for the rest of the representatives, a $1/hr increase in relative pay decreases turnover by 17%, implying a retention elasticity of 0.12. This suggests that *selective retention* may be an important driver of increased productivity.

In the second panel of Table 6, we see that the greater responsiveness of top performers is driven by men. Among male high-performers, each dollar of relative pay reduces monthly turnover by 5pp (85%), implying a retention elasticity of 0.84 for these workers which is significantly higher than the retention elasticity among other male representatives (see Appendix Table **??**). By contrast, among female top performers, the retention elasticity is only 0.09, which is not meaningfully different from that of other women (as shown in Appendix Table **??**.

The gender difference in the interaction between baseline productivity and relative pay in determining retention is further evidence that women are less likely to seek out alternative job opportunities, but instead often leave the retailer for non-work concerns. Indeed, suppose this were the case and women were less likely to seek out alternative job opportunities. This would manifest in a lower elasticity of retention with respect to the retailer's relative wage. If male and female representatives were both more likely to be able to get job offers when they are more productive, then this would also attenuate the interaction between productivity and relative pay for female representatives, since highly productive women would be less likely to actualize better job opportunities. Thus, both lower retention elasticities overall and especially the lower retention elasticities of highly productive women support the idea that women are less likely to search for alternative jobs while in the current one.

7.D Higher Pay has Negative Spillovers for 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? An intimately related question is whether higher pay is zero-sum between firms?

We address how higher wages may spillover to other firms in the context of High Roller, since High Roller hires through the staffing agency only during the holiday season when they need more loader-movers, and who pays all their workers $17/hour (see Section **??** for an introduction to High Roller). 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 High Roller is hiring to see if High Roller causes other firms to have unexpected turnover. We then look at whether High Roller 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.

Empirical Approach. 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 born 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.

We define a treated commuting zone to have at least one month in which High Roller hires more than 45 individuals in that month and to have had at least 20 hires outside the High Roller. We require that control months have at least 20 job placements outside High Roller. The entire sample is limited to warehousing jobs. We further restrict the sample to the three months just before High Roller seeks workers to account for potential seasonality. The sample is described in the first three columns of Table **??**. The commuting zones where High Roller 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 High Roller locates, which might be due to having job that are expected to last longer.

An alternative control group includes all the warehouse jobs outside the treated commuting zones in the state. The same requirements that a treated commuting zone must have more than 45 High Roller hires in a month and at least 20 hires outside that firm and that control months have at least 20 job placements outside High Roller still hold. This set of commuting zones features lower-paying, shorter jobs than the commuting zones where High Roller hires. There are fewer excellent evaluations and more bad endings in High Roller's locations (see Table **??**).

We use a simple difference-in-differences approach:

$$Y_{it} = \alpha_0 \mathbb{1}^{cz}_{ijt} + \alpha_1 + \mathbb{1}^{month}_{ijt} + \beta(T \cdot \mathbb{1}^{cz}_{ijt} \cdot \mathbb{1}^{month}_{ijt} \cdot \overline{\text{Pay}}_{ijt}) + \mu_T + \epsilon_{ijt} \tag{6}$$

Since our data includes several periods of hiring, we estimate the treatment effect for each year separately, and include μ_T fixed effects for the year. We keep the years separate to ensure that we do not put negative weight on any of our comparisons, in keeping with the current literature critiquing two-way fixed effects models (e.g., Goodman-Bacon, 2018; Abraham and Sun, 2018; Imai and Kim, 2019; de Chaisemartin and D'Haultfoeuille, 2020). 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, we interact the treatment effect with $\overline{\text{Pay}}_{ijt}$, the average pay differential in the commuting zone - season pair. β thus captures how much an additional dollar of relative pay offered by High Roller impacts job outcomes at rival firms in each of our treated years.

To assess the parallel trends assumption, we plot in Figure **??** 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-year fixed effects, as Pischke (2019) recommends, and we find no significance.

Higher Pay impacts performance among local rival firms. When High Roller is hiring in one dollar higher than the outside option, workers at rival firms are between 16 and 34 percent less likely to complete the job off a base job completion rate of 25 percent in control commuting zones and months, and are between 9.6 and 18 percent more likely to have a bad ending off a base of 45 percent.

If selection is at work, we may see not only increased turnover at rival firms, but also lower quality workers placed in those jobs. Workers hired into rival firms when High Roller his 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 High Roller, one could imagine workers at the agency leaving their existing gigs in order to take higher-paid positions.[15] To see whether the higher pay at High Roller 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 High Roller is hiring to the job endings in other locations where High Roller's rivals locate. This is distinct from the above analysis, where we were comparing workers *placed* in jobs at the same time; here we examine those jobs that end in the months when High Roller is hiring.

We consider all warehouse jobs in commuting zones that High Roller or its rivals locate in. 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 High Roller is hiring at a pay differential of \$1/hour more than the outside option.

Table A.2 shows that there is not an uptick in staffing workers leaving their job or otherwise having an unsatisfactory end in order to take jobs at the central firm.

[15] 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.

8 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. We estimate that 20-50 percent of the improvement arises from sorting better workers to higher-paying positions, but that the remainder is attributable to workers' behavioral responses to higher pay.

This paper is relevant to many firms seeking to set wages for their routine workers. It suggests that in some cases seeking to trim the bottom line may in actuality increase costs by increasing turnover or dampening on-the-job productivity. Especially since many warehousing jobs occur in states with relatively low minimum wages, including Texas, Kentucky, Ohio and Georgia, higher wages for warehouse workers will only occur through firms' wage-setting decisions.

This paper is relevant to those local governments that are inclined to use their policy levers to bolster pay of low-wage workers. In particular, the elasticities of labor supply with respect to wages that we calculate help inform whether a minimum wage boost or an increase in the EITC is more likely to help workers. When labor-supply is highly elastic and labor-demand is highly inelastic, work subsidies like the Earned Income Tax Credit will not benefit employees: workers will see their wages fall in lockstep with the increases in government benefits. Conversely, in such a world, increases in the minimum wage will benefit workers, who will see wages rise with little change in available opportunities. Likewise, if firms keep wages low but are inelastic in demand, then an increase in the minimum wage will increase take-home pay without a significant effect on the number of individuals hired. Thus, our elasticity estimates are key in guiding policy in this changing economy.

If policymakers are intent on raising the minimum wage as a way to address monopsony power of employers to potentially depress wages, our result suggest that this may be even more crucial in female-dominated industries. We find that male workers have greater retention elasticities than female workers. Thus when the concentration of firms is used as a measure of monopsony power, we may underestimate firms' power to set female wages. Thus focusing on female-dominated industries when considering wage-setting regulations may be sensible.

To the extent that our results are often measuring the difference between a firm's pay and workers' outside options, minimum wage changes also matter. If a minimum wage increase compresses the wage distribution, workers who were paid above the minimum wage will have less difference between their wages and their outside options. In that case, raising wages may help firms capture lower turnover and higher productivity. Thus our paper shows why firms may seek to raise wages in the wake of a minimum wage change, even for workers who were not paid the minimum wage.

Our results do have limitations and leave a number of questions unanswered. Our results that explore the spillovers of high-paying firms on local rival firms can only document the effect in terms of worker-retention, 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.

While we provide suggestive evidence the mechanisms underpinning our results, we cannot perfectly nail 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 answer general equilibrium questions such as what might happen if wages were raised universally in a given geography. If the effect we document is coming from greater effort, then all firms might see an increase in productivity when all workers' pay is raised. If higher pay generates greater productivity because better workers gravitate toward higher-paying firms, then a global increase in pay will not induce greater productivity since no sorting would occur. Alternatively, 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 extremely and high-value for informing policy.

Our work does not completely capture involuntary unemployment, a key component of efficiency wage theory. If firms raise wages to reduce turnover or increase productivity, they also create involuntary unemployment and cuing for jobs. In our current data we are unable to test this component of the theory. Ideal data might be able to show whether there are higher application rates jobs in areas where the relative pay is higher.

While there is ample room for additional research, this paper contributes by (a) estimating extensive and intensive labor supply elasticities in several contexts, (b) bringing objective productivity metrics to bear on this question, (c) providing suggestive evidence about the relative contributions of sorting and workers' behavioral responses to pay, and (d) investigating the spillovers of higher pay on worker retention.

References

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

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.

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.

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

Frantz, Roger S, *X-efficiency: Theory, evidence and applications*, Vol. 2, Springer Science & Business Media, 2013.

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.

Kahn-Lang, Ariella and Kevin Lang, "The promise and pitfalls of differences-in-differences: Reflections on 16 and pregnant and other applications," *Journal of Business & Economic Statistics*, 2020, *38* (3), 613–620.

Katz, Lawrence F, "Efficiency Wage Theories: A Partial Evaluation," *NBER Working Paper*, 1987, p. 62.

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.

Leibenstein, Harvey, "Allocative efficiency vs." X-efficiency"," *The American Economic Review*, 1966, *56* (3), 392–415.

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.

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

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.

9 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	13.40	9.30	5.15	4.40	6.59	4.31
% Quit	10.66	7.50	4.05	3.57	5.55	3.50
% Fired	2.02	1.24	0.70	0.62	0.83	0.58
% Turnover 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 Hours				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	140,664.00	5,701.00	32,009.00
# Assignments	222,904.00	6,664.00	45,454.00
n_cz	374.00	83.00	83.00
n_firms	3,950.00	1.00	1,448.00

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: Customer Service: Effect of Relative Pay on Recruitment

Panel A: Pay's Effects on Recruitment

	# 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	No	No	No	Yes	Yes
Retailer Non-CSR Counts	No	No	No	No	Yes
Mean Recruits/MSA	0.73	0.73	0.73	0.73	0.73
# MSAs	920	920	920	920	920
F	138.83	35.4	74.4	54.95	43.41
R^2	0.232	0.289	0.297	0.300	0.388

Panel B: Pay's Effects on Recruitment by Sex

	# 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	Quartic	Quartic	Quartic	Log
Retailer Non-Csr Presence	No	No	No	Yes	Yes
Retailer Non-Csr Counts	No	No	No	No	Yes
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

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 3: Higher Pay Enhances Productivity

Panel A: Turnover Effects In the Retailer's Warehouse

	First Stage		Monthly Retention		Quits		Fires	
Post x Treated	1.755***	2.152***						
	(0.072)	(0.093)						
$1/hour			−2.504***	−3.780***	−2.270***	−3.797***	0.061	0.245
			(0.937)	(1.207)	(0.822)	(1.075)	(0.405)	(0.487)
Elasticity			3.03***	4.57***	3.45***	5.77***	0.49	1.97
			(1.13)	(1.46)	(1.25)	(1.63)	(3.26)	(3.91)
Base Mean	16.2	16.2	13.4	13.4	10.66	10.66	2.02	2.02
Comparison		In-State		In-State		In-State		In-State
Workers	514	1557	514	1557	514	1557	514	1557
F	599.62	535.15	7.31	3.52	7.81	4.22	0.02	2.75
Observations	50,478	149,656	50,478	149,656	50,478	149,656	50,478	149,656

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	✓	✓	✓
F	1.93	1.79	0.99
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	1.709**	−0.523***	−0.631***
	(0.801)	(0.202)	(0.212)
Elasticity	0.09	0.4	0.45
	(0.04)	(0.15)	(0.15)
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	0.078	0.055	0.032

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

Table 4: Retention Elasticities by Worker Group

Panel A: Pay's Effects on Retention by Sex

	Customer Service			Warehouse		
	Female	Male	Δ	Female	Male	Δ
Effect on Turnover	-0.91	-1.87**	0.96*	-2.92	-2.01*	-0.91
	(0.6)	(0.6)	(0.54)	(1.79)	(1.03)	(2.06)
Elasticity of Turnover	3.47	6.63	-3.16	5.06	2.45	2.6
Mean Turnover	4.17	4.57	-0.78***	9.32	13.25	-3.93
			(0.23)			(1.62)
Mean Pay	15.94	16.19	-0.07	16.14	16.21	-0.07
			(0.07)			(0.22)
# Workers	2097	901		99	395	
# MSAs	39	23		1	1	

Panel B: Pay's Effects on Retention by Age

	Customer Service			Warehouse		
	Age <30	Age ≥30	Δ	Age <30	Age ≥30	Δ
Effect on Turnover	-2.4**	-0.39	-2**	-5.19**	-0.57	-4.61**
	(1.13)	(0.45)	(0.87)	(1.81)	(0.99)	(2.06)
Elasticity of Turnover	7.33	1.74	5.59	4.42	1.05	3.37
Mean Turnover	5.19	3.66	1.78***	18.61	8.89	9.72**
			(0.35)			(1.5)
Mean Pay	15.88	16.12	-0.3***	15.86	16.39	-0.53***
			(0.07)			(0.16)
# Workers	1312	1688		201	301	
# MSAs	25	39		1	1	

Panel C: Pay's Effects on Retention by Worker Type

	Customer Service			Warehouse		
	Start Top Third	Start Bottom	Δ	Night Shift	Day Shift	Δ
Effect on Turnover	-2.13**	-0.77**	-1.36	0.22	-4.97***	5.2**
	(1.06)	(0.33)	(1.05)	(0.93)	(1.45)	(1.72)
Elasticity of Turnover	6.68	2.76	3.92	0.76	4.08	-3.32
Mean Turnover	4.98	4.39	0.32	4.87	19.42	-14.55***
			(0.44)			(2.29)
Mean Pay	15.62	15.78	-0.27***	16.56	15.95	0.6***
			(0.04)			(0.16)
# Workers	615	1207		132	382	
# MSAs	17	25		1	1	

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

We calculate responsiveness to the retailer's (relative) wages, separately for workers with different demographics. Differences in means among customer service representatives are calculated base on gregressions that include date and MSA fixed effects and have standard errors clustered at the MSA-level. Differences for warehouse workers are calculated from regressions with standard errors clustered at the individual level.

Table 5: Higher Pay Enhances Productivity

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

	First Stage		Boxes/Hr		Boxes/Moving Hr		Moving/Total Hrs	
Post x Treated	1.755***	1.622***						
	(0.072)	(0.026)						
\$1/hour			0.348***	0.375***	0.316***	0.557***	0.018***	0.001
			(0.015)	(0.021)	(0.014)	(0.031)	(0.001)	(0.002)
Productivity ϵ			1.15***	1.23***	0.67***	1.17***	0.46***	0.03
			(0.05)	(0.07)	(0.03)	(0.07)	(0.02)	(0.05)
Base Mean	16.2	16.2	4.92	4.92	7.69	7.69	0.64	0.64
Comparison		Twin		Twin		Twin		Twin
Workers	514	1733	514	1733	514	1733	514	1733
F	567.46	3826.99	567.46	176040.6	535.98	98323.11	563.13	31671.89
Observations	50,478	187,861	50,478	187,861	50,478	187,861	50,478	187,861

Panel B: Pay's Effects Among Customer Service Representatives

	Daily Call Volume	Satisfaction (out of 5)	% of Absences Unapproved
Entry Relative \$1/hr	3.331***	0.007	−12.080***
	(0.954)	(0.006)	(4.408)
Elasticity	2.1***	0.022	2.822***
	(0.6)	(0.019)	(1.029)
FE: Date	✓	✓	✓
F	68.9	1.9	12.81
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. Figure **??** 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 6: Customer-Service Retention Effects By Initial Productivity

By First Month Call Volume	Top 1/3rd	Bottom 2/3rds	Δ
Effect of Relative $/hr on Retention	2.15**	0.75**	1.40*
	(0.83)	(0.36)	(0.75)
Elasticity of Retention	0.354	0.124	0.230
Mean Monthly Retention	95.11	95.58	-0.20
			(0.44)
Mean $/hr	15.65	15.77	-0.19***
			(0.04)
# Reps	611	1211	
# MSAs	17	25	

By Gender: Top 1/3rd First Month Call Volume	Female: Top 1/3rd	Male: Top 1/3rd	Δ
Effect of Relative $/hr on Retention	0.53	5.01***	-4.49***
	(1.01)	(0.57)	(1.09)
Elasticity of Retention	0.087	0.840	-0.753
Mean Monthly Retention	95.50	94.13	1.03
			(0.63)
Mean $/hr	15.60	15.77	0.22*
			(0.12)
# Reps	440	171	
# MSAs	15	9	

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

We leverage the stickiness of the retailer's wage to calculate retention elasticities across workers with different baseline productivities. We assess productivity according to representatives' daily call volumes in their first month of calls after formal training relative to the average number of calls taken in those days at the retailer. We ask whether turnover rises more sharply for more productive workers in those MSAs where the retailer's sticky wage loses more ground to the local alternatives. Differences in means are calculated based on regressions that include date and MSA fixed effects. All standard errors are clustered at the MSA-level. We find that higher pay is particularly effective at retaining top-performers, who may be better able to convert more lucrative outside options into attractive job offers. This finding suggests that *selective retention* may be an important driver of increased productivity. The second panel demonstrates that the greater responsiveness of high performers is driven by men.

Table 7: Customer-Service Productivity Effects By Initial Productivity

By First Month Call Volume	Top 1/3rd	Bottom 2/3rds	Δ
Effect of Relative $/hr on Calls	-0.18	1.87***	-2.05***
	(0.52)	(0.43)	(0.43)
Elasticity of Calls	-0.040	0.387	-0.427
Mean Calls	31.71	23.80	6.35***
			(0.38)
Mean $/hr	15.61	15.73	-0.19***
			(0.04)
# Reps	557	1119	
# MSAs	16	24	

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 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.

Table 8: Customer-Service Productivity Effects Within and Across Representative

	Daily Call Volume		Satisfaction (out of 5)		% of Absences Unapproved	
	Across	Within	Across	Within	Across	Within
Entry Relative $1/hr	1.848**	1.189	0.007	0.012	−12.080***	−4.429*
	(0.924)	(1.221)	(0.006)	(0.007)	(4.408)	(2.326)
% of Full Effect		64.3%		177.5%		36.7%
FE: Date	✓	✓	✓	✓	✓	✓
FE	MSA	Rep	MSA	Rep	MSA	Rep
F	77.05	71.16	1.9	2.41	12.81	20.81
Mean $/hr	15.88	15.88	15.88	15.88	16.11	16.11
Dependent Mean	25.68	25.68	4.89	4.89	68.31	68.31
MSAs	41	41	41	41	41	41
Workers	2598	2598	2598	2598	2782	2782

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

Since the stickiness of the retailer's pay generates variation in its relative pay over time, we can utilize this design to investigate how productivity changes *within* a representative as well as across representatives. Although this substantially reduces the power of our design, the point estimates suggest that about two-thirds of the call effect and a third of the reliability effect can be attributed to changes within a worker rather than across workers.

Table 9: Incentive vs Selection Effects in Staffing Agency Warehouse Jobs

Panel A: Cross-worker regressions

	Job Completed	Quits	Bad Ending	Excellent Eval
Hourly Pay	2.067***	−1.954***	−1.283***	0.467*
	(0.352)	(0.277)	(0.291)	(0.244)
Elasticity	0.44	0.95	0.67	0.29
	(0.07)	(0.14)	(0.15)	(0.15)
Occ-CZ Fixed Effects	✓	✓	✓	✓
Ind-CZ-Time Fixed Effects	✓	✓	✓	✓
Controls	Days Quartic	Days Quartic	Days Quartic	Days Quartic
Mean	55.432	24.251	22.529	19.019
Workers	23,161	23,161	23,161	23,161
Observations	84,123	84,123	84,123	84,123
R^2	0.558	0.321	0.317	0.593

Panel B: Within-worker regressions

	Job Completed	Quits	Bad Ending	Excellent Eval
Hourly Pay	1.249***	−1.710***	−0.335	0.335*
	(0.476)	(0.349)	(0.318)	(0.177)
Elasticity	0.27	0.83	0.18	0.21
	(0.1)	(0.17)	(0.17)	(0.11)
Individual Fixed Effects	✓	✓	✓	✓
Occ-CZ Fixed Effects	✓	✓	✓	✓
Ind-CZ-Time Fixed Effects	✓	✓	✓	✓
Controls	Days Quartic	Days Quartic	Days Quartic	Days Quartic
Mean	55.432	24.251	22.529	19.019
Workers	23,161	23,161	23,161	23,161
Observations	84,123	84,123	84,123	84,123
R^2	0.774	0.660	0.667	0.742

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

[This Table will be put in the format of Table 8 and made as the second panel in that table.] In the warehouse context, incentives make up 60 percent of the increase in job completion, as measured by the change seen *within* workers and 87 percent of the change in quits is attributable to incentives. Within worker, there is no statistically significant change in bad endings, but 71 percent of the increase in excellent evaluations occurs within the same worker

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 10: Spillovers from High Roller's Hiring on other Local Firms

	Job Completed		Quits		Bad Ending		Excellent Eval	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated Month	−2.449	5.104***	−2.970	−3.438***	−9.687**	−4.348***	−9.692**	20.370***
	(3.861)	(0.820)	(8.255)	(0.747)	(4.789)	(0.730)	(4.168)	(0.718)
Treated CZ	4.204	−28.970***	6.680***	14.530***	5.587	15.230***	−1.718	−20.000***
	(2.836)	(0.786)	(2.325)	(0.716)	(4.151)	(0.700)	(4.745)	(0.688)
$ Diff. x Treat x 2016	−0.734	−1.036***	3.031**	1.128***	3.522***	−0.040	2.459***	−5.107***
	(0.940)	(0.316)	(1.488)	(0.288)	(0.726)	(0.281)	(0.822)	(0.276)
$ Diff. x Treat x 2017	0.496	2.010***	1.247	−0.462	1.658**	−0.700*	2.621**	−4.174***
	(0.573)	(0.411)	(1.858)	(0.374)	(0.800)	(0.366)	(1.203)	(0.359)
$ Diff. x Treat x 2018	1.577	−0.329	1.299	0.802	1.734*	2.396***	1.325	−3.059***
	(0.964)	(0.566)	(2.101)	(0.516)	(0.946)	(0.504)	(1.173)	(0.496)
Year Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Mean	15.6	63.4	36.6	20.7	39.1	19.5	7.4	31.7
Workers	2,170	13,693	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	2,439	24,050
R^2	0.007	0.116	0.019	0.041	0.018	0.040	0.024	0.167

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

[These yearly point estimates will be aggregated soon.] We perform a continuous difference-in-differences exercise to see how High Roller's hiring affects firms hiring in the same local labor market in the same month. When pay at High Roller is one dollar higher than the outside option, workers hired into other local firms are less likely to complete the job, more likely to quit or otherwise have a bad ending. 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.

10 FIGURES

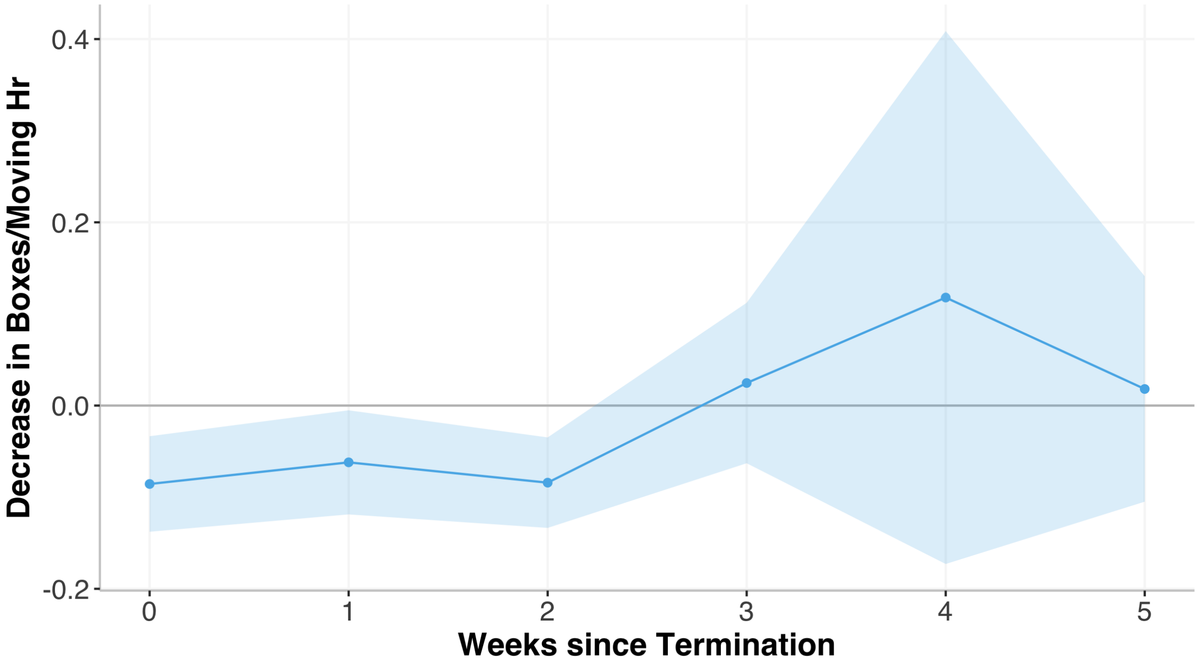

Figure 1: Boxes moved per moving-hour decreases for the three weeks after a worker has left a warehouse. The baseline number of boxes moved per moving hour is 8.83, so a reduction of -0.086 boxes reflects a 1 percent decrease in productivity. These estimates come from regressing a terminations and five weeks' worth of lagged terminations on the number of boxes moved, with standard errors clustered at the warehouse level.

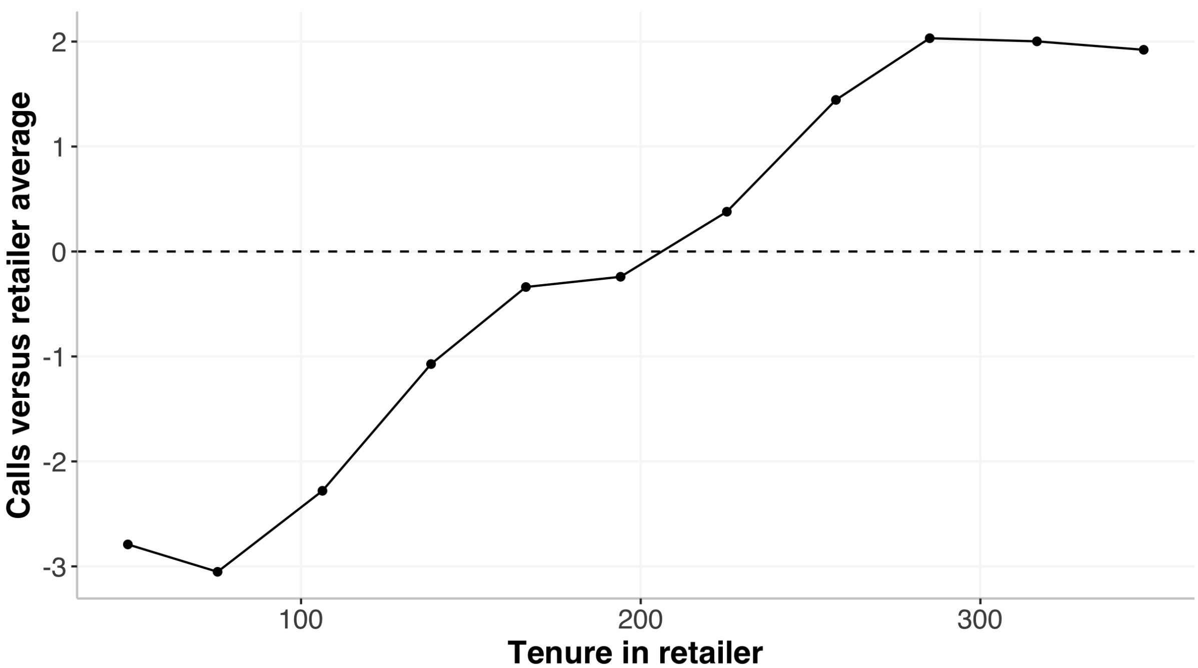

Figure 2: We plot the daily call volume of workers vis-a-vis a typical customer service representative in same time-zone taking calls on the same day. We exclude the initial three weeks when workers are in formal training so the number of calls they handle may not reflect their personal productivity. After training, we see that new hires make about 3 calls fewer per day than the average representative. By the time the representative has been with the retailer for about a year, productivity levels off at about 2 calls over the mean. Taken together this suggests that each month, the representative improves by about half a call or shaves off approximately 30 seconds per call. This development of worker capacity suggests that turnover is costly for the firm because workers take away accumulated human capital.

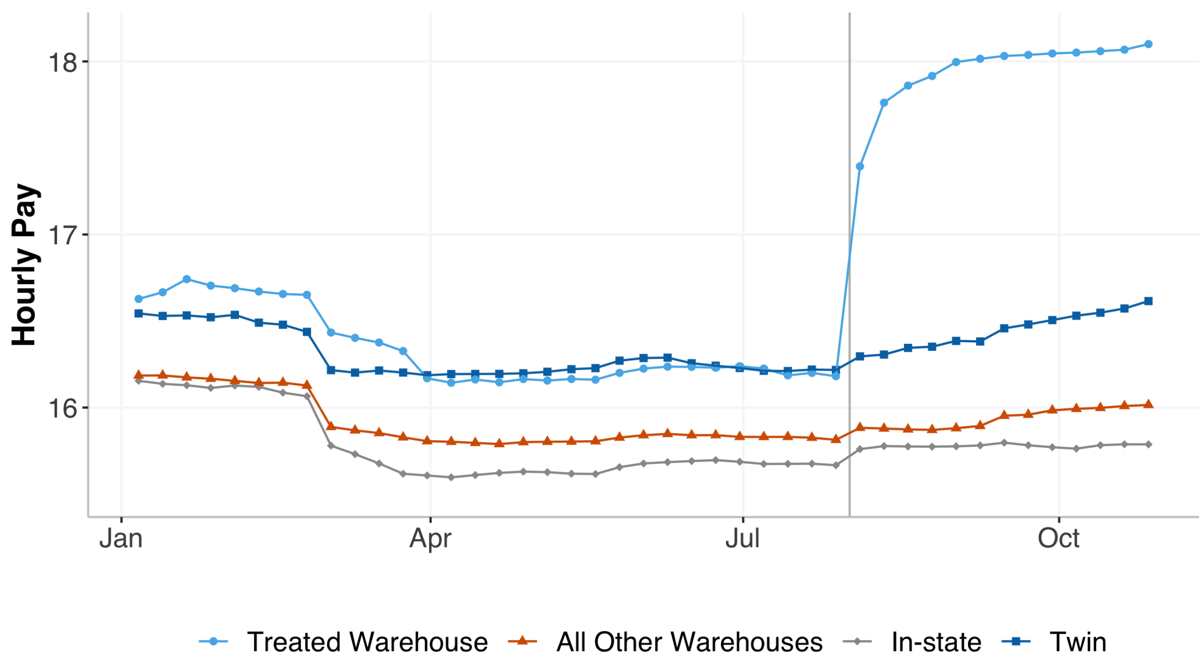

Figure 3: In August 2019, Lauren & Co increased pay at a single warehouse, keeping pay at all other warehouses (orange), warehouses in the same state as the treated warehouse (grey), and "twin" warehouses that handle the same type of package (dark blue) relatively constant. The pay increase arose owing to systemic problems with higher levels of turnover in the treated warehouse, which was in turn due to handling heavier packages. As such a "material handler" surcharge was added to these workers' hourly pay.

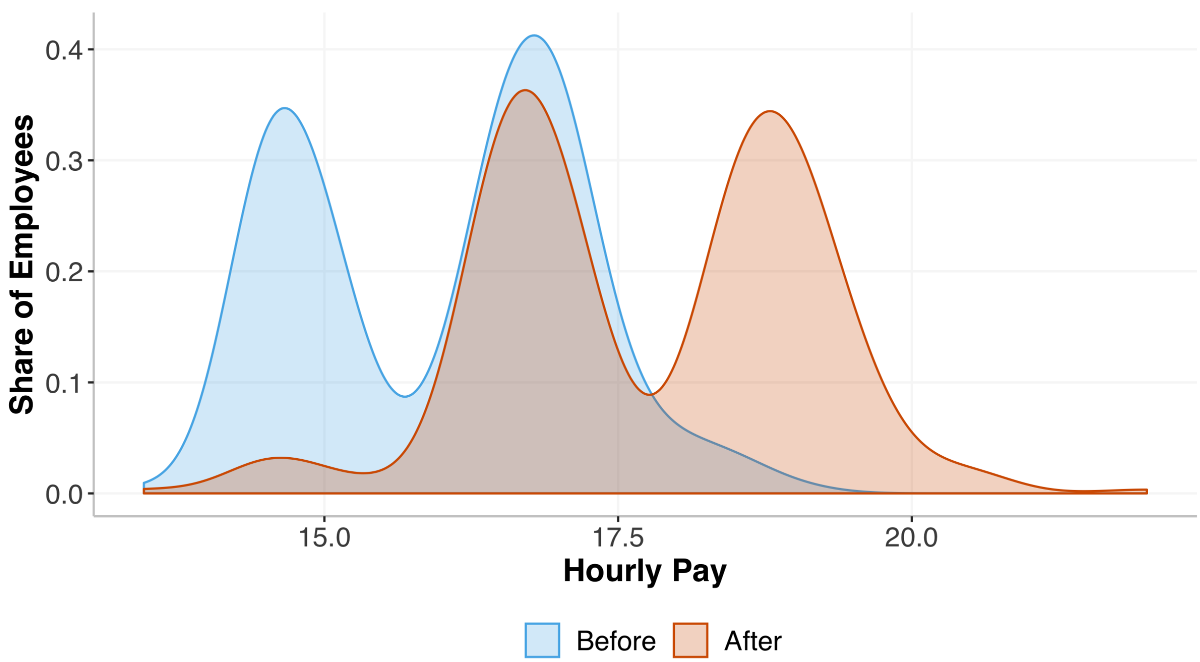

Figure 4: The distribution of pay amongst Level-1 warehouse workers remained fairly constant throughout the pay increase: before the pay change, the standard deviation of hourly pay was 1.18, afterward it was 1.21. Thus, the incentives to be promoted likely remained similar, which would not have been the case had there been greater compression in the wages. Note that there were 34 people whose pay remained constant throughout this period, but with an extended bandwidth, most of these people's pay also increased.

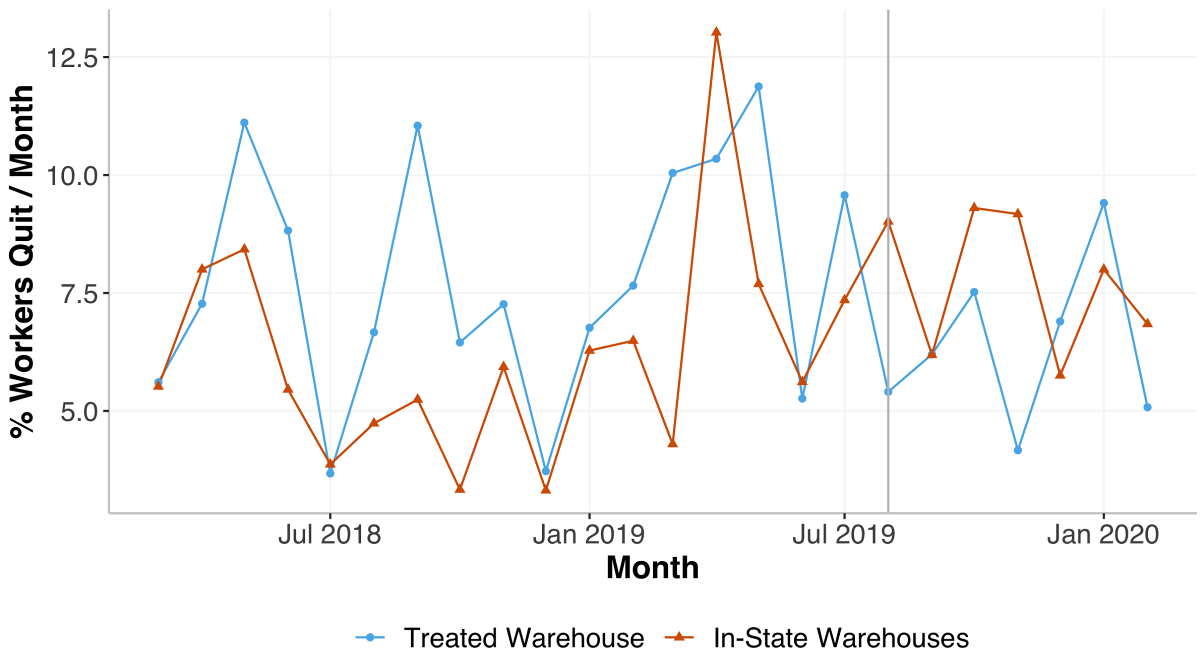

Figure 5: Rates of monthly quits are high and variable in the treated warehouse (blue). In other in-state warehouses, quit rates are likewise variable and also rather high. This reflects the hyper-competitive local labor market for warehouse workers. These warehouses are located right by a major cargo airport and many large companies, including Amazon have extensive warehouse real estate in close proximity. In contrast, the turnover rate is lower at the remaining warehouses out of state. Interestingly, both the in-state and out-of-state warehouses seem to have a gradual increase in quits over the course of the year around the pay jump in the treated warehouse. Given the jagged nature of the treated warehouse's turnover, it is hard to discern large trends.

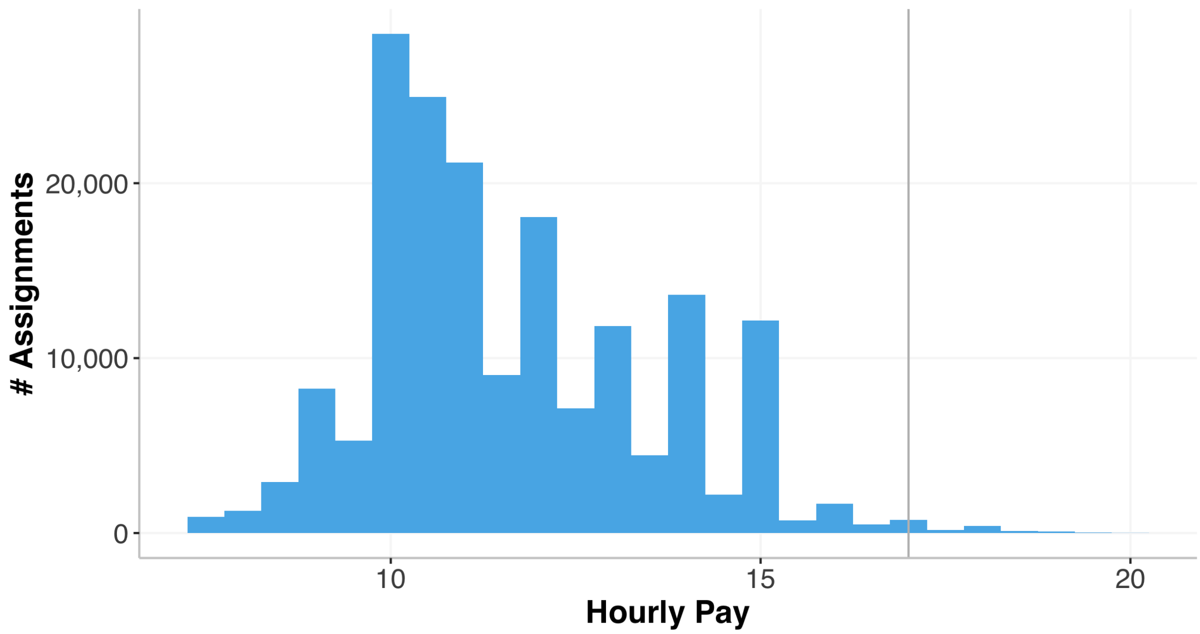

Figure 6: High Roller pays all of it's workers the same high rate relative to the going rate in the geographic area. We leverage the variation in the local hourly rate for loader-movers to look at the effect of relative pay on performance. Above we plot the wages paid to loadter-movers in other firms in the same industry in the same commuting zones as High Roller locates. The grey line shows the wages paid at High Roller.

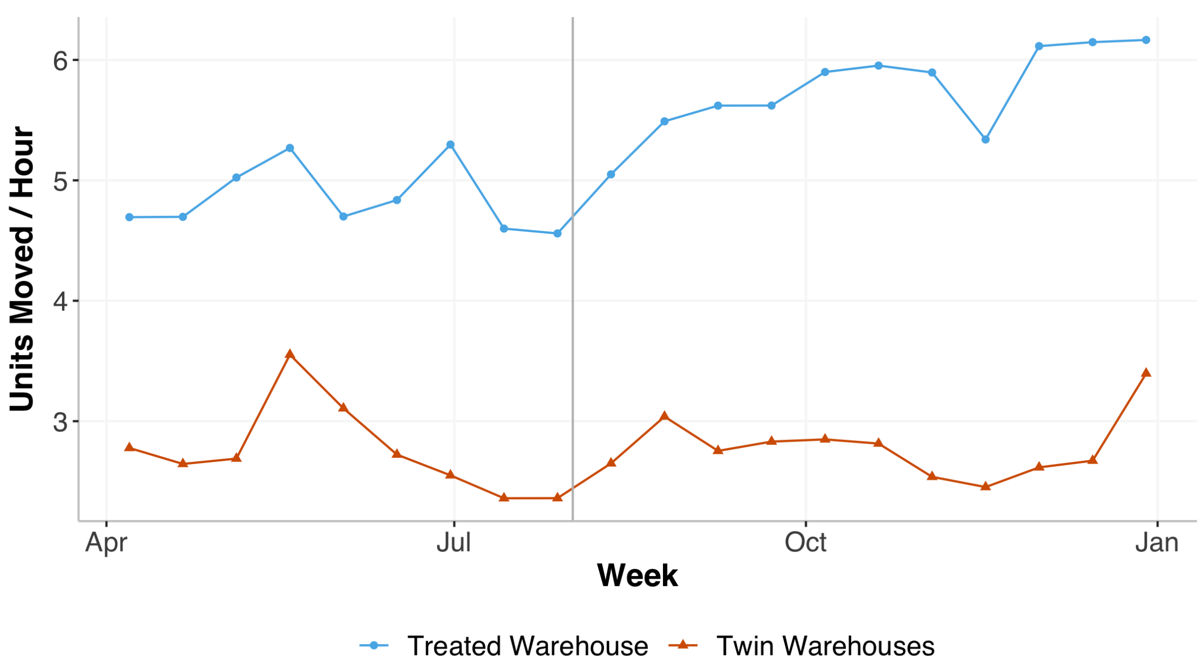

Figure 7: After the pay change, we see a gentle trend upward in boxes moved per hour at the treated warehouse and no such move at the twin warehouses.

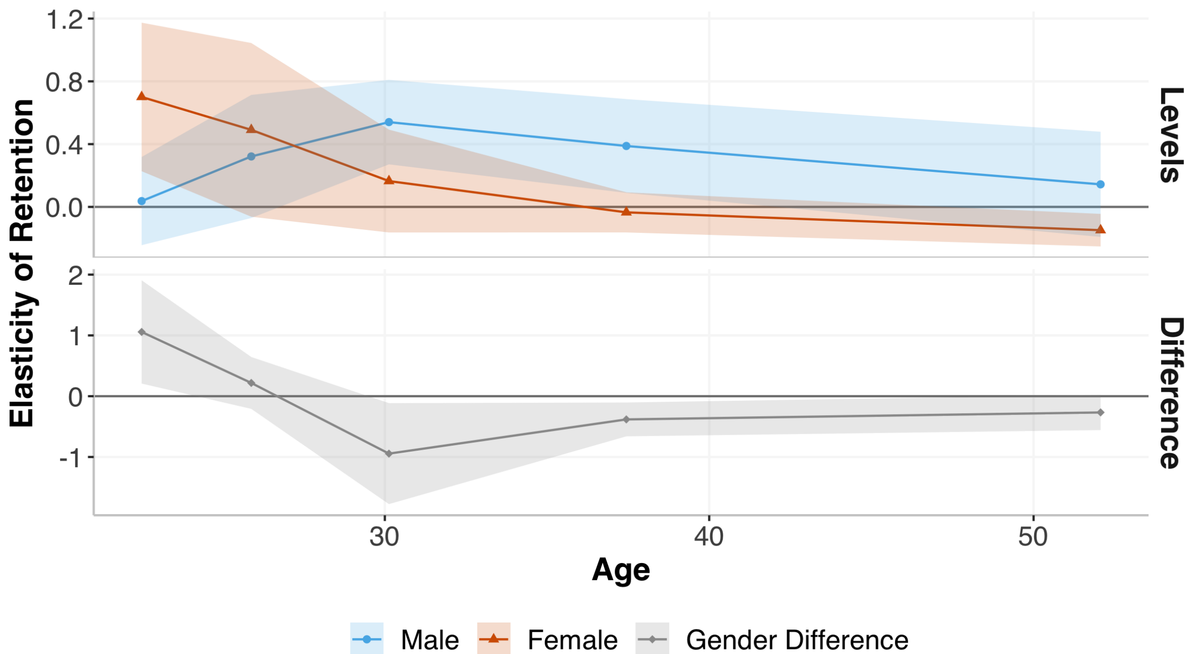

Figure 8: We utilize the stickiness of the retailer's wage to isolate variation in the retailer's relative pay based on the wage growth in the local MSA. We leverage this quasi-random variation in the retailer's relative pay within an MSA to identify the elasticity of retention with respect to relative pay for male and female representatives of varying ages. Specifically, we consider male and female representatives in the various quintiles of age. While, on average, the retention of male representatives is more sensitive to the relative wage than the retention of female representatives, this belies considerable variation over the life-cycle Early on, female representatives are significantly more responsive to the relative pay then men. As men enter their child-rearing ages, they become more responsive to pay. At the same time, female representatives become less responsive to pay, perhaps as they put more weight on the other attributes of the job, such as the potential to work remotely.

Appendix A: Additional Tables and Figures

Table A.1: Worker Quality Spillovers from High Roller'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: *p<0.1; **p<0.05; ***p<0.01

When pay at High Roller 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.2: Job Endings in Months and Locations where High Roller 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

[The yearly estimates will be aggregated.] To see whether the higher pay at High Roller 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 High Roller is hiring to other locations where High Roller'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: A Dynamic Model

We first show the intermediate steps deriving the results in Section 1 and then explore a dynamic model.

Intermediate Steps

Firms thus seek to minimize the cost of L units of effective units of labor:

$$\min_w wN + TNq(w,z) \text{ s.t. } \phi(w,z)N = L$$

We can then use the constraint of $\phi(w,z)N = L$ to substitute in for N. This gives us:

$$\min_w w\frac{L}{\phi(w)} + T\frac{L}{\phi(w)}q(w,z).$$

Taking the first order condition with respect to w gives us:

$$\begin{aligned}
\frac{L}{\phi(w)} - w\frac{L}{\phi(w)^2}\phi'(w) - T\frac{L}{\phi(w)^2}\phi'(w)q(w,z) + T\frac{L}{\phi(w)}q'(w,z) &= 0 \\
1 - \frac{w}{\phi(w)}\phi'(w) - T\frac{1}{\phi(w)}\phi'(w)q(w,z) + Tq'(w,z) &= 0 \\
w - \epsilon_{\phi,w}w - Tq(w,z)\epsilon_{\phi,w} + Tq'(w,z)w &= 0
\end{aligned}$$

If we define the elasticity of quits such that when wages go up, quits go down, then we flip the sign of $\epsilon_{q,w}$ and arrive at:

$$w = \frac{Tq(w,z)(\epsilon_{\phi,w} + \epsilon_{q,w})}{1 - \epsilon_{\phi,w}}$$

Dynamic Model

As above, assume that the firm sets the wage to minimize the discounted costs, C, of obtaining L units of effective labor. The firm discounts the future at rate δ. And the costs of procuring labor depend on the proportion of people retained from last period, $(1 - q(w_{t-1}, z))N$, which in turn depends on the pay rate in the previous period. We begin by assuming that the firm needs a constant amount of effective labor, L, from one period to another.[16]

$$C\left((1 - q(w_{t-1}, z))N_{t-1}\right) = \min_{w_t} w_t N_t + T\left(N_t - (1 - q(w_{t-1}))N_{t-1}\right) + \delta C\left((1 - q(w_t, z))N_t\right)$$

[16] If a firm is uncertain about whether it will need as much labor next period, it will be comfortable keeping wages lower since not all workers who leave will need to be replaced.

Again, we substitute using $L = \phi(w)N$:

$$C\left((1-q(w_{t-1},z))\frac{L}{\phi(w_{t-1})}\right) = \min_{w_t} w_t\frac{L}{\phi(w_t)} + T\left(\frac{L}{\phi(w_t)} - (1-q(w_{t-1}))\frac{L}{\phi(w_{t-1})}\right) + \delta C\left((1-q(w_t,z))\frac{L}{\phi(w_t)}\right)$$

We take the first order condition with respect to w_t:

$$\frac{L}{\phi(w_t)} - w_t\frac{L\phi'(w_t)}{\phi^2(w_t)} - T\frac{L\phi'(w_t)}{\phi^2(w_t)} - \delta C'q'(w_t,z)\frac{L}{\phi(w_t)} - \delta C'(1-q)L\frac{\phi'(w_t)}{\phi^2(w_t)} = 0$$

By the envelope theorem, $C' = -T$:

$$\begin{aligned}
\frac{L}{\phi(w_t)} - w_t\frac{L\phi'(w_t)}{\phi^2(w_t)} - T\frac{L\phi'(w_t)}{\phi^2(w_t)} + \delta Tq'(w_t,z))\frac{L}{\phi(w_t)} + \delta T(1-q)L\frac{\phi'(w_t)}{\phi^2(w_t)} &= 0\\
\frac{L}{\phi(w_t)}\left[1 - w_t\frac{\phi'(w_t)}{\phi(w_t)} - T\frac{\phi'(w_t)}{\phi(w_t)} + \delta Tq'(w_t,z) + \delta T(1-q)\frac{\phi'(w_t)}{\phi(w_t)}\right] &= 0\\
w_t - w_t\epsilon_{\phi,w} - T\epsilon_{\phi,w} + \delta Tq'(w_t,z)w_t + \delta T(1-q)\epsilon_{\phi,w} &= 0
\end{aligned}$$

We allow the sign for the elasticity of quits with respect to wages to be negative such that quits increase as wages decrease:

$$\begin{aligned}
w_t[1-\epsilon_{\phi,w}] - T\epsilon_{\phi,w} - \delta T\epsilon_{q,w}q(w_t) + \delta T(1-q)\epsilon_{\phi,w} &= 0\\
w_t[1-\epsilon_{\phi,w}] - \delta T\epsilon_{q,w}q(w_t) - T(1-\delta(1-q))\epsilon_{\phi,w} &= 0\\
T\frac{\delta q\epsilon_{q,w} + (1-\delta(1-q))\epsilon_{\phi,w}}{1-\epsilon_{\phi,w}} &= w_t
\end{aligned}$$

Note that if $\delta = 1$, this collapses to the static version. We can also rewrite this in terms of the effect of δ like so:

$$Tq(w_t)\left[\frac{\epsilon_{q,w} + \epsilon_{\phi,w}}{1-\epsilon_{\phi,w}}\right] - (1-\delta)T\left[\frac{q\epsilon_{q,w} - (1-q)\epsilon_{\phi,w}}{1-\epsilon_{\phi,w}}\right] = w \tag{7}$$

$$\frac{\partial w}{\partial \delta} = T\left[\frac{q\epsilon_{q,w} - (1-q)\epsilon_{\phi,w}}{1-\epsilon_{\phi,w}}\right]$$

The first term of (1) reflects the fact that patient firms will care more about the fact that lower wages will lead to greater quits and thus more training costs in future periods. This force will be especially strong when the quit rate is high in levels and especially responsive to the wage. The second term pushes in the opposite direction and suggests that more patient firms may have some incentive to pay their current workers less. The intuition for this is that paying the workers less today means the firm will employ more workers and thus have a larger pool of workers to draw upon in future periods. If the elasticity of the quit rate with respect to the wage is high and/or the quit rate is high in levels, then impatience will lead firms to pay less.

An alternative interpretation of $(1-\delta)$ is as a measure of uncertainty. If firms are not certain about the labor they'll need, they pay lower wages. Lower pay suits them because if they may need less labor in future periods, then slightly elevated quits would not make them pay higher training costs.

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.