

IZA DP No. 9149

**Fairness and Frictions:
The Impact of Unequal Raises on Quit Behavior**

Arindrajit Dube
Laura Giuliano
Jonathan Leonard

June 2015

Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior

Arindrajit Dube

University of Massachusetts Amherst and IZA

Laura Giuliano

University of Miami

Jonathan Leonard

University of California, Berkeley

Discussion Paper No. 9149

June 2015

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior^{*}

We analyze how quits responded to arbitrary differences in own and peer wages using an unusual feature of a pay raise at a large U.S. retailer. The firm's use of discrete pay steps created discontinuities in raises, where workers earning within 1 cent of each other received new wages that differed by 10 cents. First, we estimate a regression discontinuity (RD) model based on own wages; we find large causal effects of wages on quits, with quit elasticities less than -10. Next, we address whether the overall quit response reflects the impact of comparisons to market wages or to the wages of in-store peers. Here we use a multi-dimensional RD design that includes both a sharp RD in the own wage and a fuzzy RD in the average peer wage. We find that the large quit response mostly reflects relative-pay concerns and not market comparisons. After accounting for peer effects, quits do not appear to be very sensitive to wages – consistent with the presence of significant search frictions. Finally, we find that the relative-pay effect is nonlinear and driven mainly by workers who are paid less than their peers – suggesting concerns about fairness or disadvantageous inequity.

JEL Classification: J00, J31, J42, J63

Keywords: turnover, search frictions, fairness, quits, peer effects, regression discontinuity

Corresponding author:

Arindrajit Dube
Department of Economics
Thompson Hall
University of Massachusetts Amherst
Amherst, MA 01003
USA
E-mail: adube@econs.umass.edu

^{*} We thank Joshua Angrist, David Card, Christian Dustmann, Ethan Kaplan, Suresh Naidu and seminar participants at UC Berkeley, LSU, MIT, UMass Amherst, and the 2014 NBER Summer Institute for Personnel Economics for helpful comments and suggestions.

1 Introduction

When making decisions about job mobility, do workers compare their pay with that of their co-workers? Or do they mainly compare their pay with the outside market? If peer comparisons do matter, do they reflect other-regarding concerns such as status or fairness, or self-regarding behavior such as learning about the prospect for future raises? While economists have long been interested in these questions,¹ the endogeneity of wages has made causal evidence hard to establish.

We analyze how quits responded to arbitrary differences in own and peer wages among sales employees at a large U.S. retailer with hundreds of stores nationwide (henceforth “the firm”). Our research design is based on exogenous variation in wages that resulted from an unusual feature of the firm’s wage policy. In response to the federal minimum wage increases in 1996 and 1997, the firm implemented raises above the minimum wage that followed a uniform corporate policy irrespective of any local or individual characteristics other than initial wage. Moreover, this policy set the new wage as a step function of the initial wage, grouping employees into 15-cent wide bands and assigning them to the bottom of the same new pay step. This resulted in a set of 10-cent, discontinuous jumps in the new wage (one jump at the threshold for each new pay step). Thus workers on either side of a threshold whose initial wage differed by one cent received new wages that differed by ten cents.

We begin by estimating the elasticity of quits with respect to own wages using a sharp regression discontinuity (RD) design that exploits the multiple discontinuities in the new wage as a function of the initial wage. Our RD estimates suggest quits are highly responsive to wages—with elasticities of -21, -13 and -10 for 3, 6 and 9 months after the raise. However, while these estimates show a strong causal effect of wages, the interpretation of this effect is ambiguous. If workers make only market comparisons, then such a strong sensitivity of quits to own wages would indicate a highly competitive labor market. But the firm’s policy caused wages to change both relative to the market and relative to peers within the store. Workers just above a pay-step threshold ended up on average with a higher wage than their peers, while the opposite was true for those just below a threshold. If relative-pay concerns matter, then the RD estimates may reflect not only market comparisons but also peer comparisons.

In order to distinguish between these two explanations, we expand our model of quits to incorporate a peer-wage effect by exploiting discontinuities in the wages of one’s peers. Since most workers have multiple peers, we model quits as a function of the average peer wage; we then instrument this average with the wage discontinuity of a given peer using a fuzzy RD design.² Our main results are based on a multi-dimensional RD (MRD) model that

¹E.g. See Akerlof & Yellen (1990) for a review of the literature on relative pay both within and outside of economics; also, Luttmer (2005) discusses the interest by economists beginning with Adam Smith (1759).

²We define peers generally as co-workers earning broadly similar initial wages; our preferred specification uses wages within 20 cents of one’s own initial wage. The large number of stores ensures sufficient identifying

incorporates both the sharp RD in the own-wage (as before) and the fuzzy RD for the average peer wage. In this model, the own-wage estimate continues to give the total effect of wages on quits. But now the peer-wage estimate lets us recover the effect of a change in wages relative to one’s peers—i.e., the effect of peer comparisons. And by netting out the latter from the former, we get an estimate of the effect of a raise that is common across peers—i.e., the effect of market comparisons.

Overall, we find strong evidence that peer comparisons matter. Workers were much more likely to quit when their peers received bigger raises because the peers were just above a pay-step threshold. The MRD estimates imply that the 3-, 6- and 9-month quit elasticities with respect to the average peer wage are 20, 11 and 8. The MRD estimates of the own-wage elasticities, which are similar to the RD estimates, are -23, -15 and -11. If we compare the own-wage and peer-wage estimates, we see that the peer-wage elasticities are only slightly smaller in magnitude. Indeed the estimates imply that relative-pay concerns account for roughly 70-90% of the total effect of own wages on quits. On the other hand, market wage comparisons are relatively unimportant. The estimates suggest that when a raise is common across peers (e.g., an across-the-board raise in a firm), the elasticities of 3-, 6-, and 9-month quit rates are -4, -3, and -3, and are not statistically different from zero.

The MRD estimates imply that the overall effect of wages on quits is driven mostly by peer comparisons and not market comparisons. To check this implication, we estimate the own-wage RD model separately for two samples formed on the basis of whether the majority of one’s peers are on the same side or opposite side of a pay-step threshold. Consistent with the MRD results, these estimates show that the effects of own wages on quits are small in the same-side sample where the raises are similar across peers, but are large in the opposite-side sample where workers receive raises that differ from those of their peers.

Finally, we find that the relative-pay concerns are nonlinear. Specifically, we find that a worker’s quit behavior is highly responsive to the size of a peer’s raise only if the worker ends up being paid less than the peer. In contrast, when workers end up with wages that are either similar to or higher than those of their peers, we find no significant quit response. We interpret this nonlinear response as suggesting that the quit behavior is motivated by concerns about fairness. In particular, the concentration of the quit response among lower-paid workers is consistent with Fehr and Schmidt’s (1999) model of fairness as aversion to disadvantageous inequity, in which workers who earn less than their peers care most about relative pay.

Our paper relates to the literatures on both relative pay and search frictions. On the former, there are two broad theories about why peer wages may affect job mobility. Rational learning models assume workers are mainly self-regarding and have no direct concern about

variation across workers in the initial distribution of peer wages (i.e., in their positions relative to the nearest pay-step threshold). The small number of peers in most stores ensures a strong first-stage effect of each peer’s wage on the peer average.

the pay of their peers, but may use peer wages to help predict their own future pay. The standard learning model suggests that having peers with higher pay sends a positive signal about the opportunity for promotion within the firm, and hence that lower-paid workers will have higher job satisfaction and lower quit rates.³ Alternatively, other-regarding models assume that workers care directly about how their pay compares to that of their peers, and that lower-paid workers will have lower job-satisfaction and higher quit rates.

In the empirical literature, non-experimental studies have produced valuable evidence that supports both theories: some are consistent with rational learning (Clark et al. 2009; Galizzi and Lang 1998; and Pfeifer and Schneck 2012), and others are consistent with other-regarding preferences (Clark and Oswald 1996; Rege and Solli 2014). However, the endogeneity of wages makes it difficult for these studies to identify causal effects. Further, it is also hard for them to distinguish self-regarding motives from other-regarding concerns since peer wages may convey information about own wages. Such difficulties have encouraged experimental approaches, and our findings are perhaps most closely related to a field experiment by Card et al. (2012). This study looks at the effect of disclosing information about peer wages on job satisfaction and job search intentions among employees at the University of California. It finds the information had no effect on workers earning more than their peers, but led to lower job satisfaction and increased likelihood of job search for peers earning less. The authors thus find a causal effect of relative pay. They also find a nonlinear effect consistent with our results and with Fehr and Schmidt's model of fairness.⁴ Arguably, though, the study's findings could reflect rational learning as well as fairness concerns.⁵ Furthermore, while the study does look at actual separations, its evidence here is weaker.

Our paper also relates to the literature on search frictions. Models of monopsonistic competition (Burdett and Mortenson 1998; Manning 2003) show that if mobility decisions are based only on market comparisons, then the effect of own wages on separations can be used to assess the degree of frictions in the labor market. Using payroll data, several studies have estimated separation elasticities between -2 and zero, implying large search frictions (Ransom and Oaxaca 2010; Hirsch, Shank and Schnabel 2010; Sorenson and Depew 2013; Webber 2013). However, others have focused on pay variation due to collective bargaining agreements or policy inter-

³An alternate learning-based hypothesis is proposed by Buntrock (2014), who interprets a negative correlation between relative pay and quits as evidence that low relative pay can signal poor match quality. In his data workers in lower pay percentiles not only quit more, but also experience wage gains from a move.

⁴Nonlinear responses to relative pay have also been found in other experimental contexts. Cohn et al. (2014) use a field experiment to study productivity responses to randomly assigned wage cuts, and find that being paid less than one's teammate reduces a worker's productivity, but being paid relatively more has no effect. Similarly, in a laboratory experiment, Bracha, Gneezy and Loewenstein (2015) find that unequal pay affects labor supply (along the intensive margin) by reducing the labor supplied by lower-paid workers. However, the experimental literature on worker effort has also produced some contrary findings. In lab experiments, both Charness and Kuhn (2007) and Goreg, Kube and Zultan (2010) find that effort provision is very sensitive to the subjects' own wages but is not systematically affected by the wages of coworkers or team members.

⁵The authors discuss the difficulty of ruling out rational learning on p. 2984 (fn. 8) and p. 2997.

ventions, and have found elasticities ranging from -2 or -3 (Ransom and Sims 2010; Falch 2011) to -9 (Mas 2014).⁶ Again, a major challenge in this literature is the endogeneity of wages, and the difference in estimates may be partly due to different strategies for addressing this issue. But another potential factor is that the effect of own wages on separations may reflect not only market comparisons but also peer comparisons. If so, then for separation elasticities to be informative about search frictions, they must be identified using wage variation that is not only exogenous, but also common across all relevant peers.

In this paper, we use a unique source of variation in wages to identify the effects of both market competition and peer comparisons on quit behavior. In our context, we find that quits are highly responsive to wage increases and that this behavior is driven largely by relative-pay concerns. After accounting for peer effects, quits do not appear to be very sensitive to wages—consistent with the presence of significant search frictions. Finally, we find that the relative-pay effects are nonlinear and are driven mainly by workers who are paid less than their peers—suggesting that the quit behavior is motivated by concerns about fairness or disadvantageous inequity.

We advance the literature on relative pay in three ways. While recent field experiments have found causal effects of relative pay in the context of information shocks (Card et al. 2012) and nominal wage cuts (Cohn et al. 2014), we show that such effects generalize to the context of unequal *raises* in pay. Moreover, in support of earlier studies that found a link between relative pay and separations (e.g., Card et al. 2012; Rege and Solli 2014), we provide definitive evidence that relative pay can have an effect on job turnover. Finally, our finding of a nonlinear response is consistent with recent studies and with models of fairness. But further, as we discuss in section 7, an advantage of our setting is that the arbitrary nature of the firm’s pay policy makes it unlikely that our findings reflect any type of rational learning. Overall, we thus provide perhaps the best evidence to date that when workers make decisions about job mobility, they can be motivated by concerns about fairness.

Our findings also have implications for the wage-setting behavior of the firm. On one hand, the modest response of quits to an across-the-board wage increase is consistent with the presence of search frictions that give the firm significant monopsony power. On the other hand, fairness concerns among employees may serve as an important constraint on a firm’s wage-setting behavior beyond the constraints of the market. In particular, the concentration of the quit response to relative pay among lower-paid workers suggests that a more equitable

⁶Other labor market regulations, including minimum wages, have been used to assess labor market frictions. Though estimates of the employment response to a minimum wage increase are mixed, recent analyses of separations and hiring responses produce evidence consistent with search frictions in low-wage labor markets such as the retail setting we study here (Giuliano 2013; Dube, Lester and Reich 2014). Evidence of substantial monopsony power has also been found in studies of legislated wage changes for registered nurses (Staiger, Spetz and Phibbs 2010) and changes in mobility restrictions for migrant workers (Naidu, Nyarko and Wage 2015). In contrast, however, Matsudaira (2014) finds no evidence of monopsony power in a study of minimum staffing regulation for nurses aids.

pay distribution can lead to an overall reduction in turnover. The results also suggest that even modest pay differences can have big effects if they appear unfair.

The rest of the paper is structured as follows. In section 2, we extend the job ladder model to incorporate relative-pay concerns. In section 3, we discuss the institutional setting and our payroll data. In section 4, we provide the regression discontinuity estimates of the total effect of own wages on quit behavior. Section 5 extends the analysis to estimate peer effects, and assesses the relative importance of market competition and peer comparisons in determining the total own-wage effect. Section 6 presents a number of falsification tests, and section 7 concludes.

2 Theoretical Framework

To begin, we extend the canonical “job ladder” model of on-the-job search to include relative-pay concerns. We then show how the total effect of a raise on quits can be decomposed into: (1) a response to a common wage increase, and (2) a response to an increase in the wage gap between a worker and her peers.

Starting with Burdett and Mortensen (1998), the job ladder model has provided a useful way to model worker mobility. In this model, separations occur either as exogenous transitions to non-employment or as endogenous transitions to jobs offering wages that exceed the worker’s current wage w . The separation rate is given by $S(w) = \delta + \lambda [1 - F(w)]$. Here δ is the exogenous separation rate and λ is the offer arrival rate; search frictions are captured by lower values of λ .⁷ If we define “quits” as all job-to-job transitions plus a share θ of the exogenous separations, then the quit rate is $Q(w) = \theta \cdot \delta + \lambda [1 - F(w)]$.

A key assumption of this model is that wages vary across firms but not within firms. The quit response to a wage increase therefore depends only on market comparisons, and is given by $\frac{dQ}{dw} = \frac{dS}{dw} = -\lambda f(w)$. Hence in a context where wages vary only across firms, the quit response can be used to assess the degree of competition in the labor market. And under a stationarity assumption, the separation elasticity can be used to derive the extent of wage-setting (monopsony) power that arises from search frictions (Manning 2003).⁸

We expand the job ladder model to allow for internal wage variation and social comparisons by introducing a reference wage w_p , which is a function of the wages earned by one’s peers.

⁷Here we make the simplifying assumption that individuals are similar in terms of their offer arrival rates and wage offer distributions; i.e, that $\lambda_i = \lambda$, $F_i(\cdot) = F(\cdot) \forall i$. The corresponding assumption in our empirical model is that λ and $F(\cdot)$ do not change discontinuously at the firm’s pay-step thresholds.

⁸Specifically, Manning (2003) shows that if recruited and separating workers face the same offer wage distribution, then the labor supply elasticity facing the firm is -2 times the separation elasticity; he also derives the conditions under which this assumption holds.

We now assume that a worker's job satisfaction U depends not only on her own wage, but also on the gap between her own wage and the reference wage:⁹

$$U(w, w_p) = w - \alpha w_p = (1 - \alpha)w + \alpha(w - w_p) = (1 - \alpha)w + \alpha(w_g).$$

Here the parameter α reflects the strength of the relative-pay concerns and we assume that $0 \leq \alpha \leq 1$. When $\alpha = 0$, workers do not care about relative pay and the model reverts to the case with self-regarding preferences. And when $\alpha = 1$, workers care *only* about relative pay: an equal raise in both w and w_p that keeps the gap w_g constant does not improve the worker's welfare.

How does a worker choose between her current job and a new offer? There is no obvious rationale for a worker to expect her peers at the new job to be systematically paid more or less than herself. Therefore we assume the expected wage of peers at the new job is equal to the offered wage: $w'_p = w'$.¹⁰ Job-to-job transitions are now based on a comparison of U and U' , so the worker leaves when:

$$U = w - \alpha w_p = (1 - \alpha)w + \alpha w_g < (1 - \alpha)w' = U'.$$

Quits are now a function of both own wage w and the peer wage w_p :

$$Q(w, w_p) = \theta\delta + \lambda \left[1 - F \left(\frac{1}{1 - \alpha} (w - \alpha w_p) \right) \right]. \quad (1)$$

Differentiating equation (1) with respect to w and w_p gives the partial effects of own and peer wage increases:

$$\begin{aligned} \frac{\partial Q(w, w_p)}{\partial w} &= -\lambda \cdot f \left(\frac{1}{1 - \alpha} (w - \alpha w_p) \right) \cdot \left(\frac{1}{1 - \alpha} \right) \\ \frac{\partial Q(w, w_p)}{\partial w_p} &= \lambda \cdot f \left(\frac{1}{1 - \alpha} (w - \alpha w_p) \right) \cdot \left(\frac{\alpha}{1 - \alpha} \right) \end{aligned} \quad (2)$$

These two quit responses are estimated directly in our empirical analysis. Armed with these estimates, we can recover the parameter α measuring the strength of relative-pay concerns:

$$\alpha = - \frac{\frac{\partial Q(w, w_p)}{\partial w_p}}{\frac{\partial Q(w, w_p)}{\partial w}}. \quad (3)$$

⁹Our approach is similar to that used elsewhere in the literature on peer comparisons in the workplace. This formulation corresponds most closely to the symmetric preference case considered in Charness and Kuhn (2007). That paper, along with Fehr and Schmidt (1999) and Card et al. (2012), also allows for nonlinear effects when $w > w_p$, which we consider later in this section.

¹⁰While the assumption $w'_p = w'$ simplifies the exposition, it is not necessary for our key results on the identification of α or the quits elasticities. These rely only on the weaker assumption that any change in w or w_p does not affect the peer wage w'_p at the new job. To see this, note that workers move when $w - \alpha w_p < w' - \alpha w'_p$. If we interpret $F(\cdot)$ as the distribution function of the net outside wage ($w' - \alpha w'_p$), the expressions for the quit function (equation 1) remains the same, as do the expressions for all the subsequent quit elasticities.

We can also rearrange the terms in equation (1) to express quits as a function of the own wage w and the wage gap w_g :

$$Q(w, w_g) = \theta\delta + \lambda \left[1 - F \left(w + \frac{\alpha}{1-\alpha} w_g \right) \right]. \quad (4)$$

This formulation is useful because it lets us decompose the total effect of a wage increase on quits into two conceptually distinct responses—one based on market comparisons and the other based on peer comparisons. Totally differentiating equation (4) with respect to w gives $\frac{dQ}{dw} = \frac{\partial Q(w, w_g)}{\partial w} + \frac{\partial Q(w, w_g)}{\partial w_g} \frac{dw_g}{dw}$ where:

$$\begin{aligned} \frac{\partial Q(w, w_g)}{\partial w} &= -\lambda \cdot f \left(w + \left(\frac{\alpha}{1-\alpha} \right) w_g \right) \\ \frac{\partial Q(w, w_g)}{\partial w_g} \frac{dw_g}{dw} &= -\lambda \cdot f \left(w + \left(\frac{\alpha}{1-\alpha} \right) w_g \right) \cdot \left(\frac{\alpha}{1-\alpha} \right) \frac{dw_g}{dw}. \end{aligned} \quad (5)$$

The first term in equation (5) holds the wage gap constant and thus shows the impact of a wage increase that is common across peers. This “gap-constant” effect represents the quit response that is due to market comparisons. As in the standard model without peer effects, this response is stronger in more competitive markets with higher values of the offer arrival rate λ and the wage offer density $f(\cdot)$. The second term shows the impact of an increase in the wage gap—and thus represents the response that is due to peer comparisons. For a given level of market competition, this “relative pay” response is larger the greater is the concern about relative pay, α , and the larger the increase in the wage gap w_g .

This decomposition highlights the importance of accounting for peer effects when interpreting own-wage quit elasticities. We can see from the second term in equation (5) that if wage increases vary across peers, relative-pay concerns could cause the quit elasticity to be large even if λ is relatively small. Consequently, the total quit response to a change in own wage may not provide an accurate measure of monopsony power that arises from search frictions. In contrast, the “gap-constant” quit response is useful for assessing search frictions because it switches off the relative-pay channel.

Because raises in our setting generally cause wages to vary relative both to the market and to one’s peers, the partial derivatives in equation (5) are not directly estimable. However, they can be recovered from the estimates of the derivatives in equation (2). First, the relative-pay effect is identified as the negative of the peer-wage effect: $-\frac{\partial Q(w, w_p)}{\partial w_p}$.¹¹ From here, we can construct the gap-constant quit response by subtracting the relative wage effect from the total own-wage effect, or equivalently by summing the own-wage and peer-wage effects:

$$\frac{\partial Q(w, w_g)}{\partial w} = \frac{\partial Q(w, w_p)}{\partial w} + \frac{\partial Q(w, w_p)}{\partial w_p} \quad (6)$$

¹¹In section 5.4, we also construct a Wald estimate for the relative-pay (peer-wage) effect. This estimate is based on the difference in RD estimates for the total own-wage quit response in two sub-samples—one where raises differ across peers and one where the raises are more similar.

So far we have assumed away nonlinearities in the effect of peer wages on quits. But if workers are averse to disadvantageous inequity—as in Fehr and Schmidt (1999)—those with $w < w_p$ exhibit strong reactions from relative wage reductions, while those with $w \geq w_p$ exhibit smaller reactions to relative wage gains. We can incorporate such nonlinearities by modeling preferences as: $U(w, w_p) = v_0 w + v(w - w_p)$. Now the quit function can be written as $Q(w, w_p) = \theta\delta + \lambda[1 - F(v_0 w + v(w - w_p))]$. And the partial effects of own and peer wage increases are as follows:

$$\begin{aligned}\frac{\partial Q(w, w_p)}{\partial w} &= -\lambda f(v_0 w + v(w - w_p)) [v_0 + v'(w - w_p)] \\ \frac{\partial Q(w, w_p)}{\partial w_p} &= \lambda f(v_0 w + v(w - w_p)) [v'(w - w_p)].\end{aligned}$$

As before, when relative-pay concerns matter, $v'(w_g) > 0$ and quits respond to peer wages. When $v_0 = 1 - \alpha$ and $v'(\cdot) = \alpha$, we revert to the linear model. But if instead $v''(w_g) < 0$, then workers care more about relative pay when $w_g < 0$ than they do when $w_g \geq 0$. In section 5.6, we test for nonlinear preferences by splitting worker-peer pairs into four groups based on exogenous variation in the *ex post* wage gap, and by separately estimating workers' quit responses to peer raises in each of these groups.

Our theoretical framework does not allow for the possibility that peer wages affect quit behavior by providing a signal about one's own future wage. While such a learning mechanism may be relevant in other contexts, we think it is unlikely to explain the quit behavior in our case. We address this point more fully in section 7.

3 Data and Institutional Setting

3.1 The firm and its compensation policy

Our data is constructed from personnel records spanning the 30-month period from February 1, 1996, to July 31, 1998. The firm operated more than 700 retail stores nationwide during this period, and employed an average of 33 workers per store. We analyze the quit behavior of employees in an entry-level sales job that accounts for 90 percent of the firm's retail workforce. These employees all have the same job description, which involves customer service and support duties. The job requires only basic skills and employees receive cursory on-the-job training. Hourly wages are the main form of compensation, and employees do not receive commissions or performance-based bonuses. These are relatively low-wage jobs with little expected wage growth. Promotions are rare. Among those who remain employed, less than 5% are promoted to a higher-paid job within a year of being hired. The main opportunity for wage growth in the firm is through merit raises, which are given annually. All those employed for at least 90 consecutive days are eligible for the annual merit raise, and approximately 80% of eligible employees receive one. These raises are determined by store managers and average 2.2%.

Our analysis focuses on a non-standard set of raises that the firm implemented in response to increases in the federal minimum wage. The minimum wage rose twice during our sample period—from \$4.25 to \$4.75 on October 1, 1996, and then to \$5.15 on September 1, 1997. Unlike decisions about starting wages and merit raises, wage adjustments following the minimum wage increases were set centrally by a corporate policy that was uniformly implemented at all stores. The share of hourly retail employees who earned less than the new minimum was roughly 5% in 1996 and 10% in 1997. However, the firm’s policy increased wages substantially more than was necessary to comply with the law, extending raises to the 30th percentile of the wage distribution in 1996 and to the 40th percentile in 1997.¹² These new wages were applied to all hourly employees irrespective of any local or individual characteristics other than initial wage.

A key feature of the firm’s policy—and the source of the arbitrary variation exploited in our analysis—is the discontinuous nature of the formula used to implement the raises. For a worker with wage w_{0y} before the minimum wage increase in year y , the scheduled raise, Δw_y , was calculated as:¹³

$$\Delta w_y = w_y - w_{0y} = \begin{cases} (MW_{1y} - w_{0y}) + 0.10 \times \text{int} \left(\frac{w_{0y} - MW_{0y}}{0.15} \right) & \text{if } w_{0y} \in [MW_{0y}, \bar{w}_{0y}) \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

Here MW_{0y} is the initial minimum; MW_{1y} is the new minimum; and \bar{w}_{0y} represents the maximum initial wage for which there is a raise, and is equal to \$5.45 (1996) or \$5.65 (1997). In each year, the resulting new wage schedule is a step function with new pay steps at 15-cent intervals of w_{0y} within the indicated range. These wage schedules are illustrated in Figure 1, which shows scatter plots of new (“day after”) wages on initial (“day before”) wages for all hourly employees who had at least one month tenure and a wage less than \$1.00 above the new minimum on the day before the increase.¹⁴

As Figure 1 illustrates, the policy created multiple discontinuities in the relationship between the initial wage and the new wage—one at the bottom of each new pay step. The number and location of the thresholds, T_y^k , varies by year: there are seven thresholds in 1996 and five in 1997. But in all cases, the initial wages for employees on either side of a threshold

¹²There is little direct evidence on the extent to which minimum wage increases result in wage spillovers, and measurement error makes it hard to quantify spillovers using household data such as the Current Population Survey (Autor, Manning and Smith 2015). It is thus noteworthy that our firm implemented sizeable spillovers as a matter of corporate policy, giving raises to workers earning as much as 15 percent above the new minimum.

¹³We refer to the raise determined by equation (7) as the “scheduled raise” to distinguish it from the actual raise. The actual raise may be different if, for example, the employee receives a promotion on the same day. In practice, however, fewer than .5% of raises differ from the scheduled raise, so the scheduled raise predicts the actual raise very closely (see Figure 2 below).

¹⁴In practice, employees hired less than a month before the minimum wage increase were often paid starting wages based on the new pay scale and hence they did not receive subsequent raises. We exclude these new hires from our estimation sample.

differed by one cent, while their new wages differed by ten cents. The raises, therefore, differed by \$0.09 or roughly 2% of the typical wage. These arbitrary differences in raises are the basis for our RD design.

3.2 Sample construction and descriptive statistics

In the job we analyze, there were 10,390 workers at wages scheduled for raises on October 1, 1996, and there were 13,548 such workers on September 1, 1997. Our estimation sample is a subset of these employees that meets two conditions. First, we condition on $w_{0y} \in [MW_{0y} + .08, \bar{w}_{0y} - .08]$ and thus exclude wages near the endpoints of the range in which there was a scheduled raise. This ensures a consistent bandwidth of \pm \$.07 around each discontinuity threshold in the RD model.

Second, we exclude workers who are still earning their starting wages, and thus restrict attention to those who had received a merit raise during the previous raise cycle. This is important because our RD design rests on the assumption that wages are exogenously determined in the vicinity of the discontinuity thresholds. Since starting wages for new hires are always multiples of \$0.05, and since step thresholds occur only at multiples of \$0.05, starting wages can be located at pay-step thresholds but are never located just below a threshold. If we were to include starting wages in the sample, this would lead to sharp discontinuities in employee tenure and other associated characteristics.¹⁵

The final estimation sample consists of 6,691 scheduled raises. Wages and scheduled raises are summarized in Panel A of Table 1. Employees in our sample earned an average of \$4.99 and \$5.28 before the 1996 and 1997 raises, and received average raises of \$0.21 and \$0.18. For the pooled sample, the average initial wage is \$5.15 and the average raise is \$0.19.

Apart from wages, the data contains an employee’s age, race, gender, and full or part time status. For employment spells that begin or end within our sample period, we also observe dates of hire or termination; and for terminations, we observe the reason. Panel B shows the characteristics of our estimation sample. The sample is largely female (81%) and white (76%), and is relatively young—the mean age is 23 and about half are teenagers. Less than 1% work full time. Since we don’t observe hire dates before Feb. 1, 1996, and since sample employees are hired before April 1st of each year, our measure of tenure is censored at 8 months for 86% of those employed on the date of the first increase (October 1, 1996). Of those employed on the second date (September 1, 1997), tenure is censored for only 16% and the median tenure is 11.7 months. Because the sample is limited to employees who have both a scheduled raise and a prior merit raise, it consists of relatively low-wage earners with relatively high job tenure.

¹⁵Since merit raises are given annually at the end of June and eligibility requires at least 90 days of tenure, this sample restriction excludes employees who were hired after April 1st of the year of each minimum wage increase. It also excludes about 15% of the remaining sample because despite being eligible, these employees did not receive the most recent merit raise.

Each employee in the estimation sample is linked to two other sets of employees that are important for our analysis: coworkers and peers. Coworkers are defined as those who are in both the same job and the same store as the sample employee on the day of the minimum wage increase. A typical employee in our sample has 27 coworkers, who on average are slightly older than sample employees (24.2 vs. 22.6) and earn somewhat higher wages (\$5.55 vs. \$5.15). Coworker characteristics are summarized in Panel C of Table 1. We use these characteristics as controls in certain model specifications to demonstrate the robustness of our results.

We define a worker’s peers as the subset of all coworkers whose initial wage is: (1) in the range for a scheduled raise, and (2) within a fixed distance from the worker’s own initial wage. Our preferred specification uses a \pm \$.20 wage band to define peers. Based on this definition, panel D of Table 1 summarizes the characteristics of peer groups for our sample. On average, employees have 7.2 such peers, and the average peer wage (\$5.15) and average peer raise (\$.20) are both very close to the means for sample employees.

Our outcome of interest is the probability of quitting within a window following one of the minimum wage increases. We examine windows of 1, 2, 3, 6 and 9 months.¹⁶ Quits are defined as voluntary separations that occur for job-related reasons (leaving for another job, dissatisfaction, or simply not showing up); quits comprise 45% to 55% of all separations within the time frames considered. For employees who separate for other reasons (including moving, returning to school, transferring to another store, or being fired), quit decisions are treated as censored. Panel E of Table 1 shows the quit rates among the non-censored observations in our sample. The probability of quitting ranges from 6% within one month of a wage increase to 36% within 9 months.

4 Regression Discontinuity Estimates of the Quit Response to an Increase in Own Wage

This section presents an RD analysis of the effect of own-wage increases on quit decisions. We will build on this framework in the next section when we present the MRD model that incorporates peer effects.

4.1 RD estimation framework

The RD design exploits the variation in wages that results from discontinuities in the scheduled raise formula shown in equation (7). We therefore begin by examining the relationship between the scheduled raise and the actual raise received on the date of each minimum wage increase.

¹⁶Although our data set extends through July 1998 (11 months after the 1997 minimum wage increase), we restrict attention to intervals of 9 months or less, because the relationship between the scheduled raise and wages is attenuated by merit raises given in June 1997 and in June 1998 (9 and 10 months, respectively, after the minimum wage increase).

For each date, Figure 2 plots the mean observed raise against the initial wages of employees in the estimation sample. It also shows fitted values from regressions of the observed raise on the scheduled raise. The scheduled raise predicts the actual raise very well. Consistent with the formula, the actual raise declines linearly with the initial wage except for the positive 10-cent discontinuity at each 15-cent interval, and the plotted points deviate very little from this pattern.

Since scheduled and actual raises are highly correlated, we use the “scheduled wage” (obtained by applying the scheduled raise to the initial wage) as a proxy for the actual wage, and we treat the own-wage RD design as “sharp.” A potential concern with this approach is that actual wages may increase over time due to raises from promotions (or, in one case, state law). If so, then our estimates will reflect the quit response to wages on the day of the minimum wage increase, but may understate the quit response to current wages. To investigate this issue, we estimated discontinuities in current wages for those still employed between 1 and 9 months later. The impact of the scheduled raise is quite persistent; even after six months, a \$.10 discontinuity in the scheduled raise predicts a discontinuity of \$.095 in the actual wage. By 9 months, the wage discontinuity shrinks by about 20% but it is still highly significant.¹⁷

To gain power when estimating the discontinuities in quit rates, we use a parametric framework and pool data from the 1996 and 1997 raises. In particular, we estimate models of the form:

$$Q_{iy}^m = \beta \times w_{iy} + f_y(w_{0iy}) + X_{iy}\Gamma + \lambda_{z(i)} + \epsilon_{iy}, \quad (8)$$

where the dependent variable, Q_{iy}^m , is an indicator for whether employee i quits within m months of the year y raise; $w_{iy} = w_{0iy} + \Delta w_{iy}$ is the scheduled new wage; and $f_y(w_{0iy})$ is a smooth function of the initial wage.¹⁸ Because the scheduled wage is a close proxy for the actual wage, we can interpret β as the effect of a wage increase on the quit decision. If the specification for f_y is sufficiently flexible, then β is estimated using only the discontinuities in the scheduled wage function. Our baseline model uses a linear specification for f_y and allows a different intercept and slope in each year; we also document robustness using quadratic and cubic functions. Equation (8) also includes a set of fixed effects, $\lambda_{z(i)}$, for 3-digit ZIP codes

¹⁷As noted above, our largest quit window of 9 months does not include the subsequent merit raise cycle; the wage increases after 9 months are due mainly to promotions and to the California minimum wage increases in March 1997 and 1998, which affect about 5% of the sample. We have also estimated wage discontinuities at 10 months (after the merit raise cycles in both years), and these are attenuated by another 10%. However, we find no evidence that managers adjust merit raises to compensate employees for previous “bad luck”; the scheduled raise does not predict the size of the future merit raise.

¹⁸Our research design is similar to Angrist and Lavy (1999) and Angrist, Battistin and Vuri (2014). These studies use the discontinuous Maimonides’ rule as an instrument for actual class size to estimate the impact on educational outcomes. Like their Maimonides-predicted class size, our scheduled wage is a discontinuous function of a running variable with multiple thresholds. And like them, we use the discontinuous treatment variable as a regressor while controlling for a smooth function of the running variable (the initial wage in our case, total enrollment in theirs).

based on the store’s location (hereafter “ZIP codes”) and a vector of additional controls, X_{iy} , that varies by specification.¹⁹ These control variables are discussed below in section 4.3. The standard errors in all specifications are clustered by store.²⁰

4.2 Addressing threats to internal validity

4.2.1 Manipulation of wages

In an RD framework, identification requires that assignment of the running variable be as good as random in a window around the discontinuity threshold. In our context, this means that employees with initial wages just above and below a pay-step threshold must not differ systematically with respect to their latent propensities to quit.

One potential concern is that wages could be precisely manipulated by managers. For example, it would lead to bias if managers anticipated the raise schedule and “topped up” the merit raises of their most valuable employees to ensure they would be bumped up to a higher pay step when the minimum wage increased. Such manipulation would result in bunching of wages at the thresholds for the new pay steps (Lee and Lemieux 2010; McCrary, 2008).

Appendix Figure 1 shows the frequency distribution of wages in each year of our analysis sample. Even though the sample excludes employees who are still earning their starting wages, the distribution is still quite lumpy with bunching that appears mainly at multiples of \$.05. To assess this bunching, we pool all the 15-cent wide segments that are centered on a pay step threshold in either year, and in Appendix Figure 2 we plot the frequency distribution of the distance from the nearest threshold. There are clear spikes not only at the thresholds, but also at wages that are \pm \$.05 from the thresholds. Since manipulation would cause spikes at the thresholds but not elsewhere, this pattern is inconsistent with manipulation as the main cause of the bunching.

Appendix Figure 3 provides evidence against manipulation of merit raises in particular. The figure shows the average size of the merit raise by the distance from the threshold. The average merit raise is very similar on both sides of the threshold, and thus it does not appear that merit raises have been topped up in order to ensure that favored employees receive bigger raises.

¹⁹We use 3-digit ZIP codes to control for market heterogeneity while also maintaining sufficient within-market variation in peer groups to estimate peer effects with reasonable precision. Within-ZIP code variation includes both variation across stores that share 3-digit ZIP codes and within-store variation across the two years and multiple peer groups. Our analysis sample contains 375 3-digit ZIP codes and each ZIP has an average of 3.3 stores. We obtain broadly similar results from models that use wider (e.g., region) or narrower (e.g., 5-digit ZIP code) geographic controls.

²⁰The standard errors change very little if we define clusters more conservatively at the level of the ZIP code or state, or if we cluster by the discrete values of the running variable.

4.2.2 Bunching for other reasons

While manipulation of wages seems unlikely, the bunching in the initial wage distribution is nevertheless inconsistent with random assignment. Though the reasons for the bunching are not clear, the pattern suggests that merit raises are often given in \$.05 increments. This tendency could lead to bias in our RD estimates if it is correlated with characteristics of the employee, store, or labor market that determine quit rates. We explore this issue in Appendix Table 1, which shows discontinuity estimates from regressions of several employee and coworker characteristics on the scheduled wage. Column 1 reports estimates from specifications that control only for a linear function of the initial wage. Here we find significant discontinuities in employee age and tenure and in several coworker characteristics, confirming that the assignment of wages to either side of a threshold is not random.

To avoid potential bias from such non-random sorting, our baseline specification of equation (8) controls for two key sources of heterogeneity. First, unobserved labor market characteristics are captured by the ZIP code fixed effects, λ_z . Second, to directly control for variation associated with the bunching of wages, we include a dummy variable for initial wages that are multiples of \$.05. The estimated discontinuities in observable characteristics from this baseline specification are reported in column 2 of Appendix Table 1. Reassuringly, the discontinuities in employee age and in coworker characteristics are eliminated, and only a small discontinuity in tenure remains.²¹

To assess the possibility of remaining bias, we present models that include flexible controls for tenure plus other employee and coworker characteristics. Our results are generally robust to these controls.²² We also demonstrate robustness using “donut-hole” specifications that exclude all initial wages that are multiples of \$.05 (Barreca et al., 2011). And in section 6, we present two types of falsification exercise that provide additional validation of our results.

4.2.3 Differential attrition

Another concern is the possibility of differential sample attrition due to endogenous competing risks. When analyzing the probability of quitting within a given window, our analysis sample excludes employees who left the store for another reason. The estimated effects on quitting would therefore be biased if the scheduled wage affects the non-quit separation rate. We

²¹The table reports coefficients on the scheduled wage; these must be rescaled to obtain the change associated with a \$0.10 discontinuity in the wage. For example, the coefficient of 2.98 from the tenure regression corresponds to a discontinuity of .3 months or about 9 days.

²²Specifically, we include a dummy for each month of tenure. Employee characteristics are described in Panel B of Table 1 and include: age and age-squared, gender and race dummies, an indicator for full-time status, size of the most recent merit raise, and the median household income in the employee’s residential ZIP code. Coworker characteristics (see Table 1, Panel C) include: total number of entry-level employees on the day of the minimum wage increase, average employee age, average employee wage, the fraction who received a scheduled raise, the fraction who received a merit raise in July of the same year, and the fraction whose initial wage is a multiple of \$.05.

address this issue by estimating models for the probability of being in the analysis sample for each of our five quit windows. The results, reported in Appendix Table 2, show no evidence of selective attrition. We conclude that the censoring of quit outcomes for employees with non-quit separations does not pose a concern.

4.3 Estimates from RD models of quit behavior

Table 2 presents the estimated values of β (the effect of own wage on the quit decision) from various specifications of equation (8). Each column corresponds to a different window for the quit rate, while each row reports a different model specification. The first row shows a specification that controls linearly for the initial wage but does not include any additional controls. These estimates are all negative; they imply that a \$.10 wage increase is associated with a 1.4 - 4.4 percentage point fall in the quit rate; and in all but the first column (the model for 1-month quits) they are at least marginally significant. However, these estimates may be biased due to heterogeneity associated with the observed bunching of wages.

In the next two specifications, we first add ZIP code fixed effects (row 2) and then a dummy for initial wages that are multiples of \$.05 (row 3). The coefficients from these models are again negative, but are larger in magnitude than those in row 1 – confirming that these controls capture important sources of heterogeneity. The model in row 3 serves as our baseline model, and the estimates here imply that a \$.10 increase in the wage results in a 2.4 – 7.0 percentage point reduction in the quit rate—depending on the time window. Notably, the coefficients double in magnitude between the 1 and 2-month windows (from .24 to .51), but they increase by only another 40% between 2 and 9 months. This pattern implies that most of the quit behavior caused by the wage increase occurs within 2 months of the raise.

Figure 3 presents a graphical analysis of the baseline model. While there are 12 different wage thresholds, our baseline regression controls linearly for the running variable w_{0iy} and estimates a single β by pooling across the different thresholds. To simplify the visual presentation, we can therefore use a “representative” 15-cent interval and a normalized running variable, r_{iy} , that is defined as the distance from the representative threshold.²³ For the representative interval, we plot the fitted relationship between the normalized running variable and the residuals from a regression of quits on our baseline controls.²⁴ Since the quit rate varies across different wage intervals and different quit windows, we normalize the fitted values so they are equal to zero at the left limit at the threshold. This normalization has the desirable property that the right limit at the threshold represents the effect of a \$.10 discontinuity in the

²³Formally, the normalized running variable is defined as $r_{iy} = w_{0iy} - T_{iy}^k$, where $T_{iy}^k = \arg \min_{T_y^k} |T_y^k - w_{0iy}|$ is the wage threshold nearest to w_{0iy} . Workers are considered “above” the threshold when $w_{0iy} \geq T_y^k$, or equivalently $r_{iy} \geq 0$. They are considered “below” when $r_{iy} < 0$.

²⁴We use a slightly modified version of the “baseline” model in row 3 of Table 2 that constrains the slope on the initial wage to be the same across the two years—i.e., $f_y(w_{0iy}) = f(w_{0iy})\forall y$. However, this restriction makes very little difference to the estimation of β as the slopes never differ significantly between the two years.

wage.²⁵ Additionally, we plot the mean residuals from the baseline quit regressions, averaged across all 12 intervals, for each value of r_{iy} . These mean residuals help assess the overall fit of the data to the linear model. Aside from deviations at $-\$0.04$ (the most thinly populated bin), Figure 3 shows that the data fit the linear model reasonably well. And except for the 1-month quit rate, which shows an anomalous deviation just below the threshold, the plotted points show clear evidence of discontinuities.

Returning to Table 2, rows 4-7 report various sensitivity tests for the baseline model in row 3. In row 4 we show that the estimates are nearly identical when we use a quadratic specification for the initial wage function $f_y(w_{0iy})$. We also obtain very similar results when using a cubic specification (results not shown). Row 5 adds a dummy for each month of tenure, and other controls for employee and coworker characteristics. The estimates grow slightly in magnitude – suggesting that if anything the baseline estimates are slightly biased toward zero due to remaining heterogeneity.²⁶

We can control for all fixed heterogeneity at the store level by replacing the ZIP-code fixed effects in row 5 with store fixed effects. Row 6 shows the estimates from such a specification, and these are even larger than those in row 5—increasing in magnitude by about 12-32%. One interpretation of this finding is that even with detailed controls for employee and coworker characteristics as in row 5, the estimates are biased toward zero due to omitted variables. An alternate interpretation, which is more consistent with the findings in the next section, is that the estimates in row 6 are identified using within-store wage variation, and quit behavior is especially responsive to within-store variation because of relative-pay concerns.

Finally row 7 shows the “donut-hole” specification, which is similar to the baseline model except that instead of including a dummy variable for initial wages that are $\$.05$ multiples, it excludes these wages from the estimation sample. The estimates are again quite similar to those in the baseline model, though the standard errors are 25-40% larger. Overall, then, the estimates in Table 2 imply significant negative effects of own wages on quit rates.

5 Peer Comparisons: The Effects of Unequal Raises

In the absence of relative-pay concerns, the large effects of own wages on quit behavior would suggest that workers are very responsive to employer differences in wages – consistent with a highly competitive labor market. However, we have seen that the quit responses are stronger

²⁵This is equivalent to the scheduled wage coefficient scaled by .10 (the size of the wage discontinuity).

²⁶Because our measure of tenure is censored at 8 months for 86% of the 1996 sample, we perform additional tests to assess whether unobservables related to tenure are likely to impart bias. We estimate quit models separately by year and compare models with and without controls for tenure. Tenure is censored for only 16% of the observations in 1997; hence if tenure is an important confounder, the estimates should be more sensitive to the model specification in the 1997 subsample. We find that the estimates are similarly robust in both years: in all subsamples, excluding the tenure controls from the model causes the coefficients to change by .01 percentage point or less.

when estimated using models with store fixed effects—i.e., when the identifying variation comes from employees who work in the same store. This suggests that the quit behavior may be driven at least partly by wage comparisons among in-store peers. In this section we assess the “peer effect” channel directly by asking how quit behavior responds to arbitrary variation in peer wages, and we assess the relative importance of market competition and peer comparisons in determining quit behavior.

5.1 Measuring peer wages

To move from our theoretical model with unspecified reference wage w_p to an empirical model of peer effects, we must specify the definition of a peer and the function used to aggregate peer wages. We define a worker’s peers as the subset of coworkers whose initial wages were both in the range for a scheduled raise and within a wage band defined by a fixed distance from the worker’s own initial wage. For each worker, we then construct the average scheduled wage of the worker’s peers. Though we give equal weight to all peers in the group, we determine the appropriate size of the peer group on the basis of an analysis that allows the size to vary (see section 5.5). Cornelissen, Dustmann and Schönberg (2013) show that choosing a definition that is too exclusive (e.g., a wage band that is too narrow) leads to attenuation bias in the peer effect estimate. Using a more inclusive peer measure (e.g., a wider wage band) reduces this bias. But an increase in the number of peers also reduces identifying variation in the average peer wage due to the law of large numbers, and hence larger peer groups lead to less precision. Based on a comparison of estimates for peer groups of different sizes, our preferred specification uses a wage band of \pm \$.20 (about 5%) of one’s own wage. Thus defined, peer groups average 7 peers per worker, and they include 60% of coworkers with scheduled raises and 25% of all coworkers in the store.

5.2 Identifying peer effects

To identify the effect of peer wages on quit behavior, we again exploit discontinuities in the scheduled raise formula in equation (7). For intuition, suppose that for each worker i in each year y , we randomly select a single peer j with initial wage w_{0jy} . If we apply the scheduled raise formula to w_{0jy} , this creates discontinuities in the peer’s scheduled wage, w_{jy} (just as it does in the own wage). Our original RD model can be extended to include a peer effect component as follows:

$$Q_{iy}^m = \beta \times w_{iy} + \delta \times w_{jy} + f_y(w_{0iy}) + g_y(w_{0jy}) + X_{ijy}\Gamma + \lambda_{z(i)} + \epsilon_{ijy} \quad (9)$$

Equation (9) is a multi-dimensional RD (MRD) that uses smooth functions f_y and g_y of initial own wage, w_{0iy} , and initial peer wage, w_{0jy} , respectively, as running variables. As in equation (8), a flexible specification for f_y ensures that the identifying variation used to estimate β

comes only from the discontinuities in w_{iy} at the pay-step thresholds. Similarly, controlling for a flexible function $g_y(w_{0jy})$ ensures that the coefficient δ on the peer scheduled wage is identified using only the discontinuities in w_{jy} . In this single-peer case, the scheduled peer wage w_{jy} is a proxy for the actual peer wage and the peer-wage RD is sharp. The coefficient δ thus represents worker i 's quit response to the wage increase of a randomly selected peer.

As this “single peer” example makes clear, discontinuities in single peer wages provide the identifying variation we need to estimate the peer effect. However, most workers in our sample have multiple peers, and our behavioral model assumes that workers respond to the wages of all peers. To aggregate the peer effects, we model the quit decision as a function of the average peer wage, $\bar{w}_{p(i,y)y} = \frac{\sum_{j \in p(i,y)} w_{jy}}{N_{p(i,y)}}$, where $N_{p(i,y)}$ is the number of peers of worker i . We then instrument $\bar{w}_{p(i,y)y}$ with the discontinuity in w_{jy} by replacing the sharp RD for w_{jy} in equation (9) with a fuzzy RD for $\bar{w}_{p(i,y)y}$. The two-stage specification is as follows:²⁷

$$\begin{aligned} \text{1st stage : } \quad & \bar{w}_{p(i,y)y} = \gamma \times w_{iy} + \gamma_p \times w_{jy} + \tilde{f}_y(w_{0iy}) + \tilde{g}_y(w_{0jy}) + X_{ijy}\tilde{\Gamma} + \tilde{\lambda}_z(i) + \tilde{\epsilon}_{ijy} \\ \text{2nd stage : } \quad & Q_{iy}^m = \beta \times w_{iy} + \beta_p \times \bar{w}_{p(i,y)y} + f_y(w_{0iy}) + g_y(w_{0jy}) + X_{ijy}\Gamma + \lambda_z(i) + \epsilon_{ijy} \end{aligned} \tag{10}$$

The first-stage coefficient γ_p gives the effect of a single peer-wage discontinuity on the average peer wage. A strong first-stage relationship is ensured by the small number of peers in most peer groups. The second-stage equation contains our two parameters of interest. First, as in equation (8), β continues to represent the total effect of the own wage on quit behavior. The assumed exogeneity of w_{iy} implies that the estimate for β should be robust to the inclusion of peer wages and the other peer controls, and should therefore be similar across the RD and MRD models. Second, the MRD model also gives us an estimate for β_p —the quit response to an increase in the average wage of one’s peers. With estimates for β and β_p , we can use equation (3) to construct our measure of the weight placed on relative pay in the utility function: $\alpha = -\frac{\beta_p}{\beta}$. We can also sum β and β_p to estimate the “gap-constant” quit response in equation (6)—and thereby evaluate the role of market comparisons.

While equation (10) is estimable using a randomly selected peer for each worker, we can gain efficiency by using all of a worker’s peers. To do this, we stack the data by pairing each worker i in year y with all possible peers $j \in p(i, y)$. The stacked dataset replicates each observation in the original dataset $N_{p(i,y)}$ times; hence, to ensure the results are representative of our initial estimation sample, we weight each observation by $\frac{1}{N_{p(i,y)}}$. By clustering the

²⁷Our two-stage MRD design is most similar to Feigenbaum (2014), who studies the effect of unions on political outcomes at the county level and instruments the number of unions by close NLRB elections using elections stacked by year. Methodologically, our paper adds to the growing literature on aggregating across multiple discontinuities (e.g., Papay et al., 2011; Reardon and Robinson, 2012).

standard errors at the store level, we account for the impact that repeated observations may have on our statistical inference.²⁸

In our baseline specification for equation (10), we continue to include a set of ZIP-code dummies and a dummy for workers earning multiples of \$.05; we also include a 5-cent wage dummy for peer j and a dummy for the peer having received a merit raise. We document the robustness of our results to the inclusion of additional covariates X_{ijy} , which are composed both of individual characteristics of worker i and peer j and also of store-wide coworker characteristics.

5.3 Peer effects main estimates and decomposition of own-wage effect

Table 3 reports results from the two-stage MRD model.²⁹ Row 1 shows the baseline specification. Row 2 adds controls for individual worker characteristics and store-averaged coworker characteristics; and row 3 adds peer characteristics and higher-order terms in peer wages. We continue to find strong effects of own wages on quits across all three specifications. The estimates of β from the models for 2-, 3-, and 6-month quits range from -0.53 to -0.75 and all are significant at the 1 percent level. As expected, these results are similar to the estimates from the univariate RD (Table 2). But in addition, the MRD estimates also show sizable effects of the peer average wage. The estimates of β_P are only a little smaller in magnitude than the estimates of β , but are the opposite in sign—implying that workers respond to higher peer wages with a substantially higher probability of quitting. The coefficients from the 2-, 3- and 6-month quit models range from 0.44 to 0.77; most are statistically significant at the 5 percent level, and all are statistically significant at the 10 percent level. The 9-month coefficients are similar to the 6-month estimates but are less precise. All coefficients are robust to the inclusion of employee and coworker characteristics (row 2), and are even larger with controls for peer characteristics and a quadratic function of the peer’s initial wage (row 3).

Figure 4 provides graphical confirmation of these results. As in Figure 3, we show representative intervals of own initial wage and peer initial wage; and we use normalized running

²⁸As an alternative to stacking the data, we could use worker-level data to estimate an MRD model that includes each of the peer’s initial wages as a separate running variable. Under the assumption that the function of the running variable has a polynomial form, these peer running variables can be aggregated. However, when the number of peers varies by worker, so do the coefficients on the aggregated running variables. In our data $N_{p(i,y)}$ takes on 36 different values. Hence this alternative approach requires estimating 36 different functions of initial wages, which substantially reduces statistical power.

²⁹The first-stage and the reduced-form estimates are reported in Appendix Table 3. The first-stage estimates suggest a very strong relationship between the wage of a specific peer and the average peer wage, with a coefficient of around 0.44 and a t-statistic exceeding 40. We also test for discontinuities in predetermined worker characteristics at the peer wage thresholds. These results are reported in Appendix Table 4. Reassuringly, most covariates do not show any discontinuities. We do find a small but statistically significant negative estimate for the size of the previous merit raise, and a positive estimate for the fraction of coworkers in the store who received a merit raise. However, as we show in Table 3 Panel A, controlling for employee and coworker covariates—including the size of one’s merit raise and the proportion of coworkers receiving a merit raise—does not appreciably affect the peer-wage coefficient in the MRD model.

variables, r_{iy} and r_{jy} , defined as the distance from the respective initial wage (w_{0iy} or w_{0jy}) to the nearest discontinuity threshold.³⁰ Against these normalized running variables, we plot the mean residuals from regressions of quits on the baseline controls along with the fitted linear projections. Panel A plots the residuals against r_{iy} ; as expected, these results look similar to the plot from the univariate RD in Figure 3. The residuals in Panel B, plotted against r_{jy} , are from reduced-form models like equation (9) with the baseline set of controls (i.e., they are not scaled by the inverse of the first-stage coefficient). The mean residuals in Panel B are somewhat more scattered than in Panel A, reflecting lower precision in the peer wage estimates. However, the figure shows evidence of discontinuities—quit rates are generally higher to the right of the discontinuity than to the left. Except for some sparse cells like -0.06 for 6-month quits, the linear projections fit the data reasonably well.

Table 4 reports estimates of the own-wage and peer-wage quit elasticities, which are calculated using the baseline estimates from the univariate RD model (panel A) and the two-stage MRD model (panel B). For the MRD model, we also report the implied values of the preference parameter α , which quantifies the importance of relative pay in the quit decision; and we report the estimates of the “gap-constant” own-wage quit elasticity, which measures the quit response to market wage comparisons.

Both models imply large total quit elasticities with respect to the own wage. In Panel A the estimated elasticity is largest at 2 months (-24.8) and declines with longer quit windows, but is still quite large (-9.9) at 9 months. The estimates in Panels B show a similar pattern, and are if anything slightly larger in magnitude than those in Panel A. If these responses reflected only market comparisons, they would suggest a highly competitive labor market with negligible search frictions. Yet even before decomposing these estimates into market and peer-based components, we note that the estimates exhibit a pattern that is hard to reconcile with purely market-driven behavior. Specifically, they show that the quit response diminishes over time. Given that the number of outside offers can only increase over time, quit elasticities that are driven by outside offers should be larger in the longer-run (9 months) than in short-run (1-3 months).³¹

In Panel B, evidence from the MRD model suggests that the large quit response to own wages mainly reflects relative-pay concerns and not market comparisons. Here, the peer-wage elasticities are the opposite sign of the own-wage elasticities and are only slightly smaller in magnitude. The peer-wage elasticities (like the own-wage elasticities) are strongest in the first

³⁰As before, the normalized running variable for own wage is $r_{iy} = w_{0iy} - T_{iy}^k$, where $T_{iy}^k = \arg \min_{T_y^k} |T_y^k - w_{0iy}|$ is the wage threshold closest to the worker’s own initial wage, w_{0iy} . Similarly, the normalized running variable for peer wage is defined as $r_{jy} = w_{0jy} - T_{jy}^k$, where $T_{jy}^k = \arg \min_{T_y^k} |T_y^k - w_{0jy}|$ is the nearest wage threshold to the peer’s wage, w_{0jy} .

³¹As noted in section 4.1, the 6- and 9-months estimates may reflect a small attenuation bias due to subsequent raises that dilute the impact of the scheduled raise discontinuities over time. However, accounting for this would increase the 6-month elasticities by only 5% and the 9-month elasticities by about 20%; hence they would both remain substantially smaller than the elasticities at 1-3 months.

three months, and then become smaller at 6 and 9 months; this implies that workers are most responsive to the wages of their peers in the first three months after the raise. The estimates of α , calculated as the ratio of the magnitude of the peer effect to the own-wage effect (equation 3), are consistent with these patterns. They range from 0.67 to 0.95, and are largest at 1-3 months and smallest at 9 months.

The “gap-constant” quit elasticity, which is obtained by netting out the peer effect from the total own-wage effect (equation 6), measures how quit behavior responds when one’s wage increases relative to the market but *not* relative to one’s peers. Our estimates of the gap-constant elasticities (shown in the last rows of Panels B) are much smaller than the total own-wage elasticities, and range from -2.4 to -5.0 across the five quit windows. Though less precise, these estimates are consistent with the moderate elasticities found in some of the recent monopsony literature, and they imply that the firm has significant monopsony power due to search frictions. In particular, under the stationarity assumptions discussed by Manning (2003), the 9-month estimate from our baseline model implies a labor supply elasticity to the firm of around 4. This in turn suggests that labor market frictions allow the employer to reduce wages by roughly 25% below the value marginal product of labor.³²

5.4 Additional validation: Split-sample RD

The estimates from the peer effects models suggest that the sensitivity of quits to the discontinuities in own wages is driven mainly by relative-pay concerns. If this is true, then the quit discontinuities at the own-wage thresholds should be relatively small among workers whose peers receive similar raises. To test this implication, we can return to the univariate RD framework and estimate our baseline model separately for two sub-samples of workers. In the first sample, the majority of a worker i ’s peers are on the opposite side of a pay-step threshold;³³ as a result, the average peer raise is large (small) when the worker’s raise is small (large). The second sample is the complement of the first set. Here the worker and the majority of the peers are on the same side of the threshold and thus the average peer raise is similar to the worker’s raise.

³²If we convert the gap-constant quit response to a separation elasticity using the mean separation rate of .51 (instead of the quit rate of .36), we obtain an elasticity of -2.1. Following Manning (2003), the labor supply elasticity faced by the firm is then calculated as -2 times the separation elasticity. The separation elasticity is also useful for comparing our estimates to those of other studies that report this parameter instead of a quit response. Interestingly, our 9-month estimate is comparable both to the one-year separation elasticity of roughly -2 found by Ransom and Sims (2010) and to the one-year quit elasticity of -3.5 found by Falch (2011). These studies consider labor markets that are different from ours and they do not explicitly claim to rule out peer effects. Nevertheless, the wage increases they consider are shared by teachers in a school or district, and thus are common across what is arguably the relevant group of workplace peers. As such, their estimates are conceptually similar to our gap-constant elasticity.

³³Either worker i is above a threshold ($w_{0iy} \geq T_{iy}^k$) while the majority of peers are below ($\sum_{j \in p(i,y)} \mathbf{1}\{w_{0jy} < T_{jy}^k\} \geq \sum_{j \in p(i,y)} \mathbf{1}\{w_{0jy} \geq T_{jy}^k\}$), or the worker is below ($w_{0iy} < T_{iy}^k$) while the majority of peers are above ($\sum_{j \in p(i,y)} \mathbf{1}\{w_{0jy} \geq T_{jy}^k\} \geq \sum_{j \in p(i,y)} \mathbf{1}\{w_{0jy} < T_{jy}^k\}$).

Figure 5 shows, for each of the sub-samples, the linear projections for residuals of own wage and peer average wage, the wage gap, and quit rates. As in Figure 3 above, residuals are taken after partialing out the baseline set of controls and they are plotted against the normalized running variable $r_{iy} = w_{0iy} - T_{iy}^k$. In the first sample, where the worker and the majority of the peers are on opposite sides of the threshold, the own wage rises sharply at the threshold while the average peer wage falls—leading to a large jump in the *wage gap* at the threshold (Panel A). In this sample we find large negative discontinuities in the quit rates (Panel C). In the second sample, a majority of the peers are on the same side of the threshold as the worker, so the own wage and the average peer wage both rise at the threshold—leading to a smaller discontinuity in the wage gap (Panel B). In this sample the discontinuities in quit rates are also much more muted (Panel D). Importantly, while the discontinuity in the *wage gap* differs in the two samples, the discontinuity in the *own wage* does not. The differential quit response therefore supports our interpretation of the MRD results—that the effect of wages on quit behavior operates mainly through the effect on relative pay.

The split-sample RD estimates can also be used to construct an alternative estimate for the peer effect represented by β_p in our MRD model. If we let $\left[\frac{dw_g}{dw}\right]_S$ denote the effect of the own raise on the wage gap in a given sample S , then the overall effect of the wage increase on quits can be written as: $\left[\frac{dQ}{dw}\right]_S = \frac{\partial Q(w, w_g)}{\partial w} + \frac{\partial Q(w, w_g)}{\partial w_g} \left[\frac{dw_g}{dw}\right]_S$, where the second term is the effect of peer wage comparisons. Since $\frac{\partial Q(w, w_g)}{\partial w_g} = -\frac{\partial Q(w, w_p)}{\partial w_p}$, we can also write the total quit response as: $\left[\frac{dQ}{dw}\right]_S = \frac{\partial Q(w, w_g)}{\partial w} - \frac{\partial Q(w, w_p)}{\partial w_p} \left[\frac{dw_g}{dw}\right]_S$. Consider two samples S and S' where $\left[\frac{dw_g}{dw}\right]_S \neq \left[\frac{dw_g}{dw}\right]_{S'}$ (as is true of our “opposite-side” and “same-side” sub-samples). Under the additional assumption that both the gap-constant effect, $\frac{\partial Q(w, w_g)}{\partial w}$, and the peer-wage effect, $\frac{\partial Q(w, w_p)}{\partial w_p}$, are constant across the two samples, we can construct a Wald estimator for the peer effect as follows:

$$\frac{\partial Q(w, w_p)}{\partial w_p} = \frac{\left[\frac{dQ}{dw}\right]_{S'} - \left[\frac{dQ}{dw}\right]_S}{\left[\frac{dw_g}{dw}\right]_S - \left[\frac{dw_g}{dw}\right]_{S'}}.$$

Appendix Table 5 reports the RD estimates of $\frac{dw_g}{dw}$ and $\frac{dQ}{dw}$ for each sub-sample, and the Wald estimates of the peer effect. The estimates of $\frac{dw_g}{dw}$, shown in column 1, confirm that the discontinuity in the wage gap is much larger in the “opposite-side” sample (1.67) than in the “same-side sample” (0.54). The RD estimates for $\frac{dQ}{dw}$ are also consistent with the graphical evidence in Figure 5; they are sizable and statistically significant in the first sample, and are smaller and statistically insignificant in the second. The Wald estimates of the peer effect, reported in the last row, range from .39 (1 month) to .88 (3 months). These estimates are somewhat larger than the estimates of β_p from the baseline MRD model (reported in row 1

of Table 3) but are similar in magnitude to the estimates from the most fully specified model (Table 3, row 3).³⁴

We stress that the sources of variation used to estimate peer effects in the MRD and split-sample RD are distinct. In the MRD, identification comes solely from the response of quits to peer wages. In the split-sample RD, identification comes from comparing the quit response to *own-wage* increases across samples where the corresponding increases in the wage gap differed in size. The Wald estimates thus provide independent validation of the importance of relative-pay concerns.

5.5 Varying definitions of peer group

As discussed in section 5.2, a definition of peers that is too narrow will lead to an attenuation of the estimated peer effect, and too broad a definition will tend to make the estimate imprecise by reducing identifying variation. Since we do not know the appropriate definition *a priori*, our choice of a +/- \$.20 wage band to define peers is informed by an analysis that allows the definition to vary. We consider groups of coworkers initially earning within \$.05, \$.10, \$.15, \$.20, \$.25, or \$.30 of the worker—groups that comprise between 26 and 73 percent of all coworkers who qualify for a scheduled raise; we also consider a specification that includes all such coworkers in the store. For each peer group definition, we estimate the two-stage MRD models of quits with the baseline set of controls.

Figure 6 shows the estimated peer effects and the 95 percent confidence intervals for peer groups based on successively wider wage bands. Consistent with attenuation bias from the exclusion of relevant peers, estimates based on narrow wage bands are relatively small. The \$.05 wage bands produce no evidence of peer effects, and the estimates based on \$.10 and \$.15 bands, while positive, are statistically insignificant. The estimates generally increase with the size of the wage band, especially as it expands from +/- \$.05 to +/- \$.20. The point estimates are relatively stable in the \$.20-\$.30 range and are statistically significant for the 2-3 month quit windows; however, the standard errors tend to increase with the size of the band—leading us to prefer the \$.20 band. Finally, while the estimates based on a store-wide definition of peers are the least precise, they are reasonably consistent with those from our preferred \$.20 definition. Overall, the patterns in Figure 6 support our use of a +/- \$.20 wage band to define peers, and also show that our findings are robust to moderate variation in the peer definition.

³⁴Because the validity of the Wald estimate requires a constant effect of peer wage in the two subsamples, nonlinearity in the peer effect may result in an upward bias. In section 5.6 we present evidence that the peer effect is indeed nonlinear; hence, such a bias may account for the difference between the Wald estimates and the baseline MRD estimates of β_p . Specifically, we find that the peer effect is large when workers are paid distinctly less than their peers, but it is not significant when workers earn close to the average peer wage. This nonlinearity suggests that in the same-side sample (where the wage gap is small on both sides of the threshold), the true peer effect may be zero and the quit discontinuity may reflect only the response to market comparisons. If we assume that the peer effect is zero in this sample and re-calculate the Wald estimate for the opposite-side sample, we obtain estimates that range from .26 to .60 and are very similar to the baseline MRD estimates.

5.6 Nonlinear peer effects

So far, our empirical models have assumed that the quit response to a change in relative pay is linear in the wage gap. We now ask whether this response differs depending on the worker’s location in the peer wage distribution. Theory suggests that the shape of the response function may depend on workers’ underlying motives. For example, if workers are especially concerned about being treated fairly, then the response may be nonlinear—and stronger when workers are paid less than their peers (Fehr and Schmidt 1999, Card et al. 2012).

To test for such patterns in our data, we divide our sample of worker-peer pairs into four groups based on plausibly exogenous variation in the wage gap ($w_{iy} - \bar{w}_{p(i,y)y}$), and we compare the causal effect of peer wages on quit rates across these four groups. To construct the four groups, we first split workers into two groups based on whether a worker’s initial wage is below or above the nearest pay-step threshold (i.e., $w_{0iy} < T_y^k$ or $w_{0iy} \geq T_y^k$).³⁵ This will allow us to compare the peer-wage response of workers who themselves received arbitrarily small raises to the response of those who received large raises. Next, we divide each of the worker’s *peers* into two sub-groups based on whether the peer’s initial wage is above or below the initial wage of the worker.³⁶ Thus, for each of the two worker types (i.e., those with small and large own-wage increases), we estimate separate quit responses to the wage increases of higher-paid and lower-paid peers. In each case, we use the baseline specification of the two-stage MRD model to estimate the peer-wage effect β_P , and the treatment variable in the second stage ($\bar{w}_{p(i,y)y}$) is the average wage of the worker’s peers in that group.³⁷

In Figure 7, we first show a box plot of the *ex post* wage gap over the four groups of worker-peer pairs (Panel A). The leftmost group is composed of pairs in which the worker initially earned less than the peer and the worker herself got a small raise. After the raise, most workers still earned less than their peers in this group—with a median gap of $-\$10$. The rightmost group is composed of pairs in which the worker initially earned more than the peer and the worker herself got a large raise. In this group, workers earned more than their peers after the raise—with a median gap of $\$10$. The middle two groups consist of workers who earned less than their peers and got a large raise, and workers who earned more than their peers and got a small raise. In these groups, the wage gaps are mostly zero or very small; the median gaps are $-\$0.02$ and $\$0$, and roughly half the workers in each group end up earning the same wage as their peers.

For each of the four groups, Panel B of Figure 7 shows the estimates of β_P and the 95% confidence intervals. The estimates are sizable and positive only for the leftmost group – where workers were initially earning less than the peers and the workers themselves got arbitrarily

³⁵Recall that $T_y^k = \arg \min_{T_y^k} |T_y^k - w_{0iy}|$ is the threshold nearest to w_{0iy} .

³⁶Peers who are initially paid the same as the worker are grouped with those who were paid less; hence the groups are defined by $w_{0iy} + 0.20 \geq w_{0jy} > w_{0iy}$ and $w_{0iy} - 0.20 \leq w_{0jy} \leq w_{0iy}$.

³⁷Since each sub-sample contains either workers to the left of a threshold or those to the right, but not both, the own-wage effect is not identified.

small raises. For this group, the 2- and 3-month estimates are statistically significant at the 5 percent level and the 6-month estimate is significant at the 10 percent level. In contrast, the point estimates for the other three groups are small and never statistically significant.

Our results indicate that a worker’s quit behavior responds strongly to differences in peer raises only when the worker ends up earning less than the average wage of her peers. In contrast, when workers end up with wages that are either similar to or higher than those of their peers, they do not appear sensitive to the raises of their peers. As long as a worker’s location above or below the nearest pay-step threshold is random, this nonlinearity in the peer effect is unlikely to be driven by differences other than one’s location in the peer wage distribution. Hence these findings are consistent with a model of fairness as an aversion to disadvantageous inequity, in which workers who earn less than their peers care most about relative pay.

6 Falsification Tests

This section presents two falsification tests that lend further support to the causal interpretation of our estimates. First, we test for discontinuities in quit rates at the actual pay-step thresholds for wages and peer wages, but at a date prior to the raises. Second, we test for discontinuities at the actual date, but at wages other than the true pay-step thresholds.

6.1 Tests for discontinuities prior to implementation of raises

If our estimates were biased due to a correlation between threshold wages and the latent propensity to quit, then this correlation should cause discontinuities in quit rates even before the raises were implemented. We test for such “pre-treatment” effects in the three months before the implementation date of Oct. 1, 1996. We use a sample that consists of all employees who received a merit raise during the week that annual merit raises were given, June 30 - July 7, 1996; who were still employed on July 8, 1996; and whose wages on July 8 meet the range restrictions for the main estimation sample. We estimate models of quit rates within 1, 2, and 3 months of July 8 as functions of the scheduled wage that took effect on October 1. Using the baseline specifications, we estimate both the univariate RD and the MRD models.

Table 5 shows the results. In Panel A, the univariate RD estimates suggest a small positive but insignificant association between the future scheduled raise and quit rates within 1-2 months after the merit raises; at three months, the coefficient is negative but close to zero. Panel B shows the MRD results. The own-wage coefficients are similar to those in Panel A. The peer-wage estimates are small and positive at one month but are negative at 2 and 3 months, and in no case are they significantly different from zero. In sum, there is generally no evidence of a pre-treatment effect of own wages or peer wages on quits.

6.2 Tests for discontinuities at non-threshold wages

Our second test looks for discontinuities in quit rates at wage values where there is no associated discontinuity in the scheduled wage. We are especially interested in values \$.05 above and below each threshold. As discussed in section 4.2, bunching at these values raises the concern that workers whose own wages are multiples of \$.05 have different latent quit propensities than those who do not. This would result in non-random sorting and spurious quit discontinuities at the own-wage thresholds (which occur only at \$.05 multiples). A similar argument also applies to peer wages. If the fraction of peers at \$.05 multiples is correlated with store or market characteristics that predict quits, this could lead to spurious quit discontinuities at the peer-wage thresholds. Either type of confound, however, would lead to discontinuities in quit rates not only at the pay-step thresholds but also at wage values \$.05 above and below each threshold. We can therefore assess these potential threats by testing for discontinuities at values other than the true thresholds.

For the univariate RD, we implement this test by replacing the initial wage w_{0iy} (which enters equation (8) both directly and indirectly through the scheduled wage w_{iy}) with a “shifted” wage $w_{0iy}^s = w_{0iy} + s$ for $s \in \{-.07, -.06, \dots, +.07\}$. For the MRD, we replace w_{0iy} with w_{0iy}^s in equation (10) and similarly we replace each peer wage w_{0jy} with a shifted peer wage w_{0jy}^s .

In Appendix Figure 4, Panel A plots the t-statistics for the estimated quit discontinuities from the univariate RD for each value of s . We see evidence of a significant negative discontinuity at the true threshold (indicated by $s = 0$), and also when wages are shifted by +\$.01.³⁸ But importantly, we do not see discontinuities elsewhere in the wage distribution, including at $w_{0y} + $.05 or $w_{0y} - $.05. Panel B shows t-statistics for the peer-wage estimates from the wage-shifted versions of the two-stage MRD model. The pattern here is similar to Panel A in that the estimates at $s = 0$ and $s = 1$ stand out as the most significant. Again, the estimates at $.05 above and below the true thresholds are statistically insignificant (and are opposite in sign from the true discontinuity estimates). In sum, this test helps confirm that our findings are not driven by unobserved heterogeneity in the quit propensity associated with either own wages or peer wages at $.05 multiples.$$

7 Discussion

We use a unique source of variation in wages to identify the effects of both market competition and peer comparisons on quit behavior. In our context, we find that quits are highly responsive to wage increases and that this behavior is driven largely by relative-pay concerns. After accounting for peer effects, quits do not appear to be very sensitive to wages—consistent with

³⁸The discontinuity at $s=1$ occurs because bins just above the true threshold are more heavily populated than bins just below it; as a result, shifting wages by +\$.01 has little impact on the predicted value just above the threshold.

the presence of significant search frictions. Finally, we find that the relative-pay effects are nonlinear and are driven mainly by workers who are paid less than their peers—suggesting that the quit behavior is motivated by concerns about fairness.

Is it possible that the response to peer wages reflects some type of rational learning? We think this is unlikely. In the standard learning model, workers use peer wages to help predict their own future pay within the firm. The model predicts that an exogenous increase in peer wages causes workers to revise their expected future wage upwards and thus makes them less likely to quit. Since the direction of the peer effects we find is the opposite of that predicted by the model, this type of rational updating clearly cannot explain our results. In principle, such learning could have a countervailing effect on our estimates. But in our context, this seems doubtful because the peer raises are based on arbitrary cutoffs and should not predict future raises at all. We confirm this lack of predictive power in Appendix Figure 5. We examine future wage growth, which is defined as the change in own wage from the day after the raise through the next merit raise cycle 10 months later. The figure plots future wage growth against the normalized peer-wage running variable and shows there is no discontinuity at the thresholds. It thus confirms that peer wages cannot provide a basis for rational learning about future wages within the firm.

Could peer wages instead be providing workers with a signal about market wages? If higher peer wages signaled that there were better paid jobs available in the market, then workers might increase their search intensity and be more likely to leave the job. However, the true wage offer distribution is unlikely to change at the firm’s arbitrary pay-step thresholds. In our context, therefore, peer wages cannot provide a basis for rational learning about the market.

Overall, the evidence suggests our results are best explained by a model of fairness in which workers who are paid less than their peers care most about relative pay. If so, we should ask what else might characterize the fairness concerns in our setting. One point which merits consideration is that the wage differentials that elicit the quit response are relatively small (roughly 2% of the typical wage). Why would quits respond to such modest differences in relative pay? We suggest two possibilities, which could be mutually reinforcing. First, it may be that workers care about pay inequality *per se* (even if the pay differences are small). Card et al. (2012) offer some support for this hypothesis by showing that job satisfaction is more closely related to pay rank than to actual pay differentials. Second, recent laboratory evidence suggests that pay differences elicit stronger responses when they are arbitrary than when they are justified (Bracha et al. 2015). Hence the arbitrary nature of the pay differences at our firm may have made workers more likely to respond to moderate inequities.

Finally, our findings have implications for the wage-setting behavior of the firm. On one hand, the small estimates of the gap-constant quit elasticity suggest that the firm has significant monopsony power. Our study thus supports a growing body of evidence that search frictions may be significant even in low-wage labor markets. On the other hand, the large response of

quits to relative pay suggests that peer comparisons can create internal constraints on wage-setting. In particular, the concentration of this response among workers with low relative pay suggests that firms have an incentive to compress wages in order to reduce turnover. The results also indicate that even modest pay differences may have big effects if they appear unfair.

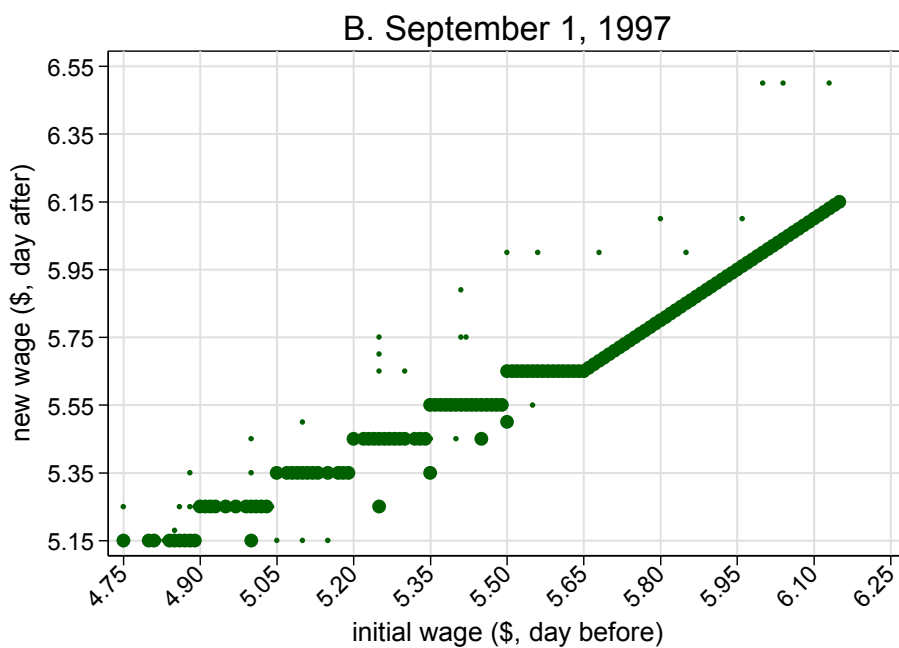
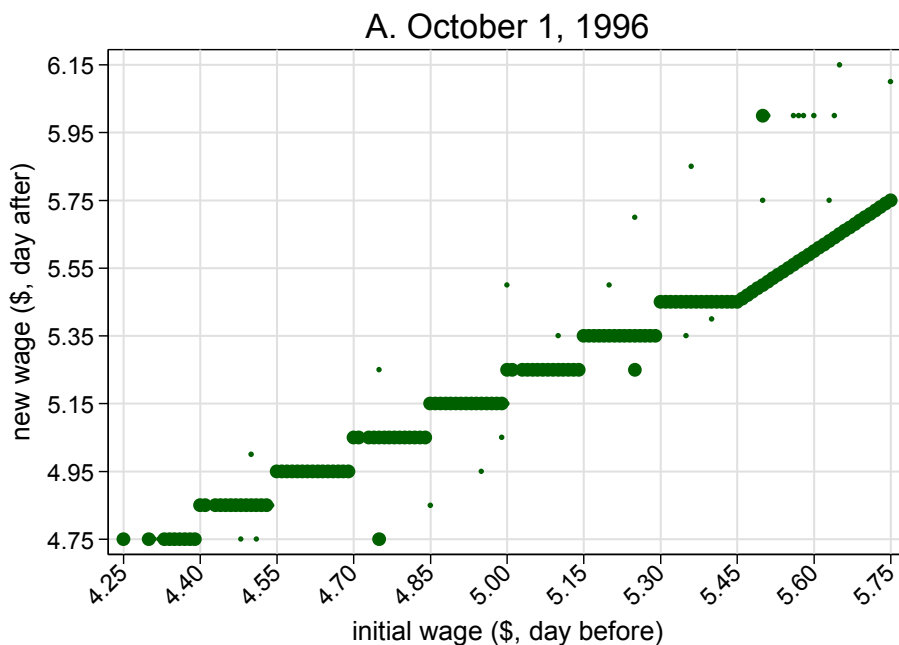
References

- [1] Akerlof, George A., and Janet L. Yellen. 1990. "The fair wage-effort hypothesis and unemployment." *Quarterly Journal of Economics* 105(2): 255-283.
- [2] Angrist, Joshua D., and Victor Lavy. 1999. "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement." *The Quarterly Journal of Economics* 114(2): 533-575.
- [3] Angrist, Joshua D., Erich Battistin, and Daniela Vuri. 2014. "In a Small Moment: Class Size and Moral Hazard in the Mezzogiorno." NBER Working Paper No. w20173.
- [4] Autor, David, Alan Manning, and Christopher L. Smith. 2015. "The Contribution of the Minimum Wage to U.S. Wage Inequality over Three Decades: A Reassessment." MIT Working Paper.
- [5] Barreca, Alan I., Jason M. Lindo, and Glen R. Waddell. 2011. "Heaping-Induced Bias in Regression-Discontinuity Designs." NBER Working Paper No. w17408.
- [6] Bracha, A., Gneezy, U. and Loewenstein, G. 2015. "Relative Pay and Labor Supply." *Journal of Labor Economics* 33(2): 297-315.
- [7] Buntrock, Evan. 2014. "Comparison-Group Position and Job Search: Information or Irritation?" Mimeo.
- [8] Burdett, Kenneth and Dale Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39: 257-273.
- [9] Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review* 102 (6): 2981-3003.
- [10] Charness, Gary, and Peter J. Kuhn. 2007. "Does Pay Inequality Affect Worker Effort? Experimental Evidence." *Journal of Labor Economics* 25: 693-723.
- [11] Clark, Andrew, Nicolai Kristensen and Niels Westergaard-Nielsen. 2009. "Job Satisfaction and Co-Worker Wages: Status or Signal?" *The Economic Journal* 119(536): 430-447.
- [12] Clark, Andrew and Andrew J. Oswald. 1996. "Satisfaction and Comparison Income." *Journal of Public Economics* 61(3): 359-381.
- [13] Cohn, Alain, Ernst Fehr, Benedikt Herrmann, and Frédéric Schneider. 2014. "Social Comparison and Effort Provision: Evidence from a Field Experiment." *Journal of the European Economic Association*. Forthcoming.

- [14] Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2013. "Peer Effects in the Workplace." IZA Discussion Paper No. 7617.
- [15] Depew, Briggs, and Todd A. Sørensen. 2013. "The Elasticity of Labor Supply to the Firm over the Business Cycle." *Labour Economics* 24: 196-204.
- [16] Dube, Arindrajit, T. William Lester, and Michael Reich. 2014. "Minimum Wage Shocks, Employment Flows and Labor Market Frictions." *Journal of Labor Economics*. Forthcoming.
- [17] Falch, Torberg. 2011. "Teacher Mobility Responses to Wage Changes: Evidence from a Quasi-Natural Experiment." *The American Economic Review* 101(3): 460-465.
- [18] Feigenbaum, James. 2014. "From the Bargaining Table to Ballot Box: The Effects of Labor Unions on Elections." Mimeo.
- [19] Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics* 114(3):817-68.
- [20] Galizzi, Monica and Kevin Lang. 1998. "Relative Wages, Wage Growth, and Quit Behaviour." *Journal of Labor Economics* 16 (2): 367-391.
- [21] Giuliano, Laura. 2013. "Minimum Wage Effects on Employment, Substitution, and the Teenage Labor Supply: Evidence from Personnel Data." *Journal of Labor Economics* 31(1): 155-94.
- [22] Goerg, Sebastian J., Sebastian Kube, and Ro'i Zultan. 2010. "Treating Equals Unequally: Incentives in Teams, Workers' Motivation, and Production Technology." *Journal of Labor Economics* 28(4): 747-772.
- [23] 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* 28(2): 291-330.
- [24] Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48(2): 281-355.
- [25] Luttmer, Erzo F. P. 2005. "Neighbors As Negatives: Relative Earnings And Well-Being." *Quarterly Journal of Economics* 120(3): 963-1002.
- [26] McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142(2): 698-714.
- [27] Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton: Princeton University Press.


- [28] Mas, Alexandre. 2014. “Does Transparency Lead to Pay Compression?” NBER Working Paper No. w20558.
- [29] Papay, John P., John B. Willett, and Richard J. Murnane. 2011. “Extending the Regression-Discontinuity Approach to Multiple Assignment Variables.” *Journal of Econometrics* 161(2):203-207.
- [30] Pfeifer, Christian and Stefan Schneck (2012): Relative Wage Positions and Quit Behavior: Evidence from Linked Employer-Employee Data. *Industrial and Labor Relations Review* 65(1): 126-147.
- [31] Reardon, Sean F. and Joseph P. Robinson. 2012. “Regression Discontinuity Designs With Multiple Rating-Score Variables.” *Journal of Research on Educational Effectiveness* 5(1):83-104.
- [32] Ransom, Michael R., and David P. Sims. 2010. “Estimating the Firm’s Labor Supply Curve in a ‘Mew Monopsony’ Framework: Schoolteachers in Missouri.” *Journal of Labor Economics* 28(2): 331-355.
- [33] Ransom, Michael R., and Ronald L. Oaxaca. 2010. “New Market Power Models and Sex Differences in Pay.” *Journal of Labor Economics* 28(2): 267-289.
- [34] Rege, Mari, and Ingeborg F. Solli. 2014. “Lagging Behind the Joneses: The Impact of Relative Earnings on Job Separation.” Mimeo.
- [35] Smith, Adam. 1759. *The Theory of Moral Sentiments*. (London: A. Millar; Edinburgh: A. Kincaid and J. Bell).
- [36] Webber, Douglas A. 2013. “Firm Market Power and the Earnings Distribution.” IZA Discussion Paper No. 7342.

Figure 1. Wages on days before and after each minimum wage increase




Note: Sample includes all hourly employees who were present on the day of the federal minimum wage increase on October 1, 1996 (Panel A) or September 1, 1997 (Panel B), who had at least one month of tenure, and whose “before” wage was no more than \$1.00 above the new minimum. Sample size is 15,557 (21,274) observations for Panel A (Panel B). Small dots represent cells with fewer than five observations.

Figure 2. Raise by initial wage on day of each minimum wage increase, average actual raise and scheduled raise




Note: Circles are mean values of raise received on day of minimum wage increase for each value of the “before” wage. Lines are fitted values from regressions of actual raise on the scheduled raise defined by the corporate rule for wage adjustments. Regressions are fit separately for employees in the analysis sample on each of the two days that the minimum wage increased.

Figure 3. Quit rates in months following raises for a representative interval of initial wage




Note: The figure shows residuals from the RD model of quits with baseline controls (as in Table 2, row 3) for a representative interval of the running variable (initial own wage). For visual simplification, the running variable is normalized as the distance to the nearest pay-step threshold (see text for details). The lines show the fitted relationship between residualized quit rates and the normalized running variable. For each value of the normalized running variable, the data points are constructed by adding back to fitted values the mean of the residuals taken across all 12 intervals. Marker size is scaled by the number of observations at each value. For all series, the intercepts are normalized to be zero at the left limit of the threshold, so the value at the right limit is the estimated effect of the \$.10 discontinuity in the wage. Estimation samples are as in Table 2.

Figure 4. Quit rates in months following raises for a representative interval of initial own wage or peer wage




Note: The figure shows residuals from the reduced-form MRD model of quits with baseline controls (as in row 1 of Table 3 and Appendix Table 3) for representative intervals of the two running variables—initial own wage (Panel A) and initial peer wage (Panel B). The running variables are normalized as the distance to the nearest pay-step threshold (see text for details). The lines show the fitted relationship between residualized quit rates and the normalized running variable. The points are constructed by adding back to fitted values the mean residuals for each value of the normalized running variable; marker size is scaled by the number of observations at each value. For all series, the intercepts are normalized to be zero at the left limit of the threshold, so the value at the right limit is the estimated effect of the \$.10 discontinuity in the own wage (Panel A) or the peer wage (Panel B). Estimation samples are as in Table 3.

Figure 5. Discontinuities in wages and quit rates for a representative interval, by whether employee is on opposite or same side of pay threshold as majority of peers



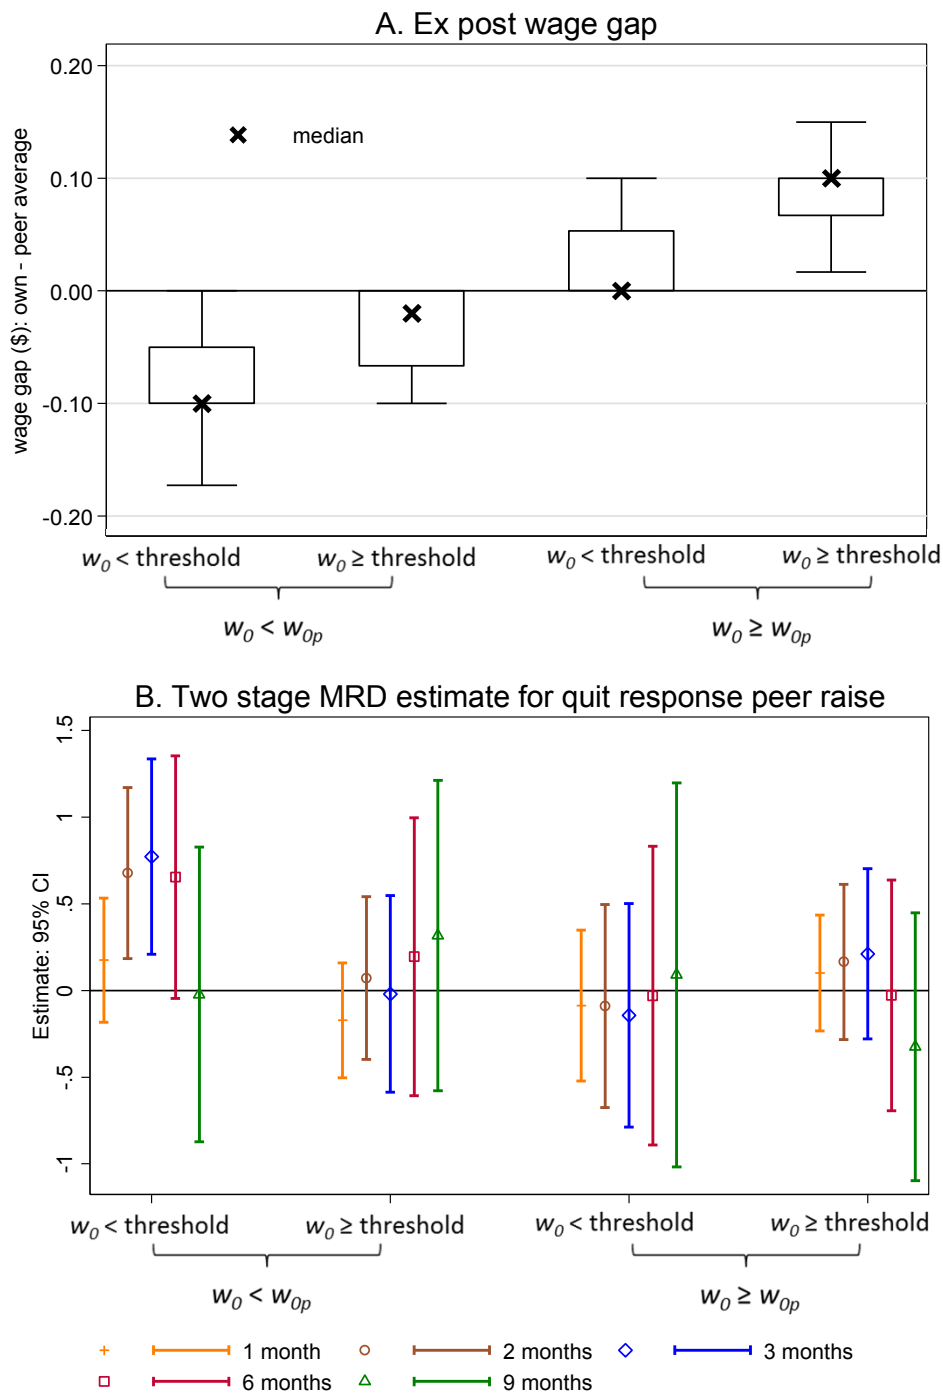
Note: The figures show fitted linear relationships between residuals from baseline RD models (as in Table 2, row 3) and a normalized running variable, for a representative interval of the running variable (own initial wage). Panels A and B show models for own wages, peer wages, and the own-peer wage gap. Panels C and D show models for quits. Panels A and C use a sample in which a majority or peers are on the opposite sides of a pay-step threshold from the worker. Panels B and D use a sample where the workers and peers are on the same side of a threshold. See text for more details.

Figure 6. Estimated effects of peer wage on quit rate, varying size of peer group





Note: Figure plots estimates and 95% confidence intervals from two-stage MRD models of the quit response to an increase in the average peer wage for peer groups defined by various wage bands. Models include baseline set of controls as in Table 3, model 1.

Figure 7. Nonlinear quit response to an increase in peer wage




Note: Four groups are constructed by splitting the full sample of worker-peer pairs, first by whether a worker's initial wage (w_0) was above or below the nearest pay-step threshold, and then by whether the peer's initial wage (w_{0p}) was above or below the worker's initial wage. For each of these groups, Panel A shows a box plot of the *ex post* wage gap (i.e. the difference between the worker's wage and the average peer wage after the raise). Panel B plots two-stage MRD estimates and 95% confidence intervals for a worker's quit response to an increase in the average wage of the relevant peer group.

Appendix figure 1. Distribution of wages on day before each minimum wage increase, estimation sample




Note: Frequency distribution of initial wages for the estimation sample, including 3,009 workers employed on October 1, 1996 and 3,682 workers employed on September 1, 1997.

Appendix figure 2. Histogram of distance from nearest discontinuity threshold




Note: Figure plots frequency distribution of the distance of initial wages in our sample from the nearest pay-step threshold. Based on estimation sample of 6,691 initial wages.

Appendix figure 3. Size of most recent merit raise (%), by distance from nearest discontinuity threshold




Note: Figure plots the average merit raise received during the merit raise cycle in the previous June by the distance of the initial wage from the nearest pay-step threshold. Based on estimation sample of 6,691 initial wages.

Appendix figure 4. Placebo test for RD and MRD models of quit behavior; placebo raises constructed using initial wage +/- \$.07



Note: T-statistics for estimated discontinuities from RD model of quits on own wage (Panel A) and MRD model of quits on peer wage (Panel B). Each model is estimated on a data set in which the distribution of wages (or peer wages) is “shifted” by the amount indicated on the x-axis, so that estimates at all values other than 0 are estimated discontinuities in quits at wages for which there is no discontinuity in the scheduled raise. See text for details.

Appendix figure 5. Estimated effect of peer raise on future wage growth



Note: Dependent variable is the change in own wage from the day after the minimum wage increase to 10 months later (after the merit raise cycle in the following year). The figure shows residuals from a reduced-form MRD model of this variable, with baseline controls (as in row 1 of Table 3 and Appendix Table 3) for a representative interval of the running variable (initial peer wage). As in Figure 4B, the running variable is normalized as the distance to the nearest pay-step threshold. The lines show the fitted relationship between residualized wage growth and the normalized running variable. The points are constructed by adding back to fitted values the mean residuals for each value of the normalized running variable. The estimated discontinuity is $-\$0.007$ (standard error 0.021). The sample is all workers in the main estimation sample who remained employed throughout the 10-month period over which wage growth is measured. Sample size is 3,574.

Table 1. Summary statistics for estimation sample

	mean	std. dev.
A. <u>Wages and raises by date of minimum wage increase:</u>		
<u>October 1, 1996:</u>		
initial wage	\$4.99	(\$0.29)
scheduled raise	\$0.21	(\$0.10)
<i>number of observations</i>	3,009	
<u>September 1, 1997:</u>		
initial wage	\$5.28	(\$0.21)
scheduled raise	\$0.18	(\$0.07)
<i>number of observations</i>	3,682	
<u>Pooled sample</u>		
initial wage	\$5.15	(\$0.29)
scheduled raise	\$0.19	(\$0.08)
<i>number of observations</i>	6,691	
B. <u>Employee characteristics:</u>		
initial wage is multiple of \$.05 (%)	39%	
fulltime status (%)	0.4%	
female (%)	81%	
white (%)	76%	
black (%)	10%	
hispanic (%)	7%	
mean of ZIP code median household income (\$1000s)	51.3	
(std. dev.)	(18.6)	
age category:		
16-17 years old (%)	24%	
18-19 years old (%)	28%	
20-22 years old (%)	21%	
23-29 years old (%)	14%	
30 years and older (%)	12%	
mean age (years)	22.6	
(std. dev.)	(7.8)	
tenure as of October 1, 1996:		
median tenure (months)	≥ 8	
tenure is censored at 8 months (%)	86%	
tenure as of September 1, 1997:		
median tenure (months)	11.7	
tenure is censored at 19 months (%)	16%	

(continued)

Table 1. cont'd.

C. Coworker (store-level workforce) characteristics:		
	mean	std. dev.
mean number of coworkers	27.3	14.90
average wage	\$5.55	\$0.46
average age	24.2	2.70
% who got scheduled raise	63%	23%
% who got a merit raise in past year	44%	15%
% who got both a scheduled raise and a merit raise	24%	14%
% with initial wages at multiple of \$.05	70%	16%
D. Peer characteristics:		
	mean	std. dev.
number of peers	7.2	5.6
average initial wage	\$5.15	\$0.28
average scheduled raise	\$0.20	\$0.08
average age	21.9	4.6
% who got a merit raise in past year	47%	32%
% with initial wages at multiple of \$.05	69%	30%
E. Quit rates by time window from minimum wage increase		
	%	# obs. (non-censored)
% quit within 1 month	6%	6,259
% quit within 2 months	11%	6,086
% quit within 3 months	14%	5,989
% quit within 6 months	25%	5,507
% quit within 9 months	36%	5,154

Note: Panels A-C based on the full estimation sample of 6,691 scheduled raises given on the day of one of the minimum wage increases (either September 1, 1996 or October 1, 1997) among employees who had previously received a merit raise (see Section 3.2 of text for explanation of sample restrictions). Coworkers are defined as those in entry-level jobs who are employed in the same store on the day of the minimum wage increase. Peers are defined as coworkers who got a scheduled raise and whose initial wage is +/- \$.20 of own initial wage; peer characteristics in Panel D are shown for the sub-sample of 6,385 employees (95% of full estimation sample) who have at least one peer by this definition. Quit rates in Panel E are shown for sub-samples of employees who do not leave the store for another reason (returning to school, transferring to another store, moving, or being fired) within the indicated time window.

Table 2. Regression discontinuity estimates of the effect of own wage on quit rates

Model specification:	<i>Dependent variable = 1 if quit within:</i>				
	(1) 1 month	(2) 2 months	(3) 3 months	(4) 6 months	(5) 9 months
1 control only for year-specific linear trends in initial wage	-0.14 (0.10)	-0.31* (0.15)	-0.29† (0.16)	-0.37† (0.21)	-0.44† (0.25)
2 add ZIP code fixed effects	-0.24* (0.11)	-0.50** (0.15)	-0.52** (0.17)	-0.58* (0.23)	-0.70** (0.27)
3 add dummy for initial wage = multiple of \$.05	-0.24* (0.11)	-0.51** (0.16)	-0.55** (0.18)	-0.62** (0.23)	-0.70* (0.28)
4 add quadratic in wage	-0.24* (0.11)	-0.51** (0.16)	-0.55** (0.18)	-0.62** (0.23)	-0.70* (0.28)
5 add employee and coworker controls	-0.24* (0.11)	-0.52** (0.16)	-0.57** (0.18)	-0.66** (0.23)	-0.73** (0.28)
6 add store fixed effects	-0.28* (0.13)	-0.58** (0.18)	-0.65** (0.20)	-0.75** (0.27)	-0.96** (0.33)
<i>number of obs. (rows 1-6)</i>	<i>6259</i>	<i>6086</i>	<i>5989</i>	<i>5507</i>	<i>5154</i>
7 "donut hole:" exclude \$.05 wages	-0.20 (0.14)	-0.57** (0.20)	-0.56* (0.23)	-0.56† (0.32)	-0.77* (0.36)
<i>number of obs. (row 7)</i>	<i>3792</i>	<i>3692</i>	<i>3632</i>	<i>3333</i>	<i>3123</i>

Note: Row 1 reports regression coefficients on scheduled (new) wage from linear probability models of quits, controlling only for a linear function of the initial wage with a separate intercept and slope in each year. Rows 2-3 add the "baseline" controls comprised of 3-digit store ZIP code fixed effects as well as a dummy for the initial wage being a multiple of \$.05. Row 4 replaces the linear function of initial wage with a quadratic function. Rows 5-6 add controls for employee and coworkers characteristics. Employee controls are: a dummy for each month of tenure, age and age-squared, gender and race dummies, an indicator for full-time status, size of the most recent merit raise, and the median household income in the employee's residential ZIP code. Coworker controls are: total number of entry-level employees on the day of the minimum wage increase, average employee age, average employee wage, the fraction who received a scheduled raise, the fraction who received a merit raise in July of the same year, and the fraction whose initial wage is a multiple of \$.05. The estimation sample in rows 1-6 includes all employees who received a scheduled raise on the day of the minimum wage increase, who had received a merit raise in July of the sample year, and who did not terminate employment for a reason other than quitting within the specified interval (i.e. non-quit separations are treated as censored). The model in row 7 is similar to that in row 3, but is estimated using a sample that excludes observations for which the initial wage is a multiple of \$.05. Parentheses contain robust standard errors clustered by store. † significant at 10%; * significant at 5%; ** significant at 1%.

Table 3. Two-stage MRD estimates of quit response to changes in own wage and peer average wage

Model specification:	Dependent variable = 1 if quit within:				
	(1) 1 month	(2) 2 months	(3) 3 months	(4) 6 months	(5) 9 months
1 baseline controls (see note)					
own wage	-0.25* (0.11)	-0.55** (0.16)	-0.61** (0.17)	-0.71** (0.23)	-0.75** (0.27)
peer average wage	0.22 (0.14)	0.45* (0.20)	0.51* (0.22)	0.55† (0.29)	0.54 (0.35)
2 control for employee characteristics and store average coworker characteristics					
own wage	-0.25* (0.11)	-0.55** (0.15)	-0.63** (0.17)	-0.75** (0.23)	-0.77** (0.27)
peer average wage	0.22 (0.14)	0.44* (0.20)	0.51* (0.22)	0.57† (0.30)	0.56 (0.36)
3 additional controls for peer characteristics; plus quadratic in peer initial wage					
own wage	-0.23* (0.11)	-0.53** (0.15)	-0.62** (0.17)	-0.73** (0.23)	-0.75** (0.27)
peer average wage	0.36* (0.16)	0.63** (0.22)	0.66** (0.25)	0.77* (0.33)	0.77† (0.40)
<i>number of observations</i>	42,549	41,180	40,470	36,961	34,695

Note: Peers are defined as coworkers who got a scheduled raise and whose initial wage is +/- \$.20 of own initial wage. Estimation sample consists of all worker-peer pairs for which the worker is in the RD analysis sample (see Table 2 note) and the peer fits our definition. Baseline controls in all models include linear functions of own initial wage and peer initial wage; dummies for initial wages that are \$.05 multiples, initial peer wages that are \$.05 multiples, and peers that got merit raises; and ZIP code fixed effects. In model 2, employee characteristics are: a dummy for each month of tenure, age and age-squared, gender and race dummies, an indicator for full-time status, size of the most recent merit raise, and the median household income in the employee's residential ZIP code. Store-level coworker characteristics are: total number of entry-level employees on the day of the minimum wage increase, average employee age, average employee wage, the fraction who received a scheduled raise, the fraction who received a merit raise in July of the same year, and the fraction whose initial wage is a multiple of \$.05. In model 3, peer characteristics are the same variables included for the employee in model 2. Parentheses contain robust standard errors clustered on store. † significant at 10%; * significant at 5%; ** significant at 1%.

Table 4. Implied quit elasticities and importance of relative wage concerns

	<i>Dependent variable = 1 if quit within:</i>				
	(1)	(2)	(3)	(4)	(5)
	1 month	2 months	3 months	6 months	9 months
A. RD model for total effect					
own-wage elasticity	-22.5 (10.3)	-24.8 (7.8)	-21.0 (6.9)	-12.7 (4.7)	-9.9 (4.0)
B. Two-stage MRD model for own and peer wage effects					
own-wage elasticity	-23.4 (10.3)	-26.7 (7.8)	-23.3 (6.5)	-14.6 (4.7)	-10.6 (3.8)
peer-wage elasticity	20.6 (13.1)	21.9 (9.7)	19.5 (8.4)	11.3 (6.0)	7.6 (5.0)
importance of relative wage concern (α)	0.88	0.82	0.84	0.77	0.72
gap-constant own-wage elasticity (net of peer effect)	-2.4 (17.2)	-5.0 (12.2)	-3.8 (11.0)	-3.4 (7.7)	-2.9 (6.5)
sample mean of dependent variable	0.06	0.11	0.14	0.25	0.36

Note: Elasticities are calculated using the coefficients on own wage (estimates of β) and peer average wage (estimates of β_p) from baseline model specifications as in Table 2, row 3 (for Panel A) and Table 3, row 1 (for Panel B); the sample means shown in last row of this table; and the sample average wage of \$5.15. Importance of relative wage concern is calculated as $\alpha = -\beta_p/\beta$ (see equation 2) using the estimates reported here. The gap-constant own-wage elasticity is calculated by first summing the estimates of β and β_p (see equation 6), and then dividing by the sample mean quit rate and multiplying by \$5.15. Parentheses contain robust standard errors clustered by store.

Table 5. Estimated effects of October 1, 1996 scheduled wages on quits in months before raises were implemented

	<i>Dependent variable = 1 if quit within:</i>		
	(1)	(2)	(3)
from July 8:	+ 1 mo.	+ 2 mo.	+ 3 mo.
from Oct. 1:	- 2 mo.	- 1 mo.	0 mo.
A. RD model for total effect			
own wage	0.17 (0.13)	0.12 (0.20)	-0.07 (0.24)
<i>number of observations</i>	5,043	4,307	4,056
B. Two-stage MRD model for own and peer wage effects			
own wage	0.12 (0.12)	0.02 (0.19)	-0.18 (0.23)
peer average wage	0.13 (0.17)	-0.19 (0.27)	-0.38 (0.30)
<i>number of observations</i>	30,282	25,623	24,083

Note: Estimation samples are constructed using employees who received a merit raise during the week of June 30 - July 7, 1996 (the week that annual merit raises were given), who were employed on July 8, 1996, and whose wages on July 8 are in the same range as the October 1, 1996 wages in the main estimation sample. Scheduled wage are calculated by applying the corporate rule for wage adjustments made on October 1, 1996 to employees present on July 8, 1996. Peers are coworkers employed on July 8, 1996 whose wage on that day was +/- \$.20 from the worker's own wage. Models include "baseline" controls as in Table 2, row 3 (for Panel A) and Table 3, row 1 (for Panel B). Parentheses contain robust standard errors clustered on store.

Appendix Table 1. Estimates from RD models for predetermined characteristics

Dependent variable	Model	
	(1) no controls	(2) baseline controls
employee characteristics:		
1. age (years)	-10.31** (3.39)	-2.10 (3.54)
2. tenure (months)	-4.17** (1.03)	-2.98** (1.01)
3. female	-0.08 (0.18)	-0.06 (0.19)
4. white	0.09 (0.20)	0.03 (0.18)
5. log (ZIP median income)	-0.79 (0.72)	-0.82 (0.76)
6. size of last merit raise (%)	-0.02† (0.01)	-0.01 (0.01)
coworker characteristics:		
7. log(# employees)	0.60* (0.25)	-0.12 (0.15)
8. % got scheduled raise	-0.33** (0.10)	-0.06 (0.06)
9. % got merit raise	-0.06 (0.08)	0.03 (0.06)
10. average age	0.10 (1.34)	-0.77 (0.94)
11. average wage	0.58** (0.19)	0.09 (0.12)
12. % with 5-cent wage	0.48** (0.08)	0.03 (0.06)
<i>number of observations</i>	<i>6691</i>	<i>6691</i>

Note: Table entries in column 1 are regression coefficients on scheduled wage from models that control only for a linear function of the initial wage. Entries in column 2 are from models with the baseline set of controls as in table 2, row 3. Models for tenure (row 2) include a dummy variable for employees whose tenure is censored. Parentheses contain robust standard errors clustered on store. † significant at 10%; * significant at 5%; ** significant at 1%.

Appendix Table 2. Estimates from RD models for probability of being in quits analysis sample

	<i>In estimation sample for quitting within:</i>				
	(1)	(2)	(3)	(4)	(5)
	1 month	2 months	3 months	6 months	9 months
1 no controls (as in table 2, row 1)	-0.05 (0.11)	0.01 (0.13)	0.02 (0.13)	0.15 (0.18)	0.07 (0.18)
2 controls as in table 2, row 3	-0.14 (0.12)	-0.02 (0.14)	0.02 (0.15)	0.10 (0.19)	0.02 (0.20)
3 controls as in table 2, row 5	-0.11 (0.12)	0.00 (0.14)	0.03 (0.15)	0.13 (0.19)	0.05 (0.20)
<i>number of observations</i>	6691	6691	6691	6691	6691

Note: Table entries are regression coefficients on the scheduled raise from models for the probability of being in the sample (i.e. not separating for a reason other than quitting). Row 1 shows models that control only for a linear function of the wage (as in table 2, row 1). Entries in row 2 are from models that control for 3-digit ZIP code fixed effects and a dummy for wages that are a multiple of \$.05 (as in table 2, row 3). Entries in row 3 are from models that include controls for employee and coworker characteristics (as in table 2, row 5). Parentheses contain robust standard errors clustered on store.

Appendix Table 3. MRD first-stage and reduced-form estimates of quit response to changes in own and peer wages

	<i>Dependent variable = 1 if quit within:</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	peer avg. wage	1 month	2 months	3 months	6 months	9 months
1 baseline controls (see note)						
own wage		-0.26* (0.11)	-0.56** (0.16)	-0.62** (0.17)	-0.73** (0.23)	-0.76** (0.27)
peer wage	0.44** (0.01)	0.10 (0.06)	0.20* (0.09)	0.22* (0.10)	0.24† (0.13)	0.23 (0.15)
2 control for employee characteristics and store average coworker characteristics						
own wage		-0.25* (0.11)	-0.57** (0.16)	-0.65** (0.17)	-0.77** (0.23)	-0.79** (0.27)
peer wage	0.44** (0.01)	0.10 (0.06)	0.19* (0.09)	0.22* (0.10)	0.25† (0.13)	0.24 (0.15)
3 additional controls for peer characteristics and quadratic in peer initial wage						
own wage		-0.25* (0.11)	-0.56** (0.16)	-0.65** (0.17)	-0.76** (0.23)	-0.79** (0.27)
peer wage	0.38** (0.01)	0.14* (0.06)	0.24** (0.08)	0.25** (0.10)	0.29* (0.12)	0.28† (0.15)
<i>number of observations</i>	45,672	42,549	41,180	40,470	36,961	34,695

Note: First-stage (column 1) and reduced-form (columns 2-6) estimates for MRD models of peer effects as in Table 3. See also Table 3 note. Parentheses contain robust standard errors clustered on store. † significant at 10%; * significant at 5%; ** significant at 1%.

Appendix Table 4. Estimates from MRD models for predetermined characteristics

Dependent variable	Running variable:	
	(1) own wage	(2) peer wage
employee characteristics:		
1. age (years)	-2.65 (3.48)	-2.41 (1.94)
2. tenure (months)	-3.15** (1.01)	0.30 (0.55)
3. female	-0.04 (0.19)	0.06 (0.11)
4. white	0.04 (0.18)	0.12 (0.10)
5. log (ZIP median income)	-0.48 (0.73)	-0.45 (0.49)
6. size of last merit raise (%)	-0.01 (0.01)	-0.01** (0.00)
coworker characteristics:		
7. log(# employees)	-0.09 (0.15)	-0.15 (0.12)
8. % got scheduled raise	-0.08 (0.06)	-0.04 (0.05)
9. % got merit raise	0.06 (0.06)	0.11* (0.05)
10. average age	-0.61 (0.90)	0.30 (0.73)
11. average wage	0.16 (0.11)	-0.05 (0.09)
12. % with 5-cent wage	0.01 (0.05)	-0.08 (0.05)
<i>number of observations</i>	45,672	45,672

Note: Table entries are regression coefficients from reduced-form MRD models with baseline controls as in table 3, model 1. Model for tenure (row 2) includes a dummy variable indicating employees whose tenure is censored. Parentheses contain robust standard errors clustered on store. † significant at 10%; * significant at 5%; ** significant at 1%.

Appendix Table 5. Split-sample RD estimates of effect of own scheduled wage on quit rates

Sample (same or opposite side of pay step threshold):	<i>Dependent variable = 1 if quit within:</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	own - peer wage	1 month	2 months	3 months	6 months	9 months
<u>opposite side:</u>	1.67** (0.05)	-0.44* (0.19)	-0.91** (0.27)	-1.17** (0.30)	-1.05* (0.43)	-1.22* (0.53)
<i>number of obs.</i>	2426	2279	2222	2197	2035	1904
<u>same side:</u>	0.54** (0.04)	0.00 (0.17)	-0.17 (0.22)	-0.16 (0.26)	-0.27 (0.35)	-0.33 (0.43)
<i>number of obs.</i>	3930	3664	3558	3493	3195	2988
Difference (opposite side - same side)	1.13** (0.06)	-0.44† (0.25)	-0.74* (0.33)	-1.01** (0.39)	-0.78 (0.55)	-0.88 (0.68)
Wald estimate of peer effect		0.39	0.65	0.88	0.69	0.78

Note: Entries are regression coefficients on own wage from linear probability models of quits with controls as in table 2, row 3. Sample split by whether the majority of one's peers are on the same side or opposite side of a discontinuity threshold. Estimates and standard errors for the difference are from a model in which own wage and all covariates are interacted with a dummy for the opposite-side sample. Wald estimate of peer effect is estimated using two-stage least squares. (See text for derivation.) Parentheses contain robust standard errors clustered on store. † significant at 10%; * significant at 5%; ** significant at 1%