

NBER WORKING PAPER SERIES

FAIRNESS AND FRICTIONS:
THE IMPACT OF UNEQUAL RAISES ON QUIT BEHAVIOR

Arindrajit Dube
Laura Giuliano
Jonathan S. Leonard

Working Paper 24906
<http://www.nber.org/papers/w24906>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2018

We thank Joshua Angrist, Emily Breza, David Card, Christian Dustmann, Ethan Kaplan, Suresh Naidu, Todd Sørensen and seminar participants at UC Berkeley, UC Irvine, UC Merced, UC Santa Cruz, LSU, MIT, UMass Amherst, UNR, UVA, and the NBER Summer Institute for Personnel Economics for helpful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Arindrajit Dube, Laura Giuliano, and Jonathan S. Leonard. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior
Arindrajit Dube, Laura Giuliano, and Jonathan S. Leonard
NBER Working Paper No. 24906
August 2018
JEL No. D9,D91,J01,J3,J42,J63

ABSTRACT

We analyze how separations responded to arbitrary differences in own and peer wages at a large U.S. retailer. Regression-discontinuity estimates imply large causal effects of own wages on separations, and on quits in particular. However, this own-wage response could reflect comparisons either to market wages or to peer wages. Estimates using peer-wage discontinuities show large peer-wage effects and imply the own-wage separation response mostly reflects peer comparisons. The peer effect is driven by comparisons with higher-paid peers—suggesting concerns about fairness. Separations appear fairly insensitive when raises are similar across peers—suggesting search frictions and monopsony are relevant in this low-wage sector.

Arindrajit Dube
Department of Economics
University of Massachusetts
Crotty Hall
412 N. Pleasant Street
Amherst, MA 01002
and NBER
adube@econs.umass.edu

Jonathan S. Leonard
Haas School of Business
University of California
Berkeley, CA 94720
leonard@haas.berkeley.edu

Laura Giuliano
Department of Economics
University of California, Santa Cruz
1156 High Street
Santa Cruz, CA 95064
lgiulian@ucsc.edu

A data appendix is available at <http://www.nber.org/data-appendix/w24906>

1 Introduction

When making decisions about job mobility, do workers mainly compare their pay with the outside market? Or do they also make comparisons with the pay of co-workers in their own firm? If workers are very sensitive to wage differences across firms as in the standard competitive labor-market model, then employers have little scope to set wages. But even if a firm’s wage-setting is not fully constrained by market competition, employers could still be constrained by employee concerns about relative pay within the firm. Such internal constraints could lead to wage compression or to measures that inhibit comparisons across coworkers—such as pay secrecy and outsourcing of jobs. Estimates of how employee turnover responds to external and internal wage comparisons are therefore central to understanding the wage-setting behavior and other employment practices of firms. Yet the endogeneity of wages in natural employment settings has made causal evidence hard to establish.

In this paper, we analyze how job separations responded to arbitrary wage differences among sales employees at a large U.S. retailer with hundreds of stores nationwide (henceforth “the firm”). Using quasi-experimental variation that resulted from a rule-based formula for pay raises, we examine how separations were affected both by the size of one’s own raise and by the average raise of one’s peers within the store. We make three main contributions. First, we advance the literature on relative pay in the workplace by establishing that changes in peer wages can have large, causal effects on job separations and on quits in particular. Second, we use the peer-wage estimate to decompose the effect of own wage changes on separations into a response that is due to relative pay comparisons within the workplace and a response due to comparisons with the market. The latter response allows us to measure the wage-setting power of the employer. Third, we present evidence that the separation response to peer wages cannot be explained by rational learning about own future wage growth; instead, the patterns we find suggest concerns about fairness.

In response to the federal minimum wage increases in 1996 and 1997, our firm implemented a policy that raised wages above the new minimum. Two features of this policy provide the basis for our research design. First, the firm raised wages by applying a uniform rule nationwide irrespective of any local or individual characteristics other than the initial wage. Second, this rule 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 wages differed by only one cent, received new wages that differed by ten cents—or roughly 2% of the average wage.

We begin by estimating the response of separations with respect to own wages using a sharp regression discontinuity (RD) design that exploits the discontinuities in the new wage as a function of the initial wage. These estimates show a strong causal effect of wages in reducing separations. However, the interpretation of this effect is ambiguous. If workers make only market comparisons, then a large separation elasticity is indicative of a highly competitive labor market. But if workers also compare their own wages with those of their peers, then the size of the separation response reflects not only how much workers react to the market but also how much they care about relative pay within the workplace.

To distinguish between these two channels, we estimate a multi-dimensional RD (MRD) model of separations that incorporates both a sharp RD in the own-wage (as before) and a fuzzy RD that uses peer-wage discontinuities to instrument for the average peer wage.¹ In this model, the own-wage estimate continues to give the total effect of wages on separations. 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 holds relative pay constant—i.e., the effect of market comparisons.

Overall, we find strong evidence that peer comparisons matter. Our estimates suggest separations are highly responsive to increases in the average peer wage—with elasticities of 20, 9 and 3 for 3, 6 and 9 months after the raise. For the median worker in our sample, who separates roughly 9 months after the raise, the estimates imply that a \$0.10 increase in the average raise among peers reduces tenure by about a month. Compared to the own-wage responses, the estimates of peer-wage separation responses are broadly similar in magnitude but opposite in sign—which implies that relative-pay concerns account for much of the total effect of own wages on separations. In contrast, market wage comparisons appear relatively unimportant in explaining the separation response. The estimates suggest that when a raise is uniform across peers (so that the gap between own and peer wage is held constant), the elasticities are not statistically distinguishable from zero.

¹We define peers generally as coworkers in the same job and store who earn broadly similar initial wages. Our preferred specification uses wages within 30 cents of one’s own initial wage (roughly 70% of the low-wage coworkers who receive a wage adjustment). Section 5.6 explains our data-driven procedure for choosing the peer definition and presents estimates based on alternative definitions.

To check this interpretation, 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 opposite side or same side of a pay-step threshold. These estimates show that the effects of own wages on separations are large in the opposite-side sample where workers receive raises that differ from those of their peers, but are small in the same-side sample where the raises are similar across peers. Consistent with the MRD results, these results also imply that the overall effect of wages on separations is driven mostly by peer comparisons.

Further analysis sheds light on the motives underlying the peer-wage effect. We find that the effect is asymmetric: it is driven by comparisons with higher-paid peers. This asymmetry suggests workers are averse to unequal treatment that is to their disadvantage, and we interpret this aversion to disadvantageous inequity as evidence of concerns about fairness.² An alternate theory of peer-wage effects that does not rely on other-regarding preferences is that workers use peer wages to help predict their own future pay.³ However, the peer-wage response in our setting cannot be easily explained by such rational learning. We find no evidence that peer raises predict future wage growth—which is consistent with the arbitrary nature of the raises we study. Moreover, even under very conservative assumptions about selective attrition, the signal conveyed by peer wages about future raises can explain no more than 20% of the separation effect we find.

Our findings on relative pay complement a set of recent field experiments that find negative effects of pay disparities in the workplace on morale, job satisfaction, and effort (Breza, Kaur and Shamdasani 2017; Card et al. 2012; Cohn et al. 2014). We advance the literature by documenting the causal effects of relative pay in a non-experimental setting and by showing that the effects extend to job separations.⁴ Our paper is most closely related to Breza, Kaur and Shamdasani (2017), who randomly vary wage increases across worksites in India and find strong evidence that wage inequality causes reductions in

²See Fehr and Schmidt (1999) on aversion to disadvantageous inequity as a notion of fairness.

³The standard learning model posits that having peers with higher pay sends a positive signal about future wage growth. Consistent with this hypothesis, some non-experimental studies have found that low relative pay is associated with higher job satisfaction and lower quit rates (Clark, Kristensen and Westergård-Nielsen 2009; Galizzi and Lang 1998; Pfeifer and Schneck 2012); however, others have found patterns that are the opposite in sign and consistent with our results (Clark and Oswald 1996; Rege and Solli 2013). The mixed evidence in this literature is hard to interpret, partly because the endogeneity of wages raises questions about causality. It is also likely that signaling value of peer wages varies across settings. In some cases, having higher-paid peers may signal negative wage growth due to poor match quality (e.g., Buntrock 2014); while in other cases (like ours) there is no signaling value.

⁴Card et al. (2012) examine actual separations, but their estimates are too imprecise to reach a definite conclusion and are likely attenuated by the diffusion of their information treatment over time to employees in their control group.

attendance and productivity. Our results also complement those of Card et al. (2012), who find that employees at the University of California reported lower job satisfaction and increased likelihood of job search after learning they earned less than their peers. Despite differences across settings and outcomes, the magnitude of our estimates are broadly consistent with those implied by these two papers (see Section 7).

Our findings also have implications for the wage-setting power of the firm. Consistent with models of monopsonistic competition (Burdett and Mortensen 1998; Manning 2003), the modest response of separations to an across-the-board wage increase indicates the presence of search frictions that give the firm significant monopsony power. While prior empirical studies of labor market monopsony have reached similar conclusions, only a few have used quasi-experimental variation to estimate separation responses to firm-specific wage changes, and these studies focus on a narrow range of settings.⁵ Moreover, the evidence for low-wage labor markets is particularly scarce and the empirical relevance of monopsony power in these markets remains a subject of debate.⁶ Our results thus complement the extant literature and provide novel evidence that considerable wage-setting power exists in a setting characterized by low wages and high turnover. At the same time, they also suggest that employee concerns about fairness may serve as an important constraint on a firm’s wage-setting behavior beyond the constraints of the market. Such internal constraints could help explain why the rise in inequality has occurred mainly between firms (Song et al. 2015) and why firms have increasingly re-structured employment in ways that limit interaction between workers who are paid very differently (Weil 2014).

⁵Falch (2011) and Ransom and Sims (2010) both use plausibly exogenous wage variation in the market for teachers; they find separation elasticities of -3.5 and -1.8 respectively. Mas (2017) estimates an elasticity of -11 using quasi-experimental wage variation for city managers. Using observational data, several studies have estimated separation elasticities between -2 and zero, implying large search frictions (Depew and Sørensen 2013; Hirsch, Schank and Schnabel 2010; Ransom and Oaxaca 2010; Webber 2015); however, these estimates are potentially biased toward zero due to the endogeneity of wages (Manning 2003). Others have studied employer responses to labor market regulations to assess labor market frictions. Evidence of substantial monopsony power has 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 Wang 2016). Using both experimental and machine learning evidence, Dube et al. (2018) find evidence of strong monopsony power in online labor markets, with labor supply elasticities smaller than 0.5. In contrast, however, Matsudaira (2014) finds no evidence of monopsony power in a study of minimum staffing regulation for nurses aides.

⁶Much theoretical work on labor market frictions has been stimulated by empirical research showing that minimum wages do not always have negative effects on employment as predicted by a competitive model of low-wage labor markets (Card and Krueger 2016). Though estimates of the employment response to a minimum wage increase are mixed, recent studies have produced evidence consistent with search frictions in low-wage labor markets such as the retail setting we study here (e.g., Giuliano 2013; Dube, Lester and Reich 2016).

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 present the regression discontinuity estimates of the total effect of own wages on separations. 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 additional falsification tests, and Section 7 concludes.

2 Theoretical Framework

In the canonical “job ladder” model of on-the-job search (Burdett and Mortensen 1998), separations occur either as exogenous transitions to non-employment, or as endogenous transitions (quits) to jobs offering wages that exceed the worker’s current wage w . The separation rate is given by $S(w) = \delta + \lambda [1 - F(w)]$, where $F(w)$ is the wage offer distribution, δ is the exogenous separation rate, and λ is the offer arrival rate; search frictions are captured by lower values of λ .⁷

A key assumption of this model is that wages vary across firms but not within firms. The separations (or quits) response to a wage increase therefore depends only on market comparisons, and is given by $\frac{dS}{dw} = -\lambda f(w)$. Hence in a context where wages vary only across firms, the separation 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 model to allow for internal comparisons by introducing a reference wage, w_p , which is a function of the wages earned by one’s peers. We now assume that a worker’s job satisfaction U depends not only on her own wage, w , but also on the *gap* between her own wage and the reference wage, $w_g = w - w_p$:⁹

$$U(w, w_p) = v_0 w + v (w - w_p) = v_0 w + v (w_g).$$

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.

⁸Manning (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.

⁹Our approach is similar to that used elsewhere in the literature on peer comparisons in the workplace (e.g., Charness and Kuhn (2007), Fehr and Schmidt (1999), and Card et al. (2012)).

When $v'(\cdot) = 0$, workers are only self-regarding, and we revert to the standard job ladder model. However, when $v'(w_g) > 0$, workers also care about their pay relative to the reference wage; and in the case where $v_0 = 0$, 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. This formulation also allows for asymmetries depending on whether workers are earning more or less than their peers; in particular, when $v''(w_g) < 0$, workers care more about relative pay when $w_g < 0$ than they do when $w_g \geq 0$, as in Fehr and Schmidt's (1999) model of fairness.

How does a worker choose between her current job (with wage w) and a new wage offer w' ? Since 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, 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:

$$U = U(w, w_p) = v_0 w + v(w - w_p) = v_0 w + v(w_g) < v_0 w' = U'.$$

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

$$\begin{aligned} S(w, w_p) &= \theta\delta + \lambda \left[1 - F\left(w + \frac{v(w-w_p)}{v_0}\right) \right] \\ &= \theta\delta + \lambda \left[1 - F\left(w + \frac{v(w_g)}{v_0}\right) \right]. \end{aligned} \quad (1)$$

Differentiating equation (1) with respect to w and w_p gives the partial effects of own and peer wage increases, which can be expressed in terms of the wage gap w_g as:

$$\text{and} \quad \begin{aligned} \frac{\partial S(w, w_p)}{\partial w} &= -\lambda \cdot f\left(w + \frac{v(w_g)}{v_0}\right) \cdot \left(1 + \frac{v'(w_g)}{v_0}\right) \\ \frac{\partial S(w, w_p)}{\partial w_p} &= \lambda \cdot f\left(w + \frac{v(w_g)}{v_0}\right) \cdot \left(\frac{v'(w_g)}{v_0}\right). \end{aligned} \quad (2)$$

These two separation responses are estimated directly in our empirical analysis. Importantly, $\frac{\partial S(w, w_p)}{\partial w}$, the total effect of an own wage increase on separations, reflects two conceptually distinct responses—one based on market comparisons and the other based on peer comparisons. Totally differentiating equation (1) with respect to w gives the decomposition $\frac{dS}{dw} = \frac{\partial S(w, w_g)}{\partial w} + \frac{\partial S(w, w_g)}{\partial w_g} \frac{dw_g}{dw}$ where:

¹⁰While the assumption $w'_p = w'$ simplifies the exposition, it is not necessary for our key results on identification of the separation 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 $v_0 w + v(w - w_p) < v_0 w' + v(w' - w'_p)$. If we interpret $F(\cdot)$ as the distribution function of the utility from the outside offer scaled by $\frac{1}{v_0}$, i.e., $w' + \frac{v(w' - w'_p)}{v_0}$, the expressions for the separation function (equation 1) remain the same, as do the expressions for all the subsequent separation responses.

$$\begin{aligned}
& \text{and} \quad \frac{\partial S(w, w_g)}{\partial w} = -\lambda \cdot f\left(w + \frac{v(w_g)}{v_0}\right) \\
& \frac{\partial S(w, w_g)}{\partial w_g} \frac{dw_g}{dw} = -\lambda \cdot f\left(w + \frac{v(w_g)}{v_0}\right) \cdot \left(\frac{v'(w_g)}{v_0}\right) \frac{dw_g}{dw}.
\end{aligned} \tag{3}$$

The first term in equation (3) 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 separation 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 more workers care about relative pay and the larger the increase in the wage gap w_g .

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 (3) 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: $-\frac{\partial S(w, w_p)}{\partial w_p}$.¹¹ Second, we get the gap-constant separation response by subtracting the relative wage effect from the total own-wage effect, or equivalently by summing the own-wage and peer-wage effects:

$$\frac{\partial S(w, w_g)}{\partial w} = \frac{\partial S(w, w_p)}{\partial w} + \frac{\partial S(w, w_p)}{\partial w_p}. \tag{4}$$

The ratio of the two separation responses is the marginal rate of substitution between relative and absolute pay:

$$\alpha(w - w_p) = \frac{\frac{\partial S(w, w_p)}{\partial w_p}}{\frac{\partial S(w, w_p)}{\partial w}} = \frac{v'(w_g)}{v_0 + v'(w_g)}, \tag{5}$$

which can be interpreted as the compensating differential for a higher peer wage. For example $\alpha = 0.5$ means the worker will accept a \$0.50 wage reduction to avoid a \$1.00 increase in the peer wage.

¹¹In Section 5.5, we also construct a Wald estimate for the relative-pay effect, which is based on the difference in own-wage separation responses in two sub-samples—one where raises differ across peers and one where the raises are more similar.

Our empirical analysis begins by assuming that $v(w_g)$ is linear and symmetric in w_g ; in Section 5.7, we relax this assumption by allowing a worker’s separation response to vary with the *ex ante* wage gap. We are particularly interested in whether the response is asymmetric – as would be the case if workers are especially averse to disadvantageous inequity.

We consider another potential source of nonlinearity in Online Appendix B, where we extend the theoretical model to allow search intensity—and thus the offer arrival rate λ —to depend on own and peer wages. As before, the total effect of an increase in own wage is the sum of the gap-constant effect and the peer-wage effect, and the ratio of the own and peer-wage separation responses still recovers the compensating differential for an increase in the peer wage. But now, both responses also depend on the sensitivity of the offer arrival rate to search activity, which opens the door to further non-linearities. With a fixed search cost, for example, small reductions in w_g may produce large increases in separations while further reductions have smaller marginal effects. The possibility of such non-linearities cautions against extrapolating the results from small wage increases to much larger ones.

Our theoretical framework does not include an information channel whereby peer wages provide 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 separation behavior in our setting. We address this point more fully in Section 5.8.

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. As a matter of policy, this chain pursues uniform benefits, working conditions, and job duties across all stores. We analyze the separation behavior of employees in a single, entry-level sales job that accounts for 90 percent of the firm’s retail workforce. This job involves customer service and various support duties; it requires only basic skills and employees receive cursory on-the-job training. This is a relatively low-wage job in which hourly wages are the main form of compensation and there is little expected wage growth. Employees do not receive commissions or performance-based bonuses; and

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 that are given annually. All those employed for at least 90 consecutive days are eligible for the annual merit raise, and roughly 80% of eligible employees receive one. These raises are determined by store managers and averaged 2.2% across all workers.

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. On each of these dates, the firm applied a uniform rule to all hourly employees nationwide irrespective of any local or individual characteristics other than initial wage. This policy increased wages substantially more than was necessary to comply with the law. Whereas the share of hourly retail employees who earned less than the new minimum was roughly 5% in 1996 and 10% in 1997, the firm extended raises to the 30th percentile of the wage distribution in 1996 and to the 40th percentile in 1997.¹² The average raise was \$0.21 or 4.1%.

A key feature of the firm’s policy—and the source of the arbitrary wage 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

$$\Delta w_y = w_y - w_{0y} = \begin{cases} (MW_{1y} - w_{0y}) + 0.10 \times \text{int} \left(\frac{w_{0y} - MW_{0y}}{0.15} \right) & \text{if } w_{0y} \in [MW_{0y}, \bar{w}_{0y}) \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

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 in 1996 and \$5.65 in 1997. In both years, 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

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, 2016). It is thus noteworthy that our firm implemented sizable 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 (6) 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 0.5% of raises differ from the scheduled raise, so the scheduled raise predicts the actual raise very closely (see Figure 2).

(“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 firm’s wage policy created multiple discontinuities in the relationship between the initial wage and the new wage. 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 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 provide the identifying variation in both own and peer wages that we use to estimate separation responses.

A prerequisite condition for relative pay to affect separations is that coworkers learn about each other’s wages. While we have no direct evidence about what our employees knew about coworker wages, several features of our context may have promoted information sharing. First, the raises were highly salient - they were relatively large (4.1% on average vs. 2.2% for a typical merit raise) and were given simultaneously in the context of a federal minimum wage increase.¹⁵ Further, since workers who got these raises likely assumed that their similarly paid coworkers did as well, they may have felt relatively comfortable discussing their raises with their peers.¹⁶ Finally, it is worth noting that since our identifying variation in the average peer wage, \bar{w}_p , is due to differences in the fraction of peers above or below the nearest threshold, we need not assume that workers learn the exact wage of each peer. We would obtain the same results if workers instead learn the fraction of peers that gets “lucky”.

3.2 Sample construction and descriptive statistics

In the job we analyze, there were 10,390 workers scheduled for raises on October 1, 1996, and 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 $w_{0y} \in [MW_{0y} + 0.08, \bar{w}_{0y} - 0.08]$, which ensures a consistent bandwidth of \pm \$0.07 around each of the twelve pay-

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.

¹⁵In another context where large raises were given simultaneously to all workers in a production unit, Breza, Kaur and Shamdasani (2017) find direct evidence of substantial learning about coworker wages. Although managers maintained pay secrecy, 87% of workers were able to accurately report the wages of both coworkers in their 3-person production unit within 3 weeks of the raises.

¹⁶We have no information on what managers told employees about the raise schedule, and this information may have differed across stores. But it seems likely that workers were aware of the federal minimum wage increase and were aware that the raises resulted from this increase.

step thresholds. Second, we exclude workers who are still earning their starting wages by restricting 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. Including starting wages in the sample would therefore lead to sharp discontinuities in employee tenure and other associated characteristics.¹⁷

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 each 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.

Panel C shows summary statistics for several variables describing an employee’s coworkers; these variables serve as controls in some of our models. Coworkers are defined as all those who work in the same, entry-level job and in the same store on the day of the minimum wage increase—including workers who do not meet the conditions for being in the estimation sample. A typical employee in our sample has 27 such 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).

¹⁷Since 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.

Because estimation of peer wage effects requires exogenous variation in peer wages, our analysis focuses on wage increases among coworkers in the same job whose initial wages were in the range that received a scheduled raise. We consider all such coworkers *potential* members of an employee’s peer group—including coworkers who are not themselves in the estimation sample because they have not received a merit raise.¹⁸ The typical sample employee has roughly 13 such “potential peers”; characteristics of these peers are summarized in Panel D. Because some potential peers may be more relevant than others, we treat the definition of one’s “peer group” as an empirical question (see Section 5.6). Panel E of Table 1 summarizes peer-group characteristics using our preferred definition: potential peers whose initial wages are within \pm \$0.30 of the employee’s own initial wage. On average, employees in our estimation sample have 9 such peers, and the average peer wage (\$5.13) and average peer raise (\$0.21) are both close to the means for sample employees themselves.

Our primary outcome of interest is the probability of separating from the workplace within a window following one of the minimum wage increases. We examine windows of 1, 2, 3, 6 and 9 months.¹⁹ The separation rate ranges from 12% within one month of a wage increase to 51% within 9 months (Panel F). Over 90% of separations in our sample are voluntary, and roughly 55% of those who leave within 9 months report job-related reasons for leaving (e.g., dissatisfaction with the job or finding a better one). Thus the majority of separations correspond closely to the “quits” or job-to-job transitions in our theoretical framework. Other reasons for leaving including returning to school (21%), moving or transferring to another store (16%), and being fired (8%). To guard against bias due to potentially endogenous competing risks, most of our analysis does not distinguish between reasons for separating; however, in Section 5.4, we show that our main conclusions do not change if we exclude “non-quit” separations from the analysis.

¹⁸Our main conclusion that separations respond to higher peer wages is robust to excluding coworkers who had not received a merit raise. However, estimates based on the narrower peer-group definition are somewhat attenuated, which suggests that the restriction excludes relevant peers (see Section 5.6).

¹⁹Although 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 diluted by merit raises given in June 1997 and in June 1998 (9 and 10 months, respectively, after the minimum wage increase).

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

4.1 RD estimation framework

To analyze the effect of own wage increases on separation decisions, we use an RD design that exploits discontinuities in the scheduled raise formula shown in equation (6). We begin by examining the relationship between the scheduled raise and the actual raise received on the date of each 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 ($R^2 = 0.998$); both decline linearly with the initial wage except for the positive 10-cent jump at each 15-cent interval. 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.”²⁰

Figure 3 presents visual evidence that the positive wage discontinuities seen in Figure 2 are associated with negative discontinuities in separations. After partialing out ZIP-code based fixed effects,²¹ we plot residual separation rates by initial wage and include a non-parametric fitted line for each 15-cent interval corresponding to a new pay step. The figure shows visible negative discontinuities at 7 of the 12 thresholds and a visible positive discontinuity at only one of the steps.

While Figures 2 and 3 provide motivation for our research design, Figure 3 also reveals a considerable amount of noise in separation rates – which limits our power to estimate each discontinuity separately. To gain power, we pool the data from the 1996 and 1997 raises and estimate parametric models of the form:

$$S_{iy}^m = \beta \times w_{iy} + f_y(w_{0iy}) + X_{iy}\Gamma + \lambda_{z(i)} + \epsilon_{iy}. \quad (7)$$

The dependent variable, S_{iy}^m , is an indicator for whether employee i leaves the store within m months of the year y raise; $w_{iy} = w_{0iy} + \Delta w_{iy}$ is the scheduled new wage;

²⁰The impact of the scheduled raise on wages tends to be diluted over time due to wage growth from promotions; however, its impact is quite persistent over the windows we examine. Even after six months, a \$0.10 discontinuity in the scheduled raise predicts a discontinuity of \$0.095 in actual wage growth; after 9 months the equivalent discontinuity is \$0.078 (see Appendix Table A1). The discontinuity is diluted further (to 0.075)–but is not offset–by subsequent merit raises.

²¹For reasons discussed later in this section, we use fixed effects for the store’s 3-digit ZIP code to limit comparisons to similar labor markets. Section 4.2 discusses the robustness of our results to alternative specifications.

$f_y(w_{0iy})$ is a smooth function of the initial wage; and X_{iy} is a vector of employee and store characteristics that is used for sensitivity tests and varies by specification.²² Because the scheduled wage is a close proxy for the actual wage, we can interpret β as the effect of a \$1.00 wage increase on the probability of separating.

If the specification for f_y is sufficiently flexible, then β is estimated using only the discontinuities in the scheduled wage function. Our simplest specification controls linearly for the initial wage and allows a different intercept and slope in each year. However, when pooling across multiple thresholds, potential misspecification of f_y is of particular concern—and in our context, a linear function may be insufficient. One reason is that the global relationship between separations and wages reflects both variation across workers in the same market and variation across different labor markets. In our preferred specification, we control flexibly for the cross-market heterogeneity using fixed effects, $\lambda_{z(i)}$, for 3-digit ZIP codes based on the store’s location (hereafter “ZIP codes”).²³ We then use a linear function to control for variation within markets, though we also document robustness using quadratic and cubic functions.²⁴

As an alternative to the global RD model in equation (7), we also estimate a “stacked” RD model in which we stack the 15-cent wage intervals that are centered at a pay threshold and control for the distance from w_{0iy} to the nearest threshold. Formally, 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} .²⁵ This specification more closely resembles the standard RD setup, which makes it useful for visual representation of the data. However, unlike the traditional RD setup, the “above” and “below” side of the discontinuity

²²Our approach is similar to Angrist and Lavy (1999) and Angrist, Battistin and Vuri (2016). 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 these studies, we use the discontinuous treatment variable as a regressor while controlling for a smooth function of the running variable (total enrollment in their case, the initial wage in ours).

²³Within-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. We also obtain qualitatively similar results from models with store fixed effects (see Tables 2 and 3).

²⁴The decision to use a low-order polynomial with a separate intercept for each labor market (instead of simply adding higher-order terms to f_y) is theoretically motivated but also has practical advantages. For one, the small number of bins and the pooling across many thresholds limits the usefulness of higher-order functions in our context. Gelman and Imbens (2016) describe other reasons why reliance on higher-order polynomials should be avoided in RD designs.

²⁵Workers are considered “above” the threshold when $w_{0iy} \geq T_y^k$, or equivalently $r_{iy} \geq 0$. They are considered “below” when $r_{iy} < 0$.

are not ordered in terms of the wage w_{0iy} (since the wage intervals are stacked), and this prompts us to constrain the slopes on the running variable to be the same on both sides. (We show that this increases power without significantly affecting the results.) The stacked model estimates are similar to the global RD estimates and are also robust to the use of narrower bandwidths (less than \$0.07) around each threshold.

In both variants of our model, the key identifying assumption for a causal interpretation of β is that the latent propensity to separate evolves smoothly across the thresholds. Of potential concern here is the fact that even after excluding starting wages, the distribution of wages in our sample is still quite lumpy with bunching that appears mainly at multiples of \$0.05 (Appendix Figure A1). When bunching occurs at threshold values, it often raises concerns about precise manipulation of the running variable, which can invalidate the smoothness assumption (Lee and Lemieux 2010; McCrary 2008). In our setting, a concern might be that managers anticipated the raise schedule and “topped up” the merit raises of their most valuable workers to ensure they were at a higher pay step when the minimum wage increased. However, as we show in Appendix C, the evidence in our case is not consistent with such manipulation. First, the bunching occurs not only at the pay-step thresholds but also at non-threshold wages – and there is no excess mass at the pay-step thresholds. Second, the bunching does not lead to discontinuities in the size of the merit raise. Instead, the bunching pattern is similar to the “round number” heaping found in other payroll data.²⁶

A more pertinent concern is that even in the absence of precise manipulation, heaping at round numbers can result in bias if the heaping is not accounted for (Barreca, Lindo and Waddell 2016). We employ three complementary methods to remove heaping-induced bias. Our preferred approach is to include a dummy variable in equation (7) for wages that are multiples of \$0.05. For robustness, we also present two alternatives: a “donut-hole” specification that excludes all initial wages just at or below a threshold (i.e., $r_{iy} = 0$ and $r_{iy} = -1$) and another that uses *only* wages that are multiples of \$0.05. These two alternatives both rely on extrapolation using data away from the thresholds, but they use different variation to guard against mixing \$0.05 and non-\$0.05 multiples.

Models that satisfy the conditions for valid RD-based inference should show no evidence of discontinuities for predetermined covariates. In Appendix Table A2, we test this implication for various specifications of our model, using predicted separations as

²⁶Dube, Naidu and Manning (2017) find that bunching of wages at round numbers is pervasive in administrative hourly wage data from the Unemployment Insurance records of several states. They also present evidence suggesting that this bunching is consistent with employer optimization frictions and monopsony power.

the outcome.²⁷ Column 1 shows that the simplest specification, which controls only for a global linear function of the initial wage, produces statistically significant positive discontinuities.²⁸ As discussed above, there are two likely sources of misspecification in this model. First, a linear wage function may not sufficiently control for the global relationship between wages and separations. Column 2 shows that the inclusion of ZIP-code fixed effects reduces the discontinuities by about one half. Second, there may be some non-random sorting associated with the bunching of wages at round numbers. Column 3 controls directly for this variation by including a dummy variable for initial wages that are multiples of \$0.05. Once we control for these ZIP-code fixed effects and the 5-cent dummy, the discontinuities in predicted separations are mostly eliminated (though a small and marginally significant discontinuity remains for 3-month separations). Reassuringly, the estimates from the stacked model are also very small and statistically insignificant (column 4),²⁹ and we find similarly small discontinuities from specifications that include higher order terms (columns 5 and 6), replace ZIP-code fixed effects with store fixed effects (column 7), or exclude donut-hole observations (column 8). Only the specification estimated using 5-cent wages (column 9) produces evidence of moderate positive bias; however, the standard errors from this model are also relatively large, consistent with the smaller sample. In our main analysis we estimate models that correspond to all of the specifications in Appendix Table A2; we also demonstrate robustness using models that control for all of the predetermined covariates (including both individual and coworker characteristics). Two falsification tests, presented in Section 6, lend further support to the causal interpretation of our estimates. First, separation rates do not change discontinuously at the 1996 pay-step thresholds in the months *before* the raises were implemented. Second, there are no discontinuities at wages other than the pay-step thresholds, including non-threshold multiples of \$0.05. In all specifications, the

²⁷To construct the predicted separations, we estimate a linear probability model that includes our covariates along with dummies for each value of the new wage w_{1iy} . The covariates are the employee characteristics described in Panel B of Table 1, including a dummy for each month of tenure, and the store-level coworker characteristics described in Panel C of Table 1. The inclusion of the wage dummies fully controls for any influence of wages on separations and ensures that the covariate coefficients from this auxiliary regression are not biased due to a correlation with wages. We then use these coefficients to predict the separation rate.

²⁸The table reports coefficients on the scheduled wage; these must be re-scaled to obtain the change associated with a \$0.10 discontinuity in the wage. For example, the coefficient of 0.12 for 3-month separation in column 1 corresponds to a discontinuity of 0.012 while the mean 3-month separation rate is 0.23 (Table 1, panel F).

²⁹Although not shown in the table the estimated discontinuities for the stacked model are small and not statistically significant even without ZIP FE and dummies for \$0.05 multiples.

standard errors are clustered by store; but our conclusions are robust to clustering at the level of the ZIP code or state, or by the discrete values of the running variable.³⁰

4.2 RD estimates of own wages on separations

Table 2 presents the estimated values of β , the effect of own wage on the separation decision. Each row corresponds to a different window for the separation rate and each column reports a different model specification. The first column controls linearly for the initial wage but does not include any additional controls. These estimates are all negative and imply that a \$0.10 wage increase is associated with a 0.8-4.2 percentage point fall in the separation rate; the 2, 6 and 9-month separation estimates are marginally significant. Recall, however, that the covariate smoothness test in Appendix Table A2 suggests this simple linear specification is misspecified and produces a positive bias.

In the next two specifications, we first add ZIP code fixed effects (column 2) and then a dummy for initial wages that are multiples of \$0.05 (column 3). The coefficients from these models are again negative, but are larger in magnitude than those in column 1 – due primarily to the inclusion of ZIP fixed effects.³¹ The model in column 3 serves as our baseline model, and the estimates here imply that a \$0.10 increase in the wage results in a 0.9 – 6.1 percentage point reduction in the separation rate depending on the time window. Notably, the coefficients increase sharply in magnitude between the 1 and 2-month windows (from 0.09 to 0.47), but they increase by at most another 32% between 2 and 9 months. This pattern implies that most of the separations caused by the wage increase occur within 2 months of the raise.

Column 4 includes the same set of controls as in column 3 but uses the stacked model, which controls linearly for distance from the nearest threshold (the running variable, r_{iy}). The global and stacked models produce very similar estimates; for example, the 3-month coefficients in columns 3 and 4 are -0.52 and -0.50. A graphical representation of the stacked model for 3-month separations is shown in Figure 4. We plot the mean residuals for each value of r_{iy} , along with the fitted relationship between r_{iy} and the residuals. We

³⁰If anything, the standard errors tend to be smaller when we use these alternative methods for inference; hence the standard errors presented in the main tables are conservative (see Appendix Table A3 for details).

³¹Some change in the coefficient is expected given the evidence of covariate imbalance seen in Appendix Table A2. But in addition, this increase in the coefficient is consistent with the own-wage effect being driven largely by peer-wage comparisons, since the ZIP code fixed effects cause the estimates to be identified off of ZIP codes that have greater within variation in raises. In ZIP codes with no such variation (i.e., where raises are similar across all workers and thus all peers in each store) the own-wage estimate (based on between-ZIP variation) identifies a pure market response.

normalize the fitted values so they are equal to zero at the left limit at the threshold; the right limit at the threshold thus represents the effect of a \$0.10 discontinuity in the wage, and the shaded area at $r_{iy} = 0$ shows the confidence interval of $\hat{\beta}$.³² Aside from deviations at -\$0.04 (the most thinly populated bin), Figure 4 shows that the data fit the linear model reasonably well and the plotted points show clear evidence of discontinuities.

Additional results for the stacked model are presented in the Online Appendix. Focusing on 3-month separations, Appendix Table A4 shows that the estimates are generally robust (though less precise) when we allow the slope to change at the threshold. Appendix Table A5 shows how the stacked model estimates vary when we use bandwidths of less than \$0.07 above and below the threshold. The RD estimates are quite stable for bandwidths of \$0.06 and \$0.05. When we narrow the bandwidth to \$0.04 (or smaller), the estimates typically get larger, but the standard errors also increase substantially.

Returning to Table 2, columns 5 and 6 show that the estimates are nearly identical when we replace the linear specification for $f_y(w_{0iy})$ with a quadratic or cubic. Column 7 adds controls for employee and coworker characteristics, including a dummy for each month of tenure, and column 8 replaces the ZIP-code fixed effects with store fixed effects. The estimates from these specifications are slightly bigger than the baseline estimates: for example, the 3-month estimates are -0.56 and -0.57 instead of -0.52. All but the 1-month estimates remain statistically significant at conventional levels.

Column 9 shows the “donut-hole” specification, which is similar to the baseline model except that the estimation sample excludes wages that are just at or just below a pay-step threshold. While the estimates from this model are less precise (the standard errors are 33-40% larger than those from the baseline model), the point estimates remain negative and sizable, ranging from -0.16 and -0.42. Finally, as an alternative way to avoid comparing observations that are multiples of \$0.05 with those that are not, column 10 shows the estimates for a sample that includes *only* the 5-cent wages. Again, while the standard errors are nearly twice as large, the point estimates (which now range from -0.34 to -0.63) are negative and large in magnitude.

Overall, then, the estimates in Table 2 imply significant negative effects of own wages on separation rates. These do not appear to be driven by heterogeneity associated with round-number bunching of the running variable.

³²This is approximately equal to the scheduled wage coefficient scaled by 0.10 (the size of the wage discontinuity).

5 Peer Comparisons: The Effects of Unequal Raises

The finding that workers who receive arbitrarily larger wages are less likely to separate is consistent with two broad explanations. First, the labor market might be highly competitive, causing workers to respond to small differences when comparing their wages to outside offers from other employers. Second, workers might be reacting to changes in how their own pay compares to the wages of others who work in the same store. In this section, we assess the “peer effect” channel directly by asking how separations respond to arbitrary variation in peer wages, and we assess the relative importance of market competition and peer comparisons in driving the observed behavior.

5.1 Defining the peer group

A worker’s peer group potentially includes all coworkers in the same entry-level job in the same store who qualified for a scheduled raise. However, some of these coworkers may be more relevant than others. When making wage comparisons, workers may pay more attention to coworkers whose initial wages were similar to their own. Further, if workers assumed that coworkers with similar wages got similar raises, they might have felt more comfortable discussing their raises with these peers. Or, if workers who are similar on other dimensions (e.g., age, tenure or residential proximity) are more likely to interact with one another, then these other criteria could help define a worker’s peer group.

While the broadest set of coworkers in the same job and same store is a natural starting point for defining peers, we can improve on this choice if some “potential peers” are, in fact, irrelevant. In general, a definition that is too inclusive will reduce identifying variation in the peer average wage—and this, in turn, will reduce the precision of the estimated peer effect. On the other hand, using a definition that is too narrow will attenuate the estimated peer effect, similar to classical measurement error (Cornelissen, Dustmann and Schönberg, 2017).³³ The choice of peer group therefore involves a trade-off between variance and bias.

To select a group of coworkers that best approximates the set of relevant peers, we begin by focusing on subgroups whose initial wages are within a fixed distance from the worker’s own initial wage (see Section 5.6.1). We show that among these wage-based definitions, the trade-off between variance and bias appears best resolved by a wage

³³Cornelissen, Dustmann and Schönberg (2017) show this formally; they also show that in a model with covariates, defining the peer group as too large can lead to attenuation bias as well.

band of \pm \$0.30 (or about 6% of the worker’s own initial wage).³⁴ Thus defined, peer groups average 9 peers per worker and include roughly 70% of coworkers who qualified for scheduled raises. In Section 5.6.2, we compare the estimates from wage-based definitions to those based on alternative, non-wage criteria for defining peers—including age and tenure similarity and distance between residential 5-digit ZIP codes. We show that definitions based on non-wage criteria generally lead to attenuated estimates that are similar to those obtained from excluding coworkers at random. In light of this evidence, we present our main results using our preferred definition of peers as those within \pm \$0.30 of the worker’s initial wage.

5.2 Identifying peer effects

To identify the effect of peer wage comparisons, we extend the RD model in equation (7) to include a discontinuous function of the reference wage w_p . For intuition, suppose that for each worker i in each year y , we randomly select a single peer j with initial wage w_{0jy} and scheduled wage, w_{jy} , and we extend the model as follows:

$$S_{iy}^m = \beta \times w_{iy} + \delta \times w_{jy} + f_y(w_{0iy}) + g_y(w_{0jy}) + X_{ijy}\Gamma + \lambda_{z(i)} + \epsilon_{ijy} \quad (8)$$

Equation (8) 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 (7), a flexible specification for f_y ensures that the 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 estimated 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 separation 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. To aggregate the peer effects, we model the separation 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

³⁴A formal justification for choosing the \$0.30 wage band, based on a cross-validation technique, is presented in Appendix E.

w_{jy} by replacing the sharp RD for w_{jy} in equation (8) with a fuzzy RD for $\bar{w}_{p(i,y)y}$.³⁵ The two-stage specification is as follows:

$$\begin{aligned} 1st: \quad & \bar{w}_{p(i,y)y} = \gamma \times w_{iy} + \gamma_p \times w_{jy} + \tilde{f}_y(w_{0iy}) + \tilde{g}_y(w_{0jy}) + X_{ijy}\tilde{\Gamma} + \tilde{\lambda}_{z(i)} + \tilde{\epsilon}_{ijy} \\ 2nd: \quad & S_{iy}^m = \beta \times w_{iy} + \beta_p \times \bar{w}_{p(i,y)y} + f_y(w_{0iy}) + g_y(w_{0jy}) + X_{ijy}\Gamma + \lambda_{z(i)} + \epsilon_{ijy} \end{aligned} \quad (9)$$

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 (7), β continues to represent the total effect of the own wage on separations. 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, we now have an estimate for β_p —the separation response to an increase in the average wage of one’s peers. And we use the latter to decompose the own-wage response (β) into the effect of an increase in relative pay ($= -\beta_p$) and the effect of a “gap-constant” pay increase that is uniform across peers ($= \beta + \beta_p$; see equation 4).

While equation (9) is estimable using a randomly selected peer for each worker, we 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 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 standard errors at the store level, we account for the impact that repeated observations may have on our statistical inference.

In our baseline specification for equation (9), we continue to include a set of ZIP-code dummies and a dummy for workers earning multiples of \$0.05. Now, we also include a 5-cent wage dummy for peer j and a dummy for the peer having received a merit raise.³⁶ Again, the heaping of wages at 5-cent multiples suggests a potential threat to internal validity; for example, excess heaping of peer wages may be associated with heterogeneity

³⁵Note that at the discontinuities, variation in $\bar{w}_{p(i,y)y}$ reflects variation in the fraction of peers who are just above (vs. just below) the threshold; hence $\bar{w}_{p(i,y)y}$ is a natural way to aggregate the peer wages. Our approach is similar to the more familiar linear-in-means model for peer effects (see also Section 5.3 and Appendix Table A8) but it involves less stringent parametric assumptions about the aggregability of the peer wage. Methodologically, our paper adds to the growing literature on aggregating across multiple discontinuities (e.g., Papay, Willett and Murnane 2011; Reardon and Robinson 2012).

³⁶Recall that while our estimation sample is restricted to workers with a history of at least one merit raise, we do not impose this restriction when defining a worker’s potential peers (except as a special case—see Section 5.6.2).

in store or market characteristics that predict separations. We document the robustness of our baseline results to a set of alternative specifications analogous to the specifications used for the RD model, and we again test for the smooth evolution of covariates at the thresholds by estimating models for predicted separations.

Finally, as an alternative to the two-stage MRD model estimated with worker-peer data, we also estimate a simpler model that uses worker-level data and controls for a function of the average initial peer wage, $\bar{w}_{0p(i,y)}$. This model aggregates the peer running variables, which requires that separations are a linear function of w_{0jy} . This “linear-in-means” model provides yet another check on our results.

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

Table 3 reports results from the two-stage MRD model.³⁷ Consistent with the exogeneity of the identifying wage variation, the coefficients on own wage are similar to those from the univariate RD models in Table 2 (where we do not control for peer wages), and we continue to find strong effects of own wages on separations. In the baseline specification (column 3) the estimates of β for all but 1-month separations range from -0.46 to -0.66 and all are statistically significant.

The estimates in Table 3 also show sizable effects of the peer average wage. The estimates for β_P are often quite similar in magnitude to the estimates of β , but are positive instead of negative—implying that workers respond to higher peer wages with a substantially higher probability of separating. In the baseline specification (column 3) the coefficients range from 0.29 (at one month) to 0.89 (at three months).

While the estimates vary somewhat across the other specifications, they generally imply sizable, positive effects of peer wages on separations. Column 4 shows estimates from a stacked model in which the running variables, r_{iy} and r_{jy} , are defined as the distance from the own or peer initial wage (w_{0iy} or w_{0jy}) to the nearest discontinuity threshold.³⁸ The estimates from this model are similar, though typically a little larger, than those in column 3. Returning to the global model, columns 5 and 6 show that

³⁷The first-stage and the reduced-form estimates for all of the specifications shown in Table 3 are reported in Appendix Table A6. As expected, the first-stage estimates suggest a strong relationship between the wage of a specific peer and the average peer wage, with coefficients of between 0.25 and 0.42 and t-statistics between 15 and 25 .

³⁸As before, the normalized running variable for own wage is $r_{iy} = w_{0iy} - T_{iy}^k$, where $T_{iy}^k = \arg \min_{T_y^\kappa} |T_y^\kappa - w_{0iy}|$ 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 $T_{jy}^k = \arg \min_{T_y^\kappa} |T_y^\kappa - w_{0jy}|$ is the nearest wage threshold to the peer’s wage, w_{0jy} .

the estimates are less stable when we use quadratic or cubic functions of the running variables; however, these specifications also produce relatively large standard errors. Importantly, the remaining columns show that the finding of sizable, positive peer-wage effects is robust to controlling for additional worker and peer covariates (column 7); to the inclusion store fixed effects (column 8); and to models that either exclude “donut-hole” observations around the peer-wage threshold (column 9) or use only peers with initial wages at \$0.05 multiples (column 10). Across these specifications, the 3-month estimates of β_P range from 0.66 to 0.82; most are statistically significant at the 5 percent level, and all are statistically significant at the 10 percent level.

Smooth evolution of covariates. In Appendix Table A7, we test for smooth evolution of predetermined covariates in the MRD specifications by replacing the dependent variable with predicted separations as we did in the RD model.³⁹ The global model that includes only linear controls for running variables (column 1) produces significant coefficients consistent with positive bias in the peer effects estimates (as well as in the own-wage estimates, similar to what we saw in Appendix Table A2). Inclusion of ZIP code fixed effects substantially reduces this bias (column 2), and evidence of discontinuities is mostly eliminated in column 3 where we control for initial own and peer wages that are multiples of \$0.05 and a dummy for peers that received a merit raise. The remaining columns of Appendix Table A7 corroborate the robustness of our alternative approaches for isolating exogenous variation in peer wages. The peer-wage coefficients from the stacked model (column 4) are all small and statistically insignificant, and the global specifications in columns 5-10 produce a mix of small positive and small negative coefficients that are again mostly close to zero and insignificant. In the model for 3-month separations, for example, the estimated effects on predicted separations in columns 3-10 range from -0.04 to 0.10. Even the positive estimates are an order of magnitude smaller than the estimated positive effects on observed separations.

Additional specifications and graphical representation. Appendix Table A8 provides additional tests of robustness for the peer-wage estimates. Focusing on models for 3-month separations, columns 1-5 show estimates from “linear in means” models estimated using worker-level data (instead of worker-peer pairs). The table shows coefficients on peer average wage, and the models control for functions of the peer average initial wage plus other aggregated control variables that differ across specifications. Column 2 in Panel B is the closest to our baseline model, and the estimate of 0.82 is very

³⁹Predicted separations are constructed from a model that includes all worker, peer and store-level controls included in columns 7 and 8 of Table 3 (see section 4.1 and Table 3 note for details).

similar to our baseline estimate of 0.89. The estimates are also fairly robust to the exclusion of direct controls for heaping at 5-cent wages (column 1) and to the inclusion of controls for worker characteristics and the average characteristics of peers and store-level coworkers (column 3).

The specifications in columns 6-12 of Appendix Table A8 show variants of the stacked MRD model in which the running variables measure own and peer distance from the nearest threshold. These include specifications that control for quadratic functions of the running variable, relax the restriction of equal slopes above and below the threshold, and/or exclude ZIP code fixed effects. The coefficients from these specifications all generally fall within the range of estimates from the 3-month separation models in Table 3, although allowing for different slopes results in some loss of precision. Finally, we show results from stacked MRD models that use narrower bandwidths for the peer’s distance from the threshold in columns 5-8 of Appendix Table A5. Reducing the bandwidth by \$0.01 or \$0.02 tends to increase the estimates only slightly (by 13% at most). But estimates from models with bandwidths of \$0.04 (or less) are uninformative due to a substantial loss of precision.

Figure 5 is graphical representation of the stacked model with baseline controls and symmetric slopes for the peer-wage effect on 3-month separations (as in column 4 of Table 3). As in Figure 4 (for the own-wage effect), the fitted values are normalized to equal zero at the left limit at the threshold, and the right limit at the threshold represents the effect of a \$0.10 discontinuity in the wage of a representative peer. Though the mean residuals are somewhat more scattered than in Figure 4, reflecting lower precision in the peer-wage estimates, there is a visible jump in separation rates at the threshold.

Timing and magnitude of effects. The estimates in Table 3 show that as we vary the time horizon for separating, both the own-wage and peer-wage coefficients display a similar pattern: both increase the most between 1 and 3 months. This pattern is confirmed in Figure 6, which shows the timing of the effects in more detail. Using 10-day intervals, the dark line plots the fraction of workers who remain employed at their store against time elapsed since the raises. Also shown are the mean survival rates at each interval, adjusted for the effect of a \$0.10 increase in either the own wage or the peer average wage. Both the own and peer-wage effects are close to zero in the first 30 days, increase quickly over the next two months, and are then relatively stable. The one-month lag in the response might partly reflect the time it took for workers to learn about

their coworkers' raises.⁴⁰ Interestingly, the pattern of stabilization after three months is inconsistent with market-driven separations in the framework of our model, where the number of outside offers continues to increase over time. But it can easily be reconciled with preferences over relative pay. For example, if preferences are heterogeneous across workers, this would cause the aggregate response to weaken, and the cumulative effect to stabilize, as those with the strongest preferences are selected out of the sample.⁴¹

Figure 6 also confirms that after the first month, the own-wage and peer-wage effects are consistently opposite in sign. And although the peer-wage effect is estimated with less precision and is more variable than the own-wage effect, the two effects are generally similar in magnitude. To interpret these magnitudes, we can use the information in Figure 6 to infer the effects of wage increases on the duration of employment. The median worker in our sample stays employed for slightly less than nine months following the raise. The rightward shift of the survival curve when adjusted for the own-wage response implies that all else equal, a raise of \$0.10 increases tenure by about one month—so that the median employee leaves after 10 months instead of 9. Similarly, an employee whose peers all receive an additional \$0.10 (the maximum jump at the thresholds) is predicted to leave roughly one month sooner than she otherwise would have. It is also useful to consider what the estimates imply about the compensating differential for a decline in relative pay. The fact that a \$0.10 own-wage increase roughly offsets the impact on separations of a \$0.10 peer-wage increase suggests that workers are willing to forego roughly \$0.10 (or 2% of pay) in order to avoid a \$0.10 (or 2%) reduction in relative pay.⁴²

Decomposition of own-wage effect. The offsetting effects illustrated in Figure 6 suggest that if own and peer wages rise by the same amount (e.g., by \$0.10), leaving relative pay unchanged, then the resulting change in separations is close to zero. Formal estimates of this “gap-constant” separation response are obtained by netting out the peer

⁴⁰Workers were paid every other Friday for the two-week period ending the previous Saturday; hence there was a 2-3 week lag each year between the effective date of the raise and first pay date that reflected the raise.

⁴¹As noted in section 4.1 (footnote 20), the 6- and 9-month estimates for the own-wage effect may reflect a small attenuation bias due to subsequent raises that, over time, dilute the impact of the scheduled raise discontinuities on own wages. 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 2-3 months. The impact of scheduled raises on *peer* wages may be diluted more quickly due to the turnover of one's peers. This could help explain why the peer-wage separation response appears to stabilize (and also become less precise) after around 3 months.

⁴²Depending on the time horizon (between 2 and 9 months), the implied compensating differential for a \$.10 increase in peer wage is between \$0.06 and \$0.16.

effect from the total own-wage effect (equation 4). In Table 4, we report the separation elasticities implied by our estimates of β and β_P from the baseline MRD model, and we decompose the total effect of an own-wage increase on separations into the effect of an increase in relative pay and the effect of a gap-constant wage increase. At two and three months, the own-wage and peer-wage elasticities are both large at around -13 and 20 respectively. Consistent with the relative stability of the coefficients after 2 months, both elasticities diminish in magnitude over time, but even at 9 months the total own-wage elasticity remains significant at -5.9.

The gap-constant elasticities, while imprecise (especially at 1-3 months), imply separation responses that are much closer to zero than those implied by the total own-wage elasticities. The gap-constant elasticity becomes more negative and also more precisely estimated as the separation window lengthens, but at 9 months it is still only -2.3 and the 95% confidence interval rules out elasticities more negative than -11.1. Following Manning (2003), we can multiply the estimated separation elasticity by -2 to infer that the firm faces a labor supply elasticity of around 4.6. This, in turn, implies that the firm can reduce wages by roughly 20% below the value marginal product of labor. If we focus instead on the lower bound of our 95% confidence interval, we can rule out potential wage mark-downs of less than 4%.

In short, our estimates suggest that the overall effect of own wages on separations is largely driven by peer wage comparisons, and that after accounting for this peer effect, the effect of market comparisons is relatively modest. In Section 7, we compare the magnitudes of our estimates to those found elsewhere in the literature and we discuss further implications of our findings.

5.4 Separations versus Quits

Our analysis thus far has examined separations for any reason—including quitting, returning to school, moving, transferring to another store, or being fired. While any one of these outcomes might be affected by wage changes, workers arguably have more discretion over the decision to quit than they do over other types of separations. Quit behavior is also more consistent with the behavior considered in our theoretical framework, where endogenous separations represent job-to-job transitions and the separation elasticity is informative about the extent of labor market frictions.

If the observed separation response to changes in own and peer wages were driven only by quits—and if separations for other reasons were strictly exogenous—then we might

improve precision without affecting the estimates by excluding “non-quit” separations from the estimation sample since this would reduce the error variance. On the other hand, to the extent that non-quit separations are also responsive to wages, censoring these observations may lead to discontinuities in sample attrition and might therefore bias the estimated quit response. Appendix Table A9 shows estimates from models in which the dependent variable is an indicator for quitting – with quits defined as voluntary separations that occur for job-related reasons (leaving for another job, dissatisfaction, or simply not showing up). For employees who separate for other reasons within the relevant window (roughly half of the sample), quit decisions are treated as censored. Compared to the estimates in Table 3 based on all separations, the estimates from the models of quits are sometimes smaller in magnitude; however the differences between the two sets of results are neither systematic nor large—suggesting that any bias due to endogenous competing risks is small.⁴³ Further, the estimates from the quits models are generally more precise—suggesting that censoring non-quit separations eliminates a fair amount of noise.

The gain in precision is evident in Appendix Figure A2, which shows a graphical representation of the stacked model in column 4 of Appendix Table A9. Panel A shows the relationship between quit rates and own initial wage (normalized as distance to the nearest threshold) and is analogous to Figure 4 (where the model is based on all separations). Similarly, Panel B shows how quit rates vary with distance from peer wage to the nearest threshold and is analogous to Figure 5. The improved fit is especially visible in Panel B, where the data now fit the linear model reasonably well and the plotted points show clear evidence of a discontinuity. Aside from the improved precision, the figures based on quits look similar to those based on all separations. In fact, the point estimates from the stacked models are very close (compare the own and peer-wage coefficients of -0.72 and 0.86 in Appendix Table A9 to the coefficients of -0.73 and 0.95 in Table 3). This evidence provides further support for our empirical model and also strongly suggests that the separation responses we estimate are driven mainly by voluntary job transitions or “quits.”

⁴³Appendix Table A10 presents formal tests for differential sample attrition at the thresholds; it shows estimates from MRD models for the probability of being in the quits analysis sample. These results are inconclusive: while most of the coefficients are not statistically significant, some of the peer-wage coefficients are sizable and negative, suggesting that the censoring of quit outcomes for employees with non-quit separations may induce some negative bias. Reassuringly, however, tests for covariate smoothness that use predicted quits as the outcome show no evidence of discontinuities in either the global or stacked versions of our baseline MRD model (see Appendix Table A11).

5.5 Additional validation: Split-sample RD

If the own-wage separation response is driven mainly by relative-pay concerns and the response to a “gap-constant” wage increase is relatively small, then separation discontinuities at the *own-wage* thresholds should be small among workers whose peers receive similar raises. To test this prediction, we now 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 the closest pay-step threshold;⁴⁴ 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 her peers are on the same side of a threshold and thus the average peer raise is similar to the worker’s raise.

Appendix Figure A3 presents visual evidence consistent with our prediction. The two panels show plots based on stacked RD models estimated separately for the “opposite side” and “same side” sub-samples. Aside from using different samples, these plots are constructed in the same manner as Figure 3 (see Section 4.1). In the first sample, where the worker and the majority of her peers are on opposite sides of a threshold, the discontinuity at the own-wage threshold is large, negative, and clearly significant (Panel A). In the second sample, where the majority of peers are on the same side of a threshold as the worker, there is no visible evidence of a discontinuity in separations (Panel B). Importantly, the discontinuity in the *own wage* does not differ across samples—it is always \$0.10. Hence, the only difference in treatment between the two samples is the change in *relative pay* that is associated with the change in own wage. The differential separation response therefore supports our interpretation of the MRD results—that the effect of wages on separation behavior operates mainly through the effect on relative pay.

We present a more formal “split-sample” analysis in Table 5. The top row shows estimates from the baseline MRD model for the effect of an increase in the own wage on the wage gap, $\frac{dw_g}{dw}$, in each sample. These estimates confirm that the discontinuity in the wage gap is much larger in the “opposite-side” sample (1.46) than in the “same-side” sample (0.69). The table also shows the estimated own-wage separation responses for each window and for each sample. Consistent with the graphical evidence, the coefficients are all sizable and statistically significant in the first sample and are smaller and statistically insignificant in the second.

⁴⁴Either worker i is above the closest 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 ($\sum_{j \in p(i,y)} \mathbf{1}\{w_{0jy} \geq T_{jy}^k\} \geq \sum_{j \in p(i,y)} \mathbf{1}\{w_{0jy} < T_{jy}^k\}$).

We also use these split-sample RD estimates to construct alternative estimates for the peer effect represented by β_p in our MRD model. If we let $\left[\frac{dw_g}{dw}\right]_A$ denote the effect of an increase in the own wage on the wage gap in a given sample A , then the total effect of the wage increase on separations can be written as: $\left[\frac{dS}{dw}\right]_A = \frac{\partial S(w, w_g)}{\partial w} + \frac{\partial S(w, w_g)}{\partial w_g} \left[\frac{dw_g}{dw}\right]_A$, where the second term is the effect of relative pay comparisons. Since $\frac{\partial S(w, w_g)}{\partial w_g} = -\frac{\partial S(w, w_p)}{\partial w_p}$, we can also write the total separation response as: $\left[\frac{dS}{dw}\right]_A = \frac{\partial S(w, w_g)}{\partial w} - \frac{\partial S(w, w_p)}{\partial w_p} \left[\frac{dw_g}{dw}\right]_A$. Consider two samples A and B where $\left[\frac{dw_g}{dw}\right]_A \neq \left[\frac{dw_g}{dw}\right]_B$ (as is true of our “opposite-side” and “same-side” sub-samples). Under the additional assumption that both the gap-constant effect, $\frac{\partial S(w, w_g)}{\partial w}$, and the peer-wage effect, $\frac{\partial S(w, w_p)}{\partial w_p}$, are constant across the two samples, we can construct a Wald estimator for the peer effect as follows:

$$\frac{\partial S(w, w_p)}{\partial w_p} = \frac{\left[\frac{dS}{dw}\right]_B - \left[\frac{dS}{dw}\right]_A}{\left[\frac{dw_g}{dw}\right]_A - \left[\frac{dw_g}{dw}\right]_B}.$$

The Wald estimates of the peer effect, reported in the final column of Table 5, range from 0.05 (9 months) to 1.39 (3 months). These estimates are less precise and more variable than the estimates of β_p from the baseline MRD model (Table 3, column 3); however, the 2 and 3-month estimates are both large and statistically significant. On the whole, the Wald estimates, like the MRD estimates, show a pattern of positive separation responses to changes in relative pay that are largest in the 2-3 months following the raise.

We stress that the two estimators for the peer-wage effect are distinct in that they are identified off of different parts of the joint distribution of own and peer wages. While the MRD uses worker-peer pairs that are near a peer-wage threshold, the split-sample RD uses pairs near an *own*-wage threshold (see Online Appendix D for a detailed explanation). For causal inference, the MRD is preferred because it is identified off of the discontinuities in peer wage, while the Wald estimator based on the split-sample RD requires the extra assumption that peer wages are as good as random near the *own*-wage threshold. The split-sample approach is nevertheless appealing because it provides a simple and transparent validation of our conclusions.

5.6 Varying definitions of peer group

5.6.1 Wage-based definitions and choice of wage band

As discussed in Section 5.1, a definition of peers that is too narrow will lead to attenuation of the estimated peer effect. In contrast, too broad a definition will tend to make the estimate imprecise by reducing identifying variation. Figure 7 illustrates the

bias-variance trade-off when defining peers as coworkers with similar initial wages. It presents the estimated peer effects and 95 percent confidence intervals from models of 3-month separations, for peer groups based on all “potential peers” (coworkers in the same job and store who qualify for a scheduled raise), and on subsets of these potential peers who initially earned within \$0.10, \$0.20, \$0.30,...,\$0.80 of the worker’s own initial wage. (The estimates for all separation windows are reported in Appendix Table A12.) Consistent with attenuation bias from the exclusion of relevant peers, estimates based on narrow wage bands are relatively small. In particular, the \$.10 wage bands produce little evidence of peer effects. The estimates generally increase with the size of the wage band, especially as it expands from \$0.10 to \$0.30. Beyond that point, they are relatively stable; in fact, the estimates based on the store-wide definition of peers (0.96) is only slightly larger than the estimate of 0.89 based on peers within \$0.30. However, the standard errors tend to increase with the size of the band. For example, the standard error is 0.55 for the store-wide estimate but only 0.29 for the \$0.30 estimate. Hence, while estimates based on wider wage bands are less likely to be biased due to the exclusion of relevant peers, they are also less precise.

In short, the visual evidence in Figure 7 suggests that the \$0.30 wage band does a good job of resolving the bias-variance trade-off. Appendix E presents a more formal justification for using the \$0.30 wage band to define peers. Taking the store-wide estimate (based on all potential peers) as a benchmark, we show that a cross-validation procedure selects the \$0.30 wage band as the one that minimizes the (out-of-sample) mean squared error when predicting the store-wide estimate.

5.6.2 Other peer definitions

Criteria other than proximity of initial wages may also be useful for defining peers. For example, coworkers may be more likely to interact in the workplace if they are similar in tenure or age or if they live near each other. Importantly, however, even if all coworkers in a given subgroup are relevant peers, an estimate based on the subgroup will be biased toward zero if it excludes relevant peers, and the bias will be greater the more relevant peers it excludes. To evaluate how well various criteria identify relevant peers, we must therefore take into account the size of the peer group identified by those criteria.

To do this, we compare the estimates based on each alternative peer definition to a benchmark based on a peer group of the same size that is selected at random from all potential peers. In Figure 8 the locus traces out estimated peer-wage effects on 3-month

separations based on randomly selected peer groups of different sizes. The figure also plots estimates based on alternative peer group definitions against the share of potential peers who are included in each definition.⁴⁵ If a peer definition is better at keeping relevant peers than a randomly selected subset of equal size, the estimate will tend to lie above the “random-peer” locus. If the definition loses more relevant peers than a randomly selected subset of equal size, it will tend to lie below the locus. Appendix Table A13 reports 1, 2, 3, 6 and 9 month separation responses for the peer definitions in Figure 8 and several others.

To start, we note that our baseline definition using initial wages within \$0.30 includes around 72 percent of potential peers, but has an estimate substantially above the random-peer locus. This provides further validation of our baseline definition, and suggest that the coworkers excluded from this definition are disproportionately irrelevant. Next, we consider peers defined as the 5, 10, 15 or 20 closest coworkers in terms of initial wages. The estimates for the closest 5 or 10 workers tend to be smaller than our baseline, while the point estimates based on the closest 15 or 20 workers are quite similar in magnitude but are less precise than our baseline estimate (Appendix Table A13). These results therefore suggest that a definition based on all peers within a fixed wage band tends to perform better than a definition based on a fixed number of nearby peers.

Turning to non-wage definitions, we consider peer groups based on age, tenure, past merit raise status and geographic proximity of residences.⁴⁶ These definitions reduce the share of potential peers substantially, retaining between 42 and 67 percent of them depending on the definition. The point estimates based on similar merit status (for which the peer group includes all potential peers who received a merit raise in the past) is marginally statistically significant and somewhat higher than would be expected from dropping an equivalent share of peers at random, but it is still smaller and less precise than our baseline estimate. The estimates using similarity in age or tenure or geographic proximity to define peers are either similar to or smaller than their random-

⁴⁵For each share $s \in \{10, 20, \dots, 90\}$ of potential peers, we define peers as a randomly selected s percent of all eligible coworkers and we estimate the peer-wage effect in this sub-sample using the baseline MRD specification. For each s we repeat this 500 times and report the average estimate. The rightmost point with 100 percent of potential peers is the full-store estimate of 0.96.

⁴⁶The definitions are as follows: age difference between the worker and the peer was under 5 years; both the worker and the peer were employed for more than 8 months or both less; both the worker and the peer received a merit raise; and the distance between the worker’s and peer’s residential 5-digit ZIP code was under 9 miles. In Appendix Table A13 we also report 6 and 12 mile cutoffs.

peer counterparts. And again, the standard errors are all larger than those from our baseline model.

In sum, the smaller size and relative imprecision of the estimates in Appendix Table A13 suggests that these alternative non-wage peer definitions exclude large shares of relevant peers. Moreover, in all cases but one (merit raise status), the estimates are no larger than what would be expected from randomly dropping the share of coworkers excluded by these definitions. Overall, these findings provide little support for peer definitions based on non-wage criteria.

5.7 Heterogeneous peer effects by initial wage gap

So far, we have estimated the average effect of peer wages on separations. But theory suggests that if workers are especially concerned about being treated fairly, then the separation response may be asymmetric—and stronger when workers are paid less than their peers (Card et al. 2012; Fehr and Schmidt 1999). In our setting, since the wage schedule is weakly monotonic (and so raises are weakly rank preserving), larger peer raises tend to reduce existing wage gaps vis-à-vis lower-paid peers (without ever leading to rank-reversals), but to widen existing gaps with peers who were initially paid more. Fairness concerns might therefore lead workers to be especially sensitive to raises received by their higher-paid peers. To test for such patterns in our data, we allow the separation response to differ depending on the initial wage gap between the worker and the peer. Using our baseline MRD model, we estimate separate models for six groups of worker-peer pairs defined based on the *ex ante* wage gap between own and peer wage, $\Delta w_{0ip(i)y} = w_{0iy} - w_{0p(i)y} : < -\$0.30, [-0.30, -0.16], [-0.15, -0.01], [0.01, 0.15], [0.16, 0.30], > 0.30$. (Note that the four interior groups comprise our baseline definition of peers making within \$0.30 of the worker, while the other two include peers who fall outside of this definition.)⁴⁷ We then compare the causal effect of peer wages on separation rates across these six groups.

Figure 9 plots the 3-month separation response for each of the six groups, along with the group’s share of all potential peers. We find a clear asymmetry in the response: the positive separation response to a peer-wage increase occurs when the peers in question initially earned more than the worker (and the initial gap is negative). The 3-month separation estimates for the $[-0.30, -0.16]$ and $[-0.15, -0.01]$ groups are 0.95 and 0.70, respectively, and both are statistically significant at the 5 percent level (see Appendix

⁴⁷We exclude cases where $\Delta w_{0ip(i)y} = 0$ because there is no independent variation in $w_{0p(i)y}$ among these observations.

Table A14 for the estimates). The estimate for the leftmost group where the peers earned > 30 cents more than the worker is also sizable, but imprecise. In contrast, for the three groups where the peers were initially earning less than the worker, the separation responses are small in magnitude and statistically indistinguishable from zero. If we compare the $[-0.30, -0.16]$ group with the $[0.16, 0.30]$ group, the point estimate for the former lies outside of the 95 percent confidence interval for the latter. Similarly, the estimate for the $[-0.15, -0.01]$ group lies outside of the confidence interval for the $[0.01, 0.15]$ group. It is important to note that the effects are much larger in the groups with negative initial wage gaps even though these groups have slightly smaller numbers of peers than their counterpart groups with a positive wage gap.

Appendix Table A14 shows the estimates from models of 1, 2, 3, 6 and 9-month separations; there is a similar pattern of asymmetry across all separation windows. The table also reports estimates using more aggregated categories. If we simply classify peers as initially earning either more or less than the worker, we find that the 3-month separation estimate is 1.39 and statistically significant for the former, but only 0.16 and not significant for the latter. The pattern is similar if we restrict our sample of peers to only those whose initial wages are within \$0.30 (our baseline definition). Here the 3-month separation estimate is 0.94 when the initial peer wage is above the worker's wage, but only 0.15 when it is below. For both the unrestricted and within-\$0.30 peer definitions, the 2, 3, and 6-month responses are statistically significant at least at the 10 percent level when we consider peers initially earning more than the worker, but close to zero otherwise.

In sum, we find that a worker's likelihood of separating increases sharply in response to larger peer raises only when the peers getting the raises were initially earning more than the worker, in which cases larger peer raises tend to widen existing gaps. In contrast, worker separations do not appear sensitive to the size of raises received by peers who were initially earning less. These results suggest that workers are averse to unequal treatment that is to their disadvantage—which we interpret as evidence of concerns about fairness.

5.8 Do peer effects reflect learning about future wages?

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 unexpected

increase in peer wages causes workers to revise their expected future wage upwards and thus makes them less likely to separate. 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 general, it is also possible that a higher peer raise (and a reduction in relative pay) sends a negative signal about one’s future wage. For example, it might suggest employer discrimination or signal a mismatch between the worker and the firm. *A priori*, these explanations also seem doubtful because the variation in peer raises is based on arbitrary cutoffs and should not predict future raises at all.

In Appendix F, we confirm that peer wages lack predictive power - and thus have no signaling value - for own future wage growth. Appendix Table F1 presents MRD models for future wage growth, defined as the change in log own wage from the day after the raise through the next merit raise cycle 10 months later. Among workers who remain in the sample at least 10 months, the peer-wage coefficient is small and positive (0.074) and not statistically significant; hence, there is no indication of a decline in observed future raise at the peer-wage discontinuity. Since fewer than half of the workers in our sample are still at the firm after 10 months we also estimate bounds for the peer-wage effect that account for non-random attrition.⁴⁸ The lower bound estimate is -0.063 and not statistically different from zero, implying that if all of one’s peers received a \$0.10 larger raise, the worker’s future raise would be at most 0.6 percentage points smaller. The upper bound is quite large at 0.48; however, it is not meaningful in our context since it cannot explain why workers whose peers got bigger raises were more likely to separate.

Setting aside the fact that it is not statistically distinguishable from zero, can the lower bound estimate of a 0.6 percentage point reduction in one’s own future raise explain our estimated separation response to a larger peer raise? If we take the 3-month separation response to own wage of -0.55 (Table 3, column 3) and multiply by the implied future wage increase of -\$0.32 ($= -0.063 \times \5.15) at the peer-wage discontinuity, we get an implied separations response of 0.18 – which is only $1/5$ the size of the actual separations response of 0.89. Therefore, even under the most pessimistic assumption about selective attrition, rational updating about future wages cannot plausibly explain the large separation response to an increase in peer wages.

The fact that peer raises do not predict own future raises not only implies that the arbitrary variation in peer wages has no signaling value; it also rules out other potential

⁴⁸Similar to Dong (2017) and Kim (2016), we modify Lee’s (2009) procedure for application to an RD context. See Appendix F for details.

explanations that rely on rationally updated expectations. For example, higher peer raises could in principle make peers more likely to remain with the firm and could thereby increase competition for promotions; however, our analysis suggests that any such effect is too small to explain the peer-wage effects on separations. The high rate of turnover and low promotion rates in the jobs we study also suggest that our findings are unlikely to be explained by career concerns.⁴⁹

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. Again, this seems unlikely because peer raises are based on arbitrary pay-step thresholds, and the true distribution of wage offers from the market is unlikely to change at these arbitrary cutoffs. Hence, in our context, peer wages cannot provide a basis for rational learning about the market. However, we cannot rule out updating based on non-rational beliefs, including mistaken inferences about the market.⁵⁰

6 Falsification Tests

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 separate, then this correlation should cause discontinuities in separation 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. Here we focus on employees who received a merit raise during the week that annual merit raises were given, June 30 - July 7, 1996; who were employed on July 8, 1996; and whose wages on July 8 would have made them eligible for a scheduled raise on October 1 if they were still employed. For this pre-treatment placebo sample, we use our baseline RD and MRD specifications to model separations within 1, 2, and 3 months of July 8 as a function of the scheduled wage that would take effect on October 1. Because this sample is drawn from only one date, it is smaller than our main estimation sample, which consists of two

⁴⁹Career-related explanations are also inconsistent with the fact that in-store peers are not the main source of competition for managerial jobs; rather, the majority of new managers in our data are transferred from other stores or hired from outside the firm.

⁵⁰Similarly, we cannot rule out non-rational beliefs that make expected future wages discontinuous in the *own* wage. For example, if in stores where everyone got arbitrarily small raises workers believed they would be compensated during the next merit raise cycle, then such beliefs might help explain the small gap-constant separation response. However, own-wage RD models for future merit raises, estimated for stayers, show no evidence of discontinuities in the size of the next merit raise.

such increases (October 1, 1996 and September 1, 1997). So for comparison, we also re-estimate our main models using the subsample of employees who were present on the day of the first minimum wage increase (October 1, 1996).

Table 6 shows these results. In Panel A, we continue to find significant negative effects of own wages and significant positive effects of peer wages even when we restrict attention to the October 1, 1996 subsample. In contrast, we find no evidence of a pre-treatment effect of either own or peer wages in the placebo sample. In Panel B, the placebo estimates of own-wage effects are generally small and not statistically different from zero, while the peer-wage estimates are all negative and statistically insignificant. Appendix Figure A4 provides visual confirmation of these results, focusing on the 3-month separation rates. Panels A and C reproduce the “stacked” RD and MRD plots from Figures 4 and 5 using the October 1, 1996 subsample. As before, these plots show clear discontinuities in separation rates at the own and peer-wage thresholds. The analogous plots for the placebo sample in Panels B and D show no evidence of discontinuities at the threshold wages in the months before the raises were implemented, which lends further support to the causal interpretation of our estimates.

6.2 Tests for discontinuities at non-threshold wages

Our second test looks for discontinuities in separation 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 since bunching at these values raises the concern that workers whose own wages are multiples of \$0.05 have different latent separation propensities than those who do not. By testing directly for discontinuities at these non-threshold values, we provide yet another test for heaping-induced bias.

For the univariate RD, we implement this test by replacing the initial wage w_{0iy} (which enters equation 7 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 (9) and similarly we replace each peer wage w_{0jy} with a shifted peer wage w_{0jy}^s . In Appendix Figure A5, Panel A plots the t-statistics for the estimated separation 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. The discontinuity at $s=0.01$ 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. But importantly, we do not see discontinuities elsewhere in the wage distribution, including at $w_{0y} + \$0.05$ or $w_{0y} - \$0.05$. Panel B shows t-statistics for the peer-wage estimates from the wage-shifted versions of the two-stage MRD model. Here we see significant estimates at $s = 0$ and also at $s \in \{.01, .02, .03, .04\}$, the latter being driven by the fact that these bins are sparsely populated compared to $s = 0$. But 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 separation propensity associated with either own wages or peer wages at \$.05 multiples.

7 Discussion

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

Our findings complement the existing evidence on both the wage-setting power of employers and the effects of relative pay in the workplace. Despite differences across settings, our estimates are broadly similar in magnitude to prior estimates based on related treatments and outcomes. In the literature on labor market monopsony, for example, our estimate of -2.3 for the 9-month “gap-constant” separation elasticity is similar to the estimates of two studies that use exogenous wage variation for teachers. These studies find one-year separation elasticities of -1.8 (Ransom and Sims 2010) and -3.5 (Falch 2011), and their estimates imply potential wage markdowns of 14%-27%. Our findings suggest that similar wage-setting power exists in a retail labor market marked by low wages and high turnover. Our point estimate implies that the firm can reduce wages by roughly 20%, and we can reject potential wage mark-downs of less than 4% at the 95% confidence level.

On relative pay, our findings are most easily compared to those of Breza, Kaur and Shamdasani (2017), who randomize Indian manufacturing workers to pay units where pay raises resulted in either compressed or unequal wage structures. In unequal workplaces, the wage gap was roughly 7.4% between the highest and lowest-paid worker, or 5.8% between the lowest-paid worker and the average of the other two. During the one-month experiment, absenteeism responded very strongly to the presence of the pay

gap; the study estimates an elasticity of around 26 and a willingness to forego 7.1% - 9.3% of earnings to avoid the unequal workplace.⁵¹ Our results are remarkably similar. We estimate a peer-wage elasticity of 20.4 in the second month after treatment, and like Breza, Kaur and Shamdasani (2017), we find a compensating differential for unequal pay that is roughly equal to the size of the pay gap (6-7% in their case and about 2% in ours).

It is also useful to compare our results to those of Card et al. (2012), whose outcomes include measures of job satisfaction and job search. They find that for workers earning below the median in their pay unit and occupation, learning about coworker salaries leads to a 40% - 144% increase in the probability of being both “dissatisfied” and “very likely” to search for a new job shortly after the treatment.⁵² Since the treatment in Card et al. (2012) involves information about peer wages and not a change in those wages *per se*, we cannot quantify how sensitive workers in their setting are to pay inequality without knowing the extent of “surprise” associated with the treatment. However, under the plausible assumption that the information treatment reduced the posterior mean relative pay by 5-10% for those workers, the elasticities implied by Card et al. (2012)’s estimates would lie between 4 and 29—broadly similar to the range of peer-wage separation elasticities we find (see Table 4).⁵³

Other patterns in our results bear qualitative similarities to those found in the recent literature. Notably, the parallels between our findings and those of Breza, Kaur and Shamdasani (2017) go beyond the large responses to relative pay; like us, they also find relatively small effects of a common wage increase. Also, our finding of asymmetry in

⁵¹We use the estimates in Breza, Kaur and Shamdasani (2017) Table 4 to compute the elasticity for the lowest paid worker with wage W_L as: $\frac{\Delta \text{Absentee Rate}}{\Delta(\bar{W}_P - W_L)} \frac{W_L}{\text{Absentee Rate}}$, where Δ measures the difference between treatment and control units. The percentage point difference in the absentee rate, $\Delta \text{Absentee Rate}$, is equal to -0.117 (their Table 4), the baseline $\bar{\text{Absentee Rate}}$ is 0.061, and the wage gap as a share of base wage, $\frac{\Delta(\bar{W}_P - W_L)}{W_L}$, is 0.058. The authors calculate a willingness to pay between 7.1% and 9.3% of earnings to avoid the unequal workplace.

⁵²The range of estimates reflects different assumptions regarding heterogeneity in the “first stage” effect of the information treatment on workers’ use of a website to learn their peer wages (see Card et al. 2012, pp. 2994-2995).

⁵³There are other differences between our setting and Card et al. (2012) that affect the comparability of our estimates. First, since the Card et al. (2012) estimates are based on intent to search rather than actual turnover, the effects on actual turnover in their setting are likely to be smaller than the effects we find. Second, the wage differences we study were clearly arbitrary, but in their setting even unexpected differences in pay may have been viewed as at least partly justified by productivity differences, which could have moderated the effects.

the peer-wage response adds to a growing body of evidence that workers are averse to disadvantageous inequity (Breza et al. 2017; Card et al. 2012; Cohn et al. 2014).⁵⁴

While our conclusions are broadly consistent with prior research, there are limits to their generalizability, and the sensitivity of quit behavior to seemingly small pay gaps (at most \$.10 or 2% of the typical wage) may have been mediated by context-specific factors. In particular, given the arbitrary location of the thresholds, it is likely that the inequities were perceived as unjustified. Recent experiments have found that arbitrary pay gaps elicit stronger responses than gaps that can either be justified by productivity differences (Breza, Kaur and Shamdasani 2017) or at least rationalized in some way (Bracha, Gneezy and Loewenstein 2015). We also caution against the linear extrapolation of our estimates to infer how workers might respond to much larger wage gaps since the response function is likely to be nonlinear. For example, Card et al. (2012) find that rank appears to matter more for job satisfaction than the size of the pay gap—suggesting a diminishing marginal effect of inequity. One mechanism that could create such nonlinearities is a fixed cost to job search. As we show in Online Appendix B, if we endogenize search intensity in our model and allow for a fixed search cost (so that wage disparities may generate separations in part by provoking search), then small relative wage changes may produce large increases in separations, but further changes may have smaller marginal effects.

With those caveats in mind, it is useful to consider how peer comparisons may affect wage-setting practices when wages are not fully dictated by the market. Absent internal constraints, monopsony power creates an incentive for firms to suppress the wages of workers whose labor supply to the firm is the least elastic. But relative pay concerns like those we identify will raise the cost of wage discrimination among close peers, and in turn could lead firms to pursue strategies that redefine firm boundaries and produce a fissured workplace (Weil 2014). More generally, if firms are motivated to circumvent internal equity constraints, this could help explain why rising wage inequality has been accompanied by greater sorting of workers into high or low-wage firms (Song et al. 2015), and by a rise in domestic outsourcing and other types of employment restructuring (Katz and Krueger 2016) that often result in lower wages for outsourced workers (Dube and Kaplan 2010, Goldschmidt and Schmieder 2017). Evidence linking

⁵⁴Asymmetric responses to relative pay have also been found in some lab experiments. Bracha, Gneezy and Loewenstein (2015) show that unequal pay can reduce the effort and labor supply of lower-paid workers with no effect on higher-paid workers. On the other hand, both Charness and Kuhn (2007) and Goerg, 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.

pay equity concerns to specific wage policies and employment relations is an important area for future research.

References

- Angrist, Joshua D, Erich Battistin, and Daniela Vuri. 2016. "In a Small Moment: Class Size and Moral Hazard in the Italian Mezzogiorno," *American Economic Journal: Applied Economics*.
- 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.
- Autor, David H, Alan Manning, and Christopher L Smith. 2016. "The contribution of the minimum wage to US wage inequality over three decades: a reassessment," *American Economic Journal: Applied Economics*, 8(1): 58–99.
- Barreca, Alan I, Jason M Lindo, and Glen R Waddell. 2016. "Heaping-Induced Bias in Regression-Discontinuity Designs," *Economic Inquiry*, 54(1): 268–293.
- Bracha, Anat, Uri Gneezy, and George Loewenstein. 2015. "Relative pay and labor supply," *Journal of Labor Economics*, 33(2): 297–315.
- Breza, Emily, Supreet Kaur, and Yogita Shamdasani. 2017. "The morale effects of pay inequality," *The Quarterly Journal of Economics*.
- Buntrock, Evan. 2014. "Comparison-Group Position and Job Search: Information or Irritation?", Mimeo.
- Burdett, Kenneth and Dale T Mortensen. 1998. "Wage differentials, employer size, and unemployment," *International Economic Review*: 257–273.
- Card, David and Alan Krueger. 2016. *Myth and measurement: The New Economics of the Minimum Wage, Twentieth Anniversary Edition*: Princeton, NJ: Princeton University Press.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at work: The effect of peer salaries on job satisfaction," *The American Economic Review*, 102(6): 2981–3003.
- Charness, Gary and Peter Kuhn. 2007. "Does pay inequality affect worker effort? Experimental evidence," *Journal of Labor Economics*, 25(4): 693–723.

- Clark, Andrew E, Nicolai Kristensen, and Niels Westergård-Nielsen. 2009. “Job satisfaction and co-worker wages: Status or signal?” *The Economic Journal*, 119(536): 430–447.
- Clark, Andrew E and Andrew J Oswald. 1996. “Satisfaction and comparison income,” *Journal of Public Economics*, 61(3): 359–381.
- 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*, 12(4): 877–898.
- Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2017. “Peer effects in the workplace,” *The American Economic Review*, 107(2): 425–456.
- Depew, Briggs and Todd A Sørensen. 2013. “The elasticity of labor supply to the firm over the business cycle,” *Labour Economics*, 24: 196–204.
- Dong, Yingying. 2017. “Regression discontinuity designs with sample selection,” *Journal of Business & Economic Statistics*(just-accepted).
- Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. 2018. “Monopsony in Online Labor Markets,” NBER Working Paper No. 24416.
- Dube, Arindrajit and Ethan Kaplan. 2010. “Does outsourcing reduce wages in the low-wage service occupations? Evidence from janitors and guards,” *ILR Review*, 63(2): 287–306.
- Dube, Arindrajit, T William Lester, and Michael Reich. 2016. “Minimum wage shocks, employment flows, and labor market frictions,” *Journal of Labor Economics*, 34(3): 663–704.
- Dube, Arindrajit, Suresh Naidu, and Alan Manning. 2017. “Monopsony and Employer Mis-optimization Account for Round Number Bunching in the Wage Distribution,” Unpublished manuscript.
- Falch, Torberg. 2011. “Teacher mobility responses to wage changes: Evidence from a quasi-natural experiment,” *The American Economic Review*, 101(3): 460–465.
- Fehr, Ernst and Klaus M Schmidt. 1999. “A theory of fairness, competition, and cooperation,” *The Quarterly Journal of Economics*, 114(3): 817–868.



- Galizzi, Monica and Kevin Lang. 1998. “Relative wages, wage growth, and quit behavior,” *Journal of Labor Economics*, 16(2): 367–390.
- Gelman, Andrew and Guido Imbens. 2016. “Why High-order Polynomials Should Not Be Used in Regression Discontinuity Designs,” NBER Working Paper No. 20405.
- Giuliano, Laura. 2013. “Minimum wage effects on employment, substitution, and the teenage labor supply: Evidence from personnel data,” *Journal of Labor Economics*, 31(1): 155–194.
- 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.
- Goldschmidt, Deborah and Johannes F Schmieder. 2017. “The rise of domestic outsourcing and the evolution of the German wage structure,” *The Quarterly Journal of Economics*, 132(3): 1165–1217.
- Hirsch, Boris, Thorsten Schank, and Claus Schnabel. 2010. “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.
- Katz, Lawrence F and Alan B Krueger. 2016. “The Rise and Nature of Alternative Work Arrangements in the United States,” NBER Working Paper No. 22667.
- Kim, Bo Min. 2016. “Do Developmental Mathematics Develop Mathematics Proficiency? Bounding their Effectiveness in RDD with the Presence of Dropouts.”
- Lee, David S. 2009. “Training, wages, and sample selection: Estimating sharp bounds on treatment effects,” *The Review of Economic Studies*, 76(3): 1071–1102.
- Lee, David S and Thomas Lemieux. 2010. “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 48(2), p. 281.
- Manning, Alan. 2003. *Monopsony in motion: Imperfect competition in labor markets*: Princeton University Press.
- Mas, Alexandre. 2017. “Does transparency lead to pay compression?” *Journal of Political Economy*, 125(5): 000–000.

- Matsudaira, Jordan D. 2014. "Monopsony in the low-wage labor market? Evidence from minimum nurse staffing regulations," *Review of Economics and Statistics*, 96(1): 92–102.
- McCrary, Justin. 2008. "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 142(2): 698–714.
- Naidu, Suresh, Yaw Nyarko, and Shing-Yi Wang. 2016. "Monopsony power in migrant labor markets: evidence from the United Arab Emirates," *Journal of Political Economy*, 124(6): 1735–1792.
- 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.
- Pfeifer, Christian and Stefan Schneck. 2012. "Relative wage positions and quit behavior: Evidence from linked employer-employee data," *ILR Review*, 65(1): 126–147.
- 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.
- Ransom, Michael R and David P Sims. 2010. "Estimating the firm's labor supply curve in a "new monopsony" framework: Schoolteachers in Missouri," *Journal of Labor Economics*, 28(2): 331–355.
- 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.
- Rege, Mari and I Solli. 2013. "Lagging behind the joneses: The impact of relative earnings on job separation," Mimeo.
- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. 2015. "Firming up inequality," NBER Working Paper No. 21199.
- Staiger, Douglas O, Joanne Spetz, and Ciaran S Phibbs. 2010. "Is there monopsony in the labor market? Evidence from a natural experiment," *Journal of Labor Economics*, 28(2): 211–236.

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


Weil, David. 2014. *The Fissured Workplace: Why Work Became So Bad for So Many and What Can Be Done to Improve It*. Harvard University Press.

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



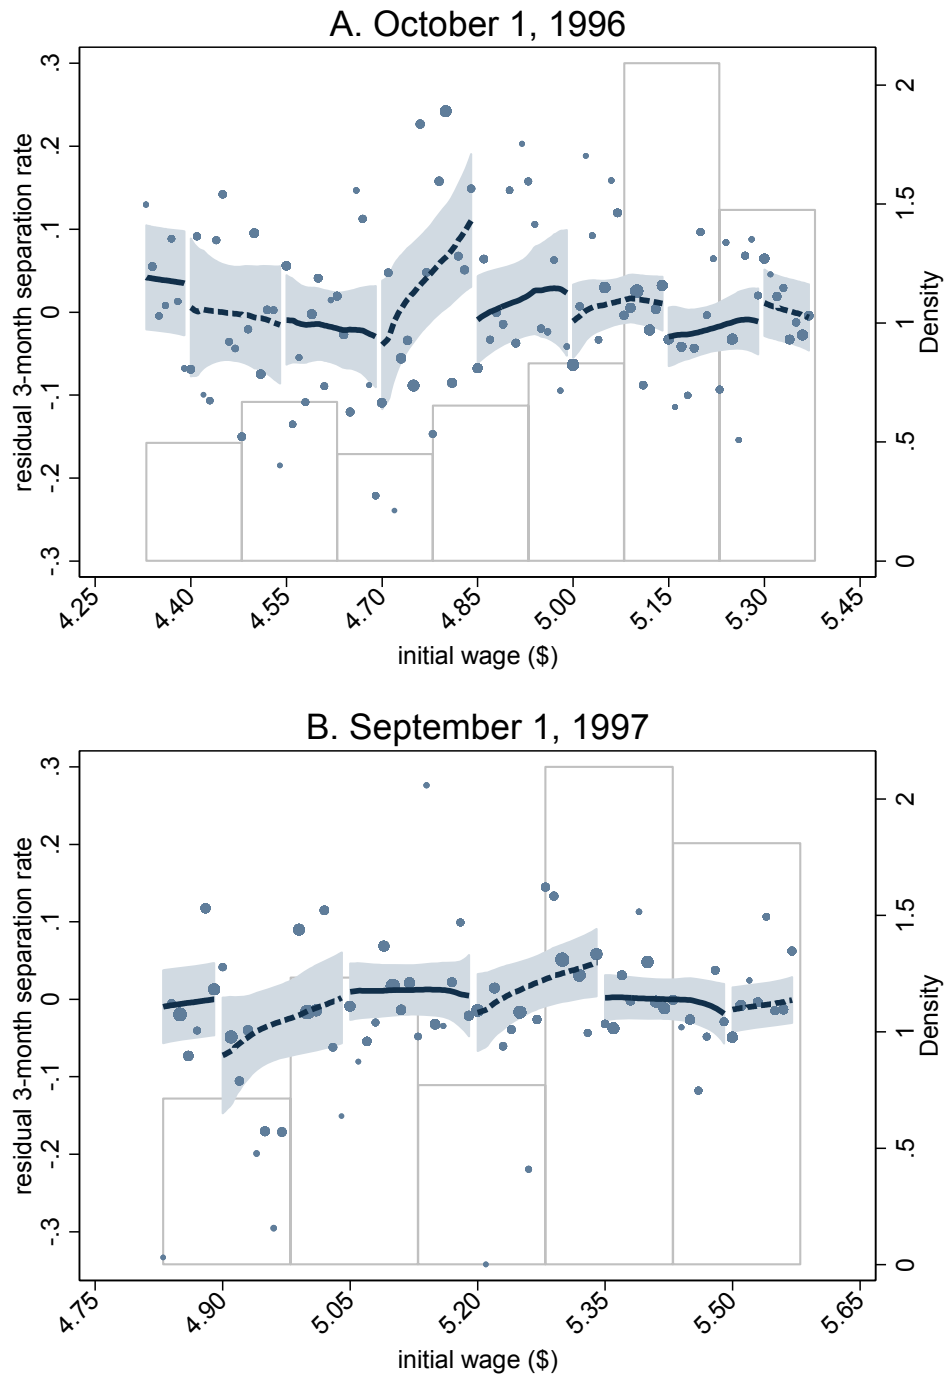
Note: Sample includes all hourly employees who were present on the day of the federal minimum wage increase on October 1, 1996 (Panel A) or September 1, 1997 (Panel B), who had at least one month of tenure, and whose “before” wage was no more than \$1.00 above the new minimum. Sample size is 15,557 (21,274)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

Figure 2. Raise by initial wage on day of 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. Three-month separation rate by initial wage on day of minimum wage increase




Note: Plotted points are mean values of residualized 3-month separation rates from a model that controls for ZIP-code based fixed effects (see text for details). Lines are local linear smoothers (bandwidth=.04) estimated within each 15-cent interval corresponding to a new pay step. Marker size is scaled by the proportion of observations at each value within each pay step. The figure also shows the overall density of initial wages in each year. Shaded areas indicate 95% confidence intervals.

Figure 4. Three-month separation rate by distance from own initial wage to nearest pay-step threshold



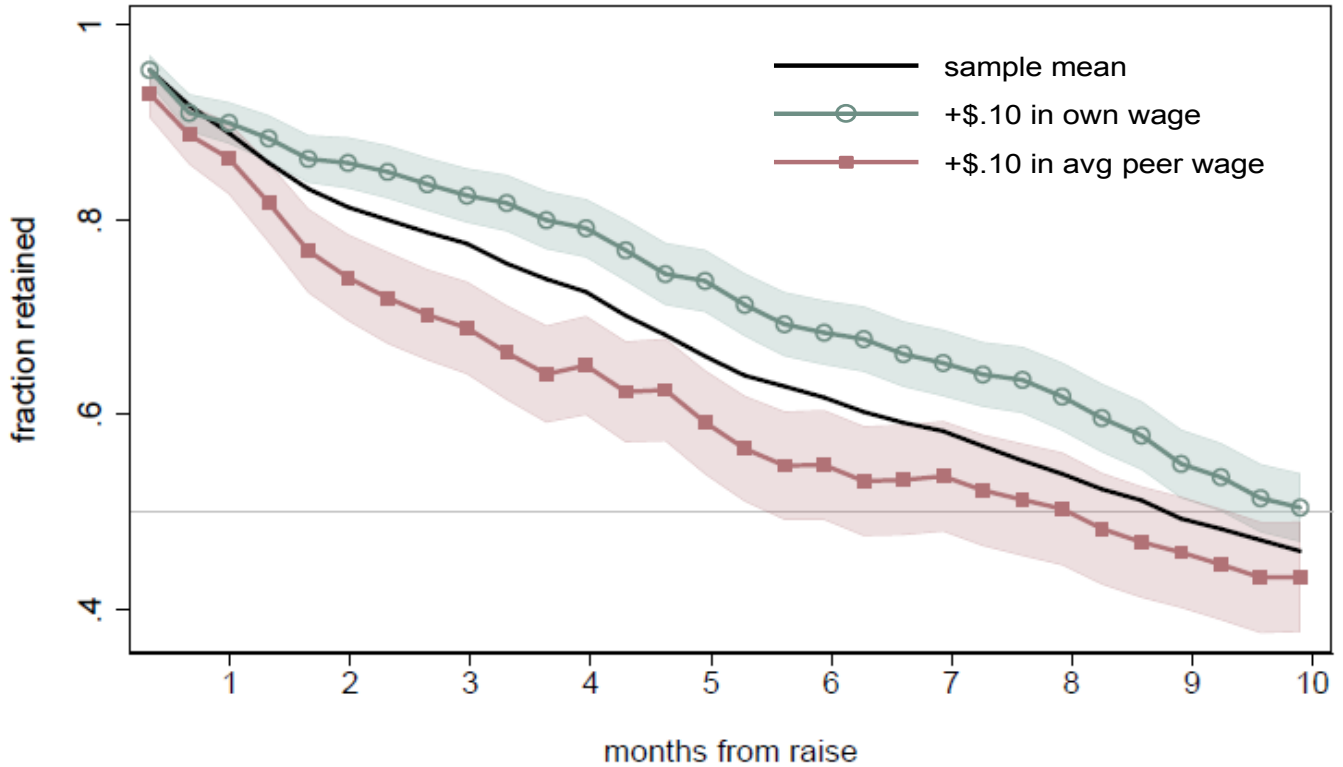
Note: The figure shows residuals from the stacked RD model of 3-month separations with baseline controls (as in Table 2, column 4). The running variable is the distance from own initial wage to the nearest pay-step threshold (see text for details). The markers show the mean residuals for each value of the running variable; marker size is scaled by the number of observations at each value. The lines show the fitted relationship between the running variable and the residualized separations. The intercepts are normalized to be zero at the left limit of the threshold, so that the value at the right limit is the estimated effect of the \$.10 discontinuity in the wage, and the shaded area at the right limit shows the 95% confidence interval for this estimate. Estimation samples are as in Table 2.

Figure 5. Three-month separation rate by distance from peer initial wage to nearest pay-step threshold




Note: The figure shows residuals from the stacked MRD model of 3-month separations with baseline controls (as in Table 3, column 4). The running variable is the distance from a representative peer's initial wage to the nearest pay-step threshold (see text for details). The markers show the mean residuals for each value of the running variable; marker size is scaled by the number of observations at each value. The lines show the fitted relationship between the running variable and the residualized separations. The intercepts are normalized to be zero at the left limit of the threshold, so that the value at the right limit is the estimated effect of the \$.10 discontinuity in the wage of a representative peer, and the shaded area at the right limit shows the 95% confidence interval for this estimate. Estimation samples are as in Table 3.

Figure 6. Survival function, adjusted for the effects of own and peer wage increases




Note: Dark line plots the fraction of the estimation sample that remains in the store as a function of the time elapsed since the raise. The other two lines show the survival rates adjusted for the estimated effect of a \$.10 increase in either own or peer average wage, based on models of separations for each interval shown, using the baseline MRD specification in column 3 of Table 3. The shaded areas show 95% confidence intervals.

Figure 7. Estimated effects of average peer wage on 3-month separation rate, varying wage band used to define peer group



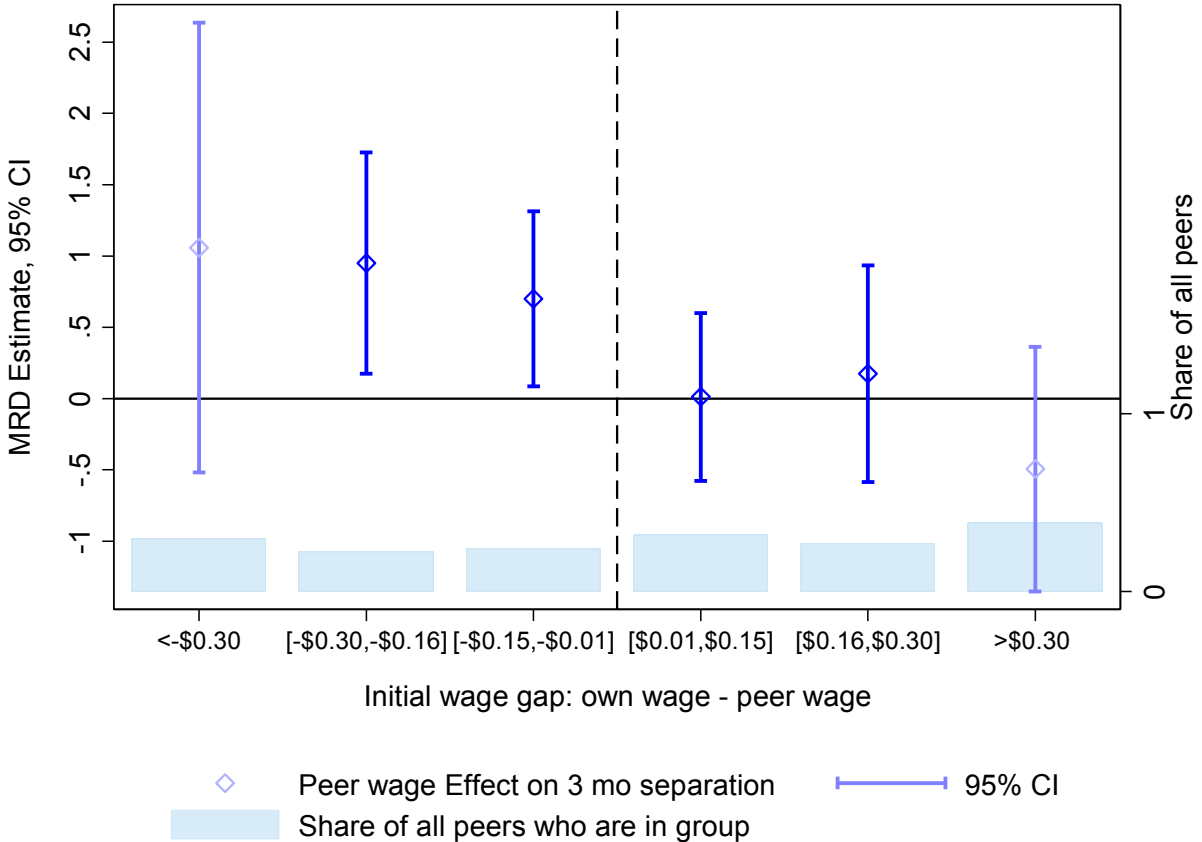
Note: The figure plots estimates and 95% confidence intervals of the 3-month separation response to an increase in the average peer wage from two-stage MRD models, for peer groups defined by various wage bands. For each wage band, peers include all co-workers who were eligible for a scheduled raise and whose initial wage is within the given dollar distance from the worker's own initial wage. Models include baseline set of controls as in Table 3, model 3.

Figure 8. Estimated effects of average peer wage on 3-month separation rate, alternative peer groups vs. random subsets of potential peers



Note: Scatter plots are estimates for two-stage MRD models of the separation response to an increase in the average peer wage for peer groups defined using alternative peer definitions (vertical axis), by the share of potential peers included when using that definitions (horizontal axis). Peer definitions are based wages, age, tenure, merit pay raise status, and geography, as indicated; see the text for more details. The straight line plots the locus of the 3-month separation response from a MRD model using a random subset consisting of the indicated share of all potential peers. All models include baseline set of controls as in Table 3, model 3.

Figure 9. Heterogeneous effects of average peer wage on 3-month separation rate, by initial wage gap (own – peer)



Note: The figure plots estimates and 95% confidence intervals from two-stage MRD models of the separation response to an increase in the average peer wage for six peer groups. These six groups were constructed by partitioning all potential worker-peer pairs by the initial wage gap between the worker and the peer. Models include baseline set of controls as in Table 3, model 3. The bars at the bottom of the figure indicate the share of all potential peers contained in a particular peer group, and are plotted using the right vertical axis.

Table 1. Summary statistics for estimation sample

	<u>mean</u>	<u>std. dev.</u>
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.

	<u>mean</u>	<u>std. dev.</u>
C. <u>Coworker (store-level workforce) characteristics:</u>		
mean number of coworkers	27.3	14.90
average wage	\$5.55	\$0.46
average age	24.2	2.70
% who got scheduled raise (potential peers)	63%	23%
% who got a merit raise in past year	44%	15%
% who got both scheduled and merit raise (estimation sample)	24%	14%
% with initial wages at multiple of \$.05	70%	16%
D. <u>Potential peers (coworkers with a scheduled raise)</u>		
number of possible peers	12.7	8.0
average initial wage	\$5.10	\$0.25
average scheduled raise	\$0.22	\$0.07
average age	21.7	3.2
% who got a merit raise in past year	42%	23%
% with initial wages at multiple of \$.05	72%	22%
E. <u>Peers within +/- \$.30 of own initial wage (preferred definition)</u>		
number of peers	9.0	6.5
average initial wage	\$5.13	\$0.28
average scheduled raise	\$0.21	\$0.08
average age	21.9	4.2
% who got a merit raise in past year	46%	28%
% with initial wages at multiple of \$.05	70%	27%
F. <u>Separation rate, by time from minimum wage increase</u>		
	<u>% who separate</u>	
within 1 month	12%	
within 2 months	19%	
within 3 months	23%	
within 6 months	38%	
within 9 months	51%	
G. <u>Reason for separation (within 9-months)</u>		
	<u>% of all separations</u>	
quit (job related)	55%	
return to school	21%	
move/transfer	16%	
fired	8%	

Note: Panels A-C and E 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 (Panel C) are defined as all employees in the same, entry-level job who were employed in the same store on the day of the minimum wage increase. Potential peers (Panel D) are defined as the subset of coworkers who got a scheduled raise. Relevant peers (Panel E) are the subset of potential peers whose initial wage is +/- \$.30 of own initial wage (see section 5.6 of text for explanation); peer characteristics in Panel E are shown for the sub-sample of 6,528 employees (97.6% of full estimation sample) who have at least one peer by this definition.

Table 2. Regression discontinuity estimates of separation response to change in own wage

Separation window:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1 month	-0.08 (0.14)	-0.11 (0.15)	-0.09 (0.15)	0.01 (0.15)	-0.09 (0.15)	-0.07 (0.16)	-0.12 (0.15)	-0.13 (0.16)	-0.16 (0.21)	-0.34 (0.27)
2 months	-0.29† (0.17)	-0.45* (0.18)	-0.46* (0.18)	-0.42* (0.17)	-0.46* (0.18)	-0.47* (0.19)	-0.50** (0.18)	-0.51** (0.19)	-0.42† (0.25)	-0.58† (0.33)
3 months	-0.28 (0.18)	-0.49* (0.19)	-0.52** (0.20)	-0.50** (0.19)	-0.52** (0.20)	-0.52** (0.20)	-0.56** (0.20)	-0.57** (0.20)	-0.39 (0.27)	-0.61† (0.35)
6 months	-0.42† (0.22)	-0.59** (0.23)	-0.61** (0.23)	-0.60** (0.22)	-0.61** (0.23)	-0.66** (0.24)	-0.66** (0.23)	-0.66** (0.24)	-0.40 (0.32)	-0.63 (0.40)
9 months	-0.38† (0.23)	-0.57* (0.24)	-0.57* (0.24)	-0.59* (0.24)	-0.57* (0.24)	-0.58* (0.25)	-0.62* (0.24)	-0.70** (0.25)	-0.41 (0.32)	-0.42 (0.44)
Number of observations	6691	6691	6691	6691	6691	6691	6691	6691	5503	2624
Linear own wage control	Y	Y	Y		Y	Y	Y	Y	Y	Y
ZIP code fixed effects		Y	Y	Y	Y	Y	Y		Y	Y
5-cent wage dummy			Y	Y	Y	Y	Y	Y	Y	
Stacked model (see note)				Y						
Quadratic own wage					Y	Y				
Cubic own wage						Y				
Worker + store controls							Y	Y		
Store fixed effects								Y		
Donut hole (see note)									Y	
5-cent wages only (see note)										Y

Note: Entries are regression coefficients from linear probability models of separation within 1, 2, 3, 6 or 9 months from the day of the minimum wage increase. All entries except those in column 4 are coefficients on own wage (as scheduled on the day of the minimum wage increase) from models that control for a smooth function of own initial wage with a separate intercept and slope in each year. Column 4 shows coefficients on a dummy for own wages at or above the nearest pay-step threshold from a "stacked" model that controls linearly for the distance from own initial wage to the nearest threshold. In columns 7 and 8, worker 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; and store 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. See section 4.1 for explanation of the controls in other models. The estimation sample in columns 1-8 includes all employees who received a scheduled raise on the day of the minimum wage increase and who had received a merit raise in July of the sample year (see section 3.2 for details). In column 9, the sample excludes wages that are just at or \$.01 below a pay threshold. In column 10 the sample includes only wages that are multiples of \$.05 (see section 4.1 for explanations). Parentheses contain robust standard errors clustered by store. † significant at 10%; * significant at 5%; ** significant at 1%.

Table 3. MRD estimates of separation response to changes in own wage and peer average wage

Separation window:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1 month										
own wage	-0.08 (0.14)	-0.11 (0.15)	-0.08 (0.15)	-0.02 (0.16)	-0.06 (0.15)	-0.07 (0.16)	-0.11 (0.15)	-0.11 (0.16)	-0.08 (0.16)	-0.05 (0.16)
peer average wage (2SLS)	0.36† (0.19)	0.24 (0.21)	0.29 (0.23)	0.17 (0.19)	0.56* (0.27)	0.30 (0.74)	0.23 (0.22)	0.25 (0.23)	-0.21 (0.39)	0.46† (0.26)
2 month										
own wage	-0.33† (0.18)	-0.47** (0.18)	-0.46* (0.18)	-0.58** (0.19)	-0.43* (0.18)	-0.47* (0.18)	-0.50** (0.18)	-0.48* (0.19)	-0.43* (0.20)	-0.44* (0.19)
peer average wage (2SLS)	0.82** (0.23)	0.58* (0.25)	0.74** (0.27)	0.76** (0.23)	1.11** (0.32)	0.23 (0.48)	0.63* (0.26)	0.54* (0.27)	0.54 (0.45)	0.74* (0.31)
3 month										
own wage	-0.34† (0.19)	-0.53** (0.19)	-0.55** (0.19)	-0.73** (0.21)	-0.52** (0.20)	-0.56** (0.19)	-0.59** (0.19)	-0.56** (0.21)	-0.47* (0.21)	-0.54** (0.21)
peer average wage (2SLS)	0.98** (0.25)	0.75** (0.26)	0.89** (0.29)	0.95** (0.24)	1.23** (0.35)	0.33 (0.38)	0.79** (0.28)	0.66* (0.29)	0.82† (0.48)	0.80* (0.33)
6 month										
own wage	-0.45* (0.22)	-0.65** (0.22)	-0.66** (0.23)	-0.80** (0.24)	-0.63** (0.23)	-0.68** (0.24)	-0.71** (0.23)	-0.65** (0.24)	-0.46† (0.25)	-0.78** (0.24)
peer average wage (2SLS)	0.84** (0.30)	0.50 (0.31)	0.65† (0.35)	0.85** (0.28)	0.93* (0.40)	-0.05 (0.99)	0.53 (0.33)	0.56 (0.34)	0.57 (0.56)	0.68† (0.39)
9 month										
own wage	-0.39† (0.23)	-0.58* (0.23)	-0.58* (0.24)	-0.73** (0.26)	-0.56* (0.24)	-0.58* (0.25)	-0.62** (0.24)	-0.68** (0.26)	-0.45† (0.26)	-0.72** (0.26)
peer average wage (2SLS)	0.50† (0.30)	0.14 (0.31)	0.35 (0.35)	0.62* (0.29)	0.55 (0.41)	-0.01 (1.22)	0.25 (0.34)	0.32 (0.34)	0.43 (0.56)	0.38 (0.41)
Number of employees	6528	6528	6528	6528	6528	6528	6528	6528	6294	6220
Number of observations	58660	58660	58660	58660	58660	58660	58660	58660	38401	43216
Linear own & peer initial wage	Y	Y	Y		Y	Y	Y	Y	Y	Y
ZIP code fixed effects		Y	Y	Y	Y	Y	Y		Y	Y
5-cent & peer merit dummies			Y	Y	Y	Y	Y	Y	Y	
Stacked model (see note)				Y						
Quadratic own & peer initial wage					Y	Y				
Cubic own & peer initial wage						Y				
Worker, peer + store controls							Y	Y		
Store fixed effects								Y		
Donut hole, peer wage (see note)									Y	
5-cent only, peer wage (see note)										Y

Note: Entries are estimated effects of increases in own wage and in average peer wage on the probability of separation within 1, 2, 3, 6 or 9 months from the day of the minimum wage increase. In all but column 4, estimates are from global MRD models that control for smooth functions of own and peer initial wages. Estimates in column 4 are from a stacked MRD model that controls linearly for distance from own and peer initial wages to the nearest threshold. See section 5.2 for more detail on the MRD model and alternative specifications. Peers in all models are defined as coworkers who got a scheduled raise and whose initial wage is +/- \$.30 of own initial wage. In columns 7 and 8, worker and store controls are the same as in Table 2 (see Table 2 note); peer controls are the same variables that are included for the worker. The estimation sample in columns 1-8 includes all worker-peer pairs for which the worker is in the RD analysis sample (see Table 2 note) and the peer fits our definition. In column 9, the sample excludes peers with initial wages that are just at or \$.01 below a pay threshold. In column 10 the sample includes only peers with initial wages that are multiples of \$.05. Parentheses contain robust standard errors clustered by store. † significant at 10%; * significant at 5%; ** significant at 1%.

Table 4. Separation elasticities implied by MRD estimates

	Dependent variable = 1 if leave within:				
	(1)	(2)	(3)	(4)	(5)
	1 mo	2 mos	3 mos	6 mos	9 mos
own-wage elasticity	-3.6 (6.7)	-12.7 (5.0)	-12.5 (4.3)	-8.9 (3.1)	-5.9 (2.4)
peer-wage elasticity (= – relative-wage elasticity)	12.9 (10.2)	20.4 (7.4)	20.3 (6.6)	8.7 (4.7)	3.5 (3.5)
gap-constant elasticity (net of peer effect)	9.3 (12.4)	7.7 (9.3)	7.8 (8.4)	-0.1 (5.7)	-2.3 (4.5)
sample mean of dependent variable	0.12	0.19	0.23	0.38	0.51

Note: Elasticities are calculated using the coefficients on own wage (estimates of β) and peer average wage (estimates of β_p) from baseline MRD model as in Table 3, column row 3; the sample means shown in last row of this table; and the sample average wage of \$5.15. The gap-constant own-wage elasticity is calculated by first summing the estimates of β and β_p (see equation 4) to obtain the gap-constant response to an increase in own-wage. Parentheses contain robust standard errors clustered by store (calculated using delta method).

Table 5. Split-sample RD estimates of separation response to change in own wage

	(1)	(2)	(3)	(4)
<i>Dependent variable:</i>	opposite side	same side	difference	Wald estimate of peer effect
wage gap (own - peer)	1.46** (0.06)	0.69** (0.06)	0.77** (0.08)	
Separation within:				
1 month	-0.52* (0.26)	0.18 (0.26)	-0.70† (0.37)	0.91† (0.50)
2 months	-0.98** (0.30)	0.05 (0.31)	-1.03* (0.44)	1.34* (0.60)
3 months	-1.13** (0.32)	-0.05 (0.34)	-1.07* (0.47)	1.39* (0.64)
6 months	-0.80* (0.39)	-0.34 (0.40)	-0.46 (0.57)	0.60 (0.74)
9 months	-0.71† (0.41)	-0.67 (0.43)	-0.04 (0.59)	0.05 (0.77)
Number of observations	2906	3485	6391	6391

Note: Entries are own-wage coefficients from baseline RD models as in Table 2, row 3 (see Table 2 note). Sample is split by whether the worker and the majority of her peers are on the same side or opposite sides of their nearest discontinuity threshold. Estimates and standard errors for the differences (column 3) are from a model in which own wage and all covariates are interacted with a dummy for the opposite-side sample. Wald estimates of peer effect are two-stage least squares estimates from the interacted model (see section 5.5 for derivation). Parentheses contain robust standard errors clustered on store. † significant at 10%; * significant at 5%; ** significant at 1%

Table 6. Effects of October 1, 1996 scheduled wages on separations, estimates from RD and MRD models using treated and placebo samples

Separation window:	RD (own-wage only)			MRD (own and peer)		
	(1) 1 mo.	(2) 2 mo.	(3) 3 mo.	(4) 1 mo.	(5) 2 mo.	(6) 3 mo.
A. Treated sample (October 1, 1996)						
own wage	-0.19 (0.20)	-0.60* (0.25)	-0.60* (0.28)	-0.18 (0.19)	-0.57* (0.24)	-0.63* (0.27)
peer average wage				0.18 (0.22)	0.65* (0.28)	0.60† (0.33)
Number of employees	3009	3009	3009	2901	2901	2901
Number of observations	3009	3009	3009	23172	23172	23172
B. Placebo sample (July 8, 1996)						
own wage	0.23 (0.17)	-0.05 (0.26)	-0.10 (0.28)	0.18 (0.17)	-0.10 (0.25)	-0.14 (0.27)
peer average wage				-0.15 (0.23)	-0.25 (0.35)	-0.38 (0.36)
Number of employees	4399	4399	4399	4269	4269	4269
Number of observations	4399	4399	4399	32764	32764	32764

Note: Treated sample consists of observations in the main estimation from the date of the first minimum wage increase (October 1, 1996). Placebo sample is 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. Placebo wages are calculated by applying the corporate rule for wage adjustments made on October 1, 1996 to initial wages on July 8, 1996. Peers are coworkers employed on July 8, 1996 whose wage on that day was +/- \$.30 from the worker's own wage. Estimates are from baseline models as in column 3 of Table 2 (RD) and Table 3 (MRD). Parentheses contain robust standard errors clustered on store. † significant at 10%; * significant at 5%; ** significant at 1%.