NBER WORKING PAPER SERIES

HOW COSTLY IS TURNOVER? EVIDENCE FROM RETAIL

Peter J. Kuhn
Lizi Yu

Working Paper 26179
http://www.nber.org/papers/w26179

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2019

We thank Achyuta Adhvaryu, Clement de Chaisemartin, Dan Hamermesh, Mitch Hoffman, John Horton, Chris Stanton, Steve Trejo, and seminar participants at Royal Holloway University of London, the UCLA Trans-Pacific Pacific Labor Seminar, and the IZA Conference on Exploring the Breadth of Labor Economics in honor of Dan Hamermesh for helpful comments. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Peter J. Kuhn and Lizi Yu. 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.

How Costly is Turnover? Evidence from Retail
Peter J. Kuhn and Lizi Yu
NBER Working Paper No. 26179
August 2019, Revised May 2020
JEL No. J31,J63,J64

ABSTRACT

We estimate turnover costs in small retail sales teams using daily sales data and an advance notice requirement to address endogeneity concerns. In addition to short-staffing and onboarding costs, we identify two less familiar sources of turnover costs: incumbent workers' recruitment activities, and reductions in team morale after a departure is announced. Our estimates of total turnover costs are relatively modest, however: Ten percent higher turnover is about as costly as a 0.6% wage increase. We attribute these low costs to a set of complementary personnel policies which ensure that only 25 percent of departures result in a short-staffing spell.

Peter J. Kuhn
Department of Economics
University of California, Santa Barbara
2127 North Hall
Santa Barbara, CA 93106
and NBER
pjkuhn@econ.ucsb.edu

Lizi Yu
liziyu@ucsb.edu

1 Introduction

Annual turnover rates in the U.S. are high, and both employers and management experts frequently bemoan turnover costs.[1] Despite this, well-identified estimates of the cost of turnover to employers remain elusive, due in part to the well-known possibility of reverse causation: turnover and productivity may be correlated not because turnover reduces productivity, but because low productivity causes workers to leave. To address this issue, a number of researchers have recently studied the effects of unexpected worker exits, such as those due to sudden worker deaths (Jones and Olken, 2005; Azoulay et al., 2010; Jäger, 2016) and sudden political change (Waldinger, 2011; Borjas and Doran, 2012). While this approach arguably solves the endogeneity issue – because the deaths in question are not likely caused by low productivity– these studies mostly focus on influential but highly atypical groups of workers like CEOs and star scientists. Because regular workers may be easier to replace than leaders, and because organizations may use advance knowledge of 'normal' departures to mitigate their costs, these studies seem likely to overestimate the costs associated with the everyday turnover that occurs in most workplaces.

Well-identified studies of the costs of 'normal' turnover, on the other hand, are scarce, perhaps because they face an additional identification problem related to anticipatory behaviors caused by advance knowledge of the departure. Whenever such behaviors occur, some of the causal effects of a worker's departure are realized before the worker leaves, complicating efforts to test or adjust for reverse causation. A related challenge is the lack of information about when members of a team become aware that a member will leave. Finally, progress has also been hampered by a lack of granularity in the time dimension of most data: Without precise knowledge of when a worker departs and when replacements are hired, the exact mechanisms via which turnover affects productivity can be hard to isolate.[2]

Motivated by these challenges, we estimate the effect of employee turnover on the productivity of front-line retail sales teams. In addition to comprising a substantial share of national employment, these jobs are arguably representative of other high-turnover, low-paid service occupations, such as customer service, hospitality and restaurants.[3] We measure the effect of the ongoing, voluntary, employee-initiated turnover that occurs at a relatively high rate in many

[1] According to the Bureau of Labor Statistics, the annual separation rate (total separations divided by mean annual employment) in 2017 was 43.0% overall, and 53.0% in retail trade (U.S. Bureau of Labor Statistics, 2018d). Articles in the business press have argued that high turnover rates are very costly in retail settings (Ton, 2012).

[2] Early turnover cost studies like Glebbeek and Bax (2004) and Siebert and Zubanov (2009) measured turnover and productivity on an annual level; without an instrument this makes it very hard to determine whether departures that occur during the course of a year are the cause or consequence of low productivity in that year. More recent studies use monthly data, which mitigates these problems and provides valuable additional insights about mechanisms (Ton and Huckman, 2008; Bartel et al., 2014; Drexler and Schoar, 2014). We are not aware of any turnover cost studies that measure productivity more frequently than monthly, or that utilize information on when teams became aware of an impending departure to distinguish between anticipatory behavior and reverse causation.

[3] In December 2018, retail workers accounted for 10.8 percent of total U.S. employment; accommodation and food service workers comprised another 9.2 percent. (U.S. Bureau of Labor Statistics, 2018a,c). The 2017 annual turnover rate in Accommodation and Food Services is even higher than in retail trade, at 72.5 percent (U.S. Bureau of Labor Statistics, 2018d).

such establishments, using daily sales and employment records from 118 menswear stores in 2015 and 2016. These stores were operated by the same firm ("Firm A"), and are mostly located in Guangdong Province, China. Four features of our data and setting allow us to improve on existing estimates of turnover costs: the employer's policy of no layoffs or dismissals; the very small team size of 2-7 workers; our (unique) access to daily data on team productivity, worker departures, and hires; and the fact that all worker departures are pre-announced in the workplace we study.

In more detail, the employer in our context, Firm A, has a policy of not laying off or dismissing sales employees unless it shuts down a store, and has adhered to this policy since 2008. While this does not rule out the possibility that under-performing workers could be induced to quit in various ways –in fact we argue that this is a component of Firm A's personnel strategy– it simplifies our analysis by ensuring that all departures are made via the same formal mechanisms and are subject to the same employment regulations. Second, the small size of our sales teams allows us to study the effects of a discrete event –the departure of a single employee–, and makes it easier to detect the effects of that event on the productivity of the entire team. Third, in addition to daily observations on employee departures and team productivity, we also have daily observations on the hiring of replacements, linked to each individual departure; as a result all the changes in team composition in our data occur in the seams between periods. This greater precision allows us to learn more about the sources and timing of turnover costs, for example by distinguishing the effects of short-staffing (operating with fewer workers until a replacement is found) and on-boarding (productivity losses associated with integrating a new employee into the team) from other possible sources of turnover costs such as employee involvement in recruiting activity and changes in team morale.

Finally, the fact that all employees of Firm A submit an official departure notice two weeks in advance, combined with the fact that departure is essentially assured after notice has been given, shuts down the possibility of reverse causation during the two-week notice period: After the notice date, the worker has already decided to leave, so her actual departure can no longer be caused (or prevented) by shocks to team productivity.[4] Shutting down reverse causation during this period gives us a clean estimate of anticipatory behavior – i.e. of pre-departure changes in team performance that are caused by the worker's impending departure. Such estimates are scarce because information on notice dates is not typically available. Absent this information, we cannot tell whether dips in team productivity before the departure are the economic shocks that caused the departure, or behavioral responses to the expected exit.[5]

Our main findings are as follows. First, we see robust and statistically significant productivity

[4]Firm A's enforceable notice contract prevents workers from leaving before two weeks have passed. A small share of departing workers (under 5 percent) agree to stay a few days beyond two weeks to help with the transition to a new employee. Otherwise, Firm A reports that all workers depart exactly two weeks after they give notice.

[5]We use two complementary approaches to rule out reverse causation *before* the notice date – i.e. the possibility that workers give notice in response to previous productivity shocks. One is the standard approach of checking for pre-trends during that interval. Our more novel approach invokes the fact that notice dates are strongly bunched exactly two weeks before the end of a calendar month. This behavior –which is consistent with a (rational) desire by workers to synchronize their departure with Firm A's pay cycle– strongly suggests that departures are planned several weeks in advance of the notice date.

losses at four distinct times during the departure process: Around the time the worker is required to give notice (AN); just before she departs (BD); just after she leaves, if she hasn't already been replaced (short-staffing, SS); and just after the replacement worker joins the team (on-boarding, OB). Second, while the mechanisms behind the on-boarding and short-staffing effects seem clear, we attribute the AN and BD effects, respectively, to two mechanisms that have received little attention in the turnover literature: the involvement of incumbent employees in recruiting activities, and short-timer effects, i.e. reductions in worker effort associated with a change in the repeated game between the departing worker, her colleagues, and her employer.

Third, the four sources of productivity losses surrounding a departure (AN, BD, OB and SS) each contribute about equally to the total. Thus, only 27 percent of the lost output associated with turnover is associated with short-staffing. This is because (a) advance notice of the departure makes short-staffing spells quite rare; (b) trained, temporary workers are sometimes used to bridge the gap between the departure and the hire of a permanent replacement; and (c) when they occur, short-staffing spells are brief and lead only to short productivity reductions.

Fourth, almost two thirds (63%) of all productivity losses associated with turnover are incurred before the departing worker leaves. This is because the employee recruiting activities and morale reductions associated with AN and BD effects occur before the departure, and because a substantial share of on-boarding happens while the departing worker is still present. While it might be tempting to view this result as a peculiar feature of China's labor laws (which make it relatively easy for employers to enforce contracts requiring employee notice), it is important to note that a large number of countries mandate employee advance notice, and that two weeks of employee notice is customary in the United States.[6]

Fifth, after accounting for administrative and wage costs of turnover, and for changes in variable costs associated with turnover, we calculate that the turnover of a single employee reduces profits by an amount that represents 9.4 days of per-employee net sales, or 1.1% of a worker's net sales over a 2.3-year career. In terms of per-employee wages, a departure costs the equivalent of 63 days of wages, or 7.6% of wages over a 2.3-year career.[7] Expressed a different way, we estimate that permanently reducing its turnover rate by ten percent would allow Firm A to permanently raise its wage payments to workers by 0.6%.

Our results contribute to a number of literatures in labor and personnel economics. One of these is the literature on firm wage effects – i.e. the fact that some firms consistently pay higher wages than others in the same location and industry. While some authors (Burdett and Mortensen, 1998; Ton, 2012) have argued that higher wages can pay for themselves by reducing turnover costs, our

[6]Based on a web search of employer and employee notice requirements, as of May 2018, 29 of the 33 OECD countries –all but Greece, South Korea, Mexico and the United States– required quitting employees to give advance notice; mandated amounts ranged from one week to six months depending on country, years of service and salary level. In the U.S., business websites like Lifehacker.com and Marketwatch.com refer to two weeks as the "general" or "tried and true" standard for employee notice (Henry, 2011; Jagannathan, 2018). Survey evidence (Klotz and Bolino, 2016) indicates that 60% of U.S. employees have formally or briefly given an advance notice when quitting a job.

[7]The difference between these output- and wage-equivalent estimates stems from the low share of wages in Firm A's total costs, in which rental payments for retail space figure very prominently.

estimates suggest this is very unlikely.[8] Specifically, the absolute value of the quit-wage elasticity would need to exceed 16.8, a number that substantially exceeds all estimates we know of, including our own. Thus, other explanations of firm wage effects, such as rent-sharing by employers (Hildreth and Oswald, 1997), or an effect of wages on worker quality (Giuliano, 2013) are needed.

Second, our estimates of total turnover costs –which are based on voluntary departures– imply a strict threshold for worker retention: With turnover costs this low, Firm A would be better off if all workers whose productivity was more than 2.3 percent below average departed voluntarily. Importantly, this does not imply that Firm A should dismiss all workers who fall below this threshold. A number of factors, including legally required severance pay, noisy worker productivity, and greater morale costs may make firm-initiated dismissals a poor substitute for separation decisions made at the team level. Instead, we argue that Firm A has adapted to its high-turnover environment by delegating a significant share of the personnel selection and retention process to sales team members. This approach is part of an HRM system where team-based pay incentivizes workers' effort *and* personnel selection decisions.

Third, our analysis of on-boarding costs –productivity losses associated with the addition of a new employee– relates to a large literature on training and firm-specific skills. While it is frequently argued that new employees are on average less productive than experienced ones because they lack industry- or firm-specific skills, our findings suggest an important team-specific component as well. Specifically, we find that *total* team productivity falls for a period of time after a new member is added. This reduction is observed even when the new worker has previous experience in the firm and industry. In conjunction with a similar result for nursing teams in Bartel et al. (2014), our results therefore suggest that integration costs –the costs of learning to work with a new person– as distinct from skills acquisition *per se* play an important role in turnover costs.

Fourth, while a number of recent papers have considered the advantages and disadvantages of using employee referrals to recruit workers, those papers focus primarily on potential benefits of such programs in identifying and hiring better workers (Pallais and Sands, 2016; Friebel et al., 2019). We believe our paper is the first to document the direct productivity costs of involving incumbent employees in the recruiting process. Specifically, we find that the allocation of incumbent employees' time and attention to filling a vacancy is just as costly to the firm as the on-boarding or the short-staffing costs associated with a typical departure.

Fifth, our estimates of short-timer effects contribute to a number of literatures that emphasize the 'shadow of the future' in disciplining current interactions between workers, their peers, and their employer in enforcing efficient exchange. This shadow plays a central role in models of relational contracting (Kreps et al., 1982), of dismissal threats as employee motivation (Shapiro and Stiglitz, 1984), and of mutual monitoring in teams (Kandel and Lazear, 1992). To the best of our knowledge, this paper provides the first well-identified, quantitative evidence of the effests of

[8] Burdett and Mortensen (1998)'s influential model derives an equilibrium wage distribution purely from the tradeoff between wages and turnover. Ton (2012) argues that turnover costs help account for the co-existence of high- and low-wage employers like Costco and Walmart in the retail sector.

(removing) expected future interactions on current productivity in a real workplace.[9] Sixth, our estimates contribute to a small but influential lab and field literature documenting peer effects on worker effort (Falk and Ichino, 2006; Mas and Moretti, 2009). In particular, we show that short-timer effects on team performance increase in magnitude with team size, and with the share of the work day during which team members' hours overlap, patterns that are hard to explain without reference to contagion across team members.

Seventh, our results contribute to the literature on the role of advance notice in managing worker separations. While effects of advance notice requirements for firms who wish to dismiss workers have been studied (Ruhm, 1992; Jones and Kuhn, 1995), advance notice requirements for workers who wish to quit have received relatively little attention. Thus, our finding that employers find replacement workers during the employee's notice period parallels existing results showing that workers find new jobs during the employer's notice period (Jones and Kuhn, 1995).

Eighth, the timing of our estimated turnover costs provides a methodological lesson regarding the interpretation of event studies (Dobkin et al., 2018). While testing for pre-trends has become a popular diagnostic tool in these studies, our results illustrate the underappreciated fact that tests for pre-trends are joint tests for spurious relationships (such as reverse causation and 'correlated shocks') and for anticipatory behavior, which can be thought of as a causal effect of the event (Hendren, 2017; Abraham and Sun, 2019). In our case, strong and robust productivity declines are observed before an employee's departure date. Because a worker's departure is essentially locked in after she announces, we can be confident that productivity losses between the announcement and departure are a result, not a cause of the impending departure. In our context, this is critical because 63 percent of the costs associated with turnover are incurred during that notice period.

Finally, we contribute to literatures on systems of HRM practices (Ichniowski et al., 1997), agility in organizations (Doz and Guadalupe, 2019), employment protection laws (Acemoglu and Angrist, 2001; MacLeod and Nakavachara, 2007), the labeling of separations (McLaughlin, 1991; Kuhn and Sweetman, 1998), and team-based pay (Hamilton et al., 2003). Rather than passively responding to employee departures, we argue that Firm A has established a set of complementary personnel policies that have allowed it to adapt to a high-turnover environment in which quits are less costly than layoffs. These policies include stipulating and enforcing an advance notice requirement, using that requirement to hire early enough to eliminate most short-staffing spells, arranging access to an experienced, contingent workforce to fill many of the remaining staffing gaps, involving incumbent workers in the recruitment process, delegating many employee discipline problems to sales teams by announcing a no-layoffs policy, and incentivizing team members to manage themselves using team-based commissions. In sum, our analysis shows that firms can adapt in a number of creative ways to a high turnover environment; ignoring these adaptations could lead investigators to substantially over-estimate the costs of turnover to firms.

[9] Moskos (1975) provides anecdotal evidence of 'short-timer' effects on work effort among U.S. soldiers during the Vietnam War. Friebel et al. (2016) find reductions in output in response to sales and closure announcements for small retail establishments, but these can also be interpreted as workers' reciprocal responses to firms' actions. Brown et al. (2004) show how relational contracting can emerge and improve efficiency in a laboratory environment.

2 Setting and Data

All our data were supplied by Firm A, a large manufacturer and retailer of men's clothing in China. This Section provides a brief overview of the firm and the data available to us; additional details and supporting documentation are provided in Appendix A.1. During the analysis period (January 1, 2015 through December 31, 2016), Firm A operated 164 retail stores, mostly in Guangdong Province. In this period, we observe the exact dates of hires and departures from 118 stores, whose sales employees are centrally hired and paid by Firm A. We observe 12 store openings and 17 store closures among the 118 stores. Dropping observations that are within 30 days of store openings, closures or remodeling, we are left with 75,801 daily observations of team-level sales. The company does not maintain individual employee sales data, nor does it use them in setting pay. Instead, all employees receive a base salary that varies by location and seniority, plus a bonus that is determined by the monthly sales of their store-level team.[10]

2.1 Team Production, Pay and Store Sales

Firm A operates retail stores in two types of locations: department stores and shopping malls, henceforth "host institutions". Each of Firm A's stores has a target size of 2-7 employees that is set annually based on expected sales. Around half of the stores in our analysis have a target size of three employees, and stores are at their target size almost all the time. According to Firm A, essentially all the salespeople in our sample are female and between the ages of 20 and 55. While their duties are mostly similar to U.S. retail workers', Firm A's salespeople play a more active role, since most items must be retrieved from storage to be tried on. Since customers place a high value on fit and speed of service, it is widely believed that effort has a substantial effect on sales. Production also requires co-operation among sales workers; for example customers can be served faster and better when employees share the tasks of fetching items and interacting with the customer.

Co-operation is also important in setting work schedules, a process that is co-ordinated by one of the salespeople who also acts as the store manager: Effective scheduling requires all team members to make efficiency-enhancing compromises. In addition to deciding when each worker takes her one rest day per week, workers must be allocated to all the six-hour shifts that allow the stores to operate for 12 hours per day, 7 days a week. The most important feature of these shifts for our purposes is that the amount of overlap between workers' shifts increases with store size: workers never overlap in two-person stores (each working for six of the store's twelve opening hours), with more overlap in larger stores.[11]

[10] Team-based pay is a large and growing feature of compensation: According to Lawler and Mohrman (2003), the share of Fortune 1000 companies using work-group or team incentives for more than a fifth of their workers more than doubled, from 21 to 51 percent, between 1990 and 2002. Examples include Continental Airlines (Knez and Simester, 2001), Microchip (Adamson et al., 2014), steel minimills (Boning et al., 2007) and apparel manufacturing (Berg et al., 1996; Hamilton et al., 2003). Examples from retail include German retail establishments(Friebel et al., 2017) and tip sharing at restaurants (Scudder, 2017). We study the incentive effects of Firm A's commission scheme in a related project (Kuhn and Yu, 2019).

[11] Based on our conversations with Firm A, it is unlikely that rest days account for the separation-related productivity

Consistent with the co-operative nature of production, Firm A's retail sales workers are paid based on total team performance. Employees' monthly compensation consists of a base payment and a commission component. The base payment, about $273-$360 per month, differs somewhat across stores to reflect different living costs in different cities. Within stores, base pay varies with the employee's firm tenure.[12] Commissions are based on the team's total monthly performance; commission rates per employee vary across stores between 0.7% and 1.2% of gross sales. Commission rates are identical within a team, but are lower in larger teams in order to yield similar total pay in teams of different sizes. Monthly pay is deposited into employees' bank accounts on the 20th of the following month. Overall, Firm A pays at- or above-market wages: During the analysis period, its sales workers earned an average of $538 per month, compared to an average monthly compensation for retail salespeople of $510 in Guangdong province.

Sales at Firm A's stores are subject to large day-of-week, holiday, and seasonal effects; most of these patterns are relatively predictable and our analysis will correct for them with 731 calendar day fixed effects; one for each day in our two-year observation window. Daily store-level sales are, however also subject to large random shocks, in part due to the 'lumpy' nature of the product: on a typical day, a store sells 11.4 relatively high-value items, though days with zero sales and days with much higher sales are also common.

2.2 Departure and Hiring Procedures

As noted, Firm A has a policy of not dismissing employees and has honored this policy since at least 2008, except when it decides to shut down a store. Therefore the 186 departures we observe are all employee-initiated, voluntary departures. At 34 percent, the quit rate among Firm A's sales workers is comparable to the industry average in both China and the U.S.. Firm A's salespeople on average have 3.45 years of firm tenure, similar to the 3.0 years of median tenure in the U.S. retail industry in 2018; the mean tenure of employees who leave Firm A is about 2.27 years.

The departure process from Firm A begins when the worker provides two weeks' advance notice of her departure, in writing. This is required by the labor contract signed by every front-line sales employee of Firm A at the time of hire; compliance with this requirement is very high for two main reasons. First, Firm A is entitled to claim losses from employees' breach of the contracted amount of notice; such claims are easily collected by withholding wages because employees are not paid until the 20th of the following month. Second, employees need departure paperwork from their original employer to start a new job and transfer their social security benefits, so they have a strong interest in maintaining good relations with their previous employer.[13]

Hiring is done by Firm A's 11 regional managers, with the assistance of the store's employees.

reductions we observe. As we detail in Section 2.2, even among about-to-depart workers, managers have considerable leverage to ensure workers are present when they are needed.

[12]The store manager receives a small additional allowance for coordinating work schedules.

[13]Firm A's two-week contractual worker notice requirement is actually less strict than the 30 days of notice stipulated by Chinese labor law (China Labor Contract Law 2007), but Chinese courts generally defer to the terms of written employment contracts that deviate from standard legal requirements, when the deviations are modest and when they favor workers.

As soon as regional managers receive the departure notice, they solicit employee referrals from the departing worker's store. Employees are happy to assist, not least because the team compensation formula rewards the hiring of productive and co-operative team-mates. According to Firm A's managers, the involvement of sales employees in the hiring process takes place mostly right after the departure has been announced, and may pick up again near the departure date if a replacement has not been identified. Finally, since departures from stores that are not at their target size are rare and have different effects, our analysis will focus only on departures from target-size stores.[14]

Figure 1 provides a flow chart of the departure process, showing its four possible outcomes:

- **Early Refills** (40% of departures) occur when the replacement employee starts to work before the departure; thus the team briefly rises above its target size during the replacement process.
- **On-time Refills** (17% of departures) occur when the replacement worker starts on the day after the departing employee's last day of work; here the team remains at its target size throughout the replacement process.
- **Temporary Replacements** (18 % of departures) are used when smaller teams (target size of 2 or 3) cannot find a permanent replacement on time. These temporary replacement workers are typically re-assigned from the host institutions; thus they have (very) local sales experience and may have worked at the store before. As in the two previous cases, these stores never experience a period of short-staffing.
- **Late Refills** (25% of departures) occur in larger stores (4+ employees) when a replacement is not hired on time; here the team falls below its target size following the departure, resulting in a period of short-staffing.

In Figure 2, we provide more precise details on the timing of replacement hires by plotting the cumulative share of departures that have been replaced as a function of elapsed time since notice. During the two-week notice period, the probability a new worker arrives is relatively constant at about 3 percent per day, cumulating to 40 percent by the leaver's last day of work. The modal hiring day is the day immediately after the leaver's last day; these *on-time* replacements bring the share of departures that have been replaced up to 57 percent. After that, replacements continue to arrive at about three percent per day for about a week, then at a slower rate thereafter. Overall, about 89 percent of departures are replaced within 30 days after the departure date, and 94 percent are replaced within three months.[15]

[14]We do not observe any departures from below-target-size stores. Above-target-size stores that lose a worker are coded as not experiencing a departure. A small number of cases where two employees left on the same day are coded as one single departure (event).

[15]In the small minority of cases where recruitment proves difficult after a month or two, Firm A sometimes abandons its recruitment activities and takes the alternative approach of raising commission rates and letting the store remain at a smaller size.

3 Results

3.1 Econometric Approach

We begin our analysis by conducting a non-parametric event study (Dobkin et al., 2018) of team productivity trends in the 61 days surrounding a representative departure, without conditioning on when the departure is replaced.[16] This gives us a clean summary of the productivity changes surrounding a typical departure from Firm A, while imposing the fewest possible assumptions about the timing and magnitude of those changes. To isolate the causal effects of turnover, these trends are measured relative to store-day observations from a large control group of store-day cells that are more than 30 days from a departure, more than 30 days from a hire, and when the store is at its target size.[17] As noted, workers' two-week advance notice requirement allows us to cleanly interpret productivity changes during the notice period as anticipatory behavior that is caused by the impending departure, since reverse causation is not possible during this period (the decision to leave has already been made).

Throughout our analysis, the key time points for each departure are defined as follows:

- P_0 is the day of departure, i.e. the last day the departing member works at the store.
- P_{-14} is the *required* day of notice corresponding to this departure.[18]

To keep our specification as flexible as possible, we estimate 30 lead and lag terms for 2-day and 3-day bins in this 61-day interval using a modified event study approach:[19]

$$\begin{aligned} S_{it} =& \alpha + \beta_{-30,-29} \cdot P_{-30,-29} + ... + \beta_{-14,-13} \cdot P_{-14,-13} + ... + \beta_{-1,0} \cdot P_{-1,0} + ... + \beta_{29,30} \cdot P_{29,30} \\ & + \gamma_1 \cdot D_t + \gamma_2 \cdot I_i + \epsilon_{it}, \end{aligned} \tag{1}$$

where *i* indexes stores, *t* indexes days, and S_{it} is the store's daily sales. We combine P_0 and P_{-1} into a 2-day binned indicator variable, denoted $P_{-1,0}$, which takes a value of 1 if today is the departing employee's last two days of work, and 0 otherwise. Similarly, we combine P_{-14} and P_{-13} into a 2-day binned indicator variable, identifying the *first* two days to which the notice mandate applies (i.e. the first day of the required notice period and the following day). All the other days are then

[16] Friebel et al. (2018) is the only field experiment we know of that attempts to manipulate turnover rates. While their intervention did reduce turnover, it also caused a re-allocation of middle managers' time that directly affected productivity. This makes it hard to infer turnover costs from their treatment effects.

[17] Appendix A.2 shows that the results are robust to both shorter and longer windows than 61 days.

[18] Firm A has no data on the exact dates when it receives the notice, but it asserts that in almost all cases, notice is given either exactly two weeks before the departure, or a day or two before that. Therefore, if an employee's last day of work is a Friday, then we assume that the worker submits her notice during the 14th day prior to the departure, also on a Friday, or on the preceding Wednesday or Thursday.

[19] In a standard event study, one pre-event time indicator is excluded as a reference category. In our setting, we do not have a strong prior of a specific time point when the impact should start to occur, and a small group of reference observations will not be sufficient to identify our many store and day fixed effects. So we estimate coefficients for all periods during the 61-day period surrounding the departure, with the observations outside of that window (control period observations) acting as the reference category. See Balasubramanian and Sivadasan (2011) for a similar specification.

grouped into 2-day binned variables with one exception: Halfway between the advance notice and departure date, we use a 3-day bin, $P_{\text{-8,-6}}$. The lead coefficients prior to the announcement day ((-30, -29) through (-16, -15)) provide a check for pre-trends, since this may be when the employee is deciding to quit (though some workers may have given notice during days (-16, -15).) The coefficients between the notice and departure date ((-14, -13) through (-1, 0)) examine productivity changes after an employee has notified the team of her departure but while she is still working at the team, while the lag terms ((1,2) through (29,30)) measure the change in team productivity after the departure has occurred. In all cases, the coefficient estimates should be interpreted as productivity deviations from the control period when there are no personnel disruptions going on, net of the other regression controls in equation 1.

Motivated by the strong day-of-week, seasonal and holiday effects documented in Appendix A.1, the remaining regression controls in equation 1 are time and store fixed effects. Specifically, D_t is a vector of 731 dummies for each day in 2015 and 2016, which captures time-varying effects that are common to all stores. Notably, in addition to common seasonal and weekday effects, these dummies also capture holidays that occur on different dates in different years, plus daily fluctuations in business conditions and weather that are common to Firm A's stores.[20] Finally, I_i is a vector of 118 store dummies that control for time-invariant unobserved heterogeneity across stores, such as physical location and store display. Robust standard errors are clustered at the store level. In addition we report significance levels that correct for multiple hypothesis testing using the Bonferroni correction procedure (Benjamini and Hochberg, 1995).

3.2 Full Sample Results

Estimates of equation 1 are presented in Table 1. As noted, estimation sample includes all departures from Firm A, plus the much larger control sample of store×day cells that are more than 30 days distant from a hire or departure. None of the estimates include controls for when the replacement worker arrived.

A first observation from Table 1 is the absence of pre-trends in team sales before the announcement date (with the exception of days -16 and -15, on which some employees may have delivered their notices). Second, when we do not condition on when the replacement was hired, Table 1 indicates that store sales differ significantly from control-period sales at only three points during our treatment window: the four days surrounding the employee's announcement of her departure, her last four days on the job, and the first two days immediately following her departure. In the four days surrounding the employee's departure announcement, there is an average productivity loss of (78 + 128)/2 = 103 dollars per day, which amounts to 17.4% of average daily sales. The productivity loss near the departure is of a similar magnitude, leading to about a 20.9% reduction of daily performance in the last four days of the departing employee's work. As already noted, we refer to these output reductions –which turn out to be highly robust– as the around-notice (AN) and before-departure (BD) effects. Following the departure, there is a sizable but short-lived

[20] Weather conditions across locations in Guangdong are generally very similar.

productivity loss of $131 or 22.2% per day, but it only lasts for two days and no significant losses are identified thereafter.

An interesting feature of these results is that most of the turnover-related costs they identify occur before the worker leaves. Although this may seem counter-intuitive at first, it is consistent with the fact that 57 percent of departing workers are replaced on or before the day following the departure, plus the fact that small stores that fail to hire 'on time' fill the vacancy with temporary replacement workers. Thus, short-staffing is rare, which may also account for the small and temporary productivity losses right after the departure. For the same reason, most of the on-boarding at Firm A occurs before the actual departure, which could account for some of the pre-departure losses we see.[21]

Despite the advantages of our advance-notice-based approach, at least two remaining weaknesses might still affect the interpretation of the AN and BD effects estimated in Table 1. The first is that the AN effect (the productivity decline around a worker's *notice* date) may have caused her departure announcement, because these two events are relatively contemporaneous. Second (and less plausibly) it is possible that the employee was able to foresee the productivity decline that occurred around her *exit* at least two weeks in advance, and decided at that early date to announce her decision to end her career at Firm A in order to avoid this small decline.[22]

We offer three pieces of evidence to address these concerns. The first is that the productivity declines we detect around both the announcement and departure dates are small and very typical of the short-term productivity changes in Firm A's stores. Due to strong holiday, weekend and idiosyncratic shocks, plus the lumpy nature of sales, 55.4% of four-day$\times$store bins in our data have an average output that is more than 25 percent below their store's mean. It is hard to imagine employees deciding to end their career at a Firm in response to such common, transitory output fluctuations. Second, consistent with the idea that workers' departure decisions should depend more on predictable, longer-term productivity trends than on temporary shocks with low information content, Appendix Figure A.1.6(a) shows strong seasonal patterns in both sales and departures from Firm A. The illustrated trends are consistent with the idea that rational employees of Firm A time their quits to avoid working at Firm A during low-sales months like March – April and July – October. Put another way, the team productivity information that is most relevant to when most workers should leave Firm A is already available well in advance, from the seasonal cycle.[23]

Our final piece of evidence that the sales declines around the mandated notice and departure

[21] Whether these pre-departure losses represent on-boarding versus, for example, reductions in employee morale depends on whether they coincide with hiring. We explore this in the next two sections.

[22] Another reasonable concern is that departing employees might choose to take rest-days near the departures, and that will make the actual departure dates occur earlier than our data indicate. This concern does not apply in our case because the date of departure observed is technically the last day when the departing employee works at the site.

[23] A more surprising feature of Appendix Figure A.1.6(a) is that many of the departures occur in December, even though monthly performance is relatively high. Store managers we interviewed stated that this is due to geographical moves back to employees' hometowns from the cities in which most of Firm A's retail outlets are located. Many news stories have documented that migrant workers return home for the Chinese New Year and how labor shortages after the New Year are getting worse in recent years. For example, see http://usa.chinadaily.com.cn/epaper/2013-02/01/content_16194037.htm

dates are not the cause of the decision to quit is the strong bunching of departures from Firm A on the last day of the calendar month. This is illustrated in Appendix Figure A.1.6(b), which plots histograms of the day-of-the-month on which departures occur, along with the average daily sales per store. More than half of the departures occur on the last day of a month, implying that announcement dates are highly bunched in the middle of a month. This pattern strongly suggests that both departures *and* announcements are planned well in advance, rather than being responses to daily productivity fluctuations. Indeed, if employees were choosing a specific day to depart based only on daily sales patterns, it would make considerably more sense to depart a few days after the first of the month, to take advantage of the higher-than-normal sales that are typical at that time. Instead, store managers attribute the concentration of departures at the end of a month to employees' desires to synchronize their departure date with the store's monthly pay cycle.[24]

3.3 Disaggregated Results

To explore the mechanisms underlying the productivity losses identified in Table 1, Table 2 replicates that analysis for sub-samples of departures defined by when the departing worker was replaced.[25] To isolate short-staffing effects, column 1 restricts the sample to the 25 percent of departures that result in *late refills*, which are the only departures that generate a period of short-staffing. Here we observe a significant productivity loss associated with short-staffing that is larger and longer-lasting than in the full sample. Specifically, in the first six days following the departure, team performance declines by about $(263 + 164 + 148)/3 = \$192$ per day, which amounts to a 32.4% reduction.[26] Perhaps surprisingly, these short-staffing costs are quite similar in size to our estimated AN and BD effects, which occur before the departure while the team composition remains intact. Finally, starting in the second week after the departure, we find that teams do not perform worse than in the control period, even though (by construction) all teams in this estimation sample are still operating with one less employee. We discuss how such short-staffed teams might be able to produce at full-staffed levels for short periods of time in Section 3.4.

To isolate on-boarding costs, column 2 focuses on the 17 percent of departures that are replaced exactly "on time", i.e. where the new employee joins the team the day after the departing employee's last day of work. As noted, this is by far the modal day – relative to the departure – that refills took place. These post-departure productivity trends represent pure on-boarding effects because there are no short-staffing costs, and all of the on-boarding all takes place *after* the AN and BD periods. In the six days right after these departures, column 2 shows an acute loss of

[24]Recall that Firm A's workers are paid on the 20th of each month for the *previous* calendar month. Therefore, any work done after the last day of a month becomes part of a new pay cycle, which will not appear in a paycheck for at least 48 days.

[25]Motivated by Table 1's results, Table 2 displays only the coefficients for days -16 through 6. Of the 57 coefficients that are outside this time range, just one is statistically significant in conventional tests, and none are significant after adjusting for multiple hypothesis testing.

[26]To remove the influence of on-boarding costs on these estimates, our estimation sample discards all post-departure observations which occur after a replacement worker is hired. Thus, by construction team productivity levels after an employee's departure are associated with short-staffing only.

team productivity (about 33.2% of team sales) that last about a week. Overall, these estimated on-boarding costs are surprisingly similar to the costs of operating short of one team member, identified in column 1. In addition to lower firm-specific skills possessed by the new hires, these costs could also include a team-specific component: It may simply take some time for a team to learn how to work with a new member (Bartel et al., 2014).[27]

Finally, to isolate productivity changes that are unrelated to concurrent changes in team composition, column 3 focuses on the 60 percent of departures in which the replacement worker arrived strictly after the departure occurred, (i.e. the union of groups *b*, *c* and *d* in Figure 1). These teams remain intact during the entire pre-departure period, with no arrivals or departures occurring. Strikingly, column 3 continues to show strong reductions in team productivity during that period. Specifically, there is a productivity loss of about (100 + 138)/2 = $119, or 20.1% in the four days surrounding the announcement (AN). In the employee's last six days on the job (BD), team productivity declines by about 21.7%. In Section 4 we explore some possible sources of these effects, including time lost due to recruitment activities by the incumbent workers, and short-timer effects on employee effort before the departure date.[28]

3.4 Parametric Event Study

The results in Table 2 suggest the presence of short-staffing costs during a brief period after a departure, of on-boarding costs during a similarly brief period after a new employee joins the team, and of potentially morale-based output reductions that are unrelated to any concurrent changes in team composition. Motivated by those findings, we now conduct a parametric event study in our entire sample of departures that regresses daily team output on explicit indicators for whether a new employee is on-boarding, whether the team is operating short-staffed, plus indicators for periods around the notice (AN) and before the departure date (BD). Our estimation framework –which restricts the potential duration of on-boarding and short-staffing effects to specific lengths– allows all these processes to occur simultaneously. We therefore estimate:

$$\begin{aligned} S_{it} =& \beta_1 \cdot AN + \beta_2 \cdot BD + \beta_3 \cdot OB + \beta_4 \cdot LOB + \beta_5 \cdot SS + \beta_6 \cdot LSS \\ &+ \beta_P \cdot P \\ &+ \gamma_1 \cdot D_t + \gamma_2 \cdot I_i + \epsilon_{it}, \end{aligned} \tag{2}$$

[27]Notably, the on-boarding costs that are identified in column 2 refer specifically to the cost of on-boarding a new worker *in the absence of the departing worker*, who has already left. On-boarding costs could, of course, be different if the old and new workers overlap for some time at the firm, for example because useful information may flow from the former to the latter. In Section 4, we present on-boarding cost estimates that distinguish between these two situations.

[28]Also of interest in this sample is the fact that we detect no productivity declines after the departure. To make sense of this, recall that short-staffing is rare even in this sample: 17 percent of departures are replaced exactly *on time*, and 18 percent of departures are from teams of fewer than four workers, where vacancies are filled with *temporary replacements*. Thus, only 25/60 = 42 percent of these departures resulted in any short-staffing. In addition, on-boarding periods are spread throughout the post-departure period, making their effects hard to detect.

where the dependent variable S_{it}, fixed effects D_t and I_i, and error term ϵ_{it} are all defined as in equation 1. AN is an indicator for the *M* days starting on the day before the required notice date (day -15, to capture early announcements), while BD is an indicator for the departing worker's last *M* days of work (-(*M-1*) through 0). OB is an indicator variable identifying the first *M* days in which a new employee was added to the team; SS identifies the first *M* days when a team is operating short-staffed with vacancies unfilled.[29] Coefficients on these two variables capture immediate on-boarding and short-staffing effects. To examine the persistence of on-boarding and short-staffing effects, we include two more indicator variables: a late on-boarding variable LOB identifies the (*M*+1)th through 14th days after the hiring, and a late short-staffing variable LSS identifies the (*M*+1)th through 30th days that a team is operating short one member. Finally, the coefficient vector $\beta_P \cdot P$ divides the portions of the 61-day treatment window that are not included in the AN and BD effects into bins, most of which are four days in length.[30] These period effects will measure any causal effects of turnover that are not captured by our AN, BD, OB and SS effects.

In Table 3, we present estimates of equation 2 using *M*=4, with period effects controlled in column 2. Estimates of the on-boarding coefficients indicate that team productivity falls by about 76 dollars or 12.9% in the first four days following a new employee's entry, but this loss does not extend beyond those first four days. Similarly, the estimates indicate that short-staffing is costly at first, leading to around a 36.9% reduction in team output in the first four days, but it also dissipates quickly. This suggests that teams quickly adapt to being short-staffed, and find a way to at least temporarily maintain their normal output, even while operating with one less employee.[31] In the four days surrounding the notice (AN) and in the four days before the actual departure (BD), we see significant reductions in team output, amounting to about 19.0%. As noted, these AN and BD coefficients measure changes in team performance that cannot be linked to current or recent changes in composition of the sales team. Finally, none of the remaining four-day bin variables in column 2 are statistically or quantitatively significant, suggesting that our AN, BD, SS and OB coefficients capture all relevant effects of turnover on team productivity. While these costs are statistically and

[29] A potential concern in estimating equation (2) is that our estimates of short-staffing costs apply only to subset of departures that failed to find a replacement worker before the departure date, which is an endogenous outcome. For example, suppose that certain key workers are harder to replace, and that teams perform especially poorly when operating short of such workers. Then our estimates of short-staffing costs will overestimate the costs that would occur if worker replacement occurred randomly (because the short-staffing spells that appear most frequently in the data are the most costly ones). Notably, however, our goal is not to identify the causal effect of a randomly-assigned short-staffing spell, but to measure the costs of the short-staffing spells *that actually result from a typical departure*. For our purposes, it would therefore be inappropriate to control for non-random selection into short-staffing.

[30] When *M* =4, the 16 day 'notice' interval between days -15 and 0 divides evenly into four periods of four days each; thus all the bins in *P* are four days long, except at the boundaries of the treatment window (where they are truncated at two or three days). When *M* takes on other values we shorten or lengthen the two bins in the middle of the notice interval to accommodate the longer or shorter AN and BD effects. See Table A.2.4 for details.

[31] We can think of three reasons why short-staffed teams return so quickly to baseline productivity levels even without hiring replacement workers. First, Firm A informs us that in longer-term short-staffing situations, the remaining employees will temporarily expand their work shifts and postpone their rest days to ensure the store's normal operations. Second, in this team-based pay environment, remaining employees will experience immediate income losses if staff shortages reduce store sales, so they may increase their efforts to avoid such losses. A third possibility is endogenous timing of replacements: The short-staffing spells that last the longest might be the ones that are least costly to the team.

quantitatively significant, they are however quite short in duration.

4 Robustness and Heterogeneity

We conducted a variety of robustness and heterogeneity analyses of our main results. This Section provides a brief overview of these analyses; additional details and supporting documentation are provided in Appendix A.2 and Appendix A.3 respectively.

4.1 Robustness to Specification Changes

One modeling choice that could affect our results is the width of the window around the departure date within which we allow departures to affect team productivity. In all the results presented so far, this window included 30 days on either side of the departure date. We replicated our aggregate analysis (Table 1) for windows of 50, 90 and 120 days, and found very similar results: Significant productivity losses only occur surrounding the required announcement, before the departure, and shortly after the departure. In addition, we find no pre-trends in team performance even with the widest window, which supports our main identifying assumption– that the causal effects of a departure are confined to a relatively short time period surrounding the departure.[32]

A second modeling choice refers to the assumed duration of the four main effects estimated in Table 3's parametric analysis. There we set a duration, M, of four days for the AN, BD, early OB, and early SS effects. To assess the consequences of this assumption, we replicated Table 3 for M=2, 3, 5 and 6. Overall, the estimated effects on daily productivity are numerically larger if we assume a shorter duration, and shrink in size if we impose a longer duration, as we would expect if the underlying causal effects were relatively short-lived. Otherwise, the effects are statistically robust and exhibit similar patterns to the main specification. An alternative way to represent the time structure of team productivity trends during the notice period is with a continuous quartic in time. Consistent with our main, binned approach, this yields substantial and statistically significant productivity losses near the notice and departure dates, and small, insignificant losses in the middle of the notice period.

Since more than half of the departures in our data occur on the last day of a calendar month, we were also concerned that our estimated announcement (AN) and before departure (BD) effects might be picking up sales patterns associated with those parts of the month. While our day fixed effects should already account for such influences, day fixed effects constrain the baseline time pattern of sales to be the same across all stores. To address this issue a different way, we estimated AN and BD effects in a sub-sample that excludes departures in the last three days of a calendar month. They were very similar to our main estimates. We also conducted a placebo test, which codes the same day of the third month *before* the actual departure date as a placebo departure. The goal is to test whether low sales just happen to occur on the days of the month when notices are

[32] In Appendix Table A.2.4, we push even further by doubling the length of the pre-treatment window associated with our W=120 experiment from 60 to 120 days. (In other words, we look at 120 days before and 60 days after that departure date.) We find no pre-trends even if we trace team performance to two months before the departure.

typically given and when departures typically occur. The results showed no estimated effects of these placebo dates.

Finally, to assess the possibility that different day-of-week, seasonal and holiday patterns in different stores could somehow be affecting our results, we re-estimated Table 3, replacing the 731 day effects by 24 month effects (to capture aggregate trends). In addition, we interacted seven day-of-the week effects, ten holiday effects, and twelve month-of-the-year effects with the store fixed effects, to allow each store to have its own pattern of sales over the week and year, and to be differently affected by holidays from other stores. While this reduced the precision of some of our estimates, the results were mostly unchanged.

4.2 Heterogeneous Effects

To shed additional light on the mechanisms underlying the four types of turnover-related productivity losses (AN, BD, OB and SS), we explored how the magnitude of these effects varies with the characteristics of the departure and the sales team. We began by asking whether any of these effects depended on *who is leaving* the team. For example, it may be harder for teams to adjust to the departure of a manager or a highly experienced colleague than another team member. We found that none of these productivity reductions were statistically different when the leaver had more tenure at Firm A, but that the departure of a manager led to a much larger productivity decline right after the departure announcement than did the departure of other workers. To the extent that employees need more time to find a new manager than a new worker, this supports our interpretation of the AN effect as due to recruitment activities of the incumbent workers. We also found a large but imprecisely-estimated additional cost of short-staffing when a manager is leaving ($222 or 38 percent, p=0.11).[33] In sum, at least in our context, the loss of a manager appears to be more costly to a team than the loss of a senior member of unspecified rank. Related, we also asked whether any of these productivity effects depend on *who is staying*, i.e. with the tenure of the workers or manager who are left behind. None of these factors had a detectable effect.

Next, we asked whether a team's previous experience with handling turnover events might alleviate (or exacerbate) turnover costs. To do this, we first classified stores according to the total turnover they experienced during our sample period, and found that teams with an above-median turnover rate had significantly lower short-staffing costs than other stores ($101 versus $381), while the other costs were not significantly different.[34] The same pattern was found for stores that experienced a turnover during the past three months. In this case, the team seems to produce just as much sales while it is short-staffed than while operating with a full complement of workers. Together, these results suggest that teams can learn to manage turnover costs better with experience, in particular by better coordinating work schedules to prevent sales losses while the team is short-staffed.

[33] The remaining interactions are small and statistically insignificant.

[34] During our analysis period, the median store experienced two turnovers, so stores that experienced more than 2 departures are identified as high-turnover stores.

Another store characteristic that might influence turnover costs is how busy the store happens to be. Indeed, all the sources of turnover costs we have described –shirking, recruitment activities, training replacement workers, and operating with fewer workers– should have a lower opportunity cost when there are few, or no customers are available to serve. To see if this is the case, we asked whether high-performing stores (ones that consistently had above-average sales for their size) experienced higher turnover costs, and whether turnover costs within a store are higher on days that are usually much busier than others (specifically Fridays, Saturdays, Sundays and holidays). Consistent with our hypothesis, we found much higher costs in busy stores, and on days when stores are most likely to be full of customers. These additional costs were experienced across all four cost components (AN, BD, OB and SS). Together, these results speak to a measurement problem that arises in many field studies of worker and team performance: how to account for variations in *slack time* during which there is no work available to be done. Our results suggest that using plausibly exogenous variation in work flow associated with location and timing of business activity can help isolate the circumstances in which the team's available human capital and work effort do, indeed, affect its performance.[35]

As noted, our estimates of on-boarding costs could represent a lack of firm- or industry-specific skills among the new hires, or pure adjustment or integration costs associated with learning to work with a different person. To distinguish these scenarios, we asked whether on-boarding costs were higher when the replacement worker came from outside Firm A, compared to trained, temporary replacements from the host institutions or internal hires from other retail stores of Firm A. Somewhat surprisingly, we find that on-boarding effects are no different for external hires. This result mirrors a finding in Bartel et al. (2014), who show that merely changing the identity of a teammate – even when she has prior experience with the same employer – generates a temporary reduction in team productivity. Thus, at least in our relatively low-wage context, the integration of a new worker into a team (regardless of her qualifications) and not the acquisition of new occupation- or firm-specific skills appears to be the main source of on-boarding costs.

To examine how new workers learn their jobs at Firm A, we asked whether it is less costly to on-board a new worker before versus after the departing worker has left the firm. If so, then requiring workers to provide advance notice has an additional benefit beyond reducing the number of short-staffing spells: on-boarding in the presence of the departing employee could allow the leaver to transfer her knowledge and skills to the newcomer. We do not detect a difference in these costs, however, suggesting that the knowledge embodied in a single team member is not especially valuable to a new retail sales employee, even in our context of very small sales teams. The other remaining employees apparently have enough institution-specific knowledge to make the continued presence of the departing worker irrelevant.

Based on our discussions with Firm A, we have interpreted the productivity declines just before the departure (BD effects) as short-timer, or morale effects. It is conceivable, however, that this decline represents a final burst of recruiting activity by the incumbent workers. To address this

[35] See Lazear et al. (2016) for an interesting discussion of the slack time in a service context that is similar to ours.

concern, we interacted the BD coefficient (an indicator for the last four days the leaver is still working) with an indicator for whether her replacement had already arrived (i.e. the replacement started before day -4). While the point estimates suggest that the pre-departure productivity loss is slightly smaller in those circumstances, we cannot reject that the BD effect is the same regardless of whether a replacement has already been hired. More importantly, we find strong, significant BD effects in teams who replaced their departing worker more than four days before the end of the notice period. This suggests that morale-based factors, and not last-minute recruiting activity, are the primary drivers of the productivity losses we observe during the departing employee's last few days on the job.

A final source of information about the mechanisms underlying our estimated AN and BD effects is their interaction with team size. For example, if the short-timer effects captured by our BD coefficient represent withdrawal of a fixed amount of effort by the departing employee only, that coefficient should not vary with team size. If shirking by the departing worker is contagious to her teammates, however, short-timer effects can increase with team size.[36] Similar arguments apply to the AN effects: If –as Firm A believes– the AN effects represent the costs of recruiting a single new employee, it is not immediately obvious that these costs should increase with team size. However, recruiting costs could increase with team size if there were diseconomies of scale in recruiting. For example, if all employees must be consulted (for example, if they all must attend the same meeting to discuss candidates and co-ordinate schedules) then interactions between team members would cause recruiting costs to rise with team size. To distinguish these scenarios, we study the effects of team size on AN and BD effects in Table 4.[37]

The results in Table 4 suggest that both the AN and BD effects increase with team size. In teams of two employees, no AN or BD effects are detected. In teams of three employees, effort reductions are sizable and statistically significant, and the magnitudes are even larger in teams of four or more employees. Both these patterns suggest strategic complementarities between the actions of team members: contagion in shirking, and co-ordination costs in recruiting.[38] Notably, both these processes are consistent with the shift arrangements at Firm A, described in Appendix Figure A.1.4: In stores of two employees, the coworkers do not interact during operating hours; in stores of three employees, on average the departing employee overlaps with another employee for $\frac{2}{3}$ of her working time, while in stores of four or more employees, the departing employee will always be paired with at least one other employee.

[36] Previous research that detects effort contagion in small teams includes Mas and Moretti (2009) and Ichino and Maggi (2000). Another way that shirking by the departing worker can be magnified in larger terms is if other team members (optimally) reduce their monitoring of a teammate when they learn she is about to depart Kandel and Lazear (1992).

[37] Since actual team size varies during the turnover process, we classify teams by their target size.

[38] A simpler reason why productivity reductions might be greater in larger teams would be if per-worker productivity was higher in larger teams. The loss of one member's time –due to either shirking or recruiting activity– would then reduce output more in larger than smaller teams. Per-worker sales, however, are actually somewhat smaller in larger teams (see Panel C of Appendix Table A.1.1).

5 Assessing Cost Magnitudes

To assess the economic significance of the turnover-induced productivity declines we have estimated, we start by multiplying each of the four effects on daily sales (OB, SS, AN and BD) by its expected duration and summing the results. This gives us the total lost sales caused by a single departure from Firm A; using our preferred regression specification (column 2 of Table 3), this equals $1,654.[39] Interestingly, all four sources of productivity losses (on-boarding, short-staffing, around-notice and before-departure) contribute about equally to this total (at $354, $400, $448 and $452 respectively).

Two aspects of the composition of these total losses deserve some comment. One is that – in large part because short-staffing spells are so rare– only 24 percent of the lost output associated with a departure ($400/$1654) is associated with short-staffing. This finding may have implications for equilibrium search models: By conceptualizing the cost of posting a lower wage as a lower probability that a job is occupied at any particular moment, Burdett-Mortensen-type search models implicitly assume that short-staffing is the only source of turnover costs. Second, almost two thirds (63%) of the productivity losses associated with turnover are incurred before the departing worker leaves: these comprise AN and BD costs, plus 40% of OB costs (the share of replacements that arrive strictly before the departing worker leaves). As noted, this finding illustrates the value of having information on announcement dates in turnover cost studies: Without this information, it would be impossible to know whether dips in team productivity before the departure are the economic shocks that caused the departure, or –as is the case here– behavioral responses to an anticipated departure.

Having quantified the magnitude and makeup of the output losses associated with turnover, our next task is to assess their overall economic significance: According to our estimates, how important is employee turnover to Firm A's bottom line? To answer this question, we must first convert our estimated reductions in gross sales into net sales, then incorporate any additional costs of turnover that do not take the form of team productivity reductions. Turning first to the gross-to-net sales conversion, we consider a situation where Firm A experiences one extra departure in a particular year, leading to a $1,654 decline in gross sales. After that, the store returns to its original team size and sales. According to Firm A, the main variable costs that would be saved by this sales decline involve manufacturing and transporting the sold items to the store. Since these costs are 35 percent of gross sales, we calculate the net revenue costs of a single departure as $\$1654 \times 0.65 = \$1075$.[40]

[39] Detailed calculations underlying all the results in this Section are provided in Appendix A.4; here we provide an overview of the main steps and intuition only. Appendix A.4 also assesses the robustness of our total lost sales estimate ($1654) to alternative assumptions about effect duration, M. The largest estimate, for M=6, is 7.8 percent higher than our baseline estimate.

[40] While clothes are (obviously) produced before they are sold, essentially all items produced by Firm A's factories are eventually sold at its retail outlets (sometimes at discounts, which are reflected in our sales statistics.) Thus, a turnover-induced shortfall in sales in one month is eventually reflected in fewer shipments from the factories in a later month.

We next account for three additional components of total turnover costs: recruiting time spent by Firm A's regional managers, the extra salary costs that are incurred while the departing and replacement employee overlap in the firm, and the *savings* in wage costs while a team is operating short-staffed. Combining these components –all of which are small– with our net sales reductions yields a total cost of $1,132 per departing employee, which works out to 9.4 days' worth of lost per-employee net sales or 1.1 percent of an employee's net sales over a typical 2.3-year with Firm A. In terms of wages, $1,132 corresponds to 63 days of a typical employee's pay, or 7.6% of wages over the career. Importantly, the dramatic difference between net sales-denominated and wage-denominated turnover costs does not necessarily imply that Firm A is exploiting its sales employees: Instead, the main source of the difference is the low share of wages in total costs, due in large part to the high costs of renting retail store space.

A final way to express the magnitude of our estimated turnover costs is via the trade-off between the rate of turnover and the firm's (permanent) wage rate. Specifically, by how much could Firm A afford to raise wages (keeping profits constant) if it was able to cut turnover by 10 percent? According to our estimates, Firm A should be willing to raise its total compensation per worker by up to 0.6 percent of its annual wage bill for this turnover reduction. Put another way, a 10 percent increase in turnover is as costly as a 0.6 percent increase in wages.

Comparing our estimates of total turnover costs to other cost estimates is challenging, given the paucity of existing evidence. For example, descriptive employer survey data (Barron and Bishop, 1985) indicate that U.S. employers spent an average of 9.87 hours of incumbent employee time to hire a single applicant; this number was 7.25 and 10.60 hours in the retail industry and sales occupations respectively. These numbers are broadly similar to our estimate of regional managers' recruiting time per hire of 1.5 days, but our estimates also show that recruiting activities are only a small fraction of total turnover costs. Similar survey results in Barron and Black (1997) indicate that employer-reported days of training per new hire vary from 6 to 19, but we do not know how productive both the trainer and trained employees are during those days, or how long the new employees are likely to remain with the firm.

Bartel et al. (2014) conduct a study similar to ours using monthly data on the performance of registered nursing teams. One of their most intriguing findings is that replacing a nurse with another equally qualified one from within the same hospital reduces team productivity. Thus there appear to be costs associated with team *disruption* (changes in the identity of a member) that are unrelated to the team's measurable human capital. While direct comparisons between our cost estimates and Bartel et al.'s are hard to make –their main team performance measure is very different from ours– we can more confidently compare the *relative* costs of team disruptions versus human capital shortages between two papers, in terms of each paper's own performance measure. In Appendix A.4.6, we estimate that replacing ten percent of a team's workforce with a different employee from within Firm A each month (with no short-staffing spells) is about 23 percent as costly as *permanently* operating with ten percent fewer workers. In Bartel et al.'s nursing environment, that number is about 35 percent. Thus, both studies underscore the role of team disruption in turnover costs: Regularly changing the membership of workplace teams has sizable and broadly similar effects on team performance in these two very different workplaces.

We conclude by exploring the implications of Firm A's turnover costs for two key strategic decisions it faces. The first is to ask whether these turnover costs could be large enough to make raising wages a profitable strategy. In this regard, it follows directly from our cost estimates that a wage increase at firm A would pay for itself (in terms of reduced turnover costs) only if the elasticity of quits with respect to wages, $\eta < 0$, was at least .100/.006 = 16.8 in absolute value.[41] This seems extremely unlikely: Manning's (2003) median estimate of η is -0.5, and the largest well-identified estimates of η we know of is -3.5, from studies of schoolteachers by Falch (2010). In a study of a large U.S. retailer using quasi-experimental wage variation, Dube et al. (2018)'s estimate is statistically indistinguishable from zero.[42] Our own best estimates of η at Firm A, identified from fluctuations in total monthly pay generated by the seasonal sales cycle, yield an elasticity of about minus one. We conclude that turnover costs, on their own, cannot explain the substantial firm-wage effects that have been documented within industries, including the retail sector Ton (2012). Other factors, such as rent-sharing by employers, differences in labor quality, and efficiency wage effects (i.e. causal effects of higher wages on worker *productivity*) are needed to explain these effects.

The second strategic decision we consider is the firm's reservation worker quality: Given the total costs of replacing one worker with another (C = \$1,132), how much less productive must a worker be to make it profitable for Firm A to replace her? To answer this question, we consider an employee working at an average-size team, in which every worker makes average net sales of q = \$120. At the midpoint of an employee's average (829-day) completed career, the firm realizes that the worker is d percent less productive than an average worker. The value of d at which Firm A should be indifferent between keeping and replacing her then satisfies $C = Tqd$, where T = 414 days is her remaining tenure at Firm A if she is not replaced now. It follows that d = 2.3 percent. In other words, with these turnover costs (which we have calculated from voluntary departures), it would pay to replace all mid-career employees whose productivity was more than 2.3 percent below the average.[43]

While this very strict retention threshold illustrates the low magnitude of our estimated turnover costs, we caution that it does not imply that Firm A should *dismiss* all mid-career workers who fall below it. In part, this is because a number of factors –including legally required severance pay (which raises d to 4.4 percent), noisy worker productivity in team production, and the potential

[41] We estimate that a 10 percent increase in turnover is as costly as a 0.6 percent increase in wages. Thus the critical elasticity at which a wage increase pays for itself is given by -.100/.006 = -16.8.

[42] Manning's estimates of the quit-wage elasticity come from a comprehensive study using data from the PSID, NLSY, BHPS and LFS. His estimates range from -.156 to -1.010 and are centered around 0.5. Portugal and Cardoso (2006) and Dube et al. (2011) also estimate small negative effects of *minimum* wage increases on separations, but these identify a different (market equilibrium) parameter from the number we need here.

[43] We assess how our critical absolute quit elasticity, $|\eta^*|$, and employee retention threshold, d^* would differ in other types of labor markets in Appendix A.4.9. We consider two alternative scenarios: tighter labor markets (where short-staffing spells are twice as common), and more skilled jobs where on-boarding costs are twice as high. Neither change has large effects: Tighter labor markets reduce $|\eta^*|$ from 16.8 to 14.8, and raise d^* from 2.3 percent to 2.6 percent. In the higher-skills scenario, $|\eta^*|$=13.9 and d^*=2.7 percent.

extra morale costs of dismissals– could make firm-initiated separations considerably more costly than worked-initiated ones. More importantly, this is because Firm A already has a better way to shed less-prodictive workers: Firm A's pay policy, which has a large variable component based on team sales, provides substantial incentives for under-performing salespeople to quit, and for their team-mates to encourage them to quit. This policy, combined with Firm A's announced policy of no firm-initiated layoffs, effectively delegates most employee discipline and selection problems to the workers and teams themselves.

6 Summary and Discussion

This paper has estimated the effects of employee turnover on team productivity in a retail sales context. Our unique access to daily productivity and staffing data and the employer's advance notice requirement for quits enable us to address well-known concerns regarding endogenous turnover, and allow us to identify causal effects of turnover that occur before as well as after the employee leaves the team. We find robust evidence of productivity losses associated with the on-boarding of new employees and with operating short-staffed, as well as around the time a worker gives her notice, and just before she leaves. We attribute the latter two losses to the involvement of sales team members in recruiting the replacement worker, and short-timer effects on worker morale. While all these estimated productivity effects are highly statistically significant and represent large shares of daily sales, they are, however, very short in duration.

The speed with which our retail sales teams adjust to the departure of a member means that, overall, turnover is not very costly in this environment: a ten percent increase in the turnover rate is as costly to our firm as a 0.6 percent increase in wages. These low costs should perhaps not be surprising given the relatively routine nature of retail sales jobs, and given the fact that employers can mitigate turnover costs in a number of ways we have documented in this paper. Indeed, we argue that turnover costs are low at Firm A, in part because it has developed an HRM system of complementary personnel policies that allow it to operate effectively in a high-turnover environment. These policies include an advance notice requirement for workers that allows Firm A to hire early enough to eliminate most short-staffing spells, and access to an experienced, contingent workforce to fill many of the remaining staffing gaps. Firm A's strategy also includes a recruitment process that engages incumbent workers in hiring, and a no-layoffs policy that effectively delegates many employee discipline problems to the sales teams themselves. These team activities, in turn, are incentivized by Firm A's team-based commission scheme, which motivates team members to recruit and train high-quality workers quickly, and to encourage underperforming workers to leave when necessary. Taken together, these mutually-reinforcing policies may help us understand how many firms can operate successfully in labor markets characterized by high turnover, including retail, customer service, hospitality, restaurants and similar sectors.

Table 1: Non-Parametric Event Study – Full Sample

Dependent variable: daily sales (in US $)					
Before the Departure:			**After the Departure:**		
	β	SE		β	SE
$P_{-30,-29}$	70	(97)	$P_{1,2}$	-131*	(76)
$P_{-28,-27}$	10	(80)	$P_{3,4}$	-59	(63)
$P_{-26,-26}$	6	(70)	$P_{5,6}$	-58	(48)
$P_{-24,-23}$	-56	(42)	$P_{7,8}$	27	(64)
$P_{-22,-21}$	-59	(39)	$P_{9,10}$	-38	(51)
$P_{-20,-19}$	-4	(63)	$P_{11,12}$	57	(69)
$P_{-18,-17}$	-0	(62)	$P_{13,14}$	54	(63)
$P_{-16,-15}$	-78**	(37)	$P_{15,16}$	-8	(46)
$P_{-14,-13}$	-128***	(31)†††	$P_{17,18}$	-4	(47)
$P_{-12,-11}$	-54	(45)	$P_{19,20}$	-18	(67)
$P_{-10,-9}$	47	(63)	$P_{21,22}$	-10	(68)
$P_{-8,-6}$	-43	(32)	$P_{23,24}$	-47	(65)
$P_{-5,-4}$	-35	(40)	$P_{25,26}$	-4	(47)
$P_{-3,-2}$	-94***	(31)††	$P_{27,28}$	-16	(42)
$P_{-1,0}$	-153***	(35)†††	$P_{29,30}$	12	(102)
N			68238		
H_0: $P_{-30,-29}$ = ... = $P_{-24,-23}$ = 0					F(4, 117): 1.17, p-value: 0.33
H_0: $P_{19,20}$ = ... = $P_{29,30}$ = 0					F(6, 117): 0.28, p-value: 0.95
H_0: $P_{-30,-29}$ = ... = $P_{-24,-23}$ = $P_{19,20}$ = ... $P_{29,30}$ = 0					F(10, 117): 0.84, p-value: 0.59

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$); † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$.

Notes: This table presents results from estimating equation 1. $P_{-14,-13}$ identifies the first two days of the required notice period; $P_{-1,0}$ identifies the departing employees' last two days of work. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Table 2: Non-Parametric Event Study – Sub-samples

	Dependent variable: daily sales (in US $)			
	(1) Late refills	(2) On-time refills	(3) Intact pre-dep teams	
$P_{-16,-15}$	-156***	-70	-100**	**Around Notice (AN)**
	(59)	(89)	(40)	
$P_{-14,-13}$	-165**	-143*	-138***	
	(69)	(84)	(40)†††	
$P_{-12,-11}$	-43	-34	-94	
	(108)	(89)	(58)	
$P_{-10,-9}$	235	-49	97	
	(199)	(61)	(95)	
$P_{-8,-6}$	-80	-69	-66	
	(73)	(83)	(44)	
$P_{-5,-4}$	-56	-39	-89*	
	(85)	(85)	(50)	
$P_{-3,-2}$	-59	-144*	-93**	**Before Departure (BD)**
	(68)	(85)	(45)	
$P_{-1,0}$	-200***	-263**	-203***	
	(63)†	(103)	(45)†††	
$P_{1,2}$	-263*	-203**	-132	**After Departure**
	(135)	(98)	(116)	
$P_{3,4}$	-164**	-200**	-82	
	(64)	(97)	(74)	
$P_{5,6}$	-148*	-186*	-101	
	(87)	(96)	(70)	
N	60496	59984	62787	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$); † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$.

Notes: Columns (1)–(3) in this table present results from estimating equation 1 for late refills, on-time refills, and intact pre-departure teams, respectively. Two-day bins for days -30 to -16 and days 6-30 are included in the regression. Column (1) includes departures that result in at least one day of short-staffing, and we only include observations that occur *before* replacement workers are hired. Column (2) includes only departures that are replaced on the day after the departing employee's last day of work. The sample in column (3) excludes departures that are refilled early, so the team size and team composition remain intact before the actual departure occurs. In all columns, regression is estimated with 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Table 3: Parametric Analysis of Team Daily Output (*M*=4)

	Dependent variable: daily sales (in US $)	
	(1)	(2)
Around Notice_4 ($P_{-15,-12}$)	-111^{***}	-112^{***}
	$(31)^{\dagger\dagger\dagger}$	$(32)^{\dagger\dagger\dagger}$
Before Departure_4 ($P_{-3,0}$)	-114^{***}	-113^{***}
	$(25)^{\dagger\dagger\dagger}$	$(26)^{\dagger\dagger\dagger}$
Early On-boarding_4	-79^{**}	-76^{**}
	$(32)^{\dagger}$	(36)
Late On-boarding_4	-2	-5
	(36)	(35)
Early Short-staffing_4	-240^{***}	-218^{*}
	$(86)^{\dagger\dagger}$	(111)
Late Short-staffing_4	-38	-41
	(54)	(55)
Other Period Effects		Yes
N	75801	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$

Notes: This table presents results from estimating equation 2. *Around* $\textit{Notice}_4$ identifies a 4-day period including the required notice date. *Before* $\textit{Departure}_4$ identifies the last 4 days the departing employee spending at the team. *Early* $\textit{On-boarding}_4$ identifies the first 4 days a new employee is present, and *Late* $\textit{On-boarding}_4$ identifies the 5th day through two weeks after the hiring. *Early* $\textit{Short-staffing}_4$ identifies the first 4 days following an unfilled departure, and *Late* $\textit{Short-staffing}_4$ identifies other days following *Early* $\textit{Short-staffing}_4$ with the departure unfilled through 30 days following the departure. "Other Period Effects" refer to (mostly) four-day bins (P) during the notice period other than the Around-Notice and Before-Departure periods. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Table 4: Heterogeneity Examination – *Team Size*

	Dependent variable: daily sales (in US $)
	(1)
size=3	11
	(19)
size≥4	15
	(57)
$AN_4 \times$ (size=2)	32
	(74)
$AN_4 \times$ (size=3)	-123**
	(50)††
$AN_4 \times$ (size≥4)	-154***
	(47)†††
$BD_4 \times$ (size=2)	3
	(56)
$BD_4 \times$ (size=3)	-121***
	(43)††
$BD_4 \times$ (size≥4)	-151***
	(43)†††
N	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$
Notes: This regression includes controls for early and late on-boarding effects, and for early and late short-staffed effects, as defined in equation 2. We interact the AN and BD effects with team size. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Figure 1: Flow Chart of Notice, Departure, and Replacement Procedures

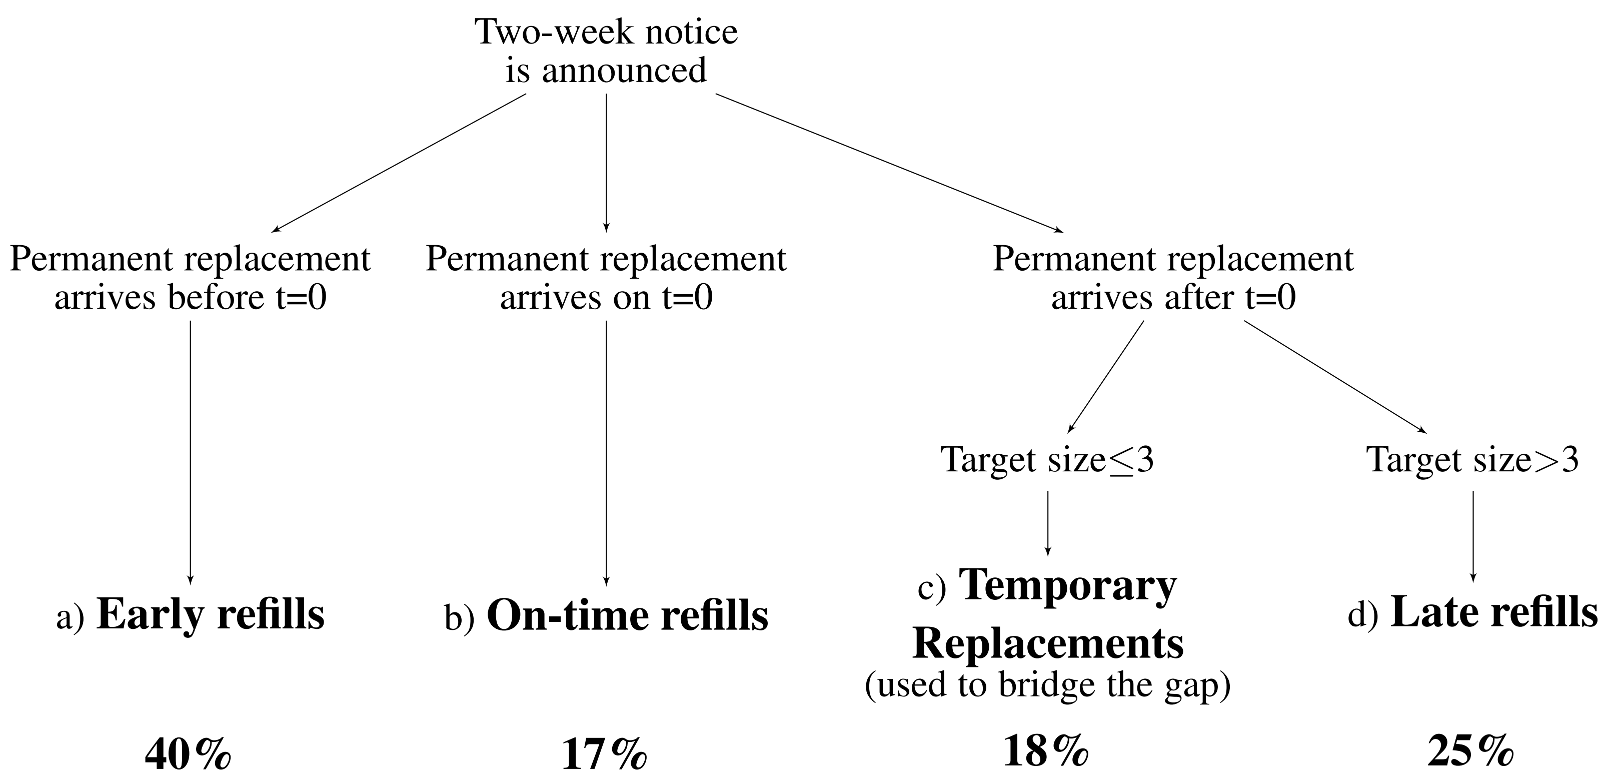

Notes: The procedures in this figure apply to departures from stores that are at their target size. In these stores, group (d) is the only one that experiences short-staffing. In groups (b), (c) and (d), the entire sales team remains unchanged throughout the two-week notice period.

Figure 2: Share of Departures Replaced as a Function of Time since Notice

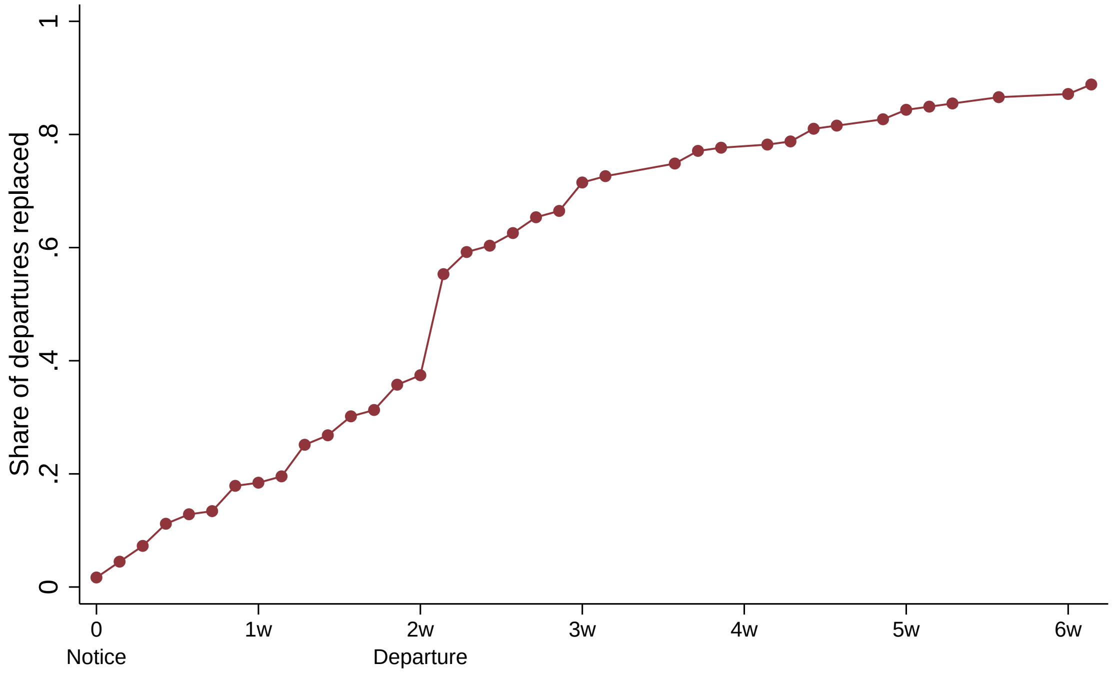

Notes: This figure plots the share of departures that have been replaced as a function of elapsed time since the required notice date. The sample is restricted to stores that were at their target size on that date. The 17 percentage point jump on the day after the departure represents *on-time refills*, and is the modal replacement date (relative to the departure date) by a large margin.

A Appendix

Spreadsheets that compute the main results reported in this Appendix can be downloaded at: https://sites.google.com/view/peter-kuhn/home/data/how-costly-is-turnover. Data and do-files are available on request from the authors. Access is limited to qualified researchers, for academic research purposes only.

A.1 Setting and Data

Table A.1.1: Descriptive Statistics

	Mean	**SD**	**Median**	**N**
Panel A: Product price (in US $)	51.70	49.79	30.94	437
Accessories	16.41	9.77	20.00	20
Shirts and Polos	29.08	10.05	27.81	142
Pants	27.87	10.84	26.25	129
Sweaters	41.63	16.06	40.31	19
Jackets	79.87	34.22	77.81	91
Suits	180.04	53.94	180.94	36
Panel B: Daily sales (in US $)	591	1246	270	75801
Target size=2	330	603	166	14470
Target size=3	616	1388	265	40217
Target size≥4	724	1260	381	21114
Panel C: Monthly compensation (in US $)	538	159	501	7,650
Target size=2	500	105	480	857
Target size=3	542	176	500	4,036
Target size≥4	544	145	513	2,763
Panel D: Average team tenure (in years)	3.45	2.45	2.75	75801
Target size=2	2.60	2.10	1.97	14470
Target size=3	3.36	2.33	2.78	40217
Target size≥4	4.20	2.65	3.41	21114
Before the departure	2.86	2.21	2.23	186
After the departure	2.96	2.39	2.21	186
Departing employees' tenure	2.27	2.69	1.37	186
Before the hiring	3.08	2.37	2.41	218
After the hiring	2.42	2.00	1.83	218

Notes: Product prices in Panel A are from a sample of items sold in September, 2016. Target sizes in Panels B and C are taken from the annual sales plan, which is filed at the store-year level. For newly-opened stores whose target size is not available in the current year, we use actual team size 30 days after the opening instead. Monthly compensation includes a base salary and a commission component based on team performance, along with the social security payments. Monthly compensation is missing for 6% of employee-month observations. Tenure in Panel D measures the team's average tenure with Firm A in years.

Given the average product price of $52, a store on average sells 11.4 items per day. This 'lumpiness', plus large holiday and day-of-week effects, and sizable idiosyncratic daily sales shocks, accounts for the high variance of sales: The standard deviation ($1246) is 2.1 times the mean and 4.6 times the median.

Firm A in Relation to its Labor Market

As noted, Firm A pays at- or above-market wages: During the analysis period, its sales workers earned an average of \$538 per month, compared to an average monthly compensation for retail salespeople of \$510 in Guangdong province. The Guangdong average is based on 1273 retail salespeople's salary reviews from `Kanzhun.com`, which is a Glassdoor-equivalent employee review website in China. The rest-day schedule at Firm A is also at or above the market standard for retail: see a retail industry report at `http://sta.doumi.com/src/vip/report/doumi_final01.pdf`. Chan et al. (2014) provide additional information on compensation at department stores in China.

Turnover at Firm A, in Context

The annual quit rate (number of annual quits / average annual employment) for front-line salespeople in our data is 33.7% and 34.2% in 2015 and 2016 respectively. According to Li et al. (2016), the annual quit rate in China's retail industry was around 30-40% in 2012, and the annual quit rate in the U.S. retail industry in 2017 was 35.4% (U.S. Bureau of Labor Statistics, 2018d).

Tenure at Firm A, in Context

According to Table A.1.1, Firm A's salespeople have an average of 3.45 years of firm tenure; the mean tenure of employees who leave Firm A is about 2.27 years. Median tenure in the U.S. retail industry was 3.0 years in 2018 (U.S. Bureau of Labor Statistics, 2018b).

More on Hiring at Firm A

According to our interviews with Firm A, the regional managers' role in personnel selection is not especially resource-intensive. Hiring one worker requires about 12 hours of a regional manager's time. The process consists of identifying possible candidates (including those suggested by store workers), collecting basic paperwork (proof of identity, high-school or above diploma, and health certificate) and a short interview.

Figure A.1.1: Map of Stores in the Analysis

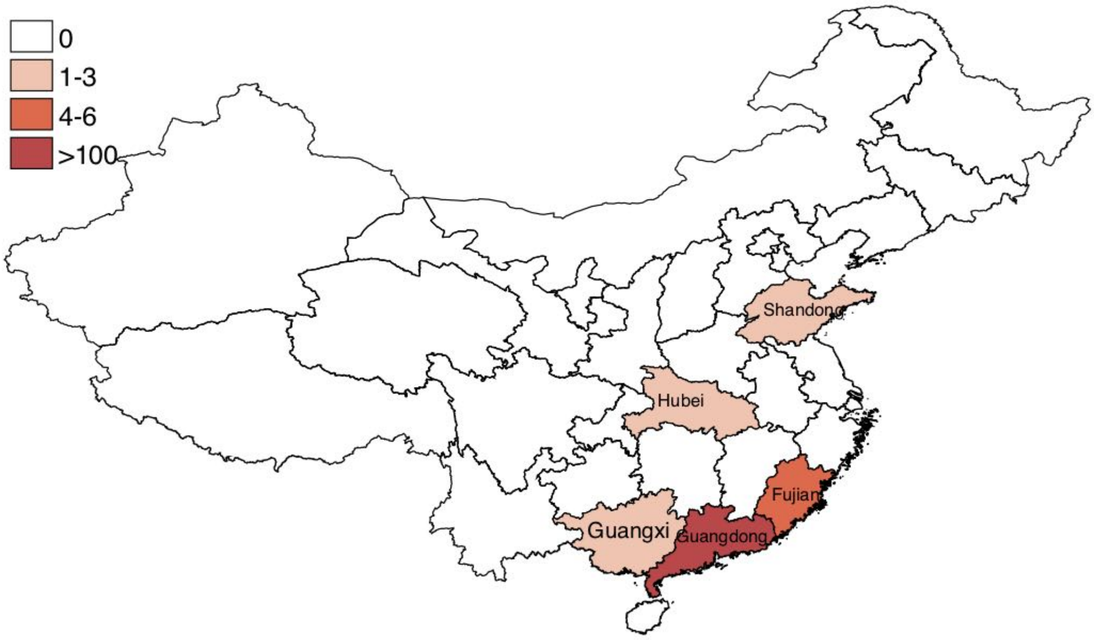

(a) Map of stores in China

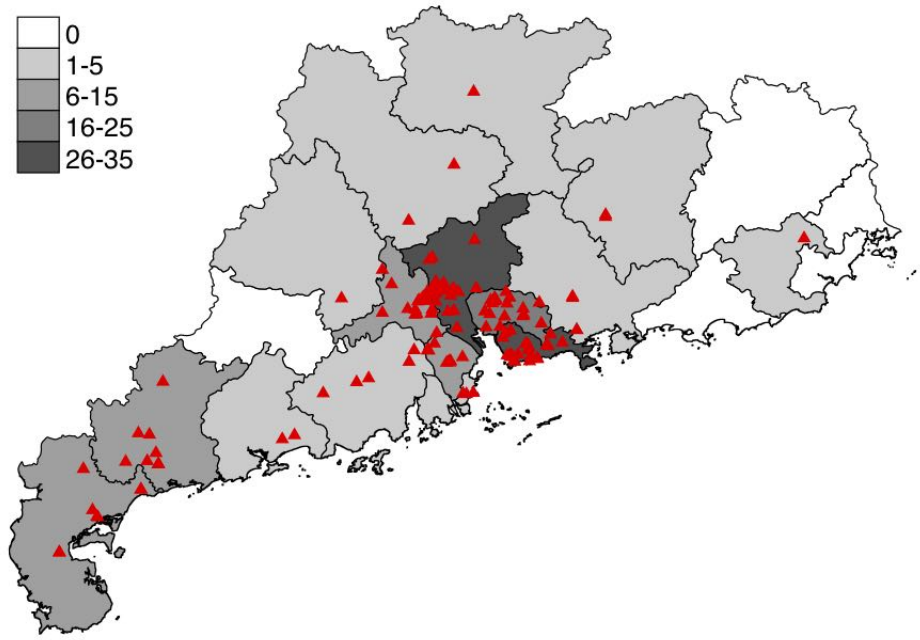

(b) Map of stores in Guangdong

Figure A.1.2: Photos of Stores

(a) A typical store of 2-3 employees (inside a department store)

(b) A typical store of 4 or more employees (inside a shopping mall)

Figure A.1.3: Histogram of Store Sizes

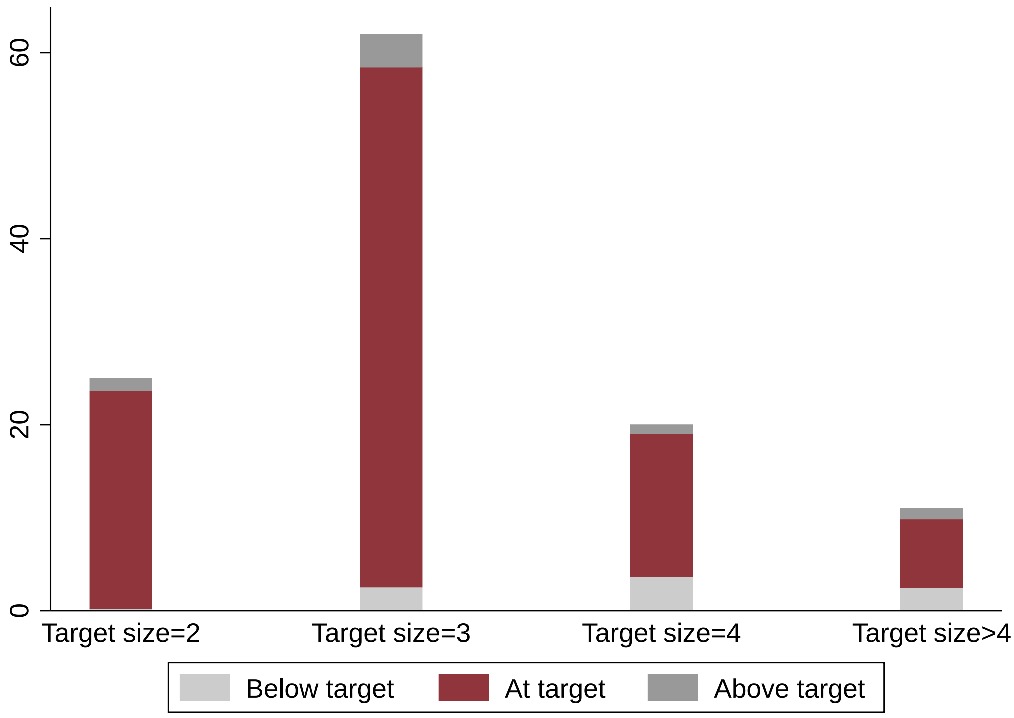

Notes: The histogram categorizes stores according to their target size at the beginning of the analysis period. (For newly-opened stores whose target size is not available in their first year of operation, we use actual team size 30 days after the opening.) The vertical axis indicates the number of stores with each target size. The share of store-month cells at which the stores in each category are at, above, or below their target size are indicated by different colors.

Stores are below their target size if they have an unfilled vacancy. They can be above their target size, for example, if Firm A recently closed a store and reallocated those employees to stores that were already at their target size. The vast majority of the time, however, stores are at their target size.

Shifts and Schedules

Chinese retail employees typically work according to two major staffing schedules. One is the *zuo-yi-xiu-yi*, i.e. an employee is required to work a 12-hour shift on one day and take the next day off. The other one is the *zuo-liu-xiu-yi*, i.e. an employee works a 6- to 8-hour shift per day and takes one rest day in every workweek. Firm A uses the second practice in all its stores, which operate from 9am to 9pm every day, and all the sales employees work exclusively six-hour shifts. These shifts are arranged as follows: For stores of two employees, one works from 9am to 3pm and the other works from 3pm to 9pm; thus there are no overlapping hours between the team members. For stores of three employees, in addition to these two separate shifts there is a third employee working from 12pm to 6pm, since afternoon is the busiest time in a day. For stores with more than three employees, every shift is covered by at least two employees. As a result, the share of the day that a worker overlaps with at least one other team member is 0 percent in 2-person teams, two thirds in three-person teams, and 100 percent in teams with four or more members.

In stores of two or three employees, no two employees can take the same day off. In addition, host institutions (i.e. the malls or department stores in which Firm A's stores are located) help small stores cover rest days by lending them their own, experienced salespeople on a daily basis in return for a share of sales revenue. These reassigned employees earn a flat amount of $15.60 per shift, paid by Firm A, and receive no commissions on store output. This pool of host-institution workers is also the source of the *temporary replacement workers* described in Figure 1, who fill in when smaller teams have not hired a permanent replacement by the departure date.

Figure A.1.4: Shift Arrangements

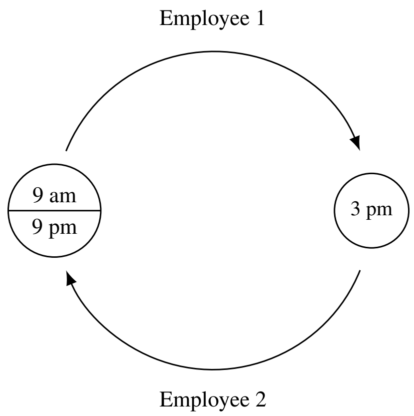

(a) Target size=2

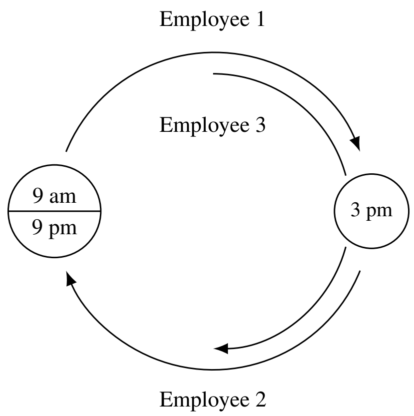

(b) Target size=3

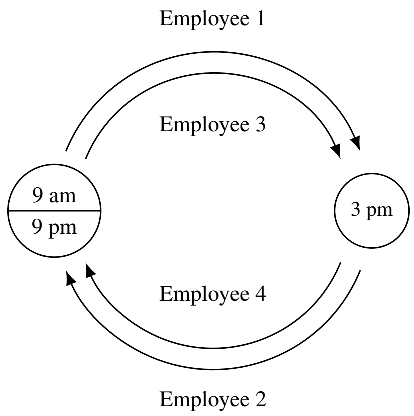

(c) Target size=4

Notes: Arrows indicate the hours each employee works in stores of different sizes. In two-person stores, the employees never overlap. In four-person stores, two workers are always present.

Figure A.1.5: Average Daily Sales per Store by Calendar Day

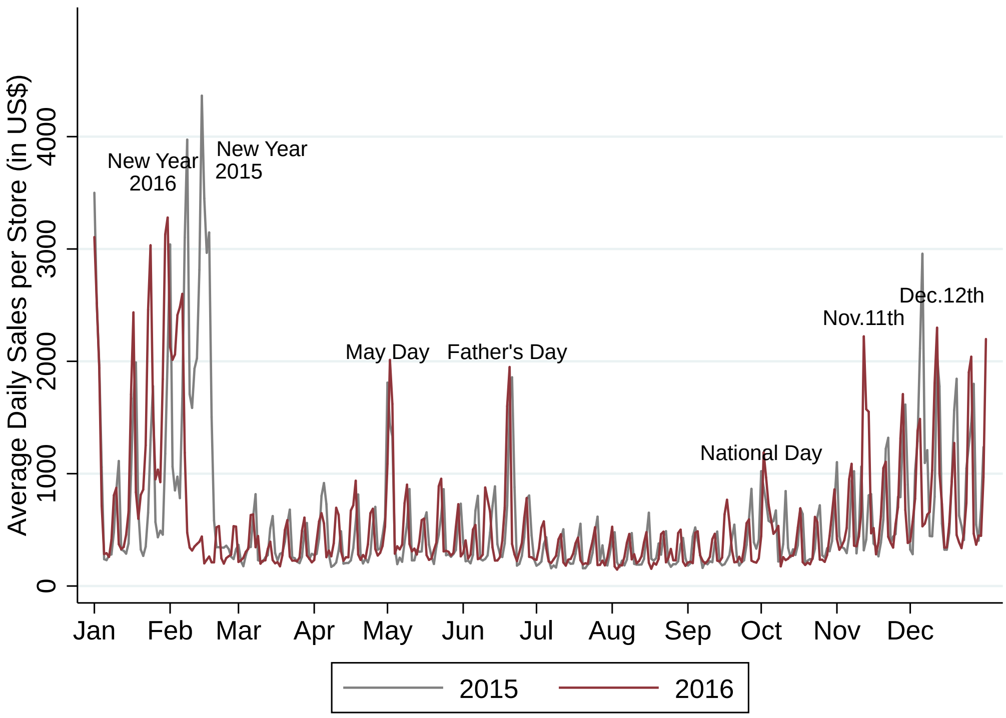

Notes: This figure plots average daily sales per store on every calendar day in 2015 and 2016. The labeled spikes correspond to holidays or major shopping events. New Year denotes the Chinese Lunar New Year. November 11th and December 12th are major shopping events in China, similar to Black Friday or Cyber Monday.

Figure A.1.5 plots the average daily sales per store on every calendar day in 2015 and 2016. Total sales are higher in winter, in part because winter items cost more than summer ones. The smaller cycles represent weekly fluctuations, with higher daily sales occurring on weekends. The largest spikes are labelled, and correspond to major holidays when people shop for menswear heavily, such as Chinese New Year and Father's Day. Significantly, the seasonal and holiday effects are very similar in 2015 and 2016, suggesting that they are predictable well in advance.

Figure A.1.6: The Timing of Departures

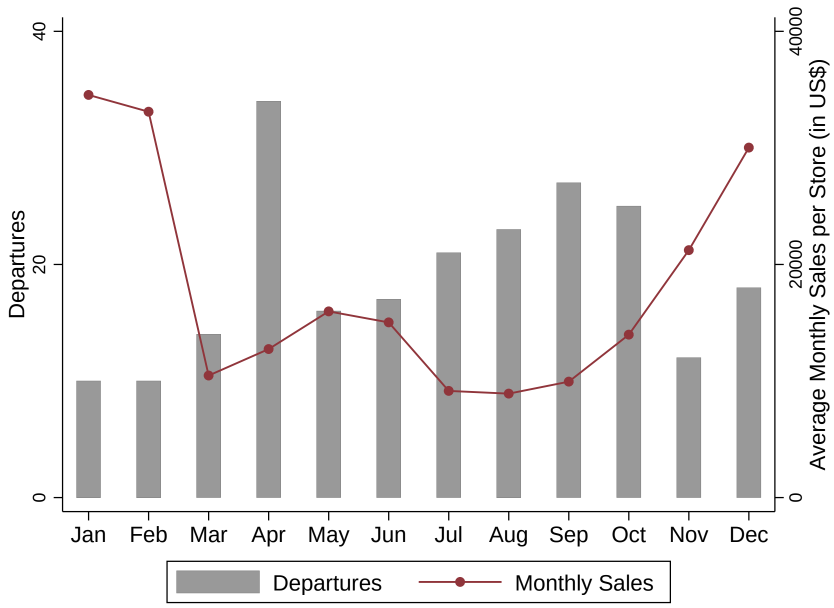

(a) Month of Departures

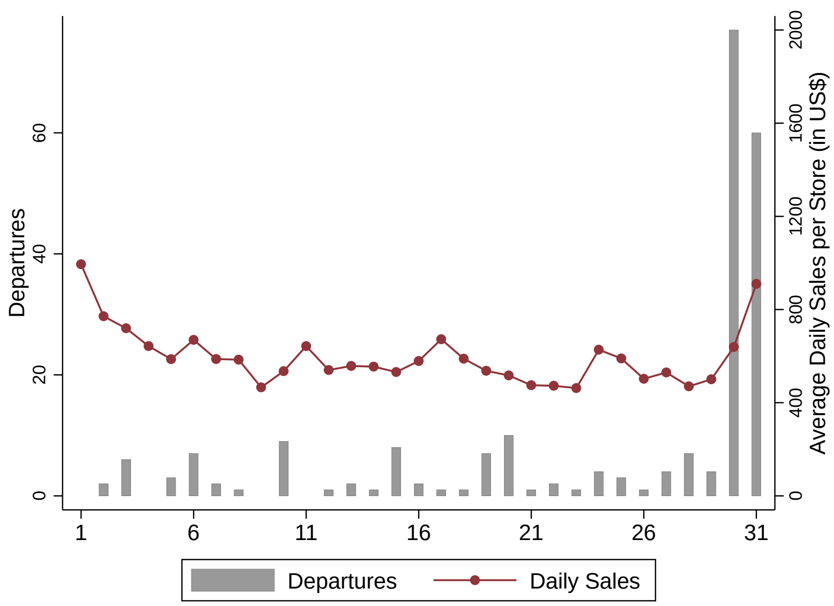

(b) Day of Departures

Notes: The left vertical axis is the number of voluntary departures observed in each calendar month, or day of the month. The right vertical axis is the average monthly or daily sales per store.

Figure A.1.6(a) shows an inverse relationship between a store's monthly sales and its quit rate over the course of a year; we use this variation to estimate a wage-quit elastcity in Section A.4.8. Figure A.1.6(b) shows that departures from Firm A are highly bunched on the last days of calendar months.

A.2 Robustness Checks

Tables A2.1-A2.3 replicate Table 1 for three alternative treatment window widths, *W*=50, 90 and 120.

A.2.1 Treatment Window Width: *W*=50

Table A.2.1: Aggregate Productivity Trends (*W*=50)

Dependent variable: daily sales (in US $)					
Before the Departure:			**After the Departure:**		
	β	SE		β	SE
$P_{-25,-23}$	-28	(49)	$P_{1,2}$	-129*	(76)
$P_{-22,-21}$	-61	(39)	$P_{3,4}$	-59	(63)
$P_{-20,-19}$	-7	(63)	$P_{5,6}$	-58	(48)
$P_{-18,-17}$	-3	(62)	$P_{7,8}$	27	(64)
$P_{-16,-15}$	-79**	(37)	$P_{9,10}$	-41	(51)
$P_{-14,-13}$	-130***	(31)†††	$P_{11,12}$	54	(69)
$P_{-12,-11}$	-52	(46)	$P_{13,14}$	52	(63)
$P_{-10,-9}$	45	(63)	$P_{15,16}$	-9	(46)
$P_{-8,-6}$	-44	(32)	$P_{17,18}$	-5	(47)
$P_{-5,-4}$	-38	(41)	$P_{19,20}$	-19	(67)
$P_{-3,-2}$	-98***	(31)††	$P_{21,22}$	-11	(68)
$P_{-1,0}$	-158***	(36)†††	$P_{23,25}$	-26	(58)
N				66671	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$); † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$.

Notes: This table presents results from estimating a specification similar to equation 1. $P_{-14,-13}$ identifies the first two days of the required notice period; $P_{-1,0}$ identifies the departing employees' last two days of work. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The *q*-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

A.2.2 Treatment Window Width: *W*=90

Table A.2.2: Aggregate Productivity Trends (*W*=90)

Dependent variable: daily sales (in US $)					
Before the Departure:			**After the Departure:**		
	β	SE		β	SE
$P_{-45,-43}$	-14	(57)	$P_{1,2}$	-127*	(76)
$P_{-42,-41}$	-99**	(45)	$P_{3,4}$	-57	(62)
$P_{-40,-39}$	-68	(75)	$P_{5,6}$	-55	(47)
$P_{-38,-37}$	42	(62)	$P_{7,8}$	28	(63)
$P_{-36,-35}$	51	(89)	$P_{9,10}$	-36	(53)
$P_{-34,-33}$	22	(53)	$P_{11,12}$	56	(68)
$P_{-32,-31}$	-39	(44)	$P_{13,14}$	57	(62)
$P_{-30,-29}$	74	(98)	$P_{15,16}$	-6	(46)
$P_{-28,-27}$	17	(79)	$P_{17,18}$	-5	(47)
$P_{-26,-26}$	9	(71)	$P_{19,20}$	-16	(67)
$P_{-24,-23}$	-53	(43)	$P_{21,22}$	-9	(68)
$P_{-22,-21}$	-58	(39)	$P_{23,24}$	-46	(65)
$P_{-20,-19}$	-5	(63)	$P_{25,26}$	-2	(48)
$P_{-18,-17}$	1	(62)	$P_{27,28}$	-13	(43)
$P_{-16,-15}$	-76**	(36)	$P_{29,30}$	15	(102)
$P_{-14,-13}$	-126***	(31)†††	$P_{31,32}$	-64	(67)
$P_{-12,-11}$	-54	(45)	$P_{33,34}$	83	(75)
$P_{-10,-9}$	45	(64)	$P_{35,36}$	30	(85)
$P_{-8,-6}$	-41	(32)	$P_{37,38}$	38	(76)
$P_{-5,-4}$	-35	(41)	$P_{39,40}$	-49	(50)
$P_{-3,-2}$	-93***	(31)†	$P_{41,42}$	-59	(58)
$P_{-1,0}$	-150***	(35)†††	$P_{43,45}$	-54	(44)
N				69343	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$); † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$.

Notes: This table presents results from estimating a specification similar to equation 1. $P_{-14,-13}$ identifies the first two days of the required notice period; $P_{-1,0}$ identifies the departing employees' last two days of work. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The *q*-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

A.2.3 Treatment Window Width: *W*=120

Table A.2.3: Aggregate Productivity Trends (*W*=120)

Dependent variable: daily sales (in US $)					
Before the Departure:			**After the Departure:**		
	β	SE		β	SE
$P_{-60,-59}$	-46	(72)	$P_{1,2}$	-127*	(77)
$P_{-58,-57}$	-58	(51)	$P_{3,4}$	-56	(62)
$P_{-56,-55}$	65	(75)	$P_{5,6}$	-55	(47)
$P_{-54,-53}$	-25	(56)	$P_{7,8}$	29	(65)
$P_{-52,-51}$	-7	(62)	$P_{9,10}$	-37	(53)
$P_{-50,-49}$	-18	(59)	$P_{11,12}$	54	(69)
$P_{-48,-47}$	55	(101)	$P_{13,14}$	56	(63)
$P_{-46,-45}$	-51	(49)	$P_{15,16}$	-7	(46)
$P_{-44,-43}$	11	(68)	$P_{17,18}$	-8	(48)
$P_{-42,-41}$	-103**	(46)	$P_{19,20}$	-17	(67)
$P_{-40,-39}$	-72	(75)	$P_{21,22}$	-9	(68)
$P_{-38,-37}$	41	(62)	$P_{23,24}$	-44	(65)
$P_{-36,-35}$	49	(90)	$P_{25,26}$	-2	(47)
$P_{-34,-33}$	22	(53)	$P_{27,28}$	-15	(44)
$P_{-32,-31}$	-41	(45)	$P_{29,30}$	12	(100)
$P_{-30,-29}$	70	(98)	$P_{31,32}$	-64	(67)
$P_{-28,-27}$	14	(80)	$P_{33,34}$	83	(74)
$P_{-26,-26}$	7	(71)	$P_{35,36}$	30	(84)
$P_{-24,-23}$	-53	(44)	$P_{37,38}$	37	(75)
$P_{-22,-21}$	-58	(39)	$P_{39,40}$	-51	(50)
$P_{-20,-19}$	-4	(63)	$P_{41,42}$	-59	(59)
$P_{-18,-17}$	-0	(62)	$P_{43,44}$	-36	(58)
$P_{-16,-15}$	-76**	(36)	$P_{45,46}$	-82**	(34)
$P_{-14,-13}$	-131***	(34)†††	$P_{47,48}$	8	(39)
$P_{-12,-11}$	-69	(50)	$P_{49,50}$	25	(90)
$P_{-10,-9}$	37	(63)	$P_{51,52}$	-24	(96)
$P_{-8,-6}$	-40	(33)	$P_{53,54}$	-27	(51)
$P_{-5,-4}$	-33	(41)	$P_{55,56}$	83	(141)
$P_{-3,-2}$	-94***	(32)†	$P_{57,58}$	76	(77)
$P_{-1,0}$	-151***	(35)†††	$P_{59,60}$	-41	(67)
N				70281	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$); † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$.

Notes: This table presents results from estimating a specification similar to equation 1. $P_{-14,-13}$ identifies the first two days of the required notice period; $P_{-1,0}$ identifies the departing employees' last two days of work. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

A.2.4 Stricter Test for Pre-Trends

In this table, we push even further to test for any presence of pre-trends before the notice date. Doubling the length of the pre-treatment window associated with our W=120 experiment, we look at 120 days before and 60 days after the departure date. To make this tractable, we changed our bin width from two to four days (except for one 5-day bin from days (-8,-4) to bridge the other bin variables). As presented in Table A.2.4, even in the widest window, we find no pre-trends in productivity before the required notice date.

Table A.2.4: Stricter Test for Pre-Trends

Dependent variable: daily sales (in US $)					
Before the Departure:			**After the Departure:**		
	β	SE		β	SE
$P_{-120,-117}$	7	(56)	$P_{1,4}$	-91	(59)
$P_{-116,-113}$	38	(62)	$P_{5,8}$	-16	(48)
$P_{-112,-109}$	119	(91)	$P_{9,12}$	6	(48)
$P_{-108,-105}$	26	(49)	$P_{13,16}$	22	(45)
$P_{-104,-101}$	-57	(38)	$P_{17,20}$	-11	(42)
$P_{-100,-97}$	34	(44)	$P_{21,24}$	-24	(58)
$P_{-96,-93}$	96	(63)	$P_{25,28}$	-9	(42)
$P_{-92,-89}$	61	(79)	$P_{29,32}$	-29	(70)
$P_{-88,-85}$	75	(80)	$P_{33,36}$	52	(73)
$P_{-84,-91}$	43	(53)	$P_{37,40}$	-10	(45)
$P_{-80,-77}$	99	(67)	$P_{41,44}$	-50	(54)
$P_{-76,-73}$	150*	(77)	$P_{45,48}$	-40	(26)
$P_{-72,-69}$	31	(52)	$P_{49,52}$	1	(89)
$P_{-68,-65}$	2	(62)	$P_{53,56}$	30	(88)
$P_{-64,-61}$	66	(79)	$P_{57,60}$	15	(56)
$P_{-60,-57}$	-47	(56)			
$P_{-56,-53}$	21	(62)			
$P_{-52,-49}$	-12	(52)			
$P_{-48,-45}$	6	(58)			
$P_{-44,-41}$	-45	(52)			
$P_{-40,-37}$	-11	(53)			
$P_{-36,-33}$	37	(59)			
$P_{-32,-29}$	18	(60)			
$P_{-28,-25}$	9	(58)			
$P_{-24,-21}$	-60	(37)			
$P_{-20,-17}$	-1	(49)			
$P_{-16,-13}$	-103***	(30)			
$P_{-12,-9}$	-11	(45)			
$P_{-8,-4}$	-37	(32)			
$P_{-3,0}$	-123***	(27)			
N				71176	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$); † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$.

Notes: This table presents results from estimating a specification similar to equation 1. We double the length of the pre-treatment window associated with our W=120 experiment from 60 to 120 days. $P_{-16,-13}$ identifies the four days around the notice; $P_{-3,0}$ identifies the departing employees' last four days of work. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

A.2.5 Effect Duration (M=2,3,5,6)

Table A.2.5: Robustness Check – *Effect Duration (M=2,3,5,6)*

	Dependent variable: daily sales (in US $)			
	(1) M=2	**(2) M=3**	**(3) M=5**	**(4) M=6**
Around Notice$_M$	-87**	-114***	-93***	-82***
	(41)	(35)†††	(33)††	(31)†
Before Departure$_M$	-143***	-121***	-102***	-84***
	(37)†††	(29)†††	(25)†††	(26)††
Early On-boarding$_M$	-81**	-75**	-47	-28
	(36)	(35)	(36)	(39)
Late On-boarding$_M$	-13	-8	-14	-26
	(32)	(33)	(37)	(40)
Early Short-staffing$_M$	-241*	-204	-211**	-197**
	(140)	(126)	(94)†	(84)†
Late Short-staffing$_M$	-53	-45	-33	-29
	(52)	(52)	(55)	(55)
N	75801	75801	75801	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$

Notes: *Around Notice$_M$* identifies an M-day period surrounding the notice period. *Before Departure$_M$* identifies the last M days the departing employee spends at the team. *Early On-boarding$_M$* identifies the first M days a new employee is present, and *Late On-boarding$_M$* identifies the (M+1)th day through two weeks after the hiring. *Early Short-staffing$_M$* identifies the first M days following an unfilled departure, and *Late Short-staffing$_M$* identifies other days following *Early Short-staffing$_M$* with the departure unfilled through 30 days following the departure. As in column 2 of Table 3, all regressions include a full set of fixed effects for (mostly) four-day bins in the treatment window (W) that are not included in the AN and BD effects. Non-four-day bins are at the outer bounds of the treatment window (W), and in the middle of the 16-day 'notice' interval (days -15 through 0. When M= 5 or 6 (thus extending the AN and BD effects towards the middle of the 16-day notice interval) the two middle bins in this interval shorten to three or two days each; when M=2 or 3 they length to five and six days each. Regressions also include 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, presented by the standard errors, correct the false discovery significances following Benjamini-Hochberg procedure.

This table replicates column 2 of Table 3 for different assumed durations, M of the AN, BD, early OB, and early SS effects. (The durations of the late OB and SS effects are determined residually.) For the most part, the estimated effects on daily sales decline in magnitude with the assumed effect duration, suggesting that the true effect is of relatively short duration.

A.2.6 A Polynomial Approach to the Timing of Pre-departure Productivity Losses

An alternative way to model the team productivity trend during the fourteen-day advance notice period specifies team productivity as a polynomial function of time during that period (while allowing for different productivity intercepts before, during and after the notice period). The estimates and 95% confidence interval from this approach are plotted in Figure A.2.6, for both cubic and quartic specifications of the polynomial. As in our main, binned approach, the productivity losses during the notice period are highly concentrated at the two ends of the period, and are insignificant in the middle.

Figure A.2.6: Polynomial Estimates of the Productivity Trend During the Notice Interval

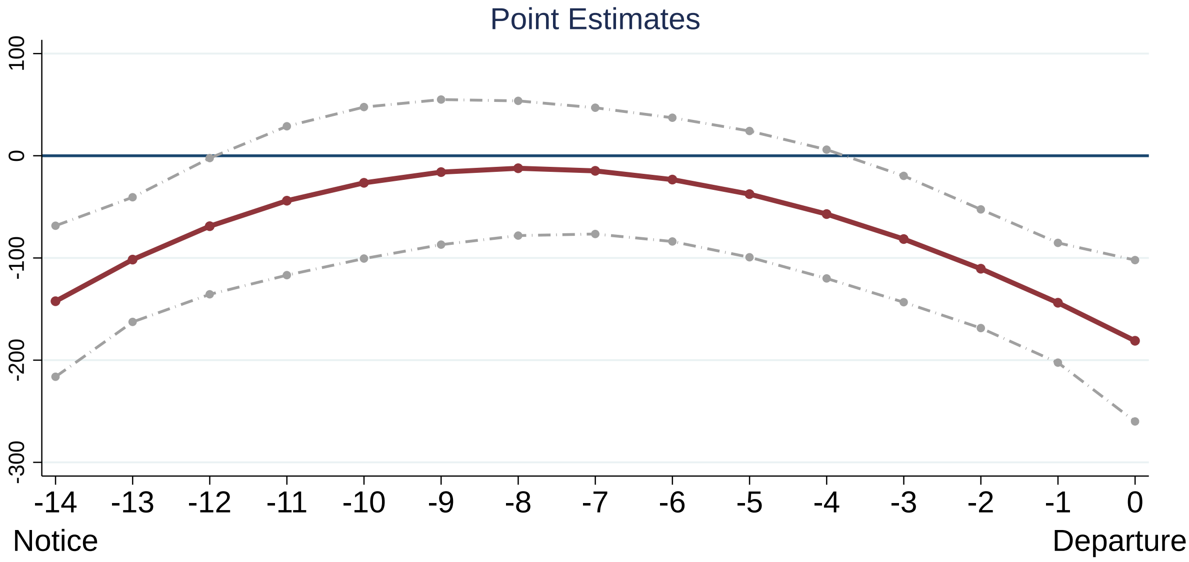

(a) Cubic

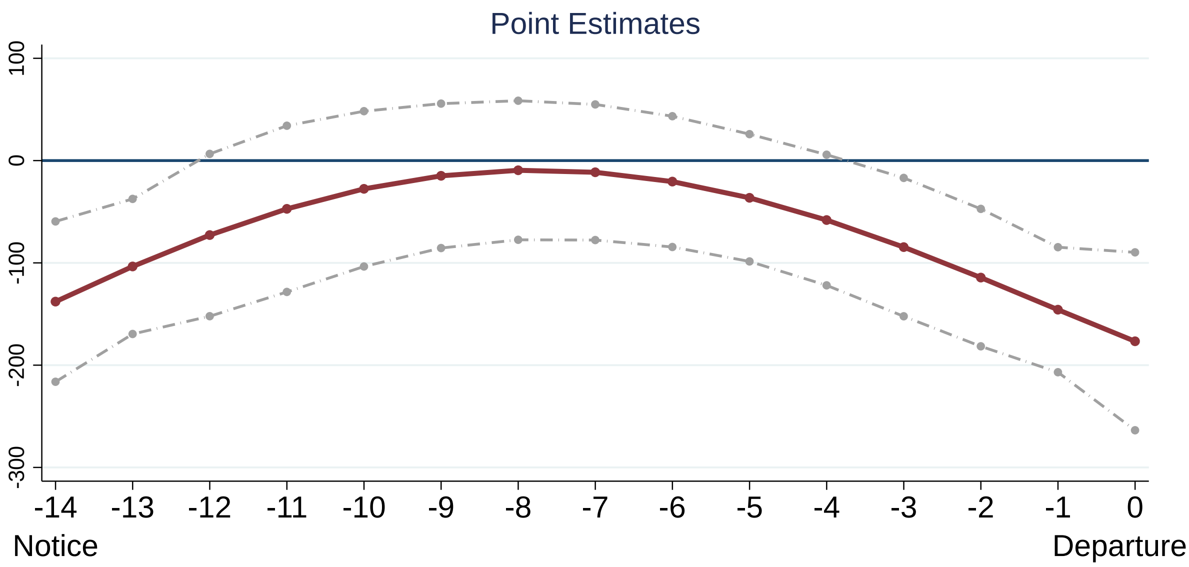

(b) Quartic

A.2.7 Assessing the Role of End-of-Month Departures

As noted, almost half the departures in our sample occur on the last day of a calendar month. In this Section, we explore whether our results might somehow be driven by any special features of the timing of departures that are not accounted for by our calendar-day fixed effects. In Table A.2.7 we report the results of two alternative specifications of Table 1 that address this issue. As a baseline for these two exercises, column 1 of Table A.2.7 displays the main coefficients from Table 1.

In column (2) of Table A.2.7, we exclude from our sample all departures that happen during the last three days of a calendar month. Like our baseline estimates, column (2) shows output declines around the time of notice and departure of a roughly similar timing and size as in the full sample.

In column (3) of Table A.2.7 we estimate the effects of placebo departures with the same distribution across days of the month as the actual departures, but in months where no departure occurred. Specifically, we code the same day of the third month *before* the actual notice date in a store as a placebo departure from that store. For example, if the actual departure day is on September 22^{nd}, we use June 22^{nd} instead. The goal is to test whether low sales just happen to occur on the days of the month when departures (and notices) are typically given. Here, we see no estimated productivity declines around the placebo departure and notice dates, suggesting that there is nothing special about these days of the month that might be generating spurious 'effects' of those events.

Table A.2.7: Assessing the Role of End-of-Month Departures

	Dependent variable: daily sales (in US $)		
	(1)	(2)	(3)
	Full Sample	Excluding end-of-month Departures	Placebo departures
$P_{-22,-21}$	-59	-87	-13
	(39)	(64)	(90)
$P_{-20,-19}$	-4	80	125
	(63)	(145)	(104)
$P_{-18,-17}$	-0	-15	164
	(62)	(114)	(97)
$P_{-16,-15}$	-78**	-124	-1
	(37)	(86)	(68)
$P_{-14,-13}$	-128***	-144**	-32
	(31)†††	(56)	(47)
$P_{-12,-11}$	-54	-93	-64
	(45)	(87)	(51)
$P_{-10,-9}$	47	108	-38
	(63)	(148)	(49)
$P_{-8,-6}$	-43	-43	13
	(32)	(72)	(44)
$P_{-5,-4}$	-35	10	31
	(40)	(86)	(64)
$P_{-3,-2}$	-94***	-43	50
	(31)††	(50)	(88)
$P_{-1,0}$	-153***	-118**	88
	(35)†††	(60)	(71)
$P_{1,2}$	-131*	-4	86
	(76)	(127)	(118)
$P_{3,4}$	-59	-90	120
	(63)	(77)	(94)
$P_{5,6}$	-58	-101	104
	(48)	(86)	(89)
N	68238	61805	60895

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$

Notes: Regression specification is identical to equation 1, although coefficients before day -22 and coefficients after day -6 are omitted in this table. In each column, regression is estimated by including days being identified by the independent variables and the control sample. Column 2 includes only departures that did not occur in the last three days of the month, plus the same control group of observations where no hires or separations occurred. Column 3 uses the same day in the past third month plus the same control group of observations to run a placebo test. All regressions include 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

A.2.8 Controlling for Store-Specific Trends

In this table, we replicate Table 3, replacing the 731 day effects by 24 month effects (to capture aggregate trends). In addition, we interacted seven day-of-the week effects, ten holiday effects, and twelve month-of-the-year effects with the 118 store fixed effects, to allow each store to have its own pattern of sales over the week and year, and to be differently affected by holidays from other stores.[44]

Table A2.8: Controlling for Store-Specific Trends

	Dependent variable: daily sales (in US $)	
	(1)	(2)
Around Notice$_4$ (P$_{-15,-12}$)	-78***	-78***
	(23)†††	(24)††
Before Departure$_4$ (P$_{-3,0}$)	-122***	-128***
	(33)†††	(37)††
Early On-boarding$_4$	-41	-40
	(42)	(45)
Late On-boarding$_4$	28	36
	(30)	(29)
Early Short-staffing$_4$	-128	-151
	(81)	(102)
Late Short-staffing$_4$	-77	-74
	(51)	(51)
Other Period Effects		Yes
N	75743	75743

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$

Notes: This table presents results from estimating equation 2. "Other Period Effects" refer to (mostly) four-day bins (P) during the notice period other than the Around-Notice and Before-Departure periods. The regression includes a vector of 24 year×month dummies, a vector of 1,336 store×month-of-year dummies, 826 store×day-of-week dummies, and a vector of 1,230 store×holiday dummies. 58 singleton observations are dropped from this regression. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

[44] Holiday is a vector of 10 dummies, identifying the new year, Chinese new year, May Day, Father's Day, National Day, November 11, December 12, and three other traditional Chinese holidays.

A.3 Heterogeneity Analysis

This Section examines how various components of turnover costs (AN, BD, OB and SS) vary with characteristics of the departure and characteristics of the team. Motivated by Table 3's finding that late on-boarding and short-staffing effects are absent, all these regressions focus on the short-term OB and SS effects only. In all other respects, the regression specifications are identical to Table 3.

Table A.3.1: Heterogeneity Examination – *Leaver's Rank and Seniority*

	Dependent variable: daily sales (in US$)	
	(1)	(2)
Around Notice$_4$	-91***	-116***
	(34)††	(32)†††
Before Departure$_4$	-108***	-109***
	(24)†††	(26)†††
Early On-boarding$_4$	-77**	-75**
	(34)†	(35)†
Early Short-staffing$_4$	-217**	-267**
	(91)†	(115)†
i.(Leaver=Manager)	-16	
	(45)	
AN$_4$ $\times$ i.(Leaver=Manager)	-133**	
	(61)†	
BD$_4$ $\times$ i.(Leaver=Manager)	-26	
	(78)	
OB$_4$ $\times$ i.(Leaver=Manager)	16	
	(72)	
SS$_4$ $\times$ i.(Leaver=Manager)	-222	
	(138)	
i.(Leaver=Senior)		16
		(49)
AN$_4$ $\times$ i.(Leaver=Senior)		19
		(94)
BD$_4$ $\times$ i.(Leaver=Senior)		-36
		(61)
OB$_4$ $\times$ i.(Leaver=Senior)		-40
		(93)
SS$_4$ $\times$ i.(Leaver=Senior)		101
		(161)
N	75801	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$
Notes: In Column (1), we interact all short-run effects with an indicator variable identifying the departing employee being a store manager. In Column (2), we interact all short-run effects with an indicator variable identifying whether the departing employee's firm tenure is above the average in Firm A. All regressions include 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Table A.3.1 shows that productivity losses near the announcement (AN effects) are much larger when a manager quits than when other workers quit. Compared with a regular employee's announcement – which reduces productivity by \$91 or 15% –, a store manager's announcement adds an extra loss of \$133 – more than doubling the AN effect associated with a normal employee. Operating without the manager (SS) also appears to entail substantial additional costs, though these are statistically insignificant. Interaction effects with the seniority of the departing worker are mostly small, and are statistically insignificant in all cases.

Table A.3.2: Heterogeneity Examination – *Stayers' Seniority*

	Dependent variable: daily sales (in US$) (1)	(2)
Around Notice$_4$	-160**	-198***
	(70)	(67)††
Before Departure$_4$	-80	-105**
	(48)	(44)†
Early On-boarding$_4$	-108*	-78
	(57)	(52)
Early Short-staffing$_4$	-35	-137
	(219)	(152)
Manager's Tenure	5	
	(9)	
AN$_4$ $\times$ Manager's Tenure	12	
	(11)	
BD$_4$ $\times$ Manager's Tenure	-8	
	(8)	
OB$_4$ $\times$ Manager's Tenure	4	
	(8)	
SS$_4$ $\times$ Manager's Tenure	-32	
	(29)	
Stayers' Tenure		-1
		(13)
AN$_4$ $\times$ Stayers' Tenure		28
		(20)
BD$_4$ $\times$Stayers' Tenure		-3
		(11)
OB$_4$ $\times$ Stayers' Tenure		2
		(11)
SS$_4$ $\times$ Stayers' Tenure		-23
		(29)
N	74221	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$

Notes: In column (1), we interact all short-run effects with the store manager's firm tenure. To not confound with the costs of a manager's departure, we exclude 30 days before or after a manager's departure in estimating this model. Stayers' Tenure in column (2) is the average firm tenure of employees who stay at the team, excluding the departing employee and the replacement employee. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Table A.3.2 studies whether stayers with more firm tenure can alleviate the costs of a departure. In column (1), we focus on the tenure of the store manager. To exclude any direct effects of a store manager's departure, this column excludes the 30 days before or after a manager's departure. In column (2), we consider the average firm tenure of all team members who remain at the team. The departing employees and the new employees who arrived during the 61-day window around the departure are not included. Results in both columns indicates that none of the four components of turnover costs are significantly affected by these characteristics of the stayers.

Table A.3.3: Heterogeneity Examination – *Previous Experience with Turnover*

	Dependent variable: daily sales (in US$)	
	(1)	**(2)**
Around Notice$_4$	-148***	-108***
	(49)††	(33)††
Before Departure$_4$	-141***	-135***
	(43)††	(30)†††
Early On-boarding$_4$	-59	-38
	(51)	(35)
Early Short-staffing$_4$	-381**	-302***
	(146)†	(114)††
AN$_4$ $\times$ High-Turnover Team	70	
	(66)	
BD$_4$ $\times$ High-Turnover Team	51	
	(57)	
OB$_4$ $\times$ High-Turnover Team	-26	
	(65)	
SS$_4$ $\times$ High-Turnover Team	280*	
	(156)	
Recent Turnover		-14
		(27)
AN$_4$ $\times$ Recent Turnover		88
		(61)
BD$_4$ $\times$ Recent Turnover		69
		(55)
OB$_4$ $\times$ Recent Turnover		-12
		(70)
SS$_4$ $\times$ Recent Turnover		352***
		(123)††
N	75801	66318

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$
Notes: High-turnover store is an indicator variable, taking a value of 1 if the store has experienced above-median number of turnover events (>2) during the analysis period. (This variable alone is omitted since we control for store fixed effects.) Recent turnover is an indicator variable, taking a value of 1 if there is at least one departure occurring at the same store in the past 3 months. Observations in January-March 2015 are excluded in column (2), since previous turnovers are not observed. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

This table studies whether a team's experience in dealing with turnover may reduce the costs of turnover. Columns (1) and (2) indicate that previous experience with turnover significantly reduces sales reductions during periods of short-staffing, but not the other costs. In fact, these short-staffing costs appear to be absent if a store experienced a turnover during the past three months. A likely mechanism for these cost reductions is improvements in the co-ordination of the remaining members' work schedules during short-staffing spells, a consideration which is highly salient to the workers.

Table A.3.4: Heterogeneity Examination – *Productive Teams and Productive Days*

	Dependent variable: daily sales (in US$)	
	(1)	(2)
Around Notice$_4$	9	-18
	(42)	(33)
Before Departure$_4$	-6	-79***
	(32)	(24)†††
Early On-boarding$_4$	-9	-37
	(37)	(31)
Early Short-staffing$_4$	-178	-80
	(149)	(99)
AN$_4$ $\times$ Productive Team	-258***	
	(60)†††	
BD$_4$ $\times$ Productive Team	-233***	
	(54)†††	
OB$_4$ $\times$ Productive Team	-130*	
	(66)	
SS$_4$ $\times$ Productive Team	-67	
	(151)	
AN$_4$ $\times$ Productive Day		-189***
		(57)††
BD$_4$ $\times$ Productive Day		-85*
		(50)
OB$_4$ $\times$ Productive Day		-74
		(58)
SS$_4$ $\times$ Productive Day		-254*
		(147)
N	75801	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$
Notes: In column (1), productive team is an indicator variable, taking a value of 1 if a store's average sales output is above the median sales at its size. (This variable alone is omitted since we control for store fixed effects.) In column (2), productive day is an indicator variable, taking a value of 1 if the focal day is a Friday, Saturday, or Sunday, or one of the ten holidays when Chinese people shop heavily. (This variable alone is omitted since we control for date fixed effects.) The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Table A.3.4 shows that turnover costs are mostly from stores with above-average sales for their size, and from days when teams can potentially perform better than other days. These additional costs were experienced across all four cost components (AN, BD, OB and SS), consistent with the idea that the opportunity costs of shirking, employee recruiting activities, training new workers, and operating with one less worker are all lower when storese are not busy, i.e. when there are fewer customers available to serve.

Table A.3.5: Heterogeneity Examination – *Different Types of On-Boarding*

	Dependent variable: daily sales (in US$)	
	(1)	(2)
Around Notice$_4$	-110***	-110***
	(31)†††	(31)†††
Before Departure$_4$	-112***	-120***
	(25)†††	(26)†††
Early On-boarding$_4$	-76**	-98**
	(34)†	(42)†
Early Short-staffing$_4$	-238***	-238***
	(86)††	(86)††
Trained	3	
	(38)	
OB$_4$ $\times$ Trained	-23	
	(109)	
Leaver's working		23
		(51)
OB$_4$ $\times$ Leaver's working		70
		(67)
N	75801	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$

Notes: In Column (1), we interact the on-boarding effect with the type of the new employee – new outside hires versus trained, temporary replacements from the host institution or internal hires from other retail stores of Firm A. In Column (2), we interact on-boarding with an indicator variable which equals 1 if the departing employee is still working at the store, and 0 otherwise. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

Column 1 of Table A.3.5 shows that on-boarding costs are statistically no different when the new worker has very localized retail sales experience ("trained") than when she does not. "Trained" means that the replacement previously worked at another of Firm A's stores, or in the current store's host institution. Column 2 shows that on-boarding costs are statistically no different when the departing worker is still present when on-boarding occurs, than when she has already departed.

Table A.3.6: Heterogeneity Examination – *Hiring Urgency*

	Dependent variable: daily sales (in US \$)
	(1)
Around Notice$_4$	-118***
	(33)†††
Before Departure$_4$	-134***
	(31)†††
Early On-boarding$_4$	-84**
	(33)††
Early Short-staffing$_4$	-238***
	(86)††
Vacancy Filled	25
	(45)
BD$_4$ $\times$ Vacancy Filled	48
	(70)
N	75801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$; † $q < 0.1$, †† $q < 0.05$, ††† $q < 0.01$
Notes: Vacancy Filled is an indicator variable identifying the teams that have hired replacement workers four or more days before the departure. The regression includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors in parentheses are clustered at the store level. The q-values, represented by daggers (†), correct for false discovery following the Benjamini-Hochberg procedure.

This table focuses on the four days before the departure (the BD period), and compares the productivity of teams who are still looking for a replacement with those who have already hired one. While the point estimate on the interaction term (BD$_4$ $\times$ Vacancy Filled) suggests that eliminating the pressure to recruit mitigates team productivity losses just before the departure, the coefficient is not statistically significant.

A.4 Detailed Calculations of Turnover Costs

A.4.1 Total Lost Sales Associated with a Departure, and their Composition

In this Section, we calculate the total reduction in team sales associated with a departure by multiplying estimates of each of the four sources of lost sales from column (2) of Table 3 (AN, BD, OB and SS) by the duration of each of those sales declines.[45] Our results are as follows:[46]

$$\begin{aligned}
\textbf{On-boarding Costs} &= 4 \times \text{Early on-boarding costs } + 10 \times \text{Late on-boarding costs} \\
&= 4 \times \$76 + 10 \times \$5 = \textbf{\$354} \\
\textbf{Short-staffing Costs} &= \text{Average early short-staffed days} \times \text{Early short-staffing costs} \\
&\quad + \text{Average late short-staffed days} \times \text{Late short-staffing costs} \\
&= 0.97 \times \$218 + 4.59 \times \$41 = \textbf{\$400} \\
\textbf{Around-Notice Costs} &= 4 \times \$112 = \textbf{\$448} \\
\textbf{Before-Departure Costs} &= 4 \times \$113 = \textbf{\$452}
\end{aligned}$$

$$\textbf{Total Lost Sales} = \$354 + \$234 + \$448 + \$452 = \textbf{\$1,654}$$

Turning to the composition of output losses, we note that only 24 percent of the lost output associated with turnover costs ($400/$1,654) is due to short-staffing. In addition, we find that almost two thirds (63%) of all productivity losses associated with turnover are incurred *before the departing worker leaves*. This comprises AN and BD costs, plus 40% of OB costs.

A.4.2 Robustness of Total Lost Sales to Assumed Effect Durations

Here we provide total lost sales estimates under alternative assumptions for the duration of the AN, BD, early-short-staffing and early-on-boarding periods, which were computed using M = 4 days in the preceding analysis, based on the Table 3 regressions. The following cost estimates are based on estimates for M = 2,3,5, and 6 from Appendix Table A.2.5.

[45] In Table 3, most effect durations are assumed to be 4 days (alternative assumptions are explored in Appendix A.2.) Expected short-staffing durations, however, are averages computed from the data. For example, average early short-staffed days = [0×(number of early refills + number of on-time refills + number of temp workers)+ 1×(number replaced on the 2nd day) + 2×(number replaced on the 3rd day) + 3×(number replaced on the 4th day) + 4×(number not replaced by the 4th day)] ÷ (total number of departures). The average number of late short-staffed days is calculated in a similar manner.

[46] Motivated by the fact that 94% of the departures in our data were replaced within three months following the departure, these calculations assume that every departure is eventually replaced. In other words, we assume that a single on-boarding event is associated with every departure.

Table A.4.2: Robustness of Total Lost Sales to Assumed Effect Durations

	$M = 4$	$M = 2$	$M = 3$	$M = 5$	$M = 6$
On-boarding Costs	354	318	313	361	376
Short-staffing Costs	400	389	368	397	411
Around-Notice Costs	448	174	342	465	492
Before-Departure Costs	452	286	363	510	504
Total Lost Sales	**1,654**	**1,167**	**1,386**	**1,733**	**1,783**

Thus, the largest estimate ($1783 for *M*=6) is 7.8 percent higher than our baseline estimate of $1654. Estimates for effect durations greater than *M*=6 become statistically imprecise.

A.4.3 Gross versus Net Sales

In all our econometric analysis, team output is measured by the gross value of retail sales, which is the metric Firm A uses to track store performance and pay its workers. When a turnover occurs, however, Firm A's profits do not necessarily decline by the fall in gross revenues; to calculate the effects on profits we need to adjust for any declines in variable costs that results from lower sales.

To address this issue, we asked representatives of Firm A to consider which variable costs would be saved if an additional departure occurred within a store in a year; their answer was the cost of manufacturing the clothes and transporting them to the store, which amount to about 35 percent of the clothes' selling price. All remaining costs, including rent paid for retail space, would remain unaffected. Thus, our baseline estimate of the decline in *net* sales (and hence in profits) associated with a departure equals 65 percent of the decline in gross sales, i.e. $1654 $\times$0.65 = $1075.

Two comments on this approach may be worth noting. First, while derived in the context of a single departure, this 65% adjustment factor also applies to the effects of any permanent change in a store's turnover rate, as long as that store's long run size (i.e. rental space and target team size) are unaffected by the change in turnover. We shall use this property to estimate the effects of permanent changes in Firm A's wage policy below. Second, since 65% is an imprecise estimate, we can bound the impact of turnover costs on profits by considering the most extreme alternatives. At one extreme, all of Firm A's costs are fixed with respect to changes in turnover. In that case, –where none of the clothing that is unsold due to a turnover ever produces revenues– turnover costs will be 100/65 -1 = 54 percent larger than our baseline estimates. At the other extreme, all of Firm A's costs would be variable with respect to changes in turnover. In this case, true turnover costs will be smaller than our baseline estimates, consisting only of the (very small) administrative costs of processing a turnover discussed below.

A.4.4 Other Turnover Costs

In addition to the reductions in team sales performance documented in our econometric analysis, turnover also generates additional administrative costs and leads to temporary changes in wages paid to salespeople. Our estimates of these quantities (for a single, representative departure, based on discussions with Firm A) are presented below:

$$\begin{aligned}
\textbf{Administrative Costs} &= \text{Average time regional managers spend on hiring} \\
&\quad \times \text{Regional managers' daily wage} \\
&= 1.5 \times \$72 = \textbf{\$108} \\
\textbf{Additional Salary When Workers Overlap} &= \text{Average early refilled days} \\
&\quad \times \text{Salespeople's average daily wage} \\
&= 2.89 \times \$21 = \textbf{\$61} \\
\textbf{Wage Savings While Short-Staffed} &= \text{Average short-staffed days} \\
&\quad \times \text{Salespeoples' daily wage} \\
&= 5.32 \times \$21 = \textbf{\$112}
\end{aligned}$$

$$\begin{aligned}
\textbf{Total Non-Sales Turnover Costs} &= \text{Administrative Costs} \\
&\quad + \text{Additional Salary When Workers Overlap} \\
&\quad - \text{Wage Savings While Short-Staffed} \\
&= \textbf{\$57}
\end{aligned}$$

Note that wage savings while short-staffed only apply to teams of four or more employees. Smaller teams never experience short-staffing because the vacant slot is filled by a temporary replacement worker, whose wage costs are approximately equal to what the departed worker would have been paid. Finally, combining our estimated \$1,075 reduction in net sales with \$57 in non-sales costs of turnover yields our estimate that a single departure reduces Firm A's profits by **\$1,132.** This is our most comprehensive estimate of the cost of a single departure from Firm A. Table A.4.4 summarizes these calculations (which are for assumed effect durations of M=4) and shows how they change when M=2, 3, 5 or 6.

Table A.4.4: **Total Turnover Costs (for Each Assumed Duration *M*)**

	$M = 4$	$M = 2$	$M = 3$	$M = 5$	$M = 6$
On-boarding Costs	230	207	203	235	244
Short-staffing Costs	260	253	239	258	267
Around-Notice Costs	291	113	222	302	320
Before-Departure Costs	294	186	236	332	328
a) Total Lost Net Sales	**1,075**	**759**	**901**	**1,126**	**1,159**
b) Total Other Costs	**57**	**57**	**57**	**57**	**57**
Total Turnover Costs	**1,132**	**816**	**958**	**1,183**	**1,216**

A.4.5 Expressing Turnover Costs in Terms of Career Sales and Wages

The preceding section's estimate of total turnover costs –$1,132 of net sales– does not allow for easy comparisons with other firms and industries. To facilitate such comparisons, we now express these costs in terms of days' worth of employee net sales, days of wages, and as fractions of total sales and wages over a typical employee's career. Turning first to net sales, at $120 in net sales per employee per day, $1,132 translates to 9.43 days of per-employee sales. Using a mean completed tenure of 2.27 years, the cost of a single turnover is equivalent to 1.14% of an employee's expected career net productivity. Turning to wage-based measures, at average wage of $18 per day, $1,132 translates to 63 days of per-employee wages. Thus, the cost of a single turnover can also be expressed as 7.6 percent of an employee's career wages at Firm A. As noted in the paper, the very low share of sales employees' pay in their net sales (18/120 = 15 percent) does not necessarily imply a high rate of exploitation by Firm A. Instead, the main reason for this difference is a very high share of fixed costs –primarily rent for retail space– in Firm A's costs. Sales revenues must at least cover these fixed costs in order for Firm A to remain in business.

A final way to express the magnitude of our estimated turnover costs is via the trade-off between the rate of turnover and the firm's (permanent) wage rate. Specifically, by how much could Firm A afford to raise wages (keeping profits constant) if it was able to cut turnover by 10 percent? With 378 sales employees and a baseline annual turnover rate of 34%, a 10% decline in turnover would eliminate 12.84 departures per year. According to our estimates, Firm A should therefore be willing to pay up to 12.84 $\times$ $1,132 = $14,533, or 0.6 percent of its annual wage bill for this turnover reduction. Put another way, a 10 percent increase in turnover is as costly as a 0.6 percent increase in wages.

A.4.6 Comparison of Turnover Cost Estimates to Bartel et al. (2014)

This Section briefly describes how we construct comparable estimates of human capital versus team 'disruption' effects on team performance in Bartel et al. (2014) and in our own context.

Additional details are available in the on-line spreadsheets for the paper.

Background

Bartel et al. (2014) estimate the effects of three types of month-to-month transitions – a departure without a hire, a hire without a departure, and a departure plus a hire – on the productivity of registered-nurse (RN) teams with an average of nine members.[47] Direct comparisons of their estimated turnover costs with ours are not very meaningful because their performance measure –patients' residual length of hospital stay (LOS)– is very different, and is much less sensitive to team size or composition than ours (net daily sales). We can, however, compare our estimates of the *relative* costs of frequent changes in team membership (switching the identity of a team member) versus the costs of permanently operating with fewer workers, to the same relative costs in Bartel et al. In both cases, note that our calculations capture only the reductions in team performance associated with turnover. Other costs, such as recruiting activities by persons not on the team, or salary costs associated with overlapping old and replacement workers, are excluded for comparability between the two papers.

Calculations for Firm A

To measure the cost of *disruptions* to team membership in Firm A we start with our estimate of the total decline in team productivity associated with a single departure ($1654) and subtract from it short-staffing costs ($400) to obtain productivity losses that are not related to human capital shortages of $1254. We then express this loss as a share of mean monthly team sales, and adjust for the fact that turnover of 1 member of a 3.2-person team represents a departure of 31.3 percent of the workforce. This gives us the estimated cost of continuously turning over ten percent of the sales team each month. This cost, which also represents the cost of losing one member of a 3.2 person team about every three months, equals 2.3 percent of a team's monthly gross sales.

Since we do not have direct measures of the effect of larger sales teams on productivity, we assume that in the long run a ten percent smaller sales team would sell ten percent less. Combining these two estimates means that at Firm A, the continuous turnover of ten percent of the work force is about 23 percent as costly as a permanent workforce reduction of ten percent.

Calculations for the Nursing Teams in Bartel et al. (2014)

To measure the cost of disruptions to team membership in Bartel et al.'s nursing teams we start with the authors' estimated effects of three events –a departure without a hire, a hire without a departure, and a departure plus a hire– on the team's log(residual length of stay) (LOS) from column 1 of their Table 4. Notably, each of these disruptions has a very similar effect on log LOS (of about .0075). These estimated effects control for team human capital levels, so there is no need to explicitly subtract out the short-staffing effects of these events. We then adjust for the fact that turnover of 1 member of a 9-person nursing team (the mean in their sample) represents a bit more than a ten percent turnover rate, and express these cost increases as a share of the mean LOS of 5.92

[47] Bartel et al. also calculate the effect of raising average tenure on a nursing unit on hospital costs; they estimate that costs fall since the reduction in patient bed days would outweigh the increased salary costs. Almost all of this effect, however, stems from the fact that experienced nurses are more productive than novice nurses, which is not the case in our sales context.

days among these teams. We find that the estimated cost of continuously turning over ten percent of a nursing team each month works out to 0.1 percent of a team's mean length of patient stay, a much smaller number than the effect of turnover on sales at Firm A. Intuitively, this is simply because Bartel et al's team productivity measure is much less sensitive to team characteristics than is our measure of sales.

Finally, we compute the effects on LOS of a ten percent permanent reduction in team labor input using the estimated effect of a one-hour increase in RN time on log(LOS) of .0348 (from column 1 of Bartel et al.'s Table 4). Using their reported data on mean hours and mean LOS we then compute that a ten percent reduction in team size would raise mean LOS by about 0.3 percent. Combining these two estimates means that in Bartel et al.'s nursing teams, the continuous turnover of ten percent of the work force is about 35 percent as costly as a permanent workforce reduction of ten percent. Thus, the relative importance of the *disruption* costs of turnover versus the level of human capital input is surprisingly similar in the two papers, and suggests a relatively prominent role for disruption effects.

A.4.7 Implications of Turnover Costs for Reservation Worker Quality

Another strategic decision that may be impacted by turnover costs at Firm A is its reservation worker quality. Specifically, if replacing a salesperson by a new hire costs Firm A a total of C = \$1,132, how bad must an underperforming worker be to make replacing them profitable? To provide a concrete answer to this question, we consider a single employee in a sales team of average size (3.2 workers) who has arrived at the midpoint of an average (829-day) completed career at Firm A; thus she is expected to leave the firm voluntarily in $T = 414$ days. She has just been discovered to be d percent less productive than an average Firm A salesperson (who sells q=\$120 per day). If replacing her results in the hire of a new worker of average quality, what is the threshold level of d at which this replacement is profitable? This threshold level satisfies: $C = Tqd$; solving for d yields $d = 2.3$ percent. This number is so low because turnover costs are only moderate in size, and are very short in duration compared to the cost of operating with an underperforming worker for the remainder of her career.

While these calculations suggest that Firm A should impose a very strict retention threshold, several caveats must be considered before drawing this conclusion. First, these calculations assume zero productivity growth –for example from getting better at the job– over the course of the worker's career. If there was a significant amount of human capital acquisition, a large majority of mid-career workers could, in fact, be above-average producers by the time they reach the middle of their careers at Firm A. That said, our finding that new workers hired from outside Firm A are just as productive as hires from within the firm (Section 4.2) suggests that such growth may be minimal.[48] Second, our estimates of turnover costs are based on voluntary departures only. According to our calculations,

[48]To shed additional light on this question, we also estimated a tenure-productivity profile directly by regressing monthly team sales on the mean tenure of team members with a full set of store and month fixed effects. We did not find any significant relationship, though the standard errors were large.

the severance payments mandated by China's labor law for an average, mid-career worker at Firm A would approximately double the cost of turnover, raising d from 2.3 to 4.4 percent.[49] Potentially more important, the morale costs of employer-initiated dismissals could be much greater than those of voluntary departures.[50] Finally, and perhaps most important, our calculations assume that productivity is an immutable worker characteristic that is completely revealed by mid-career (i.e. after 414 days). But if productivity is imprecisely measured and its underlying mean evolves over time (Kahn and Lange, 2016), replacing workers whose ability currently appears to be only slightly below average may not be a profitable policy.

In sum, if workers' sales abilities have converged to a known and fixed number by the middle of their short careers at Firm A, our calculations imply a very strict underperformance threshold of 2.3 percent, *if* underperforming workers can be induced to leave Firm A voluntarily. Underperformance thresholds for firm-initiated dismissals are 4.4 percent at the very least, and potentially much higher if such dismissals lead to larger morale costs. Together, these patterns may help explain why Firm A chooses to refrain from firm-initiated layoffs, despite the low turnover costs we measure: Firm A's team-based commission policy incentivizes teams (and individual workers) to make these separation choices in a more effective and less costly way.

A.4.8 Estimating Quit-Wage Elasticities at Firm A

To obtain a credible estimate of the causal effect of a higher wage on the share of workers who quit firm A, we need a plausible source of wage variation that is exogenous to Firm A's workers. In our view, the best candidate we have is the month-to-month variation in total compensation per worker that is generated by seasonal fluctuations in the demand for Firm A's product. To use this variation, we collapse the data to the store×month level, then compute total monthly compensation per team member in each store×month cell. Finally, we regress the store's quit rate (share of workers who quit in that month) on the log of per-worker compensation in that month. As in our main analysis, we control for store fixed effects. We also add a fixed effect for the second year of our data (2016) but do not include month effects, because we want to use this seasonal variation. As shown in Table A.4.8, the resulting elasticity estimate was almost exactly equal to minus one, and was statistically significant at the five percent level (for additional details, please see the online spreadsheets that accompany the paper).

[49] Under the Labor Contract Law of the PRC (2007), legally required dismissal compensation is half a month's salary for workers with less than six months of tenure, and a full month for workers with six to twelve months of tenure. Beyond that, one additional month's salary is required per year of firm tenure. Mean tenure of Firm A's retail salespeople is 3.45 years, and the half-month minimum compensation becomes vested at the end of a worker's 15-day probationary period.

[50] A more minor consideration is the fact that our back-of-the-envelope calculations ignore the probability that the newly-hired worker might also need to be replaced. This has two opposing effects on the estimated value of the replacement worker: it overstates her value because it ignores the expected costs of replacing her; but it understates her value because it ignores the option value associated with a new worker of unknown quality (Lazear, 1995)

Table A.4.8: Quit-Wage Elasticities at Firm A

	Dependent Variable: Quit rate in month *t*		
	(1)	(2)	(3)
Log(wage$_t$)	-0.0118**	-0.0117**	-0.0103*
	(0.0055)	(0.0056)	(0.0058
Year FE (2 dummies)	✓	✓	✓
Store FE (118 dummies)	✓	✓	✓
Mean team tenure in month t		✓	✓
Log(wage$_{t-1}$)			✓
N	2022	2022	1911

A.4.9 Sensitivity of Optimal Firm Strategies to Labor Market Conditions

Scenario 1: Tighter labor markets - we double the incidence of short-staffing spells

Table A.4.9 (a): Sensitivity of Optimal Firm Strategies to Tighter Labor Markets

	$M = 4$	$M = 2$	$M = 3$	$M = 5$	$M = 6$
On-boarding Costs	230	207	203	235	244
Short-staffing Costs	520	506	478	516	534
Around-Notice Costs	291	113	222	302	320
Before-Departure Costs	294	186	236	332	328
a) Total Lost *Net* Sales	**1,335**	**1,011**	**1,140**	**1,385**	**1,426**
Administrative Costs:	108	108	108	108	108
Extra Salary when overlap	61	61	61	61	61
Salary Savings when SSed	-224	-224	-224	-224	-224
b) Total Other Costs	**-55**	**-55**	**-55**	**-55**	**-55**
Total Turnover Costs	**1,280**	**956**	**1,085**	**1,330**	**1,371**

For M=4, total turnover costs in the tighter labor market scenario ($1,280) are equivalent to 10.63 days per-employee net sales, or 1.3 percent of career net sales. In terms of wages, this translates to about 71 days of a typical employee's daily wage, or 8.6 percent of her career wages.

With these higher turnover costs, the minimum absolute quit elasticity needed for a wage increase to pay for itself falls from 16.8 to 14.8. The minimum level of employee underperformance, d at which it pays to replace a worker rises from 2.3 percent to 2.6 percent (or from 4.4 to 4.7 percent if mandated severance pay for firm-initiated departures is included.)

Scenario 2: Higher skill requirements - we double the on-boarding costs

Table A.4.9 (b): Sensitivity of Optimal Firm Strategies to Higher Skill Requirements

	$M = 4$	$M = 2$	$M = 3$	$M = 5$	$M = 6$
On-boarding Costs	460	413	407	469	489
Short-staffing Costs	260	253	239	258	267
Around-Notice Costs	291	113	222	302	320
Before-Departure Costs	294	186	236	332	328
a) Total Lost *Net* Sales	**1,335**	**1,011**	**1,140**	**1,385**	**1,426**
Administrative Costs:	108	108	108	108	108
Extra Salary when overlap	61	61	61	61	61
Salary Savings when SSed	-112	-112	-112	-112	-112
b) Total Other Costs	**57**	**57**	**57**	**57**	**57**
Total Turnover Costs	**1,362**	**1,022**	**1,161**	**1,418**	**1,460**

For M=4, total turnover costs in the higher skill scenario ($1,362) are equivalent to 11.35 days of per-employee net sales, or1.4 percent of career net sales. In terms of wages, this translates to about 76 days of a typical employee's daily wage, or 9.2 percent of her career wages.

With these higher turnover costs, the minimum absolute quit elasticity needed for a wage increase to pay for itself falls from 16.8 to 13.9. The minimum level of employee underperformance, d at which it pays to replace a worker rises from 2.3 percent to 2.7 percent (or from 4.4 to 4.9 percent if mandated severance pay for firm-initiated departures is included.)

References

Abraham, S. and Sun, L. (2019). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Unpublished paper, Massachusetts Institute of Technology*.

Acemoglu, D. and Angrist, J. D. (2001). Consequences of employment protection? the case of the americans with disabilities act. *Journal of Political Economy*, 109(5):915–957.

Adamson, B., Dixon, M., and Toman, N. (2014). Why individuals no longer rule on sales teams. *Harvard Business Review*, January 9.

Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar extinction. *Quarterly Journal of Economics*, 125(2):549–589.

Balasubramanian, N. and Sivadasan, J. (2011). What happens when firms patent? New evidence from U.S. economic census data. *Review of Economics and Statistics*, 93(1):126–146.

Barron, J. M. and Bishop, J. (1985). Extensive search, intensive search, and hiring costs: New evidence on employer hiring activity. *Economic inquiry*, 23(3):363–382.

Barron, John M., M. C. B. and Black, D. A. (1997). Employer search, training, and vacancy duration. *Economic Inquiry*, 35:167–192.

Bartel, A. P., Beaulieu, N. D., Phibbs, C. S., and Stone, P. W. (2014). Human capital and productivity in a team environment: Evidence from the healthcare sector. *American Economic Journal: Applied Economics*, 6(2):231–59.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300.

Berg, P., Appelbaum, E., Bailey, T., and Kalleberg, A. L. (1996). The performance effects of modular production in the apparel industry. *Industrial Relations: A Journal of Economy and Society*, 35(3):356–373.

Boning, B., Ichniowski, C., and Shaw, K. (2007). Opportunity counts: Teams and the effectiveness of production incentives. *Journal of Labor Economics*, 25(4):613–650.

Borjas, G. J. and Doran, K. B. (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *Quarterly Journal of Economics*, 127(3):1143–1203.

Brown, M., Falk, A., and Fehr, E. (2004). Relational contracts and the nature of market interactions. *Econometrica*, 72(3):747–780.

Burdett, K. and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, pages 257–273.

Chan, T. Y., Li, J., and Pierce, L. (2014). Compensation and peer effects in competing sales teams. *Management Science*, 60(8):1965–1984.

Dobkin, C., Finkelstein, A., Kluender, R., and Notowidigdo, M. J. (2018). The economic consequences of hospital admissions. *American Economic Review*, 108(2):308–352.

Doz, Y. L. and Guadalupe, M. (2019). Escaping the 's-curve'–is the 'agile' organization the answer? *INSEAD Working Paper*.

Drexler, A. and Schoar, A. (2014). Do relationships matter? Evidence from loan officer turnover. *Management Science*, 60(11):2722–2736.

Dube, A., Giuliano, L., and Leonard, J. S. (2018). Fairness and frictions: The impact of unequal raises on quit behavior. *NBER working paper no. 24906*.

Dube, A., Lester, T. W., and Reich, M. (2011). Do frictions matter in the labor market? Accessions, separations and minimum wage effects. *IZA working paper no. 5811*.

Falch, T. (2010). The elasticity of labor supply at the establishment level. *Journal of Labor Economics*, 28(2):237–266.

Falk, A. and Ichino, A. (2006). Clean evidence on peer effects. *Journal of Labor Economics*, 24(1):39–57.

Friebel, G., Heinz, M., Hoffman, M., and Zubanov, N. (2019). What do employee referral programs do? Technical report, National Bureau of Economic Research.

Friebel, G., Heinz, M., Krüger, M., and Zubanov, N. (2017). Team incentives and performance: Evidence from a retail chain. *American Economic Review*, 107(8):2168–2203.

Friebel, G., Heinz, M., and Zubanov, N. (2016). The effect of announced downsizing on workplace performance: Evidence from a retail chain. *IZA discussion paper no. 9739*.

Friebel, G., Heinz, M., and Zubanov, N. (2018). Middle managers, personnel turnover and performance: A long-term field experiment in a retail chain. *Goethe University*.

Giuliano, L. (2013). Minimum wage effects on employment, substitution, and the teenage labor supply: Evidence from personnel data. *Journal of Labor Economics*, 31(1):155–194.

Glebbeek, A. C. and Bax, E. H. (2004). Is high employee turnover really harmful? An empirical test using company records. *Academy of Management Journal*, 47(2):277–286.

Hamilton, B. H., Nickerson, J. A., and Owan, H. (2003). Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation. *Journal of Political Economy*, 111(3):465–497.

Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, 107(7):1778–1823.

Henry, A. (2011). How much notice should I give my employer that I'm quitting my job? *Lifehacker.com*, (June 20).

Hildreth, A. K. and Oswald, A. J. (1997). Rent-sharing and wages: evidence from company and establishment panels. *Journal of Labor Economics*, 15(2):318–337.

Ichino, A. and Maggi, G. (2000). Work environment and individual background: Explaining regional shirking differentials in a large Italian firm. *Quarterly Journal of Economics*, 115(3):1057–1090.

Ichniowski, C., Shaw, K., and Prennushi, G. (1997). The effects of human resource management practices on productivity:: A study of steel finishing lines. *American Economic Review*, 86:291–313.

Jagannathan, M. (2018). The Rules: How much notice should you give your employer when quitting a job? *Marketwatch.com*, (April 15).

Jäger, S. (2016). How substitutable are workers? Evidence from worker deaths. unpublished paper, Massachusetts Institute of Technology.

Jones, B. F. and Olken, B. A. (2005). Do leaders matter? National leadership and growth since World War II. *Quarterly Journal of Economics*, 120(3):835–864.

Jones, S. R. and Kuhn, P. (1995). Mandatory notice and unemployment. *Journal of Labor Economics*, 13(4):599–622.

Kahn, L. and Lange, F. (2016). Employer learning, productivity, and the earnings distribution: Evidence from performance measures. *Review of Economic Studies*, 81(4):1575–1613.

Kandel, E. and Lazear, E. P. (1992). Peer pressure and partnerships. *Journal of Political Economy*, 100(4):801–817.

Klotz, A. C. and Bolino, M. C. (2016). Saying goodbye: The nature, causes, and consequences of employee resignation styles. *Journal of Applied Psychology*, 101(10):1386.

Knez, M. and Simester, D. (2001). Firm-wide incentives and mutual monitoring at Continental Airlines. *Journal of Labor Economics*, 19(4):743–772.

Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic theory*, 27(2):245–252.

Kuhn, P. and Sweetman, A. (1998). Unemployment insurance and quits in canada. *Canadian Journal of Economics*, pages 549–572.

Kuhn, P. and Yu, L. (2019). Kinks as goals. *work in progress, UC Santa Barbara*.

Lawler, E. and Mohrman, S. A. (2003). Pay practices in Fortune 1000 corporations. *WorldatWork Journal*, 12(4):45–54.

Lazear, E. P. (1995). Hiring risky workers. *NBER working paper no. 5334*.

Lazear, E. P., Shaw, K. L., and Stanton, C. (2016). Making do with less: Working harder during recessions. *Journal of Labor Economics*, 34(1):S333–S360.

Li, C., Qiao, S., and Wang, J. (2016). The dimension of pay satisfaction and its influence on turnover intention of the retail enterprises. *Journal of Capital University of Economics and Business*, 18(1):108–116.

MacLeod, W. B. and Nakavachara, V. (2007). Can wrongful discharge law enhance employment? *The Economic Journal*, 117(521):F218–F278.

Mas, A. and Moretti, E. (2009). Peers at work. *American Economic Review*, 99(1):112–45.

McLaughlin, K. (1991). A theory of quits and layoffs with efficient turnover. *Journal of Political Economy*, 99:1–29.

Moskos, C. C. (1975). The American combat soldier in Vietnam. *Journal of Social Issues*, 31(4):25–37.

Pallais, A. and Sands, E. (2016). Why the referential treatment? evidence from field experiments on referrals. *Journal of Political Economy*, 124(6):1793–1828.

Portugal, P. and Cardoso, A. R. (2006). Disentangling the minimum wage puzzle: An analysis of worker accessions and separations. *Journal of the European Economic Association*, 4(5):988–1013.

Ruhm, C. J. (1992). Advance notice and postdisplacement joblessness. *Journal of Labor Economics*, 10(1):1–32.

Scudder, M. D. (2017). Will the DOL rescind the tip pool rule? *The National Law Review*. October 26, 2017.

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

Siebert, W. S. and Zubanov, N. (2009). Searching for the optimal level of employee turnover: A study of a large UK retail organization. *Academy of Management Journal*, 52(2):294–313.

Ton, Z. (2012). Why ”good jobs” are good for retailers. *Harvard Business Review*, 90(1-2):124–31.

Ton, Z. and Huckman, R. S. (2008). Managing the impact of employee turnover on performance: The role of process conformance. *Organization Science*, 19(1):56–68.

U.S. Bureau of Labor Statistics (2018a). Economic news release table B-1. Employees on nonfarm payrolls by industry sector and selected industry detail.

U.S. Bureau of Labor Statistics (2018b). Employee tenure in 2018 - September 2018. Table 5.

U.S. Bureau of Labor Statistics (2018c). Industries at a glance: Retail trade: NAICS 44-45.

U.S. Bureau of Labor Statistics (2018d). Job openings and labor turnover survey - March 2018. Table 16.

Waldinger, F. (2011). Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany. *Review of Economic Studies*, 79(2):838–861.