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 Clement de Chaisemartin, Dan Hamermesh, Steve Trejo, and participants at the UCLA Trans-Pacific Pacific Labor Seminar and 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.

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

ABSTRACT

Identifying the causal effects of turnover on organizational productivity is challenging, due to data constraints and endogeneity issues. We address these challenges by using day-to-day variation in the composition and performance of small retail sales teams, and by exploiting an advance notice requirement for quits. We find robust and statistically significant productivity losses at four distinct times during the departure process: after the worker gives notice, before she departs, after she leaves, and after a new worker starts. We attribute the first two effects to a combination of recruitment activities by incumbent workers and reductions in morale, and the last two to short-staffing and on-boarding costs. Almost two thirds (63 percent) of these productivity losses occur before the departing worker leaves, and only 24 percent result from operating with an unfilled vacancy. Overall, we estimate that the costs of a ten percent increase in turnover are equivalent to a 0.6 percent wage increase; wage hikes will therefore pay for themselves (in turnover cost savings) only if the elasticity of quits to wages exceeds 16.8 in absolute value.

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
University of California at Santa Barbara
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, in this paper 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. Thus, we can rule out one potential channel of reverse causation: the possibility that under-performing workers are laid off or fired.[4] 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 has two benefits. First, it gives us a sharp date at which the sales team's beliefs that a departure is imminent might jump discretely upwards, and is a natural time to start looking for causal effects of the impending departure.[5] Second, the fact that the 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.[6] Shutting down reverse causation in this way 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, tests for pre-trends in team productivity are joint tests for reverse causation and anticipatory behavior (Abraham and Sun, 2019).

[4]Of course, under-performing workers could be induced to quit in various ways. While Firm A's management does not condone such behavior, team-mates of an under-performing worker might still engage in it. Even in these cases, however, the under-performing worker is required to give two weeks' notice, so we can rely on our advance-notice strategy to isolate causal effects of their actual departure.

[5]Hendren (2017) documents these types of effects (in the form of household consumption and spousal labor responses) in the period leading up to a worker's job loss. In lieu of an announcement date, Hendren relies on infrequently-elicited subjective probability expectations to quantify workers' beliefs; he also marshals data on pre-departure income changes to try to rule out reverse causation ("correlated shocks" in his terminology.)

[6]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.

We develop our results in three stages. In the first (Section 3), we conduct 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. 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. Relative to control days (which are more than 30 days from any departure or hire), we see a two-day productivity loss right after the departure, plus two longer-lasting productivity declines before the departure. These effects –which turn out to be highly robust– occur during the four days surrounding a single employee's required announcement date and during her last four days on the job. Throughout the paper we refer to them as the 'around notice' (AN) 'before departure' (BD) effects respectively.

In the second stage of our analysis (Section 4), we conduct the same nonparametric event study on subsamples of departures, classified by when the replacement worker joined the team. In addition to revealing whether the aggregated results mask different patterns across separation types, these estimates provide preliminary evidence on the likely sources of turnover costs. The resulting patterns are suggestive of two additional types of turnover costs: on-boarding (OB) costs associated with adding a new team member, and short-staffing (SS) costs associated with operating with fewer than the team's normal number of members. Finally, in Section 5, we examine total turnover costs using a parametric event study that is guided by the nonparametric findings, and that allows the four effects described above (AN, BD, OB, and SS) to overlap in time. For example, on-boarding could occur during the last four days before a worker leaves.

Section 6 examines the robustness of our main results and studies heterogeneity in the size of our estimated effects in order to shed additional light on the mechanisms underlying the turnover-induced production losses we observe. Section 7 aggregates all the output reductions associated with turnover, and adds to them some other costs imposed by turnover on this firm. It then assesses the magnitude of our turnover cost estimates by comparing them to other sources of Firm A's costs and revenues, and by estimating the effects of turnover on profits. Section 8 concludes.

Our main findings are as follows. First, 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 happens while a team is short-staffed. 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. Second, while the mechanisms behind the on-boarding and short-staffing effects seem clear, we attribute the 'around notice' and 'before departure' 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, 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, 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.[7] Fourth, 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.[8] 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 findings have a number of implications for research and policy, one of which is methodological: If, as we find, the majority of employee turnover costs are incurred before the departing employee leaves, these costs of turnover might not be detected by an event study that treated the worker's departure date as the event date, and did not have information on when the team was notified of the impending departure. In such a study, we would see pre-trends before the departure date (such as the sharp BD effect in our data), and might interpret them as reverse causation (i.e. as a productivity shock that precipitated the departure.) As a result we might either abandon the study, or (worse) attempt to find covariates, a matching strategy, or a synthetic control that washes out these pre-trends, leading to seriously biased estimates. Thus, our results illustrate how acquiring additional information on the timing of actors' decisions and announcements prior to an event can yield high returns in terms of the credibility of event study estimates.

Second, our findings have implications for models of firm wage effects, i.e. of the tendency for some firms to pay higher wages than others in the same location and industry. One well known explanation of these effects is based on the idea that higher wages can (partially or fully) pay for themselves by reducing turnover costs. This idea has been modeled formally by Burdett and Mortensen (1998), and has been invoked in the business literature to account for the co-existence of high- and low-wage employers like Costco and Walmart in the retail sector (Ton, 2012). Using the median estimate of the quit-wage elasticity in Manning (2003), however, we estimate that employee turnover costs can only account for 3.0 percent of within-industry firm wage effects. Indeed, we

[7] 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). In addition, 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.

[8] 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.

calculate that for wage increases to pay for themselves in terms of reduced turnover costs, the absolute value of the quit-wage elasticity would need to exceed 16.8, a number that substantially exceeds all estimates we know of. Thus, other explanations of firm wage effects, such as an effect of wages on the quality of workers who are hired and retained (Giuliano, 2013), or rent-sharing by employers (Hildreth and Oswald, 1997), are probably needed.

Third, to the extent that our estimated pre-departure output losses reflect short-timer effects, they are relevant to the theoretical and experimental literature on effort provision in teams and finitely repeated games: Our estimated BD effects may reflect a reduction in co-operation among team members towards the end of such a game.[9] Anecdotal evidence of this type of behavior includes accounts of "short-timer's syndrome" in the military during the Vietnam War (Moskos, 1975). This literature describes soldiers' noticeable withdrawal of effort and reluctance to engage in combat in the last two months of their military service.[10] Our study of the effects of anticipated employee departures may provide the first well-identified, quantitative evidence of short-timer effects from a real workplace.

Finally, this paper touches on the effects of advance notice of separations in labor markets. While effects of advance notice requirements for firms who wish to lay off or 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). Both of these results may of course be intended, efficiency-enhancing effects of advance notice. On the other hand, advance notice may simply shift the *timing* of some costs (for example moving on-boarding from after to before the departure), and may even create new costs (such as short-timer effects among about-to-depart workers). Thus, additional research is needed to distinguish whether employee advance notice is a net benefit to firms, and to measure the net costs it imposes on workers.

2 Background and Data

Firm A is a large manufacturer and retailer of men's clothing in China. During the analysis period (January 1, 2015 through December 31, 2016), it operated 164 retail stores, mostly in Guangdong Province. During 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. Figure 1 shows the geographic locations of these 118 stores, on which our analysis is based. We observe 12 store

[9]For example, Embrey et al. (2017) provide a meta-analysis of finitely repeated games and find that there is a significant decline in cooperation in the last round. For indefinitely repeated games, Dal Bó (2005) and Dal Bó and Fréchette (2018) find that cooperation is increasing in the probability of future interactions.

[10]This behavior is formalized by (Cremer, 1986) who argues that team members who are about to leave the organization provide zero efforts in any stationary equilibrium.

openings and 17 store closures in this period, 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.[11]

2.1 Store Sales and Team Production

Retail stores are operated in two types of locations: department stores, or shopping malls, referred to as "host institutions" hereafter. A typical store includes a display of products and a counter, as shown in Figure 2. Sales employees who work at the site are paid by Firm A to promote and sell products. According to Firm A, essentially all the salespeople in our sample are female and between the ages of 20 and 55. Their tasks are similar to the tasks performed by salespersons in the U.S. retail industry, such as greeting customers and recommending products based on customer needs. A major difference in our context, however, is that salespeople play a more active role when customers want to try on products. Typically only one or two items of each style are placed on the rack, so salespeople need to obtain items of the customer's size from the inventory room, which is usually located in a separate storage space at the host institutions. As customers' time constraints and desires for a good fit play an important role in their shopping decisions, both Firm A and its sales workers believe that a store's sales performance is highly associated with salespeople's efforts in quickly and actively responding to customers' needs.

Sales employees work in teams, with team sizes varying from 2 to 7 employees. Firm A sets a target team size for each store, based on how busy the store is likely to be and on the terms of its contract with each host institution. This target size is observed at store-year level, and the actual size of the store is usually consistent with its target size.[12] Figure 3 shows a histogram of the target sizes of stores at the beginning of our analysis period: Around half of the stores in our analysis have a target size of three employees. For each target size, the shares of all store-day observations when a store is actually at, above, or below its target size are shaded. As we can see, in the vast majority of observations, stores are at their target size, when there are likely no team disruptions going on.

In each store, one of the salespeople also acts as the store manager. In addition to her regular

[11] 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).

[12] Some stores do not have a recorded target size in their first year of operation. For these observations, we use the team size 30 days after the opening as the target size.

sales tasks, she also coordinates work schedules for the store.[13] Stores operate from 9am to 9pm, and all the sales employees work exclusively six-hour shifts.[14] Figure 4 shows how work shifts are arranged. 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. Stores operate seven days a week including national holidays, but employees take one day off every week.[15] For 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 temporarily reassigned employees earn a flat amount of $15.60 per shift, paid by Firm A, and receive no commissions on store output.

Figure 5 plots the average daily sales per store on every calendar day in 2015 and 2016. Reflecting the cyclical nature of the retail industry, we observe recurrent fluctuations in daily performance. Total sales are higher in winter and lower in summer, since winter items cost more than summer ones. The smaller cycles in the figure 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.

2.2 Employee Compensation and Turnover

Achieving high levels of sales performance at Firm A requires co-operation among sales workers in at least two ways. One is in setting work and rest-day schedules; doing this well requires team members to be flexible enough to make efficiency-enhancing compromises regarding when they work. The other is on the shop floor: when multiple employees are present, customers can be served more quickly and with greater satisfaction when employees smoothly share the tasks of

[13] Shift schedules are coordinated within each team, and are not observed by Firm A or by us.

[14] Retail employees typically work according to two major staffing schedules in China. 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 employs the second practice in all its stores.

[15] Like the shift schedules, rest-day schedules are coordinated by the team leader within each store, and are not observed by Firm A or by us. In principle these rest days could account for some of the productivity reductions we observe if they coincide with events like the on-boarding of a new employee or the recent departure of an existing one. Based on our conversations with Firm A, we think this is unlikely as the days in which we see productivity reductions are the days on which teams most need their full complement of members. (As detailed in Section 2.3, China's advance notice laws give employers enough leverage over departing workers to ensure they comply with the store's scheduling needs.) Also, as detailed below, in small teams (which are likely more sensitive to rest days), host institutions temporarily loan experienced employees to cover rest days.

fetching items from the store-room and interacting with the customer. Recognizing this, Firm A's retail sales workers are paid based on total team performance; individual performance is neither observed nor rewarded. 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, variations in base payment depend on employees' firm tenures. The store manager, normally the most tenured worker in the team, receives a small additional allowance for coordinating work schedules. Commissions are based on the team's total monthly performance, with commission rates per employee varying across stores between 0.7% and 1.2% of gross sales. Commission rates are identical within a store, but are lower in larger teams, in order to yield similar daily wages across teams of different sizes. Monthly compensation is directly deposited into employees' bank accounts on the 20th of the following month.[16]

Descriptive statistics are summarized in Table 1. Product prices are from a sample of items sold in September, 2016. Given the average product price of $52, a store on average sells 11.4 items per day. While sales performance is higher in stores of more employees, monthly compensation is similar across different store sizes, due to the lower commission rates in larger stores. 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.[17]

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.[18] Therefore the 186 departures we observe are all employee-initiated, voluntary departures. 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.[19] 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). Thus, the quit rate at Firm A is comparable to the industry average in both China and the U.S. As shown in Panel D of Table 1, Firm A's salespeople on average have 3.45 years of firm tenure during our analysis period, and their average tenure is slightly higher in larger teams. In comparison, median years of firm tenure was 3.0 years in the U.S. retail industry in 2018 (U.S. Bureau of Labor Statistics, 2018b). Average team tenure at the

[16] See Chan et al. (2014) for additional information on employee compensation at department stores in China.

[17] The average Guangdong monthly compensation 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 market standard for retail. For example, see a retail industry report at http://sta.doumi.com/src/vip/report/doumi_final01.pdf

[18] The general manager we interviewed gave two reasons for this. First, management believes that dismissal risk is not an effective motivator in Guangdong's tight labor market. Second, the Labor Contract Law of the PRC (2007) would require Firm A to pay dismissal compensation at a rate that Firm A believes is generally not worth the cost: 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.

[19] We use the U.S. Bureau of Labor Statistics definition of the quit rate to compare the U.S. and Chinese rates.

time when departures occur in Firm A is about 2.86 years, while the mean tenure of employees who leave Firm A is about 2.27 years.

2.3 Departure and Hiring Procedures

The departure process from Firm A begins with the worker providing 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 employees have a strong interest in maintaining good relations with their previous employer.[20]

Hiring is done by Firm A's 11 regional managers, with the assistance of the store's employees. As soon as regional managers receive the departure notice, they will start the hiring process by seeking employee referrals from the departing workers' store. Store managers and other employees are happy to assist with this process, 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 by that date. Any impact of these recruiting activities on sales should therefore be reflected in store sales at those times. In addition, management estimates that processing a departure and replacement requires about 12 hours of a *regional manager's* time; we will add these additional costs –which are not included in our estimates of sales team productivity changes– to our estimate of the total costs of turnover in Section 5.[21]

All our empirical analysis in this paper restricts attention to the vast majority of departures from Firm A which occur from stores that are initially at their target size.[22] In these cases, the departing employee is almost always replaced, though the timing varies across stores. Figure 6 provides a flow chart that establishes our terminology for the four possible re-hiring outcomes in a store that is initially at its target size. The four possible rehiring outcomes are:

[20] Firm A's two-week contractual worker notice requirement is actually less strict than the 30 days of notice stipulated by Chinese labor law, (PRC, 2007), but Chinese courts generally defer to the terms of written employment contracts that deviate from standard legal requirements.

[21] According to our interviews with Firm A, their hiring process is not very selective, as the priority is to refill the vacancy. The hiring process consists only of identifying possible candidates, collecting basic paperwork (proof of identity, high-school or above diploma, and health certificate) and a short interview.

[22] As one might expect, most departures from stores that are initially above their target size are not replaced. Stores 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. We exclude these departures from our analysis because we expect turnover costs to differ there, and because these departures are too few in number to allow for effective heterogeneity analysis.

- If the replacement employee starts to work before the actual departure, the team will briefly go above its target size but then return to the target size following the departure. This type of departure/hiring transition is categorized as an ***early refill*** in this paper; early refills comprise **40 percent** of departures from stores that are initially at their target size.

- If the replacement worker starts on the day after the departing employee's last day of work, then the team remains at its target size throughout the departure and post-departure period. We call this an ***on-time refill***; it occurs in **17 percent** of departures from target-size stores.

- For smaller teams (a target size of 2 or 3), if Firm A cannot hire a replacement on time, the host institutions –shopping malls or department stores– will temporarily assign one of their own employees to fill the vacancy; we refer to these workers as ***temporary replacements***. This situation occurs in **18 percent** of departures. Because this process is the same one used to cover rest days in small stores (see Section 2.1), these employees have local sales experience, and may have been temporarily assigned to Firm A before. In sum, smaller stores do not experience a period of short-staffing, even when they fail to hire a permanent replacement by the departure date. [23]

- In larger stores, if a replacement is not hired on time, the team size becomes smaller following the departure, resulting in a period of short-staffing. We refer to such vacancies as ***late refills***; they comprise the remaining **25 percent** of departures from target-size stores.

In Figure 7, we provide more precise detail 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. As noted, the sample is restricted to departures from teams that were initially at their target size. 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.[24]

[23]Temporary replacements are paid a flat amount of $15.60 per shift and receive no commissions on store revenue. Since this is approximately equal to the total per-shift earnings of an average regular team member, we do not include wages paid to temporary replacements in Section 5's estimates of total turnover costs.

[24]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 Aggregate Productivity Trends

3.1 Econometric Approach

As noted, the first phase of our econometric strategy conducts a non-parametric event study of team productivity trends, focusing on the 61-day window surrounding an employee's departure. 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.[25] Also as noted, workers' two-week advance notice requirement allows us to address the endogeneity of turnover because workers' departures cannot be a result of any observed productivity changes during the two weeks prior to their departure: Their decision to leave has already been made. This allows us to cleanly interpret productivity declines during the notice period as anticipatory behavior that is caused by the impending departure.

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. 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 14^{th} day prior to the departure, also on a Friday, or on the preceding Wednesday or Thursday.

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:[26]

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

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

[25] Appendix A.1 shows that the results are robust to both shorter and longer windows than 61 days.

[26] 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 the 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 only 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 in our setting, 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 cyclical patterns in Figure 5, the remaining regression controls in equation 1 are time and store fixed effects. Specifically, D_t is a vector of 731 time dummies for each calendar day in 2015 and 2016, which captures time-varying effects that are common to all stores. In addition to seasonality and day-of-the-week 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.[27] 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. The random, unobserved error is denoted as ϵ_{it}. Robust standard errors are clustered at the store level. In all three stages of our analysis, we also correct for multiple hypothesis testing using the Bonferroni correction procedure described in Benjamini and Hochberg (1995) as a second precaution.[28]

3.2 Results

Estimates of equation 1 are presented in Table 2, and graphed (along with their 95% confidence intervals) in Figure 8. As noted, in this Section the estimation sample includes all departures regardless of when the replacement occurred, 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. As discussed, the goal is to describe productivity patterns during the 61 days surrounding a typical departure while imposing as little structure as possible on when and why team output should rise or fall relative to the control period, after accounting for a

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

[28] For m hypotheses $H1, ..., H_m$ being tested in the specification, we rank the corresponding single-hypothesis p-values $p_1, ..., p_m$ in the order so that $p_{(1)} < ... < p_{(m)}$. Null hypotheses $H_{(1)}, ..., H_{(m)}$ are denoted corresponding to $p_{(1)}, ..., p_{(m)}$. The Benjamini-Hochberg procedure defines k as:

$$k = \max\{i : p_{(i)} \leq \frac{i}{m}\alpha\}.$$

This procedure controls the false discovery rate at α and allows us to reject all hypotheses $H_{(1)}, ..., H_{(k)}$. Benjamini and Yekutieli (2001) further show that this procedure also controls the false discovery rate at α when the tests have dependencies.

complete set of store and year fixed effects.

A first observation from Table 2 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). This suggests that our identifying assumption – that departures cannot have effects on productivity that are too distant in time – is valid. Second, when we do not condition on when the replacement was hired, Table 2 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 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 actually leaves*. Although this may seem counter-intuitive at first, it is consistent with what we know about when departures are replaced at Firm A: 57 percent of departing workers are replaced on or before the day following the departure, and small stores that fail to hire by the departure date fill the vacancy with temporary replacement workers. Thus, short-staffing is rare, which may help explain 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 – together with short-timer effects and recruiting activities by incumbent workers – could help account for the pre-departure losses we see.[29]

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 2. 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 date* 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.[30]

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. In our entire data set, 55.4% of four-day

[29] 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.

[30] 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. This is not true in our case because the date of departure observed is technically the last day when the departing employee works at the site.

* store bins 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, Figure 9(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.[31]

Our final piece of evidence that the sales decline around the mandated notice date is 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 Figure 9(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 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, 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.[32]

4 Disaggregated Productivity Trends

In this Section we replicate the preceding estimates for subsamples of departures defined by when the departing worker was replaced. This reveals patterns specific to the different types of replacements that might be obscured by our aggregate analysis, and provides some preliminary evidence on the likely sources of turnover costs. Specifically, we focus on three subsamples, which shed light on the roles of short-staffing costs, on-boarding costs, and all other productivity costs respectively.

[31] A more surprising feature of Figure 9(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, which is also consistent with the fact that many departures occur in stores in developed cities. 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

[32] 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.

4.1 Late Refills and Short-staffing Effects

Our first disaggregated estimates restrict 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. To remove the influence of on-boarding effects, we further discard all post-departure observations which occur after a replacement worker is hired. Thus, by construction team productivity levels after an employee's departure in this sample should be directly informative about the size and time-structure of short-staffing costs.[33]

Estimates of equation 1 for this subsample of departures are presented in Table 3(a) and plotted in Figure 10(a). Preceding the announcement, we do not detect any pre-trend in sales performance, except for the two days before the required notice date (-16, -15) which most likely represent workers who submit their notice with one or two days to spare. During the notice period, as in the aggregated analysis we observe significant AN and BD effects. Here, however, our main focus is on the post-departure period, which now represents a short-staffing spell for every departure in our sample. During this spell, we now observe a significant productivity loss associated with short-staffing costs that is larger and longer-lasting than in the aggregated analysis. 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. Perhaps surprisingly, this post-departure loss – when the team operates short-staffed – is 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) they 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 6.

4.2 On-time Refills and On-boarding Costs

To examine the size and duration of on-boarding costs, we now focus on the 17 percent of departures that are replaced exactly "on time", i.e. in which the new employee joins the team on the day immediately following the departing employee's last day of work. This is by far the modal day – relative to the departure – that refills took place in our data. In these departures, there are no short-staffing costs because the size of the team never departs from its target size. The identity of one of the workers is, however, permanently changed on day 1. Thus, team performance in this group of departures after the departure date will give us a clean picture of the size and duration of on-boarding effects.

Estimates for this group are presented in Table 3(b) and plotted in Figure 10(b). Again, as in the

[33]Note that our estimated short-staffing effects include the lost output from the departing worker *plus* any associated changes in effort from the remaining workers. These changes could either be compensatory (with the survivors 'doing the work of' the departed employee), or reinforcing, via a reduction in motivation.

aggregate analysis, we see AN and BD effects around the mandated notice date and just before the departure, with an especially large output drop during the leaver's last two days of work.[34] After the departure, there is an acute loss of team productivity in the first week, but the loss is not statistically significant after the first six days. As all the replacement workers in this sample enter the firm on the day after the departure, these results suggest on-boarding costs of about 33.2% of daily team performance that last about six days. In the next three weeks, the coefficient estimates are negative but not statistically significant, suggesting that on-boarding costs, like short-staffing costs, are quite short-lived. Overall, these estimated on-boarding costs are quite similar to the short-staffing costs estimated in Table 3(a), and are about equally short-lived.[35]

4.3 Intact Pre-departure Teams and Other Turnover Costs

Here, we estimate equation 1 for the entire set of departures in which the replacement worker arrived after the departure occurred (i.e. the union of groups *b*, *c* and *d* in Figure 6, comprising 17 + 18 + 25 = 60 percent of all departures). Together, these three groups comprise all the departures in which the sales team remains intact during the notice period, with no arrivals or departures. Thus, none of the pre-departure output changes in this sample can be attributed to on-boarding or short-staffing effects: Some other mechanism must be at work.

Estimation results for this sample are presented in Table 3(c) and plotted in Figure 10(c). Strikingly, we still we see strong reductions in team productivity around the notice (AN) and before the departure (BD), despite the fact that these teams experience no personnel changes during the advance notice period. Specifically, there is a productivity loss of about (100 + 138)/2 = $119, or 20.1% in the four days surrounding the announcement. In the employee's last six days on the job, team productivity declines by about 21.7%. In Section 5 we explore some possible sources of these effects, including time lost due to recruitment and re-scheduling activities by the incumbent workers, and short-timer effects on employee effort before the departure date.

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.

[34]We also see a somewhat puzzling pre-trend in this group: Before the notice date, these teams perform worse than the control period, although the coefficients are not statistically significant except for one.

[35]Notably, the on-boarding costs that are identified in this sample 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 6, we present on-boarding cost estimates that distinguish between these two situations.

5 Parametric Analysis of Daily Team Output

5.1 Econometric Approach

Section 4's results suggest the presence of short-staffing costs during a brief period after a departure occurs, of on-boarding costs during a similarly brief period after a new employee joins the team, and of output reductions in a small number of days surrounding the required notice date and preceding the departure date. The latter effects are present even in teams whose membership remains completely unchanged during the advance notice period. 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:

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

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 (days -(*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.[36] 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

[36] 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. (A similar concern applies to our on-boarding variables, but is quantitatively less important since the vast majority of departures are replaced during our sample period.) 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). The opposite would be true if the departures that were replaced last were for inessential positions that the firm was least motivated to refill. 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 attempt to control for non-random selection into experiencing a short-staffing spell.

in the AN and BD effects into bins, most of which are four days in length.[37] These period effects will measure any causal effects of turnover that are not captured by our AN, BD, OB and SS effects.

5.2 Results

In Table 4, 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.[38] 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 quantitatively significant, they are however quite short in duration.

6 Robustness and Heterogeneity

6.1 Treatment Window Width and Effect Duration

One modeling choice that could potentially 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 61-day window has included 30 days on each side of the departure date, which we denote as a width (W) of 60. In Appendix A.1 (a) - (c), we replicate our aggregate analysis (Table 2) for treatment window widths (W) of 50, 90 or 120. In all these cases, we

[37] 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 4 for details.

[38] 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 involves endogenous timing of replacements: It could be that the short-staffing spells that last the longest are the ones that are least costly to the team.

find 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 confirms the validity of our identifying assumption that departures do not have effects that are too distant in time from the departure date.

A second modeling choice refers to the assumed duration of the four main effects estimated in Table 4's parametric analysis. There we set a duration, M, of four days for the 'around notice' (AN), 'before departure' (BD), early on-boarding (OB), and early short-staffing (SS) effects. In Table 5, we replicate Table 4 for other values of M.[39] Overall, the estimated effects on daily productivity are numerically larger if we focus on a narrower duration, and shrink in size if we use a wider duration, as we would expect if the causal effects were relatively short-lived. Otherwise, the effects are statistically robust and exhibit similar patterns to the main specification. In Section 7, we compute total turnover costs –which depend on both the size and duration of all the component costs– for these all these values of M ; the largest (\$1840) is about 7.5 percent larger than our baseline ($M=4$) estimate of \$1711.

As a final way to summarize the time structure of team productivity trends during the notice period, Appendix A.2 parameterizes these trends in two additional ways. The first divides the notice period into three equal 5-day bins, while the second represents the time trend of productivity during the notice period by a quartic in time. Both specifications show substantial and statistically significant productivity losses near the notice and departure dates, and small, insignificant losses in the middle of the notice period.

6.2 Excluding End-of-Month Departures

As noted earlier, more than half of the departures in our data occur on the last day of a calendar month, which implies that most workers give notice in the middle of the month. While we view this bunching of announcement and departure dates as encouraging evidence that departures are planned in advance – and thus not responses to short-term productivity fluctuations – we are also concerned that our estimated announcement effects (AN) might pick up sales patterns associated with the middle of a calendar month rather than effects of a departure notice. 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 potential weakness, Appendix A.3 conducts additional robustness and placebo analyses to rule out this possibility.

Briefly, Appendix A.3 estimates AN and BD effects in a sub-sample that excludes departures

[39] Like Table 4, the Table 5 regressions include a full set of fixed effects for (mostly) four-day bins that are not included in the AN and BD effects. Bins that are not four days long are at the outer ends of the treatment window (W), and in the middle of the 'notice' interval. 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.

in the last three days of a calendar month. If AN and BD effects are still present in this sub-sample, they are not artifacts of low output in the middle and end of calendar months. We find that estimated AN and BD effects are indeed very similar in this sample of departures that occurred at other times of the month, suggesting that the estimated AN and BD effects are in fact caused by departures. Appendix A.3 also conducts 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 typically given and when departures typically occur. The results show no estimated effects of these placebo dates.

6.3 Who's Leaving? Effects of Rank and Seniority

How a team adjusts to the departure of a member may depend on who is leaving the team. For example, it may be harder for teams to adjust to the departure of a manager than another team member, or to the departure of its most experienced members. In Table 6, we explore how all four of our main estimated turnover costs (AN, BD, OB and SS) differ when the departing employee is the team's manager, or has above-average tenure with Firm A. In column 1, we interact all four of these effects with a dummy for whether the leaver is the team manager; column 2 interacts all four with the seniority of the departing worker.

As shown in column 1, productivity losses near the announcement 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. To the extent that it requires more employee time to find a new manager than a new worker, this supports our interpretation of the around-notice (AN) effect as due to recruitment activities of the incumbent workers. The point estimates indicate that the productivity reduction just before departure (BD) is also greater when a manager is leaving, but this difference is small and statistically insignificant, suggesting that our BD effects may be driven more by morale costs than by incumbent employees' recruiting activities. In addition, short-staffing costs associated with a store managerâs departure are quite large, though the estimate is not statistically significant. Whether the leaver is the store manager does not impact the integration of a replacement worker: on-boarding costs are not different when the leaver is the manager.

In column (2) of Table 6, we interact all four of our main effects with an indicator variable identifying whether the leaver's firm tenure is above the average for all employees in our sample. We find that none of these main effects vary with the seniority rank of the departing employee. Thus, at least in our context, the loss of a manager is more consequential for a team than the loss of a senior member of unspecified rank.

6.4 Who's Coming, and When Do They Start?

The ultimate impact of an employee's departure could also depend on the type of replacement that is eventually recruited, and when that person arrives. In particular, the costs of on-boarding a new worker may depend on whether that worker has relevant experience, and on whether she starts before the departing worker leaves. If she does, there may be valuable opportunities for information transfer that are otherwise unavailable. We address these questions in Table 7. In Column (1), we interact on-boarding with the type of the replacement employee – hired from outside Firm A versus 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 not statistically different for external hires, compared to internal hires or trained replacements. 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 this context, the integration of a new worker into a team (regardless of her qualifications) and not the acquisition of new job-specific skills appears to be the main source of on-boarding costs.

In Column (2), we interact on-boarding effects with a dummy for whether the departing employee is still present in the store. While we might expect the on-boarding process to be facilitated by the continued presence of the departing employee – who could then transfer her knowledge and skills to the newcomer –, column (2) indicates that this is not the case: On-boarding costs are the same in both situations. Thus, it appears that 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.

Also of interest, our column (2) result identifies a channel whereby employee notice requirements may raise overall turnover costs. To see this, note that employee notice requirements likely shift on-boarding events from after to before the departure. While our result says that on-boarding has the same productivity effects regardless of when it occurs, the salary implications of on-boarding at these two different times are not the same: Early on-boarding is more expensive, because the firm must pay two salaries while the departing and new employees overlap.

6.5 Hiring Urgency and Team Size

Some insight into the possible causes of our around-notice and before-departure effects might be available from studying their interaction with two features of the turnover event. One of these is the urgency to hire: If the around-notice or before-departure effects do indeed represent recruitment activities by team members, they should be largest when the need to hire is greatest. Here, we exploit this idea by focusing on the four days before the departure, and comparing the productivity of teams who are still looking for a replacement with those who have already hired one. To do this we replicated column 2 of Table 4, adding an interaction between BD and whether the replacement arrived before day -4. As shown in Table 8, we do not observe a statistically significant effect on the interaction term: During the BD period when the departure has already been filled, there is an insignificant productivity improvement of $57, relative to the main BD effect of $135. Thus, while the point estimate suggests that eliminating the pressure to recruit mitigates the team productivity losses that are observed just before the departure, we cannot reject that the BD effect is the same regardless of whether a replacement has already been hired. Like our results for managers' departures –which show that AN but not BD effects are larger when the departing worker is the team manager –, this suggests that morale-based factors, like a short-timer effect, are the primary drivers of the productivity losses we observe during the departing employee's last few days on the job (i.e. our estimated BD effects).

Another possible source of information about the mechanisms underlying our estimated AN and BD effects is their interaction with team size. For example, if the costs of recruiting a single employee do not vary with team size, and if the AN and BD effects represent recruiting costs, we should not expect these costs to increase with team size. Likewise, if the AN and BD effects represent a fixed amount of shirking *by the departing employee* due to short-timer effects, there is again no clear reason to expect them to increase with team size. Under some other scenarios, however, these productivity reductions might increase with team size. For example, if all members of a team must be consulted in the recruitment process, the costs of recruiting a single employee could increase with team size. And if shirking by the departing worker is 'contagious' to her teammates, short-timer effects can increase with team size as well.[40] To distinguish these scenarios, we study the effects of team size in Table 9.[41]

The results in Table 8 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. This pattern suggests that either the costs of recruiting a single worker increase with team size, or that shirking is contagious from the departing member to other team members.[42] Notably, the contagion scenario is consistent with the shift arrangements at Firm A, described in Figure 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

[40]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).

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

[42]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 Table 1).

be paired with at least one other employee.[43] In sum, while the estimated effects of team size on the AN and BD effects cannot conclusively distinguish between morale-based versus recruiting-activity-based mechanisms for those effects, the team size interactions strongly suggest that one or both of these mechanisms is magnified in larger teams. Since other evidence suggests that AN effects primarily represent recruiting activities but BD effects do not, it is possible that both sources of magnification are at work.

7 Assessing Cost Magnitudes

To assess the economic significance of the OB, SS, AN and BD effects we have estimated, we start by multiplying by the magnitude of each effect on (daily) sales 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 4), Appendix A.4.1 calculates this number as $1654.[44] 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 part because periods of short-staffing are quite rare–, only 24 percent of the lost output associated with a departure ($400/$1654) is associated with short-staffing. This finding has implications for how Burdett-Mortensen-type search models are formulated: By conceptualizing the cost of posting a lower wage as a lower probability that a job is occupied at any particular moment, these 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 dramatically illustrates the fact that pre-trends in event studies require careful interpretation: Without supplementary information on the extent to which the event was expected, pre-trends do not distinguish between spurious correlations (i.e reverse causation) and causal behavioral responses to an anticipated event.

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 do two things: Convert our estimated reductions in gross sales into net sales, and incorporate costs of turnover (such as administrative costs and temporary changes in wage costs) that do not take the form of team productivity reductions. Turning first to the gross-to-net sales conversion, we consider

[43] In the three-employee case, if the departing employee works in the afternoon, she interacts with the another employee for 100% of the time. If she is doing the morning or evening shift, then she interacts with another employee for 50% of the time. Thus she has a chance of: $\frac{1}{3} \cdot 100\% + \frac{2}{3} \cdot 50\%$, i.e. $\frac{2}{3}$, to work with another employee.

[44] Appendix Appendix A.4.2 assesses the robustness of this number to alternative assumptions about effect duration, M. The largest estimate, for M=6, is 7.8 percent higher than our baseline estimate.

a situation where Firm A experiences one extra departure in a particular year, leading to a $1654 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 .65 * $1654 = $1075.[45]

We next account for three additional components of total turnover costs: Recruiting time spent by Firm A's regional managers ($108), the extra salary costs that are incurred while the departing and replacement employee overlap in the firm ($61), and the savings in wage costs while a team is operating short-staffed (-$112). Combining these costs 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 sales or 1.1 percent of an employee's gross 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 necessary 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.

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.

We can, however, make some comparisons with Bartel et al. (2014), which is arguably the best-identified study of 'normal' turnover costs to date. In registered-nurse (RN) teams with an average of nine members, 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 log of residual patient length of stay (LOS). Importantly, these estimates control for team staffing levels, so they do not include any short-staffing effects. When we use these estimates to calculate the effects of turning over ten percent of a nursing team each month on a unit's total costs (which are proportional to LOS) we find very small effects of about one tenth of one percent.[46] Effects of a ten percent turnover rate on Firm A's revenues, however, are much larger

[45]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.

[46]Additional details on these and the following calculations are provided in Appendix A.5.

at 2.3 percent. Intuitively, the main reason for the difference is simply that Bartel et al's measure of team performance (mean length of hospital stay) is not very sensitive to nursing team size or composition.

Abstracting from this overall difference in the elasticity of performance measures to team characteristics, we can however compare our estimates of the relative importance of *disruptions* to teams (changes in membership) versus steady-state team size, (*human capital* levels) to Bartel et al.'s. In our data, we make this comparison by calculating the sales costs of turnover excluding short-staffing effects and comparing them to the sales effects of permanently adding another employee (assuming constant returns). For Bartel et al.'s setting, we compare their regression estimates of disruption costs to their estimates of additional RN hours when both are included in the same regression. Interestingly, here our estimates are much more similar: In our data, replacing ten percent of a team's workforce with a different employee each month 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 workforce *disruption* in turnover costs: Just changing the identity of a given share of a team's members each month has effects on team performance that are 23 to 35 percent as large as permanently shrinking the team by the same fraction of its workforce.[47]

A final way to express the magnitude of our estimated turnover costs is via the trade-off between turnover costs and wages. For example, 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 x \$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.

This trade-off between wages and turnover raises the question of whether a wage increase in Firm A's industry could pay for itself in terms of lower turnover costs. To answer this question, we consider a worker earning a per-period (e.g. annual) wage of w. She quits (and is replaced) within the period with probability q. Finally, if she quits she imposes a total cost of T on the employer. The impact of a one-dollar increase in w on the employer's expected turnover costs associated with this worker's position can then be written as:

$$\frac{d(qT)}{dw} = q \times \frac{T}{w} \times \eta \qquad (3)$$

[47]Bartel et al. also calculate that raising average tenure on a nursing unit would result in a net savings in hospital costs, since the estimated 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.

where $\eta < 0$ is the elasticity of quits with respect to wages. Inserting values of *w*, *q* and *T* from our data and Manning's (2003) median estimate of η yields:[48]

$$\frac{d(qT)}{dw} = 0.34 \times 0.175 \times -0.5 = -0.03 \tag{4}$$

Thus, using Manning's median estimate, a one-dollar wage increase reduces turnover costs by only three cents because (a) turnover is not a very common event (q=.34), (b) turnover costs are a small share of wages (T/w =.175), and (c) turnover is not very sensitive to wages (η=-0.5). Put another way, turnover costs can account for 3.0 percent of firm wage effects. If, in contrast, we use the largest well-identified estimates of η we know of (η = -3.5, from studies of schoolteachers by Falch (2010) and Clotfelter et al. (2008)) this number rises to 21 percent.

We conclude from these calculations that wage increases cannot reduce turnover costs enough to pay for themselves at Firm A, even under the most optimistic assumptions concerning the quit-wage elasticity. Indeed (from inverting equation 4), a full 'pay-for-itself' would require an elasticity of -16.8, a magnitude which exceeds all estimates we are aware of. Thus, factors other than turnover costs, 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 why some firms persistently pay more than others in the retail sector.

8 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 sizable as a share of daily sales, they are, however, quite 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

[48] 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.

can mitigate turnover costs in a number of ways we have documented in this paper. Indeed, this cost mitigation is facilitated by the advance notice workers give before leaving. Finally, because these costs are so low, we calculate that raising wages to reduce turnover is not a profitable strategy in these jobs for any plausible value of the quit-wage elasticity. In sum, one reason why turnover is high (and wages are low) in front-line retail jobs is the simple fact that turnover is not very costly to employers in that environment.

Figure 1: Map of Stores in the Analysis

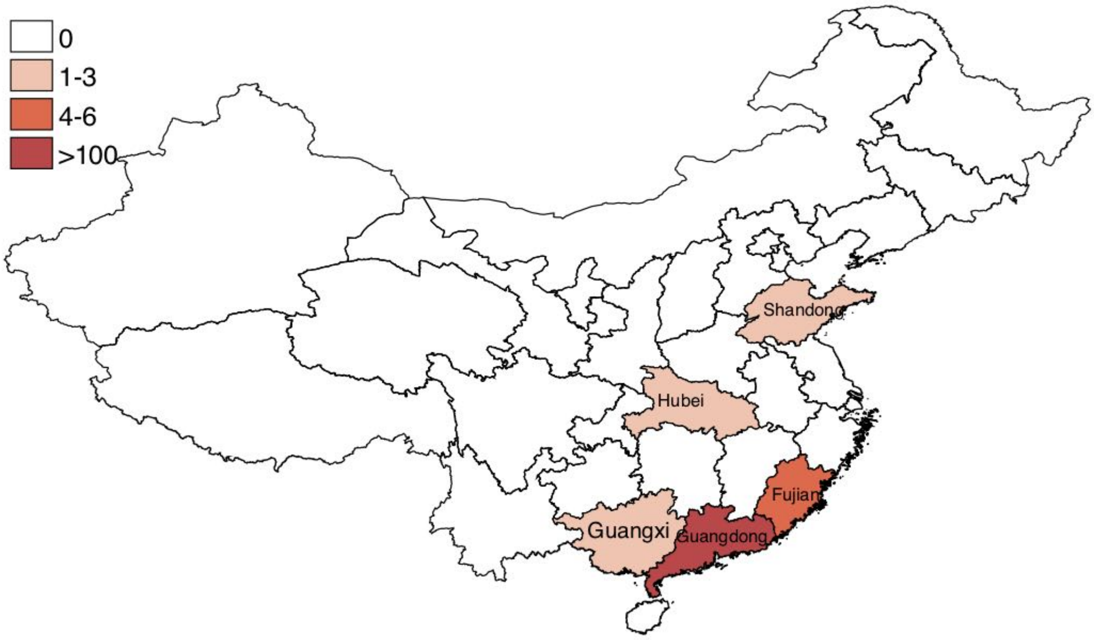

(a) Map of stores in China

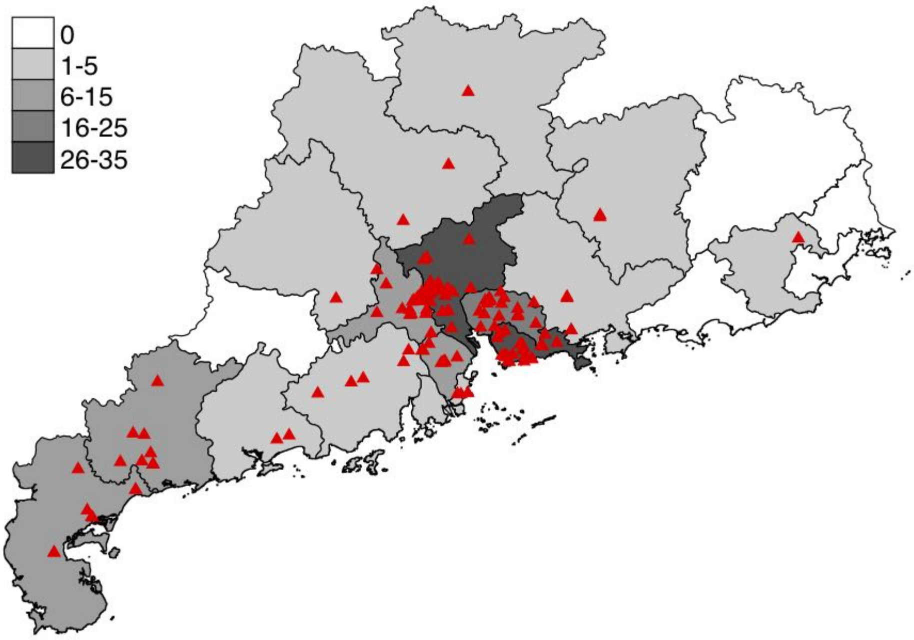

(b) Map of stores in Guangdong

Figure 2: Photos of Stores

(a) A typical store of 2-3 employees

(b) A typical store of 4 or more employees

Figure 3: Histogram of Store Sizes

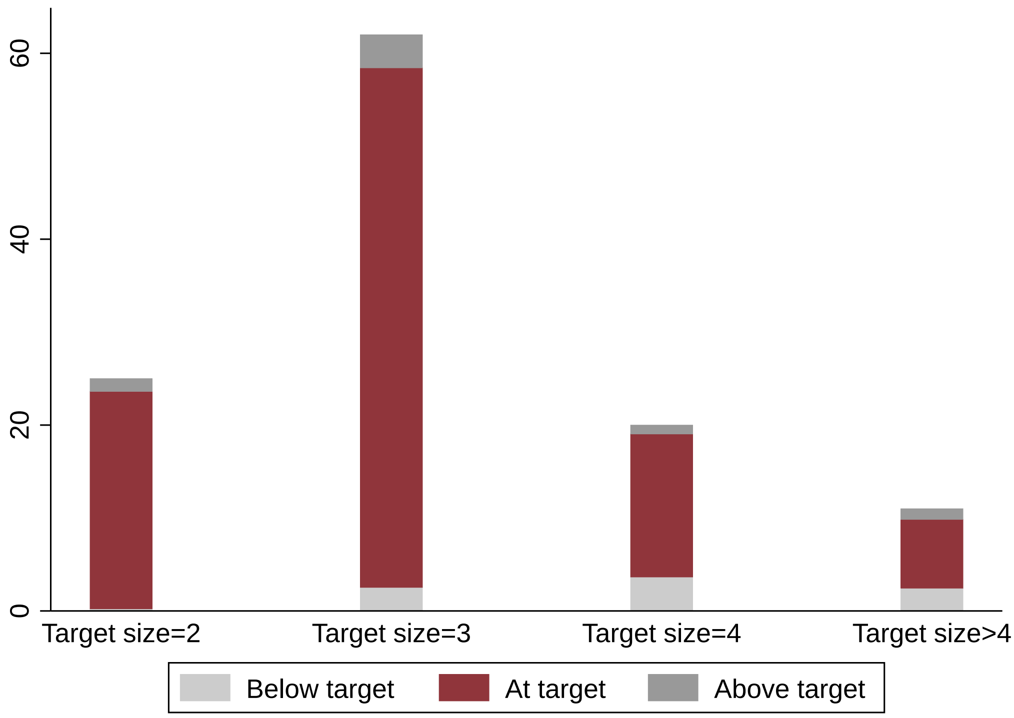

Notes: The histogram categorizes stores according to their target size at the beginning of the analysis period. The vertical axis indicates the number of stores with each target size. The share of stores that are above, at, or below their target size are indicated by different colors. Target size is observed from the annual sales plan, at the store-year level. 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 instead.

Figure 4: Shift Arrangements

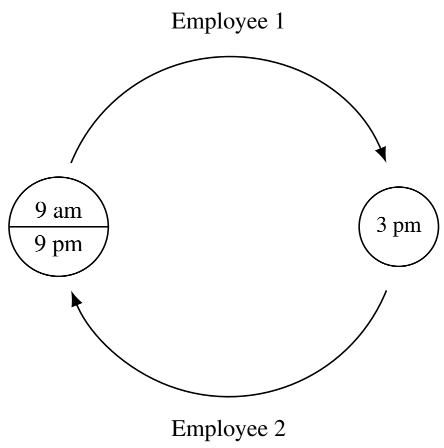

(a) Target size=2

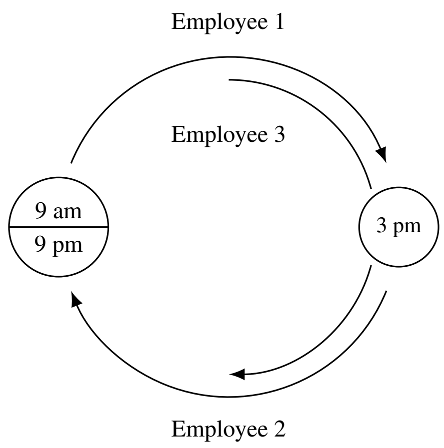

(b) Target size=3

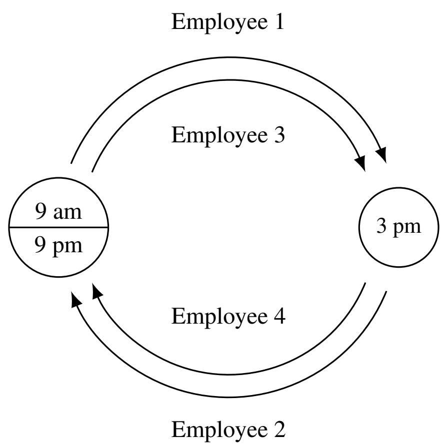

(c) Target size=4

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 5: Average Daily Sales per Store by Calendar Day

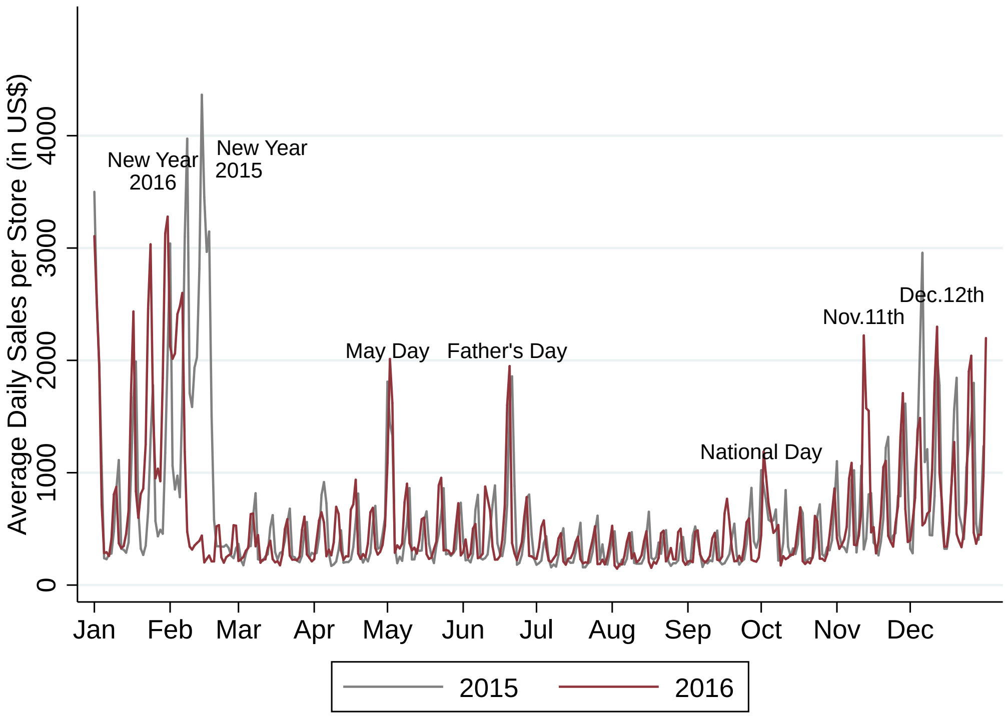

Notes: This figure plots average daily sales per store on every calendar day in 2015 and 2016. The labeled spikes correspond to holidays or major shopping events. New Year denotes the Chinese Lunar New Year. November 11th and December 12th are major shopping events in China, similar to Black Friday or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

Figure 6: Flow Chart of Notice, Departure, and Replacement Procedures

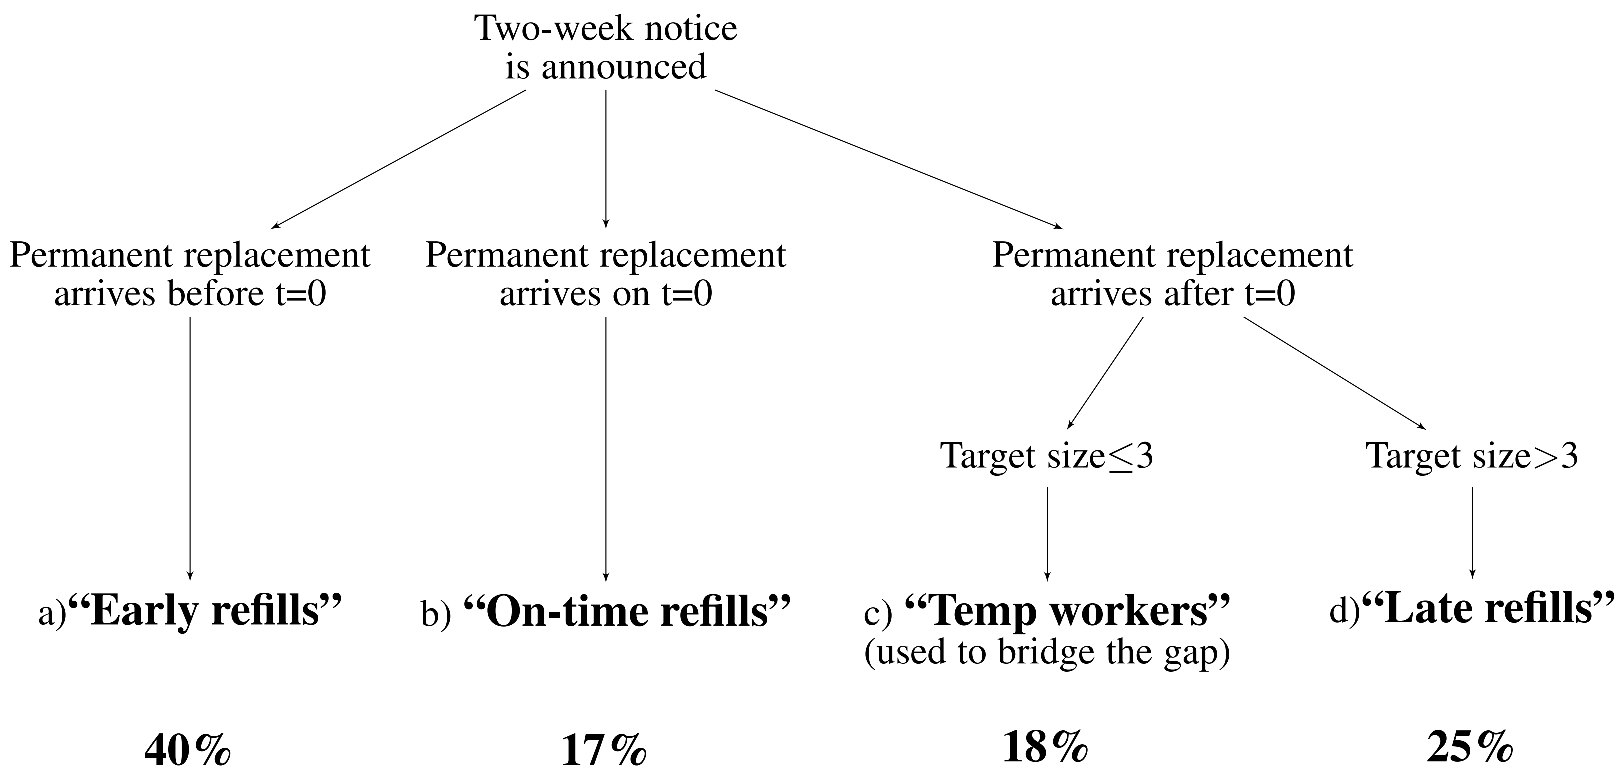

Notes: The procedures in this figure apply to departures from stores that are at their target size. In these stores, group (d) is the only one that experiences short-staffing. In groups (b), (c) and (d), the entire sales team remains unchanged throughout the two-week notice period.

Figure 7: Share of Departures Replaced as a Function of Time since Notice

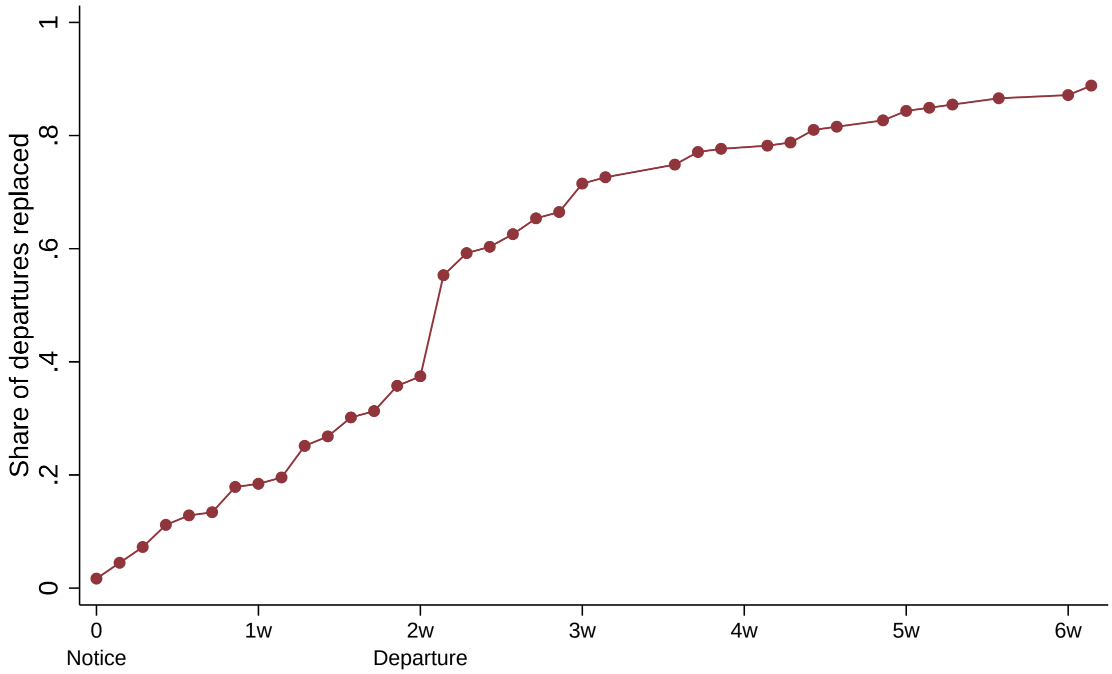

Notes: This figure plots the share of departures that have been replaced as a function of elapsed time since the required notice date. The sample is restricted to stores that were at their target size on that date. The 17 percentage point jump on the day after the departure represents *on-time refills*,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

Figure 8: Aggregate Productivity Trends

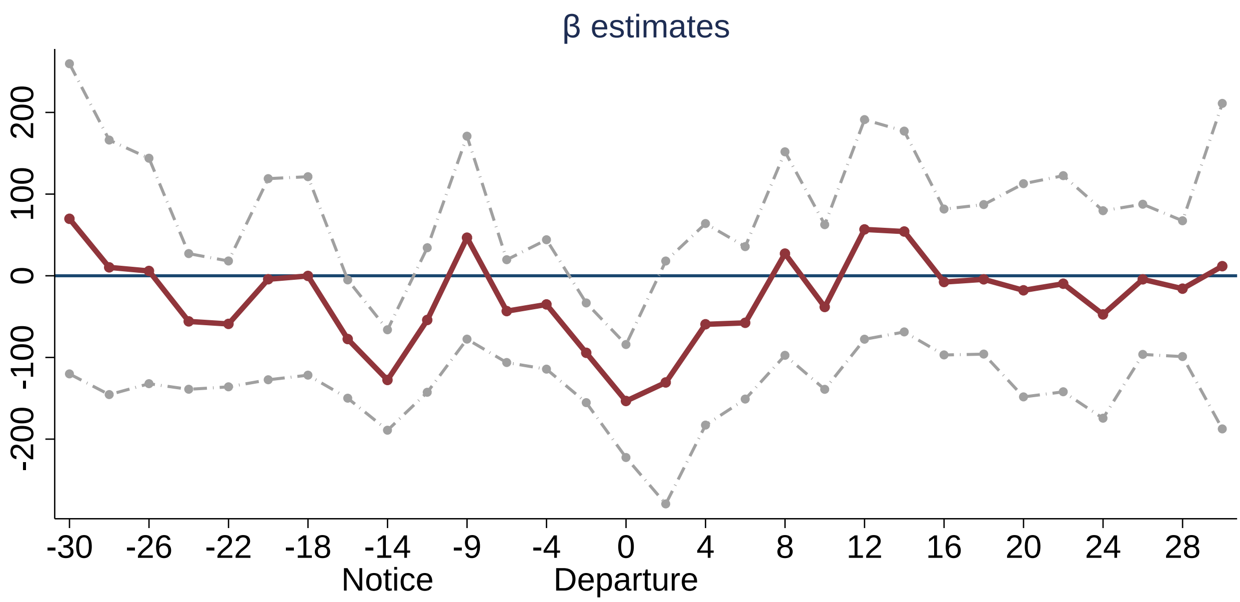

Notes: This figure plots coefficients of the 30 lead and lag terms in equation 1, along with 95 percent confidence intervals. It is based on the regressions reported in Table 2. The "Departure" bin denotes the departing employee's last two days of work, and the "Notice" bin denotes the first two days of the required notice period. The specification controls for 731 calendar day dummies and 118 store fixed effects. Confidence intervals are constructed from robust standard errors, clustered at the store level.

Figure 9: The Timing of Departures

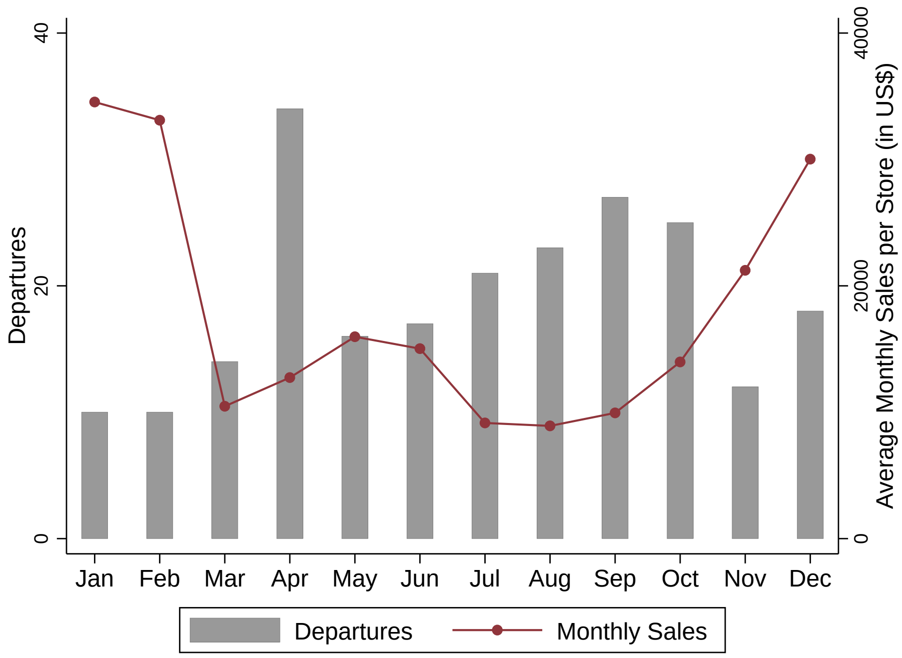

(a) Month of Departures

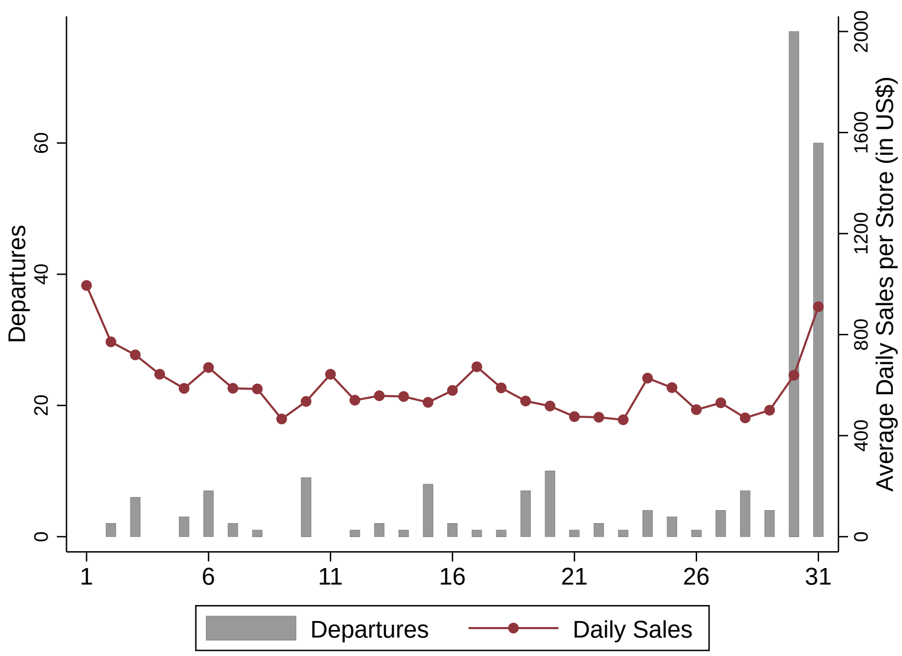

(b) Day of Departures

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

Figure 10: Disaggregated Productivity Trends

(a) Late Refills and Short-staffing Effects

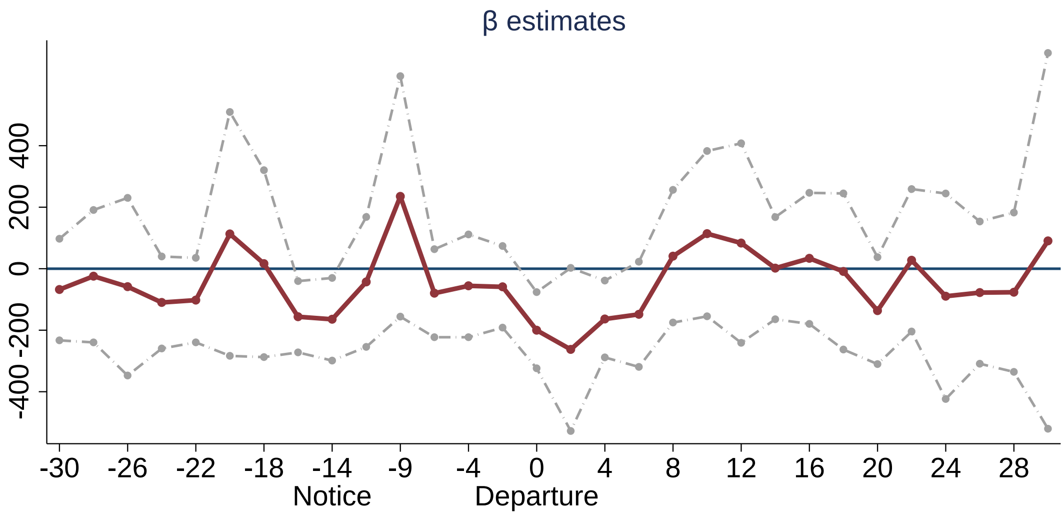

Notes: This figure plots coefficients from a regression similar to equation 1, for the late refills sample. This sample comprises departures that result in at least one day of short-staffing following the departure, and we only include observations that occur *before* replacement workers are hired. The figure is based on the regression results in Table 3(a), which include 731 calendar day dummies and 118 store fixed effects. Robust standard errors are clustered at the store level.

(b) On-time Refills and On-boarding Costs

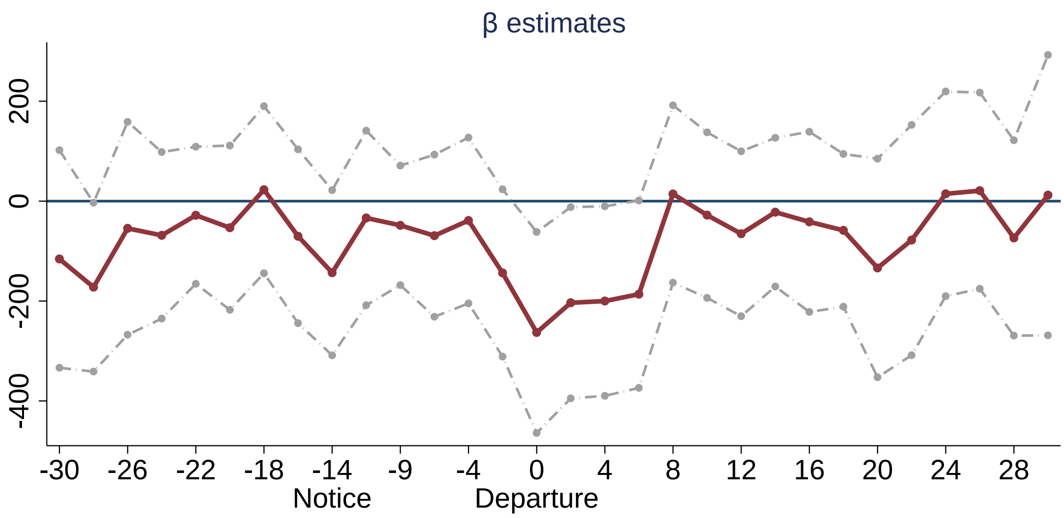

Notes: This figure plots coefficients from a regression similar to equation 1, for the on-time refills sample. This sample includes only the departures in which replacement workers are hired on the day after the departing employee's last day of work. The figure is based on the regression results in Table 3(b), which include 731 calendar day dummies and 118 store fixed effects. Robust standard errors are clustered at the store level.

Figure 10: Disaggregated Productivity Trends – Continued

(c) Intact Pre-departure Teams and Other Turnover Costs

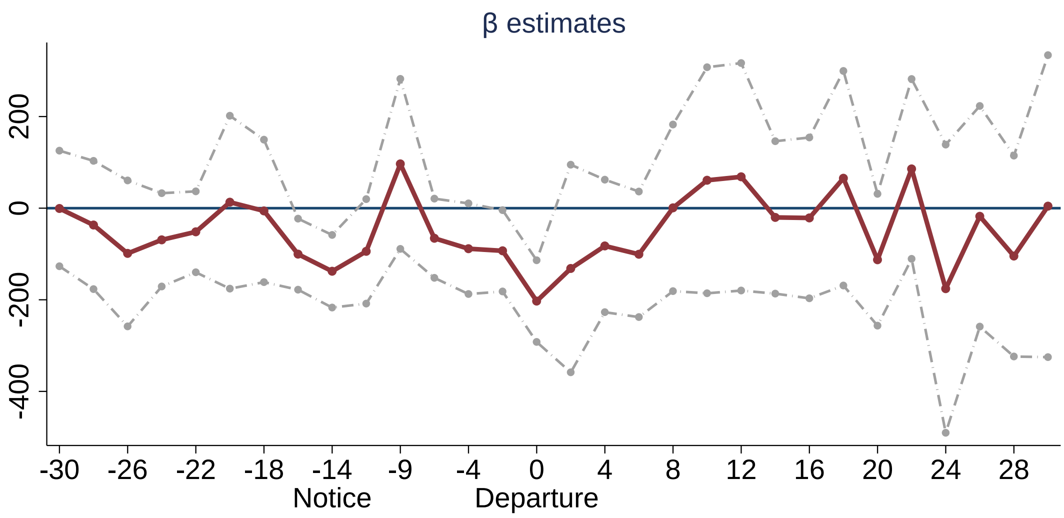

Notes: This figure plots coefficients from a regression similar to equation 1, for intact pre-departure teams. This sample excludes departures that were refilled early, so there were no changes in team size or composition before the departure date. The figure is based on the regression results in Table 3(c), which includes 731 calendar day dummies and 118 store fixed effects. Robust standard errors are clustered at the store level.

Table 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 size in Panels (B) and (C) is 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.

Table 2: Aggregate Productivity Trends

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, with estimated coefficients plotted in Figure 8. $P_{-14,-13}$ 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 3(a): Disaggregated Productivity Trends – *Late Refills*

Dependent variable: daily sales (in US $)					
Before the Departure:			After the Departure:		
	β	SE		β	SE
$P_{-30,-29}$	-68	(84)	$P_{1,2}$	-263*	(135)
$P_{-28,-27}$	-24	(110)	$P_{3,4}$	-164**	(64)
$P_{-26,-26}$	-58	(147)	$P_{5,6}$	-148*	(87)
$P_{-24,-23}$	-110	(76)	$P_{7,8}$	41	(110)
$P_{-22,-21}$	-102	(70)	$P_{9,10}$	114	(137)
$P_{-20,-19}$	113	(202)	$P_{11,12}$	83	(166)
$P_{-18,-17}$	16	(155)	$P_{13,14}$	2	(85)
$P_{-16,-15}$	-156***	(59)	$P_{15,16}$	34	(109)
$P_{-14,-13}$	-165**	(69)	$P_{17,18}$	-9	(130)
$P_{-12,-11}$	-43	(108)	$P_{19,20}$	-136	(89)
$P_{-10,-9}$	235	(199)	$P_{21,22}$	27	(118)
$P_{-8,-6}$	-80	(73)	$P_{23,24}$	-90	(170)
$P_{-5,-4}$	-56	(85)	$P_{25,26}$	-78	(118)
$P_{-3,-2}$	-59	(68)	$P_{27,28}$	-77	(132)
$P_{-1,0}$	-200***	(63)†	$P_{29,30}$	90	(312)
N				60496	

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

Notes: This table presents results from estimating equation 1 for late refills, with estimated coefficients plotted in Figure 10(a). This sample includes departures that result in at least one day of short-staffing, and we only include observations that occur *before* replacement workers are hired. 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 3(b): Disaggregated Productivity Trends – *On-time Refills*

	Dependent variable: daily sales (in US $)				
Before the Departure:			After the Departure:		
	β	SE		β	SE
$P_{-30,-29}$	-116	(111)	$P_{1,2}$	-203**	(98)
$P_{-28,-27}$	-172**	(86)	$P_{3,4}$	-200**	(97)
$P_{-26,-26}$	-54	(109)	$P_{5,6}$	-186*	(96)
$P_{-24,-23}$	-68	(85)	$P_{7,8}$	14	(91)
$P_{-22,-21}$	-28	(70)	$P_{9,10}$	-28	(85)
$P_{-20,-19}$	-53	(84)	$P_{11,12}$	-65	(84)
$P_{-18,-17}$	23	(85)	$P_{13,14}$	-22	(76)
$P_{-16,-15}$	-70	(89)	$P_{15,16}$	-41	(92)
$P_{-14,-13}$	-143*	(84)	$P_{17,18}$	-59	(78)
$P_{-12,-11}$	-34	(89)	$P_{19,20}$	-134	(112)
$P_{-10,-9}$	-49	(61)	$P_{21,22}$	-78	(118)
$P_{-8,-6}$	-69	(83)	$P_{23,24}$	15	(105)
$P_{-5,-4}$	-39	(85)	$P_{25,26}$	21	(100)
$P_{-3,-2}$	-144*	(85)	$P_{27,28}$	-74	(100)
$P_{-1,0}$	-263**	(103)	$P_{29,30}$	12	(143)
N			59984		

$^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$); $^{\dagger}\ q < 0.1$, $^{\dagger\dagger}\ q < 0.05$, $^{\dagger\dagger\dagger}\ q < 0.01$.

Notes: This table presents results from estimating equation 1 for on-time refills, with estimated coefficients plotted in Figure 10(b). This sample includes only departures that are replaced on the day after the departing employee's last day 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 3(c): Disaggregated Productivity Trends – *Intact Pre-departure Teams*

	Dependent variable: daily sales (in US $)				
Before the Departure:			After the Departure:		
	β	SE		β	SE
$P_{-30,-29}$	-1	(64)	$P_{1,2}$	-132	(116)
$P_{-28,-27}$	-37	(71)	$P_{3,4}$	-82	(74)
$P_{-26,-26}$	-99	(81)	$P_{5,6}$	-101	(70)
$P_{-24,-23}$	-69	(52)	$P_{7,8}$	1	(93)
$P_{-22,-21}$	-52	(45)	$P_{9,10}$	61	(126)
$P_{-20,-19}$	13	(96)	$P_{11,12}$	68	(127)
$P_{-18,-17}$	-6	(79)	$P_{13,14}$	-20	(85)
$P_{-16,-15}$	-100**	(40)	$P_{15,16}$	-21	(90)
$P_{-14,-13}$	-138***	(40)†††	$P_{17,18}$	65	(120)
$P_{-12,-11}$	-94	(58)	$P_{19,20}$	-113	(73)
$P_{-10,-9}$	97	(95)	$P_{21,22}$	86	(100)
$P_{-8,-6}$	-66	(44)	$P_{23,24}$	-176	(160)
$P_{-5,-4}$	-89*	(50)	$P_{25,26}$	-18	(123)
$P_{-3,-2}$	-93**	(45)	$P_{27,28}$	-105	(112)
$P_{-1,0}$	-203***	(45)†††	$P_{29,30}$	4	(168)
N			62787		

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

Notes: This table presents results from estimating equation 1 for intact pre-departure teams, with estimated coefficients plotted in Figure 10(c). This sample excludes departures that are refilled early, so the team size and team composition remain intact before the actual departure occurs. 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: 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)†††	(32)†††
Before Departure$_4$ (P$_{-3,0}$)	-114***	-113***
	(25)†††	(26)†††
Early On-boarding$_4$	-79**	-76**
	(32)†	(36)
Late On-boarindg$_4$	-2	-5
	(36)	(35)
Early Short-staffing$_4$	-240***	-218*
	(86)††	(111)
Late Short-staffing$_4$	-38	-41
	(54)	(55)
Period Effect Controls:		
P$_{-30,-28}$		35
		(86)
P$_{-27,-24}$		15
		(53)
P$_{-23,-20}$		-62*
		(33)
P$_{-19,-16}$		-5
		(49)
P$_{-11,-8}$		2
		(39)
P$_{-7,-4}$		-15
		(40)
P$_{1,4}$		-24
		(70)
P$_{5,8}$		5
		(48)
P$_{9,12}$		27
		(46)
P$_{13,16}$		33
		(45)
P$_{17,20}$		-5
		(46)
P$_{21,24}$		-10
		(67)
P$_{25,28}$		1
		(43)
P$_{29,30}$		17
		(101)
	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 Notice*$_4$ identifies a 4-day period including the required notice date. *Before Departure*$_4$ identifies the last 4 days the departing employee spending at the team. *Early On-boarding*$_4$ identifies the first 4 days a new employee is present, and *Late On-boarding*$_4$ identifies the 5th day through two weeks after the hiring. *Early Short-staffing*$_4$ identifies the first 4 days following an unfilled departure, and *Late Short-staffing*$_4$ identifies other days following *Early Short-staffing*$_4$ with the departure unfilled through 30 days following 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 ($\dagger$), correct for false discovery following the Benjamini-Hochberg procedure.

Table 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)
Period Effect Controls:				
$P_{-30,-28}$	35	35	34	34
	(86)	(86)	(86)	(86)
$P_{-27,-24}$	14	14	16	15
	(53)	(53)	(54)	(53)
$P_{-23,-20}$	-62*	-62*	-62*	-61*
	(33)	(33)	(33)	(33)
$P_{-19,-16}$	-4	-4	-4	-4
	(49)	(49)	(49)	(49)
$BW2_M$	-45	-19	8	19
	(34)	(36)	(46)	(57)
$BW1_M$	-39	-30	-0	-6
	(31)	(35)	(43)	(46)
$P_{1,4}$	-45	-39	-28	-35
	(60)	(64)	(68)	(67)
$P_{5,8}$	9	6	19	28
	(51)	(50)	(50)	(52)
$P_{9,12}$	32	29	29	31
	(48)	(48)	(49)	(49)
$P_{13,16}$	37	35	35	37
	(49)	(49)	(49)	(49)
$P_{17,20}$	-1	-3	-4	-3
	(48)	(48)	(48)	(48)
$P_{21,24}$	-6	-7	-9	-10
	(69)	(69)	(69)	(69)
$P_{25,28}$	4	3	3	2
	(44)	(44)	(44)	(44)
$P_{29,30}$	21	20	19	19
	(100)	(100)	(100)	(100)
N	75801	75801	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 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 Table 5, all regressions include a full set of fixed effects for (mostly) four-day bins 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, denoted as *BW2$_M$* and *BW2$_M$*). 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 includes 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.

Table 6: 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)†
Manager	-16	
	(45)	
AN$_4$ $\times$ Manager	-133**	
	(61)†	
BD$_4$ $\times$ Manager	-26	
	(78)	
OB$_4$ $\times$ Manager	16	
	(72)	
SS$_4$ $\times$ Manager	-222	
	(138)	
Seniority		16
		(49)
AN$_4$ $\times$ Seniority		19
		(94)
BD$_4$ $\times$ Seniority		-36
		(61)
OB$_4$ $\times$ Seniority		-40
		(93)
SS$_4$ $\times$ Seniority		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 7: Heterogeneity Examination – *On-Boarding Effects*

	Dependent variable: daily sales (in US$) (1)	(2)
Early On-boarding$_4$	-78**	-101**
	(34)†	(42)†
Late On-boarding$_4$	1	-6
	(39)	(38)
Trained	13	
	(47)	
OB$_4$ $\times$ Trained	-30	
	(114)	
LOB$_4$ $\times$ Trained	-58	
	(100)	
Leaver's working		23
		(50)
OB$_4$ $\times$ Leaver's working		71
		(68)
LOB$_4$ $\times$ Leaver's working		4
		(67)
N	75801	75801

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

Notes: All regressions include controls for around-notice effects, before-departure effects, and early and late short-staffing effects, as defined in equation 2. 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.

Table 8: **Heterogeneity Examination – *Hiring Urgency***

	Dependent variable: daily sales (in US $)
	(1)
Around Notice_4	-120***
	(33)†††
Before Departure_4	-135***
	(32)†††
Early On-boarding_4	-87**
	(34)††
Late On-boarding_4	-12
	(38)
Early Short-staffing_4	-240***
	(86)††
Late Short-staffing_4	-38
	(53)
Vacancy Filled	27
	(45)
BD_4 $\times$ Vacancy Filled	57
	(77)
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.

Table 9: 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.

A Appendix

A.1 Robustness Check: Treatment Window Width *W*=50, 90, 120

A.1.1 *W*=50

Table A.1.1: 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.1.2 *W*=90

Table A.1.2: 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.1.3 *W*=120

Table A.1.3: 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^{***}	$(34)^{\dagger\dagger\dagger}$	$P_{47,48}$	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^{***}	$(32)^{\dagger}$	$P_{57,58}$	76	(77)
$P_{-1,0}$	-151^{***}	$(35)^{\dagger\dagger\dagger}$	$P_{59,60}$	-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 Alternative Specifications of the Pre-departure Productivity Loss

The time points in this Appendix are defined the same way as in the paper:

- P_0 is the day of departure, which is the last day the departing employee working at the team.
- $P_{\text{-}14}$, is the last day on which the departing employee is allowed to notify Firm A of her departure. Firm A has no data on the exact dates when it receives the notice, but it asserts that notice is submitted exactly two weeks in advance of the departure, or one or two days earlier.

We start by estimating the following specification

$$S_{it} = \alpha + \beta \cdot PreDep + \gamma_1 \cdot D_t + \gamma_2 \cdot I_i + \epsilon_{it}, \tag{A.2.1}$$

where $PreDep$ is a binary variable, which takes a value of 1 if the current observation is within the 14-day notice period, and takes a value of 0 elsewhere. The D_t are the 731 day fixed effect dummies, and the I_i are the 118 store fixed effect dummies, both as defined in the main analysis.

The baseline sample is as defined in the main analysis in the paper. Thus β estimates the productivity loss during the notice period, relative to the control observations which are more than 30 days distant from an employee departure.

A.2.1 A Binned Approach

To examine which portions of the notice period exhibit the greatest productivity losses, here we create three 5-day binned variables and estimate the following specification:

$$\begin{aligned} S_{it} = {} & \alpha + \beta_3 \cdot Dep_{-3} + \beta_2 \cdot Dep_{-2} + \beta_1 \cdot Dep_{-1} \\ & + \gamma_1 \cdot D_t + \gamma_2 \cdot I_i + \epsilon_{it}, \end{aligned} \tag{A.2.2}$$

where Dep_{-1} is a 5-day binned variable, taking a value of 1 if the current day is within the last five days that the departing employee is still working at the store. Dep_{-3} is a 5-day binned variable, which takes a value of 1 if the current day is within the five days following the required notice day. Finally, Dep_{-2} identifies the five days between Dep_{-1} and Dep_{-3}.

In Table A.2.1, Column (1) presents estimates of equation A.2.1, and Column (2) presents estimates of equation A.2.2. On average, column (1) shows that during the entire notice period, team sales performance is reduced by an average of \$70, 11.8% per day, relative to control days. However, as column (2) shows, this loss is quite concentrated at the two ends of the period. During the middle of the notice period (i.e. when Dep_{-2} =1), team productivity is not significantly different from its baseline levels.

Table A.2.1

	Dependent variable: daily sales (in US$)	
	(1)	(2)
$PreDep$	-71***	
	(25)	
Dep_{-3}		-81***
		(29)†††
Dep_{-2}		-16
		(32)
Dep_{-1}		-116***
		(31)†††
N	60801	60801

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Notes: This table presents results from estimating equation A.2.1 and A.2.2. $PreDep$ identifies the 15 days within the advance notice to departure period. Dep_{-3} identifies the five days following the advance notice; Dep_{-1} identifies the departing employee's last five days of work; and Dep_{-2} identifies the five days between Dep_{-1} and Dep_{-3}. 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.2.2 A Polynomial Approach

An alternative way to model the productivity trend during the advance notice period treats time more continuously, using the following specification:

$$\begin{aligned} S_{it} = \alpha + \beta \cdot PreDep + \psi_1 \cdot d + \psi_2 \cdot d^2 + \psi_3 \cdot d^3 + \psi_4 \cdot d^4 \\ + \gamma_1 \cdot D_t + \gamma_2 \cdot I_i + \epsilon_{it}, \end{aligned} \tag{A.2.3}$$

where *d* equals one if the bin falls within the notice period and zero otherwise. After estimating equation A.2.3, we then compute point estimates for every single day during this period. The estimates and 95% confidence interval are plotted in Figure A.2.1, for both cubic and quartic specifications of the polynomial in equation A.2.3. As in the 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 of the period.

Figure A.2.1: All Departures

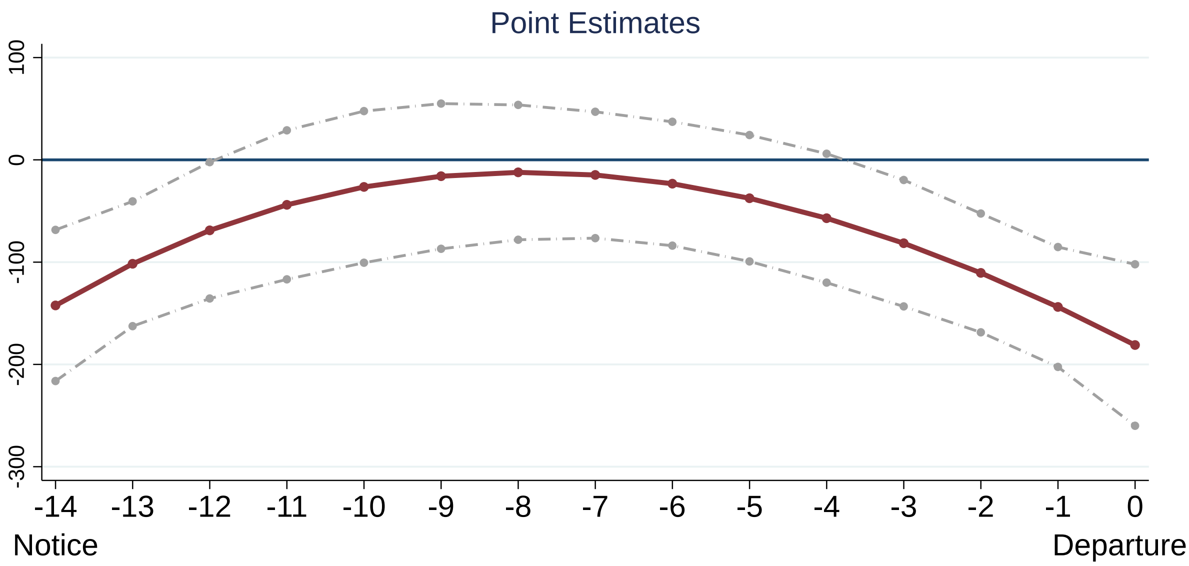

(a) Cubic

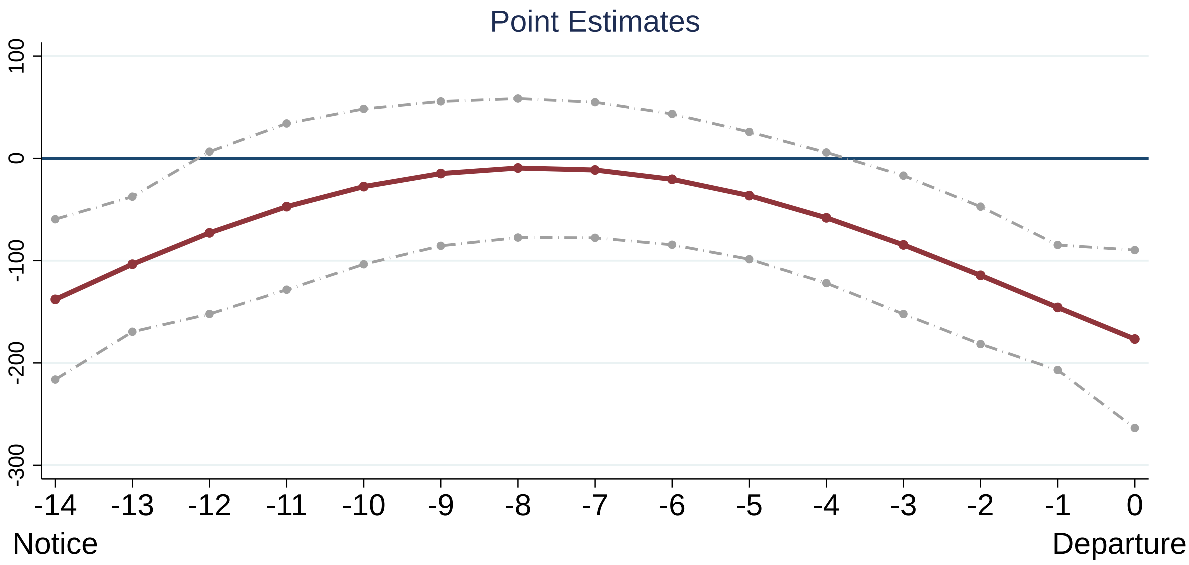

(b) Quartic

A.3 Assessing the Role of End-of-the-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 main 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.3.1 we report the results of two alternative specifications that address this issue. As a baseline for these two exercises, column 1 of Table A.3.1 displays the main coefficients from our baseline estimates of equation (1), taken from Table 2.

In column (2) of Table A.3.1, 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.3.1 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.3.1: Effect of Peer's Advance Notice on Team Productivity

	Dependent variable: daily sales (in US $)		
	(1)	(2)	(3)
	All	Robustness check	Placebo test
$P_{-22,-21}$	-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
	$(31)^{\dagger\dagger\dagger}$	(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
	$(31)^{\dagger\dagger}$	(50)	(88)
$P_{-1,0}$	-153^{***}	-118^{**}	88
	$(35)^{\dagger\dagger\dagger}$	(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.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 4 (AN, BD, OB and SS) by the duration of each of those sales declines.[49] Our results are as follows:[50]

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

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

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 5 regressions. The following cost estimates are based on estimates for M = 2, 3, 5 and 6 from Table 6.

[49] In Table 4, most effect durations are assumed to be 4 days (alternative assumptions are explored in Appendix A.1.) 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.

[50] 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.

- M=2:
  OB Costs = $2 \times 81 + 12 \times 13 = \$318$
  SS Costs = $0.50 \times 241 + 5.07 \times 53 = \$389$
  AN Costs = $2 \times 87 = \$174$
  BD Costs = $2 \times 143 = \$286$
  Total Lost Sales = $(318 + 389 + 174 + 286) =$ **\$1167**

- M=3:
  OB Costs = $3 \times 75 + 11 \times 8 = \$313$
  SS Costs = $0.74 \times 204 + 4.83 \times 45 = \$368$
  AN Costs = $3 \times 114 = \$342$
  BD Costs = $3 \times 121 = \$363$
  Total Lost Sales = $(313 + 368 + 342 + 363) =$ **\$1386**

- M=5:
  OB Costs = $5 \times 47 + 9 \times 14 = \$361$
  SS Costs = $1.20 \times 211 + 4.35 \times 33 = \$397$
  AN Costs = $5 \times 93 = \$465$
  BD Costs = $5 \times 102 = \$510$
  Total Lost Sales = $(361 + 397 + 465 + 510) =$ **\$1733**

- M=6:
  OB Costs = $6 \times 28 + 8 \times 26 = \$376$
  SS Costs = $1.48 \times 197 + 4.12 \times 29 = \$411$
  AN Costs = $6 \times 82 = \$492$
  BD Costs = $6 \times 84 = \$504$
  Total Lost Sales = $(376 + 411 + 492 + 504) =$ **\$1783**

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. 0.65 times \$1654 = \$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 = 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.

A.4.5 Assessing Cost Magnitudes

In this Section we express the above cost estimates –which are denominated in dollars of net revenues– in terms that can be more easily compared to other firms and industries: *days'* worth of employee net sales, and days of employee wages. 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. Finally, turning to wage-based measures, at $18 in net sales per employee per day, $1,132 translates to 63 days of per-employee sales. Thus, the cost of a single turnover can also be expressed as 7.6 percent of an employee's career wages at Firm A.

It is noteworthy that 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. 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.5 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. A spreadsheet containing our detailed calculations is available from the authors. Note that all these calculations measure 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 contexts.

A.5.1 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.

A.5.2 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 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.

References

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

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*, 23:363–382.

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.

Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, 29(4):1165–1188.

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.

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.

Clotfelter, C., Glennie, E., Ladd, H., and Vigdor, J. (2008). Would higher salaries keep teachers in high-poverty schools? Evidence from a policy intervention in North Carolina. *Journal of Public Economics*, 92(5-6):1352–1370.

Cremer, J. (1986). Cooperation in ongoing organizations. *Quarterly Journal of Economics*, 101(1):33–49.

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review*, 95(5):1591–1604.

Dal Bó, P. and Fréchette, G. R. (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1):60–114.

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

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

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*.

Embrey, M., Fréchette, G. R., and Yuksel, S. (2017). Cooperation in the finitely repeated prisoner's dilemma. *Quarterly Journal of Economics*, 133(1):509–551.

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

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.

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. and Oswald, A. (1997). Rent-sharing and wages: Evidence from company and establishment panels. *Journal of Labor Economics*, 15(2):318–37.

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.

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.

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.

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.

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.

Manning, A. (2003). *Monopsony in motion: Imperfect competition in labor markets*. Princeton University Press, Princeton, NJ.

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

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

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.

PRC (2007). Labor Contract Law of the People's Republic of China. Article 37.

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.

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.