"I Quit":
The Role of Schedule Volatility in Employee Turnover*

Alon Bergman Guy David Hummy Song

August 23, 2021

Abstract

We examine how employer-driven volatility in workers' schedules impacts their decision to voluntarily leave their job. Using novel time-stamped work log data of salaried home health nurses, we construct a measure of schedule volatility that can be easily implemented by firms. This measure may be endogenous to the worker's decision to quit. We instrument for schedule volatility using time off taken by other workers at the same office, which exogenously increases schedule volatility for all remaining workers. We find that a single 30-day period of high schedule volatility (defined as one standard deviation above the average level of schedule volatility) increases the average worker's probability of quitting that year by approximately 25%. Our results are robust to different measures of schedule volatility and model specifications. Through policy simulations, we illustrate how schedule volatility, and in turn, employee turnover, could be mitigated.

*Alon Bergman (corresponding author): Postdoctoral Scholar, University of Pennsylvania, The Wharton School; alonberg@wharton.upenn.edu. Guy David: Professor, University of Pennsylvania, The Wharton School; gdavid2@wharton.upenn.edu. Hummy Song: Assistant Professor, University of Pennsylvania, The Wharton School; hummy@wharton.upenn.edu.

1 Introduction

Changes to the organization of service delivery and production processes have reduced workers' control over their schedules at a time when the demand for flexibility in the workplace is on the rise. Only 56% of full-time workers report having flexible work hours (Council of Economic Advisors, 2010), with less than 20% of salaried workers having a flexible schedule in which they were able to frequently vary the times they began and stopped working based on their own preferences (U.S. Bureau of Labor Statistics, 2019). A significant contributing factor is the widespread use of shift work, temporary workers, and just-in-time scheduling by employers (Katz and Krueger, 2019), which improves matching of supply to demand in real time, and in turn raises both worker productivity and firm profitability.

While attractive to employers, this practice has the potential to make work schedules more volatile and to disrupt both pay dynamics and work-life balance. Reports show that nearly half of workers in the U.S. regularly experienced employers deciding their schedules without their input (Golden, 2015; U.S. Bureau of Labor Statistics, 2019) and 17% of workers had unstable work shift schedules, with the lowest income workers facing the most irregular work schedules (Golden, 2015; Yu et al., 2021). Little control over their own work schedules lowers workers' attachment to a given firm, raises dissatisfaction, heightens work-family conflict, and subsequently may increase the likelihood of job switching (Barmby et al., 2002; Golden, 2015; Henly and Lambert, 2014). In recent years, we find some workers have chosen to exit the traditional labor market altogether in favor of the gig economy, where workers can have more control over their own work schedules (Hall and Mueller, 2018; Chen et al., 2019). The preference for employee-driven flexibility and aversion to employer-driven schedule volatility is particularly salient for women, for whom time conflicts between work and family responsibilities are more pronounced (Landers et al., 1996; Mas and Pallais, 2020; Henly and Lambert, 2014). While accommodating worker preferences can be costly to firms, its absence may lead to undesirable outcomes for the employer as well. A recent field experiment illustrates that less stable work schedules are linked to lower levels of labor productivity (Kesavan et al., 2021). High levels of employee dissatisfaction may also lead to higher levels of employee turnover, which highlights an important tradeoff for employers between the stability of the labor force and its efficiency.

Given the pervasive nature of shift-based work, the evolution of just-in-time scheduling, and the growing emphasis on work-life balance, understanding how the volatility in workers' schedules may lead to job quitting is timely. In this paper, we directly measure the effect of schedule volatility on the probability that an employee would voluntarily separate from the employer (in other words, quit), by using a novel and detailed time-stamped dataset of salaried home health nurses' schedules. Our data allow us to construct a novel measure of schedule volatility, make causal inference to its effect on turnover, and provide counterfactual analyses to assess the impact of different scheduling policies. By focusing on salaried employees, whose wages are independent of the number of hours actually worked during a given pay period, we are able to study schedule volatility holding pay constant.

When turnover arises from employees voluntarily choosing to leave the firm, as opposed to the firm terminating their employment, high levels of turnover pose a significant problem for firms (Jovanovic, 1984; McLaughlin, 1991). A sizable body of research documents the negative effects of turnover on performance (Hom et al., 2017; Griffeth et al., 2000), including substantial productivity losses (Black and Lynch, 1996), operational disruption (Kozlowski et al., 1999), and organizational forgetting (David and Brachet, 2011). Turnover also necessitates hiring and retraining, both of which are costly to the firm (Deere, 1987). There are also negative downstream effects from the demoralization of the remaining employees, who may in turn exhibit lower levels of productivity and turnover themselves (Staw, 1980).

To measure the effect of employer-driven schedule volatility on employee turnover, we identify a setting that exhibits both high levels of employee turnover and externally-driven volatility in work schedules, reflecting employer-dictated schedules that are only made available on short notice as opposed to employee-driven flexible work hours. Our setting is that of full-time salaried nurses working in home health care. This service setting has seven key unique characteristics that allow us to address our research question. First, full-time nurses in this setting are salaried, their pay ungoverned by the number of hours they work in a given time period. This allows us to study the effects of schedule volatility on turnover independently from any potential effect of volatility in pay.[1] Second, we are able to observe the detailed schedules and work logs of nurses, as nurses log their activities in real time using tablets. This allows us to index schedule volatility for each nurse. Third, since home health nurses work independently, whereby each visit to a patient's home is conducted by a single nurse at a time, we can attribute the outcome of interest (i.e., voluntary turnover) to the individual nurse's schedule volatility at the same level of observation. This is in contrast to many other service settings where employees work in groups or teams and adhere to a common schedule that varies at the team level and not at the individual level. This also reduces the influence of peers in attrition decisions. Fourth, by design, visits are assigned to nurses' schedules the evening of the previous day (i.e., less than 24 hours in advance). The lack of control over their schedule makes it such that variability in schedules captures employer-driven schedule volatility rather than a preference for flexibility. Fifth, while some of the schedule volatility occurs naturally due to the stochastic arrival of patients, it is also influenced by the various objectives used by the firm when scheduling nurses. This allows us to simulate various counterfactual scheduling policies that prioritize different objectives. Sixth, in part due to the high rates of turnover in this setting, we are able to restrict our analyses to nurses who start their employment within our study period. As a result, we are able to observe a complete history of each nurse's employment spell. Seventh, the detailed human resources data allow us to separate nurses who quit from those who were laid off. In other words, our outcome is able to specifically capture the turnover that is undesirable to the firm. Note, while our analyses are contained to the home health care setting, we expect our findings to be highly generalizable to other occupations and industries that rely on shift-based work. We discuss the

[1] See (Conroy et al., 2021) for a recent study on the effect of pay volatility on employee turnover.

generalizability of our results to other economic sectors in section 7.

Using our unique data, we propose a novel approach for measuring schedule volatility and then quantify its causal impact on voluntary employee turnover. We use more than 3 years of highly granular data from one of the largest home health agencies in the United States, in which we are able to observe exactly at what time and for how long each day nurses spent visiting patients, driving between patient visits, and doing other work (e.g., documenting notes), in addition to observing whether the employee separation was voluntary or involuntary. We define schedule volatility as it relates to a short term (28-day) variability in key scheduling characteristics, which include the number of hours spent conducting different work activities (visiting patients, driving, or other assignments); total day length and total daily hours of work; work start and end times; and two characteristics that are unique to the health care setting, the number of daily home visits, and the daily average of time spent per patient. We measure this variability using a coefficient of variation calculated over daily characteristic values over a 28-day period. Our results are qualitatively robust to extending or narrowing the measurement period length.

Causal identification of the effect of schedule volatility on voluntary separation is threatened by both omitted variable bias (e.g., variation in scheduling policies or diverse patient populations may explain both volatile schedules and high rates of turnover) and reverse causality (e.g., workers who are (unobservably) about to quit, may be phased out by having fewer cases assigned to them, leading to a mechanical correlation between a planned attrition and more volatile schedules). The directionality of the bias generated by such behaviors is unclear, since these might increase or decrease volatility depending on how short- or long-lived they are.

To recover causal estimates for the effect of schedule volatility on workers' voluntary turnover, we use an instrumental variable approach. We instrument for workers' schedule volatility using Paid Days Off (PDOs) taken by other workers in the same branch; this exogenously increases the focal worker's schedule volatility because it is the branch that is ultimately responsible for meeting patient demand. Specifically, we find that one additional PDO taken at a branch increases a focal worker's total work time that day by an average of 5.2%. Using this instrument, and controlling for branch and month-year fixed effects and additional controls, we find that a one standard deviation (SD) increase in schedule volatility, defined as the coefficient of variation of the daily number of visits conducted over the past four weeks, would increase the average worker's propensity to quit on a given day by more than three-fold. We find similar effects when using several other measures of schedule volatility. Translated into yearly terms, a single 30-day period of high schedule volatility (defined as one SD above the average level of schedule volatility) increases the average worker's probability of quitting that year by 25.45% (SE: 5.08 p.p.).

Incorporating these estimates into simulations, we quantify the mitigating effect that different scheduling policies could have in reducing voluntary turnover by lowering schedule volatility. We study three distinct policies: (1) a policy in which visits (and therefore hours) are redistributed equally among workers on duty that day; (2) a policy that forces workers to redistribute

their week's visits equally across all days on which they are on duty; and (3) a policy in which "excess" visits, defined as the number of visits above the median count of visits per worker-day at the branch that quarter, are offloaded to part-time workers at the branch, if any are available on the branch's roster and have the capacity for additional visits. While all three policies may have potential unintended consequences, each is successful in reducing schedule volatility, and therefore, is able to create more predictable and stable work schedules. We calculate the counterfactual schedule volatility resulting from each policy under consideration, and find the potential impact on worker retention to be substantial, reducing daily schedule volatility by between 11% and 34% and daily worker propensity to quit by 5% to 16%, depending on the proposed scheduling policy.

This paper contributes to our understanding of the extent to which employees value control over their own work schedules and are averse to unpredictable work schedules that are dictated by employers. Consequently, it relates to the literature on the value of non-pecuniary job attributes and labor supply preferences, including hours-wage tradeoffs and compensating differentials for shift work (Altonji and Paxson, 1988; Kostiuk, 1990). Hall and Krueger (2018) show that non-pecuniary job attributes vary substantially across jobs and that they comprise an important role in the process of matching employees to employers. When it comes to alternative work arrangements in particular, much of the literature has found that workers report that they are willing to take a significant pay cut to avoid working in jobs where employers dictate workers' schedules (Mas and Pallais, 2017; Maestas et al., 2018; Wiswall and Zafar, 2018; Chen et al., 2020). Some of this prior work also documents meaningful differences by gender, where women are willing to pay more for control of their work schedule.[2] Seen through the lens of this literature, our findings could be interpreted to suggest that workers continuously update their beliefs of their aversion to schedule volatility, or their knowledge of the schedule volatility inherent in their job, sometimes concluding that the value of the outside option exceeds the value of their current position.

The rest of the paper is organized as follows. Section 2 describes the data sources, defines the analysis sample, and presents key summary statistics. Section 3 offers details on how we define and measure schedule volatility. We describe the empirical strategy in section 4, and present results in section 5. In section 6, we present a series of counterfactual analyses by simulating alternate scheduling policies. Section 7 discusses the generalizability of our results to other economic sectors, and section 8 concludes.

[2] Maestas et al. (2018) estimate that workers value the ability to set their own schedule at approximately 9%, with minor differences by gender. Wiswall and Zafar (2018) find that women are more willing to take a larger pay cut when going from a job with no part-time option to one that does have one (7.3% for women compared to 1% for men). In an experimental study in the employment process for a national call center, Mas and Pallais (2017) find that workers have a strong dislike of jobs that permit employer discretion in scheduling, where the average applicant would rather take a 20% wage cut to avoid such positions. Leveraging a natural field experiment at Uber, Chen et al. (2020) estimate that the average driver would require a 21% wage increase to prefer a set schedule position over a flexible schedule position, with women drivers requiring a higher wage increase then men (38% vs. 19%).

2 Data and Analysis Sample

2.1 Data

The proprietary data we use for our analyses come from one of the five largest national home health care agencies in the United States and spans a period of January 2016 to March 2019. By the end of this period, this home health agency was operating more than 250 branches in 36 states and employing approximately 2,100 full-time registered nurses and 1,000 part-time registered nurses.

Our unique data resulted from merging two sources of data: human resources records and time-stamped activity logs. The human resources records capture the start and end dates for each worker's employment with the firm, which is critical to measuring turnover. Importantly, these records also indicate whether a worker left voluntarily or was terminated by the firm involuntarily. In addition, the data specify the position held by the worker, their contractual status (i.e., full time or part time), job discipline, branch of employment, promotions during the employee's tenure with the firm, any leaves of absence that were taken, and demographic details. We also observe the two-week pay periods over which the firm tracks worker productivity.

The provider visit logs are specific to employees who conduct home visits, as opposed to those with administrative roles that do not involve field work. The visit logs reflect all data captured via the provider's tablet during a home visit, including the start and end time for each visit, the time spent driving between visits, and other nonvisit activities such as training and documentation. Using these data, we are able to construct a complete record of the daily work schedule for each field worker, including the number of visits conducted, the time spent conducting visits, the time spent driving, and the time spent conducting other nonvisit activities. We can also observe the time at which the worker began and ended their work day, which allows us to calculate the length of the entire work day as well. By extension, we are able to calculate similar measures at the weekly level, including the number of days worked per week and the time spent conducting visits per week.

2.2 Sample definition

To answer our research question, we impose several restrictions on the data. First, we limit our sample to full-time nurses who conduct home visits. Doing so allows us to focus on the consequences of schedule volatility that are employer-driven; observed volatility in the schedule of part-time nurses may reflect the flexibility that is desirable to and requested by those workers. Second, we only consider nurses for whom we observe their entire tenure at the firm, which excludes any worker who began their employment prior to January 2016. This helps us avoid potential issues arising from censoring. Third, we exclude nurses who changed their position, job title, or contractual status while being employed by the firm, and nurses who were previously employed by the firm prior to our study period. Doing so, we are able to ensure that nurses are comparable to one another throughout their tenure with the company. Fourth, we restrict our

sample to nurses who work in a branch that has at least one other full-time nurse, as this is a factor that is necessary for our identification strategy (see section 4).

For the workers remaining in the sample, we impose two additional worker-day level exclusions. Specifically, we exclude the first 90 days of employment from the analysis in order to avoid the spurious relationship between volatility and turnover during this period. In the first 90 days of employment, schedule volatility is high by construction as the firm gradually increases each nurse's assigned workload before plateauing at the level expected of a seasoned worker (Bergman et al., 2021). At the same time, nurses are particularly likely to quit during this time, as they assess their fit with the firm and their satisfaction with the job. We further exclude periods lasting 4 weeks or longer during which the worker was not working (i.e., no visits conducted and no nonvisit activities recorded), typically due to a leave of absence.

Our final analysis sample includes 1,182 full-time nurses who collectively worked 276,554 worker-days conducting 820,860 patient visits.

2.3 Summary statistics

Table 1 provides summary statistics of the analysis sample. In our sample, the average worker is 42.15 years old (SD: 10.76) at hiring, and 92.90% are female. On an average work day, each nurse conducts 4.18 visits (1.85), which results in a total of 3.73 hours (1.61) of in-home activity. On average, 1.58 hours (1.22) are spent driving to and from patient homes. An additional 0.66 hours (1.27) are spent performing nonvisit activities. The average length of a work day is 7.14 hours (2.91), which is longer than the 5.98 hours (2.48) spent visiting patients, driving, or performing nonvisit activities; this suggests there are gaps between visits and breaks over the course of the day.

On average, a worker works 4.75 days (1.43) per week, spending a total of 17.70 hours (6.77) in patient homes and 28.38 hours (10.85) actively doing work-related activities. When accounting for the start and end times of each work day, on average each worker spends 33.88 hours (12.96) at work each week.

Of the 1,182 nurses in our sample, 540 (45.69%) separated from the firm during our analysis period. Of those, a majority (479 or 89%) separated voluntarily (i.e., quit), while the remainder were separated involuntarily (i.e., were fired). The average worker tenure at separation, of either type, was approximately 9 months.

3 Measuring Schedule Volatility

Despite the ubiquitous nature of shift-based work and just-in-time scheduling, limited prior research exists on how to measure and quantify the volatility of workers' schedules. Earlier research on the topic has typically used qualitative interviews (Henly et al., 2006) and survey instruments (Schneider and Harknett, 2019). Of note, researchers at the University of Chicago's Employment Instability, Family Well-being, and Social Policy Network (EINet) have developed

Table 1: Summary Statistics of Analysis Sample

		Mean	Std. Dev.
Worker-day level	Number of visits conducted	4.18	1.85
	Time spent conducting visits (hours/day)	3.73	1.61
	Time spent conducting nonvisit activities (hours/day)	0.66	1.27
	Time spent driving (hours/day)	1.58	1.22
	Total time (hours/day)[a]	5.98	2.48
	Average time per visit (hours/day)	0.84	0.25
	Work day start time	09:03	01:53
	Work day end time	16:09	02:19
	Length of work day (hours/day)	7.14	2.91
Worker-week level	Number of days worked	4.75	1.43
	Time spent conducting visits (hours/week)	17.70	6.77
	Total time (hours/week)	28.38	10.85
	Length of work day (hours/week)	33.88	12.96
Separations	All separations (% of workers)	45.69	49.83
	Voluntary separations (%)	40.52	49.11
	Average tenure at voluntary separation (months)	8.72	5.46
	Involuntary separations (%)	5.16	22.13
	Average tenure at involuntary separation (months)	9.68	7.42
Demographics	Age at hiring	42.15	10.83
	Female (%)	93.06	25.42
Sample size	Number of workers	1,182	
	Number of worker-days	276,554	
	Number of visits	820,860	

Notes: [a]Total time is the sum of time spent conducting visits, time spent conducting nonvisit activities, and driving.

and tested survey questions that measure "precarious work schedules", such as unpredictable and fluctuating hours, non-standard work timing, and employee control over work schedules (Lambert and Henly, 2014).

Rather than relying on interviews and survey instruments, which rely on retrospective input from the workers, we propose a method for measuring worker schedule volatility that utilizes detailed time-stamped work logs that are collected in real-time during the course of regular operations as the basis for compensation, plan of care tracking, and compliance with regulatory authorities. Given there are several aspects of a worker's schedule that are relevant when it comes to thinking about volatility, we set forth a single definition for measurement and provide volatility estimates using this definition for various aspects of a worker schedule. Specifically, we consider workers' schedule volatility with respect to the number of hours worked in a day and the length of the work day. Given the specific setting from which we obtained our data, we also examine schedule volatility vis-à-vis the number of visits conducted in a day, the number of hours spent conducting visits in a day, the number of hours spent driving in a day, and the average number of hours spent conducting each visit that day.

To measure schedule volatility, we capture the variability of the specific scheduling aspect (e.g., the length of the work day) over a fixed period of time. We use the coefficient of variation (CV) to measure schedule volatility. The CV is the standard deviation divided by the mean value of the distribution, and therefore is independent of the unit in which the measurement has been taken. We use CV rather than the standard deviation, which can only be understood in the context of the mean of the data. In our main specification, the CV is calculated over a 28-day moving window.

To formalize this discussion, consider worker i who experiences a scheduling aspect m over a lookback window ranging from time $t-\bar{t}$ to time $t-1$, $M_{it} \equiv \{m_{it-1}, m_{it-2}, \ldots, m_{it-\bar{t}}\}$. We set m_{it} to zero for any day t on which the worker is not on duty. Let D_{it} be the number of days in which the worker was on duty during this period, $D_{it} \equiv \sum_{h=1}^{\bar{t}} 1[m_{ih-1} > 0]$. Then, for each scheduling aspect m, schedule volatility is given by the following:

$$\mathcal{V}(M_{it}) \equiv \frac{\sqrt{\frac{1}{D_{it}-1}\sum_{h=1}^{\bar{t}}\left(m_{it-h} - \frac{1}{D_{it}}\sum_{l=1}^{\bar{t}} m_{it-l}\right)}}{\frac{\sum_{h=1}^{\bar{t}} m_{it-h}}{D_{it}}} \tag{1}$$

Note that we only consider days in which the worker was on duty when calculating Equation 1. This is to ensure that our measure of schedule volatility is not inflated by the values of 0 arising on off-duty days. In addition, we use $W = 28$ throughout our analyses, which corresponds to two pay periods (each pay period is 14 days). In other words, we consider volatility as it applies to the previous 28 days in the worker's employment. We use a time period equal to two pay periods because it captures a relatively short time window over which volatility in one's schedule could be expected to be associated with one's subsequent decision to quit (i.e., we would expect the link to become less strong as the window becomes much longer), while still being long enough

Table 2: Distributions of Measures of Schedule Volatility

	(1)	(2)	(3)
Measure	μ	σ	σ/μ
Daily visit count	4.30	1.46	0.35
	(0.90)	(0.40)	(0.09)
Daily visit time (hours/day)	3.83	1.23	0.33
	(0.78)	(0.62)	(0.12)
Daily drive time (hours/day)	1.67	0.75	0.46
	(0.58)	(0.76)	(0.18)
Daily total time (hours/day)	5.97	1.98	0.34
	(1.20)	(1.05)	(0.14)
Time per visit (hours/day)	0.84	0.20	0.24
	(0.14)	(0.07)	(0.07)
Day length (hours/day)	7.12	2.41	0.35
	(1.30)	(1.13)	(0.13)

Notes: Each observation is at the worker-day level. Each measure of schedule volatility is constructed using a 28-day lookback window. $N = 276,549$ observations, with the exception of daily drive time ($N = 274,749$). Standard deviations in parentheses.

that we can meaningfully capture day-to-day variability. In our setting, it is important to set W to be a multiple of the two-week pay period because there are biweekly productivity targets that affect workers' schedules. Defining W in multiples of two weeks ensures that our estimations are not biased by the potential within-pay-period smoothing. As a robustness check, we examine the sensitivity of our results to the choice of the length of the lookback window in section 5.2.2.

To illustrate how we operationalize each of the measures of schedule volatility, in Table 2 we show the mean (μ), standard deviation (σ), and CV (σ/μ) over the previous 28 days for each worker-day level observation. To provide insight into how each of these measures is distributed, we report the standard deviation of each of these measures in parentheses. In column (1), we see that the average for each worker-day's 28-day lookback window is nearly identical to the sample average reported in Table 1. This is consistent with what we would expect since the average of the 28-day lookback window preceding each worker-day is equal to the sample average multiplied by 28.[3] In contrast, the standard deviation for each worker-day's 28-day lookback window (column (2)) is lower than what is observed in Table 1, as the variability observed over a 28-day lookback window for a given worker is lower than overall variability observed in the sample. This is because the latter is calculated not only over a longer period of time for each worker but also across different workers. The CV (σ/μ) for each worker-day's 28-day lookback window, reported in column (3), shows how the previous two measures come together. In addition, it shows that each of the measures of schedule volatility captures different aspects of variability inherent in workers' schedules.

[3]The small discrepancies between the means in Table 1 and those in Table 2 arise from (a) excluding values of 0 (e.g. 0-visit days are excluded from the calculation, even if nonvisit activities occured on that day); and (b) using out-of-sample worker-days to calculate the beginning-of-panel volatility measures. For example, we calculate schedule volatility for workers with 90 days of tenure using the prior 28 work days, but those 28 work days are excluded from the analysis sample since the worker has fewer than 90 days of tenure for each of these work days.

Schedule volatility varies considerably not only across workers, but also within each worker's period of employment. Figure 1 plots, for each measure of schedule volatility, the distribution of the average level for each worker during her career (i.e., between-worker schedule volatility, in blue) and the average distribution for each worker (i.e., within-worker schedule volatility, in red). Overall, volatility *across* workers follows a similar distribution to volatility *within* workers. This suggests that, on average, workers are exposed to different levels of schedule volatility, but each worker will experience a wide range of schedule volatility over the course of her employment. Comparing across the six measures, we find similarities in both the between- and within-worker volatility measure distributions of daily visit count, daily visit time, daily total time, and day length. Conversely, we find higher overall volatility in daily drive time. Finally, we document significantly lower volatility in time per visit, suggesting that time spent per visit is relatively constant both between workers and within a worker's career.

4 Empirical Strategy

The granular nature of our data enable us to calculate schedule volatility, using a rolling look-back window, for each day in a worker's employment period and to directly observe voluntary separations at the worker-day level. As such, a direct approach to studying the effect of schedule volatility on voluntary separation would be to estimate, using the following linear probability model, the conditional expectation of quitting on a given day (a binary variable) as a function of schedule volatility calculated over the previous 28 days:

$$1(Quit_{it}) = \beta + \theta \mathcal{V}(M_{it}) + X_{it}\delta + \epsilon_{it} \tag{2}$$

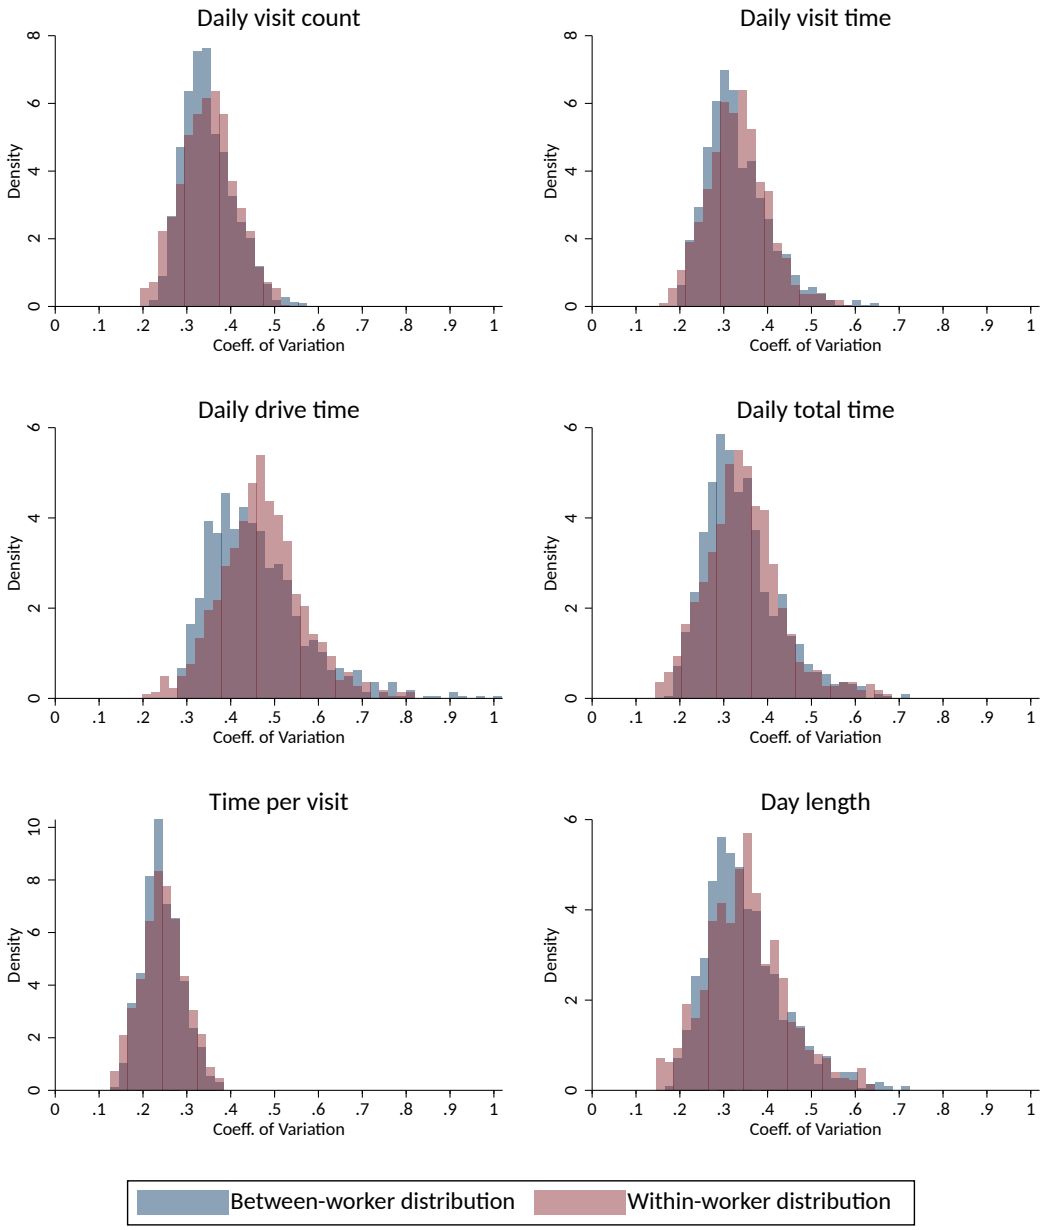
where X_{it} is a set of additional controls. However, there are two main empirical challenges to identifying and estimating this effect: identifying the timing of a worker's decision to quit and addressing the potential endogeneity of schedule volatility.

4.1 Timing of the decision to quit

First, it is not possible to observe how far in advance a worker has made up her mind to quit prior to actually resigning. A worker could decide to quit, but continue to stay on in her position while searching for a new job, only separating from the company once a new job has been secured. Alternatively, she could decide to quit and immediately give the firms' standard two-weeks' notice to the company. In some extreme cases, the worker could decide to quit without giving any notice to the company, thus leaving her position immediately. The agency from which we obtained data states that this last scenario rarely occurs and that the previous scenarios are much more common.

We address this particular challenge by allowing for a leading positive outcome, where the length of the lead is two weeks. Specifically, the dependent variable (Y_{it}) takes on a value of 1 on

Figure 1: Between- and Within-Worker Distributions of Measures of Schedule Volatility



the day the worker quits and in the 14 days leading up to the separation.[4]

Employing a leading positive outcome is important, as schedule volatility directly affects the decision to quit, and only subsequently the act of quitting. Therefore, our approach allows us to proxy for when the worker made the decision to quit as opposed to merely capturing the last day of the worker's employment. We examine how sensitive our results are to this choice of the length of the lead in section 5.2.1.

4.2 Endogeneity of schedule volatility

Second, schedule volatility leading up to separation is likely endogenous, as other worker-level or branch-level changes may affect both the worker's schedule volatility and her inclination to quit. For example,

- After a worker has given notice to quit, the firm may decide to re-assign cases to other workers, leading to a reduction in workload for the exiting worker, which would affect her schedule volatility.
- A worker might attempt to reduce or otherwise change her workload prior to quitting, which might affect her schedule volatility.[5]
- An interpersonal conflict between a branch manager and a worker may lead the branch manager to poorly allocate work assignments to the worker, which would increase schedule volatility. Simultaneously, the conflict may drive the worker to decide to quit. Other unobservable channels operating over time and within branches may affect both schedule volatility and voluntary turnover.

Of course, we cannot address each potential scenario that results in endogeneity. For this reason, we employ an instrumental variable approach in estimating the causal relationship between schedule volatility and voluntary turnover.

We instrument for a worker's schedule volatility by using Paid Days Off (PDOs) taken by *other* workers at the same branch as the focal worker. Because each home health visit is conducted by an individual nurse, the absence of one nurse on a given day has a direct and measurable impact on the schedules of other nurses working at the same branch on that day, as they must now perform the visits which would have otherwise been assigned to the nurse taking the PDO.[6]

[4]Namely, define k_i to be the date on which worker i quits, for workers who quit during the study period. Then, for these workers, $Y_{it} = 1$ when $t \geq k_i - 14$. For workers who do not quit during the study period, $Y_{it} = 0$ for any t.

[5]In our setting, workers have little control over their schedules, and are expected to conduct a pre-determined number of visits in each 2-week pay period, but a worker who has decided to quit may not feel bound by such expectations.

[6]This is in contrast to other work environments where workers operate in teams, where the absence of one worker on a given day does not necessarily affect the schedule of the other workers; rather, it may increase the intensity of work carried out during the work day or result in work being delayed until a future date. Consider, for example, an absent floor worker in a department store. Tasks that would have been performed by this worker can be performed by the other workers on duty, which may increase the intensity of work but does not result in longer shifts for the on-duty workers.

In other words, PDOs taken by other workers at the same branch are expected to increase the number of visits that the focal worker conducts that day, which in turn may affect other schedule-related aspects, such as the time spent driving to and from visits and the total time spent "on duty" that day.

We assess the relevance of this instrument by estimating the following regression for each measure of schedule volatility, m_{it}:

$$\ln(m_{it}) = \alpha + \gamma \sum_{-i \in B_{it}} PDO_{-it} + Branch_{it} + Month \times Year_t + \epsilon_{it} \tag{3}$$

where $\sum_{-i \in B_{it}} PDO_{-it}$ is the number of PDOs taken by other full-time workers at the focal worker's branch on day t, $Branch_{it}$ are branch fixed effects, and $Month \times Year_t$ are month-year fixed effects. Here, γ captures the effect of other workers' PDOs on the focal worker's schedule volatility. To aid with interpretability and comparability, we standardize each measure of schedule volatility by subtracting the mean from the observed value and dividing by the standard deviation.

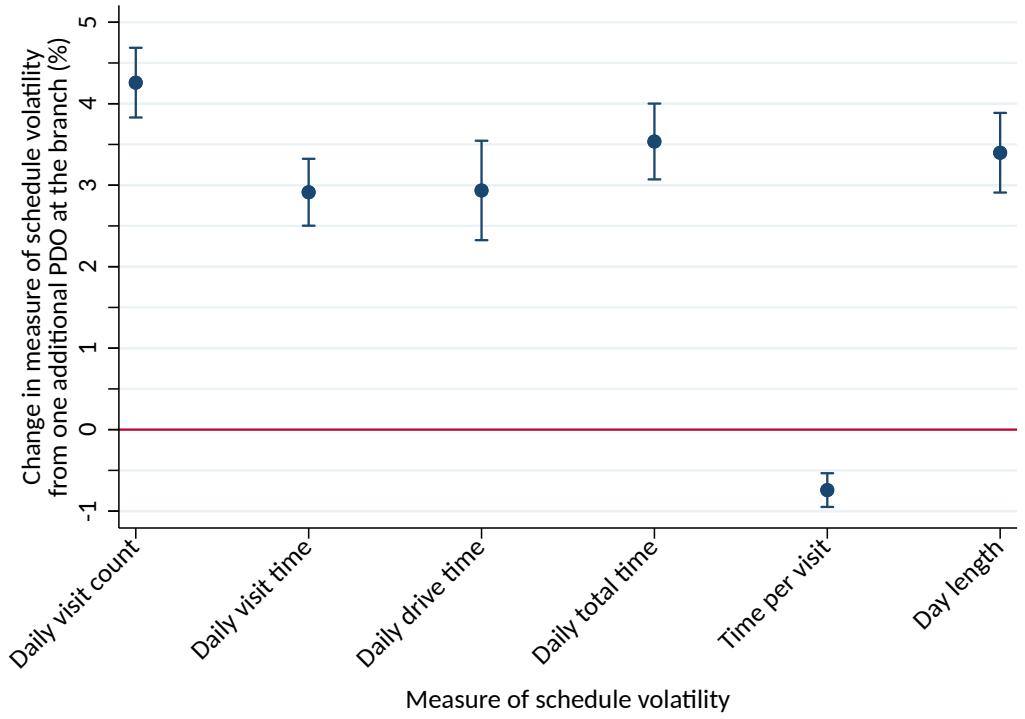
Figure 2 plots the point estimate of γ and its corresponding 95% confidence interval (CI) for each measure of schedule volatility.[7] For all measures excluding time per visit, we find that the focal worker's schedule volatility increases when there are more PDOs taken by other workers at the same branch; this is in line with what we would expect since the remaining workers must carry out the visits that would have otherwise been conducted by the worker taking the PDO. Specifically, an additional PDO taken by others at the branch increases the focal worker's daily visit count by an average of 4.26% (SE: 0.22 percentage points (p.p.)), her daily visit time by 2.9% (0.21 p.p.), her daily drive time by 2.94% (0.31 p.p.), her total time spent on duty that day by 3.54% (0.24 p.p.), and the length of her work day by 3.41% (0.25 p.p.). With additional visits to perform, workers spend less time per visit, with an average decrease of 0.86% (0.11 p.p.) in time spent per visit.

In addition to assessing its relevance, we also examine whether the proposed instrument would satisfy the exclusion restriction. This assumption would be violated if PDOs taken by others at the branch were correlated with the focal worker's decision to quit through some unobservable variable. With the inclusion of branch fixed effects, such an unobservable variable would have to vary within a branch, excluding any constant branch-level characteristic that might affect both PDOs and turnover. With month-year fixed effects, the unobservable variable would also have to vary within a month-year, excluding any effects relating to either agency-wide temporal shocks, like a new corporate policy that affects both PDOs and turnover, as well as any seasonality effect. Given our inclusion of both branch fixed effects and month-year fixed effects, we expect the exclusion restriction to be satisfied.

For our estimation, we employ a set of instruments that relate to the PDOs taken by other workers at the same branch. Each of these variables is constructed based on PDOs taken in the

[7] Complete estimation results are presented in Table A1 in the Appendix.

Figure 2: Effect of Other-Worker Paid Days Off in the Branch on Focal Worker's Schedule



Notes: This figure shows the percent change in each measure of schedule volatility given an additional Paid Day Off (PDO) taken by another worker at the same branch on the same day.

28 days prior to the focal day, which is the same lookback window we use for defining schedule volatility (see section 3). The variables are designed to succinctly capture key aspects of the PDO distribution over this time period. Specifically, the instruments are: (a) an indicator variable that receives value when any PDOs were taken by other workers at the same branch during the lookback window; (b) the average number of PDOs taken per day by other workers at the same branch during the lookback window; and (c) the standard deviation of PDOs taken per day by other workers at the same branch during the lookback window. The first two variables capture the *level* of PDOs taken by other workers at the branch over the lookback window, while the third variable captures the *volatility* in PDO activity. For the latter two variables, we include both linear and squared terms, respectively, bringing our set of instruments to five. For all five PDO measures, we standardize the instruments prior to estimation.

4.3 Two-stage estimation approach

Our main empirical specification is therefore a two-stage least squares (2SLS) model:

$$\mathcal{V}(M_{it}) = \alpha + Z_{it}\beta + X_{it}\gamma + Branch_{it} + Month \times Year_t + u_{it} \tag{4}$$

$$Y_{it} = \eta + \theta\mathcal{V}(M_{it}) + X_{it}\kappa + Branch_{it} + Month \times Year_t + \nu_{it} \tag{5}$$

where we instrument for schedule volatility, $\mathcal{V}(M_{it})$, using Z_{it}, the vector of instruments described in section 4.2. For ease of interpretation and as described in section 4.2, we standardize the measure of schedule volatility $\mathcal{V}(M_{it})$ prior to estimation. X_{it} are a set of exogenous controls described below; $Branch_{it}$ and $Month \times Year_t$ are branch and month-year fixed effects, respectively; and u_{it} and v_{it} are the remaining unobservables.

The causal effect of schedule volatility on voluntary separation is captured by θ. We estimate this model separately for each of the measures of schedule volatility described in Table 2. In our estimations, we calculate heteroskedasticity- and autocorrelation-robust standard errors. Moreover, our results are robust to a number of additional specifications (see section 5.2.)

We include several controls in the model (X_{it}) to help isolate the effect of interest. First, we control for the number of full-time workers and number of part-time workers on roster at the same branch as the focal worker, since the presence of additional workers at the branch would determine the extent to which a single worker's absence (via a PDO) may impact the focal worker's workload. We include both the linear and squared terms of each of these measures in order to account for potential non-linearities in these effects. Second, we account for the worker's tenure with the firm, given nurses' likelihood to quit varies significantly based on her tenure with the firm (Bergman et al., 2021). Third, we include the number of days for which the worker was on duty in the previous 28 days, both as a linear term and as a logged term, to control for any mechanical correlation between the number of days worked and volatility. We further include branch fixed effects to capture time invariant branch characteristics that might affect norms around taking PDOs, schedule volatility, and turnover. Finally, we include month-year fixed effects, which allows us to account for both seasonality and any firm-wide shocks to schedule volatility and turnover.

5 Results

5.1 Effects of schedule volatility on turnover among full-time workers

Table 3 summarizes the estimation results for the 2SLS regressions described in Equations 4-5. Each column represents a separate 2SLS regression, where the column heading indicates which measure of schedule volatility was used as the treatment variable. For each 2SLS regression, we report the first-stage coefficients for each of the instruments (β), the second-stage coefficient of the measure of schedule volatility (θ), a weak instrumentation Cragg-Donald F-test statistic, and the number of observations used in the estimation.

Our regression sample includes 276,549 worker-days for workers across 208 branches between 2016 and 2019. We find qualitatively similar first-stage coefficient estimates for daily visit count, daily visit time, daily drive time, and time per visit. Our estimates show that schedule volatility has a (marginally decreasing) positive correlation with "volatility" in PDOs taken by other workers at the branch, as measured by the standard deviation in the number of daily PDOs taken by other workers over the lookback window. By contrast, an increase in the average number of daily

PDOs taken by the other workers at the same branch, holding the standard deviation of those daily PDOs constant, is negatively correlated with schedule volatility. For two other measures of schedule volatility, daily total time and day length, we find similar qualitative patterns—diminishing positive correlation with the SD of others' PDOs and negative correlation with the average of others' PDOs—but with a smaller magnitude. A weak instrumentation test, carried out using the Cragg-Donald F-test, suggests our instruments are sufficiently strong regardless of the scheduling aspect used to measure volatility, ranging from 21.09 (when using daily total time) to 61.51 (when using daily visit time).

Table 3: 2SLS Estimation Results: Effect of Schedule Volatility on Voluntary Separation

	(1) Daily visit count	(2) Daily visit time	(3) Daily drive time	(4) Daily total time	(5) Time per visit	(6) Day length
First stage						
1(SD of others' PDOs > 0)	-0.0265*	-0.0329**	0.0214	-0.0153	0.0106	-0.00612
	(0.0114)	(0.0115)	(0.0117)	(0.0118)	(0.0119)	(0.0116)
SD of others' PDOs	0.0888**	0.0679**	0.0534**	0.0404**	0.0398**	0.0377**
	(0.0105)	(0.0107)	(0.0108)	(0.0114)	(0.0107)	(0.0110)
SD of others' PDOs, squared	-0.0534**	-0.0559**	-0.0538**	-0.0256*	-0.0615**	-0.0310**
	(0.00950)	(0.00989)	(0.0104)	(0.0106)	(0.00983)	(0.0102)
Average of others' PDOs	-0.0127**	-0.0125**	-0.00261	-0.00709*	-0.00884**	-0.00737*
	(0.00262)	(0.00288)	(0.00256)	(0.00305)	(0.00271)	(0.00291)
Average of others' PDOs, squared	0.000545	-0.00145	-0.00246	-0.00245	0.00612*	-0.00196
	(0.00278)	(0.00307)	(0.00282)	(0.00315)	(0.00297)	(0.00306)
Second stage						
Volatility measure	0.0473**	0.0362**	0.0594**	0.0700**	0.0386**	0.0632**
	(0.0103)	(0.00850)	(0.0117)	(0.0172)	(0.0124)	(0.0145)
IV F-stat	55.13	61.51	42.75	21.09	35.98	26.31
Observations	276,549	276,549	274,749	276,549	276,549	276,549
Dep. var. mean	0.0130	0.0130	0.0130	0.0130	0.0130	0.0130

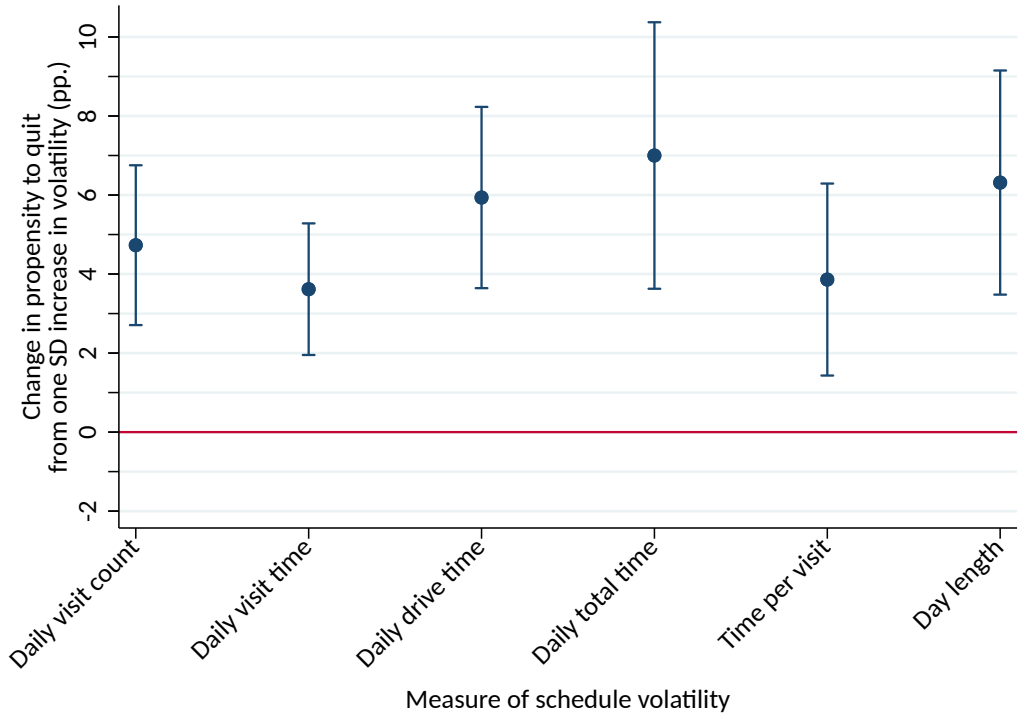
Notes: Key coefficient estimates from the LPM 2SLS model described in Equations 4-5. Newey-West (heteroskedasticity and autocorrelation robust) standard errors are in parentheses. Standard errors are calculated using 2 lags. ** $p<0.01$, * $p<0.05$.

Figure 3 visualizes our results by plotting the second-stage coefficients appearing in Table 3. Overall, we find that schedule volatility has a positive and statistically significant impact on workers' propensity to quit. Focusing on schedule volatility as measured using the volatility in daily visit count (column (1)), our results show that a one standard deviation increase in volatility would increase a worker's propensity to quit by 4.73 p.p. (SE: 0.103 p.p.); this corresponds to a 364% increase in the average worker's propensity to quit in a day.[8]

We find similar results when using several other measures of schedule volatility. Specifically, the increase in propensity to quit associated with a one standard deviation increase in volatility is 3.62 p.p. (SE: 0.85 p.p.) when using daily visit time (column (2)), 5.94 p.p. (1.17 p.p.) when using daily drive time (column (3)), 7.00 p.p. (0.17 p.p.) when using daily total time (column (4)), 3.86 p.p. (1.24 p.p.) when using time per visit (column (5)), and 6.32 p.p. (1.45 p.p.) when using day length (column (6)).

[8]The coefficients are described in units of the dependent variable, as defined in 4.1. The average propensity to quit on a given day is approximately 1/14th that value, at 0.0009148 (SD: 0.0302324). Adjusting to the "0-day lead" level (dividing by the dependent variable mean and multiplying by the average propensity to quit on a given day), we can report a coefficient value of 0.333 p.p. (SE: 0.002 p.p.).

Figure 3: 2SLS Coefficients Summary: Effect of Schedule Volatility on Voluntary Separation



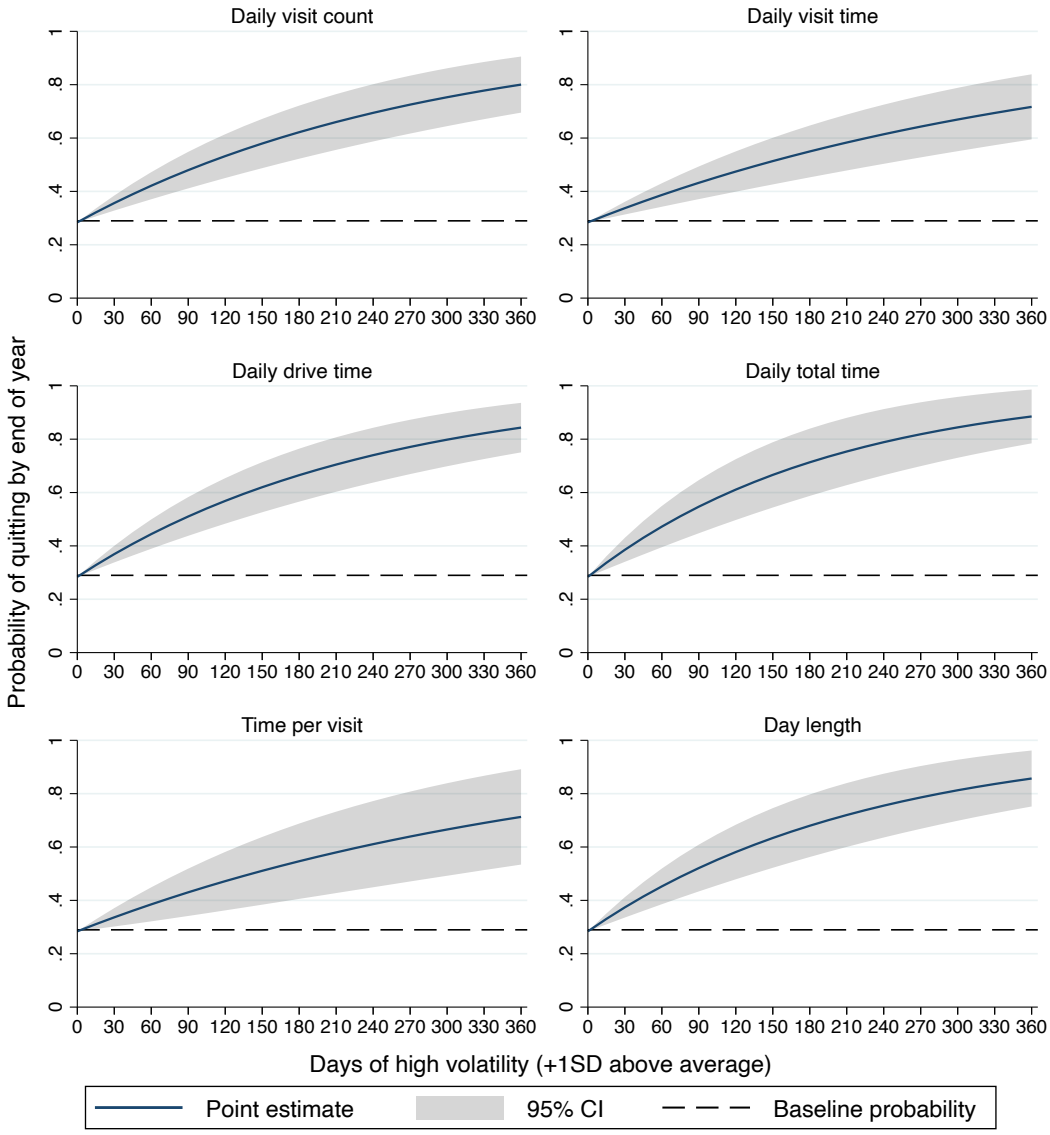
Notes: This figure plots the percentage point increase in a worker's propensity to quit given a one standard deviation (SD) increase in volatility for each of the six measures of schedule volatility, using the 2SLS linear probability model described in Equations 4-5. The error bars correspond to the 95% confidence intervals calculated using Newey-West standard errors.

To contextualize these results, we calculate the impact of schedule volatility on the probability that a worker will quit within a full year. Figure 4 plots, for each measure of schedule volatility in our analysis, the estimated probability of a worker quitting by the end of the year as a function of the number of days of high volatility. Since workers in our sample have an average daily propensity to quit of 0.0914%, their probability of quitting over the course of a full year is 28.4% on average.[9] Focusing first on daily visit count as our measure of schedule volatility, we find that introducing a single 30-day period of high volatility (where high volatility is defined as one standard deviation higher than the sample average) would increase a worker's probability of quitting in that year to 35.62% (SE: 1.44 p.p.), which corresponds to a 25.45% (5.08 p.p.) increase in likelihood over the baseline. In a similar exercise, a worker experiencing 180 days of high volatility would be 119% (17.89 p.p.) more likely to quit within a year than the average worker, and a worker experiencing a full year of high volatility would be 182% (18.89 p.p.) more likely to quit within a year than the average worker.

Similar patterns emerge when using the other measures of schedule volatility. A single 30-

[9]Since $\Pr[\text{Quit that year}] = 1 - (1 - \Pr[\text{Quit that day}])^{365}$. For this exercise, we standardize the coefficient estimates by a factor of $\mu_{1[Quit]} / \mu_Y$, thereby expressing the coefficient in terms of daily probability of quitting rather than in units of $Y_i t$, which includes termination leads.

Figure 4: Probability of Voluntarily Separation by End of Year as a Function of Schedule Volatility



Notes: This figure presents, for each of the six measures of schedule volatility, a worker's probability of voluntary separation over the course of a year as a function of the number of days during which the worker experienced high levels of volatility. A high level of volatility is defined as one standard deviation (SD) above the average level of schedule volatility. The shaded region represents the 95% confidence bands calculated using the delta method.

day period of high volatility would increase a worker's probability of quitting within a year by anywhere from 18.49% (SE: 6.19 p.p.), when measuring volatility using daily time per visit, to 35.67% (8.09 p.p.), when using daily total time. A worker experiencing a full year of high volatility would have a 71.30% chance of quitting by the end of the year when measuring schedule volatility using time per visit, a 71.69% chance when measured using daily visit time, a 80.04% chance using daily visit count, a 84.34% chance using daily drive time, a 85.69% chance using

day length, up to an 88.53% chance when measuring schedule volatility using daily total time.

5.2 Robustness tests and sensitivity checks

By using several measures of schedule volatility, our initial analyses already demonstrate that our results are robust to the specific scheduling aspect used to measure the level of volatility in a worker's schedule. That said, our choices of the length of the lead prior to quitting (14 days) and the duration of the lookback window (28 days) are, to some degree, arbitrary, and could influence our results. In this section, we present how our coefficients of interest change when we perturb each of these two parameters.

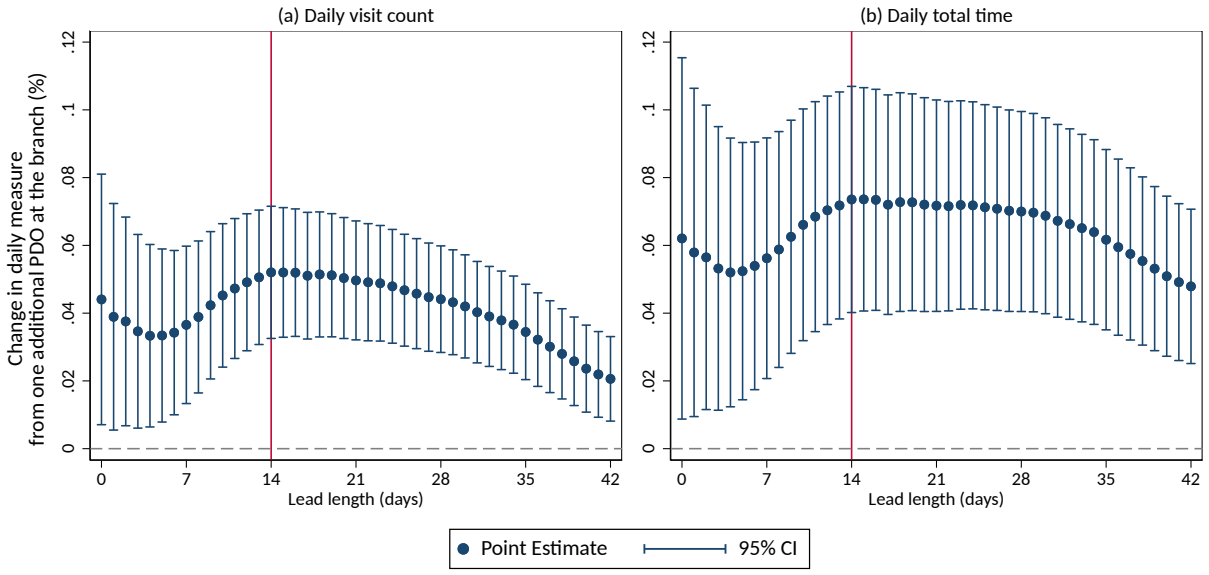
5.2.1 Length of the lead prior to quitting

In section 4.1, we describe the challenge of identifying the specific day on which a worker makes the *decision to quit*, which typically occurs sometime prior to the *actual date of quitting* (k_i). We address this issue by allowing the dependent variable (Y_{it}) to take on a value of 1 on the day of separation and in the L days leading up to the separation. We must choose a length of the lead, L, because the true date on which the worker makes the decision to quit is unobservable. In our main analyses, we choose to set $L = 14$ because employees at this firm are expected to give two weeks' notice of their separation prior to quitting.

Misspecifying L is expected to bias our results, although the directionality of this bias is unclear. If L is too short, we would miss capturing the true date on which the worker decided to quit. If L is too long, the estimated impact of schedule volatility on a worker's propensity to quit would be diluted. To understand how the choice of the lead length may impact our results, without imposing additional priors, we repeat our estimation of Equation 5 using lead lengths ranging from 0 days (i.e., $Y_{it} = 1$ only on the day of worker separation) to 42 days prior to separation. In Figure 5, we show the results of doing this for two measures of schedule volatility, daily visit count (panel a) and daily total time (panel b). Extending the lead length increases the mean of the dependent variable by construction, which in turn has a mechanical scaling effect on the coefficient of interest. We remove this mechanical bias by normalizing the coefficient—i.e., dividing its value by the average value of the dependent variable (for the given lead length), and then multiplying it by the average value of the dependent variable for a 14-day lag, our baseline lead length. We find similar patterns of results across the two measures of schedule volatility. The impact of schedule volatility on the probability of separation decreases slightly as we extend the length of the lead from 0 to 7 days, and then increases until reaching its peak when the lead length is set at 14 days, after which the effect gradually decreases.

The patterns in Figure 5 suggest that there are two key points at which workers may be deciding to quit. The main one seems to occur at 14 days prior to separation, which coincides with the firm's statement that most employees give two weeks' notice prior to voluntary separation. Importantly, this supports our use of a 14-day lead length in our main analyses. Another decision

Figure 5: Effect Size and Length of the Lead Prior to Quitting



Notes: This figure plots the results of repeating the estimation of our main results using lead lengths ranging from 0 days to 42 days prior to separation, scaling the coefficient by dividing its value by the average value of the dependent variable (for the given lead length), and multiplying it by the average value of the dependent variable for a 14-day lag. Panel (a) uses daily visit count as the measure of schedule volatility, and panel (b) uses daily total time as the measure of schedule volatility. The vertical red line at 14 days shows the length of the lead used in our main analyses. The error bars correspond to the 95% confidence intervals calculated using the delta method.

point seems to occur on the day of separation, i.e., some workers seem to be quitting without giving prior notice to the employer. That said, the large 95% confidence intervals when using very short lead lengths (e.g., $L < 5$) suggest that there is a substantial fraction of workers who most likely made their decision to quit two or more weeks earlier.

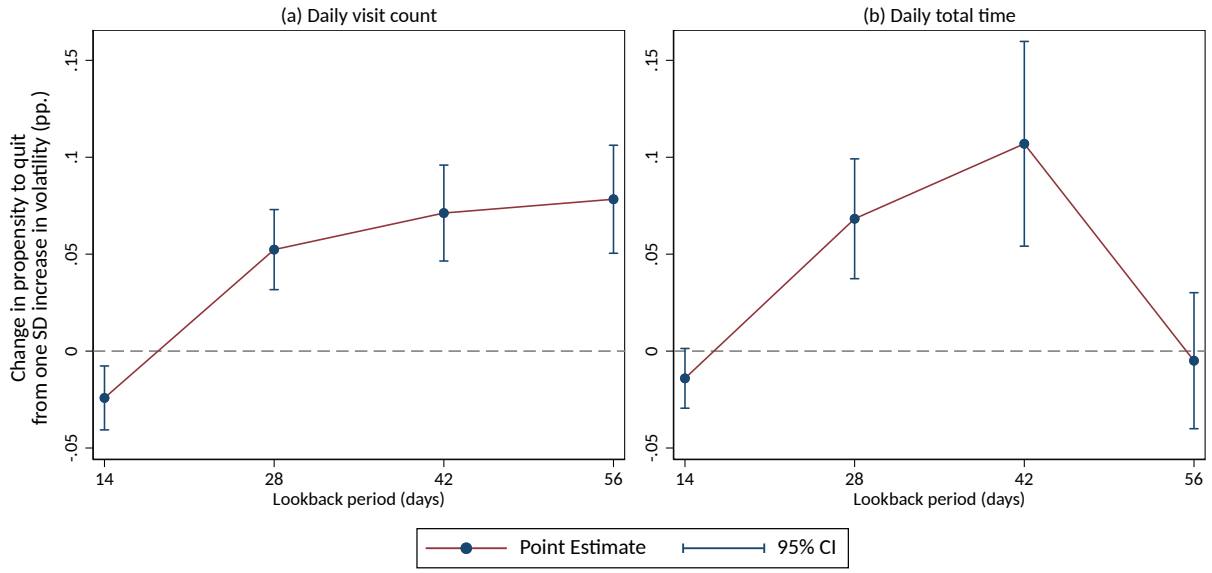
As an alternate specification, we present the estimation results when using a 28-day lead length in Appendix Table A3. We find our main results to be robust to this alternate specification of the length of the lead.

5.2.2 Length of the lookback window in measuring schedule volatility

Determining the length of the lookback window involves a tradeoff between precision, salience, and relevance. When the lookback window is longer, the measure of schedule volatility is constructed using more observations; thus, it has higher precision and is less impacted by extreme workload days. However, longer lookback windows also smooth out volatility over time, which makes it difficult to pinpoint triggers for a worker's decision to quit. On the other end of the extreme, when the lookback window is very short, there is a risk of capturing observations that are not relevant to the worker's decision to quit, as they may exclude the earlier time frame during which the volatility in one's schedule led to the decision to quit. Therefore, it is important to strike a balance between these factors and determine an optimal length of the lookback window

that is neither too short nor too long.

Figure 6: Effect Size and Length of the Lookback Window



Notes: This figure plots the results of repeating the estimation of our main results using lookback windows ranging from 14 days (one pay period) to 56 days (four pay periods). Panel (a) uses daily visit count as the measure of schedule volatility, and panel (b) uses daily total time as the measure of schedule volatility. The error bars correspond to the 95% confidence intervals calculated using Newey-West standard errors.

In our main analyses, we use a lookback window of 28 days, or 4 weeks. As discussed in section 3, this is equivalent to two pay periods at the firm. To examine how this choice of $W = 28$ may influence our results, we repeat our analyses using additional values of W, each a multiple of the two-week pay period. Figure 6 shows the results when using 14, 28, 42, and 56-day lookback windows, respectively, for two measures of schedule volatility, daily visit count (panel a) and daily total time (panel b). There are two main takeaways from these results. First, setting the lookback window to a single pay period ($W = 14$) seems to miss the action that is driving the decision to quit. While we do not observe a significant reduction in precision, the relevance seems to be lost, suggesting a lookback window of a single pay period is too short. Second, our results are robust to using lookback windows of two or three pay periods ($28 \leq W \leq 42$), but begin to lose precision and salience as the lookback window grows larger. Given the optimal W ought to balance precision, salience, and relevance, these findings support our use of a 28-day lookback window in our main analyses, which is the shortest possible lookback window (for greater salience) that does not suffer from a loss of relevance.

5.2.3 Functional form of controls and instruments

In our main specification, we sought to select the functional form of controls and instruments in a way that simultaneously alleviates concerns for omitted variable bias, limits efficiency losses,

and preserves the interpretation of coefficients. In this section, we demonstrate that our regression estimates are robust to increasing the flexibility in the functional forms of both the control variables and the instrumental variables.

To allow for the most flexibility, we estimate our results using highly flexible functional forms for each of the controls and instruments in the model specified in Equations 4-5. Namely, we model each of the instruments (SD of others' PDOs and average of others' PDOs) and each of the controls (number of full-time workers in the branch, number of part-time workers in the branch, worker tenure, and the number of work days in the past 28 days) using cubic B-splines. A summary of the estimation results, describing the estimate of the treatment effect and the instruments' F-statistics, is presented in Table 4. Our point estimates remain qualitatively unchanged from the estimates reported in 3. We note that the first-stage instruments' F-statistics are smaller, which is to be expected given the loss of efficiency induced by including additional covariates.

Table 4: Second-stage Estimation Results with Fully Flexible Instruments

	(1) Daily visit count	(2) Daily visit time	(3) Daily drive time	(4) Daily total time	(5) Time per visit	(6) Day length
Volatility measure	0.0480**	0.0369**	0.0424**	0.0717**	0.0268*	0.0634**
	(0.00972)	(0.00790)	(0.0103)	(0.0155)	(0.0107)	(0.0139)
IV F-stat	38.90	44.36	27.67	15.14	25.95	17.12
Observations	234,849	234,849	233,182	234,849	234,849	234,849
Dep. var. mean	0.0127	0.0127	0.0127	0.0127	0.0127	0.0127

Notes: Newey-West (heteroskedasticity and autocorrelation robust) standard errors are in parentheses. Standard errors are calculated using 2 lags. ** $p<0.01$, * $p<0.05$.

5.2.4 Probit regression model

In this section, we rerun our main analysis using an IV probit model, estimated in a single step using maximum likelihood estimation and utilizing the same set of controls and instruments described in Equations 4-5. We employ this model later in Section 6, where we predict voluntary separation probabilities under alternate scheduling policies.

When studying binary dependent variables, researchers often specify the regression model using a probit or logit form rather than a linear probability model (LPM). The main advantage of these models is that they tend to fit the regression function better than LPM, and the resulting predictions are constrained to lie on the unit interval. In constrast, LPM predictions often extend beyond that range. Nevertheless, when estimating marginal effects, LPM produces consistent estimates that are often very similar to the marginal effects calculated from probit models (Angrist and Pischke, 2009).

We report the estimation results when using an IV probit model in Table 5. We note that the resulting average marginal effect estimates are qualitatively similar to the coefficient estimates produced by the LPM (Table 3), albeit slightly larger. For example, when studying the volatility

in daily visit count, we estimate an average marginal effect of 6.5 p.p. using the IV probit model, and a coefficient of 4.7 pp. (SE: 1.03) using the LPM.

Table 5: IV Probit Estimation Results

Dep. var.: Y_{it} (28-day lead)	(1) Daily visit count	(2) Daily visit time	(3) Daily drive time	(4) Daily total time	(5) Time per visit	(6) Day length
Volatility measure	0.949**	0.835**	0.972**	0.997**	0.815**	0.938**
	(0.057)	(0.085)	(0.052)	(0.047)	(0.113)	(0.078)
Observations	200,179	200,179	199,509	200,179	200,179	200,179
Average marginal effect	0.065	0.049	0.074	0.097	0.046	0.065

Notes: Instrumental variable maximum likelihood probit estimates. Instrument and control variable sets are identical to the ones used in estimating Equations 4-5. Robust standard errors in parentheses. ** p<0.01, * p<0.05.

6 Policy Simulations

In this section, we consider three potential scheduling policies that are aimed at reducing employees' schedule volatility. Using a simulation approach, our goal is to illustrate the extent to which schedule volatility can be lessened by introducing simple operational rules, and how much these policies may in turn reduce employee turnover. To this end, we use our analysis sample to first simulate the effect of each policy on workers' schedule volatility, and then use our estimation results from section 5 to approximate how the simulated changes in schedule volatility will reduce workers' propensity to quit.

In settings like ours, much of the schedule volatility is derived from fluctuation in demand for services. Having to respond quickly to hospital discharges, much of the schedule volatility among full-time workers arises from employer-dictated schedules that are set at the last minute rather than employees' preference for flexibility. Therefore, the solutions we propose do not involve turning down new patients, increasing the size of the available workforce, shortening visits, or increasing wait times. Our simulated solutions are limited to the reallocation of workload among existing workers, taking the patient caseload as given. As a consequence, our policy interventions would result in a reduction—not elimination—of schedule volatility.

Reducing schedule volatility for full-time workers involves increasing the consistency of the work schedule. This reduction can be achieved in many ways, and is sensitive to the way schedule volatility is measured by employers. Depending on the aspect of scheduling being measured, reducing schedule volatility could involve standardizing the length of workers' shifts; standardizing the time at which a worker starts her shift; standardizing the number of specific tasks a worker needs to accomplish on a specific work day; or fixing the days on which a worker is on duty. Such interventions are expected to incur some costs from an operational standpoint, but could prove to yield cost savings if they sufficiently improve employee satisfaction and reduce turnover.

We focus on policies that affect schedule volatility as measured by daily visit count. In other words, we study policies that directly manipulate the number of visits a full-time worker

conducts on a given work day. There are several reasons for choosing daily visit count as the scheduling aspect to focus on in our simulations. First, visits form the "building blocks" of an employee's work day in the home health setting, where visit time, drive time, total time, and day length are all affected by the number of visits the worker conducts in a day. Second, policies where we reassign visits from one worker to another are easier to conceptualize than those that regulate the length of a visit, drive time between visits, or the length of a work day. Relatedly, these aspects are easier for firms to operationalize as well. Finally, focusing on the number of visits conducted in a day allows us to ensure that the service level being provided (i.e., number of customers being served) is being held constant.

6.1 Simulated policies to reduce schedule volatility

We simulate three potential scheduling policies, which we describe in detail below. Each policy reconsiders the number of visits a worker is to conduct on a given day using a different mechanism. We enforce two operational constraints on each one of the scenarios we consider. First, we hold constant the total number of visits performed at the branch level. In other words, all patient visits that we observe in the data must still be carried out. However, we allow a visit to be reassigned to a different worker (either a full-time worker or part-time worker, depending on the policy), and to be rescheduled to a different day within the same calendar week. Similarly, the intervention cannot add additional, nonexistent visits to workers' schedules. Second, we do not change the days on which a worker is on duty, allowing workers to preserve the flexibility of choosing the days on which they wish to work. The policies we consider are:

1. **Cross-sectional smoothing through redistribution.** Visits performed by full-time workers at the branch are evenly redistributed across all full-time workers on duty that day. For example, suppose workers A and B are the only full-time workers on duty at the branch on a given day, with A performing 5 visits and B performing 3 visits. In this case, one visit would be reassigned from worker A to worker B, such that both workers are scheduled to conduct 4 visits each. When visits cannot be redistributed evenly (without assigning fractional visits to workers), we assign an additional visit to the worker who originally had the highest number of visits, or alternatively take a visit away from the worker who originally had the lowest number of visits. If a visit needs to be reassigned to one of many workers with the same number of original visits, we randomly pick a single worker to reassign a visit to/from.

2. **Inter-temporal smoothing through redistribution.** Rather than redistributing visits across all full-time workers on duty on a given day, in this policy we evenly redistribute each full-time worker's visits across her work days within a calendar week. We exclude days on which a worker is not originally on duty (e.g., weekend days) from the set of days to which visits may be redistributed. When visits cannot be evenly distributed without assigning fractional visits, we assign (remove) a visit to the work day with the highest

(lowest) number of visits. For example, if a worker conducts 8 visits on Monday, 3 visits on Tuesday, and no visits on any other work day, we would reassign 2 visits from Monday to Tuesday, leaving her with a new schedule of 6 visits on Monday and 5 visits on Tuesday.

3. **Cross-sectional smoothing through offloading.** Similar to the first policy where we smooth cross-sectionally through redistribution, this third policy also balances the number of visits conducted across workers at the branch but goes one step further by transferring "excess" visits from full-time workers to part-time workers when part-time workers are available. We define excess visits as the number of visits performed above a branch-quarter specific threshold, which we set at the median number of visits performed per day by full-time workers in a given quarter at a given branch. The reassignment of visits from full-time workers to part-time workers is restricted by the availability of part-time workers; a visit can be offloaded to a part-time worker if: (a) there is a part-time worker on roster at the branch on that day, and (b) the part-time worker is not already performing at maximum capacity that day (defined as the greater between 8 visits and the maximum observed for that specific part-time worker). Visits are offloaded to part-time workers sequentially, starting with the full-time worker who has the highest number of excess visits on a given day.[10]

We note that the operational constraints we have placed on these policies are minimal, and that firms may be subject to additional constraints. For example, workers may be unable to conduct all assigned visits on a given day, or certain visits may need to be conducted on a specific day (i.e., cannot be reassigned to another day of the week). In addition, when simulating these policies, we enjoy the foresight of knowing exactly how many visits would be performed each day, ignoring the fact that visits may get cancelled after they are assigned or that unexpected visits may shift schedules. Similarly, workers might call in sick at the last minute, hazardous weather conditions could affect branch operations, and a myriad other events could certainly introduce additional unforeseen volatility to workers' schedules.

Our simulations also ignore potential externalities generated by our policies. First, some of the existing schedule volatility may generated by the worker's preferences, despite them being full-time workers. While our interventions do not force workers to be on duty when they were not otherwise scheduled to be, they may increase the number of visits a worker is expected to conduct on the days when they are on duty. This may result in workers rushing to complete visits, which in turn may have negative quality implications for patients (Song et al., 2021). Second, while we constrain the simulations to maintain the same level of service (visits) as observed in the data, reassigning visits to different workers might negatively affect the quality of the work performed. For example, in the home health setting, redistributing visits to other workers may

[10]This exercise could be expanded to deal with volatility generated by *too few* visits on a given day, by transferring visits from part-time workers to full-time workers. We refrain from applying this "onloading" policy, as it is expected to result in a similar outcome as the first policy (cross-sectional smoothing through redistribution), whereby visits are offloaded from overworked full-time workers to part-time workers, and subsequently from the part-time worker to a different, under-worked, full-time worker.

result in patients experiencing more frequent handoffs and, in turn, care discontinuity over the course of a home health episode. This may negatively impact the quality of care received by patients (David and Kim, 2018).

6.2 Simulation results

To simulate each of the policies, we use data from our analysis sample for the first two simulated policies and extend this sample to include part-time workers for the third simulated policy. In Appendix Table A4, we report summary statistics of 115,194 visits conducted over 74,079 worker-days by 428 part-time workers and the extended analysis sample that includes both full-time and part-time workers. Using the observed data, we carry out the redistribution or offloading of visits according to the decision rules stated in section 6.1, re-calculate the schedule volatility that would have been experienced by full-time workers under each policy, and estimate the resulting change in these workers' propensity to quit.

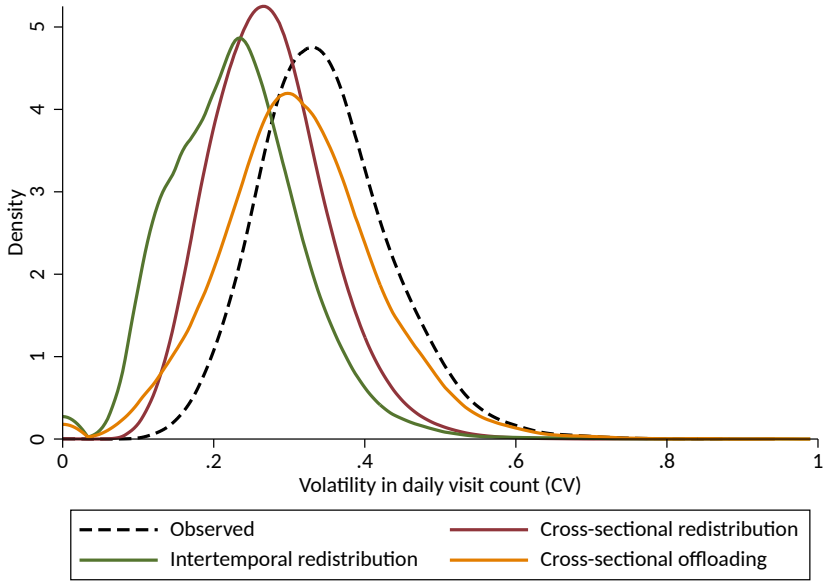
6.2.1 Simulated effects on schedule volatility experienced by full-time workers

Figure 7 plots the distribution of schedule volatility in daily visit count as observed in the data versus under the different policies we consider. All three simulated policies succeed in generally reducing schedule volatility experienced by full-time workers, though to varying degrees. The policy intervention that yields the greatest reduction in schedule volatility, on average, is inter-temporal smoothing through redistribution, where each worker's daily visits are evenly distributed across her work days within a calendar week. This policy manages to reduce the average schedule volatility in the sample of full-time workers by 34%. In contrast, cross-sectional smoothing through redistribution, where visits on a given day are evenly redistributed across on-duty full-time workers at the branch, reduces their average schedule volatility by 20%. Finally, cross-sectional smoothing through offloading visits on a given day to available part-time employees reduces the average schedule volatility of full-time workers by 11%. These effects are large, but as mentioned earlier, may compromise quality to some extent. Cross-sectional smoothing involves increasing the average number of different providers each patient sees, while inter-temporal smoothing involves changes to the timing of care, which may deviate from the original clinical care plan. Therefore, in a way, these effects should be viewed as the largest effect attainable using these policies, ignoring any other considerations.[11]

It is worth noting that the policy of cross-sectional smoothing through offloading visits to part-time workers seems to impact workers somewhat differently than the other two simulated policies. This is apparent in Figure 7, where the distribution resulting from simulating the cross-sectional smoothing through offloading policy exhibits greater variance in comparison to the

[11] In the case of the two policies that engage in cross-sectional smoothing, a small fraction of workers experience an occasional *increase* in schedule volatility. This is because the policy aims to smooth visits at the branch level across all full-time workers on duty on a given day. While the redistributive allocation can sometimes negatively impact a few workers on the margin, on average the policies yield substantial reductions in schedule volatility.

Figure 7: Distribution of Schedule Volatility Following Policy Simulations



Notes: Kernel density estimates of schedule volatility in visit count distributions for 275,322 worker-days in the sample. Observed distribution (dashed black) plotted against the simulated schedule volatility distributions resulting from each of the policies studied in Section 6.1.

other two simulated policies. Furthermore, the shape of the distribution for the other two simulated policies, both of which redistribute visits amongst full-time workers, is similar to the distribution that is observed in the data (albeit shifted to the left). To understand why this may be the case, we examine the extent to which each of the simulated policies results in a change in schedule volatility at the worker-day level.

We find that for the first two simulated policies where visits are redistributed cross-sectionally or inter-temporally, 91% and 98% of observations, respectively, see a change in schedule volatility. In contrast, for the third policy where cross-sectional smoothing is attained by offloading visits to available part-time workers, only 57% of observations experience a change in schedule volatility. In other words, for almost half of the worker-day level observations, this last policy simulation does not result in a transfer of visits from full-time workers to part-time workers when part-time workers are available. This is likely because part-time workers are not always available; there may not be any part-time workers on duty that day, the branch may not have any part-time workers on roster, or all existing part-time workers may already be performing at maximum capacity that day. Because of these operational constraints imposed by the policy, our findings may be biased towards the null in our consideration of this particular policy.

Of course, it would also be possible to combine all three policies. For example, we could start by implementing inter-temporal smoothing across full-time workers, then reallocate excess visits to part-time workers, and when this is not possible, redistribute the visits to other full-time workers. Using all three levers would make it easier for schedulers to weigh other worker- and

patient-level considerations. Furthermore, it would allow the firm to attain greater reductions in schedule volatility.

6.2.2 Simulated effects on turnover among full-time workers

Using the changes in schedule volatility resulting from each of the simulated policies, we calculate its potential impact on full-time workers' propensity to quit. To ensure that the predictions generated by our model are constrained to lie on the unit interval, we use estimates from the IV probit model discussed in section 5.2.4.

Table 6 summarizes the simulated effects of each policy on (a) the change in daily volatility and (b) the change in full-time workers' propensity to quit. We find that the average reduction in workers' propensity to quit is approximately proportional to the average decrease in their daily schedule volatility. The policy of cross-sectional smoothing through redistribution reduces workers' average daily propensity to quit by 11%, with a median decrease of 10%. With this policy, approximately 5% of the sample is negatively affected, though to a small degree. The policy of inter-temporal smoothing through redistribution yields the greatest reduction in voluntary separation, exhibiting a 16% reduction in the average daily propensity to quit among full-time workers. Furthermore, nearly all worker-days in the sample are positively affected, with the reduction in propensity to quit ranging from 4% at the 5th percentile to 32% at the 95th percentile. Finally, as expected given the muted effects on schedule volatility, the policy of cross-sectional smoothing through offloading produces the smallest change, where the average propensity to quit among full-time workers is reduced by 6%. For approximately 25% of the sample's worker-days, we observe that the policy has either a slightly negative effect or no effect on workers' likelihood of quitting.

Table 6: Effect of Policy Simulations on Schedule Volatility and Voluntary Separation

	(a) Change in daily volatility (%)		(b) Change in daily propensity to quit (%)						
Policy	Mean	SD	Mean	SD	5th Pctile.	25th Pctile.	50th Pctile.	75th Pctile.	95th Pctile.
Cross-sectional redistribution	-19.76	16.38	-11.20	9.33	0.57	-4.23	-10.04	-16.91	-28.16
Intertemporal redistribution	-33.69	18.93	-16.12	9.50	-3.60	-9.95	-15.31	-21.41	-32.43
Cross-sectional offloading	-10.61	19.40	-5.47	9.67	2.19	0.00	-0.71	-10.84	-22.49

Notes: Estimated change from enacting each policy on (a) daily volatility and (b) daily propensity to quit, expressed as a percent of baseline values for each worker-day. Predictions of baseline and counterfactual daily propensity to quit are calculated using IV probit estimates as described in section 5.2.4.

These simulation results illustrate that policy-induced changes to schedule volatility can have a measurable impact on reducing workers' propensity to quit. Furthermore, the policies described above are mechanical in nature and therefore relatively easy to implement. Nevertheless, we reiterate our caution from section 6.1 that reallocating patients using these crude policies may result in potential unintended consequences, which are beyond the scope of this simulation exercise.

7 Discussion

The implications of our findings are substantial, especially given the high costs of turnover beleaguering many firms. In the context of home health, the costs associated with turnover are estimated to be equivalent to 20% of a nurse's annual salary (Boushey and Glynn, 2012). Given an annual mean wage of $75,870 and 169,630 registered nurses working in home health as of 2020 (U.S. Bureau of Labor Statistics, 2021), we can estimate the costs associated with turnover to exceed $2.5 billion in just a single occupation within a single industry.

We expect our results to generalize to other occupations and other industries that are characterized by shift-based work, unstable work schedules, or just-in-time scheduling. While our causal analysis is limited to a specific occupation in a specific industry, we find that the association between schedule volatility and worker separation is observed in many other sectors of the economy. In panel A of Figure 8, we use data from the Bureau of Labor Statistics (BLS) to plot, for each economic sector defined by the North American Industry Classification System (NAICS), the average quarterly rate of worker separation against the share of workers in the sector with unpredictable work schedules, where unpredictability is defined as receiving notice of one's schedule less than 4 weeks in advance.[12] Because there are no general surveys (to our knowledge) that ask workers about their schedule volatility, we rely on schedule unpredictability as a proxy for schedule volatility, since workers' schedules are likely to be more volatile when they lack predictability in their schedules. We find a strong positive correlation between the quarterly rate of separation and schedule unpredictability ($\rho = 0.84$, weighted by sector size).

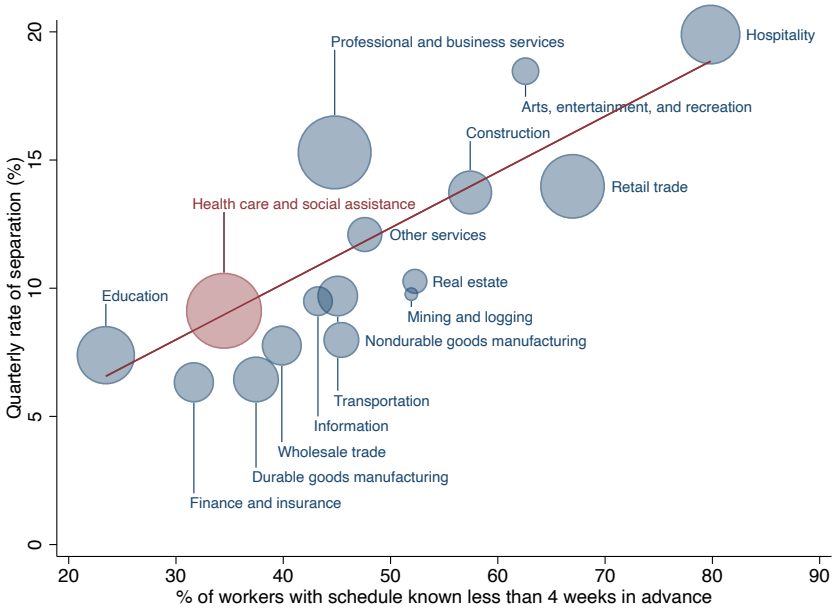
While the health care and social assistance sector, as defined by 2-digit NAICS codes, exhibits relatively low levels of schedule unpredictability and quarterly separation rates when considered in the aggregate, we observe substantial within-sector heterogeneity when we separate out the 4-digit NAICS code industry groups that comprise this sector. Panel B of Figure 8 shows, for each health care-related industry group within the health care and social assistance sector, the quarterly rates of separation against the share of workers in the industry group with unpredictable work schedules. Among the health care-related industry groups, we again find a strong positive correlation between the rate of worker separation and the unpredictability of workers' schedules ($\rho = 0.79$, weighted by industry size). In addition, we see that the home health industry exhibits the highest levels of worker separation and schedule unpredictability. We note that these measures include all occupations within the home health industry, including non-clinical administrative staff, who may have lower schedule unpredictability and lower separation rates than the nurses in our study.

It is important to note that the vast majority of home health nurses are female workers. Based on the literature on the value of non-pecuniary job attributes, we might expect female workers

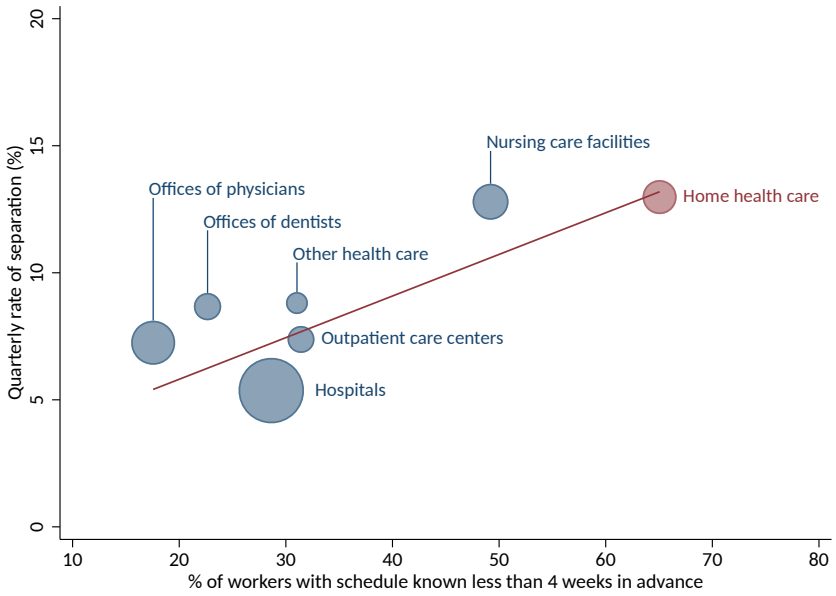
[12] We combine several BLS data sources to construct Figure 8. Data on the unpredictability of workers' schedules come from the 2017 American Time Use Survey (ATUS) Leave Module data. Data on quarterly rates of worker separation and number of employees come from the 2017 Quarterly Workforce Indicators (QWI) data. We use data from 2017 because it is the most recent available data for the ATUS Leave Module.

Figure 8: Schedule Unpredictability and Rates of Separation across Economic Sectors

Panel A: Average quarterly rate of separation and schedule unpredictability across economic sectors, 2017



Panel B: Average quarterly rate of separation and schedule unpredictability across industry groups within the health care sector, 2017



Notes: Scatterplots of the average national quarterly rate of worker separation in 2017 against the share of workers, nationally, who knew their primary work schedule less than 4 weeks in advance in 2017, across economic sectors (Panel A) and across industry groups within the health care sector (Panel B). For Panel B, we restrict to health care-related industry groups within the health care and social assistance sector in Panel A. Average quarterly separation rate is calculated derived from the BLS Quarterly Workforce Indicators (QWI) data, and is defined as, for each sector or industry group, the ratio of total separations in the quarter over the average number of employees in that quarter, averaged over Q1-Q4 of 2017. The share of workers who know their primary work schedule less than 4 weeks in advance is calculated using the BLS American Time Use Survey (ATUS) Leave Module Data. Marker size is proportional to the number of workers recorded in each sector or industry group in 2017 in the QWI data. Trend line plots the linear regression fit, weighted by sector or industry group size. The hospital industry group in Panel B includes general acute care and specialty hospitals but excludes psychiatric and substance abuse hospitals. Psychiatric and substance abuse hospitals, offices of chiropractors, and offices of optometrists are excluded from Panel B due to the small sample of ATUS respondents working in these industry groups.

to be more sensitive to schedule volatility then male workers (Mas and Pallais, 2017; Wiswall and Zafar, 2018; Chen et al., 2020). Thus, it is possible that the effects of schedule volatility on turnover may be more muted in male-dominated professions and industries.

Of course, several other factors may also impact workers' decisions to quit, such as low pay, poor management, lack of development opportunities, or geographic relocation. Nevertheless, schedule volatility seems to be an important determinant, and one that employers have the ability to influence operationally. Given what our findings show, it is crucial that employers understand and internalize the consequences of high schedule volatility. Especially in the current environment where there is a growing emphasis on work-life balance and employee-driven flexibility (Council of Economic Advisors, 2010; Goldin and Katz, 2011), finding a way to support stable schedules could be an important way for employers to attract and retain workers.

8 Conclusions

Employers across many sectors of the economy have been fast to adopt variable work scheduling policies, which seek to maximize firm profitability and reduce labor costs by "right-sizing" the workforce at the shift, or even hourly, level. The cost of this flexibility for employers is usually borne by the employees, for whom unstable work schedules create disruptions to work-life balance and the predictability of pay. In this paper, we show that volatility in work schedules can backfire for employers by leading to higher rates of voluntary separations among workers.

We use unique and detailed data of nurses conducting home health visits, employed by a large, multi-state home health agency, to construct several measures of schedule volatility at the worker-day level. Using Paid Days Off taken by other nurses in the same branch as an instrument for a focal nurse's schedule volatility, we quantify the causal impact of schedule volatility on voluntary employee turnover. Our findings show a strong positive link between the two, which holds across various measures of schedule volatility, regardless of whether the volatility is in the number of visits conducted per day, the total time spent working per day, or the length of the work day. The magnitude of the relationship is also economically significant: when measuring schedule volatility based on the number of visits conducted per day, we find that a single 30-day period of high schedule volatility (defined as one standard deviation above the average level of schedule volatility) increases the average worker's probability of quitting that year by 25.45% (SE: 5.08 p.p.).

Using our model estimates, we simulate various counterfactual policies to examine the extent to which employers may be able to reduce workers' schedule volatility and, consequently, their likelihood of quitting. Our simulations suggest that excess schedule volatility can explain a significant portion of voluntary turnover, and some allocative interventions have the potential to substantially reduce workers' daily propensity to quit. Among the policies we simulate, the policy of ensuring an even distribution of daily visits within a calendar week for each full-time worker yields the greatest reduction in average schedule volatility (34%); the resulting change

in daily propensity to quit is a 16% reduction. The implementation of the policies we consider is conceivably easy, though would require a re-prioritization of employee work-life balance over customer and firm-centered considerations.

References

Altonji, J. G. and C. H. Paxson (1988). Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs. *Journal of Labor Economics 6*(2), 254–276.

Angrist, J. D. and J.-S. Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Barmby, T. A., M. G. Ercolani, and J. G. Treble (2002). Sickness Absence: An International Comparison. *The Economic Journal 112*(480), F315–F331.

Bergman, A., H. Song, G. David, J. Spetz, and M. Candon (2021). The Role of Schedule Volatility in Home Health Nursing Turnover. *Medical Care Research and Review*, 107755872110343.

Black, S. E. and L. M. Lynch (1996). Human-Capital Investments and Productivity. *American Economic Review 86*(2), 263–267.

Boushey, H. and S. J. Glynn (2012). There Are Significant Business Costs to Replacing Employees. Technical report, Center for American Progress.

Chen, K.-M., N. Ding, J. A. List, and M. Mogstad (2020). Reservation Wages and Workers' Valuation of Job Flexibility: Evidence from a Natural Field Experiment. *SSRN Working Paper*.

Chen, M. K., J. A. Chevalier, P. E. Rossi, and E. Oehlsen (2019). The Value of Flexible Work: Evidence from Uber Drivers. *Journal of Political Economy 127*(6), 2735–2794.

Conroy, S. A., D. Roumpi, J. E. Delery, and N. Gupta (2021). Pay Volatility and Employee Turnover in the Trucking Industry. *Journal of Management*, 014920632110196.

Council of Economic Advisors (2010). Work-life balance and the economics of flexibility. Technical report, Executive Office of the President, Washington, D.C.

David, G. and T. Brachet (2011). On the Determinants of Organizational Forgetting. *American Economic Journal: Microeconomics 3*(3), 100–123.

David, G. and K. L. Kim (2018). The effect of workforce assignment on performance: Evidence from home health care. *Journal of Health Economics 59*, 26–45.

Deere, D. R. (1987). Labor Turnover, Job-Specific Skills, and Efficiency in a Search Model. *The Quarterly Journal of Economics 102*(4), 815.

Golden, L. (2015). Irregular Work Scheduling and Its Consequences. EPI Briefing Paper 394, Economic Policy Institute.

Goldin, C. and L. F. Katz (2011). The Cost of Workplace Flexibility for High-Powered Professionals. *The ANNALS of the American Academy of Political and Social Science 638*(1), 45–67.

Griffeth, R. W., P. W. Hom, and S. Gaertner (2000). A Meta-Analysis of Antecedents and Correlates of Employee Turnover: Update, Moderator Tests, and Research Implications for the Next Millennium. *Journal of Management 26*(3), 463–488.

Hall, J. V. and A. B. Krueger (2018). An Analysis of the Labor Market for Uber's Driver-Partners in the United States. *ILR Review 71*(3), 705–732.

Hall, R. E. and A. I. Mueller (2018). Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values. *Journal of Political Economy 126*(4), 1594–1637.

Henly, J. R. and S. J. Lambert (2014, July). Unpredictable Work Timing in Retail Jobs: Implications for Employee Work–Life Conflict. *ILR Review 67*(3), 986–1016.

Henly, J. R., H. L. Shaefer, and E. Waxman (2006). Nonstandard Work Schedules: Employer- and Employee-Driven Flexibility in Retail Jobs. *Social Service Review 80*(4), 609–634.

Hom, P. W., T. W. Lee, J. D. Shaw, and J. P. Hausknecht (2017). One hundred years of employee turnover theory and research. *Journal of Applied Psychology 102*(3), 530.

Jovanovic, B. (1984). Matching, Turnover, and Unemployment. *Journal of Political Economy 92*(1), 108–122.

Katz, L. F. and A. B. Krueger (2019). The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015. *ILR Review 72*(2), 382–416.

Kesavan, S., S. Lambert, J. Williams, and P. Pendem (2021). Doing Well By Doing Good: Improving Store Performance with Responsible Scheduling Practices at the Gap, Inc. *SSRN Working Paper*.

Kostiuk, P. F. (1990). Compensating Differentials for Shift Work. *Journal of Political Economy 98*(5, Part 1), 1054–1075.

Kozlowski, S. W., S. M. Gully, E. R. Nason, and E. M. Smith (1999). Developing adaptive teams: A theory of compilation and performance across levels and time. *Pulakos (Eds.), The changing nature of work performance: Implications for staffing, personnel actions, and development 240*, 292.

Lambert, S. J. and J. R. Henly (2014). Measuring Precarious Work Schedules. *A Working Paper of the EINet Measurement Group*.

Landers, R. M., J. B. Rebitzer, and L. J. Taylor (1996). Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms. *American Economic Review 86*(3), 329–348.

Maestas, N., K. Mullen, D. Powell, T. von Wachter, and J. Wenger (2018). The Value of Working Conditions in the United States and Implications for the Structure of Wages. *IZA Discussion Papers* (No. 11925).

Mas, A. and A. Pallais (2017). Valuing Alternative Work Arrangements. *American Economic Review 107*(12), 3722–3759.

Mas, A. and A. Pallais (2020). Alternative Work Arrangements. *Annual Review of Economics 12*(1), 631–658.

McLaughlin, K. J. (1991). A Theory of Quits and Layoffs with Efficient Turnover. *Journal of Political Economy 99*(1), 1–29.

Schneider, D. and K. Harknett (2019). Consequences of Routine Work-Schedule Instability for Worker Health and Well-Being. *American Sociological Review 84*(1), 82–114.

Song, H., E. Andreyeva, and G. David (2021). Time Is the Wisest Counselor of All: The Value of Provider–Patient Engagement Length in Home Healthcare. *Management Science*, mnsc.2020.3921.

Staw, B. M. (1980). The Consequences of Turnover. *Journal of Occupational Behaviour 1*(4), 253–273.

U.S. Bureau of Labor Statistics (2019). Job Flexibilities and Work Schedules — 2017-2018 Data From the American Time Use Survey. Technical Report USDL-19-1691.

U.S. Bureau of Labor Statistics (2021). May 2020 National Industry-Specific Occupational Employment and Wage Estimates. https://www.bls.gov/oes/current/naics4_621600.htm.

Wiswall, M. and B. Zafar (2018). Preference for the Workplace, Investment in Human Capital, and Gender*. *The Quarterly Journal of Economics 133*(1), 457–507.

Yu, Q., S. Mankad, and M. Shunko (2021). Evidence of The Unintended Labor Scheduling Implications of The Minimum Wage. *SSRN Working Paper*.

A Additional Results and Tables

Table A1: Effect of Other-Worker Paid Days Off in the Branch on Focal Worker's Schedule

	(1) Daily visit count	(2) Daily visit time	(3) Daily drive time	(4) Daily total time	(5) Time per visit	(6) Day length
Others' PDOs	0.0426**	0.0291**	0.0294**	0.0354**	-0.00742**	0.0340**
	(0.00218)	(0.00209)	(0.00311)	(0.00237)	(0.00106)	(0.00249)
	(0.218)	(0.209)	(0.311)	(0.237)	(0.106)	(0.249)
Observations	190,497	190,492	184,342	195,619	188,145	195,521
R-squared	0.075	0.078	0.097	0.059	0.096	0.042

Notes: PDO = Paid Days Off. Robust standard errors in parentheses. ** p<0.01, * p<0.05.

Table A2: 2SLS Estimation Results: Effect of Schedule Volatility on Full-Time Worker Separation

Panel A: First Stage Regression Estimates

Dep. var.: Volatility measure	(1) Daily visit count	(2) Daily visit time	(3) Daily drive time	(4) Daily total time	(5) Time per visit	(6) Day length
1(SD of others' PDOs > 0)	-0.0265*	-0.0329**	0.0214	-0.0153	0.0106	-0.00612
	(0.0114)	(0.0115)	(0.0117)	(0.0118)	(0.0119)	(0.0116)
SD of others' PDOs	0.0888**	0.0679**	0.0534**	0.0404**	0.0398**	0.0377**
	(0.0105)	(0.0107)	(0.0108)	(0.0114)	(0.0107)	(0.0110)
SD of others' PDOs, squared	-0.0534**	-0.0559**	-0.0538**	-0.0256*	-0.0615**	-0.0310**
	(0.00950)	(0.00989)	(0.0104)	(0.0106)	(0.00983)	(0.0102)
Average of others' PDOs	-0.0127**	-0.0125**	-0.00261	-0.00709*	-0.00884**	-0.00737*
	(0.00262)	(0.00288)	(0.00256)	(0.00305)	(0.00271)	(0.00291)
Average of others' PDOs, squared	0.000545	-0.00145	-0.00246	-0.00245	0.00612*	-0.00196
	(0.00278)	(0.00307)	(0.00282)	(0.00315)	(0.00297)	(0.00306)
Full-time workers in branch	0.0142**	0.00337	-0.0376**	-0.0225**	-0.0141**	-0.0208**
	(0.00415)	(0.00466)	(0.00492)	(0.00452)	(0.00474)	(0.00446)
Full-time workers in branch, squared	-0.000421*	0.000243	0.00224**	0.00175**	0.000633**	0.00180**
	(0.000174)	(0.000207)	(0.000223)	(0.000191)	(0.000219)	(0.000197)
Part-time workers in branch	-0.0100*	0.00936*	0.0133**	0.0157**	0.0162**	0.0135**
	(0.00433)	(0.00412)	(0.00405)	(0.00425)	(0.00426)	(0.00417)
Part-time workers in the branch, squared	0.00176**	-0.000339	-0.000406	-0.000800	-0.00260**	-0.00114**
	(0.000457)	(0.000428)	(0.000443)	(0.000453)	(0.000436)	(0.000432)
Tenure	0.000127**	4.75e-05**	5.87e-05**	-1.43e-05	-0.000106**	-0.000104**
	(1.33e-05)	(1.24e-05)	(1.35e-05)	(1.30e-05)	(1.36e-05)	(1.34e-05)
Work days in past 28 days	0.398**	0.281**	0.186**	0.293**	0.107**	0.289**
	(0.0104)	(0.00899)	(0.00842)	(0.0112)	(0.00922)	(0.0112)
ln(Work days in past 28 days)	-6.328**	-4.372**	-2.855**	-5.209**	-1.677**	-5.215**
	(0.190)	(0.162)	(0.150)	(0.203)	(0.168)	(0.204)
IV F-stat	55.13	61.51	42.75	21.09	35.98	26.31

Panel B: Second Stage Regression Estimates

Dep. var.: Y_{it}	(1) Daily visit count	(2) Daily visit time	(3) Daily drive time	(4) Daily total time	(5) Time per visit	(6) Day length
Volatility measure	0.0473**	0.0362**	0.0594**	0.0700**	0.0386**	0.0632**
	(0.0103)	(0.00850)	(0.0117)	(0.0172)	(0.0124)	(0.0145)
Full-time workers in branch	0.00225**	0.00306**	0.00542**	0.00458**	0.00407**	0.00442**
	(0.000655)	(0.000577)	(0.000687)	(0.000651)	(0.000578)	(0.000612)
Full-time workers in branch, squared	-6.63e-05*	-9.96e-05**	-0.000226**	-0.000206**	-0.000130**	-0.000200**
	(2.83e-05)	(2.51e-05)	(3.24e-05)	(3.31e-05)	(2.49e-05)	(3.04e-05)
Part-time workers in branch	-0.000233	-0.00105	-0.00161**	-0.00180**	-0.00134*	-0.00156*
	(0.000563)	(0.000546)	(0.000601)	(0.000651)	(0.000585)	(0.000609)
Part-time workers in the branch, squared	1.43e-05	0.000111*	0.000136*	0.000155*	0.000199**	0.000171**
	(5.87e-05)	(5.54e-05)	(6.03e-05)	(6.35e-05)	(6.61e-05)	(6.22e-05)
Tenure	-5.88e-06**	-1.56e-06	-3.61e-06*	1.15e-06	4.26e-06*	6.76e-06**
	(2.09e-06)	(1.62e-06)	(1.84e-06)	(1.77e-06)	(2.05e-06)	(2.29e-06)
Work days in past 28 days	-0.0244**	-0.0158**	-0.0166**	-0.0261**	-0.00978**	-0.0238**
	(0.00415)	(0.00245)	(0.00228)	(0.00510)	(0.00147)	(0.00425)
ln(Work days in past 28 days)	0.370**	0.230**	0.241**	0.436**	0.137**	0.401**
	(0.0663)	(0.0385)	(0.0357)	(0.0908)	(0.0238)	(0.0767)
Observations	276,549	276,549	274,749	276,549	276,549	276,549
Dep. var. mean	0.0130	0.0130	0.0130	0.0130	0.0130	0.0130

Notes: SD = Standard deviation. PDOs = Paid Days Off. Complete estimates from the LPM 2SLS model described in Equations 4-5. Newey-West (heteroskedasticity and autocorrelation robust) standard errors are in parentheses. Standard errors are calculated using 2 lags. ** $p<0.01$, * $p<0.05$.

Table A3: Robustness Test: 2SLS estimation results when using a dependent variable with a 28-day lead length

Dep. var.: Y_{it} (28-day lead)	(1) Daily visit count	(2) Daily visit time	(3) Daily drive time	(4) Daily total time	(5) Time per visit	(6) Day length
Volatility measure	0.0827**	0.0688**	0.0964**	0.128**	0.0597**	0.111**
	(0.0150)	(0.0126)	(0.0165)	(0.0277)	(0.0170)	(0.0223)
Full-time workers in branch	0.00397**	0.00532**	0.00934**	0.00812**	0.00704**	0.00777**
	(0.000912)	(0.000806)	(0.000964)	(0.000974)	(0.000776)	(0.000876)
Full-time workers in branch, squared	-0.000112**	-0.000168**	-0.000379**	-0.000363**	-0.000219**	-0.000345**
	(3.97e-05)	(3.57e-05)	(4.55e-05)	(5.14e-05)	(3.39e-05)	(4.48e-05)
Part-time workers in branch	-2.81e-05	-0.00151*	-0.00232**	-0.00286**	-0.00184*	-0.00235**
	(0.000773)	(0.000740)	(0.000820)	(0.000960)	(0.000775)	(0.000862)
Part-time workers in the branch, squared	-2.32e-05	0.000149*	0.000184*	0.000227*	0.000278**	0.000251**
	(7.91e-05)	(7.31e-05)	(8.03e-05)	(9.12e-05)	(8.53e-05)	(8.62e-05)
Tenure	-1.02e-05**	-2.94e-06	-5.73e-06*	2.15e-06	6.66e-06*	1.19e-05**
	(2.89e-06)	(2.20e-06)	(2.52e-06)	(2.59e-06)	(2.78e-06)	(3.43e-06)
Work days in past 28 days	-0.0399**	-0.0264**	-0.0248**	-0.0444**	-0.0135**	-0.0390**
	(0.00612)	(0.00373)	(0.00335)	(0.00825)	(0.00212)	(0.00662)
ln(Work days in past 28 days)	0.603**	0.382**	0.355**	0.747**	0.182**	0.659**
	(0.0980)	(0.0592)	(0.0532)	(0.147)	(0.0346)	(0.120)
Observations	276,554	276,554	274,754	276,554	276,554	276,554
IV F-stat	55.13	61.51	42.75	21.09	35.98	26.31
Dep. var. mean	0.0229	0.0229	0.0229	0.0229	0.0229	0.0229
Volatility measure coefficient, normalized to 14-day lead coefficient	0.0458**	0.0381**	0.0534**	0.0708**	0.0330**	0.0614**
	(0.00828)	(0.00697)	(0.00912)	(0.0153)	(0.00939)	(0.0124)

Notes: Second stage estimates from a 2SLS regression as described in Equations 4-5, where the dependent variable receives value on the day of termination and the 28 days leading to it. First stage estimates remain unchanged from our main specification presented in Table A2. Newey-West (heteroskedasticity and autocorrelation robust) standard errors are in parentheses. Standard errors are calculated using 2 lags. ** $p<0.01$, * $p<0.05$.

Table A4: Summary Statistics: Part-time Workers and All Workers

		Part-time Workers		All Workers	
		Mean	Std. Dev.	Mean	Std. Dev.
Worker-day level	Number of visits conducted	3.04	1.88	4.00	1.90
	Time spent conducting visits (hours/day)	2.67	1.61	3.56	1.66
	Time spent conducting nonvisit activities (hours/day)	0.41	1.10	0.62	1.24
	Time spent driving (hours/day)	1.00	10.03	1.49	4.19
	Total time (hours/day)	4.08	10.31	5.67	4.78
	Average time per visit (hours/day)	0.85	0.28	0.84	0.25
	Work day start time	10:10	02:30	09:14	02:02
	Work day end time	15:11	02:44	16:00	02:25
	Length of work day (hours/day)	5.11	10.47	6.81	5.04
Worker-week level	Number of days worked	3.22	1.76	4.42	1.62
	Time spent conducting visits (hours/week)	8.60	6.39	15.73	7.65
	Total time (hours/week)	13.15	20.48	25.08	14.95
	Length of work day (hours/week)	16.44	21.90	30.10	16.97
Separations	All separations (% of workers)	51.40	50.04	47.80	49.97
	Voluntary separations (%)	46.26	49.92	42.57	49.46
	Average tenure at voluntary separation (months)	8.21	5.47	8.87	5.56
	Involuntary separations (%)	5.14	22.11	5.23	22.26
	Average tenure at involuntary separation (months)	8.41	5.52	9.44	6.96
Demographics	Age at hiring	45.72	12.42	43.12	11.34
	Female (%)	93.22	25.16	92.99	25.54
Sample size	Number of workers	428		1,569	
	Number of worker-days	74,079		350,633	
	Number of visits	115,194		936,054	

Notes: Key summary statistics of the analysis sample of part-time workers (i.e., part-time registered nurses (RNs)), and all workers (i.e., part-time and full-time RNs). [a]Total time is the sum of time spent conducting visits, time spent conducting nonvisit activities, and driving.