

NBER WORKING PAPER SERIES

EARNINGS INSTABILITY

Peter Ganong
Pascal J. Noel
Christina Patterson
Joseph S. Vavra
Alexander Weinberg

Working Paper 34227
<http://www.nber.org/papers/w34227>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2025, Revised July 2026

We thank Fiona Greig for the initial conversations which inspired this paper. We thank Niels Gormsen, Simon Jaeger, Emir Kamenica, Greg Kaplan, Anil Kashyap, Erik Hurst, Thibaut Lamadon, Susan Lambert, Marta Lachowska, Bentley MacLeod, Alan Manning, Alex Mas, Evan Rose, Jon Roth, Danny Schneider, and Steven Woodbury for helpful conversations, and numerous seminar participants for suggestions. We thank the Data Science team at PayrollCompany and Daniel Sullivan and Sarah Kuehl at JPMorganChase Institute (JPMCI). We thank Rupsha Debnath, Felix Farb, Avik Garg, Koren Gila, Theodore Grayer, Sedona Jolly, Min Seo Kim, Jay Leeds, Paul Metzler, Michael Meyer, John Spence, Anna Volpp, and Madeline Zuckerman for excellent research assistance. We thank the Becker Friedman Institute and the Kathryn and Grant Swick Faculty Research Fund at the University of Chicago Booth School of Business for financial support. This research was made possible by the University of Chicago's Research Computing Center cluster, by a data-use agreement between the authors and PayrollCompany, and by a data-use agreement between the authors and JPMCI, which has created de-identified data assets that are selectively available to be used for academic research. All statistics from JPMCI data, including medians, reflect cells with multiple observations. The opinions expressed are those of the authors alone and do not represent the views of JPMorgan Chase & Co or the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Peter Ganong, Pascal J. Noel, Christina Patterson, Joseph S. Vavra, and Alexander Weinberg. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Earnings Instability

Peter Ganong, Pascal J. Noel, Christina Patterson, Joseph S. Vavra, and Alexander Weinberg

NBER Working Paper No. 34227

September 2025, Revised July 2026

JEL No. E21, J23, J31, J33

ABSTRACT

This paper uses high-frequency administrative data to show that the majority of U.S. workers experience substantial month-to-month fluctuations in pay, even within ongoing employment relationships. This earnings instability is pervasive, but it has been masked in past analysis of annual data. Moreover, this instability is unequally distributed: lower-income, hourly workers face more instability than higher-income, salaried workers. This is because earnings instability arises in large part from firm-driven fluctuations in hours. This earnings instability is a meaningful source of economic risk: we provide evidence that it increases consumption volatility and leads to greater job separations. These findings suggest that short-term earnings risk is a significant feature of the labor market and that this risk falls disproportionately on the most financially fragile workers.

Peter Ganong
The University of Chicago
Harris School of Public Policy
and NBER
ganong@uchicago.edu

Joseph S. Vavra
The University of Chicago
Booth School of Business
and NBER
joseph.vavra@chicagobooth.edu

Pascal J. Noel
The University of Chicago
Booth School of Business
and NBER
pascal.noel@chicagobooth.edu

Alexander Weinberg
The University of Chicago
weinberga@uchicago.edu

Christina Patterson
The University of Chicago
Booth School of Business
and NBER
christina.patterson@chicagobooth.edu

1 Introduction

This paper uses administrative payroll data to investigate the prevalence, causes, and consequences of high-frequency earnings changes among U.S. workers. We find that beneath the surface of stable employment lies a pattern of instability: for most workers, earnings vary substantially from month to month, even within continuing employment relationships. Although monthly instability is apparent in household surveys, it is masked in the large-scale annual administrative income data that now underpin much of the economics literature. Using high-frequency data covering millions of workers, we show that this instability is large, unequally distributed, and economically meaningful.

More specifically, we document five facts. First, consistent with survey evidence in the Survey of Income and Program Participation (SIPP) and detailed financial diaries from Morduch and Schneider (2017), we document substantial month-to-month earnings fluctuations. Importantly, these fluctuations are substantial even among workers who remain in the same job. Second, we show that this instability is concentrated among hourly workers, who make up 60 percent of the labor force and tend to have lower income than salaried workers. Although it has been well documented that *wages* are largely stable from month to month, we find that *hours* fluctuate substantially. Thus, wage stability does not translate into earnings stability for most U.S. workers. Third, we present evidence that these fluctuations in hours are driven in large part by fluctuations in firm labor demand. Fourth, we find that earnings instability leads to more volatile spending. Fifth, we show that workers are more likely to separate from high-volatility jobs. Together, these results suggest that high-frequency earnings risk is an important feature of the labor market and that this risk falls disproportionately on the most financially fragile workers.

Most of the economics literature on earnings risk has focused on annual income dynamics. That work relies either on survey data, especially the Panel Study of Income Dynamics,¹ or on large administrative datasets such as Social Security records (Guvenen, Ozkan, and Song 2014), Census data (Abowd and McKinney 2024), tax data (Pruitt and Turner 2020), and linked datasets combining information from multiple sources (Ziliak, Hardy, and Bollinger 2011; Moffitt et al. 2022). These large-scale datasets have the power to establish representative patterns and determinants of income changes, but not at high frequencies.

In the United States, direct evidence on sub-annual earnings volatility has been built primarily from surveys. Bania and Leete (2009) use the SIPP to show that monthly income volatility among low-income households rose after the mid-1990s welfare reforms. Hannagan and Morduch (2015) and Morduch and Schneider (2017) then provide especially vivid evidence on month-to-month instability by using detailed financial diaries to show how volatile earnings shape the financial lives of low- and moderate-income households over the course of a year. These studies directly capture high-frequency fluctuations, but their smaller samples make it harder to identify the forces

¹Some prominent examples using PSID data include Gottschalk and Moffitt (1994), Haider (2001), and Meghir and Pistaferri (2004). See Moffitt and Zhang (2018) for an overview of many of these PSID studies. A more recent strand of the survey-based literature leverages consumer expectations data, such as Caplin et al. (2023) and Arellano et al. (2026).

driving this instability and to evaluate how broadly this instability extends across the labor force. A nascent literature using high-frequency administrative data from outside the United States has also documented substantial monthly earnings volatility (Druehl, Graber, and Jørgensen 2023; Brewer, Cominetti, and Jenkins 2025; Andresen et al. 2025).

We extend this literature using large-scale high-frequency U.S. administrative data from both the firm side (via a payroll processor) and the worker side (via paycheck deposits into Chase bank accounts). The payroll-level data allow us to measure detailed components of individual pay, including hours, wages, and bonuses for the universe of workers at a given firm. The bank account data enable us to link earnings data with spending and liquidity and to aggregate jobs into households. Throughout, we focus on pay variation within continuing jobs.² Our data allow us both to generalize the patterns shown by the earlier survey-based literature and to go beyond that work by showing where this instability comes from and how it affects different types of workers.

Our first finding is that monthly earnings volatility is substantial. In about three quarters of months, workers' pay differs from the prior month. The median month has a change of 5 percent, and in one quarter of months the change in pay is at least 17 percent. These earnings changes are large relative to changes in wages and relative to typical liquidity holdings. We also find that these changes are fairly transitory, which makes them difficult to detect in annual data. As a result, leading heterogeneous-agent models that infer high-frequency dynamics from annual moments miss much of the within-year instability we observe directly in the data (e.g., Kaplan, Moll, and Violante 2018; Maxted, Laibson, and Moll 2025; Crawley, Holm, and Tretvoll 2026; Kaplan and Violante 2022).

Second, we find that earnings volatility is concentrated in hourly work. Earnings for hourly workers change in almost every month, and these changes are often quite large: the median change is 9 percent, and in one quarter of months earnings change by at least 21 percent. The vast majority of this earnings volatility is driven by fluctuations in hours rather than wages. Moreover, these fluctuations do not follow predictable annual patterns. In contrast, salaried workers' pay rarely varies from month to month. When it does vary, it usually varies to the upside and arises from annual wage increases and performance pay such as bonuses and commissions. Because hourly workers also tend to have lower incomes and less liquidity, this means that the workers most exposed to pay volatility are on average the ones least able to absorb it.

The concentration of earnings instability among low-wage hourly workers is consistent with a long tradition in labor economics linking job instability to job quality. Dual labor market theory emphasizes that many low-wage jobs are characterized by unstable hours, high turnover, and limited advancement, while higher-wage workers enjoy more stable employment, with firms often using more peripheral workers to absorb fluctuations in demand (Doeringer and Piore 1971; Rebitzer and Taylor 1991).³ Recent empirical work has reinforced this picture: Morduch and Schneider (2017) find that

²Volatility arising from unemployment and job transitions—which we cannot reliably measure in our data—would only amplify our conclusion that monthly pay volatility is large.

³For earlier work on labor market segmentation and the use of peripheral workers to buffer demand fluctuations, see Reynolds (1951), Myers and Shultz (1951); for a qualitative account of instability among peripheral low-wage

instability is concentrated among low-income households, and Blundell et al. (2025) show that annual earnings volatility for less-educated Black men is more than twice that of college-educated white men. Our findings are consistent with this tradition and highlight the distinction between hourly and salaried work as an important dimension of labor market heterogeneity. More broadly, our results suggest that pay volatility is an important dimension of labor-market inequality that is not captured by wages alone. This connects our findings to a growing literature on job amenities and total compensation, which has emphasized wage levels but not pay stability.⁴

Our third finding is that firms play an important role in generating workers' hours volatility. We show that firms experience substantial month-to-month fluctuations in *total* hours, likely reflecting fluctuations in labor demand. To our knowledge, this is the first paper to use high-frequency firm-level data to document labor demand shifts at this frequency. We find that these fluctuations are large enough to explain at least half of a typical worker's hours volatility.⁵ We reach this conclusion using two approaches: one based on the excess reallocation statistic in Davis and Haltiwanger (1992), and one based on movers between firms with different levels of total-hours volatility. The remaining volatility might arise through various channels like labor supply choices or idiosyncratic firm scheduling practices. Even though a substantial fraction of volatility appears involuntary, theory still suggests two reasons why such volatility might nevertheless have limited consequences for workers: First, these earnings changes are relatively transitory, and many theoretical models imply that transitory shocks should have limited effects on real outcomes. Second, if workers learn about hours changes in advance, they may be able to blunt their impact.

Our final two findings provide direct empirical evidence that monthly earnings instability nonetheless has meaningful consequences for workers. First, using bank account data, we show that earnings volatility translates into spending volatility. Our primary strategy instruments for individual earnings volatility with firm-level volatility and finds that spending volatility increases when workers move to firms with more volatile earnings. The effects are especially large for workers with low liquidity, consistent with the presence of binding budget constraints.

Finally, hourly workers are more likely to quit high-volatility jobs, suggesting that high volatility is a disamenity. Although we use both individual fixed effects and instrumental variables to address the potential correlation between worker and job characteristics, we cannot fully separate earnings volatility from other unobserved attributes of high-volatility jobs. We therefore interpret this evidence as showing that monthly earnings instability is part of a broader bundle of undesirable attributes associated with low-quality hourly jobs.

Taken together, these final two facts show that high-frequency earnings instability has meaningful consequences for households and workers. Our spending results may be surprising from the perspective of consumption-smoothing models in which transitory pay fluctuations have limited

workers, see Ehrenreich (2001).

⁴See, e.g., Maestas et al. (2023), Mas and Pallais (2017), Sorkin (2018), Humlum, Rasmussen, and Rose (2025), and the recent review by Mas (2025).

⁵Consistent with the interpretation that many hours fluctuations are involuntary, Lachowska et al. (2026) find that most workers would prefer to work more hours than their employer offers.

effects on spending, a logic that has also led parts of the inequality literature to emphasize persistent rather than transitory income changes (e.g., Moffitt and Gottschalk 2012). However, they are consistent with a large literature documenting spending responses to transitory income changes, including the tax-rebate evidence in Johnson, Parker, and Souleles (2006). For a given pass-through of income to spending, more frequent income fluctuations imply more volatile consumption. By showing that monthly earnings are highly unstable and that this instability passes through to spending, our results connect the literature on consumption smoothing with the literature measuring income volatility.

Our findings also relate to a growing cross-disciplinary literature documenting negative effects of earnings and schedule instability for workers. Lambert, Henly, and Kim (2019) show that greater volatility in weekly work hours is associated with greater perceived financial insecurity among hourly workers. Schneider and Harknett (2019) find that routine schedule instability among hourly retail workers is associated with worse psychological distress, poorer sleep quality, and lower happiness. Other work provides quasi-experimental evidence that pay volatility increases turnover in the retail and home health sectors (Kesavan and Kuhnen 2017; Bergman, David, and Song 2023). We provide systematic evidence that these relationships are not confined to specific sectors or occupations, but appear across the U.S. labor market.

This evidence raises a natural question: if instability is costly for workers, why do firms generate so much variation in hours? One reason may be that firms are making a “mistake,” in the sense that improved management practices could smooth hours for workers at little cost to firm profits. But another reason may be that variation in hours input is closely tied to variation in demand for the firm’s output. Firms may have some scope to insure workers against earnings changes (e.g., salaried workers receive the same base pay each month), but in many production settings they have less scope to smooth labor input itself. The timing of demand may simply require more hours of labor input in some months than in others. In that case, hourly contracts can help firms elicit additional labor input when demand is high, but they do so partly by shifting demand risk onto workers. This relates to the literature on implicit contracts, which shows that there is a tradeoff between insurance and incentives when labor input is difficult to monitor or specify directly.⁶

This perspective also helps to connect our results to the literature on worker insurance. That literature studies how firm shocks pass through to workers’ earnings. For example, Guiso, Pistaferri, and Schivardi (2005) show limited pass-through of transitory firm value-added shocks into workers’ annual earnings in Italy. Our findings do not contradict this result. Instead, they show that sizable residual instability remains at the monthly level after whatever insurance firms provide. Moreover, our heterogeneity results emphasize that the shocks that do get passed through to workers are disproportionately borne by a specific subset of workers: low-wage hourly employees. These are exactly the workers who are more financially fragile and who could in principle benefit from greater insurance.

⁶See Hart and Holmström (1987), Prendergast (2002), MacLeod and Parent (1999), and Lemieux, MacLeod, and Parent (2009).

2 Data Description and Measuring Volatility

We use two main datasets to study monthly earnings volatility. The primary dataset comes from a payroll processor which reports detailed information about employees’ paychecks. The second is bank account data from the JPMorganChase Institute where we can link income with spending and liquidity.

2.1 Payroll Records

Our primary data source consists of de-identified administrative earnings records from an anonymous U.S. payroll processor, hereafter referred to as the “PayrollCompany.” We work with data from 2010 to 2023, and there are between 2 and 5 million workers in the data at any point in time. Most PayrollCompany clients are small firms. Looking at a dataset where each observation is a worker-month, the median firm size is 18 employees. However, we show that earnings volatility is nearly identical when using bank account data, which captures a more representative distribution of firm sizes.

For each paycheck, we observe regular pay, bonuses, commissions, overtime, and paid leave. In addition, almost all salaried workers have a pre-populated value each pay period for regular pay and almost all hourly workers have an hourly wage rate. We refer to this pre-populated value as the “base wage.” We observe both gross pre-tax earnings and net earnings after withholding; our analysis focuses primarily on gross earnings. For a subset of workers, we also observe job title, age, gender, number of dependent children from the IRS Form W-4, and reason for separation.

High-frequency earnings variation is potentially consequential because a large share of household spending is committed and difficult to adjust at short horizons (Chetty and Szeidl 2007). We focus on monthly variation for three reasons: it is the shortest interval that accommodates the range of pay schedules of the workforce; many major household obligations, such as rent, mortgages, and car payments, recur at roughly monthly frequencies; and households themselves commonly budget at that horizon (Zhang et al. 2022).⁷ At the same time, we show that key findings are similar at quarterly frequencies, suggesting that the instability we document is not merely an artifact of measuring volatility at the monthly horizon.

Throughout, we focus on pay variation within continuing jobs. As noted above, this is a conservative choice: volatility from job transitions and unemployment would only amplify the instability we document. We adopt this focus because when earnings disappear, we cannot distinguish non-employment from employment that is not observed in our data.⁸ We then clean the data to remove two sources of monthly pay variation that do not reflect genuine within-job instability. First, for workers paid weekly or biweekly, total monthly pay fluctuates mechanically with calendar timing. To remove this source of variation, we normalize total monthly pay by the number of paychecks

⁷2023 Consumer Expenditure Survey data show that nearly 50% of household spending goes to housing and transportation—expenses that typically recur monthly.

⁸This is true in the PayrollCompany data since it covers only a subset of firms, and in the Chase data because not all jobs are paid via direct deposit into Chase bank accounts.

received in that month and define our monthly pay measure as average pay per paycheck. For example, in each calendar year, workers who are paid weekly receive four checks in roughly eight months and five checks in the remaining four months. If a worker receives the same amount in *each* paycheck but happens to get an extra paycheck in one month relative to another, this measure will correctly reflect that their weekly earnings are stable from month to month. Second, we exclude monthly volatility arising from partial months of employment. Specifically, we define a job spell as a continuous series of months with positive earnings. We then exclude the first and last month of each worker’s job spell because lower measured pay per paycheck during these months may reflect partial employment during that pay period rather than true within-job volatility.

Since much of our analysis focuses on hourly workers, we exclude workers whose contract type (hourly versus salaried) cannot be reliably classified. Specifically, we drop workers without a consistently observed positive base wage (11 percent of worker-months) and those whose base wage changes in at least half of the months they are observed (4 percent of worker-months). We implement a few additional sample restrictions which are described in Appendix A.1. For computational feasibility, our baseline sample uses a 1 percent random sample of firms, yielding a final sample of 19,893 firms.⁹

Despite focusing on small firms, the PayrollCompany data appears representative of the U.S. workforce along several dimensions. The distributions of wages (Figure A-1), quarterly hours (Figure A-2), pay frequency (Figure A-3), aggregate seasonality (Figure A-4), and monthly hire and separation rates (Figure A-5) are all similar to those in administrative benchmarks and BLS data. Both in our sample and in representative benchmarks, 60 percent of workers are hourly and 40 percent are salaried (Table A-1).

Our analysis of the sources of hours fluctuations requires observing hourly workers who move between firms. Since we can only track moves between two PayrollCompany clients, we construct a second sample focused on such transitions. We start by constructing a random sample of 1,000 firms in which we observe 8 or more such transitions.¹⁰ We then combine this set of 1,000 firms with all other firms that are directly linked to this set through at least one move by a worker who is on an hourly contract at both the origin and destination firms, resulting in a total sample of 20,319 firms and 68,995 moves. Finally, we sample *all* of the hourly workers at these 20,319 firms (i.e., including both the workers who move between PayrollCompany firms and those who do not). We use this sample of firms in the firm-level analysis in Sections 5.2 and 7.

Finally, even after the data-cleaning steps described above, some of what we measure as earnings volatility may reflect residual measurement error rather than genuine instability. For example, our procedure would still register volatility for a worker who is paid weekly but predictably works every

⁹We also investigated whether temporary-help agencies account for a meaningful share of instability in our sample. They represent a very small share of worker-month observations (both in our data and in survey data from the BLS) and do not materially affect our results.

¹⁰The minimum value of 8 selects a sample with a non-trivial number of transitions. Sampling a random subset of 1,000 firms meeting this criteria is for computational feasibility since we must draw data on all workers at firms connected by these moves.

other Friday, even though the underlying schedule is fixed.¹¹ Yet in Sections 6 and 7, we show that earnings volatility predicts both spending volatility and quit behavior. If volatility simply reflected measurement error, it should not systematically affect either outcome.

2.2 Bank Records

We supplement the PayrollCompany data with data on Chase bank customers from the JPMorgan-Chase Institute (JPMCI). We use the same sample as Ganong et al. (2025), and we refer readers to that paper for details on data construction and sample selection.

The main value of the JPMCI data is that they allow us to link income to spending and liquidity. This builds on prior JPMCI work examining links between income and expense volatility (Farrell and Greig 2015; Farrell, Greig, and Yu 2019). The JPMCI data also help address two limitations of the PayrollCompany data: they include employees at large firms and they capture income from multiple jobs, allowing us to study volatility at the household level.

Our measure of earnings comes from payments made by employers via direct deposit. We observe the amount and date of the deposit as well as the counterparty who made the deposit. This counterparty identifier enables us to identify other Chase customers paid by the same employer. Because these data capture what workers actually receive, our JPMCI earnings measure is net of withholding. We clean the JPMCI data in the same way that we clean the PayrollCompany data, constructing average monthly pay per paycheck and excluding the first and last month of each employment spell. As in the PayrollCompany data, our main measure of earnings in the bank account data is at the job level. However, since we observe all employer direct deposits into a household’s Chase bank accounts, we can combine income from various jobs and construct a measure of total household earnings. We use this household-level measure to study whether households smooth income fluctuations from one job using income from another job.

Because we do not observe the terms of job contracts, we construct an imputed indicator for whether a job is “pseudo-hourly” or “pseudo-salaried” based on the properties of its pay stream. Specifically, within each job spell, we classify a worker as pseudo-hourly if pay changes by more than 0.01 percent in at least 70 percent of months and average pay per check is below \$2,000.¹² We use this imputed classification only to define subsamples in the JPMCI analysis; we do not use it to explain earnings volatility, since the classification is itself based on variation in pay.

We also observe spending and liquidity for each account. Our main spending measure captures expenditures on nondurable goods and services and is constructed following Ganong et al. (2025). Examples include groceries, food away from home, fuel, utilities, clothing, medical co-pays, and payments at drugstores. Spending is measured using debit and credit card transactions, cash

¹¹Our paycheck-count adjustment removes variation from months with four versus five paychecks, but some months will still contain more actual work days than others. We do not observe which days are worked, but our key findings remain similar if we measure average monthly pay per calendar day rather than pay per paycheck.

¹²We can evaluate the accuracy of the prediction rule in the PayrollCompany data and we find that these thresholds accurately impute whether a worker is hourly or salaried in 86 percent of cases. These are the thresholds which maximize the accuracy of the prediction rule.

withdrawals, and electronic transactions observed in the bank account. We measure liquidity as the checking account balance at the end of each month.

2.3 Measuring Earnings Instability

Our main measure of earnings growth is the percent change in pay per paycheck:

$$\Delta_{i,t} = \frac{y_{i,t} - y_{i,t-1}}{y_{i,t-1}} \quad (1)$$

where $y_{i,t}$ is average earnings per paycheck in month t for worker i .¹³ We winsorize this variable at the 2.5th and 97.5th percentiles of non-zero changes to limit the influence of outliers.¹⁴

We begin by describing the distribution of monthly percent changes $\Delta_{i,t}$. To do so, we pool all of the monthly earnings changes together across workers and time periods to compute a cross-sectional distribution. This approach is typical in the literature studying annual earnings fluctuations (e.g., Guvenen, Ozkan, and Song 2014). Its main advantage is that it combines a large number of observations, allowing us to flexibly characterize the full distribution of changes. If workers draw earnings changes from a common distribution, then this cross-sectional distribution is informative about the earnings process faced by individual workers. However, this is a strong assumption, and the pooled approach is also not well-suited for answering questions about how worker-specific volatility affects behavior.

Therefore, when studying heterogeneity across workers, we instead summarize volatility as $Vol_i = Median(|\Delta_{i,t}|)$, where the median is taken across all monthly earnings changes for worker i within a job spell. We focus primarily on the median absolute change because it captures the typical earnings change faced by a worker and is robust to outliers. In Section 4, we discuss in more detail why we prefer this measure to the standard deviation. To distinguish between pooled measures and individual-specific measures, we index individual volatility statistics with subscript i . For example, Vol_i refers to the median absolute percent change for worker i .

3 Fact 1: U.S. workers experience substantial monthly earnings volatility

We begin by documenting that workers face substantial month-to-month earnings fluctuations. Panel (a) of Figure 1 shows the cumulative distribution function of monthly earnings changes while panel (b) shows the corresponding histogram. Panel A of Table 1 reports summary statistics for this distribution. In almost 70 percent of months, workers receive a different amount of pay than in the prior month. Moreover, these changes are often large. The median monthly change in pay

¹³By construction, $\Delta_{i,t}$ only captures changes across months with non-zero earnings. Including months with zero earnings would only strengthen the conclusion that earnings volatility is large.

¹⁴Table A-2 shows the impact of alternative winsorization choices.

Table 1: Summary Statistics of Earnings Changes

Sample	Variable	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Skewness	Kurtosis
A. Baseline							
All	Total Earnings	0.69	0.05	0.17	0.25	1.80	6.46
All	Base Wage	0.10	0.00	0.00	0.03	4.75	65.99
B. Alternative Measures of Earnings							
All	Quarterly Earnings	0.80	0.06	0.16	0.20	1.32	4.30
All	Net Earnings	0.76	0.05	0.17	0.24	1.65	5.66
JPMCI	Net Earnings	0.79	0.05	0.17	0.23	1.56	5.29
C. Hourly vs Salaried							
Hourly	Total Earnings	0.91	0.09	0.21	0.26	1.50	4.64
Hourly	Base Wage	0.12	0.00	0.00	0.03	4.58	65.33
Hourly	Hours	0.90	0.07	0.19	0.22	1.10	2.93
Salaried	Total Earnings	0.34	0.00	0.05	0.23	2.42	10.58
Salaried	Base Wage	0.07	0.00	0.00	0.04	4.86	64.24
D. Hourly Subsamples							
Full-time Hourly	Total Earnings	0.90	0.07	0.16	0.17	0.90	2.27
Prime-age Hourly	Total Earnings	0.93	0.09	0.21	0.25	1.51	4.94
No Overtime Hourly	Total Earnings	0.86	0.09	0.24	0.28	1.44	3.94

Notes: This table reports distributional statistics of percent change in pay from the prior month. “Total Earnings” is pre-tax earnings while “Net Earnings” captures pay net of withholding and deductions. “Base Wage” in row 2 is defined as base wage per hour for hourly workers and per-pay-period base salary for salaried workers. All data is from PayrollCompany except “JPMCI” which uses data on Chase bank customers from JPMorganChase Institute. PayrollCompany and JPMCI data are analyzed separately and were not merged as part of this analysis. Before computing higher-order moments (standard deviation, skew, kurtosis), the measures of change are winsorized at the 2.5th and 97.5th percentiles of non-zero changes. Table A-2 shows the standard deviation using alternative thresholds. Full-time hourly is defined as those who work an average of at least 30 hours per week and prime-age hourly is hourly workers age 25 to 54.

is 5 percent, and in one quarter of months the change is at least 17 percent.¹⁵

These patterns are not specific to the PayrollCompany data. Monthly income volatility within jobs in the JPMCI data is nearly identical to that in the PayrollCompany, as shown in Figure A-7 and panel B of Table 1. Because the JPMCI data include workers at both small and large firms, this implies that monthly pay volatility is pervasive across the firm size distribution—and indeed, Table A-3 shows similar volatility for employees of small and large firms.


Monthly income volatility is large in three regards. First, comparing the blue lines to the orange lines in Figure 1, we see that total earnings changes are an order of magnitude larger than changes in base wages. In contrast to the fact that earnings change in 70 percent of months, base wages only change in 10 percent of months. Moreover, the base wage almost never falls, consistent with previous work using administrative payroll data in Grigsby, Hurst, and Yildirmaz (2021), while monthly earnings fall in over a quarter of months. These results show that while ongoing employment relationships exhibit substantial wage rigidity, they do not exhibit the same rigidity in hours or other components of pay.

Second, monthly income changes are large relative to typical household liquidity levels. Using


¹⁵These exhibits pool observations from many workers, some of whom might have very large earnings risk and others of whom might much less risk. Figure A-6 shows the distribution of individual-level volatility Vol_i . It shows that volatility is high for the typical worker but also reveals heterogeneity across workers, which we explore later in the paper.

Figure 1: Within-Job Earnings and Wage Volatility

(a) Cumulative Distribution Function



(b) Histogram



Notes: This figure shows the within-job distribution of the change in earnings (blue) and wages (orange) from the prior month. “Total Earnings” is average pay per paycheck to abstract from pay schedule-driven fluctuations. “Base Wage” is the hourly wage for hourly workers and the per-period base salary for salaried workers.

the JPMCI data, we compare each worker’s monthly change in income to their median checking account balance. In forty percent of months, workers have an absolute dollar change in income

that exceeds 50 percent of their median checking account balance, and in one-fourth of months the absolute change exceeds the median balance.

While monthly earnings changes at a single job are large relative to household liquidity, households may have other jobs whose earnings partially offset these changes. Table A-3 analyzes this possibility using the JPMCI data and finds that such offsetting is limited. Total household earnings volatility is, if anything, slightly greater than volatility at the individual-job level. This is somewhat counterintuitive, since we might expect some decline in volatility from aggregating multiple jobs unless the jobs are perfectly correlated. However, this slight increase in volatility when moving from the job to the household level reflects the pooling of hourly and salaried jobs, which have very different earnings-change distributions, as we show in the next section.¹⁶

More relevant to our focus in the rest of the paper, we also analyze the effect of aggregating two hourly jobs, and still find high volatility at the household level. The bottom two rows of Table A-3 show that smoothing goes in the expected direction in this case, but the dampening in volatility from second jobs is very small.¹⁷ Thus, it appears that households have little ability—or little willingness—to use second jobs to offset the high degree of job-level income volatility that we document.

The third sense in which monthly income volatility is large is that the monthly income changes in the PayrollCompany data are far larger than what leading income models imply when calibrated to annual data. Figure 2 compares the observed distribution of monthly earnings changes with those implied by several income processes frequently used in macro models. The models substantially understate the frequency of large month-to-month earnings changes. For example, the 75th percentile of the absolute percent change in earnings is between 0 and 8 percent in the models, compared with 17 percent in the data, while the 90th percentile is between 10 and 16 percent in the models, compared with 39 percent in the data. We report additional details on the models and additional moment comparisons in Appendix B.

The reason that these fluctuations were missed by past research that inferred monthly earnings risk from annual data is that many of these changes are relatively transitory. To see this, we calculate the variance ratio:

$$\text{Variance Ratio: } \frac{\text{Var}(\log y_{t+k} - \log y_t)}{k \cdot \text{Var}(\log y_{t+1} - \log y_t)}$$


where y_t again denotes monthly earnings per paycheck. If earnings follow a random walk so that monthly changes are permanent, this ratio equals one for all k , since shocks accumulate over time. In contrast, if monthly earnings are white noise so that shocks are purely transitory, the variance ratio quickly converges to zero.

Figure 3 plots the empirical variance ratio. The ratio is much closer to white noise than to a

¹⁶Appendix C.2 provides a simple stylized example to show how combining jobs with two very different adjustment frequencies can lead our summary measures of volatility to increase rather than decrease.

¹⁷The amount of dampening is broadly consistent both with the smoothing implied by simulating income streams with independent within-household earnings shocks and with the amount of household smoothing documented in Norwegian administrative data by Andresen et al. (2025).

Figure 2: Earnings Risk in Monthly Data versus Models Calibrated to Annual Data



Notes: This plot compares the distribution of monthly earnings changes in PayrollCompany data to the distributions implied by benchmark models of earnings processes which are calibrated to annual data. KMV is Kaplan, Moll, and Violante (2018), KV is Kaplan and Violante (2022), MLM is Maxted, Laibson, and Moll (2025), and CHT is Crawley, Holm, and Tretvoll (2026). See Appendix B for additional details.


random walk, indicating that a large share of month-to-month earnings changes are fairly transitory. Estimating an AR(1) process to match this variance ratio yields a coefficient of 0.53, implying meaningful persistence at monthly horizons but little persistence at annual horizons ($0.53^{12} = 0.0005$). This helps explain why models disciplined only by annual data substantially understate the frequency of large monthly changes.

4 Fact 2: Earnings volatility is concentrated among hourly, low-wage workers

Our second fact is that monthly earnings volatility is a pervasive feature of hourly work and is therefore concentrated among low-wage workers, since hourly jobs are concentrated in the lower part of the wage distribution. For hourly workers, this volatility is driven by month-to-month changes in hours that pass through directly into pay. By contrast, salaried workers experience much less frequent earnings changes, and when they do experience a change, those changes largely reflect one-time payments such as bonuses, commissions, or raises. This distinction motivates our focus in the rest of the paper on the causes and consequences of earnings instability in hourly jobs.

Differences between hourly and salaried workers are apparent in the level, source, and symmetry of earnings volatility. Figure 4 shows that frequent earnings fluctuations are the norm for hourly workers but relatively rare for salaried workers. Table 1, Panel C shows that hourly workers

Figure 3: The Persistence of Monthly Earnings Changes



Notes: This figure plots the variance ratio $\frac{\text{Var}(\log y_{t+k} - \log y_t)}{k \text{ Var}(\log y_{t+1} - \log y_t)}$ for different values of k .


experience earnings changes in 91 percent of months—roughly 11 out of every 12 months. These changes are sizable: the median absolute change is 9 percent, and the 75th percentile is 21 percent. In contrast, salaried workers’ earnings are far more stable: the median absolute change is zero, and the 75th percentile is only 5 percent.

Differences in volatility between hourly and salaried workers are even larger when compared to liquidity. Using JPMCI data, we find that in over 47 percent of months, pseudo-hourly workers experience earnings changes exceeding half their median checking account balance, compared to 33 percent for pseudo-salaried workers. In 28 percent of months, the change exceeds their entire median balance, versus 18 percent for pseudo-salaried workers. Thus, pseudo-hourly workers not only face more volatility but they are also less buffered against it.

The reason hourly workers face substantial pay volatility is that they face large changes in hours that pass through directly into earnings. Indeed, Panel C of Table 1 shows that monthly fluctuations in hours are very similar to monthly fluctuations in total earnings. Changes in earnings for hourly workers are also roughly symmetric. To assess (a)symmetry in earnings changes, we measure changes relative to the median of the prior three months. This prevents a temporary spike in month t from registering as both a rise from $t-1$ to t and a drop from t to $t+1$. Table A-4 shows that hourly workers experience increases and decreases with roughly equal frequency.

Table 1, Panel D shows that volatility remains substantial for both full-time and prime-age hourly workers, indicating that instability is not simply due to weaker labor force attachment among hourly workers. Earnings volatility is also large when excluding workers who ever receive

Figure 4: Earnings Volatility for Hourly and Salaried Workers



Notes: This figure shows the within-job distribution of the change in earnings and wages from the prior month. “Total Earnings” is average pay per paycheck to abstract from pay schedule-driven fluctuations. “Base Wage” is the hourly wage for hourly workers and the per-period base salary for salaried workers.

overtime pay.¹⁸ We also find little evidence that these fluctuations are driven by unpaid leave, since unpaid leave is too infrequent and also would not generate the symmetry we find (see Appendix C.3 for details). Together, these results suggest that the hours changes we document do not simply reflect workers choosing when to work, a point to which we return later in the paper.

This volatility is also not well explained by annually recurring patterns in pay. Table A-5 shows that predictable annual patterns, whether based on a worker’s own earnings a year earlier or on regular firm-specific seasonal patterns, explain only a modest share of monthly pay changes for continuing hourly workers (see Appendix C.4 for details). These specifications capture recurring seasonal patterns—for example, a restaurant that is always busy in December. They do not capture irregular seasonal variation, such as an unusually harsh winter that reduces hours at a landscaping firm in a particular year. Such shocks would instead appear as firm-level fluctuations in labor demand, which we examine in Section 5.

Earnings changes for salaried workers, in contrast, are mostly increases off a stable base. Furthermore, these increases appear to largely reflect bonuses, performance pay, and commissions.¹⁹ Because these payments are infrequent but large, outlier-sensitive statistics such as the standard deviation, skewness, and kurtosis of monthly earnings changes are all similar or *higher* for salaried

¹⁸Overtime itself slightly increases volatility, but workers who ever receive overtime tend to be full-time and have slightly lower volatility overall. As a result, excluding these workers leaves summary statistics regarding volatility essentially unchanged.

¹⁹The less persistent variance ratio for salaried workers shown in Figure A-8 is consistent with earnings changes mostly reflecting one-time payments. Among salaried workers, volatility is highest for the highest earners, who are also the most likely to receive performance pay (Lemieux, MacLeod, and Parent 2009).

workers than for hourly workers. However, we do not view these statistics as the most informative measures of typical earnings risk. For hourly workers, earnings fluctuations are frequent and difficult to predict, whereas for salaried workers they are concentrated in infrequent, largely predictable bonus-like payments (see Appendix C.4 for details). We therefore focus on Median $|\Delta_{i,t}|$ as our primary measure of earnings instability. That said, our other findings are robust to using the standard deviation in place of the median.²⁰

Earnings instability declines with income, as shown in Table 2. This pattern is consistent with Morduch and Schneider (2017), who also find that income volatility is highest among low-earning workers. Table 2 shows that this gradient is explained in large part by the fact that lower-income workers are more likely to be paid hourly than higher-income workers. While volatility also declines with income among hourly workers, it remains large even for the highest-paid hourly workers. This pattern of declining volatility with income holds whether we measure volatility using the share of months with a pay change, the median absolute pay change, or the 75th percentile of the absolute pay change.²¹ Appendix Table A-6 uses the JPMCI data to show that earnings instability also declines with liquidity. Looking among households with similar income, those with lower liquidity tend to have greater earnings instability. For example, among low-income households, the bottom tercile of liquidity has a median monthly change of 9%, as compared to 6% for the top tercile. Because hourly workers also tend to have lower incomes and less liquidity, this pattern shows that the workers most exposed to pay volatility are also among those least able to absorb it.

Among hourly workers, volatility remains high across a wide range of worker characteristics. Table A-7 shows that volatility remains substantial even among older hourly workers. Table A-8 shows that volatility differs little by gender or by the presence of children in the household. This broad consistency across demographic groups is difficult to reconcile with explanations centered on childcare-related constraints. Such factors may matter for some workers, but they do not appear to explain why earnings volatility is so widespread among hourly workers.

This stark heterogeneity between hourly and salaried workers appears across occupations and industries, as shown in Appendix Table A-9. The distribution of contract type across occupations is bimodal: some occupations are predominantly hourly while other occupations are predominantly salaried. The top-left panel of Figure 5 shows that occupations that are usually salaried have median monthly earnings changes close to zero, while occupations that are usually hourly have median monthly earnings changes of around 10 percent. High-volatility hourly occupations include physical production (welder, mechanic), health (medical assistant), transportation (driver, warehouse), and hospitality (host, bartender, server). Occupations that are usually salaried are predominantly those with high cognitive task content, such as accountants, analysts, engineers, and teachers.

The top-right panel shows that this occupational gradient largely reflects contract type rather

²⁰Specifically, our findings on the relationships between firm-level and individual volatility (Section 5), volatility and spending (Section 6), and volatility and quits (Section 7) all hold using the standard deviation.

²¹In contrast, the standard deviation of income is highest for the top income quartile, reflecting the bonus-driven volatility discussed above. This finding is consistent with Farrell, Greig, and Yu (2019) since that paper focuses on the coefficient of variation. That paper also includes capital income and does not winsorize outliers, which further amplifies volatility for this highest-income group.

Table 2: Heterogeneity by Contract Type and Wage Level


Sample	<i>Job Characteristics</i>			<i>Earnings Changes</i>			
	Share salaried	Hours	Pay	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.
All							
Q1	15%	–	\$230	0.82	0.11	0.28	0.30
Q2	22%	–	\$569	0.79	0.05	0.14	0.20
Q3	44%	–	\$938	0.64	0.03	0.11	0.18
Q4	81%	–	\$2674	0.47	0.00	0.12	0.28
Hourly							
Q1	0%	26	\$288	0.93	0.11	0.25	0.29
Q2	0%	32	\$465	0.93	0.09	0.21	0.25
Q3	0%	35	\$655	0.92	0.08	0.19	0.24
Q4	0%	34	\$1081	0.86	0.07	0.18	0.25
Salaried							
Q1	100%	–	\$474	0.28	0.00	0.02	0.16
Q2	100%	–	\$986	0.33	0.00	0.03	0.15
Q3	100%	–	\$1525	0.34	0.00	0.04	0.18
Q4	100%	–	\$4264	0.41	0.00	0.17	0.35

Notes: In the top and bottom panels, worker-months are assigned to quartiles based on their pay per week. In the middle panel, worker-months are assigned to quartiles based on their hourly wage. Hours and pay are shown as weekly averages.

than occupation itself. When we restrict attention to hourly workers, volatility is high across occupations, including in occupations where most workers are salaried. For example, hourly accountants, analysts, engineers, and teachers have volatility much closer to that of other hourly workers than to the unconditional volatility of their occupation. This suggests that contract type accounts for much of the occupational variation in earnings volatility.

The same pattern appears across industries. The bottom-left panel shows that earnings volatility is highest in industries where hourly work is common, including those involving physical production, health, transportation, and hospitality. Volatility is lowest in industries such as Information, Finance and Insurance, and Professional and Scientific Services, where salaried work is more prevalent. The share of hourly work varies more continuously across industries than occupations, ranging from below 30 percent to above 80 percent, and median earnings volatility rises almost linearly with this hourly share. This cross-industry relationship again appears to be primarily driven by the prevalence of hourly work: the bottom-right panel shows that earnings volatility is high in every industry after restricting to hourly work. Thus, the evidence from both occupations and industries reinforces the conclusion that hourly work, rather than a narrow set of sectors or tasks, is a key determinant of earnings volatility.

Figure 5: Pay Volatility by Occupation and Industry



Notes: See Appendix Table A-9 for additional industry and occupation statistics.

However, even though the level of volatility is high for hourly workers in all occupations and industries, this does not mean that there are no differences in volatility across hourly jobs. Volatility for hourly workers is higher in customer-facing occupations such as servers, bartenders, and hosts, and lower in office-based hourly jobs such as administrative assistants. Looking across industries, volatility is lowest for hourly workers in Public Administration and highest for hourly workers in Accommodation and Food Services and Arts and Entertainment. Low volatility in Public Administration could reflect the fact that the public sector is highly unionized and collective bargaining imposes constraints on hours volatility which are not present in other sectors. High volatility for hourly workers occurs where firms are most likely to adjust labor input in response to fluctuating demand. We build on this observation systematically in the next section by studying whether firm-level labor-demand fluctuations generate worker-level hours instability.

5 Fact 3: High-frequency firm labor demand fluctuations are a key driver of hours instability

In this section, we provide evidence that firms play an important role in driving worker-level hours volatility using two distinct methodological approaches. First, we show that firms have substantial month-to-month fluctuations in total hours. We then allocate these firm-level fluctuations to individual workers, and show they are large enough to generate substantial individual volatility. Second, we use movers designs to compare the same workers across jobs at different firms. These designs show that firms have large effects on worker volatility and that most of these firm effects are associated with differences in total hours volatility.

5.1 Method 1: Accounting for Firm Total-Hours Changes

5.1.1 Magnitude of Firm Total-Hours Changes

We begin by showing that firms exhibit substantial monthly volatility in total hours, i.e., in the total number of hours worked by all of their employees in a given month. We interpret these fluctuations as reflecting shifts in labor demand by firms, rather than shifts in labor supply by workers. This interpretation is most plausible for larger firms, where idiosyncratic worker-level shocks such as illness are likely to wash out. For this reason, we restrict attention in this section to firms with a median size of at least 20 hourly workers.²² Results are also similar after controlling for industry-by-month effects.²³

Figure 6 shows that firms do not have stable total hours from month to month. The blue bars show total-hours growth, including both hours changes from continuing workers and from hires and

²²Appendix D.1 provides additional discussion of this threshold and related robustness. A threshold of 20 is chosen to reduce small sample issues while still retaining a sizable number of firms in our analysis, but results are similar when using a threshold of 50 or 100.

²³Alternative explanations based on labor supply would therefore require pervasive correlated labor supply shifts at particular firms, beyond those common to all firms in the same industry and month.

Figure 6: Firm Total-Hours Volatility



Notes: This shows the distribution of firm-month changes in total hours, for firms with a median of 20+ hourly workers. Total hours conditional on continued employment includes only hours of workers who are employed in the firm in the current and previous month while total hours includes all changes, including from hires and separations.

separations. It is not surprising that total firm hours fluctuate when employment is also changing, since firms grow and shrink over time. More notably, the orange bars show that firms also exhibit substantial monthly fluctuations in total hours worked by *continuing* workers. The median absolute percent change in total hours worked by continuing workers is 3.4 percent, and the 75th percentile is 7.3 percent. These fluctuations are large enough to drive substantial worker-level volatility, as we quantify below.

To our knowledge, this is the first paper to use high-frequency firm-level data to document labor demand shifts at this frequency. For the remainder of this section, we focus exclusively on this intensive-margin variation, since it informs our analysis of earnings volatility within employment relationships. Thus “total firm hours changes” will henceforth refer to changes in total hours worked by *continuing workers*.

What drives these changes in total firm hours? Table A-5 showed that stable annually recurring patterns explain little of monthly earnings changes for continuing workers. However, irregular fluctuations in demand — driven by forces like weather, customer traffic, project flow, event timing, or changes in broader economic conditions — may generate firm-wide changes in hours. Consistent with this interpretation, Farkas (2025) uses data from a scheduling software provider to show that firms directly pass weather-driven demand fluctuations onto workers’ schedules.

Several examples in our data illustrate the types of forces that generate hours fluctuations and how these forces differ across sectors. Although we do not disclose individual firm employment series, we observe restaurants where hours spike around large local events and fall sharply after disruptions such as severe weather events, hurricanes, or the onset of Covid-19. In manufacturing, we observe hours rising sharply at a public-safety equipment producer following the surge in demand

after the George Floyd protests in 2020. In finance, hours at a mortgage company track booms and busts in refinancing demand as interest rates rise and fall. In construction, hours increase around the timing of major project completions. These examples are illustrative rather than exhaustive. Even for these firms, many high-frequency fluctuations likely reflect more mundane and harder-to-observe forces that cannot be linked to observable events.

While the underlying shocks driving shifts in labor demand differ across sectors and across firms within sectors, the common lesson is that firms employing hourly workers can adjust labor input at high frequency. Hourly contracts provide firms with flexibility, and firms appear to use that flexibility to vary workers’ hours in response to changing conditions.

5.1.2 Allocating Total-Hours Changes to Workers

How much do these fluctuations in total hours matter for individual workers in the firm? One way to quantify the importance of fluctuations in total hours is with a reallocation statistic inspired by Davis and Haltiwanger (1992) that compares the change in total firm hours from one month to the next with the gross sum of all individual hours changes at that firm. This statistic captures the extent to which individual changes for workers i at firm j in month t are “necessary” to achieve the net change in firm hours observed in t :

$$\Delta H_{j,t}^{firm} \equiv \sum_{i \in j} \Delta h_{i,j,t} \quad (2)$$

$$\Delta H_{j,t}^{gross} \equiv \sum_{i \in j} |\Delta h_{i,j,t}| \quad (3)$$

$$\text{Necessary share}_{j,t} \equiv \frac{|\Delta H_{j,t}^{firm}|}{\Delta H_{j,t}^{gross}}. \quad (4)$$

Put differently, necessary changes are hours movements that are not offset by other workers in the firm changing hours in the opposite direction. Averaging across all firm-months, we find that 42 percent of worker hours changes are necessary to achieve the observed change in total firm hours. In an accounting sense, this implies that changes in total firm hours can “explain” a substantial share of changes in worker hours within firms. However, these comparisons of net to gross changes for the firm as a whole do not reveal the importance of changes in total firm hours for any individual worker.

To quantify the contribution of fluctuations in total firm hours ($\Delta H_{j,t}^{firm}$) to *individual* volatility, we need a counterfactual for what individual hours would be in the absence of changes in total firm hours. We proceed by allocating changes in total firm hours across individual workers using a rule which makes two key assumptions. First, the change in firm hours is allocated only among workers whose observed hours moved in the same direction as the firm. Second, the firm change is then allocated across these workers in proportion to their observed individual hours changes. The formula and additional details on computation are provided in Section D.2. We find the median

value of volatility for all workers is 9.0 percent while the median value of counterfactual volatility absent changes in total firm hours is 3.9 percent. This calculation suggests that almost 60 percent of individual volatility for the typical worker arises from fluctuations in total firm hours.

This suggests that firms play an important role in driving individual volatility. However, this result focuses on one channel, changes in total firm hours, and it relies on strong assumptions about how those changes are allocated across workers.

5.2 Method 2: Moves Between Firms

We now turn to a distinct source of evidence that firms matter for individual volatility based on worker moves between firms. This analysis addresses two questions. First, what is the overall causal effect of firms on individual volatility? Second, how much of that effect operates through firm total-hours volatility in particular? For the first question, we use a two-way fixed effects movers design to capture the role of all time-invariant firm characteristics. For the second, we use a similar movers design but replace firm fixed effects with a single firm characteristic: total-hours volatility. As in the rest of this section, this analysis of firms focuses on hourly workers.

5.2.1 Research Design

We use a fixed effects specification with movers following Abowd, Kramarz, and Margolis (1999, hereafter AKM) to study the effects of firms and workers on individual earnings volatility.²⁴ Since we only observe a small share of all firms in the economy in PayrollCompany data and our connected set of firms therefore is small, we follow Bonhomme, Lamadon, and Manresa (2019) by grouping firms and estimating the effects of moves between firm groups rather than between firms within groups. We estimate:

$$Vol_{i,j} = \mu + \alpha_i + \psi_{k(j)} + \varepsilon_{ij}, \text{ with normalization } \psi_1 = 0 \quad (5)$$


where $Vol_{i,j}$ is $Median(|\Delta_{i,j,t}|)$ for worker i 's job spell at firm j , μ is an overall intercept, $\psi_{k(j)}$ is the effect of firm j 's group (with group 1 normalized to 0), α_i is a worker fixed effect, and ε_{ij} is a residual "match-specific" effect.²⁵ We group firms into deciles of average worker-level volatility, so $\psi_{k(j)}$ measures the difference in volatility relative to the baseline decile 1.

A large literature emphasizes that causal interpretations of these estimates require strong assumptions. We estimate this specification with exactly two spells per worker, so the identification assumption can be stated in first differences: $\mathbb{E}(\Delta\varepsilon_{ij}|\psi_{k(j')} - \psi_{k(j)}) = 0$. This assumption would be violated if changes in workers' idiosyncratic volatility cause them to switch to firms with different

²⁴Section 5.1.2 focuses on the relationship between firm and individual hours, so it studies outcomes in hours space. Here, we are ultimately interested in how firms affect workers' earnings volatility through all channels. For example, earnings could potentially respond more than one-for-one to hours changes (e.g., through overtime premia) or firms might differ in the frequency of wage adjustment. In practice, however, our conclusions are very similar if we instead measure individual *hours* volatility.

²⁵Each (i, j) pair indexes a unique match. Since we estimate this specification at the spell-level, we suppress the time index t in the regression notation.

Figure 7: Firm Causal Effects by Decile Estimated from Movers Designs



Notes: This shows the results of estimating the two-way fixed effects design in Equation (5) with firm-decile fixed effects. The blue bar is the common component μ (i.e., average volatility in decile 1), the orange bar is the causal effect of the firm decile $\psi_{k(j)}$ relative to decile 1, and the green bar is the selection effect (i.e., average worker-level volatility α_i in $k(j)$ relative to decile 1). The variables on both axes are transformations of $Vol_{i,j} \equiv \text{Median}(|\Delta_{i,j,t}|)$. In defining the deciles, we define the volatility of firm j as the average of individual hourly workers' volatility $Vol_{i,j}$ weighted by each worker's spell length. The dependent variable $\overline{Vol}_j$ is the group-level average of $Vol_{i,j}$.

volatility or if the realization of the match-specific component affects which matches are actually formed. In Appendix D.3, we show that standard event-study diagnostic validations are satisfied: there are sharp changes in volatility around job transitions, no evidence of pre-trends, and increases for those moving to higher volatility firms are similar to decreases for those departing high volatility firms, suggesting that changes in worker volatility are not causing moves to different firms.²⁶

5.2.2 Main Results

Figure 7 visualizes the firm effects that arise from estimating Equation (5). For each firm decile, the orange bar shows ψ_k , the firm effect relative to decile 1. Assuming exogenous mobility (like AKM), this is the firm's causal effect on individual volatility. The green bar captures differential selection of workers (i.e., average α_i) across deciles relative to the average value of α_i in decile 1. The blue bar represents μ , the intercept, which captures volatility in decile 1. Since volatility captured by this intercept is common to all firms and workers, it cannot be attributed separately to either firm or worker characteristics.

Indeed, typical AKM regressions in wage space would not bother to report the value of this intercept. However, because we are particularly interested in understanding the average level of volatility, we include this component in the plot: adding the blue, orange and green bars then delivers the total observed level of volatility in each decile. Thus, the plot can be interpreted as a

²⁶Borovičková and Shimer (2024) argues that these diagnostics may fail to detect violations. In Appendix D.3 we argue that this is likely to be less of a concern in our context than in typical wage regressions because 1) volatility is not as easily observed as the wage at the time of potential match formation and 2) if high volatility matches are less likely to form, this would lead us to *understate* the importance of firms using our regression.

decomposition of the total observed volatility in each decile into a common intercept, firm-specific causal effects, and worker-selection effects.²⁷

Figure 7 highlights the importance of firm effects in explaining volatility for most deciles. For example, moving from a bottom-decile to a sixth-decile firm causes individual volatility to rise by 8.1 percentage points. One way to gauge the size of this effect is that the average level of individual volatility for a worker at a sixth-decile firm is 14.1 percent, which implies that around 60 percent of that individual volatility is caused by the firm. This is a lower bound, since some of the common component of volatility μ may also be driven by firms. An upper bound of 90 percent for the role of the firm is obtained if the entire common component is driven by firms.²⁸

Although these results show that firms explain a large share of volatility on average, this does not necessarily imply that firms are equally important for every worker. Figure A-6 shows that beyond the already-high volatility for the median worker, some workers experience much more extreme volatility. Estimates of Equation (5) can also inform whether firm effects explain a large share of the *variation* across workers in volatility (as opposed to whether they are big enough to explain the typical *level* of volatility we observe). Such comparisons are the focus of the existing wage AKM literature. Table A-10 shows that firm effects account for just over half of explained heterogeneity.²⁹ The workers with very high volatility play an outsized role in variance decompositions, and it is plausible that some of their largest monthly changes are worker-initiated, such as switches from full-time to part-time work.³⁰ Thus, while firm-driven fluctuations loom large for the typical worker, they may be less important for the most extreme cases of individual volatility.

5.2.3 Role of Total-Hours Volatility

The fixed effects design captures the combined effect of all time-invariant firm attributes on individual volatility. We now ask how much of that firm effect can be linked to one particular firm characteristic: total-hours volatility. This also provides a useful complement to the proportional allocation calculation in Section 5.1.2, which used a different approach to assess the role of firm-wide hours fluctuations.

Figure 8a shows that firms with more volatile total hours also have workers with more volatile pay. Figure 8b shows that this pattern also holds in a movers specification that relates changes in a worker’s pay volatility across job spells to differences in firm total-hours volatility, thereby absorbing time-invariant worker characteristics. The linear regression corresponding to Figure 8b yields a large and highly significant coefficient of 1.07. This means that moving from a firm with


²⁷By construction, the match-specific residual ε_{ij} averages to 0 in each firm group k so the decomposition in Figure 7 is exact.

²⁸This reflects the fact that the role of individuals as captured by selection effects is small.

²⁹In Appendix D.4, we also explore a number of additional robustness results that reinforce the conclusion that firm effects play an important causal role in driving worker volatility. Firm effects remain large when restricting to moves within industry or between firms with similar wages; they are not driven by any particular sector; they become slightly stronger when we expand the number of firm groups; they are very similar if we study hours volatility instead of pay volatility; and they remain large when restricting the sample to prime-age workers.

³⁰This does not imply they are desirable. Worker-driven hours changes might be welfare increasing (like choosing to work part-time after having children) or welfare decreasing (like negative health shocks).

Figure 8: Relationship Between Firm Total-Hours Volatility and Individual Worker Volatility



Notes: This shows a binscatter of the relationship between firm total-hours volatility and individual worker pay volatility. Firm volatility is measured as the firm's median absolute change in total hours. The left panel shows relationships in levels for all workers while the right panel measures the changes in individual pay volatility and firm total-hours volatility for workers who move between firms.

the median level of total-hours volatility to a firm with no total-hours volatility is associated with a 5.3 percentage point decline in individual pay volatility, which is close to the 5.1 percentage point decline in hours volatility from the counterfactual exercise in Section 5.1.2.³¹ Thus, two approaches with distinct identifying assumptions both suggest that firm-wide hours fluctuations are an important source of worker volatility.³²

In principle, firms can differ in both total-hours fluctuations and also in how they allocate those fluctuations across workers (e.g., their scheduling practices). It is therefore useful to compare the effects of moving between firms with different total-hours volatility to the firm-group movers specification above, which captures the combined effects of all time-invariant firm attributes on individual volatility. A one-standard-deviation increase in firm total-hours volatility predicts a 0.043 increase in individual pay volatility, compared with 0.050 for a one-standard-deviation increase in the estimated firm-group effect. Concretely, this means that differences in firm total-hours volatility generate nearly 90% as much cross-firm variation in pay volatility as the broader firm effects.

If differences in total-hours fluctuations explain most of the differences in pay volatility across firms, then this leaves little role for other cross-firm differences such as scheduling practices. This

³¹The median firm has total-hours volatility of 0.0494, which is moderately higher than the median *firm-month* absolute change of 0.034 reported in Section 5.1.1; $5.3\% = 1.07 \times 4.94\%$. The 5.1 percentage point decline in Section 5.1.2 comes from comparing median hours volatility in the main sample, 9.0 percent, to median hours volatility in the counterfactual without firm total-hours fluctuations, 3.9 percent. The movers sample has higher volatility than the main sample.

³²The movers design measures effects on earnings volatility, whereas Section 5.1.2 measures effects on hours volatility. Re-estimating the movers design using individual hours volatility yields an effect of 5.2 percentage points.

in turn has implications for policies that regulate workers' schedules, which we discuss in the conclusion. However, the comparison of these two coefficients should be interpreted cautiously for two reasons. First, a causal interpretation of the total-hours volatility regression requires not only the exogenous mobility assumption underlying Equation (5), but also the additional assumption that firm total-hours volatility is uncorrelated with other residual firm attributes that independently affect worker volatility. For example, firms with more volatile total hours may also have scheduling practices, staffing policies, or compensation structures that directly increase worker volatility. Second, some of the common component of volatility, μ , may itself reflect firm-driven channels that operate independently of total-hours fluctuations.

6 Fact 4: Higher earnings volatility is associated with higher spending volatility

We now turn to whether monthly earnings instability has meaningful consequences for workers.

We use the JPMCI data to investigate the link between earnings volatility and spending volatility. We measure the median absolute percent change in earnings $Vol_{i,j}^y$ and in non-durable spending $Vol_{i,j}^c$ within a job spell of worker i at firm j , applying similar filters and sample restrictions as in our baseline PayrollCompany analysis. All of our analysis focuses on pseudo-hourly workers since this is the group of workers for whom monthly earnings instability is most relevant. Most of our analysis focuses on households with one job, although we show robustness to multiple job-holding below. Appendix A.2 discusses these choices in more detail.

Much of the literature linking consumption to income focuses on the spending response to a particular income shock. Our goal is different: we ask more broadly whether unstable earnings translate into unstable spending.³³ Thus, we estimate the cross-sectional relationship between $Vol_{i,j}^c$ and $Vol_{i,j}^y$ using the following regression:

$$Vol_{i,j}^c = \alpha + \beta Vol_{i,j}^y + u_{i,j} \quad (6)$$

The estimate in column 1 of Table 3, $\hat{\beta} = 0.260$, implies that when the median monthly percent change in income increases by 10 percentage points, the median monthly percent change in spending increases by 2.6 percentage points. We interpret this magnitude below.

However, this cross-sectional relationship need not be causal. The estimate may be biased due to reverse causality (e.g., working more hours to fix a broken car), omitted variable bias (e.g., health shocks might reduce labor supply and spending, or wealth shocks might reduce labor supply but increase spending), or persistent heterogeneity (e.g., people with different risk preferences may have different volatility of both spending and income).

We argue for a causal link from income volatility to spending volatility using a combination

³³A month-to-month pass-through regression requires choices about timing, persistence, and the dynamic structure of earnings and spending responses and so is less well-suited for characterizing the broader relationship between earnings instability and spending instability.

Table 3: The Effect of Income Volatility on Consumption Volatility

	Dependent Variable: Med $ \%C $				
	OLS (1)	IV (2)	IV (3)	IV (4)	IV (5)
Med $ \%Y $	0.260*** (0.003)	0.256*** (0.010)	0.152*** (0.012)	0.346*** (0.012)	0.263*** (0.022)
Med $ \%Y $: High Checking				-0.162*** (0.018)	-0.127*** (0.016)
Implied WTP	8.77%	8.65%	5.05%	-	-
Group	All	All	Movers	All	All
Income Quintile FEs	-	-	-	-	✓
Observations	889,379	889,379	104,493	594,923	591,963

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Notes: This table estimates the effect of income volatility Med $|\%Y|$ on consumption volatility Med $|\%C|$. Each observation is one job spell. Standard errors are clustered by firm. The IV specifications instrument for individual income volatility using the average volatility of the worker's firm. Column (3) analyzes workers who switch between two firms and includes worker fixed effects. Columns (4) and (5) explore the role of liquidity, measured as the median checking account balance within a job spell. We define high vs. low checking as the top and bottom third of each household's median balance, dropping the middle tercile. Column (5) adds fixed effects for (average monthly) income quintiles and interactions of those with Med $|\%Y|$. The implied willingness to pay uses a formula based on Lucas (1987) to convert the regression coefficient into the amount that a worker with a relative risk aversion coefficient of 2 would be willing to pay to avoid the spending variance induced by median monthly income volatility. See Appendix E.1 for details. All specifications restrict to pseudo-hourly workers defined based on characteristics of the pay stream.

of instrumental variables, individual fixed effects, and heterogeneity analysis. In the first step of our identification strategy, we instrument for individual income volatility using the average income volatility at the worker's firm, $\overline{Vol}_{j(i)}^y$ and run the two-stage IV regression with second stage:

$$Vol_{i,j}^c = \alpha + \beta \widehat{Vol}_{i,j}^y + u_{i,j}, \quad (7)$$

where $\widehat{Vol}_{i,j}^y$ is the predicted value of individual income volatility from the first-stage regression:

$$Vol_{i,j}^y = \kappa + \pi \overline{Vol}_{j(i)}^y + e_{i,j}. \quad (8)$$

Since this IV regression only uses volatility that is common to all co-workers in the firm, it effectively asks whether workers at high-volatility firms have higher spending volatility than workers at low-volatility firms. Column 2 of Table 3 shows that this is indeed the case: the relationship between income volatility and spending volatility remains strong when we use firm-level volatility to remove idiosyncratic confounds. The validity of this IV specification requires that $\overline{Vol}_{j(i)}^y$ satisfies a relevance condition $\text{Cov}(Vol_{i,j}, \overline{Vol}_{j(i)}^y) \neq 0$ and exclusion restriction $\text{Cov}(\overline{Vol}_{j(i)}^y, u_{i,j}) = 0$. Given

our findings in fact 3 that firms explain a large share of individual volatility, it is unsurprising that this instrument is highly relevant, with first stage F-statistics above one thousand.

The exclusion restriction requires that firm-level volatility affect spending volatility only through its effect on individual income volatility. Because we do not observe exogenous variation in firm volatility, this condition could fail if firm volatility is correlated with other firm characteristics, such as management practices, industry, or unobserved amenities, that independently affect spending volatility. However, the most natural channel through which these firm characteristics would affect consumption volatility is through their effects on income volatility. A more plausible concern is sorting: workers with different preferences, such as different degrees of risk aversion, may systematically sort into firms with different volatility. In that case, firm-level volatility could be correlated with spending volatility through worker selection rather than through income volatility itself.

To address this concern, Column 3 of Table 3 again instruments for individual volatility using the average volatility of the firm, but now includes worker fixed effects, so α becomes α_i and κ becomes κ_i in Equations (7) and (8). These fixed effects absorb permanent worker characteristics that might differ across firms. We again find a significant positive relationship between spending volatility and income volatility. However, the coefficient is reduced relative to Column 2, suggesting that there was indeed some role for selection in those results.

Since this fixed effect IV specification relies only on within-worker variation, it is identified using workers switching between firms with different levels of average volatility. The identifying assumption now requires that *changes* in spending volatility across jobs arise only from *changes* in income volatility across jobs.³⁴ However, some life events might induce job changes at the same time that preferences for spending volatility change. For example, having a child might cause spending patterns to change without any change in income (e.g., spending on pediatricians, baby food, and diapers) and might also cause someone to change to a job with different hours and resulting income volatility. Two pieces of evidence suggest that such confounds do not drive the results.

First, in Appendix D.3, we explore difference-in-difference event study designs and show that there are sharp increases in spending volatility when workers move to jobs with higher income volatility, with no evidence of pre-trends. If life events were simultaneously shifting preferences and causing job transitions, this would likely manifest as trends in spending volatility prior to moves, since it is not clear why spending changes arising from such confounds should shift discretely at the time of a job transition.

Second, if life events were causing spending volatility and income volatility to shift in non-causal ways around job transitions, then excluding the months around job transitions from our measures of volatility should lead to different relationships. However, in Appendix D.3, we re-estimate our previous specifications by computing spell-level volatility excluding progressively larger

³⁴Unlike in AKM, endogenous job transitions do not violate the identification assumption, as long as spending volatility only changes because of the changes in income volatility. For example, suppose that an individual wants to work more volatile hours so that they can pick up their child from school and this requires switching to a more volatile job. As long as spending volatility only changes because of the change in income volatility and not for other reasons, we still recover the correct causal effect of income volatility on spending using this “endogenous” job transition.

“donuts” around the date of move. We find very similar relationships between income and spending volatility even when volatility is computed excluding several months before and after the move, again suggesting that endogenous job transitions coinciding with confounding shocks are not driving our results.


One way to gauge the magnitude of these effects is with a simple back-of-the-envelope calculation based on Lucas (1987). This procedure converts the regression coefficient into the amount a household with coefficient of relative risk aversion γ would be willing to pay to avoid the spending variance induced by monthly income volatility. Assuming a value of $\gamma = 2$, the regression estimates in columns 1 through 3 of Table 3 imply a willingness to pay for the median hourly worker of five to nine percent. This suggests that the spending volatility induced by income volatility is not just statistically significant, it is also economically meaningful. See Appendix E.1 for details.

The subsequent columns of Table 3 examine heterogeneity in ways that further support a causal interpretation of these estimates. We do so by re-estimating Equation (7) from the baseline specification and interacting income volatility with indicators for different worker characteristics. Column 4 shows that low-liquidity households exhibit a much stronger relationship between income volatility and spending volatility than high-liquidity households. This pattern is especially informative because liquidity has a clear economic link to the household’s ability to smooth income fluctuations. Column 5 shows that the same pattern remains even after conditioning on income, indicating that our liquidity measure is capturing something distinct from earnings level alone. Taken together, these results suggest that the consequences of earnings instability are particularly severe for financially fragile workers.³⁵

Table A-11 presents a series of additional robustness checks for our consumption analysis. Columns 2 and 3 split spending into work and non-work related categories following Ganong and Noel (2019). We find that non-work related spending volatility actually responds more strongly to income volatility than the rest of nondurable spending, allaying concerns that higher hours in a given month may mechanically raise work-related spending (e.g., commuting, meals away from home). Column 4 shows that we obtain similarly strong effects when we measure the volatility of total spending rather than non-durables. Column 5 shows that we find similar relationships when measuring both spending and income volatility at quarterly frequencies, which are more commonly analyzed in macroeconomic models. Column 6 shows that even when we estimate the effect of *monthly* income volatility on *quarterly* spending volatility, the results are very similar to our baseline. This shows that the high-frequency monthly earnings instability we document has meaningful consequences for spending at lower frequencies and is not merely high-frequency noise. Column 7 shows that, unsurprisingly, when we focus on two-job households, total income volatility matters more than individual job volatility for spending volatility. Column 8 shows that we find similar estimates when we limit the sample to larger firms.

³⁵While this section focuses on pseudo-hourly workers, the relationship between spending volatility and income volatility is weaker for pseudo-salaried workers. This is consistent with our earlier finding that salaried earnings changes are more predictable and less persistent.

Figure 9: Relationship between Separation Rates and Volatility



Notes: This figure shows a binscatter of the relationship between firm volatility and firm separation rates. The underlying unit of observation is a firm. Firm volatility is defined as the weighted mean of individual volatility at the firm, weighting by individual tenure. Individual volatility is defined as $Med|\Delta Y|$. The quit rate is computed for the subset of firms that record separation reasons for at least 50 percent of their workers.

7 Fact 5: Workers leave jobs with high earnings volatility

For our fifth and final fact, we return to the payroll data to show that hourly workers are more likely to quit high-volatility jobs, suggesting that volatility is a job disamenity.

We begin by documenting a strong positive relationship across firms between average individual volatility, $\overline{Vol}_j^y$, and average separation rates. We focus primarily on the total separation rate, which combines quits, layoffs, and firings, because this measure is available for all firms. Decomposing separations into quits versus other exits requires additional information that is only available for some firms in part of our sample period. Nevertheless, Figure 9 shows that for the firms that do record quits, the positive relationship between total separations and volatility is closely mirrored by the relationship between quits and volatility. For this reason, we interpret separations as largely reflecting worker choices. However, this firm-level relationship is not necessarily causal, since high-volatility firms may differ from low-volatility firms in other ways that affect quits. We return to this concern below, but first document the corresponding cross-sectional relationship at the individual level.

Table 4 presents individual-level results from a Cox proportional hazard model of separations on volatility: $H(t) = H_0(t) \times \exp[\beta_1(Med|\Delta Y_{ij}|) + \gamma'X_{ij}]$ where $H(t)$ is the hazard function at spell tenure t , $H_0(t)$ is the baseline hazard, and X_{ij} is a vector of potential controls.³⁶ The estimate

³⁶This hazard model allows for spell censoring, which occurs in the last month of our sample and also when a firm stops using PayrollCompany to process payroll. We cannot distinguish whether the latter reflects a switch in payroll providers or firm exit, so we treat these as censored spells. Results are similar if we instead re-estimate the

Table 4: The Effect of Income Volatility on Separation Rates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Med $ \% \Delta Y_{ij} $	2.85*** (0.057)	2.94*** (0.150)	2.85*** (0.204)	3.25*** (0.188)	3.09*** (0.291)	2.72*** (0.251)	2.72*** (0.238)
Implied WTP	9.6%	9.8%	9.6%	10.7%	10.2%	9.2%	9.2%
Sample	All	All	All	Movers	Full-time workers	Large firms	Exclude last quarter
Instrument?	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls?	No	No	Yes	Yes	Yes	Yes	Yes
No. Obs	213,252	213,252	213,252	52,125	135,034	176,937	174,446
No. Firms	16,725	16,725	16,725	10,989	13,965	7,469	14,630

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Notes: This table estimates the effect of worker-level earnings volatility $Med|\Delta Y_{ij}|$ on individual separation rates using Cox proportional hazard models. Standard errors are clustered by firm. Column 1 estimates this using each individual's volatility while Columns 2-7 instead instrument for individual volatility using the average individual volatility at the worker's firm (weighting individual $Med|\Delta|$ by individual worker tenure). Columns 3-7 include controls for firm average wages, firm average hours, industry fixed effects, gender, and worker age. Column 4 estimates a specification using a movers design with Γ distributed frailty random effects. Column 5 restricts to full-time workers; Column 6 restricts to firms with at least 20 workers. Column 7 re-runs results using an alternative volatility measure that can vary over time within worker and dropping the last quarter of each worker's own spell from predictions to try to rule out reverse causality. Implied willingness to pay (WTP) numbers are the percent change in wages that a worker with $Med|\Delta Y_{ij}| = 11.9\%$ would give up to move to a volatility of zero, computed using estimates of separation elasticities to wages from Lamadon, Mogstad, and Setzler (2022). Data is from PayrollCompany. See text for additional details.

of $\hat{\beta}_1$ from the simplest specification is shown in the first row of Column 1. It implies that moving from a job with constant earnings to one with median hourly worker volatility (11.9 percent) raises the separation hazard by 40 percent ($1.40 = \exp[2.85 \times 0.119]$).

To help interpret the magnitude of this empirical relationship, we compute a back-of-the-envelope willingness to pay to eliminate volatility following Gronberg and Reed (1994), who show that willingness to pay to avoid a disamenity can be inferred by comparing the elasticity of separations with respect to that disamenity to the corresponding elasticity with respect to wages. This approach, discussed in more detail in Appendix E.1, requires strong assumptions, but suggests that an hourly worker with median volatility would give up around 10 percent of wages to eliminate it.

This willingness-to-pay estimate should be interpreted cautiously. First, even if the underlying relationship is causal, it need not isolate the value workers place on fluctuating income alone. Rather, it may reflect the broader disamenity of job volatility, including both fluctuating hours and fluctuating income. This broader interpretation is consistent with survey evidence from Schneider and Harknett (2019), which suggests that workers often find unstable hours especially costly.

Second, the correlation between separations and volatility may not be causal. Subsequent columns of Table 4 therefore explore alternative specifications that address confounders that might spuriously generate this relationship. For example, serious illness could increase both hours volatil-

specifications in this table using linear probability models and treat censored spells as zeros.

ity and the probability of separating. Column 2 instruments for individual pay volatility using firm-level averages of individual pay volatility $\overline{Vol}_{j(i)}^y$, as in Section 6. The coefficient increases modestly.³⁷ This suggests that firm-wide volatility, which is less likely to be under workers' control, is more predictive of separations than idiosyncratic volatility.

In column 3, we add several firm- and worker-level controls to the IV specification (firm average wages and hours, industry fixed effects, and the worker's gender and age at job start). Unsurprisingly, adding these observable controls attenuates the relationship between volatility and separations somewhat, but the estimated effect remains quantitatively large. A causal interpretation of this specification requires a stronger exclusion restriction than in Section 6, which is that conditional on these controls, $\overline{Vol}_{j(i)}^y$ must affect separation rates only through its effect on workers' income volatility. This assumption is less innocuous in the separations setting because firms likely differ along unobserved dimensions that are correlated both with monthly income volatility and with separation rates directly. For example, a bad manager might both generate more volatile earnings and independently increase the probability that workers separate.

Although we cannot fully eliminate concerns about unobserved firm characteristics, we can partially address them by comparing hourly and salaried workers within the same firm. This comparison holds fixed firm attributes common to both groups, such as location, corporate structure, and broad management practices. Specifically, we examine the relationship between the separation rates of *salaried* workers and $\overline{Vol}_{j(i)}^y$ —the average volatility of *hourly* workers—at the same firms. Salaried and hourly workers at the same firm are exposed to some of the same unobserved firm characteristics, but salaried workers do not experience the same degree of monthly income volatility as hourly workers. If firm-level confounds were driving our results entirely, we would expect firm volatility to predict separation rates similarly for the two groups. Table A-12 shows that this is not the case.³⁸ Firm volatility increases hourly workers' separation rates by nearly three times as much as the separation rates of salaried workers at the same firm.³⁹ Any remaining confounding job characteristics would have to differ systematically between hourly and salaried workers within the same firm.

While our ability to control for unobserved firm characteristics is limited, we can do more to address worker heterogeneity and sorting. In particular, firms employing different types of workers may exhibit both greater volatility and higher separation rates even in the absence of a causal link between the two. To address this concern, column 4 estimates a shared-frailty Cox model among workers who move across firms, allowing for differences across workers in baseline separation propensities.⁴⁰ This specification yields estimates that are slightly larger, but statistically

³⁷Because the Cox model is nonlinear, we use a control function approach: we estimate a linear first-stage, and then include both predicted values and residuals in the non-linear second stage.

³⁸Here we are interested in whether *firm* volatility, $\overline{Vol}_{j(i)}^y$, has a different relationship with separation rates for hourly and salaried workers, so we estimate the reduced form rather than the IV.

³⁹The fact that $\overline{Vol}_{j(i)}^y$ still has some relationship with salaried workers' separation rates could reflect its correlation with salaried workers' own earnings volatility, the effects of *hours* instability holding earnings fixed, or remaining unobserved firm-level confounds.

⁴⁰Fixed effects cannot be implemented in the Cox hazard model, so we model worker heterogeneity as shared frailty random effects with a Γ distribution. We obtain similar results in linear probability models with worker fixed effects.

indistinguishable from the IV estimates with controls in column 3.

Since the IV with controls specification in Column 3 is simpler, has larger sample sizes and is more representative, this is our preferred specification and we use it as the baseline for remaining robustness results. Column 5 restricts to full-time workers to show that the patterns are not driven by part-time workers who may both be less attached to the labor force and have more volatile hours. Showing that the results hold in the sample of full-time workers is important because Dube, Naidu, and Reich (2022) find that among low-wage hourly workers the most highly sought-after amenity is a full-time position. Finally, we show that these patterns are also not the result of reverse causality (i.e., worker separations may cause an increase in pay volatility for coworkers who remain at the firm) by restricting only to large firms where a single worker separating has a smaller effect on average volatility (Column 6) and excluding the final quarter of each worker’s spell (Column 7).⁴¹

Across all specifications, the evidence consistently shows that workers in volatile jobs separate at higher rates. Using specifications that include firm-level instruments and controls for worker heterogeneity, we are able to address many important confounds. The comparison of hourly and salaried workers within firms suggests that the relevant disamenities are not simply firm-wide attributes shared by all workers, but are instead disproportionately concentrated in the hourly jobs within those firms. However, we are unable to separate pay volatility from other unobserved aspects of those hourly jobs that may be correlated with volatility, such as job features or managerial practices that differentially affect hourly workers relative to salaried workers within the same firm. We therefore interpret these results as showing that monthly pay volatility is part of a broader bundle of undesirable job attributes that characterize many hourly jobs.

8 Conclusion

In this paper, we document that U.S. workers face substantial month-to-month earnings instability that is largely invisible in annual data. This instability is concentrated among lower-income hourly workers and is driven in important part by firm-level labor-demand variation. This earnings instability in turn passes through into household spending instability and is associated with higher worker quit rates. Together, these findings reinforce and extend earlier research suggesting that low-wage, financially fragile workers may be particularly exposed to fluctuations in labor demand (Doeringer and Piore 1971; Rebitzer and Taylor 1991; Morduch and Schneider 2017).

Our paper raises several questions for future research. Although we find that fluctuations in firm total hours are an important driver of pay instability, our payroll data do not reveal *why* firms vary their hours so much from month to month. Our payroll data reveal the importance of the firm demand channel but cannot speak to its underlying drivers. One question we are unable to fully answer is “why does firm total labor demand vary so much from month-to-month?”

⁴¹Similar to our other analysis of volatility dynamics, this requires moving to a pooled volatility measure that has some time dimension. Since this pooled volatility measure is a slightly different instrument, we have also re-run regressions with this pooled estimate but *without* dropping the last quarter, and this produces a statistically indistinguishable coefficient of 2.79 (s.e. 0.23).

In some industries, labor generates storable output (e.g., a nail polish manufacturer), so total-hours volatility may reflect production planning, inventory adjustment, or new information about future product demand. In other industries, labor must coincide more closely with sales (e.g., a nail salon), so customer-demand fluctuations translate more directly into contemporaneous hours volatility. Understanding why firms move their total hours requires detailed information on firm sales, production, and inventories.

Second, firms may differ not only in total-hours fluctuations, but also in their scheduling practices. Our results in Section 5.2.3 suggest scheduling practices may *not* be an important source of differences in pay volatility. However, this evidence does not imply that scheduling practices are unimportant for workers. Workers may care directly about the stability and predictability of their hours, and contracts that stabilize pay need not stabilize hours. Understanding the contribution of scheduling practices to pay and hours volatility requires detailed data on workers' schedules beyond what is used in this paper.

Answering these questions matters for active policy debates about labor-market regulation. European economies rely less heavily on hourly compensation and use institutions such as temporary contracts, short-time work, or working-time accounts to manage fluctuations in labor demand. If firms in these economies face similar underlying shocks to desired labor input as in the United States, these institutions may shift more of the resulting risk away from workers and toward firms, while also raising firms' costs of adjustment (Schoefer 2025). Within the United States, several cities have adopted regulations aimed at increasing schedule predictability, reducing hours volatility, or reducing pay volatility and their consequences will depend on the answers to these same questions.⁴² Although we estimate negative consequences of instability for workers, our estimates capture individual workers' preferences in partial equilibrium, holding fixed the set of jobs. A central question for future work is whether, in general equilibrium, eliminating instability (e.g., through regulation) might lead to the existence of fewer low-wage and hourly jobs.

Our findings also have implications for structural modeling. Models of income dynamics calibrated only to annual data do not accurately capture high-frequency income risk. Because monthly fluctuations are sizable but often transitory, they are difficult to infer from annual moments alone, even though they matter for workers' financial behavior. Future income-process models should therefore be disciplined directly by the high-frequency moments we document here, rather than relying only on annual data and parametric restrictions to infer within-year dynamics. Our results also point to an important source of heterogeneity that has received little attention in such models: the sharp difference in earnings instability between hourly and salaried work. In turn, consumption models that take these income processes as inputs should aim to match not only average spending responses to specific income shocks, but also the broader relationship between earnings instability, liquidity, and spending volatility that we document in the data.

⁴²For example, San Francisco, San Jose, Berkeley, Los Angeles, Seattle, New York City, Chicago, and Philadelphia have passed ordinances about this. See Pickens and Sojourner (2026) for a discussion of the impacts of one such law. Some states, such as Arkansas, Florida, Georgia, Kansas, and Ohio, have passed statutes preventing local governments from regulating schedules.

References

- Abowd, John M. and Kevin L. McKinney.** 2024. “Mixed-Effects Methods for Search and Matching Research.” Revue économique, 75(1): 55–72.
- Abowd, John M, Francis Kramarz, and David N Margolis.** 1999. “High Wage Workers and High Wage Firms.” Econometrica, 67(2): 251–333.
- Andresen, Martin E, Andreas R Kostøl, Ross T Milton, Corina Mommaerts, and Luisa Wallossek.** 2025. “Monthly Earnings Volatility and Household Pooling.” Working Paper 34563. National Bureau of Economic Research.
- Arellano, Manuel, Orazio Attanasio, Margherita Borella, Mariacristina De Nardi, and Gonzalo Paz-Pardo.** 2026. “Subjective Earnings and Employment Dynamics.” Working Paper 35027. National Bureau of Economic Research.
- Bania, Neil and Laura Leete.** 2009. “Monthly household income volatility in the U.S., 1991/92 vs. 2002/03.” Economics Bulletin, 29(3): 2100–2112.
- Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu.** 2022. “Monopsony in Movers.” The Journal of Human Resources, 57(S): S50–S86.
- Bergman, Alon, Guy David, and Hummy Song.** 2023. ““I Quit”: Schedule Volatility as a Driver of Voluntary Employee Turnover.” Manufacturing & Service Operations Management, 25(4): 1416–1435.
- Blundell, Richard, Christopher R. Bollinger, Charles Hokayem, and James P. Ziliak.** 2025. “Interpreting Cohort Profiles of Life Cycle Earnings Volatility.” Journal of Labor Economics, 43(S1): 55–82.
- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa.** 2019. “A Distributional Framework for Matched Employer Employee Data.” Econometrica, 87(3): 699–739.
- Borovičková, Katarína and Robert Shimer.** 2024. “Assortative Matching and Wages: The Role of Selection.” Working Paper 33184. National Bureau of Economic Research.
- Brewer, Mike, Nye Cominetti, and Stephen P. Jenkins.** 2025. “What Do We Know About Income and Earnings Volatility?” Review of Income and Wealth, 71(2): e70013.
- Caplin, Andrew, Victoria Gregory, Eungik Lee, Søren Leth-Petersen, and Johan Sæverud.** 2023. “Subjective Earnings Risk.” Working Paper 31019. National Bureau of Economic Research.
- Card, David, Ana Rute Cardoso, and Patrick Kline.** 2016. “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women.” Quarterly Journal of Economics, 131(2): 633–686.

- Chetty, Raj and Adam Szeidl.** 2007. “Consumption Commitments and Risk Preferences.” Quarterly Journal of Economics, 122(2): 831–877.
- Crawley, Edmund, Martin B. Holm, and Håkon Tretvoll.** 2026. “A Parsimonious Model of Idiosyncratic Income.” International Economic Review, 67(2): 533–548.
- Davis, Steven J. and John Haltiwanger.** 1992. “Gross Job Creation, Gross Job Destruction, and Employment Reallocation.” Quarterly Journal of Economics, 107(3): 819–863.
- Doeringer, Peter B. and Michael J. Piore.** 1971. Internal Labor Markets and Manpower Analysis. Lexington, MA: Heath Lexington Books.
- Druedahl, Jeppe, Michael Graber, and Thomas H. Jørgensen.** 2023. “High Frequency Income Dynamics.” Working Paper.
- Dube, Arindrajit, Suresh Naidu, and Adam D Reich.** 2022. “Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers.” Working Paper 30441. National Bureau of Economic Research.
- Ehrenreich, Barbara.** 2001. Nickel and Dimed: On (Not) Getting By in America. New York: Metropolitan Books.
- Farkas, Hannah.** 2025. “The Economic Incidence of Schedule Unpredictability in Hourly Work.” Working paper.
- Farrell, Diana and Fiona Greig.** 2015. “Weathering Volatility.” JPMorgan Chase Institute.
- Farrell, Diana, Fiona Greig, and Chenxi Yu.** 2019. “Weathering Volatility 2.0.” JPMorgan Chase Institute.
- Ganong, Peter and Pascal Noel.** 2019. “Consumer Spending during Unemployment: Positive and Normative Implications.” American Economic Review, 109(7): 2383–2424.
- Ganong, Peter, Damon Jones, Pascal J. Noel, Fiona E. Greig, Diana Farrell, and Chris Wheat.** 2025. “Liquid Wealth and Consumption Smoothing of Typical Labor Income Shocks.” Working Paper.
- Gottschalk, Peter and Robert Moffitt.** 1994. “The Growth of Earnings Instability in the U.S. Labor Market.” Brookings Papers on Economic Activity, 1994(2): 217–272.
- Grigsby, John, Erik Hurst, and Ahu Yildirmaz.** 2021. “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data.” American Economic Review, 111(2): 428–71.
- Gronberg, Timothy and W. Reed.** 1994. “Estimating Workers’ Marginal Willingness to Pay for Job Attributes Using Duration Data.” Journal of Human Resources, 29(3): 911–931.

- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi.** 2005. “Insurance within the Firm.” Journal of Political Economy, 113(5): 1054–1087.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song.** 2014. “The Nature of Countercyclical Income Risk.” Journal of Political Economy, 122(3): 621–660.
- Haider, Steven J.** 2001. “Earnings instability and earnings inequality of males in the United States: 1967–1991.” Journal of Labor Economics, 19(4): 799–836.
- Hannagan, Anthony and Jonathan Morduch.** 2015. “Income Gains and Month-to-Month Income Volatility: Household Evidence from the US Financial Diaries.” Working Paper.
- Hart, Oliver and Bengt Holmström.** 1987. “The Theory of Contracts.” In Advances in Economic Theory: Fifth World Congress. Econometric Society Monographs, ed. Truman Fasset Bewley, 71–156, Cambridge: Cambridge University Press.
- Humlum, Anders, Mette Rasmussen, and Evan K. Rose.** 2025. “Firm Premia and Match Effects in Pay vs. Amenities.” Working Paper 33884. National Bureau of Economic Research.
- Johnson, David S, Jonathan A Parker, and Nicholas S Souleles.** 2006. “Household expenditure and the income tax rebates of 2001.” American Economic Review, 96(5): 1589–1610.
- Kaplan, Greg and Giovanni L. Violante.** 2022. “The Marginal Propensity to Consume in Heterogeneous Agent Models.” Annual Review of Economics, 14(1): 747–775.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante.** 2018. “Monetary Policy According to HANK.” American Economic Review, 108(3): 697–743.
- Kesavan, Saravanan and Camelia M. Kuhnen.** 2017. “Demand Fluctuations, Precarious Incomes, and Employee Turnover.” Working Paper.
- Kline, Patrick.** 2024. “Firm Wage Effects.” In Handbook of Labor Economics. Vol. 5, ed. Christian Dustmann and Thomas Lemieux, 115–181, Amsterdam: Elsevier.
- Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten.** 2020. “Leave-Out Estimation of Variance Components.” Econometrica, 88(5): 1859–1898.
- Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury.** 2022. “How Reliable are Administrative Reports of Paid Work Hours?” Labour Economics, 75: 102131.
- Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A. Woodbury.** 2026. “Work Hours Mismatch.” Econometrica, 94(3): 991–1025.
- Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler.** 2022. “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market.” American Economic Review, 112(1): 169–212.

- Lambert, Susan J., Julia R. Henly, and Jaeseung Kim.** 2019. "Precarious Work Schedules as a Source of Economic Insecurity and Institutional Distrust." RSF: The Russell Sage Foundation Journal of the Social Sciences, 5(4): 218–257.
- Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent.** 2009. "Performance Pay and Wage Inequality." Quarterly Journal of Economics, 124(1): 1–49.
- Lucas, Robert E. Jr.** 1987. Models of Business Cycles. Oxford: Wiley-Blackwell.
- MacLeod, W. Bentley and Daniel Parent.** 1999. "Job Characteristics and the Form of Compensation." Research in Labor Economics, 18: 177–242.
- Maestas, Nicole, Kathleen J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger.** 2023. "The Value of Working Conditions in the United States and Implications for the Structure of Wages." American Economic Review, 113(7): 2007–2047.
- Mas, Alexandre.** 2025. "Non-Wage Amenities." Working Paper 33643. National Bureau of Economic Research.
- Mas, Alexandre and Amanda Pallais.** 2017. "Valuing Alternative Work Arrangements." American Economic Review, 107(12): 3722–3759.
- Maxted, Peter, David Laibson, and Benjamin Moll.** 2025. "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy." Quarterly Journal of Economics, 140(1): 691–743.
- Meghir, Costas and Luigi Pistaferri.** 2004. "Income Variance Dynamics and Heterogeneity." Econometrica, 72(1): 1–32.
- Moffitt, Robert A. and Peter Gottschalk.** 2012. "Trends in the Transitory Variance of Male Earnings." Journal of Human Resources, 47(1): 204–236.
- Moffitt, Robert and Sisi Zhang.** 2018. "Income Volatility and the PSID: Past Research and New Results." AEA Papers and Proceedings, 108: 277–280.
- Moffitt, Robert, John Abowd, Christopher Bollinger, Michael Carr, Charles Hokayem, Kevin McKinney, Emily Wiemers, Sisi Zhang, and James Ziliak.** 2022. "Reconciling Trends in U.S. Male Earnings Volatility: Results from Survey and Administrative Data." Journal of Business & Economic Statistics, 41(1): 1–11.
- Morduch, Jonathan and Rachel Schneider.** 2017. The Financial Diaries: How American Families Cope in a World of Uncertainty. Princeton University Press.
- Myers, Charles A. and George P. Shultz.** 1951. The Dynamics of a Labor Market: A Study of the Impact of Employment Changes on Labor Mobility, Job Satisfaction, and Company and Union Policies. New York: Prentice-Hall.

- Pickens, Joseph and Aaron Sojourner.** 2026. “Effects of Fair Workweek Laws on Labor Market Outcomes.” Industrial Relations: A Journal of Economy and Society, 65(2): 132–155.
- Prendergast, Canice.** 2002. “The Tenuous Trade-off between Risk and Incentives.” Journal of Political Economy, 110(5): 1071–1102.
- Pruitt, Seth and Nicholas Turner.** 2020. “Earnings Risk in the Household: Evidence from Millions of US Tax Returns.” American Economic Review: Insights, 2(2): 237–254.
- Rebitzer, James B. and Lowell J. Taylor.** 1991. “A Model of Dual Labor Markets When Product Demand Is Uncertain.” Quarterly Journal of Economics, 106(4): 1373–1383.
- Reynolds, Lloyd G.** 1951. The Structure of Labor Markets: Wages and Labor Mobility in Theory and Practice. New York: Harper.
- Schneider, Daniel and Kristen Harknett.** 2019. “Consequences of Routine Work-Schedule Instability for Worker Health and Well-Being.” American Sociological Review, 84(1): 82–114.
- Schoefer, Benjamin.** 2025. “Eurosclerosis at 40: Labor Market Institutions, Dynamism, and European Competitiveness.” Working Paper 33975. National Bureau of Economic Research.
- Sorkin, Isaac.** 2018. “Ranking Firms Using Revealed Preference.” Quarterly Journal of Economics, 133(3): 1331–1393.
- Zhang, C. Yiwei, Abigail B. Sussman, Nathan Wang-Ly, and Jennifer K. Lyu.** 2022. “How consumers budget.” Journal of Economic Behavior & Organization, 204: 69–88.
- Ziliak, James P., Bradley Hardy, and Christopher Bollinger.** 2011. “Earnings volatility in America: Evidence from matched CPS.” Labour Economics, 18(6): 742–754. European Association of Labour Economists.

Online Appendix to “Earnings Instability”

Peter Ganong, Pascal Noel, Christina Patterson, Joseph Vavra, and

Alexander Weinberg

Contents

A Data Appendix	1
A.1 PayrollCompany Data	1
A.2 Chase Data	3
B Income Model Appendix	4
C Appendix with additional empirical results	6
C.1 Comparing Aggregate Seasonality in PayrollCompany to Benchmarks	6
C.2 Aggregating Volatility From Jobs to Households	6
C.3 Unpaid Leave	7
C.4 Predictable Annual Variation and Earnings Instability	8
D Appendix on the Role of Firms in Earnings Instability	10
D.1 Interpreting Fluctuations in Total Firm Hours	10
D.2 Proportional Allocation of Total Firm Hours Movements to Individual Workers . . .	11
D.3 Identification of Firm Effects on Income Volatility	13
D.4 Variance Decompositions	16
E Appendix on Spending Volatility	18
E.1 Calculating the Willingness to Pay to Eliminate Volatility	20
F Additional Appendix Figures and Tables	22

A Data Appendix

A.1 PayrollCompany Data

This appendix provides some additional sample restrictions and details of data construction for the PayrollCompany data described in Section 2.

In addition to the main data cleaning steps described in Section 2, we impose the following sample restrictions:

1. We exclude the 0.07 percent of paychecks where the sum of pay items in any namecode (e.g., base pay, bonuses, or overtime) is negative. These negative payments can occur if there are payroll mistakes which are corrected in a subsequent check.
2. We exclude 1.40 percent of payees classified as owners, since their pay primarily reflects their own decisions about when to withdraw funds from the firm.
3. We exclude 0.29 percent of workers whose earnings always fall below the federal minimum wage, \$7.25. Our measure of earnings includes non-cash tips and commissions and so these records likely reflect under-reporting of true earnings from missing cash payments.
4. We exclude the 1.06 percent of workers who are ever paid more than 400 hours in a single month because it is possible they are being paid in one month for work that they did in more than one month.

Some additional restrictions are imposed when studying the effects of individual volatility and the relationship between individual and firm-level volatility. We require that an individual be observed for at least four months to construct an individual volatility statistic. When we construct firm-level volatility, we require at least four months for each firm and separately require twenty-four hourly worker-month observations so that volatility statistics are not driven by a single worker. Since we are interested in demographic controls in many of our individual volatility restrictions, we include only individuals with non-missing information on age and gender. We also drop any firms with a median wage across all workers greater than \$100 (and also note that our worker-level filter also means we drop any firms with a median wage less than \$7.25). In our analysis of separations, we want to control for additional firm characteristics and so we further require non-missing industry information (we also require information on non-missing and reliable firm-wage and hours information, but these are already implied by the individual filters discussed above).

In Section 5.1.1 we calculate firm total-hours volatility with a sample that is slightly different from other firm total-hours volatility analyses. We select firms whose median monthly hourly employment is at least 20 workers – in line with other analyses – but instead use the 1% check-level sample within those firms to calculate volatility. This allows us to calculate volatility on the extensive margin.

In Section 4 we discuss heterogeneity in earnings instability by various worker characteristics. Some of these characteristics are available only for some observations. For example, bonuses are

incompletely reported. Evidence of this comes from the fact that both base pay and bonuses for salaried workers surge by roughly equal dollar amounts in December. As a second example, data on occupations is limited. We obtain information on workers’ job titles from a subset of firms covered by PayrollCompany, which reports free-text job titles as entered by employers. We then subset to the 100 most common free text titles in the data and require that 10 or more firms use each of these titles. Finally, we manually combine similar job titles (e.g., “sales representative” and “sales”) and remove titles that have an ambiguous interpretation (e.g., “laborer”).

When studying separations, we define a separation as the last pay period of an individual worker spell, with one exception: we treat the last period that a firm is observed in the data as a censored observation rather than a separation. That is, when a firm leaves the data set (which happens either in the last month of the data or if the firm switches payroll providers at some earlier date), this will also be the last observed month of pay for all of the firm’s workers, but we do not classify these months as worker separations. Since some of the specifications in Table 4 are computationally burdensome, for that table we draw a further 15 percent subsample of our primary sample.

For Figure 9, separation rates are observed for all firms since we can measure the end of pay streams for all workers. However, information on the reason for separation (e.g., quit vs. layoff vs. fire) is optional information that the firm does not have to record. This means that not all firms report this information and even for firms that do, they do not necessarily report it for all workers. Thus, when studying the relationship between quits and separations, we focus only on firms that report separation reasons for at least 50 percent of their separations. We also impose consistency between total separation rates and the sub-components using a proportional rescaling: among those separations with listed reasons, we measure the observed share of quits, layoffs and fires and we then multiply these shares times the total separation rate to construct the total quit rate. This means that if firms record separation reasons for all workers, the observed quit rate and the total quit rate are identical. When only some separation reasons are recorded, this procedure imputes quit rates for those with no recorded reason in the same proportion as quit rates for those with recorded reasons.

To ensure that outliers do not drive our results, we apply winsorization at various points in our analysis. We winsorize the lowest and highest 2.5 percent of earnings changes from one month to the next (i.e., we winsorize the 5 percent most extreme observations) to ensure that no single huge earnings changes drive any of our results. When we construct individual volatility: $Vol_i = Median|\% \Delta_{i,t}|$, we further winsorize the 5 percent largest values of Vol_i to ensure that no individual with extreme volatility drives our results.⁴³ We also winsorize firm-level volatility with the same 5 percent cutoff. However, one of the reasons that our preferred measure of volatility is the median change is to limit the role of outliers in our volatility measures. This means that in practice, this winsorization makes little difference for our results. In contrast, other measures like the standard deviation or other higher moments that are more sensitive to outliers do depend on winsorization

⁴³We winsorize the top 5 percent of volatility rather than imposing a symmetric winsorization because individual volatility near 0 is not an outlier.

choices (see Table A-2) which is why we do not focus on these moments.

A.2 Chase Data

This data appendix provides some additional detail on the bank account data described in Section 2.2 relevant for the analysis in Section 6. Our data analysis in terms of spending and liquidity definitions and samples follows that in Ganong et al. (2025), so we describe here only choices that are unique to our analysis. The unit of observation that we consider is a worker job spell, which we define as the string of contiguous months with direct deposits from the same employer. We restrict to job spells with at least 4 full months of employment. For Figure A-7 we include accounts with multiple jobs, but since we want to look at the effect of job transitions on spending, for most of the analysis in Section 6 we restrict to job spells where there are only direct deposits from a single job into the checking account over the entire course of the job spell.⁴⁴

Our firm identifier in the Chase data is encoded from information in these worker direct deposits. This means that changes in payroll processing can sometimes lead to changes in firm names and thus imputed firm identities. To identify spurious moves, we look for instances where a large share of workers move from the same origin firm to the same destination firm and then exclude these likely spurious moves from our analysis. We use a threshold for this share which varies with firm size, since at smaller firms even a small number of workers actually moving from the same origin to destination firm might lead to a large share of such moves. Concretely, we label moves from an origin firm to a destination firm as spurious when that origin-destination flow accounts for an unusually large share of all observed movers from that origin firm. Specifically, we drop an origin-destination cell if it accounts for at least 60 percent of observed movers from an origin with 5–9 movers, at least 50 percent from an origin with 10–19 movers, or at least 40 percent from an origin with 20 or more movers. These thresholds were chosen based on manual inspections of the text strings for various workers for a subset of movers, but results are similar if we use a common threshold or are more aggressive about removing potentially spurious moves.

Finally, we apply the same winsorization choices (± 2.5 percent symmetric on monthly changes and 5 percent asymmetric on individual and firm volatility) to income and spending as in the PayrollCompany data.

⁴⁴For the analysis of households we include accounts with multiple job spells at the same time.

B Income Model Appendix

The level of earnings risk faced over time is a key determinant of household savings decisions in modern consumption-savings models. Business cycle versions of these models are typically solved at sub-annual frequencies, which requires taking a stand on the level of within-year earnings risk.⁴⁵ However, the fact that panel data on earnings is typically only available annually means that the earnings process relevant at high frequencies is not observed directly. Instead, the literature has proceeded by specifying a parametric process for high-frequency earnings and then estimating the parameters of this process to match annual income moments from various sources.

For example, Kaplan, Moll, and Violante (2018, hereafter KMV) specifies a continuous time earnings process with two independent earnings shocks and shows that the parameters of this parametric model can be identified using the kurtosis and other higher moments of annual income changes in administrative social security data.⁴⁶

In Figure 2, we compare the monthly distribution of changes implied by this and several other income models to the data. Table A-13 provides various other moments. Since the original versions of these models are often specified at time horizons other than months, some small adjustments need to be made to make comparisons to monthly moments.

Since KMV is a continuous time model, calculating monthly moments is straightforward. We also simulate a monthly version of the discrete-time income process from Kaplan and Violante (2022, hereafter KV) and the continuous-time models in Maxted, Laibson, and Moll (2025, hereafter MLM) and Crawley, Holm, and Tretvoll (2026, hereafter CHT). The KMV and KV models include a two-shock process that arrives with Poisson probability. The KV model is quarterly. We translate this to a monthly model by rescaling the arrival rate of the two shocks and by assuming that the “transitory” shock lasts three months in expectation so that it has the same duration as their quarterly transitory shocks, and we re-estimate the size of shocks to match the same annual moments. The MLM model is a continuous Ornstein-Uhlenbeck process. We discretize this process and simulate it in increments of 1/100 of a month, aggregating the results to compute monthly moments. The CHT model is a continuous process with three different types of shocks; as before we can simulate this process, aggregate to the monthly level, and compute the resulting moments.

Table A-13 shows at least three distinct ways that predictions from existing income models differ from the patterns we see in the data. First, as already discussed in Section 3, these models imply less frequent earnings changes than what we observe in the data. How then do the models still match longer-run income patterns? In some models, this is channeled through very high kurtosis (much higher than what we see in the data) while in other models this is channeled through a lower standard deviation of income shocks.

Second, the table shows that deviations from the prior month are much more persistent in most

⁴⁵This issue is particularly salient for these types of applications that focus on higher frequency phenomenon, but the level of risk at all time horizons could also be relevant even for lower frequency choices like retirement savings.

⁴⁶KMV provides intuition: “...consider two possible distributions of annual earnings changes, each with the same mean and variance, but with different degrees of kurtosis. The more leptokurtic distribution... is likely to have been generated by an earnings process that is dominated by large infrequent shocks.”

of the models relative to the data. This is because in most of these models shocks slowly mean-revert, so one positive shock is followed by many months of small negative shocks (or one negative shock is followed by many months of positive shocks). That is, the conditional probability that a change is in the same direction as the previous change is an order of magnitude larger than the unconditional probability. For example, in the KMV model, if we observe a positive change this month there is a 95 percent chance that we then observe a positive change next month. Indeed, in their model households receive a large shock roughly every two years on average. While these large shocks are symmetric in sign, income deviations gradually and deterministically mean-revert toward zero between these infrequent events.⁴⁷ This means that sequences in which income changes in the same direction for 24+ months in a row are fairly common. In the data, there is much more rapid mean reversion: indeed even though the unconditional probability of a positive and negative income change is approximately equal, positive changes are more likely to be followed by negative changes and vice versa.

⁴⁷Note that the reason to distinguish positive from negative persistence is merely to highlight that positive changes tend to be followed by positive changes and vice versa, not to imply there is some asymmetry between positive and negative income shocks.

C Appendix with additional empirical results

C.1 Comparing Aggregate Seasonality in PayrollCompany to Benchmarks

In Section 2 we discuss the seasonality of total employment in PayrollCompany data compared to aggregate data reported by the U.S. Bureau of Labor Statistics, Current Employment Statistics (PAYNSA). To estimate the seasonality of total employment in BLS data, we regress log total employment on 12 calendar dummies plus a quadratic time-trend (so that general employment growth does not result in biasing up calendar effects later in the year): $\log(emp)_t = \sum_{k=1}^{12} \beta_k D_k + \alpha_0 t + \alpha_1 t^2 + \varepsilon_t$. In PayrollCompany, we use a set of firms which is balanced within calendar year to remove spurious effects from firms changing payroll processors over time.

We use data from 2010-2023 (dropping 2020 so that pandemic effects do not obscure general seasonal patterns). 95 percent confidence intervals are computed using heteroskedasticity robust standard errors. We note that standard errors are fairly large because we are estimating 12 calendar dummies + a quadratic time trend with 156 monthly observations.

Overall seasonal patterns of aggregate employment are similar, with the biggest deviation being that PayrollCompany data exhibits more employment growth in the summer. This could be driven by a different mix of industries in this data, or effects of removing firm entry and exit. The fact that we find slightly *larger* seasonality in PayrollCompany data suggests that the limited role of seasonality in driving individual earnings is not driven by studying a sample with too little seasonality relative to the economy as a whole.

C.2 Aggregating Volatility From Jobs to Households

In Table A-3 we show that household-level volatility is, if anything, larger than volatility at the individual job level. This increase across several of our summary statistics is largely an artifact of pooling of hourly and salaried jobs, which have very different earnings-change distributions. This section provides a simple stylized example to demonstrate how pooling can have this type of effect.

Consider a case where 50% of workers are salaried and 50% of workers are hourly, and each household has one hourly worker and one salaried worker. Further, imagine that hourly workers change earnings in 90% of months, and that salaried workers change earnings in 10% of months.

In this case, the median absolute change across all *job months* is 0 (since 50% of all job-months have a no change (10% of hourly worker months plus 90% of salaried workers-months)). However, if we aggregate jobs to the *household* level, the hourly part of household income is going to change in 90% of months while the salaried part is going to change in 10% of months. Therefore, the monthly change will only be zero in $(1-0.9)*(1-0.1)=9\%$ of months. So this means that the median absolute change in household income is going to be greater than zero, indicating *more* volatility than the median absolute change measured at the job level.

However, this same increase from job to household level does not occur when combining two hourly jobs into one household.

C.3 Unpaid Leave

In Fact 2 of the main text, we briefly note that unpaid leave appears too small quantitatively to explain much of the earnings volatility experienced by hourly workers. This appendix describes the procedure underlying that conclusion.

We ask whether unpaid leave, which may reflect vacation, medical leave, or caregiving responsibilities, could account for a meaningful share of the monthly earnings volatility observed for hourly workers. Because unpaid leave is not directly observed in the payroll data, we estimate its contribution using a three-step procedure that combines representative survey evidence with the payroll records.

First, we estimate the average amount of unpaid leave taken by U.S. workers. We focus on full-time workers, for whom the concept of unpaid leave is more clearly defined than for part-time workers. The 2017 American Time Use Survey asks workers whether they took any leave from their job in the previous seven days, how many hours of leave they took, and whether that leave was paid or unpaid. Full-time hourly workers report taking unpaid leave equal to 2.33 percent of their usual hours worked. Annualized, this corresponds to 6.53 days of unpaid leave per year for the average worker.

Second, we impute when that unpaid leave occurs in the payroll data. We use a simple algorithm that assigns unpaid leave to the pay periods with the fewest paid hours within a six-month window. Specifically, we first assume that unpaid leave occurs in the pay period with the lowest number of paid hours. If the unpaid leave budget is not exhausted by raising hours in that pay period up to the level of the second-lowest pay period, then we assign unpaid leave to the second-lowest pay period as well. We continue in this way until the unpaid leave budget is exhausted or hours are equalized across all pay periods in the six-month window. Figure A-9 illustrates the algorithm for one firm. This timing assumption is intentionally conservative in the sense of maximizing the potential role of unpaid leave in generating observed volatility. By assigning unpaid leave to the lowest-hours pay periods, the procedure attributes as much of the observed variation in hours as possible to leave-taking. If anything, this should overstate the contribution of unpaid leave to earnings volatility.

Third, having imputed unpaid leave in each pay period, we calculate the monthly earnings volatility that would have been observed if those leave hours had instead been worked. This yields a counterfactual measure of income volatility absent any volatility attributable to unpaid leave. Table A-14 shows that the resulting counterfactual volatility is very similar to observed volatility. For full-time hourly workers, the median absolute monthly change in pay falls only from 6 percent to 5 percent, and the 75th percentile falls from 15 percent to 13 percent. Thus, although unpaid leave may be important for some workers, typical amounts of unpaid leave are too small to explain more than a modest share of the month-to-month earnings instability we document.

Finally, we provide indirect validation of the timing assumption by applying the same algorithm to *paid* leave, which is observed in the payroll data. Figure A-10 plots average paid leave hours against vigintiles of the change in log monthly total hours per paycheck. Workers take some

paid leave in nearly every month, but paid leave rises systematically when hours worked are low. In months with declining hours, paid leave increases with a slope of roughly 0.4. The algorithm’s predictions for the timing of paid leave closely track this relationship, suggesting that the procedure does a reasonable job of identifying when leave is most likely to occur.

C.4 Predictable Annual Variation and Earnings Instability

In Fact 2 of the main text, we briefly report that the month-to-month earnings changes of hourly workers are not well explained by simple models of predictable annual recurrence, whereas the earnings changes of salaried workers who receive bonuses are more predicted by those models. This appendix provides the underlying methodology and results.

A natural hypothesis is that the volatility of hourly earnings reflects seasonal or other regularly recurring fluctuations in pay, such as recurring changes in labor demand or regular performance-related payments. To assess this possibility, we estimate the extent to which monthly earnings changes can be predicted using simple models of annually recurring variation. Specifically, we estimate regressions of the form

$$\log y_{i,j,t} - \log y_{i,j,t-1} = \beta X_{i,j,t} + \epsilon_{i,j,t}, \quad (9)$$

where $y_{i,j,t}$ is worker i ’s total earnings per paycheck at firm j in month t , and $X_{i,j,t}$ is a vector of predictors designed to capture predictable annual variation.

We consider two specifications. The first uses firm-by-month fixed effects, $\alpha_{j,m(t)}$, to capture recurring firm-specific calendar patterns, where $m(t)$ indexes calendar month so that, for example, January 2011 and January 2012 share the same value of $m(t)$. The second uses the worker’s own pay change from 12 months earlier, interacted with month fixed effects,

$$\alpha_{m(t)} + \beta_{m(t)} (\log y_{i,j,t-12} - \log y_{i,j,t-13}),$$

to capture individually recurring annual patterns. This second specification uses no firm-specific information; it relies only on the worker’s own pay change from a year earlier. Because it requires data for at least 13 months of pay changes, it can only be estimated for workers observed for at least 14 months. We also explored combinations of these specifications and reach similar conclusions.

Table A-5 shows that these models explain only a modest share of pay changes for continuing hourly workers. For hourly workers, the R^2 ranges from 0.03 to 0.13, depending on the specification. Salaried workers who do not receive bonuses have similarly low explanatory power. By contrast, earnings changes are more predictable for salaried workers who do receive bonuses, with R^2 values ranging from 0.24 to 0.39.⁴⁸ In unreported regressions, we verify that the bonus component of pay is the main reason why earnings changes are more predictable for this group. Specifically, when

⁴⁸We classify salaried workers as bonus recipients if they receive a bonus in more than $\frac{1}{24}$ of their months in the data. This threshold is intended to capture employees receiving annual bonuses. An employee who receives annual bonuses could receive one bonus in 23 months of work but would receive two bonuses in 24 months of work. Using this definition, 44 percent of salaried workers are classified as bonus recipients.

we instead define the prediction target as an indicator for receiving any bonus, or the amount of a bonus conditional on receipt, the R^2 exceeds 0.6.

A second pattern in Table A-5 is that the firm-by-month fixed effect specification typically delivers somewhat higher R^2 values than the specification based on 12-month lags. One interpretation is that firm-specific recurring calendar patterns contain information that cannot be recovered from a worker's own lagged pay history alone. Another is that the firm-by-month specification may partly overfit the data, since it includes a very large number of parameters. To limit this concern, we estimate that specification only for firms with an average of at least eight employees per month. Even so, the main conclusion is unchanged: predictable annual variation explains only a relatively small share of the month-to-month earnings volatility of continuing hourly workers.

We also compare the seasonality of individual earnings changes to the seasonality of broader firm-level labor adjustment. While hours changes for continuing workers are not very predictable under these annually recurring specifications, total firm hours, which also reflect hires and separations, are more predictable. This pattern suggests that recurring seasonal demand may matter more for firms' total staffing needs than for the month-to-month earnings changes of incumbent hourly workers.

D Appendix on the Role of Firms in Earnings Instability

D.1 Interpreting Fluctuations in Total Firm Hours

In Section 5.1.1 we show that firms have substantial fluctuations in total monthly hours. Since we wanted to interpret these fluctuations as arising from labor demand rather than supply, we focused on firms with a median size of at least 20, since smaller firms might have idiosyncratic labor supply shocks that spill over into total firm hours. In this appendix we show robustness to alternative thresholds for this firm-size cutoff and discuss further justification for the interpretation of these total-hours changes as firm rather than worker driven.

Figure A-11 shows the distribution of total firm hours changes for continuing workers under this baseline size cutoff of 20 is similar to that obtained when using a minimum size of 50 or of 100. While not identical, the key point is that firms with more than 100 workers also see sizable monthly changes in total hours of continuing workers. Furthermore, we have repeated all results in Section 5 for these larger firms and find very similar point estimates for all relationships. The only substantive difference is that results become somewhat noisier as the overall size of the firm sample shrinks rapidly. Our baseline sample with a cutoff of 20 includes 637 firms while we only retain 50 firms when using a cutoff of 100.⁴⁹ The fact that results are very similar when using a cutoff of 20 as when using these much larger cutoffs suggests that idiosyncratic shocks within firms are not driving these changes in total firm hours from month to month.

We have also explored two additional exercises that reinforce the conclusion that idiosyncratic shocks are not important for total firm hours once imposing this size cutoff of 20. First, in the data we have explored regressions of $\Delta hours_{i,t} = \alpha + \beta \Delta \overline{hours}_{j(-i),t} + \varepsilon_{i,t}$, where $\overline{hours}_{j(-i)}$ is the average change in hours for all of worker i 's co-workers at firm j in month t , i.e., it is the leave-self-out-mean of firm wide hours. If individual hours changes are the sum of some idiosyncratic and some common firm-component, it can then be shown that β will converge to one in the limit as firm-size goes to infinity. This is because in a very large firm, the leave-self-out mean converges to the overall mean. Indeed, this is another way of stating that in a large firm, any individual worker's hours will not drive a meaningful change in average firm hours. This implies that we can use this regression as a diagnostic to assess whether individual and co-worker hours co-move on average, as they should if firms are large enough for idiosyncratic shocks to wash out. We find strong support for this comovement: with a firm size of 20+, this regression yields a coefficient of 0.94.

Second, we have explored numerical simulations where workers draw some idiosyncratic and some firm-wide component and explored how these idiosyncratic shocks bias firm-wide inference as the size of firms changes. With a firm size of 20, this bias is minimal. This is especially true if firms are able to adjust other workers hours to offset idiosyncratic shocks. In particular, in some simulations we allow for hours to be determined in two steps: 1. Workers draw some firm shock + some idiosyncratic shock. With no firm adjustment, the total firm hours change is then the sum of

⁴⁹Note that our analysis begins with a one percent sample of firms, so overall sample sizes could likely be expanded with additional computational overhead by drawing a larger initial sample.

these shocks across all workers and so will deviate from the sum of the firm shock if idiosyncratic shocks do not add up to zero. 2. The firm can partially adjust the hours of every individual worker to try to offset these individual hours shocks and target a sum idiosyncratic change of zero. The amount of such smoothing allowed by the firm is then a simulation parameter. We find that in practice, if firms are able to adjust co-workers by even ± 1 day each month, this is sufficient to largely smooth out idiosyncratic shocks, even in smaller firms around 5-10 workers.

While restricting to 20+ size firms reduces concerns about idiosyncratic labor supply shocks, it is possible that fluctuations in total firm hours might be driven by correlated labor supply shocks, e.g., seasonal patterns or inability to hire in certain sectors during the pandemic. However, most of these potential confounds would likely occur at the industry rather than the firm level. To remove effects of any industry wide variation, we thus run a regression of firm total-hours changes on time $\times$ industry controls and then compute the distribution of the residuals from this regression. This specification removes any industry-wide labor supply fluctuations from month to month but also removes industry-wide demand shifts. In this sense it is likely controlling for some fluctuations that arise from demand and not just fluctuations that arise from labor supply. Nevertheless, we find that the distribution of these firm-specific residuals is extremely similar to the distribution of raw firm-month changes. While there are some industry-specific shifts, these have little effect on the overall distribution of firm-month hours changes: most of these movements are firm-specific within industry. Thus, industry-wide labor supply shocks seem unlikely to be driving our conclusions.

If idiosyncratic labor supply shocks and industry-specific labor supply shocks do not drive monthly changes in total hours at the firm, the only remaining confound is from correlated labor supply shocks that are firm-specific. We cannot entirely rule this out, and indeed some forces like strikes or firm-specific contagious health shocks might generate changes in total firm hours in some months. However, it seems unlikely that these types of relatively rare events could drive the frequent fluctuations shown in Figure 6 and so we think it is more plausible that these fluctuations are primarily driven by firm-driven shifts in labor demand.

D.2 Proportional Allocation of Total Firm Hours Movements to Individual Workers

Because many different counterfactuals are consistent with the same observed net and gross changes in hours, quantifying the role of firm hours changes for individual hours requires further assumptions. We proceed using a back-of-the-envelope allocation rule to construct counterfactual individual hours changes given counterfactual changes in total firm hours. Each worker spell i is associated with a unique firm $j(i)$, so to simplify notation, in this section we drop the redundant firm index j on ΔH_t^{firm} . Specifically, let $\Delta \tilde{h}_{i,t}(x)$ denote the change in monthly hours for worker i under a counterfactual change in total firm hours equal to x . We are interested in how $\text{Vol}_i(x) \equiv \text{Median}(|\% \Delta \tilde{h}_{i,t}(x)|)$ varies with the size of the firm wide hours shock x , and in particular in how large individual hours changes would be if total firm hours were held fixed ($x = 0$). To construct $\Delta \tilde{h}_{i,t}(x)$, we assume that changes in total firm hours are allocated across individual

workers according to the following proportional allocation rule:

$$\Delta \tilde{h}_{i,t}(x) = \begin{cases} \Delta h_{i,t}, & \text{if } \text{sign}(\Delta h_{i,t}) \neq \text{sign}(\Delta H_t^{\text{firm}}) \text{ or } \Delta H_t^{\text{firm}} = 0, \\ \frac{\Delta h_{i,t}}{\Delta H_t^{\text{same}}} (x - \Delta H_t^{\text{firm}} + \Delta H_t^{\text{same}}), & \text{if } \text{sign}(\Delta h_{i,t}) = \text{sign}(\Delta H_t^{\text{firm}}). \end{cases} \quad (10)$$

where we define $\Delta H_t^{\text{same}} \equiv \sum_{\{i: \text{sign}(\Delta h_{i,t}) = \text{sign}(\Delta H_t^{\text{firm}})\}} \Delta h_{i,t}$. When we set $x = \Delta H_t^{\text{firm}}$, Equation (10) returns the observed change in individual hours $\Delta h_{i,t}$. Alternative values of x instead generate counterfactual individual hours changes.⁵⁰ For example, when total firm hours are held fixed ($x = 0$), individual hours changes are proportionately reduced for all workers whose hours move in the same direction as the firm: $\Delta \tilde{h}_{i,t}(0) = \frac{\Delta h_{i,t}}{\Delta H_t^{\text{same}}} (\Delta H_t^{\text{same}} - \Delta H_t^{\text{firm}})$.

This allocation rule makes two key assumptions, corresponding to the two rows of Equation (10). First, the change in firm hours is allocated only among workers whose observed hours moved in the same direction as the firm. Second, the firm change is then allocated across these workers in proportion to their observed individual hours changes. The data reject equal allocation across all workers, which supports the assumption that those workers whose hours co-moved most strongly with total firm hours were those who absorbed the firm shock. Consistent with dual labor market theories in which some workers are insulated from firm-wide changes, we find that individual changes of exactly zero are much more common than firm-wide changes of exactly zero.

Using $\Delta \tilde{h}_{i,t}(x)$, we construct an estimate of what individual hours volatility would be if total firm hours volatility was eliminated ($Vol_i(0)$).⁵¹ Figure A-12 compares the observed CDF of individual volatility across workers to the counterfactual CDF obtained after eliminating total firm-hours volatility. The figure shows that eliminating total firm-hours volatility would substantially reduce individual volatility. For example, the median value of volatility $Vol_i(\Delta H_t^{\text{firm}})$ is 9.0 percent while the median value of $Vol_i(0)$ is 3.9 percent. The mean value of volatility falls from 11.6 percent to 6.3 percent. Thus, these calculations suggest that roughly half of individual volatility for the typical worker arises from fluctuations in total firm hours. We emphasize that these are back-of-the-envelope calculations that depend on the allocation rule. However, the broad conclusion that fluctuations in total firm hours account for around half of individual volatility is supported by independent evidence from the movers design in Section 5.2.

⁵⁰For example, suppose a three-worker firm has individual hours changes of $[-6, 6, 12]$. The total increase in firm hours is 12. Our allocation rule assumes that the workers with positive changes (+6 and +12) provided these additional hours in proportion to their share of all individual hours increases at the firm: (6/18) and (12/18). This implies that the +6 worker absorbed 4 hours of the firm-wide shock while the +12 worker absorbed the remaining 8 hours. Thus, individual hours changes would have been $[-6, 2, 4]$ without the +12 firm shock.

⁵¹Since percent changes depend on lagged hours levels, we need to cumulate the counterfactual changes into a counterfactual levels series. To construct percent changes we thus compute the full time-series of $\Delta \tilde{h}_{i,t}(0)$, cumulate these changes over time to obtain a counterfactual series of hours levels $\tilde{h}_{i,t}(0)$ for each worker, and then compute $Vol_i(0) = \text{Median}(|\% \Delta \tilde{h}_{i,t}(0)|)$.

D.3 Identification of Firm Effects on Income Volatility

In this appendix, we provide additional discussion and diagnostics supporting the empirical specification and causal interpretation of the movers design in Equation (5). In particular, we want to explore both whether the linear additive specification is reasonable as well as whether the identification assumptions are satisfied. As discussed in the main text, a causal interpretation of Equation (5) requires that there is no selection on the match-specific error term in the regression. This identification assumption could be violated if a change in a worker’s idiosyncratic volatility causes them to switch to a firm with different volatility. We begin our empirical analysis of identification by showing that volatility trends satisfy standard pre-trend diagnostics suggesting that this is not the case. The identification assumption could also be violated if the realization of the match-specific component of volatility affects which matches are actually formed, and Borovičková and Shimer (2024) argues that these standard diagnostics may fail to detect these types of violations. Thus, we next turn to a discussion of this particular identification challenge and argue that it is likely less of a concern in our context studying volatility than in more common applications studying wage determination.

A typical diagnostic in applications of the AKM approach to studying wages is the use of event study designs. For example, Card, Cardoso, and Kline (2016) looks at the dynamics of wages for those who transition from firms with low co-worker wages to those with high co-worker wages. In particular, they compute average wages for workers in different quartiles of co-worker wages over time before and after job transition. They then argue that sharp jumps at transition with no evidence of pre-trends supports the exogenous mobility assumption, and that roughly symmetric changes when transitioning from high wage to low wage as from low wage to high wage jobs supports the linear specification with additive separability.

We now explore similar diagnostics with our volatility outcome. However, we note that two issues complicate this analysis in our context. First, event studies that look at outcomes over time are more complicated in our setting, where the main outcome is individual-level volatility, $Vol_{i,j} = \text{Median}_t(|\Delta Y_{i,j,t}|)$, than in typical settings where the outcome is $wage_{i,t}$. This is because our individual-level volatility statistic necessarily requires using data from multiple observations over time, unlike individual wages, which can be calculated at each date t . That is, it is not straightforward to measure how an individual’s volatility varies over time within a spell when this volatility is itself constructed as the dispersion of the individual’s earnings changes over time.

To address this issue and construct a measure of volatility that varies within an individual spell, we construct $Vol_{\tau}^{pooled} = \text{Median}_i(|\Delta Y_{i,j,\tau}|)$, where τ indexes the month relative to a move, and i indexes different individuals. The key difference when comparing this volatility measure to our main individual-level volatility measure is that we interchange the dimension over which the median is taken. Our primary volatility measure fixes an individual and measures the volatility of their earnings over time. This alternative volatility measure instead fixes an event time and measures the cross-sectional dispersion of earnings changes across individuals at that event date. That is, it is a measure of volatility which pools across individuals at a particular event date rather than

measuring volatility across event dates for a particular individual. Thus, for short-hand we call this new volatility measure “pooled volatility” at event time τ and refer to our original volatility measure as individual volatility.

Using pooled volatility, we can then compute standard event study designs. However, using pooled volatility introduces a second issue: compositional concerns and sampling error. In particular, it is important to note that pooled volatility can change over time either because the individual volatility changes over time (which is what we are interested in) or because the composition of the pooled volatility changes over time (which we are not interested in). Essentially this pooled volatility event study asks whether workers one month from job transition have more volatile earnings changes (compared to each other) than workers two months from job transition, and so on. This means that it is important to compute pooled volatility for a balanced sample of workers over time. Furthermore, when making comparisons of volatility levels across groups instead of changes, compositional issues arising from sampling error can remain relevant even with a balanced panel. Finally, we note that larger samples are required in our context to eliminate sample noise than in the wage context, since we are measuring a second moment rather than a first moment.

While we observe many movers with a wide variety of individual spell lengths at firms pre and post move in the payroll data, once we restrict to balanced panels of workers with at least length T spells pre and post move and further cut the data by those who are transitioning between different quartiles of firm volatility, sample sizes rapidly become very small. This is especially true because firm turnover in our payroll data is substantial, with the median firm staying in the data for less than four years. Computing pooled volatility for a 4-month event study retains our entire movers sample, since we require individual workers to have at least 4 complete months of data to compute volatility, but as we move to longer event study windows sample sizes decline dramatically. Recall that we exclude the first and last (potentially incomplete) month of each individual worker job spell, so computing a 4-month event study requires workers to be observed at the first job for at least 6 months and at the second job for at least 6 months. Computing a 12-month event study would require restricting to workers observed for at least 14 months at the first job and 14 months at the second, which is quite restrictive since there is substantial firm turnover in the payroll data. Even if workers actually work this long, we are unlikely to observe them at the firms for this period of time.

For this reason, only 4-month event studies are feasible in our payroll data, and even these results are fairly noisy. With a noisy and short time-sample it is then challenging to differentiate noise from pre-trends. For this reason, we also compute a similar event study design using data from Chase, which allows us to look at longer event windows. This is because the Chase data is not subject to the same firm turnover issues and because overall we observe more movers in this data since it has a wider coverage of firms.

With these caveats in mind, Figure A-13 panels (a) and (b) show that there are sharp changes in volatility around job transitions.⁵² There is also no evidence of pre-trends. Furthermore, these

⁵²Transitions from quartile 1 to 4 and vice versa are more unusual and so there is more sampling error and resulting

effects are relatively symmetric, especially in the Chase data, which has less sampling error: workers transitioning from quartile 1 to 4 firms have volatility changes of the same magnitude but opposite sign as those transitioning from quartile 4 to quartile 1. While these results cannot conclusively validate the identification assumptions, they provide support similar to that in the more standard wage determination context.

Nevertheless, Borovičková and Shimer (2024) argues that these diagnostics may fail to detect violations arising from selection around match-formation. If matches with particular realizations of ε_{ij} are more likely to form, then this could lead to bias in estimated firm and worker fixed effects. However, while they cannot conclusively prove that our identification assumption is satisfied, two observations make us less concerned about selection on match formation than in the typical wage-AKM context.

First, violations of the identification assumption that arise from selection on match formation are likely to bias us towards finding *smaller* causal effects of firms. This is because our evidence shows that workers appear to dislike earnings volatility, suggesting that matches with unusually high volatility would be less likely to form. Let j denote the origin firm and j' denote the destination firm, and let $k(j)$ denote firm j 's volatility decile. Selection would then imply that moves from lower-volatility origin firms to higher-volatility destination firms, $k(j') > k(j)$, are more likely to occur when the worker's match-specific residual declines. That is,

$$E(\Delta\varepsilon_{ij} \mid k(j') > k(j)) < 0.$$

If this is the case, then the resulting estimate of $\psi_{k(j')} - \psi_{k(j)}$ will be biased down. Intuitively, if workers only transition from low to high causal-effect firms when the match-specific residual is low, we will observe small increases in observed volatility for these movers. This means we will understate firm causal effects, since we only observe these effects in the instances when they are offset by unusually low match-specific terms. Thus, the presence of match-formation-based selection is likely to lead us to *understate* the importance of firms. In this sense, our conclusion that firms have important causal effects on volatility is likely conservative. This is in contrast to the standard wage-AKM selection concern, where a small match-specific draw may make a match less likely to form and bias towards finding larger firm effects.

The second reason that selection from match formation is likely less of a concern for volatility than for wages is that the wage of a particular job is observed before accepting a job offer, but the earnings volatility of a job is less easily observable in advance. If volatility is not observed before a match is formed, it is less likely that there will be selection on this outcome. Of course, there is clearly a component of volatility that is observable (e.g., restaurant jobs are more volatile than IT jobs). Furthermore, even if volatility is not observed in advance, if volatility is correlated with other variables like the wage that are observable, then this might still induce incidental selection. However, to try to address these concerns, Appendix D.4 shows that results are very similar if

noise for these groups.

we restrict only to movers within industry, where workers are less likely to observe variation in volatility ex-ante, and are also very similar if we restrict to only moves between firms with similar wages.

D.4 Variance Decompositions

In this appendix we report variance decompositions of cross-spell heterogeneity into individual, firm and match-specific effects that we obtain from running the movers fixed effect specification in Equation (5). These types of decompositions are the focus of much of the AKM literature.

Before reporting results, we note that a large literature has emphasized the fact that sampling noise caused by small samples does not bias estimated fixed effects but can severely affect their variance properties and lead to misleading conclusions about variance shares. This concern is particularly acute in our setting since our outcome of interest is a second moment (earnings volatility in a spell) rather than a first moment (average wage in a spell), and is thus much more sensitive to sampling error. Sampling error is a particular concern for estimating worker effects, because individual job spells for most workers are not that long.⁵³ For this reason, we follow Kline, Saggio, and Sølvesten (2020) and implement a leave-one-out approach for estimating the variance of worker and firm fixed effects, and as suggested by Kline (2024), we focus on comparisons of the standard deviation of worker FE to the standard deviation of firm FE.⁵⁴

Table A-10 shows comparisons of firm and worker FE after implementing the leave-one-out bias correction of Kline, Saggio, and Sølvesten (2020). The baseline specification shows that the standard deviation of firm fixed effects is the same as the standard deviation of worker fixed effects. These units can also be directly compared to the level of typical individual volatility. For example, moving to a firm with a one standard deviation higher firm FE increases individual volatility by 4.8 percentage points, which is about one-third of the mean individual volatility of 14 percent.⁵⁵

Table A-10 also explores a number of robustness checks. As discussed above, the identification assumption in AKM regressions rules out match-specific sorting, where a change in worker’s preferences for volatility causes them to switch to a job with different volatility. We provided empirical evidence in Appendix D.3 to support this assumption, but also noted that it is more likely to be satisfied if firm volatility is not a characteristic that is observed by workers prior to taking a job. Since differences in volatility across firms *within* an industry are likely less observable prior to move than differences in volatility across industry, in Row 2, we redo the movers analysis restricting only to within industry moves, and this does not change the conclusion about the size of firm fixed

⁵³This is a fundamental feature of the data generating process for the U.S. economy and so this type of sampling error would be significant even with a full census of individual work histories.

⁵⁴The Kline, Saggio, and Sølvesten (2020) correction reduces sampling bias in SD(worker FE) and SD(firm FE) but does nothing to eliminate sampling bias in actual Vol_i since it simply reallocates variance from worker and firm FEs to match-specific components. This means that sampling bias is likely to lead to low values for measured values for explained variance even in a world where there are no true match-specific effects.

⁵⁵As a point of comparison, Kline (2024) summarizes a variety of past estimates of log wage effects typically finding values around 25 percent. This suggests that firm effects on earnings volatility are likely of similar or greater importance to their effect on wages.

effects. If volatility is correlated with wages, then match formation based selection on wages might induce selection on volatility even if volatility is not observed. Thus Row 3 restricts the analysis to only workers who are moving within the same wage decile and again shows similar effects.

The accommodation and food industry is somewhat over-represented within the movers sample. This industry has higher volatility and more job-churn than the typical industry, but Row 4 shows that our results are not driven by this particular sector. Redoing our results excluding moves within or between this industry and others delivers very similar conclusions.

Our baseline AKM specification groups firms into 10 groups. We must group firms rather than estimating true firm fixed effects, because the connected set of movers across firms in our data is small. However, this means that there is almost certainly heterogeneity in firm effects *within* these groups that is then missed in our baseline variance decomposition. To explore this, we re-estimate Equation (5) but splitting firms into 25 instead of 10 quantiles.⁵⁶ This shows that, indeed, estimated firm fixed effects rise very slightly when using these finer groups.⁵⁷

We are primarily interested in overall pay volatility, so we estimated all the AKM results just discussed in pay volatility space. Row 6 illustrates that conclusions are similar and we again find a large role for firms when re-estimating volatility in hours space. This is not surprising, since we showed that pay volatility is mostly driven by hours volatility, but it is useful for relating to the evidence in Section 5.1.2 that focused on volatility in hours space.

Finally, Table A-7 shows that the youngest workers have substantially greater pay volatility. Row 7 re-estimates results using only movers aged 26-65 to show that these workers, who may be more marginally attached to the labor force, do not drive results.

⁵⁶Note that only firm movers that move across firm groups contribute to our estimates. This is why the number of movers rises when we change from 10 to 25 firm groups.

⁵⁷Using even larger numbers of groups does not appear to change the results much more, although the connected set of movers drops rapidly as the number of groups is increased.

E Appendix on Spending Volatility

As in our PayrollCompany analysis, our measure of volatility of earnings ($Vol_{i,j}^y$) and of spending ($Vol_{i,j}^e$) is the median of the absolute monthly percent change within a job spell, but results are similar if we instead measure volatility as the standard deviation of monthly percent changes.

Several of our specifications in Section 6 instrument for individual income volatility using the average volatility of the firm. To construct our primary firm volatility measure we average monthly income volatility for all workers in that firm during the entire sample period (2012-2018), weighting each worker by the number of months in which they were employed at the firm, but we obtain similar results if we instead compute the median firm volatility or if we weight all workers equally instead of weighting by the number of worker-month observations. Note that firm volatility is the average of individual worker volatility like in Section 5.2.1, it is *not* the volatility of total firm hours measure used in Section 5.2.3. This is because we are interested in estimating the effects of all volatility on spending and not in the effects of firm total-hours volatility in particular. Furthermore, it is not possible to measure firm total-hours volatility in Chase data since we only observe the subset of workers in the firm with Chase bank accounts rather than all workers.

In the main text we discuss extensions of our results to allowing for heterogeneity as well as various robustness checks. We explore several covariates in the data. In addition to salaried vs. hourly, we define high and low income as the highest and lowest third of monthly pay per pay check. We measure the median checking account balance within a job spell and define high vs. low checking as the top and bottom third of these median balances. Work and non-work spending are split according to the definitions in Ganong and Noel (2019), which are based on sensitivity of spending to retirement. Total spending includes all account outflows except those which are specifically tagged as transfers to other financial accounts.

A causal interpretation of the results above and in the main text requires identification assumptions that are discussed in the main text. As usual, these are not directly testable, but we now provide two diagnostic exercises to provide some support for these assumptions. First, we explore difference-in-differences designs around job transitions to explore the role of pre-trends. Second, we explore robustness of our primary specifications to leaving out periods around job transitions.

Figure A-14 constructs a difference-in-differences event study around job transitions. The sample of movers is split into two groups: those who experience an increase in income volatility after the transition and those who experience a decrease. We then measure how the difference in consumption and income volatility between these two groups evolves in the months around the transition.

As in other panel specifications, we must work with a pooled volatility measure to allow for time variation, which introduces complications in constructing this event study. Concretely, for each worker i , we define event time τ relative to the move, where $\tau = 1$ corresponds to the second full month at the destination job (i.e., the first month where a within-job change can be computed), and $\tau = -1$ corresponds to the last full month at the origin job.

We are interested in how individual workers' volatility evolves with event time, comparing workers who move to higher-income volatility jobs to those who move to lower-income volatility

jobs. However, volatility is measured at the group level by pooling across individuals, which may introduce composition effects if the types of workers observed vary across event time. To address this, we normalize outcomes by demeaning each worker’s absolute percentage change in consumption and income relative to their own mean over all observed event times. For example, define:

$$|\Delta \tilde{y}_{i\tau}| = |\Delta y_{i\tau}| - \overline{|\Delta y_i|}$$

This normalized measure captures whether the absolute percent change in income for worker i at time τ is larger or smaller than their own average change across the event window.

We then estimate the following event study regression:⁵⁸

$$|\Delta \tilde{y}_{i\tau}| = \sum_{\ell \neq -1} \gamma_\ell \mathbf{1}\{\tau = \ell\} + \delta \text{Group}_i + \sum_{\ell \neq -1} \beta_\ell (\mathbf{1}\{\tau = \ell\} \times \text{Group}_i) + \varepsilon_{i\tau},$$

where Group_i is an indicator for whether median income volatility of worker i , $\text{Med}|\Delta y_i|$, rises or falls from the origin to the destination job, $\mathbf{1}\{\tau = \ell\}$ is a dummy for each event time (with $\tau = -1$ omitted as the baseline), and β_ℓ captures how normalized volatility evolves for the rising-volatility group relative to the falling-volatility group.

The coefficients β_ℓ then deliver a standard difference-in-differences event-study interpretation: they describe how volatility evolves before and after job transitions for households whose income volatility rises relative to households whose income volatility falls. Figure A-14 shows that income volatility jumps up after job transition for those in the rising income volatility group. This is by construction since this is how groups are defined. However, the lack of a pre-trend is not mechanical and provides validation for the identification assumptions in our main text. This pattern is also consistent with Figure A-13. More notably, we also observe a jump in the volatility of consumption at the time of transition, again with no evidence of pre-trends.

It is important to note that, for each event time ℓ , β_ℓ estimates how the pooled mean of the normalized absolute changes, $\mathbb{E}[|\Delta \tilde{y}_{i\tau}| \mid \tau = \ell]$, evolves around the transition. As such, this event study differs from our main specification in two important ways: (1) it pools volatility across households at each event time τ , rather than measuring it within households over time; and (2) it uses the mean of the absolute changes, $\text{Mean}|\Delta \tilde{y}_{i\tau}|$, rather than the median.

Thus, we also pursue a second approach for assessing whether confounding shocks around the time of job transition drive the relationship between income volatility and consumption volatility that more closely follows our preferred specification (7). If the relationship between income and consumption volatility were driven by confounding shocks that cause job transitions, then we would expect relationships between income and consumption volatility to attenuate when re-computing an individual’s volatility excluding the changes occurring close to job transitions. That is, we can simply exclude some “donut” around job transition when calculating each individual’s volatility and then re-estimate Equation (7) with these alternative volatility measures. The estimate at zero

⁵⁸We weight each household by the total number of periods they are observed in the sample. This ensures that households observed at a given event time τ contribute equally to the estimation of β_τ .

corresponds to the point estimate in our preferred specification (Column 3 of Table 3) and Figure A-15 shows that while there is a modest decline in precision, point estimates are essentially identical when dropping the 4 months before and after move from all volatility calculations. This suggests that the relationship is not driven by confounding shocks around the time of moves.

In Column 7 of Table A-11, we present additional results with a sample containing households with one or two jobs. For households with multiple jobs, we restrict to those with the same number of paychecks every month to avoid spurious volatility due to differences in the pay frequency. We match the relative shares of households with one relative to two jobs from the CPS in this sample. We measure $Vol_{k,j}^y$ using total income of household k and $\overline{Vol_{j(i(k))}^y}$ as the volatility a firm corresponding to a particular job within the household. Since we want to weight households rather than jobs equally in this regression, for multi-job households, we randomly select a single job when running regressions.

E.1 Calculating the Willingness to Pay to Eliminate Volatility

In the main text, we provide evidence that income volatility increases both spending volatility and worker quit rates. This appendix converts these empirical relationships into magnitudes which are more easily interpretable. In particular, we use off-the-shelf approaches from the literature to compute the “willingness-to-pay” to eliminate income volatility implied by these estimates. These approaches require strong parametric assumptions both about functional forms and parameter inputs, so they should not be interpreted as precise quantifications but instead as simple back-of-the-envelopes to gauge the broad magnitudes of empirical relationships.

To gauge the magnitude of the effects of earnings volatility on spending volatility, we use the standard welfare calculation in Lucas (1987):

$$\text{Willingness to Pay} = \frac{1}{2}\gamma\sigma_c^2 = \frac{1}{2}\gamma\beta_\sigma\sigma_y^2, \quad (11)$$

where γ is the coefficient of relative risk aversion, σ_y^2 is the variance of monthly earnings changes, σ_c^2 is the variance of monthly spending changes attributable to earnings volatility, and β_σ measures the causal effect of earnings variance on spending variance.

Our empirical specifications in Table 3, however, are estimated using the median absolute monthly change rather than variance. In particular, the coefficient reported in the table, which we denote β_{med} , comes from regressions relating $Med|\Delta c|$ to $Med|\Delta y|$. To apply Equation (11), we therefore need to translate this median-based coefficient into the corresponding variance-based coefficient β_σ .

Let σ_y^2 denote the variance of monthly earnings changes and σ_c^2 the variance of monthly spending changes. Suppose that spending variance is related to earnings variance according to

$$\sigma_c^2 = \sigma_{c,0}^2 + \beta_\sigma\sigma_y^2,$$

where $\sigma_{c,0}^2$ is the variance of spending for a household with zero earnings volatility. The parameter

β_σ is the object needed for the Lucas formula.

To connect this variance-based relationship to our empirical estimates, assume that monthly earnings and spending changes are approximately normal. Under normality,

$$\text{Med}|\Delta y| \approx 0.6745 \sigma_y \quad \text{and} \quad \text{Med}|\Delta c| \approx 0.6745 \sigma_c.$$

It follows that the coefficient estimated in the data satisfies

$$\beta_{\text{med}} = \frac{\Delta \text{Med}|\Delta c|}{\Delta \text{Med}|\Delta y|} = \frac{0.6745 \sqrt{\beta_\sigma \sigma_y^2 + \sigma_{c,0}^2} - 0.6745 \sqrt{\sigma_{c,0}^2}}{0.6745 \sqrt{\sigma_y^2}}.$$

Solving this expression for β_σ and substituting into Equation (11) yields

$$\text{WTP} = \frac{1}{2} \gamma \beta_\sigma \sigma_y^2 = \frac{1}{2} \gamma \left(\beta_{\text{med}}^2 \sigma_y^2 + 2 \beta_{\text{med}} \sigma_y \sigma_{c,0} \right).$$

This expression translates the median-based regression coefficient into the variance-based welfare object. In practice, the translation is straightforward because β_{med} is estimated directly from Table 3, while σ_y and $\sigma_{c,0}$ are measured in the data.⁵⁹

These welfare calculations also require a choice of risk aversion parameter γ . In the main text, we report willingness-to-pay estimates using $\gamma = 2$ and the level of earnings volatility observed for the median hourly worker, $\sigma_y^2 = 0.05$. Applying the formula above to the estimates in columns 1 through 3 of Table 3 yields implied willingness-to-pay values of 8.8%, 8.7% and 5.1%, respectively.

To gauge the magnitude of the relationship between quits and volatility, we compute a back-of-the-envelope willingness-to-pay using the framework of Gronberg and Reed (1994) which divides the elasticity of separations to some disamenity by the elasticity with respect to wages. In particular, we divide the elasticity of separations with respect to volatility by the elasticity of separations with respect to wages to get an implied willingness-to-pay to reduce volatility.

We use a conservative value from the literature of -3 for the elasticity of separations to the wage, but we note that the exact welfare costs will depend crucially on this parameter as well as on interpreting our empirical estimates as causal relationships.⁶⁰ The Column 1 coefficient implies that an hourly worker with median volatility would give up 9.6 percent of wages to eliminate it.⁶¹ Across all specifications, the implied WTP ranges from 9.2-10.7%.

⁵⁹ An alternative approach is simply to re-estimate the empirical specifications using $\text{Vol} = \sigma^2$ instead of $\text{Med}|\Delta|$. Doing so yields welfare estimates that are nearly identical to those obtained from the conversion above.


⁶⁰ Bassier, Dube, and Naidu (2022) and Lamadon, Mogstad, and Setzler (2022) use quasi-random variation in wages to estimate separation elasticities, finding values of -3.01 and -2.16, respectively. Lamadon, Mogstad, and Setzler (2022) compute total labor supply elasticities of 6.02. Following Bassier, Dube, and Naidu (2022) we divide by 2 to arrive at a separation elasticity. If we instead interpret the total labor supply elasticity as arising from the separations margin and use 6.02, the WTP numbers halve but still remain large. Using -2.16 would generate larger WTP

⁶¹ $0.096 = \left(1 - \frac{1}{\exp(2.85 \times 0.119)} \right) \frac{1}{3}.$


F Additional Appendix Figures and Tables

Figure A-1: Wage Distribution

(a) CPS: All Workers




(b) CPS: Workers at Small Firms




Notes: This figure shows the hourly wage distribution in the PayrollCompany data (green) and in the Current Population Survey (orange). We assume that salaried workers work 40 hours per week. The PayrollCompany data is the same in both panels. The Current Population Survey data uses the Earner Study to measure hourly wages and the ASEC to measure firm size. The bottom panel shows workers who report their firm size as less than 100 workers.

Figure A-2: Distribution of Quarterly Hours Worked




Notes: This figure shows the distribution of quarterly hours worked in the PayrollCompany data (pink) and in the Washington state tax data (black). The Washington state tax data series is from Figure 2.B of Lachowska, Mas, and Woodbury (2022). We take two steps to make the data series as comparable as possible. First, the Washington analysis requires full-quarter employment, meaning that it only reports data from quarters t where the employee also has positive earnings at the same employer in quarters $t - 1$ and $t + 1$. We therefore similarly require full-quarter employment in the PayrollCompany data. Second, the PayrollCompany data do not have hours for salaried workers. We assume they work 516 hours ($4.3 * 40 * 3$) which generates a point mass in the pink distribution.

Figure A-3: Distribution of Paycheck Arrival Frequency




Notes: This figure compares the distribution of pay frequency in PayrollCompany to the distribution in Current Employment Statistics (<https://www.bls.gov/ces/publications/length-pay-period.htm>).

Figure A-4: Seasonality



Notes: This figure compares seasonality of aggregate U.S. private employment data from the Bureau of Labor Statistics, Current Employment Statistics to seasonality of aggregate employment in PayrollCompany. We report the coefficients β_k from the following regression: $\log(emp)_t = \sum_{k=1}^{12} \beta_k D_k + \alpha_0 t + \alpha_1 t^2 + \varepsilon_t$ where D_k are dummy coefficients for each month. In PayrollCompany, we measure aggregate employment using a set of firms which is balanced within calendar year to remove spurious effects from firms changing payroll processors over time. Data runs from 2010-2023, dropping 2020. 95 percent confidence intervals are computed using heteroskedastic robust standard errors. See Appendix C.1 for more detail.

Figure A-5: Employee Turnover



Notes: This figure shows turnover rates in the PayrollCompany data compared to the BLS Job Opening and Labor Turnover Survey (JOLTS) for private employers from 2010-2019.

Figure A-6: Worker-Level Volatility




Figure A-7: Earnings Volatility in PayrollCompany compared to JPMCI





Figure A-8: The Persistence of Monthly Earnings Changes




Notes: This figure plots the variance ratio $\frac{\text{Var}(\log y_{t+k} - \log y_t)}{k \text{ Var}(\log y_{t+1} - \log y_t)}$ for different values of k .

Figure A-9: Predictions of Unpaid Leave Algorithm at One Firm




Notes: This figure illustrates the unpaid leave algorithm for 14 workers at one firm. We assume that workers have 26 hours of unpaid leave to allocate across a six month time horizon.

Figure A-10: Validation of Unpaid Leave Algorithm Using Paid Leave




Notes: This figure reports average amounts of actual paid leave, simulated leave, and the change in monthly hours worked by vigintile of the change in log monthly total hours per paycheck. The sample is full-time hourly workers. Most bins correspond to 5 percent of worker-months, but the bin at zero corresponds to 9.7 percent of worker months. Paid leave accounts for 2.23 percent of compensated hours in the PayrollCompany data (note that this is slightly different from the 2.33 percent of hours of *unpaid* leave used in other analysis). We therefore simulate the timing of unpaid leave assuming that workers have a budget equal to 2.23 percent of hours.

Figure A-11: Robustness of Firm Total-Hours Changes to Size Cutoff




Notes: This figure reproduces the total-hours changes conditional on continued employment from Figure 6 for alternative firm size cutoffs.

Figure A-12: Effect of Total Firm Hours on Distribution of Individual Volatility




Notes: “Worker Volatility: original” shows the CDF of individual worker volatility $Vol_i(\Delta H_t^{firm})$ observed directly in the data while “Worker Volatility: removing firm-wide hours changes” shows the distribution of individual worker volatility $Vol_i(0)$ after removing firm-hours shocks using the algorithm described in Equation (10).

Figure A-13: Volatility Event Studies for Workers Who Move Between Firms




Notes: This shows the evolution of earnings volatility before and after job transitions between firms at different quartiles of the volatility distribution. Volatility is calculated at each event time pooling across all workers at that event time, and we impose a balanced panel of workers over the whole sample period. We use a longer pre and post period in Chase data since there is less firm turnover in the sample, meaning that we observe a larger number of workers for many months before and after moves.

Figure A-14: Difference-in-Difference of Income and Consumption Around Job Transitions:
Groups are Those Increasing vs. Decreasing Income



Notes: This shows the event study difference-in-difference in income and consumption volatility around job transitions. Standard errors are clustered by household.

Figure A-15: IV Regression of Vol^c on Vol^Y Excluding Months around Job Transitions



Notes: This figure re-estimates IV Equation (7) but constructs the measure of individual spending and income volatility excluding between 0 and 4 months around the date of move. Excluding 0 months reproduces the baseline specification.

Table A-1: Contract Type in PayrollCompany versus Representative Benchmarks

	Representative benchmarks		
	PayrollCompany	All workers	Workers at small firms
Hourly	60%	58%	60%
Salaried	40%	42%	40%
All: bonus	35%	40%	36%
All: no bonus	65%	60%	64%

Notes: This table compares the distribution of contract types in the PayrollCompany data to representative national benchmarks. The representative share of workers that are hourly versus salaried comes from the Current Population Survey (CPS) using a sample of workers who respond to both the Earner Study and the ASEC. Small firms are defined for the CPS as less than 100 workers. In the PayrollCompany data, we classify workers as receiving a bonus if they receive a bonus in at least $\frac{1}{24}$ of their months in the data. The representative share of workers that are bonus eligible is from the National Compensation Survey. Small firms there are defined as working at establishments with less than 100 workers.

Table A-2: Earnings Volatility Under Different Winsorization Choices

Specification	Lower bound	Upper bound	Std. dev.
Winsorize top/bottom 2.5% of nonzero changes	-0.53	1.10	0.25
Winsorize top/bottom 0.1% of nonzero changes	-0.93	12.1	0.57
Winsorize top/bottom 0.5% of nonzero changes	-0.80	3.77	0.39
Winsorize top/bottom 1% of nonzero changes	-0.70	2.25	0.33
Winsorize top/bottom 5% of nonzero changes	-0.41	0.68	0.21
Winsorize top/bottom 1% of all changes	-0.64	1.72	0.30
Winsorize changes larger than 50%	Bottom 2% of data	Top 5% of data	0.20

Notes: This table reports the sensitivity of the standard deviation of earnings changes to alternative winsorization thresholds. The variable is the percent change in pay.

Table A-3: Earnings Volatility in Chase

Aggregation Level	Condition	SD	Share of $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $
Job	None	0.24	0.79	0.05	0.17
	Less than 10 employees	0.23	0.74	0.05	0.16
	11-100 employees	0.22	0.79	0.05	0.17
	Greater than 100 employees	0.25	0.80	0.05	0.18
Household	None	0.22	0.80	0.05	0.17
Job	Two-Jobs match population share	0.19	0.74	0.04	0.14
Household	Two-Jobs match population share	0.20	0.83	0.05	0.15
Job	Pseudo-hourly workers in two-job HHs	0.20	0.97	0.08	0.18
Household	Pseudo-hourly workers in two-job HHs	0.15	1.00	0.07	0.13

Notes: This table reports descriptive statistics for earnings volatility using JMPCI data. Firm size is measured as the number of Chase customers who are employees of the firm. In rows 6 and 7 we re-weight the sample to match population shares of multi-job households from the CPS to account for under-representation of multi-job households in Chase since we only observe jobs paid via direct deposit. Results are also similar if we focus only on households where we observe multiple income streams. In rows 5 through 9 we also restrict the sample to households in which different jobs have the same pay frequency, to avoid spuriously inflating household-level volatility by measuring pay per paycheck when the number of checks may vary differently across jobs.

Table A-4: Lagged Median Summary Statistics of Earnings Changes

Sample	Variable	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Share $\Delta > 0$	Share $\Delta < 0$
Full-time Hourly	Total Earnings	0.90	0.06	0.14	0.15	0.49	0.42
Full-time Hourly	Base Wage	0.17	0.00	0.00	0.02	0.17	0.00
Full-time Hourly	Hours	0.87	0.05	0.11	0.13	0.44	0.44
Salaried	Total Earnings	0.34	0.00	0.04	0.34	0.22	0.12
Salaried	Base Wage	0.11	0.00	0.00	0.07	0.10	0.02

Notes: This table displays summary statistics similar to Table 1 but measuring changes relative to a lagged median instead of measuring one-month changes. The percent change relative to the lagged median is defined as $\frac{y_t - \text{Median}(\{y_s\}_{s \in t-1, t-2, t-3})}{\text{Median}(\{y_s\}_{s \in t-1, t-2, t-3})}$. This statistic is useful for detecting asymmetry since its temporary changes induce mechanical symmetry when looking at only a one-month change.

Table A-5: Seasonality of Earnings Changes

Covariates	R^2 from regression		
	Hourly	Salaried	
	All	No bonus	Bonus
Firm x month FEs ($\alpha_{j,m(t)}$), including workers w/ ≤ 12 month tenure	0.09	0.10	0.33
Firm x month FEs ($\alpha_{j,m(t)}$),	0.13	0.13	0.39
Month FEs ($\alpha_{m(t)}$) + 12-month lags ($\beta_{m(t)}(\log y_{i,j,t-12} - \log y_{i,j,t-13})$)	0.03	0.04	0.24
Number of firms	2,620	942	693
Number of workers	62,849	25,799	20,668
Number of worker-months	1,053,239	765,948	647,172

Notes: This table reports R^2 s from regressions using Equation (9), these specifications regress month-to-month change in earnings on predictors. A firm is included in the regression if it is present for at least three years and has an average of at least eight employees of the relevant category (e.g., hourly, salaried) when it is present in the data. The three sample size rows include only workers who persist in the data for at least thirteen months. $m(t)$ is an integer from 1 to 12 (e.g., January 2011 and January 2012 both have $m(t) = 1$).

Table A-6: Joint Income-Liquidity Distribution

Income tercile	Liquidity tercile	Share	Median $ \Delta $
Low	Low	0.18	0.09
Low	Middle	0.09	0.08
Low	High	0.05	0.06
Middle	Low	0.11	0.07
Middle	Middle	0.14	0.05
Middle	High	0.08	0.04
High	Low	0.03	0.06
High	Middle	0.10	0.05
High	High	0.21	0.03

Notes: This table reports median earnings instability across the joint distribution of income and liquidity. Data is from JPMCI. Earnings instability is measured as the median absolute percent change in monthly earnings.

Table A-7: Heterogeneity by Age

(a) All Workers

Sample	Share of total group	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Skew	Kurtosis	Share salaried
Age <25	12%	0.91	0.12	0.28	0.31	1.26	2.70	9%
Age 25 - 35	24%	0.75	0.06	0.17	0.24	1.74	6.43	35%
Age 35 - 45	21%	0.68	0.04	0.16	0.24	1.88	7.30	43%
Age 45 - 55	20%	0.65	0.03	0.14	0.24	1.97	7.91	46%
Age 55+	22%	0.61	0.02	0.14	0.24	2.00	7.98	49%

(b) Hourly Workers

Sample	Share of total group	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Skew	Kurtosis
Age <25	18%	0.97	0.14	0.29	0.32	1.20	2.33
Age 25 - 35	25%	0.93	0.09	0.20	0.25	1.50	4.96
Age 35 - 45	20%	0.91	0.08	0.19	0.24	1.58	5.60
Age 45 - 55	18%	0.90	0.07	0.18	0.24	1.64	6.07
Age 55+	19%	0.89	0.07	0.19	0.24	1.63	5.77

Notes: This table repeats the summary stats for “Total earnings” changes similar to Table 1 but separately by age. Panel (a) computes these statistics for all workers while panel (b) computes these statistics restricting only to hourly workers. Data is from PayrollCompany.

Table A-8: Heterogeneity by Gender and Children

(a) All Workers

Sample	Share of total group	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Skew	Kurtosis	Share salaried
Men with children	27%	0.73	0.05	0.17	0.23	1.84	7.23	39%
Men without children	22%	0.81	0.08	0.20	0.26	1.57	5.03	26%
Women with children	27%	0.78	0.06	0.17	0.23	1.71	6.52	31%
Women without children	24%	0.80	0.07	0.20	0.26	1.60	5.28	25%
Men	49%	0.76	0.06	0.18	0.25	1.71	6.12	33%
Women	51%	0.79	0.06	0.18	0.24	1.66	5.91	28%

(b) Hourly Workers

Sample	Share of total group	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Skew	Kurtosis
Men with children	24%	0.94	0.09	0.19	0.23	1.55	5.65
Men without children	23%	0.96	0.10	0.24	0.27	1.38	3.87
Women with children	27%	0.94	0.09	0.20	0.25	1.51	5.11
Women without children	26%	0.93	0.10	0.23	0.27	1.43	4.13
Men	47%	0.95	0.10	0.21	0.25	1.47	4.70
Women	53%	0.94	0.09	0.21	0.26	1.47	4.61

Notes: This table repeats the summary stats for “Total earnings” changes similar to Table 1 but separately by gender and those with and without dependent children. The presence of children is measured using information from W-4’s. Panel (a) computes these statistics for all workers while panel (b) computes these statistics restricting only to hourly workers. Data is from PayrollCompany.

Table A-9: Heterogeneity by Industry and Occupation

Sample	All					Hourly			
	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Share $\Delta \neq 0$	Share salaried	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Share $\Delta \neq 0$
Industry									
Accommodation and Food Services	0.10	0.25	0.28	0.86	15%	0.13	0.27	0.30	0.97
Arts, Entertainment, and Recreation	0.09	0.26	0.31	0.74	33%	0.15	0.33	0.34	0.96
Management of Companies and Enterprises	0.07	0.22	0.27	0.72	35%	0.13	0.26	0.28	0.98
Retail Trade	0.06	0.18	0.25	0.74	32%	0.09	0.22	0.26	0.93
Administrative and Support and Waste Management and Remediation Services	0.05	0.16	0.24	0.72	35%	0.08	0.19	0.24	0.91
Construction	0.05	0.17	0.23	0.69	34%	0.09	0.20	0.24	0.89
Health Care and Social Assistance	0.05	0.17	0.25	0.74	26%	0.08	0.20	0.25	0.89
Manufacturing	0.05	0.16	0.23	0.76	32%	0.08	0.18	0.23	0.94
Transportation and Warehousing	0.05	0.17	0.24	0.74	35%	0.09	0.20	0.25	0.92
Utilities	0.05	0.17	0.24	0.76	31%	0.09	0.20	0.24	0.97
Mining, Quarrying, and Oil and Gas Extraction	0.04	0.14	0.20	0.64	38%	0.09	0.20	0.23	0.95
Agriculture, Forestry, Fishing and Hunting	0.03	0.15	0.22	0.61	55%	0.10	0.22	0.25	0.94
Educational Services	0.03	0.19	0.26	0.62	51%	0.13	0.30	0.32	0.91
Other Services (except Public Administration)	0.03	0.16	0.24	0.60	51%	0.10	0.24	0.28	0.92
Wholesale Trade	0.03	0.15	0.25	0.66	44%	0.07	0.18	0.23	0.91
Public Administration	0.02	0.12	0.21	0.60	57%	0.06	0.19	0.27	0.83
Real Estate and Rental and Leasing	0.02	0.12	0.23	0.58	41%	0.05	0.16	0.23	0.78
Finance and Insurance	0.01	0.13	0.25	0.55	66%	0.08	0.20	0.25	0.90
Professional, Scientific, and Technical Services	0.01	0.13	0.25	0.55	63%	0.08	0.20	0.26	0.91
Information	0.00	0.13	0.25	0.46	72%	0.10	0.25	0.30	0.87
Occupation									
Host	0.19	0.36	0.35	1.00	0%	0.19	0.36	0.35	1.00
Server	0.19	0.38	0.36	1.00	0%	0.19	0.38	0.36	1.00
Bartender	0.18	0.36	0.36	1.00	3%	0.18	0.36	0.36	1.00
Cook	0.11	0.23	0.27	0.99	5%	0.11	0.23	0.27	1.00
Operator	0.10	0.19	0.21	0.99	1%	0.11	0.20	0.21	1.00
Cleaner	0.09	0.20	0.24	0.99	0%	0.09	0.20	0.24	0.99
Driver	0.09	0.18	0.21	0.99	2%	0.09	0.18	0.21	1.00
Warehouse	0.09	0.18	0.19	0.99	1%	0.09	0.18	0.19	1.00
Welder	0.09	0.16	0.18	0.99	0%	0.09	0.16	0.18	0.99
Sales	0.08	0.30	0.35	0.75	64%	0.13	0.29	0.31	0.99
Mechanic	0.07	0.17	0.22	0.87	10%	0.07	0.18	0.23	0.92
Maintenance	0.06	0.13	0.19	0.89	9%	0.06	0.14	0.19	0.96
Medical assistant	0.06	0.15	0.19	1.00	1%	0.06	0.15	0.19	1.00
Porter	0.06	0.14	0.19	0.96	2%	0.06	0.13	0.18	0.96
Technician	0.06	0.17	0.24	0.83	25%	0.08	0.18	0.24	0.98
Customer service agent	0.05	0.14	0.20	0.88	19%	0.06	0.16	0.19	1.00
Administrative assistant	0.04	0.13	0.20	0.75	25%	0.05	0.15	0.21	0.87
Wage and salary administrator	0.04	0.11	0.20	0.74	31%	0.06	0.13	0.20	0.92
Analyst	0.02	0.10	0.23	0.63	82%	0.08	0.25	0.31	0.95
Accountant	0.01	0.09	0.18	0.55	97%	0.08	0.19	0.24	0.97
Consultant	0.00	0.17	0.27	0.47	87%	0.13	0.25	0.27	0.98
Engineer	0.00	0.06	0.23	0.45	85%	0.10	0.23	0.26	0.97
Financial controller	0.00	0.05	0.26	0.31	96%	0.11	0.42	0.41	1.00
Manager	0.00	0.09	0.24	0.47	83%	0.07	0.17	0.21	0.89
Teacher	0.00	0.06	0.18	0.44	97%	0.14	0.34	0.36	0.95

Notes: This table repeats the summary stats for “Total earnings” changes similar to Table 1 for the 20 largest industries and the 25 most common occupations. See Section A.1 for discussion of occupation definitions. The columns on the left compute volatility for all workers while the columns on the right restrict to only hourly workers. Data is from PayrollCompany.

Table A-10: AKM Variance Decompositions

Specification	Num Movers	Mean Vol	Median Vol	SD Vol	SD firm FE	SD worker FE	SD firm FE / SD worker FE	Cov(firm FE, worker FE)
Baseline	20319	0.14	0.11	0.11	0.048	0.048	1.00	0.00050
Only within Industry Moves	10128	0.14	0.11	0.11	0.048	0.052	0.92	0.00056
Only within Wage Decile Moves	12825	0.14	0.11	0.11	0.049	0.049	0.99	0.00045
Exclude Accom. and Food	15288	0.13	0.10	0.11	0.048	0.045	1.06	0.00048
25 Firm Groups	23675	0.14	0.11	0.11	0.049	0.047	1.04	0.00051
Hours Vol instead of Pay vol	20319	0.13	0.10	0.11	0.048	0.049	0.97	0.00046
Prime Age Only	16710	0.13	0.10	0.10	0.044	0.045	0.97	0.00037

Notes: This table reports variance decompositions using the fixed effect specification in Equation (5) with the leave-one-out sampling correction of Kline, Saggio, and Sølvesten (2020). The outcome is individual pay volatility except for the specification that looks at hours volatility instead of pay volatility. The baseline corresponds to the specification in the main text and other rows impose alternative restrictions on the estimating sample. Data is from PayrollCompany.

Table A-11: The Effect of Income Volatility on Consumption Volatility: Robustness

	Dependent Variable: Med [%C]							
	Nondurable (Baseline) (1)	Non-Work (2)	Other Nondurable (3)	Total Spend (4)	Nondurable (5)	Nondurable (6)	Nondurable (7)	Nondurable (8)
Med [%Y]	0.256*** (0.010)	0.371*** (0.020)	0.235*** (0.010)	0.265*** (0.008)	0.315*** (0.009)	0.240*** (0.012)	0.341*** (0.020)	0.260*** (0.027)
%Y Frequency	Monthly	Monthly	Monthly	Monthly	Quarterly	Monthly	Monthly	Monthly
%C Frequency	Monthly	Monthly	Monthly	Monthly	Quarterly	Quarterly	Monthly	Monthly
Jobs per household	1	1	1	1	1	1	1 or 2	1
Observations	889,379	879,777	889,375	889,379	625,496	625,496	977,734	433,910

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Notes: Data is from Chase. The unit of observation is a job-spell. Standard errors are clustered by firm. Non-Work expenses are spending categories where spending decreases less at retirement following Ganong and Noel (2019): “Discount Stores”, “Drug Stores”, “Airfare”, “Groceries”, and “Healthcare”. Column (3) uses the remaining nondurable spending excluding those categories. Total spend includes all account outflows except transfers to other financial accounts. Column (7) looks at the effect of household level income instead of job-level income. The data is sampled to match the proportion of one and two job households in the CPS and sub-sampled to keep the overall sample size close to the baseline specification. See Appendix E for definitions and details of the household level specifications. Column (8) restricts to firms with at least 50 workers.

Table A-12: The Effect of Income Volatility on Separation Rates

	Hourly	Salaried
$\overline{Vol}_{j(i)}^y$	3.01*** (0.240)	1.30*** (0.340)
No. Obs	112,836	55,925
No. Firms	7,143	7,143
Controls?	Yes	Yes

Notes: This table estimates a Cox proportional hazard model of separations on average firm volatility $\overline{Vol}_{j(i)}^y$: $H(t) = H_0(t) \times \exp \left[\beta_1 \overline{Vol}_{j(i)}^y + \gamma' X_{ij} \right]$ where $H(t)$ is the hazard function at spell tenure t relative to a baseline hazard function, $H_0(t)$ and X_{ij} includes controls for firm average wages, firm average hours, industry fixed effects and for a worker's gender and age at job start. This is the reduced form specification of the IV in Column 3 of Table 4. $\overline{Vol}_{j(i)}^y$ is average volatility of hourly workers. This regression is estimated separately for hourly and salaried workers and we only use firms that have both salaried and hourly workers.

Table A-13: Earnings Risk in Monthly Data versus Models Calibrated to Annual Data

	Monthly data	Model based on annual data			
		KMV (2018)	KV (2022)	MLM (2025)	CHT (2022)
P90 - P10 $ \Delta $	0.45	0.07	0.01	0.16	0.19
Share $ \Delta > 1\%$	0.64	0.30	0.11	0.87	0.89
Share $ \Delta > 20\%$	0.22	0.07	0.10	0.00	0.01
50th percentile $ \Delta $	0.05	0.00	0.00	0.04	0.05
75th percentile $ \Delta $	0.17	0.02	0.00	0.07	0.08
90th percentile $ \Delta $	0.39	0.13	0.16	0.10	0.12
Standard deviation	0.25	0.17	0.30	0.06	0.07
Kurtosis	6.46	31.18	13.81	3.01	3.01
Crow-Siddiqui kurtosis	14.07	102.57	434.55	2.91	2.91
Positive persistence	0.35	0.95	0.92	0.58	0.37
Negative persistence	0.32	0.95	0.83	0.58	0.37

Notes: This table shows summary statistics of monthly log earnings changes ($\Delta = \log y_t - \log y_{t-1}$) in PayrollCompany data and in several benchmark models of earnings processes which are calibrated to annual data. KMV is Kaplan, Moll, and Violante (2018), KV is Kaplan and Violante (2022), MLM is Maxted, Laibson, and Moll (2025), and CHT is Crawley, Holm, and Tretvoll (2026). Before computing higher-order moments (standard deviation, kurtosis), the measures of change in both the data and model distributions are winsorized at the 1st and 99th percentiles of nonzero changes in the PayrollCompany data. The Crow-Siddiqui Kurtosis is defined as $\frac{P_{97.5} - P_{2.5}}{P_{75} - P_{25}}$. Positive persistence is defined as the fraction of worker-months in which income increases conditional on income having increased in the prior month. Negative persistence is defined similarly.

Table A-14: Effects of Unpaid Leave on Earnings Changes

Variable	Share $\Delta \neq 0$	Median $ \Delta $	75th percentile $ \Delta $	Std. dev.	Share $\Delta > 0$	Share $\Delta < 0$
Δ: Change from First Difference						
Observed Earnings Fluctuations	0.89	0.06	0.15	0.16	0.46	0.43
Earnings Fluctuations After Removing Unpaid Leave	0.88	0.05	0.13	0.14	0.46	0.42

Notes: The “Observed Earnings Fluctuations” row shows actual earnings volatility for full-time hourly workers. The “Earnings Fluctuations After Removing Unpaid Leave” row shows what earnings volatility would be in a counterfactual where workers took no unpaid leave. We impute this counterfactual using our simulated allocation of unpaid leave hours. We assume that unpaid leave is equal to 2.33 percent of hours paid. The simulation assumes that unpaid leave is taken in the pay periods with the lowest paid hours.