

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

We thank Fiona Greig for the initial conversations which inspired this paper. We thank Niels Gormsen, Simon Jaeger, Emir Kamenica, Greg Kaplan, Anil Kayshap, Erik Hurst, Thibaut Lamadon, Susan Lambert, Marta Lachowska, Bentley MacLeod, Alan Manning, 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 Payroll- Company 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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

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 causal evidence that it increases consumption volatility and also leads to greater job separations, and we find that workers have a high willingness to pay to reduce earnings instability. These findings suggest that short-term earnings risk is a significant and previously underappreciated feature of the labor market.

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

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

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

Alexander Weinberg
University of Chicago
weinberga@uchicago.edu

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

1 Introduction

This paper uses high-frequency administrative earnings data to investigate the prevalence, causes, and consequences of monthly earnings volatility among U.S. workers. We find that beneath the surface of stable employment lies a hidden pattern of instability: for most workers, earnings vary substantially from month to month, even within continuing employment relationships. Although this monthly instability is pervasive, it is masked in the annual income data used most commonly in economics research. Moreover, this volatility is unequally distributed. Hourly workers, who tend to be lower income and more financially fragile, face much more volatility than salaried, higher-income workers. These fluctuations are in turn economically meaningful—shaping spending behavior, job mobility, and the lived experience of work for the 60 percent of the U.S. workforce that is paid hourly.

The paper establishes this conclusion in three parts. First, using comprehensive administrative data from both workers and firms, we document that monthly earnings fluctuate considerably, especially for hourly 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. Second, we explore why hours change from month to month: is this driven by workers choosing to vary hours, or is this something that is outside of their control? We find evidence that firm-driven labor demand changes are a key driver of hours fluctuations, while worker choices play a limited role. Third, we provide direct empirical evidence that these fluctuations impose substantial welfare costs on workers: we find that monthly earnings volatility causes consumption volatility and leads workers to leave their jobs in search of improved stability. Overall, our estimates suggest that a typical hourly worker would give up 4-11 percent of their income in exchange for the income stability experienced by salaried workers.

To measure earnings volatility, we use 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 allows us to measure detailed characteristics of individual pay, like hours, wages, and bonuses for the universe of workers at a given firm. The bank account data enables us to link earnings data with spending and liquidity. Throughout, we focus on pay variation within ongoing employment relationships. Including the additional volatility 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.

Our paper starts with simple descriptive patterns, with four main findings. First, we find that monthly earnings volatility is substantial. In about three quarters of months, workers receive a different amount of pay than they received 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, relative to various benchmark models of the earnings process used in the past literature, and relative to typical liquidity holdings.

Second, we find that earnings volatility differs markedly between salaried and hourly workers. 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. In contrast, salaried workers' pay rarely varies from month to month. When it does vary, it usually varies to the upside. This upside volatility is primarily driven by performance pay such as bonuses and commissions. Thus, the 60 percent of U.S. workers who are paid hourly have earnings dynamics that are completely different from the 40 percent of U.S. workers who are salaried. This crucial dimension of heterogeneity has received little attention in past studies of income dynamics.

Third, we find that there is a strong negative relationship between income levels and earnings volatility. Because low-income workers also have less liquidity, this suggests that the most financially fragile workers also face the greatest variability in pay. This income profile of earnings volatility arises primarily because low-income workers are more likely to work in hourly jobs.

Fourth, we find that while earnings changes are quite common, they are also fairly transitory. This is important for two reasons: first, optimal behavior in models of income risk will depend not just on the distribution of income changes but also on their persistence, so this is a useful target for modeling income dynamics. Second, this observation helps explain why the earnings instability prevalent in monthly data is hard to pick up in annual earnings data.

This descriptive analysis raises two natural follow-up questions which we explore in the remaining two parts of the paper: Why do hours move from month to month? Does this instability matter for worker welfare?

The second part of the paper explores the determinants of monthly hours fluctuations. We start by showing that several ex-ante plausible explanations have little support in the data. First, temporary unpaid leave (e.g. vacation, family leave, medical leave) is too small a share of overall hours to explain the high level of typical monthly changes. Second, we find no evidence that hours volatility is driven by childcare-related demands for flexibility. Third, predictable seasonal fluctuations explain little of the hours fluctuations in ongoing employment relationships.

In contrast, we provide evidence that firms play an important role in driving changes in worker hours. We show that this firm-driven volatility arises in part because firms have substantial month-to-month fluctuations in *total* hours, which we argue most likely arise from labor demand. We use two distinct methodological approaches to show that these monthly fluctuations in *total* firm hours are big enough to drive substantial volatility in hours for *individual* workers. The first builds on the excess reallocation statistic in Davis and Haltiwanger (1992) while the second estimates the effect of moving between firms with different levels of total hours volatility. These two distinct approaches both imply that around half of the volatility for a typical worker is caused by firm-level fluctuations in total hours. We also find evidence that firms contribute to worker volatility through channels that go beyond fluctuations in total labor demand. Even if total hours at the firm level were constant from month to month, there could still be significant monthly hours volatility for individual workers, and firms could contribute to this volatility (for example, through their scheduling practices). To capture the full set of pathways through which firms can affect the volatility for their workers we estimate fixed effects specifications with movers, building on the model of firm effects in Abowd,

Kramarz, and Margolis (1999). These fixed effects imply an even larger role of firms in driving individual fluctuations than when focusing on only fluctuations in total labor demand.

Overall, this evidence pointing to a large role for firms in driving worker-level earnings volatility leads us to conclude that the volatility we document in the first part of the paper does not simply reflect desired choices by workers. Instead, it appears that a meaningful share of this instability is imposed on the worker and is outside their control. However, even if all earnings changes were exogenous from the worker’s perspective, there are still two *theoretical* reasons to think they may not be welfare relevant. First, these earnings changes are relatively transitory, and many theoretical models imply that transitory shocks should not matter much for welfare. Second, if workers have information about their hours changes in advance, they may be able to blunt any potential negative impacts of this instability.

In the final part of the paper, however, we provide two pieces of *empirical* evidence that monthly earnings instability is indeed welfare relevant. First, using bank account data, we show that income volatility causes spending volatility. Our primary strategy instruments for individual earnings volatility using firm-level volatility and finds that spending volatility increases when workers move to higher-volatility firms. Heterogeneity analysis strengthens the causal interpretation: effects are larger for low-liquidity workers, consistent with binding budget constraints. We also find that income volatility for salaried workers is essentially irrelevant for their spending volatility. This contrast further highlights that the nature of monthly earnings volatility differs fundamentally between salaried and hourly work.

The second complementary piece of evidence that this volatility is costly is that hourly workers are more likely to quit high-volatility jobs. To address concerns that this correlation is driven by marginally attached workers who are more likely to quit also choosing more variable hours, we instrument for individual volatility with firm-level volatility and control for time-invariant worker characteristics. We also show that firm volatility affects quit rates of hourly workers much more than salaried workers at the same firms. This suggests that volatility itself—rather than some other firm characteristic—is the primary driver of increased quits.

Using back-of-the-envelope calculations, we translate our estimates of the effects of hours volatility on both spending and separations into an implied willingness to pay to avoid such volatility. For the spending estimates, we apply the Lucas (1987) formula to infer how much a worker would pay to eliminate the spending fluctuations induced by hours volatility. For the separations estimates, we use Gronberg and Reed (1994) and compare the elasticity of quits with respect to that disamenity (in our case, earnings volatility) to the elasticity of quits with respect to wages. In both cases, we find substantial welfare implications: a median hourly worker would be willing to forgo 4–11 percent of their income to attain the more stable income of a median salaried worker.

Because earnings instability has a significant welfare cost and it is concentrated among lower-income workers, this disamenity amplifies compensation inequality relative to wage inequality. One estimate of the importance of amenities for compensation inequality comes from Maestas et al. (2023), which estimates the willingness to pay for nine different job amenities. It finds that the

90-10 compensation differential—inclusive of amenities—is 8 log points larger than the 90-10 wage differential. We show that incorporating just one additional (dis)amenity—earnings instability—raises this gap from 8 to 13 log points.

The fact that workers adjust their spending in response to income volatility and actively leave more volatile jobs provides compelling evidence that these fluctuations are costly. Such behavior is inconsistent with monthly volatility stemming solely from measurement error, deliberate labor supply choices, or other non-risk forces that render them irrelevant. Instead, our findings reveal that these earnings fluctuations represent genuine, welfare-relevant risk that materially shapes household decisions.

1.1 Related Literature

Our paper contributes to several existing streams of research. There is a vast prior literature measuring income risk. One large branch of this literature uses data from the Panel Study of Income Dynamics (PSID).¹ Moffitt and Zhang (2018) provides an overview of many of these PSID studies. A second recent branch instead uses large administrative data sets. This work includes Guvenen, Ozkan, and Song (2014, social security data), Abowd and McKinney (2024, Longitudinal Employer-Household Dynamics data) and Pruitt and Turner (2020, IRS tax data). Other papers link various data sets together including Kniesner, Viscusi, and Ziliak (2006) and Moffitt et al. (2022). A distinguishing feature of this literature is that it mostly focuses on the volatility of *annual* income. In contrast, our study highlights the importance of *monthly* income fluctuations, which we find are largely masked in the annual data.

Although there are few direct measures of high-frequency earnings risk, such risk is often an important input to models of household behavior. For example, in leading heterogeneous agent macro models the degree of high-frequency earnings risk determines the household’s optimal level of liquid assets, which in turn drives the sensitivity of consumption to transitory shocks. Without access to panel data on within-year earnings risk for U.S. workers, the literature has been forced to infer high-frequency earnings dynamics from low-frequency annual data. Since those models are estimated using only annual data, any identification of high-frequency earnings dynamics comes from imposing parametric restrictions on the underlying data generating process. Recent examples using this methodology include Kaplan, Moll, and Violante (2018), Maxted, Laibson, and Moll (2025), Crawley, Holm, and Tretvoll (2022), and Kaplan and Violante (2022). We show that the actual level of monthly income risk we measure in the data has much more frequent sizable changes than what is inferred from annual data using this approach. Even though we show that these monthly income changes are costly, their low persistence makes them hard to detect in annual data. The high-frequency income moments we report can be used to directly calibrate future income models rather than relying on this indirect, parametric approach.²

¹Some prominent examples using PSID data include Dynan (2012), Gottschalk et al. (1994), Haider and Loughran (2001), and Meghir and Pistaferri (2004).

²We note that these parametric models are typically fully identified from annual moments, so jointly matching high and low-frequency income dynamics will require introducing new “shocks” rather than simply re-calibrating

We also show that there is substantial heterogeneity in monthly volatility. Modern macroeconomic models with earnings risk incorporate increasingly rich sources of heterogeneity, and our moments can directly inform how to model such heterogeneity. Most prominently, we show that there is a large difference in income volatility between hourly and salaried workers, suggesting that this is an important margin to include in future heterogeneous agent models.³

Although most prior empirical estimates of earnings volatility rely on annual data, we complement two types of prior studies that also analyze higher-frequency data. First, there are two studies of monthly earnings volatility using administrative data from outside the US. Druedahl, Graber, and Jørgensen (2023) estimates a monthly income model using Danish administrative data, and Brewer, Cominetti, and Jenkins (2025) documents monthly pay volatility in UK administrative data. The descriptive analysis in the first part of our paper overlaps in several ways with these papers: for example, Druedahl, Graber, and Jørgensen (2023) finds a high frequency of monthly earnings changes, although lower than we find in our U.S. data. This likely reflects the fact that there are fewer hourly workers in Denmark. Brewer, Cominetti, and Jenkins (2025) builds on an earlier working draft of our paper and finds some similar patterns of heterogeneity across industries and contract types to those that we document. However, these papers do not try to distinguish whether fluctuations are driven by workers or firms and they do not quantify the costs of these fluctuations.

Second, a number of surveys estimate monthly volatility in the US. In particular, we build most closely on the pioneering work of Hannagan and Morduch (2015) and Morduch and Schneider (2017). This work surveys 235 low- and moderate-income households on a monthly basis for one year and documents striking income instability. Motivated in part by these findings, many U.S. household surveys have recently added questions measuring the extent of month-to-month earnings volatility or week-to-week hours volatility, indicating growing interest by policymakers and statistical agencies in this topic.⁴ However, volatility is difficult to measure in these surveys because it relies on worker recall and surveys often only allow respondents to choose among a limited number of categories.

Relative to these two prior sources of sub-annual volatility estimates, we contribute the first systematic and representative estimates of monthly income volatility for U.S. workers. Furthermore, beyond documenting facts about the level and distribution of monthly volatility, we also provide causal estimates of the sources and consequences of this volatility.

In tandem with the growth of survey evidence on the extent of income volatility, there has also been work in sociology and operations research on the welfare consequences of earnings volatility. Lambert, Henly, and Kim (2019) show that for hourly workers, the magnitude of volatility in weekly

existing models. High-frequency hours shocks are a natural candidate.

³We also provide estimates of monthly income volatility in different industries and occupations, by income, by gender and family structure, and by age.

⁴See the Federal Reserve Board’s Survey of Household Economics and Decisionmaking, the Census Bureau’s Current Population Survey Underbanked Supplement, and the Bureau of Labor Statistics’ National Survey of Longitudinal Youth. For example, the Underbanked Supplement asks workers whether over the last 12 months their income varied a lot from month to month, varied somewhat from month to month, or was about the same each month.

work hours is negatively correlated with their perceptions of financial insecurity. Schneider and Harknett (2019) find that for hourly retail workers, routine schedule instability is associated with worse psychological distress, sleep quality, and happiness. Other work has found quasi-experimental evidence that pay volatility makes workers more likely to quit in the retail (Kesavan and Kuhnen 2017) and home health (Bergman, David, and Song 2023) sectors. Relative to this prior work, we estimate a causal effect of pay volatility and quantify its welfare consequences using economic models.

Although research in economics on the welfare consequences of earnings volatility is limited, labor economists have studied income volatility in the context of a literature examining how firms choose to structure workers’ compensation. This literature typically focuses on when and how to offer performance pay, which gives rise to a tradeoff between risk and incentives.⁵ Because performance pay is mostly offered to salaried workers, workers who are paid hourly are often interpreted as facing the *least* risk. For example, MacLeod and Parent (1999) treats hourly workers as having deterministic pay while salaried workers in many cases are classified as having stochastic pay. Our finding that hourly workers face meaningful income risk arising from hours variation shows that earnings risk is more complicated and in some important dimensions hourly workers have *more* risk than salaried workers.

Our estimates of the disamenity value of earnings instability contribute to a rapidly growing literature on revealed preference estimates of non-wage amenities. Researchers in this area have used discrete choice experiments, offer data, worker flows, equilibrium models, and job separations in the model of Gronberg and Reed (1994) to estimate the value of amenities.⁶ We implement the separations approach—which has been used in several prior papers—in our data. We also use a structural approach based on Lucas (1987) which is feasible only because we are studying a job amenity that affects workers’ consumption. To the best of our knowledge this latter approach is novel to the amenities literature.⁷

Our finding that earnings instability is highest among the lowest-income workers contributes to the broader literature on labor market inequality (e.g., Katz and Autor 1999). Much of this research investigates the role of firms in shaping wage inequality, often using job movers to identify firm-specific pay premiums.⁸ While wages are the most visible component of compensation, they do not fully capture all welfare-relevant aspects of employment. For example, Lachowska et al. (2023) find that most workers are “mismatched” in the sense that they would prefer to work more hours than they are offered at their current employer. One common question for this literature is how

⁵See Lazear (1986), Hart and Holmström (1987), MacLeod and Parent (1999), Prendergast (2002), Lemieux, MacLeod, and Parent (2009), and Lemieux, MacLeod, and Parent (2012).

⁶Mas (2025) reviews these approaches.

⁷However, it is possible that using consumption data could be helpful for valuing other employer-provided amenities that can be thought of in the same risk reduction framework, like health insurance. For example, Finkelstein, Hendren, and Luttmer (2019) show that one way to estimate the “pure insurance” value of health insurance is using data on non-health consumption.

⁸See Abowd, Kramarz, and Margolis (1999), Bonhomme, Lamadon, and Manresa (2019), Card, Heining, and Kline (2013), Engbom and Moser (2022), and Song et al. (2019).

incorporating amenities changes the distribution of total compensation.⁹ We find that accounting for the additional (dis)amenity of pay volatility meaningfully increases the contribution of amenities to compensation inequality.

Finally, our finding that hourly workers are exposed to a significant amount of risk that is driven by firm labor demand shocks is related to a literature on rent-sharing in labor economics. Card et al. (2018) summarize this literature as finding elasticities of annual wages to firm’s value added per worker between 0.05 and 0.15. Within this literature, we are particularly related to the analysis in Guiso, Pistaferri, and Schivardi (2005) who study the response of wages to transitory firm productivity shocks. Although we study a different shock and country, our finding that hours fluctuate much more than wages in response to monthly firm labor demand shocks is consistent with the limited pass-through of transitory firm productivity shocks to wages that they find in Italy.

2 Data Description and Measuring Volatility

We use two main datasets to study monthly earnings volatility. The primary dataset comes from a payroll processor. Its greatest strength is that we observe many details about employees’ paychecks. We also analyze bank account data from the JPMorganChase Institute where we can link income with consumption.

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 4 million workers in the data at any time. Most PayrollCompany clients are small firms, with the median worker employed in a firm with 19 employees. However, we show that earnings volatility is nearly identical when using bank account data, which captures a more representative firm size distribution.

For each paycheck, we observe regular pay, bonuses, commissions, overtime, and paid leave. In addition, almost all salaried workers have a default 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.

We perform most of our analysis at the monthly level. We do not analyze within-month pay variation for two reasons. First, different workers are paid on different schedules (weekly, fortnightly, semimonthly, or monthly) within each month. The shortest time interval which encompasses all of these different pay schedules is the month. Second, we are particularly interested in the economic

⁹See Caldwell, Haegle, and Heining (2025), Diamond (2016), Humlum, Rasmussen, and Rose (2025), Mas and Pallais (2017), Mueller, Ouimet, and Simintzi (2017), Sockin (2022), and Sorkin (2018).

consequences of pay fluctuations. Many forms of spending recur at monthly frequencies and so monthly volatility is of greater potential economic relevance. Although our main focus is on monthly pay volatility, we also analyze volatility at the quarterly frequencies more commonly studied in business cycle analysis.

Given our focus on monthly pay volatility within ongoing employment relationships, we purge the data of two types of measured monthly volatility that fall outside this definition. First, we use average pay per paycheck because total monthly pay fluctuates due to calendar timing for those workers paid weekly or fortnightly. For example, in each calendar year, workers who are paid weekly get 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 month to month. Second, we exclude monthly volatility arising from partial months of employment. Specifically, we define a job spell as the continuous series of months with positive earnings. We then exclude the first and last month of each worker’s job spell because lower pay per paycheck during these months could arise from partial employment during that pay period.

We exclude workers whose contract type (hourly versus salaried) is ambiguous. Specifically, we drop workers without base wage or salary data (11 percent of worker-months) and those whose hourly wage changes in more than half of their months worked (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 resulting in an ultimate sample of 19,893 firms.

Despite focusing on small firms, the PayrollCompany data appears to be representative of the U.S. workforce along several dimensions. In terms of wages, Figure A-1 shows that the wage distribution for PayrollCompany workers is similar to the wage distribution for all workers. Figure A-2 shows that the distribution of hours worked per quarter for PayrollCompany workers is similar to that in government administrative records. Figure A-3 shows that the distribution of pay frequency is similar to administrative benchmarks, as is the share of hourly versus salaried and the share that get bonuses (Table A-1). Both in our sample and in representative benchmarks, 60 percent of workers are hourly and 40 percent are salaried. Figure A-4 shows that seasonality of aggregate employment in PayrollCompany is similar to that of aggregate private employment in BLS data. Finally, Figure A-5 shows that the monthly separation and hire rates are also similar in PayrollCompany and JOLTS data.

In addition to the baseline sample described above, we also draw a second sample of firms in which we observe job transitions. We can follow workers who transition between two firms that are both clients of the PayrollCompany, and so we start by constructing a random sample of 1,000 firms in which we observe between 8 and 30 such transitions.¹⁰ We then combine this set of 1,000 firms with all of other firms that are directly linked to this set through at least one move by an

¹⁰The minimum value of 8 selects a sample with a non-trivial number of transitions, and the maximum value of 30 is for computational feasibility since we must draw data on all workers at firms connected by these moves. Sampling a random subset of 1,000 firms meeting this criteria is also for computational feasibility.

hourly worker, resulting in a total sample of 7,720 firms linked by at least one move. Finally, we then sample *all* of the hourly workers at these 7,720 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 4.2.3 and 5.2.

Finally, we note that even with the various data cleaning steps discussed above, it is possible that some of what we will label as earnings volatility is spurious and reflects measurement error. For example, the procedure outlined above would find volatility for a worker who is paid weekly and works every other Friday. But such a pattern would not reflect meaningful earnings instability from the worker’s perspective. However, in Section 5, we demonstrate that earnings volatility affects spending and quits behavior. Spurious earnings volatility driven by measurement error would have no effect on spending or quits, so the analysis there demonstrates that this potential measurement error is not the dominant driver of earnings volatility.

2.2 Bank Records

We supplement the PayrollCompany data with data on Chase bank customers from the JPMorganChase Institute (JPMCI). The sample that we use within the JPMCI data is the same sample used in Ganong et al. (2023) and we refer readers to that paper for details on the data construction and sample selection.

The main purpose of the JPMCI data is to link data on income with data on spending and liquidity. This work builds on prior JPMCI reports examining links between income and expense volatility (Farrell and Greig 2015; Farrell, Greig, and Yu 2019). The JPMCI data is also useful for addressing two limitations of the PayrollCompany data: the JPMCI data include employees at large firms and include income from multiple jobs, enabling us to analyze 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 an encrypted identifier for the firm. The firm identifier enables us to tag the transaction as labor income and to identify other workers with Chase bank accounts who are paid by the same employer. Because it reflects what is actually received by the worker, this dataset captures net earnings. 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. Because we do not observe details about the job contract, we predict whether the job is hourly or salaried based on characteristics of their pay stream. Specifically, we impute that a worker is paid hourly during that job spell if pay is changing more than 0.01 percent in at least 70 percent of months and average pay per week is less than two thousand dollars.¹¹

We also observe spending and liquidity for each account. Our main measure of spending captures

¹¹We 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.

spending on nondurable goods and services, constructed as in Ganong et al. (2023). Examples of nondurable spending include groceries, food away from home, fuel, utilities, clothing, medical co-pays, and payments at drugstores. Spending is measured from debit and credit card transactions, cash withdrawals, and electronic transactions captured through the bank account. We measure liquidity as the checking account balance at the beginning 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 monthly earnings per paycheck.¹² We winsorize this variable at the 2.5th and 97.5th percentile of non-zero changes to limit the role of outliers.¹³

The first goal of our paper is to describe the distribution of these monthly changes $\% \Delta_{i,t}$. To achieve this goal, 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). The main benefit is that it pools a large number of observations together, and so can flexibly capture the full distribution of changes. If workers are drawing earnings changes from some common distribution, then this cross-sectional distribution teaches us about the earnings process faced by individual workers. However, this is a strong assumption, which is not well-suited for answering questions about how individual workers' volatility affects their behavior.

Therefore, when focusing on individual heterogeneity, we instead summarize volatility as $Vol_i = Median(|\% \Delta_{i,t}|)$, where the median is taken across all of the monthly changes for a worker i at a particular job spell. We focus primarily on the median absolute change since this measure captures the typical earnings change faced by a worker and is robust to outliers. In Section 3.2, we further discuss why we prefer to measure volatility using the median instead of the standard deviation. To distinguish between pooled measures and individual-specific measures, we index individual volatility statistics with subscript i . For example, $Vol_i = Median(|\% \Delta_{i,t}|)$ refers to an individual-level statistic, the median absolute percent change for worker i .

3 Facts About Monthly Earnings Volatility

This section describes the empirical patterns of monthly earnings volatility. Our first and main finding is that monthly earnings volatility is substantial in both payroll and bank account data. Second, we find that volatility is higher for hourly workers than salaried workers. Third, we find

¹²By construction, $\% \Delta_{i,t}$ will only capture changes across months with non-zero earnings. Including months with zero earnings would only amplify the conclusion that earnings volatility is large.

¹³Table 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	0.70	0.05	0.17	0.31	3.28	18.07
All	Base Wage	0.10	0.00	0.00	0.05	8.40	168.21
B. Alternative Measures of Earnings							
All	Quarterly Earnings	0.82	0.06	0.15	0.24	2.43	12.13
All	Net Earnings	0.77	0.05	0.18	0.30	3.09	16.33
JPMCI	Net Earnings	0.74	0.05	0.17	0.30	3.11	16.78
C. Hourly vs Salaried							
Hourly	Total Earnings	0.92	0.09	0.21	0.30	2.77	14.58
Hourly	Base Wage	0.12	0.00	0.00	0.04	8.11	175.98
Hourly	Hours	0.90	0.07	0.19	0.25	1.90	7.76
Salaried	Total Earnings	0.35	0.00	0.06	0.32	4.02	22.91
Salaried	Base Wage	0.07	0.00	0.00	0.05	8.17	143.63
D. Hourly Subsamples							
Full-time Hourly	Total Earnings	0.90	0.06	0.15	0.19	1.56	6.26
Prime-age Hourly	Total Earnings	0.74	0.05	0.17	0.30	3.28	18.54

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.

that volatility is highest for low-wage workers (who are often paid hourly). Fourth, we find monthly earnings volatility is mostly transitory.

3.1 Earnings Are Volatile Within Ongoing Employment Relationships


Figure 1 shows that workers face substantial earnings fluctuations from month to month. Panel (a) shows the cumulative distribution function of monthly earnings changes while panel (b) shows the corresponding histogram. Summary statistics from this distribution are shown in panel A of Table 1. In almost three quarters of months, workers receive a different amount of pay than they received the prior month. Moreover, these earnings changes are often substantial. In the median month, the change in pay is 5 percent. In one quarter of months, the change in pay is at least 17 percent.¹⁴

Monthly income volatility within jobs in the JPMCI checking account data is nearly identical to that in the PayrollCompany. Specifically, Figure 2 shows the distribution of monthly changes in net pay across the two datasets and panel B in Table 1 summarizes the moments of the distribution.

¹⁴These exhibits combine data from many workers, some of whom might have larger earnings risk than the average worker and others of whom might have no earnings risk at all. Figure A-6 shows the distribution of individual-level volatility Vol_i . The figure shows that the level of individual volatility is high for the typical worker. It also shows clear evidence of heterogeneity across workers. We investigate this heterogeneity in more detail using a variance decomposition approach in Section 4.2.3.

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.

Since the JPMCI data includes workers at both small and large firms, this pattern demonstrates that monthly pay volatility is prevalent across the firm size distribution. Indeed, Table A-3 shows

that the volatility of earnings changes is similar across employees of small and large firms. Moreover, we find that within a household, earnings changes in one job are not offset by earnings changes at other jobs: Table A-3 shows that the volatility of earnings is almost identical at the individual job level and when summing earnings across all jobs at the household level.

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, the monthly income changes in the PayrollCompany data are also substantially larger than what the past literature has inferred from annual data. Figure 3 compares the distribution of monthly earnings changes in the data to that inferred from several high-frequency income models that are calibrated to target previously available annual income moments. We compute monthly statistics from the continuous-time Kaplan, Moll, and Violante (2018), Maxted, Laibson, and Moll (2025) and Crawley, Holm, and Tretvoll (2022) models, as well as a monthly version of the discrete-time income process from Kaplan and Violante (2022). Table A-4 provides summary statistics from these distributions. Sizable monthly earnings changes are much more common in the data than what is implied by any of these models. For example, the 75th percentile of the absolute percent change in earnings is between 1 percent and 9 percent in the models (as compared to 17 percent in the data), and the 90th percentile is between 10 percent and 13 percent in the models (as compared to 39 percent in the data). See Appendix B for additional details on these models and additional comparisons of model moments versus actual monthly data.


Third, monthly income changes are large relative to typical liquidity levels. Using the JPMCI data, for each household we compare monthly changes in income to the median value of their checking account balance. We find that in one-third of months, households have an absolute dollar change in income which exceeds 50 percent of their median checking account balance, and in one-fifth of months they have an absolute change which exceeds the median balance.

3.2 Hourly Workers Face More Volatility Than Salaried Workers

This section shows that the nature of earnings volatility differs markedly between hourly and salaried workers in terms of its level, its sources, and its (a)symmetry.


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 experience earnings changes in 92 percent of months—meaning that in 11 out of 12 months, they earn a different amount than the prior month. These changes are sizable: the median absolute change is 9 percent, and the 75th percentile is 21 percent. Table 1, Panel D shows similar patterns for full-time and prime-age hourly

Figure 2: Earnings Volatility in PayrollCompany compared to JPMCI



Notes: This figure shows the distribution of changes in net earnings (i.e. pay after taxes and other deductions) in PayrollCompany vs. JPMCI. PayrollCompany and JPMCI data are analyzed separately and were not merged as part of this analysis.


Figure 3: 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 (2022). See Appendix B for additional details.

workers, revealing that these differences are not simply due to weaker labor force attachment among hourly workers. In contrast, salaried workers' earnings are far more stable: the median absolute

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.

change is zero, and the 75th percentile is only 6 percent.

Differences in volatility between hourly and salaried workers are even larger when compared to liquidity. Using JPMCI data, we find that in over 40 percent of months, hourly workers experience earnings changes exceeding half their median checking account balance, compared to 23 percent for salaried workers. In 26 percent of months, the change exceeds their entire median balance, versus 13 percent for salaried workers. Thus, hourly workers not only face more volatility but are less buffered against it. The reason hourly workers face substantial pay volatility is because 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 hours—and changes in hours in earnings for hourly workers—are 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-5 shows that hourly workers experience increases and decreases with roughly equal frequency.

Earnings changes for salaried workers, in contrast, are mostly increases off a stable base. It appears that these earnings changes are most likely to be bonuses, performance pay, and commissions. Unfortunately, the tagging of such payments is incomplete in the PayrollCompany data so we cannot simply calculate volatility for salaried workers excluding bonus payments.¹⁵ Neverthe-

¹⁵When an employer wants to pay a bonus they can either temporarily increase base pay or they can enter a separate pay item labeled “bonus.” The tax treatment of bonuses is identical to that of base pay, so this choice of

less, four patterns are consistent with the conclusion that bonuses are driving volatility for salaried workers. First, many of these payments *are* tagged as bonuses and commissions. Second, both base pay and bonuses for salaried workers surge by roughly equal dollar amounts in December. This seasonal pattern suggests that half of bonus payments might not be labeled as such. Third, among salaried workers, volatility is highest among *high*-earning workers; Lemieux, MacLeod, and Parent (2009) shows that these are the workers most likely to receive performance pay. Fourth, volatility for salaried workers consists entirely of one-time payments with zero persistence, as we show below.

While these bonus payments are infrequent, they are also large. This means that even though pay rarely changes for salaried workers, when it does change, it changes by very large amounts. This explains why the standard deviation, skewness and kurtosis of monthly earnings changes are all moderately *higher* for salaried workers than for hourly workers. For example, a salaried worker with stable pay who receives a once-a-year bonus equal to one month’s salary has a standard deviation of 29 percent, the same as a worker who had 12 draws of monthly hours from a normal distribution with standard deviation of 29 percent. However, we do not view these outlier-sensitive measures as the most informative measure of earnings risk, because in Section 4.1 we show that these pay changes for salaried workers are fairly predictable, and in Section 5 we show that salaried workers’ spending is insensitive to this income volatility. Thus, we instead focus on Median $|\Delta_{i,t}|$ as our primary measure of earnings instability. That said, our main results are robust to using the standard deviation in place of the median.¹⁶

However, it is still important to note that even though hourly workers face much more frequent and less predictable fluctuations, this does not imply that salaried workers face *no* monthly income risk. Furthermore, both hourly and salaried workers would still face annual income risk even if their pay changed only once per year.

3.3 Volatility Is Pervasive but Is Highest for Low-Income Workers

The previous sections establish that the general level of pay volatility in the economy is high (Figure 1), and that this is especially true for hourly workers (Figure 4). In this section we show that low-income workers are much more exposed to pay volatility. Looking across all workers, Table 2 shows that volatility declines with income. This is true regardless of whether we look at the share of months with a change in pay, the median absolute change in pay, or the 75th percentile of the absolute change in pay.¹⁷ This income heterogeneity arises primarily because the share of workers that are salaried rises with income.

When focusing just on hourly workers though, we find that pay volatility is pervasive across many different dimensions of heterogeneity. For example, Table 2 shows that volatility remains

how to record bonuses is entirely at the discretion of the employer.

¹⁶Specifically, our findings on the relationships between firm-level and individual volatility (Section 4.2), volatility and spending (Section 5.1), and volatility and quits (Section 5.2) all hold using the standard deviation.

¹⁷The one exception to this overall pattern comes from highly-paid salaried workers, who have a higher standard deviation of the change in pay than any other group of workers, because they are more likely to be compensated with bonuses (Lemieux, MacLeod, and Parent 2009).

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	14%	–	\$229	0.83	0.11	0.28	0.36
Q2	20%	–	\$565	0.80	0.05	0.14	0.23
Q3	42%	–	\$925	0.66	0.03	0.12	0.20
Q4	79%	–	\$2583	0.49	0.00	0.13	0.39
Hourly							
Q1	0%	26	\$288	0.94	0.11	0.25	0.33
Q2	0%	33	\$467	0.93	0.09	0.21	0.30
Q3	0%	36	\$660	0.92	0.08	0.18	0.27
Q4	0%	35	\$1087	0.87	0.07	0.18	0.29
Salaried							
Q1	100%	–	\$479	0.29	0.00	0.02	0.19
Q2	100%	–	\$986	0.34	0.00	0.04	0.17
Q3	100%	–	\$1518	0.35	0.00	0.05	0.20
Q4	100%	–	\$4182	0.42	0.00	0.19	0.52

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. This table shows a subset of the statistics from Table 1. Statistics on persistence, skewness, and kurtosis are shown in Table A-6. Hours and pay are shown as weekly averages.

high for even the highest-paid hourly workers. Table A-7 shows that while volatility is highest for the under 25 age group and declines rapidly with age, this is because most workers under 25 are paid hourly while a substantial share of workers over 25 are salaried. When focusing just on hourly workers, Table A-7 shows that even the oldest workers have substantial volatility. The same pattern arises when we analyze heterogeneity in volatility by industry and occupation, as shown in Table A-8. When looking at all workers, Median $|\Delta|$ is highest in Accommodation and Food Services and the occupations typically associated with that industry. Volatility is lowest in Finance and Insurance, Professional and Scientific Services, and Information and the occupations typically associated with those industries. However, this heterogeneity is primarily driven by differences in the share of salaried work. When we narrow our focus to hourly workers, we find substantial volatility in all industries and occupations.

3.4 Monthly Earnings Changes Are Not Very Persistent


To explore persistence of these monthly changes, we calculate the variance ratio, defined as

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

where y is defined as total monthly earnings. If the monthly income data is described by a random walk and all monthly changes in earnings are permanent, then the variance ratio will be one as all shocks will accumulate over time. Alternatively, if the monthly income data is described by white noise and all shocks are purely transient, then the variance ratio will quickly converge to zero.

Figure 5 shows the empirical variance ratio for hourly and salaried workers. For salaried workers,

Figure 5: The Persistence of Monthly Earnings Changes



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

the variance ratio is close to white noise. This is what we would expect to see if the variance of monthly earnings for salaried workers was driven by bonuses, commissions, and other one-time payments.

Earnings changes for hourly workers are more persistent than those of salaried workers, although they still die off within the year. Estimating an AR(1) process to target this variance ratio yields a coefficient of 0.57, implying that there is some non-trivial persistence of earnings changes across months, but persistence at annual frequencies is nearly zero ($0.57^{12} = .0011$). This finding helps to explain why economic models which rely on annual data alone do not detect these monthly earnings changes.

4 Channels That Do and Do Not Explain Earnings Volatility

Having documented evidence of substantial earnings fluctuations, we now ask *why* earnings are so volatile. Since the descriptive patterns in the previous section demonstrated that volatility for salaried workers was primarily from bonuses while volatility for hourly workers was primarily from hours, we focus in the remainder of the paper primarily on hourly workers, asking specifically why hours are so volatile month to month. We first examine three specific channels which could, in principle, contribute meaningfully to volatility—predictable seasonal variation, leisure-related labor supply choices, or childcare needs—but find that these channels are quantitatively unimportant. We then explore a channel which we do find to be important: firms. They play a large role in driving monthly hours volatility both through high-frequency changes in total firm hours and

through idiosyncratic hours volatility (e.g., management practices).

4.1 Channels That Are *Not* Important for Explaining Earnings Volatility

Seasonality We explore the importance of seasonal or regularly recurring sources of earnings volatility, such as regular performance bonuses and seasonal fluctuations in labor demand. While we previously documented in Figure A-4 that there is substantial *aggregate* seasonality in PayrollCompany, in this section we are interested in measuring how much seasonality explains the *individual* earnings volatility within ongoing employment relationships shown in Section 3.1.¹⁸ To allow for the possibility that seasonal patterns differ across firms or across individuals, we run regressions of the following form:

$$\log y_{i,j,t} - \log y_{i,j,t-1} = \beta X_{i,j,t} + \epsilon_{i,j,t}. \quad (2)$$

where $y_{i,j,t}$ is worker i 's total earnings per paycheck at firm j in month t , $X_{i,j,t}$ is a vector of covariates meant to capture seasonality, including a constant term, and $\epsilon_{i,j,t}$ is an error term. We study two different sets of predictors $X_{i,j,t}$. The first definition of $X_{i,j,t}$ captures firm-specific seasonality in pay by estimating firm by month fixed effects $\alpha_{j,m(t)}$ where $m(t)$ is an integer from 1 to 12 (e.g., January 2011 and January 2012 both have $m(t) = 1$). The second definition of $X_{i,j,t}$ captures annually recurring pay by using a 12 month lag of individual pay changes interacted with month fixed effects: $\alpha_{m(t)} + \beta_{m(t)}(\log y_{i,j,t-12} - \log y_{i,j,t-13})$.¹⁹ Note that this second definition uses no firm-specific information at all, just the worker's own pay change from a year ago. To reduce the potential for overfitting in the first definition, we focus on a sample of larger firms with an average of at least eight employees per month.

Table A-9 shows that seasonality explains little of the variation in pay changes for continuing hourly workers, with R^2 estimates between 0.03 and 0.13. We find similarly small R^2 estimates for the half of salaried workers who do not receive bonuses. However, some types of pay changes are more seasonal. First, when we look at the half of salaried workers who do receive bonuses, we find R^2 estimates between 0.25 and 0.39. Second, while hours changes for continuing workers are not very seasonal, total firm employment is much more seasonal – the R^2 for total firm hours changes (i.e. adding changes in total hours coming from hires and separations) is about twice as high as the R^2 for just continuing workers. This pattern is consistent with firms responding to seasonal demand fluctuations with changes in employment but responding to demand shocks with changes in hours for existing workers, and it also helps explain why aggregate employment is highly seasonal, even though individual earnings are not. Additional detail on seasonality results and related robustness is discussed in Appendix C.1.

Unpaid leave We calculate whether unpaid leave—which could reflect vacation, medical leave, or caring for family—could meaningfully contribute to the measured earnings volatility for hourly

¹⁸The fact that we find slightly more aggregate seasonality in PayrollCompany than in BLS data suggests that our data will, if anything, overstate the role that seasonality plays in earnings fluctuations.

¹⁹We have also explored combinations of these seasonality specifications and reach similar conclusions.

workers. Since we do not directly observe unpaid leave in the payroll data, we instead use a three-step procedure which combines representative survey data with the payroll data.

First, we estimate the average amount of unpaid leave taken by US workers. We focus on full-time workers because unpaid vacation is an ill-defined concept for part-time workers. The American Time Use Survey in 2017 asks workers if they took any leave from their job over the past seven days, how many hours of leave they took, and whether it was paid or unpaid. Full-time hourly workers report taking unpaid leave equal to 2.33 percent of their usual hours worked. Aggregating annually, this implies 6.53 days per year for the average worker.

Second, we impute when this unpaid leave was taken. Specifically, we assume that unpaid leave occurs in the pay periods with the lowest amount of hours paid over a six-month window. Because we assume that the pay periods with the lowest hours are the ones when unpaid leave is taken, this assumption will, if anything, overstate the importance of unpaid leave in driving pay volatility. The methodology is described in detail in Appendix C.2.

Third, having imputed the number of hours of unpaid leave in each month, we calculate the monthly earnings volatility that would have been realized if workers had instead worked those hours. This delivers an estimate for counterfactual income volatility absent any volatility arising from unpaid leave. We find that monthly volatility absent unpaid leave is very similar to observed monthly volatility – the median monthly change in pay falls from 6 percent to 5 percent and the 75th percentile of the absolute change in pay falls from 14 percent to 12 percent (Table A-10).²⁰ Of course, unpaid leave may be quite important in explaining volatility for some workers. However, typical levels of unpaid leave are just too small for this channel to explain a meaningful share of the fluctuations in hours that occur on average in every month.

Childcare Needs Another possible reason why hours might fluctuate is because of childcare related needs. Because women with young children typically spend 19 hours per week on caregiving (Bureau of Labor Statistics 2022) and men with young children spend less time on caregiving, one might have expected women with children to sort into more flexible jobs and then work more volatile hours (Goldin 2021). One way to get a sense of the role that childcare needs play in driving pay volatility is thus to compare volatility for men and women by the number of dependent children they have.

Table A-11 shows that there is little difference in volatility by worker gender. This similarity is present whether one studies workers with children or workers without children. It holds for several measures of volatility. This similarity could be because the type of flexibility which parents with caregiving responsibilities seek might be at a higher frequency than the monthly pay volatility that we study. Alternatively, it could be because workers with childcare responsibilities actually prefer to have *more* stable hours (Bolotny and Emanuel 2022).

²⁰The baseline volatility estimates are for full-time hourly workers, as shown in Table 1 panel D.

4.2 Role of Firms in Driving Monthly Earnings Volatility

While the previous channels were quantitatively small, in this subsection, we provide evidence that firms play an important role in driving worker-level volatility. All analysis in this section focuses solely on hourly workers, since these are the workers that face pervasive pay volatility. We begin by showing that firms have fluctuations from month to month in total labor demand that are big enough to drive substantial *individual* volatility. After documenting this large role for fluctuations in total labor demand, we then show that firms also affect individual pay volatility through other channels (e.g., varying idiosyncratic schedules even when holding total hours fixed).

4.2.1 Magnitude of Total Firm Hours Changes

We begin by showing that firms as a whole exhibit substantial monthly volatility in total hours. Our assumption is that total firm hours are controlled by the firm so that any fluctuations in total firm hours arise from shifts in labor demand. While very small firms may not be able to achieve desired total hours in the face of individual shocks like worker sickness, idiosyncratic shocks are more likely to wash out for larger firms. For this reason, in this analysis, we restrict attention to firms with a median size of at least 20 hourly workers.²¹ We note that restricting to large firms does not rule out the presence of correlated labor supply shocks, but we find similar patterns after controlling for industry specific labor supply shifts using industry by month controls.

Figure 6 shows that firms do not have stable total hours from month to month. The blue bars show total hours growth including hours changes for continuing workers as well as changes in hours arising from hires and separations. It is not surprising that firm hours change frequently when including the effects coming from changes in total employment, since firms grow and shrink over time and there is substantial volatility of establishment employment. However, the orange bars show that firms also have significant monthly fluctuations in total hours worked by *continuing* workers. In particular, the median absolute percent change of total hours worked by continuing workers is 3.8 percent and the 75th percentile is 8.0 percent. These fluctuations in firm hours are big enough to represent a potentially important source of worker-level volatility. Indeed, we believe we are the first paper to use high-frequency firm-level data to document these types of high-frequency labor demand shifts.

For the remainder of our analysis of firms, 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*.

One useful way to capture the importance of movements in total firm labor demand for individual workers’ hours changes is a reallocation statistic inspired by Davis and Haltiwanger (1992) that compares the change in total firm hours from one month to the next to the gross sum of

²¹Appendix 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.

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.

all individual hours changes at that firm. This statistic captures the extent to which individual changes in a particular month were “necessary” or were in excess of that required to achieve the total net change in firm hours:

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

$$\Delta H_{j,t}^{gross} \equiv \sum_{i \in j} |\Delta h_{i,j,t}| \quad (4)$$

$$\text{Necessary share}_{j,t} \equiv \frac{\Delta H_{j,t}^{firm}}{\Delta H_{j,t}^{gross}}. \quad (5)$$

Put differently, necessary changes are those hours movements that are not canceled 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 change in total firm hours. This means that, in an accounting sense, changes in total labor demand can “explain” a significant fraction of changes in worker hours within firms. While this suggests that changes in total hours at the firm are an important source of worker fluctuations, we note that this is a comparison of net versus gross changes for the firm as a whole and not a quantification of the role of total hours changes for any individual worker.

4.2.2 Effect of Total Firm Hours Changes on Individual Volatility

To quantify the contribution of fluctuations in total firm hours (ΔH_t^{firm}) to *individual* volatility, we need to know what individual volatility would be in a counterfactual without fluctuations in total firm hours.²² Because many possible counterfactuals are all consistent with the same observed net and gross changes in hours, we need to make further assumptions in order to answer this question. We do so using a back-of-the-envelope allocation rule to construct counterfactuals under alternative values of the change in total firm hours. Concretely, let $\Delta \tilde{h}_{i,t}(x)$ be a potential outcome function that gives the change in monthly hours for worker i as a function of some counterfactual change in total firm hours x . We are interested in how $\text{Vol}_i(x) \equiv \text{Median}(|\% \Delta \tilde{h}_{i,t}(x)|)$ changes as we change the size of the firm wide hours shock x . To construct this counterfactual $\Delta \tilde{h}_{i,t}(x)$, we assume that changes in total firm hours are allocated across individual workers according to the following proportional allocation rule:

$$\Delta \tilde{h}_{i,t}(x) = \begin{cases} \Delta h_{i,t}, & \text{if } \text{sign}(\Delta h_{i,t}) \neq \text{sign}(\Delta H_t^{\text{firm}}) \text{ or } \Delta H_t^{\text{firm}} = 0, \\ \frac{\Delta h_{i,t}}{\Delta H_t^{\text{same}}} (x - \Delta H_t^{\text{firm}} + \Delta H_t^{\text{same}}), & \text{if } \text{sign}(\Delta h_{i,t}) = \text{sign}(\Delta H_t^{\text{firm}}). \end{cases} \quad (6)$$

where we define $\Delta H_t^{\text{same}} \equiv \sum_{\{k: \text{sign}(\Delta h_{k,t}) = \text{sign}(\Delta H_t^{\text{firm}})\}} \Delta h_{k,t}$. When we input the observed change in firm hours ΔH_t^{firm} into equation (6), it simply returns the observed change in individual hours $\Delta h_{i,t}$, but it delivers predictions for counterfactual hours changes when we input alternative values of x .²³ For example, in a counterfactual with no change in total firm hours, individual hours changes are proportionately reduced for all those moving 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, embedded in the separate rows of equation (6). 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 change. The data reject that firm hours changes are allocated equally across all workers, and we think it is then natural to assume that those workers whose hours co-moved most strongly with total firm hours in a particular month were those who absorbed that firm shock.²⁴

²²Each 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} .

²³For 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.

²⁴A pattern in the micro data supports the assumption that firm-level shocks are absorbed by some workers, while other workers are fully insulated from these shocks: individual changes of exactly zero occur in around 10 percent of worker-months, while total firm hours changes of exactly zero are extremely rare. These two patterns are mutually consistent only if (1) some workers are not allocated any of the firm-wide change or (2) many workers in each month happen to have idiosyncratic shocks that *exactly* offset the particular firm shock they were exposed to that month,

Figure 7: 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 (6).

After defining $\Delta \tilde{h}_{i,t}(x)$, we use this to construct an estimate of what individual hours volatility would be if total firm hours volatility was eliminated ($Vol_i(0)$).²⁵ Figure 7 compares the CDF of individual volatility across workers observed directly in the data to the counterfactual CDF after this firm volatility is eliminated. It shows that eliminating total firm hours volatility would substantially reduce individual volatility. For example, the median value of volatility $Vol_i(\Delta H_t^{firm})$ is 8.4 percent while the median value of $Vol_i(0)$ is 3.5 percent. The mean value of volatility falls from 10.9 percent to 5.9 percent. Thus, these back-of-the-envelope calculations suggest that around half of individual volatility for the typical worker arises from fluctuations in total firm hours. We emphasize that this is only one of many possible counterfactuals, so the results should be interpreted as a back-of-the-envelope calculation rather than a precise estimate.

4.2.3 Effect of Firm Switches on Individual Volatility

The results thus far suggest that firms play an important role in driving individual volatility, but these results only measure effects of changes in total labor demand and not any firm effects arising through other channels like idiosyncratic scheduling. They also rely on strong assumptions about how labor demand shocks are split between specific workers.

Thus, we now turn to a completely distinct source of evidence that firms matter for individual volatility, based on worker moves between firms. This analysis allows us to investigate two conceptually separate but related questions: (1) what is the effect of firms in general on individual volatility? and (2) what is the effect of total firm hours volatility in particular on individual

which seems implausible.

²⁵To construct percent changes we compute the entire time-series of $\Delta \tilde{h}_{i,t}(0)$, cumulate these changes over time to construct a counterfactual series of hours levels $\tilde{h}_{i,t}(0)$ for each worker, and then compute $Vol_i(0) = Median(|\% \Delta \tilde{h}_{i,t}(0)|)$.

volatility? Our broader goal is to disentangle worker-driven and firm-driven sources of earnings instability, making the first question central to our analysis. However, the second question is also informative as it sheds light on *why* firms matter. This second question parallels the one explored in Section 4.2.2, but here we answer it using a completely different identifying assumption.

While these two questions are distinct, we approach them using similar empirical strategies. To answer question 1, we use a fixed effects movers design to capture the role of all time-invariant firm characteristics that affect individual volatility. To answer question 2, we again use a movers design but instead of regressing changes in individual volatility on firm fixed effects, we regress changes in individual volatility on one particular firm characteristic: firm total hours volatility.

Fixed Effects Movers Designs We use a fixed effects specification with movers following Abowd, Kramarz, and Margolis (1999) to study the effects of firms and workers on individual earnings volatility.²⁶ Since we only observe a small share of all firms in the economy 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.²⁷ We estimate:

$$Vol_{i,j} = \mu + \alpha_i + \psi_{k(j)} + \varepsilon_{ij}, \text{ with normalization } \psi_1 = 0 \quad (7)$$

where $Vol_{i,j}$ is $Median(|\% \Delta_{i,j,t}|)$ for the job-spell of worker i at firm j , μ is an overall intercept, $\psi_{k(j)}$ indexes 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.²⁸ We group firms into deciles of average worker-level volatility in our baseline estimation, so $\psi_{k(j)}$ measures the difference in volatility relative to the baseline decile 1.²⁹ All of the analysis in this section restricts to hourly workers.

A large literature has emphasized that interpreting resulting estimates as causal is only valid under 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 a change in a worker's idiosyncratic volatility causes them to switch to a firm with different volatility or if the realization of the match-specific component affects which matches are actually formed. In Appendix D.2, we show that standard event-study diagnostic validations are satisfied: there are sharp changes in volatility around job transitions, no evidence of


²⁶Section 4.2.2 focused on the specific relationship between firm and individual hours, so for that section we focused on individual outcomes in hours space. However, 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., overtime premia may amplify the effect of an hours increase on earnings) or firms might differ in the frequency at which they vary wages. Thus, our primary outcome in this section is individual *earnings* volatility for hourly workers. However, in practice our conclusions are very similar if we instead measure individual *hours* volatility.

²⁷This grouping approach will miss the effect of firm heterogeneity within group and so it is conservative in relation to our finding that firms are an important source of volatility. In our baseline results, we group firms into ten deciles by the mean volatility of their individual earnings, but k-means clustering produces similar results.

²⁸Each combination of i, j indexes a unique match. Since we estimate this specification at the spell-level, to simplify notation we do not index the regression by t .

²⁹We define the volatility of firm j as the average of individual hourly workers' volatility $Vol_{i,j}$ weighted by each worker's spell length. We found that results are similar using other aggregation methods.

Figure 8: Firm Causal Effects by Decile Estimated from Movers Designs



Notes: This shows the results of estimating the two-way fixed effects design in equation (7) 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).

pre-trends, and these changes are similar for those moving to (from) higher (lower) volatility firms, suggesting that changes in worker volatility are not causing moves to different firms.³⁰

Figure 8 shows a visual representation of the firm effects that arise from estimating equation (7). For each firm decile, the orange bar shows ψ_k , capturing the differential firm effect relative to decile 1. Under the AKM identification assumption, 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, there is no way to determine whether it arises from firm or worker choices.

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 decomposition of the total observed volatility in each decile into a common intercept, firm-specific causal effects, and worker-selection effects.³¹

Figure 8 highlights the quantitative importance of firm effects in explaining volatility for most deciles – moving from a bottom to a middle decile firm causes individual volatility to rise by 7.2 percentage points, from 4.6 percent to 11.8 percent.³² One way to gauge the size of this effect is

³⁰Borovičková and Shimer (2024) argues that these diagnostics may fail to detect violations. In Appendix D.2 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.

³¹By construction, the match-specific residual ε_{ij} averages to 0 in each firm group k so the decomposition in Figure 8 is exact.

³²This is the sum of the common component + causal effect in decile 1 vs 6. We compare decile 1 and 6 since the

to compare it to the typical *level* of individual earnings volatility. We are particularly interested in this comparison since it tells us how important firms are in explaining the volatility of the typical worker. Comparing this 7.2 percentage point effect to the median value of Vol_i (11.9 percent) suggests that about 60 percent of volatility experienced by a worker in a typical firm is caused by that firm.³³ We further note that this is a lower bound, since some of the common component of volatility μ may also be driven by firms.

Another way to gauge the magnitudes of these firm effects is to compare them to heterogeneity across workers in volatility. That is, instead of asking whether firm effects are big enough to explain the typical *level* of volatility we observe, we can ask whether firm effects explain a large share of the *variation* across workers in volatility. Such comparisons are the focus of the existing wage AKM literature. Table A-12 shows that firm effects account for just over half of explained heterogeneity. In Appendix D.3, we also explore a number of additional related robustness results that further 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 to moves between firms with similar wages; they are not driven by one particular type of firm; they become slightly stronger when we expand the number of firm groups in our analysis; they are very similar if we study hours volatility instead of pay volatility; and they remain large when restricting to prime-age workers.

However, it is important to also note that in addition to the sizable role of firms, these same variance decompositions also imply an important role for worker fixed effects. While we focus primarily on the role of firms in explaining volatility for a typical worker with median volatility, this need not imply that firms play an equal role in driving volatility for every single worker. Figure A-6 shows that although volatility for the median worker is high, there are also some workers with much more extreme levels of volatility. These individuals play an outsize role in variance decompositions, and it is plausible that many of these very large monthly changes are indeed worker initiated (e.g., switches from full-time to part-time).³⁴ Although firm-driven fluctuations loom large for typical workers, they may not explain this kind of extreme individual volatility.

Moves Between Firms With Different Total Hours Volatility The analysis directly above finds that firms have an important causal effect on individual pay volatility using a firm fixed effects movers design. These estimates capture the combined effects of any time-invariant firm attributes that affect individual volatility. In Section 4.2.2 we found that firm labor demand shocks are one important channel through which firms drive individual volatility. To conclude this section, we show that these results are in fact tightly linked: there is substantial heterogeneity across firms in total hours volatility and these differences cause workers to differ in their individual pay volatility.


The binscatter in Figure 9a shows that the substantial variation across firms in total hours

median firm is exactly at the boundary between decile 5 and 6, but a movement from the middle of decile 1 to the middle of decile 6 is closer to a median change than a movement from 1 to 5. Moving from decile 1 to 5 produces a slightly smaller but similar effect.

³³Comparing to the volatility just of movers or just of decile 6 workers produces a similar result.

³⁴This still 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 9: 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.

volatility is highly correlated with individual worker volatility. This cross-sectional correlation may not reflect a causal relationship since firms with high total hours volatility may have workers with different individual characteristics. However, the binscatter in Figure 9b shows that the strong positive relationship persists when looking within individuals and relating changes in individual pay volatility across job spells to differences in the total hours volatility of the firms. Running a corresponding linear regression produces a large and highly significant coefficient of 0.88.³⁵ This coefficient means that moving from a firm with the median level of total hours volatility to a firm with no total hours volatility reduces individual pay volatility by 4.2 percentage points, which is similar to the decline in hours volatility of 4.3 percentage points that we obtained in Section 4.2.2.³⁶ Thus, two different approaches with strong but distinct identifying assumptions both imply that firm labor demand shocks play an important role in driving worker volatility.³⁷

We conclude this section by noting that these results imply that firm total hours volatility explains most but not all of firms' contribution to workers' volatility. We can infer this by comparing the size of firm fixed effects from Section 4.2.3 to the effects of firm hours volatility from the previous paragraph. We find that a one standard deviation increase in the volatility of total firm hours predicts an increase in individual pay volatility of 0.030, while a one standard deviation increase in overall firm fixed effects increases individual volatility by 0.048. This means that around two-thirds of the heterogeneity in firms' causal effects can be explained by firm total hours volatility.

³⁵Interpreting this regression as causal requires both the same exogenous mobility assumption underlying equation (7) and the *additional* assumption that variation in total hours volatility is uncorrelated with any other residual firm effects that also affect individual volatility.

³⁶The median firm has a total hours volatility of 0.048 (which we note is slightly higher than the median *firm-month* change of .040 reported in Section 4.2.1) and $4.2\% = 0.88 * 4.8\%$. The decline of 4.3 percent in Section 4.2.2 comes from comparing the median volatility in the sample (8.2) to the median volatility in our counterfactual without total hours fluctuations (3.9).

³⁷Note that the movers design measures effects on earnings volatility while Section 4.2.2 measures effects on hours volatility. If we re-do the movers design instead estimating the effect on individual hours volatility we get an estimated effect of 5.0 percentage points.

This suggests that there are also other important channels (e.g., management, scheduling) through which firms matter.

5 Is Monthly Earnings Instability Costly?

We have now shown that workers face significant fluctuations in monthly pay, and that firm-driven labor demand changes, rather than worker choices, play an important role in driving these movements. However, there are theoretical reasons to think these earnings fluctuations still might not matter much for welfare. First, Section 3.4 shows that they are transitory. Second, although Section 4.1 showed a lack of seasonality, these fluctuations may actually be predictable to the workers themselves.

In this section, we provide two complementary pieces of evidence that, despite these theoretical reasons for skepticism, pay volatility indeed matters for workers in practice. First, in Section 5.1 we show that month-to-month pay volatility causes month-to-month spending volatility. Second, in Section 5.2, we show that workers are more likely to quit jobs with high pay volatility. Both effects are quantitatively large, suggesting that the high-frequency earnings risk missed in annual data is indeed welfare relevant and should not be ignored.

5.1 Spending Evidence


In this section we use JPMCI data to investigate the link between monthly earnings volatility and monthly spending volatility. Appendix A.2 discusses this data in more detail, but we highlight a few points from that discussion here. Our main measure of volatility is the median absolute percent change. We measure the median absolute percent change in earnings $Vol_{i,j}^y$ and in spending $Vol_{i,j}^c$ within a job spell of worker i at firm j , applying similar filters and sample restrictions within the JPMCI data to construct these objects as we use in our baseline PayrollCompany analysis. We focus on broadly defined non-durable spending, but show robustness to various other measures of spending. Our primary analysis focuses on hourly workers, but we also compare hourly workers to salaried workers in robustness analysis.

The binscatter in Figure 10 shows that there is a strong positive relationship between $Vol_{i,j}^c$ and $Vol_{i,j}^y$ in the cross-section: in job spells in which a worker’s income is more volatile month to month, that worker’s spending is also more volatile month to month.³⁸ The first column of Table 3 estimates this cross-sectional relationship with the following regression:

$$Vol_{i,j}^c = \alpha + \beta Vol_{i,j}^y + u_{i,j} \quad (8)$$

³⁸Both series are de-measured relative to the sample average. We note that, perhaps surprisingly, the level of monthly spending volatility is actually higher than the level of income volatility. A natural explanation for spending volatility even when income is constant is the presence of spending shocks which are not driven by income changes, e.g., car/home repairs, medical expenses, insurance payments, vacations timed with calendar dates, etc. However, an alternative interpretation is that violating consumption smoothing at high frequencies is not costly. We return to this point when interpreting risk aversion in welfare calculations, but note that the *level* of spending volatility is largely irrelevant for our analysis since it plays no role for welfare under CRRA preferences.

Figure 10: Job Spell-Level Relationship between Income and Spending Volatility



Notes: This shows a binscatter of individual volatility of monthly non-durable spending vs. individual volatility of monthly income using Chase data. Individual volatility is measured as the individual's median absolute monthly percent change. See Section 2.2 for definition of non-durable spending.

Table 3: The Effect of Income Volatility on Consumption Volatility

	Dependent Variable: Med $ \%C $			
	OLS	IV	IV	IV
	(1)	(2)	(3)	(4)
Med $ \%Y $	0.232*** (0.005)	0.236*** (0.012)	0.149*** (0.037)	0.124*** (0.030)
Implied WTP	7.21%	7.34%	4.57%	3.79%
Group	All Households	All Households	Movers	Stayers
Variation	Cross-Section	Cross-Section	Panel	Panel
Observations	228,502	228,502	16,160	131,679

Notes: $*p < 0.1$; $**p < 0.05$; $***p < 0.01$. Each observation is one worker. Standard errors are clustered by firm. All specifications restrict to hourly workers. We calculate the implied willingness to pay to set median pay volatility to zero (WTP), assuming a coefficient of relative risk aversion of two. 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. Column (4) analyzes within-firm changes in volatility among workers who stay at a single firm and restricts to a sample of workers with tenure of 12+ months.

The estimate of $\hat{\beta} = 0.232$ implies that when the median monthly percent change in income increases by 10 percentage points, the median monthly percent change in spending increases by 2.32 percentage points. We interpret this magnitude in more detail below.

Of course, this cross-sectional relationship may not be causal. The OLS 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 (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 in two steps that combine an instrumental variables strategy with individual fixed effects. In the first step of our identification strategy, we instrument for individual income volatility using the average individual income volatility at firm j where individual i works, $\overline{Vol}_{j(i)}^y$ and run the two-stage IV regression with second stage:

$$Vol_{i,j}^c = \alpha + \beta \widehat{Vol}_{i,j}^y + u_{i,j}, \quad (9)$$

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

$$Vol_{i,j}^y = \kappa + \pi \overline{Vol}_{j(i)}^y + e_{i,j}. \quad (10)$$

Since this IV regression only uses volatility that is common to all co-workers in the firm, it is essentially asking: do 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: there is a similar, strong relationship between income and spending volatility when using this instrument to remove idiosyncratic confounds.³⁹ 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 the results in Section 4.2 that firms explain a large share of individual volatility, this instrument unsurprisingly satisfies the relevance condition with first stage F-statistics of over one thousand.

The exclusion restriction requires that the relationship between firm volatility and individual consumption volatility arises only through the effect of firm volatility on individual income volatility. Since we do not have exogenous variation in firm volatility, firm volatility might be correlated with various other firm characteristics like management practices, industry, or unobserved amenities. However, it is not clear why firm characteristics should have any effect on consumption volatility except through their effects on income volatility and so are unlikely to violate the exclusion restriction. In contrast, firm-level sorting poses a more pertinent potential violation of the exclusion restriction. Workers with particular preferences (e.g., low risk aversion) may systematically select into firms with different volatility, potentially violating the exclusion restriction that average firm volatility only affects consumption volatility through the effect it has on a worker's income volatility.

Thus, in the second step of our identification strategy, we again instrument for individual volatility with the average volatility of the firm, but introduce worker fixed effects so α becomes α_i

³⁹Since different confounds generate different relationships between income and spending volatility, it is not obvious whether simple OLS correlations are likely to under or overstate causal relationships.

and κ becomes κ_i in equations (9) and (10). These fixed effects control for any permanent worker characteristics that might differ across firms.

Column 3 of Table 3 runs this fixed effect IV regression and shows that there is again 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 consumption volatility across jobs arise only from *changes* in income volatility across jobs. Unlike in AKM, endogenous job transitions are not a problem for identification, as long as consumption volatility only changes because of the changes in income volatility.⁴⁰ 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 will still recover the correct causal effect of income volatility on spending using this “endogenous” job transition.

However, some life events might induce job changes at the same time that preferences for spending volatility are changing. 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. Three additional pieces of evidence on volatility dynamics suggest that such violations of the identifying assumption are not driving the results.

First, we consider an alternate implementation of equation (9) where we consider how spending volatility and income volatility change over time *within* job spells. Specifically, we split each job spell within a firm in half, creating two pseudo-spells for each worker-firm pair, and then re-estimate the IV with individual fixed effects but now looking at how an individual’s spending volatility changes as their income volatility changes across these pseudo-spells instead of true job spells.⁴¹ That is, we run a panel regression for stayers instead of movers, with variation over time coming from things like time varying firm labor demand shocks. This requires focusing on longer tenure workers so that we can measure volatility separately in these two pseudo-spells. It is unlikely that there are discrete life changes corresponding to the fairly arbitrary timing of these same sample splits, making it more likely that the identification assumption, which is that changes in firm-level income volatility only affect spending volatility through the effect on worker income volatility, holds. Column 4 of Table 3 shows that this stayers specification produces similar estimates to the movers design in Column 3.

Second, in Appendix D.2, we explore difference-in-difference event study designs and show that

⁴⁰Identification of the causal effect of average firm volatility on income volatility π instead requires exogenous mobility: $E[\Delta e_{i,j} \mid \Delta Vol_{j(i)}^y] = 0$.

⁴¹This requires using a slightly different measure of firm-volatility, since the previous measure is time-invariant. Instead, we measure the median absolute percent change in earnings using all co-worker-months that overlap with the calendar dates of the spell, as discussed more detail in Appendix A.2.

there are sharp increases in spending volatility when people move to jobs with higher income volatility, with no evidence of pre-trends. If life events were simultaneously shifting people’s 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 job transition.

Third, if life events were causing consumption volatility and income volatility to shift in non-causal ways around job transitions, we would expect that excluding the months around job transitions from our measures of volatility would lead to different relationships. However, in Appendix D.2, we re-estimate all of our baseline specifications but computing spell-level volatility excluding progressively larger “donuts” around the date of move. We find very similar relationships between income and spending volatility even when we compute this volatility excluding several months before and after the move, again suggesting that endogenous moves corresponding with confounding shocks are not driving our results. Thus, across all four specifications, we find highly statistically significant relationships between income and spending volatility, with coefficients ranging from 0.12 to 0.24.

Table A-13 explores whether the effect of income volatility on spending volatility varies with worker characteristics. To do this, we re-estimate equation (9) from our baseline results but add interactions with indicators for different group characteristics.⁴² Table A-13 Column 1 shows that there is a strong relationship between income volatility and spending volatility for hourly workers but not for salaried workers. This might be explained by the fact that monthly earnings changes are less persistent for salaried workers or that salaried workers’ earnings changes may also be more predictable.

The much stronger relationship we find for hourly workers again highlights that the distinction between hourly and salaried work is a crucial dimension of heterogeneity for understanding people’s lived experiences. Not only do lower-income hourly workers face much more earnings instability than salaried workers, this earnings instability also passes through more strongly into resulting monthly spending. This heterogeneity reinforces our focus on these workers throughout the paper and all subsequent regressions in the table restrict to hourly workers, as in Table 3.

Column 2 shows that low liquidity households have much stronger relationships between income and spending volatility than high liquidity households. This is important for two reasons. First, it suggests that low liquidity, financially fragile workers may suffer particularly large consequences of earnings instability. Second, it bolsters the credibility of our causal estimates because there is a clear mechanism through which liquidity interacts with our causal channel. Column 3 shows that low income households have slightly higher responses to income volatility than high income households, although this difference is not significant. However, we again note that low income households have a higher *exposure* to income volatility than high income households. Column 4

⁴²We focus on equation (9) rather than panel designs, since many of these demographic characteristics are absorbed by individual fixed effects. Since the firm-volatility instrument is interacted with these characteristics in the first stage, this means that workers with different characteristics have different predicted income volatility in the second stage.

shows no difference in responses by age, which is reassuring for identification, as it reduces concerns about life-cycle changes in preferences confounding our results.

Table A-14 presents a series of additional robustness checks for our consumption analysis. One concern is that higher hours in a given month may mechanically raise work-related spending (e.g., commuting, meals), creating a spurious link between earnings and consumption volatility. Columns 2 and 3 in Table A-14 address this by splitting 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 work related spending.⁴³ Column 4 shows that we obtain similar strong effects when we measure the volatility of total spending rather than non-durables. Column 5 tests whether monthly volatility matters less than quarterly volatility, which is more commonly analyzed in macroeconomic models. Repeating our analysis using quarterly changes in both income and spending, we continue to find strong effects. Finally, 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 suggests that monthly volatility has important consequences even at standard business-cycle frequencies, where there is more consensus that these spending fluctuations are welfare-relevant.

These effects are also economically meaningful. One way of gauging this is to ask how much households would be willing to pay to eliminate this volatility. We do so using the simple Lucas (1987) formula

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

where γ is the coefficient of relative risk aversion, and σ_c^2 is the variance of consumption arising from income volatility σ_y^2 , and β_σ is an estimate of this causal effect of income variance on consumption variance.⁴⁴ These welfare calculations must make an assumption about relative risk aversion (γ).⁴⁵ Table 3 reports the willingness to pay to eliminate the income variance observed for the median hourly worker ($\sigma_y^2 = 0.05$) using $\gamma = 2$. The implied willingness to pay values range from around four to seven percent across empirical specifications, suggesting that this high-frequency earnings risk has very substantial welfare costs.

We report results using $\gamma = 2$, since this is a relatively standard value in the business cycle

⁴³Many work-related expenses are relatively fixed, so it is not obvious that they should mechanically respond to fluctuations in hours. For example, a worker must pay the same commuting costs for working a 6 hour shift as a 12 hour shift, and the amounts spent on some other work-related categories like clothing are likely to be insensitive monthly fluctuations in hours worked.

⁴⁴We estimate the relationships in Table 3 using $Med|\Delta Y|$ since we focus on this volatility measure in the rest of our results, but the Lucas formula requires measuring volatility relationships in variance space, requiring a translation. Let β_{med} be the coefficient estimated in Table 3 (called β in equation 9). In Appendix E.1 we derive translation $WTP = 0.5\gamma(\beta_{med}^2\sigma_y^2 + 2\beta_{med}\sigma_y\sigma_{c,0})$, where $\sigma_{c,0}$ is the standard deviation of consumption for a household with zero income variance. Alternatively, simply re-estimating all of the empirical regressions using $Vol = \sigma^2$ instead of $Med|\Delta|$ delivers nearly identical welfare estimates to using this conversion.

⁴⁵Since all income volatility has the same effect on the budget constraint, we assume that we can extrapolate from the causal effects identified using the firm instrument to effects of other volatility and thus calculate the welfare effects of reducing *total* income volatility for the median worker.

literature. However, it is possible that γ may be lower at the monthly frequencies that we study since monthly violations may be less costly than quarterly violations of consumption smoothing. Since the willingness to pay is linear in γ , it is trivial to convert the numbers we report to estimates under alternative values of γ , and welfare costs would remain sizable even with substantially lower values of this curvature parameter. We also note that Table 2 showed that lower income workers face particularly high income volatility, meaning they have an even higher willingness to pay to eliminate volatility.

Because earnings instability is concentrated among lower-income workers and is welfare reducing, this disamenity increases compensation inequality relative to wage inequality. $Vol_{i,j}^y = \text{Median}(|\% \Delta_{i,j}|)$ is 15 percent for a worker at the 10th percentile of the income distribution while $Vol_{i,j}^y$ is only 5 percent for a worker at the 90th percentile of the income distribution.⁴⁶ We can then apply our previous willingness to pay calculation to these two different levels of volatility. Using $\hat{\beta}$ from Table 3 column (2) implies that a worker would give up 6.98 percent of income to eliminate $Vol_{i,j}^y = 0.15$ but only 2.23 percent of income to eliminate $Vol_{i,j}^y = 0.05$.⁴⁷ Thus, workers at the 10th percentile of the income distribution would pay nearly 5 percent more of their income to eliminate the typical volatility they face than workers at the 90th percentile of the income distribution.

When looking across the distribution of many different amenities, incorporating earnings instability yields a meaningful increase in measured inequality. Maestas et al. (2023) estimates the willingness to pay for nine different job amenities—although not earnings or hours stability—based on vignettes and finds that 90-10 inequality is 8 log points higher after accounting for differences in amenities.⁴⁸ Thus, incorporating just one additional (dis)amenity—earnings instability—raises the 90-10 gap from 8 to 13 log points.

5.2 Job Quits Evidence

In this section, we return to the payroll data and provide a complementary piece of evidence that workers dislike volatility: hourly workers are more likely to quit high volatility jobs. This evidence also provides an implied willingness to pay to eliminate volatility. This estimate differs in that it does not rely on knowing a worker’s risk aversion and it includes costs that arise from channels besides the household budget constraint: working variable hours over time may be costly even if households have the ability to perfectly smooth their consumption.


We begin by visualizing the strong positive relationship at the firm level between the average

⁴⁶For these inequality calculations we include *all* workers rather than focusing just on hourly workers since we want to capture our earlier evidence that salaried and hourly jobs have different levels of typical volatility and that the share of salaried jobs rises with income.

⁴⁷These differences are somewhat larger if we use estimates based on quits from the next section and somewhat smaller if we use a different point estimate from Table 3 but are always sizable.

⁴⁸The paper reports two different estimates, one where all workers have homogeneous valuation of job amenities and a second where amenity valuations are allowed to vary with worker observables. Our estimates capture the homogeneous case and are therefore most comparable to the former. If we allowed for heterogeneity in valuations of earnings instability, the results in Table A-13 suggest that workers with lower liquidity would have higher valuations.

Figure 11: 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.

individual volatility $\overline{Vol}_j^y$ of a firm j and the average separation rate at that firm. We focus primarily on the total separation rate, which sums up quits, layoffs and firings as we can measure the total separation rate for all firms. Decomposing separations into quits versus other separations requires firms to record additional information and only around half of firms consistently do so. Nevertheless, Figure 11 shows that for the firms who do record quits specifically, the positive relationship between total separations and volatility is driven almost entirely by a positive relationship between quits and volatility. For this reason, we interpret separations as largely reflecting worker choices.

Table 4 presents individual-level results using 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 relative to a baseline hazard function, $H_0(t)$ and X_{ij} is a vector of potential controls.⁴⁹ The coefficient $\hat{\beta}_1$ from the simplest version of this regression is shown in the first row of Column 1. This estimate implies that moving from a job with constant earnings to one with median hourly worker volatility (11.9 percent) raises the separation rate by 38 percent ($1.38 = \exp[2.72 \times 0.119]$).

To interpret this effect's magnitude, we compute the implied willingness to pay to eliminate volatility. Gronberg and Reed (1994) propose a model where dividing the elasticity of separations to an amenity (β_1) by the elasticity of separations with respect to wages identifies the willingness

⁴⁹This 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 occurs because they switch payroll providers or cease operation and so we treat these as censored spells. However, we obtain similar results re-estimating all of the specifications in this table instead using linear probability models, treating censored spells as outcomes of zero.

Table 4: The Effect of Income Volatility on Separation Rates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Movers Only	Full-time Only	Large Firms Only	Pooled IV wo last Q
Med $ \% \Delta Y_{ij} $	2.72*** (0.093)	3.26*** (0.253)	3.01*** (0.293)	3.29*** (0.277)	3.21*** (0.452)	2.81*** (0.350)	3.06*** (0.363)
No. Obs	86,598	86,598	86,598	15,296	55,517	74,911	71,304
No. Firms	6,199	5,793	5,793	3,262	4,841	2,634	5,115
Implied WTP	9.2% (0.3%)	10.7% (0.7%)	10.0% (0.8%)	10.8% (0.7%)	10.6% (1.2%)	9.5% (1.0%)	10.2% (1.0%)
Instrument?	No	Yes	Yes	Yes	Yes	Yes	Yes
Sample	All	All	All	Movers	All	All	All
Controls?	No	No	Yes	Yes	Yes	Yes	Yes

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. 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 instrumenting 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 to control for unobserved heterogeneity and selection. 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). See text for additional details.

to pay for the amenity. This calculation tells us how much the firm would have to pay the worker to counteract the increased separation rate brought on by the higher volatility.⁵⁰ Using a conservative value from the literature of -3 for the elasticity of separations to the wage, the Column 1 coefficient implies that an hourly worker with median volatility would give up 9.2 percent of wages to eliminate it.⁵¹

Subsequent columns of Table 4 explore alternate specifications that attempt to control for various confounders that might spuriously generate this correlation. For example, serious illness could increase both hours volatility 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

⁵⁰Bassier, Dube, and Naidu (2022) and Lamadon, Mogstad, and Setzler (2022) use quasi-random variation in wages to estimate separation elasticities, finding values of -3 and -2.16, respectively. Lamadon, Mogstad, and Setzler (2022) compute total labor supply elasticities of 6.02, and 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 used 6.02, the WTP numbers in the table would all halve but still remain large. Bassier, Dube, and Naidu (2022) report a preferred estimate of -2.1 for total separation elasticities, and this would generate a larger WTP than in the table.

⁵¹ $0.092 = \left(1 - \frac{1}{\exp(2.72 \times 0.119)}\right) \frac{1}{3}$.

5.1. The coefficient increases modestly.⁵² This suggests that common, firm-wide volatility, is especially costly, while idiosyncratic volatility may partly reflect voluntary choices.

The exclusion restriction supporting a causal interpretation of this relationship requires that $\overline{Vol_{j(i)}^y}$ only affects separation rates through its effects on worker’s income. This is likely a stronger assumption than the exclusion restriction in Section 5.1. For example, suppose that firms with low quality managers have higher income volatility. The assumption in Section 5.1 is that bad management only affects consumption volatility through its effect on income volatility, which seems plausible. However, it seems likely that a bad manager might directly make someone quit. While we cannot control for all characteristics of a firm that directly affect quit rates, our preferred specification in Column 3 adds controls for firm average wages and hours, industry fixed effects, and for a worker’s gender and age at job start to the IV specification. Incorporating controls for firm average wages is important given the prior literature on job ladders and incorporating controls for firm average hours is important given the finding in Lachowska et al. (2023) that workers prefer firms with higher average hours. Unsurprisingly, adding these controls reduces the relationship between volatility and separations a bit but it remains quantitatively strong.

To further bolster the conclusion that firm volatility indeed has a causal effect on quits rather than merely reflecting effects of unobserved firm characteristics, we look at the relationship between quit rates of *salaried* workers and $\overline{Vol_{j(i)}^y}$ (of *hourly* workers) at the same firms. Salaried and hourly workers at the same firms share the same unobserved firm characteristics. However, $\overline{Vol_{j(i)}^y}$ has a weaker correlation with salaried workers’ income volatility, so we should then also expect a weaker relationship with salaried workers’ quit rates. Table A-15 shows that this is indeed the case by estimating the reduced form effect of $\overline{Vol_{j(i)}^y}$ on quit rates separately for hourly and salaried workers at the same firms.⁵³ Firm volatility increases hourly workers’ separation rates by almost three times as much as for salaried workers at the same firm.⁵⁴ This difference reinforces the interpretation that volatility itself indeed has a causal effect on separations.

However, these specifications looking at salaried workers’ quit rates do not rule out concerns about sorting: firms with different types of workers may have both greater volatility and greater quit rates even if this relationship is not causal. To address this concern, Column 4 adds controls for fixed unobserved worker heterogeneity, identifying the effects of volatility on separation rates using workers who move between firms in our sample.⁵⁵ This delivers estimates that are slightly

⁵²Because 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.

⁵³Note that here we are interested in whether *firm* volatility ($\overline{Vol_{j(i)}^y}$) has a different effect on quit rates for hourly and salaried workers so we estimate the reduced form rather than the IV. In contrast, in Table A-13 we were interested in whether *individual* volatility (Vol_i^y) has a different effect on spending for salaried vs. hourly workers so we instead estimated the IV with these interactions.

⁵⁴The fact that $\overline{Vol_{j(i)}^y}$ has any relationship at all with salaried workers’ separation rates might reflect the fact that $\overline{Vol_{j(i)}^y}$ is also correlated with salaried workers’ earnings volatility, effects of *hours* instability holding constant earnings, or it might reflect the role of unobserved firm confounds.

⁵⁵Fixed effects cannot be implemented in the Cox hazard, so we model worker heterogeneity as shared frailty random effects with a Γ distribution. We obtain similar results estimating linear probability models with worker fixed effects.

higher but statistically indistinguishable from the IV specification with controls from 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 and shows 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 quits may cause an increase in pay volatility for coworkers who remain at the firm) by restricting only to large firms where a single worker quitting has a smaller effect on average volatility (Column 6)⁵⁶ and excluding the final quarter of each worker’s spell (Column 7).⁵⁷

Across all specifications, the evidence consistently shows that workers in volatile jobs quit at higher rates. The results appear robust and causal, and the estimated magnitudes imply that workers place substantial value on earnings stability. However, this does not on its own reveal *why* workers dislike pay volatility. In particular, it is not clear whether the quit elasticity reflects distaste for fluctuating hours (holding income constant) or distaste for fluctuating income (holding hours constant). This contrasts with the formula in equation (11) for spending volatility, which captures only the effects of fluctuating income. Survey evidence from workers in Schneider and Harknett (2019) suggests that fluctuating hours might be more important than fluctuating income.

Finally, this evidence on how pay volatility affects spending and quits also helps address two limitations discussed earlier in the paper: first, in Section 2 we discuss the possibility of spurious pay volatility arising from measurement error. Second, in Section 4.1 we discuss some evidence against worker labor supply as the main driver of earnings volatility. Measurement error-driven fluctuations or fluctuations in earnings driven by worker choices would be unlikely to generate the relationships with spending and quits behavior that we document in this section.

6 Conclusion

In this paper we showed that workers face substantial monthly earnings risk that has been missed in annual data. This risk is borne primarily by relatively low income, hourly workers. It is driven in large part by fluctuations in firms’ labor demand and these fluctuations induce substantial costs on affected workers.

Our paper raises several interesting questions for future research. Although we find that fluctuations in labor demand are important for driving pay instability, we do not explain *why* firms change workers’ hours so much. Is stable pay one component of “good” management practices that

⁵⁶We use the same size cutoff requiring a median of 20 workers in each month as in Section 4.2.1, but we find similar results when instead using an even more conservative cutoff of 100.

⁵⁷Similar 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 3.23 (s.e. 0.36).

make firms profitable (Bloom and Van Reenen 2007)? If so, what specific scheduling practices are effective? Alternatively, are these hours fluctuations necessary given the production process?

It also would be interesting to understand how institutions and policies in different economies handle firm- and worker-level volatility. For example, many U.S. cities have recently regulated firm scheduling practices. Many other developed economies have a lower share of hourly workers, but have other institutions to manage labor demand fluctuations (e.g. temporary contracts, short-time work, working time accounts). If firms face similar underlying shocks to desired labor demand in countries with stricter regulations, it is likely that firms in those countries are then absorbing more of this risk. How should policies be designed to balance workers' desires for stable income with firms' desires for flexibility and what are their efficiency and distributional consequences?

Finally, our results have direct lessons for structural modeling: models of income dynamics calibrated to annual data do not accurately capture high-frequency income risk. They should be adjusted to match the high-frequency patterns we document and to explicitly distinguish hourly from salaried work. In addition, consumption models that take these income patterns as inputs should in turn endeavor to produce implied spending patterns that also match our new evidence.

References

- Abowd, John M. and Kevin L. McKinney.** 2024. “Mixed-Effects Methods for Search and Matching Research.” Revue économique, 51(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.
- 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.
- Bloom, Nick and John Van Reenen.** 2007. “Measuring and Explaining Management Practices Across Firms and Countries.” The Quarterly Journal of Economics, 122(4): 1351–1408.
- Bolotnyy, Valentin and Natalia Emanuel.** 2022. “Why Do Women Earn Less than Men? Evidence from Bus and Train Operators.” Journal of Labor Economics, 40(2): 283–323.
- 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.
- Bureau of Labor Statistics.** 2022. “How parents used their time in 2021.” Accessed: Jul 16, 2025.
- Caldwell, Sydnee, Ingrid Haegele, and Jörg Heining.** 2025. “Firm Pay, Amenities, and Inequality.” Working Paper.
- 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.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline.** 2018. “Firms and Labor Market Inequality: Evidence and Some Theory.” Journal of Labor Economics, 36(S1): S13–S70.
- Card, David, Jörg Heining, and Patrick Kline.** 2013. “Workplace Heterogeneity and the Rise of West German Wage Inequality.” Quarterly Journal of Economics, 128(3): 967–1015.

- Crawley, Edmund, Martin Blomhoff Holm, and Håkon Tretvoll.** 2022. “A Parsimonious Model of Idiosyncratic Income.” Working Paper.
- Davis, Steven J. and John Haltiwanger.** 1992. “Gross Job Creation, Gross Job Destruction, and Employment Reallocation.” Quarterly Journal of Economics, 107(3): 819–863.
- Diamond, Rebecca.** 2016. “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000.” American Economic Review, 106(3): 479–524.
- Druehl, 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.
- Dynan, Karen.** 2012. “Is a Household Debt Overhang Holding Back Consumption.” Brookings Papers on Economic Activity, 43(1): 299–362.
- Engbom, Niklas and Christian Moser.** 2022. “Earnings Inequality and the Minimum Wage: Evidence from Brazil.” American Economic Review, 112(12): 3803–3847.
- 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.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer.** 2019. “The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment.” Journal of Political Economy, 127(6): 2836–2874.
- 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.** 2023. “Wealth, Race, and Consumption Smoothing of Typical Income Shocks.” Working Paper 27552. National Bureau of Economic Research.
- Goldin, Claudia.** 2021. Career and Family: Women’s Century-Long Journey toward Equity. Princeton University Press.
- Gottschalk, Peter, Robert Moffitt, Lawrence F. Katz, and William T. Dickens.** 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).
- 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. and David Loughran.** 2001. “Elderly Labor Supply: Work or Play?” Working Paper.
- 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.
- 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.
- Katz, Lawrence F. and David H. Autor.** 1999. “Changes in the Wage Structure and Earnings Inequality.” In Handbook of Labor Economics. Vol. 3, ed. Orley C. Ashenfelter and David Card, 1463–1555, Elsevier.
- 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.
- Kniesner, Thomas, W. Kip Viscusi, and James Ziliak.** 2006. “Life-Cycle Consumption and the Age-Adjusted Value of Life.” The B.E. Journal of Economic Analysis & Policy, 5(1): 1–36.

- 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.** 2023. “Work Hours Mismatch.” , (31205).
- 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.
- Lazear, Edward P.** 1986. “Salaries and Piece Rates.” The Journal of Business, 59(3): 405–431.
- Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent.** 2009. “Performance Pay and Wage Inequality.” Quarterly Journal of Economics, 124(1): 1–49.
- Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent.** 2012. “Contract Form, Wage Flexibility, and Employment.” American Economic Review, 102(3): 526–531.
- 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 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.
- Mueller, Holger M., Paige P. Ouimet, and Elena Simintzi.** 2017. “Wage Inequality and Firm Growth.” American Economic Review, 107(5): 379–383.
- 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.
- 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.
- Sockin, Jason.** 2022. “Show Me the Amenity: Are Higher-Paying Firms Better All Around?” Working Paper.
- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter.** 2019. “Firming Up Inequality.” Quarterly Journal of Economics, 134(1): 1–50.
- Sorkin, Isaac.** 2018. “Ranking Firms Using Revealed Preference.” Quarterly Journal of Economics, 133(3): 1331–1393.

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	2
B Income Model Appendix	4
C Appendix on Channels That Are <i>Not</i> Important for Explaining Earnings Volatility	6
C.1 Seasonality	6
C.2 Unpaid Leave	7
D Appendix on the Role of Firms in Earnings Instability	9
D.1 Interpreting Fluctuations in Total Firm Hours	9
D.2 Identification of Firm Effects on Income Volatility	10
D.3 Variance Decompositions	13
E Appendix on Spending Volatility	16
E.1 Median vs Variance Relationships	18
F Additional Appendix Figures and Tables	20

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 1.18 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 include 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 3.3 we discuss heterogeneity in earnings instability by various observable worker characteristics. Most of these are directly observable in PayrollCompany data. The one exception is information on worker occupation. 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 11, 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.⁵⁸ 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 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 5. Our data analysis in terms of spending and liquidity definitions and samples follows that in Ganong et al. (2023), 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

⁵⁸We winsorize the top 5 percent of volatility rather than imposing a symmetric winsorization because individual volatility near 0 is not an outlier.

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 2 we include accounts with multiple jobs, but since we want to look at the effect of job transitions on spending, for all of the analysis in Section 5.1 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.

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 as spurious if more than 40 percent (respectively 50 percent, 60 percent) of firm switches in a firm of size ≥ 20 (respectively 10-19, 5-9) move from the same origin to the same destination. 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.

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.⁵⁹ 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.⁶⁰

In Figure 3, we compare the monthly distribution of changes implied by this and several other income models to the data. Table A-4 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 (2022, 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-4 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 3.1, 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

⁵⁹This 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.

⁶⁰KMV 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.⁶¹ 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.

⁶¹Note 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 on Channels That Are *Not* Important for Explaining Earnings Volatility

C.1 Seasonality

Aggregate Seasonality Comparisons

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.

The Role of Seasonality for Individual Earnings Instability

In Section 4.1, we assess the degree to which monthly earnings volatility reflects predictable seasonal patterns. We estimate equation (2) using two different sets of predictors $X_{i,j,t}$ which are designed to capture firm-specific seasonality and annual recurring compensation. The first definition of $X_{i,j,t}$ captures firm-specific seasonality in pay by estimating firm by month fixed effects $\alpha_{j,m(t)}$ where $m(t)$ is an integer from 1 to 12 (e.g., January 2011 and January 2012 both have $m(t) = 1$). The second definition of $X_{i,j,t}$ captures annual recurring pay by using a 12 month lag interacted with month fixed effects: $\alpha_{m(t)} + \beta_{m(t)}(\log y_{i,j,t-12} - \log y_{i,j,t-13})$. Note that this second definition uses no firm-specific information at all, just the worker's own pay change from a year ago. However, because it requires information on the change in pay from 13 months ago to 12 months ago, it can only be constructed for workers who have been in the sample for at least 13 months.

Table A-9 shows that, across all specifications, income changes have a larger seasonal component for salaried workers who receive bonuses than for non-bonus salaried workers or hourly workers. Salaried workers with bonuses have R^2 s ranging from 0.25 to 0.39.⁶² Salaried workers without

⁶²We label 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 chosen to capture employees receiving annual bonuses. An employee who receives annual bonuses could potentially receive one bonus in 23 months of work but would receive two bonuses in 24 months of work. This approach categorizes 44 percent of salaried workers as bonus recipients.

bonuses and hourly workers have R^2 s ranging from 0.03 to 0.14. In unreported regressions we verify that the bonus component of pay is indeed the main reason why earnings fluctuations are more predictable for salaried workers with bonuses.⁶³

One other consistent pattern is that the R^2 is larger from the specifications with firm-month fixed effects than the specifications with lags. This is not driven by the 13-month tenure requirement because the R^2 is higher with firm $\times$ month fixed effects even when we drop the tenure screen (row 1). One possibility why the R^2 is higher is that this specification captures firm-specific seasonality in pay which can only be captured by estimating firm-specific coefficients. Another explanation is overfitting since the firm $\times$ month specifications include a very large number of parameters. The firm $\times$ month FE specification includes 12 coefficients for each firm (more than 30,000 coefficients in total), while the month + 12-month lag specification includes only 24 coefficients in total. Although requiring a minimum of 8 workers per firm means that there are still many more workers than parameters, there is still non-trivial risk of over-fitting. One piece of suggestive evidence that overfitting is indeed potentially driving some of this moderate R^2 is that the R^2 steadily decreases as the minimum firm size is increased from 8 to 20 to 30.

Although income changes are somewhat seasonal for salaried workers with bonuses, the main conclusion from Table A-9 is that most income changes are not driven by predictable seasonal patterns. For the hourly workers, the highest R^2 is 0.13 and some of this itself likely reflects overfitting. Put differently, only around ten percent of pay fluctuations appear related to predictable seasonal patterns.

C.2 Unpaid Leave

We use a simple algorithm to predict when workers take unpaid leave. We assume that unpaid leave occurs in the pay period with the lowest amount of hours paid over a six-month window. If the unpaid leave budget is not “exhausted” by raising the number of hours in the lowest pay period to match the number of hours in the second lowest pay period, then we assume that unpaid leave is also taken in the second lowest pay period. This algorithm continues until the unpaid leave budget is exhausted or hours are equalized across all pay period within the six-month window. Figure A-7 provides an illustration of the algorithm’s predictions for workers at one firm.

We stress that the estimates in this section are uncertain. Actual leave—paid or unpaid—typically occurs in discrete days or weeks. Our algorithm instead allocates hours of leave in a continuous fashion to the pay periods with the least hours (i.e. if someone works low hours in half of pay periods our algorithm assumes that they split their vacation time across those pay periods, when in reality it is probably more likely they took vacation in one of the pay periods and had low hours in the other pay periods for some other reason). This continuous approach likely leads us to, if anything, *overstate* the role of unpaid leave in driving earnings volatility.

Nevertheless, even though *unpaid* leave is unobserved, we provide some additional validation of

⁶³When we define the prediction target as an indicator for receipt of any bonus or as the amount of a bonus conditional any bonus receipt, we find that the R^2 is upwards of 0.6.

the timing assumption by showing that the same algorithm works well when we apply it to changes in *paid* leave, which are observable in the data. In Figure A-8 we show deciles of the distribution for the change in hours worked together with average paid leave hours for every decile. The red line in the figure shows that, on average, some paid leave is taken in every month and that workers are much more likely to take paid leave when hours worked are low. On the left side of the plot, in months where hours worked decrease, paid leave hours increase with a slope of approximately 0.4. The red line in the figure shows the joint distribution of the algorithm's prediction with the change in hours worked. It closely follows the blue line.

D Appendix on the Role of Firms in Earnings Instability

D.1 Interpreting Fluctuations in Total Firm Hours

In Section 4.2.1 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-9 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 4.2 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 100.⁶⁴ 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 0.97.

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

⁶⁴Note 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 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 (7). 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 (7) 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 that has outcome $Vol_{i,j} = median_i(|\Delta Y_{i,j,t}|)$ than in typical settings that have outcome $wage_{i,t}$. This is because our individual level volatility statistic necessarily requires using data from multiple observations over time, unlike individual wages, which can be easily calculated for 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 which is varying within an individual spell, we construct $Vol_t = median_t|\Delta Y_{i,j,t}|$ where t 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 t and i . Our primary volatility measure fixes an individual and measures the volatility of their earnings over time. This alternative volatility measure instead fixes a time period and then 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 t 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 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 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-10 panels (a) and (b) show that there are sharp changes in individual volatility around job transitions.⁶⁵ There is also no evidence of pre-trends. Furthermore, these 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. In Section 5, we provide evidence

⁶⁵Transitions from quartile 1 to 4 and vice versa are more unusual and so there is more sampling error and resulting noise for these groups.

that workers dislike earnings volatility, suggesting that matches with unusually high volatility would be less likely to form. This means that the probability of moving from a low volatility decile to a high volatility firm will go up as $\Delta\varepsilon_{ij}$ declines. That is, $E(\Delta\varepsilon_{ij}|\psi_{k'(j')} - \psi_{k(j)}) < 0$ for $j' > j$. 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 then observe small increases in observed volatility for these movers. This means we will then understate firm causal effects since we will 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 typical wage settings where a small match-specific draw will 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.3 shows that results are very similar if 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.3 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 (7). 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.⁶⁶ 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

⁶⁶This 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.

deviation of worker FE to the standard deviation of firm FE.⁶⁷

Table A-12 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 about 10 percent larger than 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.⁶⁸

Table A-12 also explores a number of robustness checks. As discussed above, the identification assumption in AKM regressions rules out match-specific sorting, where a change worker’s preferences for volatility causes them to switch to a job with different volatility. We provided empirical evidence in Appendix D.2 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 barely changes the conclusion about the size of firm fixed 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 (7) but splitting firms into 25 instead of 10 quantiles.⁶⁹ This shows that, indeed, estimated firm fixed effects rise very slightly when using these finer groups.⁷⁰

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

⁶⁷The 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.

⁶⁸As a point of comparison, estimating the same specification for log wages instead of volatility in our data delivers a 1 standard deviation effect of 20 percent. 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.

⁶⁹Note 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.

⁷⁰Using 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.

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 4.2.2 that focused on volatility in hours space.

Finally, results in Section 3.3 showed that younger 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.

E Appendix on Spending Volatility

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

Several of our specifications in Section 5 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 7, it is *not* the volatility of total firm hours measure used in Section 4.2.3. 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 Column 4 of Table 3 we instrument for individual volatility within job spell for stayers again using firm volatility. Since our baseline firm instrument is time invariant, we must modify the instrument for this “stayers” IV specification. To do so, we make two changes. First, instead of measuring the firm average of individual volatility $\overline{Vol_{j(i)}^y}$, we measure the pooled volatility of workers at the firm. That is, we measure the median changes in income across worker-months in the firm, rather than measuring the median changes at the worker level and then averaging these worker changes. Second, we measure this pooled firm volatility only using calendar months that correspond to the particular spell that a worker is at the firm, allowing for a time-varying measure. This is essentially the same volatility measure used in Appendix D.2 but where the dates pooled to construct firm volatility are those corresponding to the particular worker spell rather than months before and after a move. Concretely, this volatility measure for worker i in spell s at firm j is $Vol_{i,s} = median_{t \in s, k \in j(i)} |\Delta Y_{k,t}|$. That is, we collect the monthly changes for all worker-months in firm j that overlap with the months of the (pseudo) spell of worker i and then calculate the median absolute percent change across all of those monthly changes. This then delivers a firm volatility measure that varies across pseudo-spells within the same firm and can be used to instrument for individual worker volatility in that pseudo-spell. This instrument has very strong predictive power, just like the baseline instrument in our other regressions.

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 all of the results above as well as in the main text requires identification assumptions that are discussed in detail in the main text. As usual, these assumptions 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-11 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 around 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 before and after 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.

Ultimately, 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, since volatility is measured at the group level by pooling across individuals, this may introduce composition effects if the types of workers observed vary across event time.

To address this, we normalize the outcome 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:

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

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:⁷¹

$$|\Delta \tilde{y}_{i\tau}| = \sum_{\tau \neq -1} \beta_{\tau} \cdot \mathbf{1}\{\tau = t\} \cdot \text{Group}_i + \varepsilon_{i\tau},$$

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, $\mathbf{1}\{\tau = t\}$ is a dummy for each event time (with $\tau = -1$ omitted as the baseline), and β_{τ} captures how normalized volatility evolves for one group

⁷¹We 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} .

relative to the other.⁷²

The coefficients β_τ then deliver a standard difference-in-differences event study interpretation: it describes how volatility evolves before and after job transitions for households whose income volatility rises versus falls. Figure A-11 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-10. 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 β_τ estimates how the pooled mean of the normalized absolute changes, $\mathbb{E}(|\Delta\tilde{y}_{i\tau}|)$, 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 point in time, rather than measuring it within households over time; and (2) it uses the mean of the absolute changes, $Mean|\Delta\tilde{y}_{i\tau}|$, 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 (9). 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 (9) with these alternative volatility measures. The estimate at zero corresponds to the point estimate our preferred specification (Column 3 of Table 3) and Figure A-12 shows that while there is a modest decline in precision, point estimates are essential 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.

E.1 Median vs Variance Relationships

This appendix shows how to convert relationships estimated using a measure of dispersion based on medians to a measure based on variances necessary for willingness to pay calculations.

Let σ_c^2 be the variance of monthly spending changes, σ_y^2 be the variance of monthly income changes and assume that $\sigma_c^2 = \sigma_{c,0}^2 + \beta_\sigma \sigma_y^2$, where $\sigma_{c,0}^2$ is the variance of spending changes for a worker with no volatility of income and β_σ is the main parameter of interest, which gives the sensitivity of σ_c^2 to σ_y^2 and can be substituted directly into the willingness to pay formula. Our main empirical specifications estimate the relationship between $Median|\Delta_c|$ and $Median|\Delta_y|$ instead of estimating directly in variance space and so we must convert β_{med} to β_σ . If we assume that income shocks follow a normal distribution, then $Median|\Delta_y| \approx 0.6745\sigma_y$. This implies that moving from

⁷²Note that we cannot compute a within-job change at event time $\tau = 0$ (the first full month at the destination), so this period is excluded.

an income variance of 0 to an income variance of σ_y^2 generates the following change:

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

If we want to compute the willingness to pay to go from σ_y^2 to $= 0$, we can solve this equation for β_σ to get:


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

Alternatively, empirical relationships between σ_y^2 and σ_c^2 can be estimated directly in the data. We use medians for consistency with the rest of the paper, but we have estimated relationships directly in variance space and the resulting willingness to pay implications are very similar.


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 top 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: Pay 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 D_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




Notes: This figure shows the distribution of individual level volatility $Vol_i = Med|\Delta|$ in PayrollCompany data. The left panel shows the CDF and the right panel shows a histogram.

Figure A-7: 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-8: 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 earnings 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 2.24 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 2.24 percent of hours.

Figure A-9: 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-10: 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-11: Difference-in-Difference of Income and Spending Around Job Transitions: Groups are Those Increasing vs. Decreasing Income



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

Figure A-12: IV Regression of Vol^c on Vol^Y Excluding Months around Job Transitions



Notes: This figure re-estimates IV equation (9) 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	62%	58%	60%
Salaried	38%	42%	40%
All: bonus	37%	40%	36%
All: no bonus	63%	60%	64%

Notes: 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 more than $\frac{1}{24}$ 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 1% of all changes	-0.66	1.83	31%
Winsorize top/bottom 0.1% of nonzero changes	-0.93	13	60%
Winsorize top/bottom 0.5% of nonzero changes	-0.81	4	41%
Winsorize top/bottom 1% of nonzero changes	-0.71	2.39	34%
Winsorize top/bottom 5% of nonzero changes	-0.42	0.71	22%
Winsorize changes larger than 50%	Bottom 2% of data	Top 5% of data	20%

Notes: 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	0.30	0.74	0.05	0.17
	Less than 10 employees	0.27	0.69	0.03	0.15
	11-100 employees	0.28	0.75	0.04	0.16
	Greater than 100 employees	0.32	0.75	0.05	0.18
Household	None	0.29	0.78	0.05	0.18

Notes: Firm size is measured as the number of Chase customers who are employees of the firm.

Table A-4: Earnings Risk in Monthly Data versus Models Calibrated to Annual Data

	Monthly data	Model based on annual data			
		KMV (2018)	KV (2022)	MLM (2021)	CHT (2022)
P90 - P10 Δ	0.44	0.07	0.01	0.16	0.19
Share $ \Delta > 1\%$	0.65	0.30	0.11	0.87	0.89
Share $ \Delta > 20\%$	0.22	0.07	0.09	0.00	0.01
50th percentile $ \Delta $	0.05	0.00	0.00	0.04	0.05
75th percentile $ \Delta $	0.17	0.02	0.01	0.07	0.09
90th percentile $ \Delta $	0.39	0.13	0.12	0.10	0.12
Standard deviation	0.26	0.17	0.30	0.06	0.07
Kurtosis	6.96	31.1	14.1	3.01	3.01
Crow-Siddiqui kurtosis	4.98	104	421	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. KVM 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 (2022). 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-5: Lagged Median Summary Statistics of Earnings Changes

Sample	Variable	Share $\Delta \neq 0$	Median $ \Delta $	75th p $ \Delta $	Std. dev.	Share $\Delta > 0$	Share $\Delta < 0$
Hourly Full Time	Total Earnings	0.90	0.06	0.14	0.18	0.49	0.42
Hourly Full Time	Base Wage	0.17	0.00	0.00	0.03	0.17	0.00
Hourly Full Time	Hours	0.88	0.05	0.11	0.14	0.44	0.44
Salaried	Total Earnings	0.35	0.00	0.04	0.49	0.22	0.12
Salaried	Base Salary	0.12	0.00	0.00	0.09	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-6: Additional Differences in the Distribution of Earnings Changes

Sample	Positive persistence	Negative persistence	Skew	Kurtosis
All				
Q1	0.32	0.42	2.34	10.39
Q2	0.36	0.34	3.59	26.51
Q3	0.37	0.27	3.27	28.13
Q4	0.34	0.19	3.30	13.82
Hourly				
Q1	0.38	0.37	2.44	11.21
Q2	0.38	0.35	2.77	15.10
Q3	0.38	0.34	3.02	18.21
Q4	0.36	0.33	3.12	17.89
Salaried				
Q1	0.22	0.28	3.66	41.83
Q2	0.23	0.21	2.90	38.86
Q3	0.24	0.18	2.80	26.42
Q4	0.29	0.14	2.54	6.76

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.

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	13%	0.91	0.12	0.28	0.37	2.34	9.59	9%
Age 25 - 35	24%	0.75	0.06	0.17	0.29	3.22	18.73	34%
Age 35 - 45	21%	0.70	0.04	0.16	0.30	3.47	20.40	41%
Age 45 - 55	20%	0.67	0.04	0.15	0.30	3.61	21.39	44%
Age 55+	22%	0.63	0.03	0.14	0.30	3.61	20.99	46%

(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	0.38	2.23	8.78
Age 25 - 35	26%	0.93	0.09	0.20	0.29	2.80	15.70
Age 35 - 45	20%	0.92	0.08	0.19	0.28	2.93	17.30
Age 45 - 55	18%	0.90	0.07	0.18	0.27	3.05	18.49
Age 55+	18%	0.89	0.07	0.19	0.29	3.03	17.52

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.

Table A-8: Heterogeneity by Industry and Occupation

	All					Hourly			
Sample	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	0.11	0.25	0.33	0.88	14%	0.13	0.27	0.34	0.97
Arts and Entertainment	0.10	0.27	0.38	0.76	31%	0.15	0.33	0.42	0.96
Management	0.08	0.22	0.34	0.74	33%	0.13	0.26	0.32	0.99
Manufacturing	0.06	0.16	0.27	0.78	30%	0.08	0.18	0.26	0.95
Retail Trade	0.06	0.18	0.30	0.76	29%	0.09	0.21	0.31	0.93
Transportation and Warehousing	0.06	0.17	0.30	0.76	34%	0.09	0.20	0.29	0.92
Administrative and Support	0.05	0.16	0.30	0.74	33%	0.08	0.19	0.28	0.91
Construction	0.05	0.17	0.29	0.71	32%	0.08	0.19	0.27	0.90
Health Care and Social Assistance	0.05	0.17	0.30	0.76	24%	0.08	0.19	0.29	0.89
Utilities	0.05	0.17	0.31	0.77	30%	0.09	0.19	0.29	0.97
Educational Services	0.04	0.19	0.32	0.63	50%	0.13	0.30	0.39	0.91
Mining	0.04	0.15	0.23	0.67	36%	0.09	0.20	0.26	0.94
Wholesale Trade	0.04	0.15	0.33	0.68	42%	0.07	0.17	0.27	0.91
Agriculture	0.03	0.15	0.27	0.62	54%	0.10	0.22	0.29	0.94
Other Services	0.03	0.16	0.29	0.62	49%	0.10	0.24	0.34	0.92
Finance and Insurance	0.02	0.14	0.33	0.57	64%	0.08	0.20	0.30	0.90
Professional and Scientific Services	0.02	0.13	0.32	0.57	61%	0.08	0.20	0.30	0.91
Public Administration	0.02	0.12	0.26	0.61	56%	0.06	0.19	0.34	0.82
Real Estate	0.02	0.13	0.30	0.60	38%	0.05	0.16	0.27	0.78
Information	0.00	0.13	0.33	0.47	71%	0.10	0.25	0.35	0.87
Occupation									
Server	0.19	0.38	0.43	1.00	0%	0.19	0.38	0.44	1.00
Bartender	0.18	0.35	0.45	1.00	3%	0.18	0.35	0.45	1.00
Host	0.18	0.34	0.40	1.00	0%	0.18	0.34	0.40	1.00
Cook	0.10	0.21	0.29	0.99	5%	0.11	0.21	0.29	1.00
Operator	0.10	0.19	0.23	0.99	1%	0.11	0.19	0.23	1.00
Cleaner	0.09	0.20	0.27	0.98	0%	0.09	0.20	0.27	0.99
Driver	0.09	0.18	0.23	0.99	3%	0.09	0.18	0.23	1.00
Warehouse	0.09	0.18	0.19	0.99	1%	0.09	0.18	0.19	1.00
Sales	0.08	0.30	0.50	0.75	64%	0.12	0.29	0.37	0.99
Welder	0.08	0.16	0.19	0.99	0%	0.08	0.16	0.19	0.99
Maintenance	0.06	0.13	0.22	0.89	9%	0.06	0.14	0.22	0.96
Mechanic	0.06	0.15	0.21	0.87	10%	0.07	0.16	0.21	0.92
Medical assistant	0.06	0.15	0.21	1.00	1%	0.06	0.15	0.21	1.00
Porter	0.06	0.14	0.22	0.96	2%	0.06	0.14	0.21	0.95
Technician	0.06	0.16	0.28	0.82	25%	0.08	0.18	0.29	0.98
Customer service agent	0.05	0.14	0.23	0.87	19%	0.06	0.16	0.21	1.00
Administrative assistant	0.04	0.13	0.23	0.76	24%	0.05	0.15	0.24	0.87
Wage and salary administrator	0.04	0.11	0.24	0.75	29%	0.06	0.13	0.23	0.92
Analyst	0.02	0.10	0.27	0.63	82%	0.08	0.25	0.37	0.95
Accountant	0.01	0.09	0.20	0.55	97%	0.08	0.19	0.26	0.97
Consultant	0.00	0.17	0.38	0.47	87%	0.13	0.25	0.31	0.98
Engineer	0.00	0.06	0.34	0.44	85%	0.10	0.22	0.26	0.96
Financial controller	0.00	0.08	0.36	0.35	93%	0.11	0.42	0.63	1.00
Manager	0.00	0.09	0.33	0.48	83%	0.07	0.16	0.23	0.89
Teacher	0.00	0.06	0.20	0.44	97%	0.14	0.35	0.43	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.

Table A-9: 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	0.08	0.10	0.33
Firm x month FEs ($\alpha_{j,m(t)}$),	0.13	0.14	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	0.05	0.25
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 (2). 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-10: 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	0.14	0.18	0.46	0.43
Earnings Fluctuations After Removing Unpaid Leave	0.88	0.05	0.12	0.16	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.

Table A-11: 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	0.16	0.28	3.50	21.83	38%
Men without children	22%	0.81	0.07	0.20	0.31	3.03	16.31	26%
Women with children	27%	0.78	0.06	0.17	0.27	3.10	19.04	30%
Women without children	24%	0.80	0.07	0.20	0.31	2.95	15.75	25%
Men	49%	0.77	0.06	0.17	0.29	3.27	19.02	33%
Women	51%	0.79	0.06	0.18	0.29	3.04	17.45	27%

(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	0.08	0.18	0.26	2.97	18.88
Men without children	23%	0.96	0.10	0.22	0.32	2.73	14.03
Women with children	27%	0.94	0.09	0.20	0.28	2.78	15.98
Women without children	26%	0.93	0.09	0.22	0.32	2.67	13.29
Men	47%	0.95	0.09	0.20	0.29	2.87	16.35
Women	53%	0.94	0.09	0.21	0.30	2.74	14.65

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.

Table A-12: 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	5654	0.14	0.11	0.11	0.048	0.045	1.07	6.32e-04
Only within Industry Moves	2744	0.14	0.11	0.11	0.050	0.047	1.06	5.92e-04
Only within Wage Decile Moves	3254	0.14	0.11	0.11	0.053	0.047	1.12	5.12e-04
Exclude Accom. and Food	4691	0.13	0.10	0.11	0.048	0.044	1.09	5.51e-04
25 Firm Groups	6639	0.14	0.11	0.11	0.049	0.045	1.08	6.59e-04
Hours Vol instead of Pay vol	5654	0.13	0.09	0.11	0.048	0.046	1.05	6.87e-04
Prime Age Only	4877	0.13	0.10	0.10	0.045	0.042	1.08	4.53e-04

Notes: This table reports variance decompositions using the fixed effect specification in equation (7) with the leave-one-out sampling correction of Kline, Saggio, and Sølvsten (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.

Table A-13: The Effect of Income Volatility on Consumption Volatility: Heterogeneity

	Dependent Variable: Med $ \%C $			
	Hourly vs. Salaried	Low vs. High Checking	Low vs. High Income	
	(1)	(2)	(3)	(4)
Med $ \%Y $	0.236*** (0.012)	0.325*** (0.017)	0.213*** (0.021)	0.245*** (0.020)
Med $ \%Y $: Salaried	-0.210*** (0.019)			
Med $ \%Y $: High Checking		-0.145*** (0.022)		
Med $ \%Y $: High Income			-0.036 (0.025)	
Med $ \%Y $: Age 35-50				0.018 (0.024)
Med $ \%Y $: Age > 50				-0.036 (0.026)
Excluded Group: Observations	Hourly 396,830	Lowest Checking Tercile 153,097	Lowest Income Tercile 153,097	Age 18-35 224,097

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Notes: Standard errors are clustered by firm. Columns 2 and 3 drop the middle tercile and compare the highest and lowest terciles. All specifications except Column 1 restrict to hourly workers.

Table A-14: The Effect of Income Volatility on Consumption Volatility: Robustness

	Dependent Variable: Med $ \%C $					
	Nondurable (Baseline)	Work	Non-Work	Total Spend	Nondurable	Nondurable
	(1)	(2)	(3)	(4)	(5)	(6)
Med $ \%Y $	0.236*** (0.012)	0.081*** (0.017)	0.302*** (0.021)	0.235*** (0.020)	0.300*** (0.016)	0.205*** (0.010)
%Y Frequency	Monthly	Monthly	Monthly	Monthly	Quarterly	Monthly
%C Frequency	Monthly	Monthly	Monthly	Monthly	Quarterly	Quarterly
Observations	228,502	219,587	220,170	228,502	169,422	169,422

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Notes: Standard errors are clustered by firm. Work and Non-Work expenses are categorized based on those which are most sensitive to retirement, following Ganong and Noel (2019). Total spend includes all account outflows except transfers to other financial accounts.

Table A-15: The Effect of Income Volatility on Separation Rates

	(1)	(2)
Med $ \% \Delta Y_{ij} $	3.19*** (0.331)	1.21** (0.536)
No. Obs	46,240	23,145
No. Firms	2,575	2,575
Reduced Form?	Yes	Yes
Instrument?	Yes	Yes
Sample	Hourly	Salaried
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.