Earnings Instability*

Peter Ganong, Pascal Noel, Christina Patterson, Joseph Vavra, and Alexander Weinberg
University of Chicago and NBER

March 21, 2024

Abstract

We analyze monthly earnings volatility using administrative payroll data. While it is well-documented that wages are largely stable, we find that this wage stability does not translate into earnings stability for most U.S. workers. Even within stable employment relationships, and even when wages are constant, many workers nevertheless face substantial monthly earnings volatility. The standard deviation of monthly earnings changes is 28%, while the standard deviation of base wage changes is only 2%. There is substantial heterogeneity in this volatility, with much higher volatility for hourly workers than for salaried workers. This degree of volatility is far higher than what is implied by benchmark models of earnings processes which are calibrated to previously-available annual data and used as inputs for leading macro models. To understand the welfare consequences of pay volatility, we estimate the amenity value of volatility using worker quits in a model of a frictional labor market. We find that workers have a high willingness to pay to reduce earnings volatility. Overall, this analysis shows that high-frequency labor market shocks are an important source of risk and fragility which has been masked by past studies of annual earnings.

*We thank Fiona Greig for the initial early conversations which inspired this paper. We thank Niels Gormsen, Simon Jaeger, Emir Kamenica, Greg Kaplan, Anil Kayshap, Thibaut Lamadon, Susan Lambert, Marta Lachowska, Alan Manning, Evan Rose, Jon Roth, Danny Schneider, and Steven Woodbury for helpful conversations, and seminar participants at the U.S. Census Bureau, CEPR/INET, the Chicago Fed, Chicago Booth, Chicago Harris, the Cleveland Fed, CU Boulder, GW, LSE, the St. Louis Fed, and USC for suggestions. We thank the Data Science team at PayrollCompany and Daniel Sullivan and Sarah Kuehl at JPMorganChase Institute (JPMCI). We thank Rupsha Debnath, Avik Garg, Theodore Grayer, Jay Leeds, Michael Meyer, John Spence, and Madeline Zuckerman for excellent research assistance. We thank Marta Lachowska and Stephen Woodbury for discussions about and analysis of the sources of pay volatility in the Washington state data. 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.

1 Introduction

This paper provides the first estimates of monthly earnings volatility from administrative data on U.S. workers. Our main finding is that workers face substantial fluctuations in earnings from one month to the next. These fluctuations are frequent and large, even within stable employment relationships, and even when wages are constant.

This is relevant for two reasons. First, it is well-known that wages are largely stable and, when they do adjust, almost always adjust upwards (Grigsby, Hurst and Yildirmaz, 2021). Furthermore, at an annual level earnings volatility arises mostly from transitions across rather than within jobs (Guvenen et al., 2021). Our results suggest that wage stability does *not* translate into earnings stability for most U.S. workers, and that intensive-margin fluctuations are an important source of within-year earnings risk.

Second, the level of within-year earnings risk is a crucial determinant of household behavior in modern heterogeneous agent macro models. However, because of a lack of high-quality data directly measuring such risk, the literature has been forced to infer it based on annual data. We find that benchmark models of the earnings process which are calibrated to previously-available annual data substantially underestimate the degree of monthly fluctuations actually faced by workers.

To measure earnings volatility we use administrative data from both the firm side (via a payroll processor) and the worker side (via deposits into Chase bank accounts). Both datasets are representative of U.S. workers as a whole on several different metrics of labor market arrangements.

In the first part of the paper, we use this data to show that monthly earnings volatility is substantial. The standard deviation of monthly earnings changes is 28% *within* employment spells. In about three quarters of months, workers receive a different amount of pay than they received the prior month. In one quarter of months the change in pay is at least 17%.

This magnitude of earnings volatility is large relative to two natural comparison points from the prior literature. First, monthly earnings are far more volatile than monthly wages. The standard deviation of monthly earnings changes is ten times as large as the standard deviation of monthly wage changes. This is both because wages change much less often than earnings and because, conditional on a change occurring, wages usually change by a small amount while earnings change by a large amount. Second, monthly earnings are also more volatile than what is implied by benchmark models of the earnings process which have been calibrated to annual earnings data.

The second part of the paper shows that the nature of earnings volatility differs markedly between salaried and hourly workers in terms of its level, (a)symmetry, and sources. For salaried workers, pay rarely varies from month to month. When pay does vary, it usually varies to the upside. Moreover, much of this upside volatility appears to be driven by performance pay such as bonuses and commissions.

In contrast, hourly workers face substantial monthly earnings volatility. These workers

account for 60% of the U.S. workforce and rarely have the same earnings from one month to the next. On average, in 11 out of 12 months each year hourly workers' earnings change relative to the prior month. These changes are often quite large: in one quarter of months earnings change by at least 21%. The changes are also symmetric, such that earnings vary as much to the downside as they do to the upside. Virtually all of this earnings volatility is driven by fluctuations in hours. Thus, although hourly workers' contracts usually feature a stable base hourly wage, this contract does not translate into stable earnings from month to month.

The third part of the paper examines the relationship between volatility and risk and finds that most earnings *volatility* appears to be *risk* from the worker's perspective. We take two steps to try to purge earnings volatility which does not reflect risk.

First, we examine the extent to which monthly earnings fluctuations reflect predictable seasonal patterns. We do this by studying the R^2 from regressions seeking to explain individual workers' monthly income changes using variables designed to capture firm-specific seasonality or worker-specific annual recurring compensation patterns. For the 17% of workers who are salaried and receive regular bonuses, we find that between one-third and one-half of volatility is seasonal. However, for the other 83% of workers, earnings fluctuations are almost entirely unrelated to predictable seasonal patterns.

Second, we quantify the extent to which downward fluctuations for hourly workers reflect endogenously-chosen unpaid leisure rather than fluctuations imposed on the workers. If unpaid leisure is chosen by the worker then it is a source of reduced compensation but not a source of income risk. In order to examine endogenous labor supply decisions we develop an algorithm to predict vacation timing. The algorithm is able to reproduce the joint distribution of hours worked and paid vacation hours. We then use it to predict when workers are likely to have taken unpaid vacation and recompute volatility after excluding this unpaid vacation. As with predictable fluctuations, we find that unpaid leisure accounts for a minimal fraction of the overall volatility we document. Combined, these analyses suggest that earnings volatility represents meaningful risk from the worker's perspective.

One limitation of our analysis in this part of the paper is that we have only a narrow lens for distinguishing between a number of competing explanations for the substantial amount of remaining non-predictable, non-leisure driven earnings volatility. It could arise for worker-side reasons such as child care needs, sick leave, medical leave, and family leave. It could also arise because of firm-side reasons such as fluctuations in the firm's labor demand, indivisibility of shifts, lack of managerial capacity to aggregate worker preferences, and downstream effects of worker-side absences. Based on limited information from worker surveys and from what human resources departments record in the PayrollCompany software, we believe that about half of the variation is driven by firms and half is driven by workers. We view distinguishing between these explanations for volatility as a crucial topic for future research. However, what these explanations have in common is that they all look like income risk from the worker's

perspective.

In the final part of the paper we examine the economic implications of monthly earnings volatility. In principle, earnings volatility could be an amenity if it reflects job flexibility or a disamenity if it indeed reflects risk. We use revealed preference based on job transitions to infer how workers value the bundle of characteristics at one job relative to another. Presumably if workers quit one job and accept an alternative one this means that they prefer the utility bundle offered by the new job. We therefore estimate the (dis)amenity value of pay volatility using worker quits in a standard model of a frictional labor market.

We find that workers have a high willingness to pay to reduce earnings volatility. Two pieces of evidence support this conclusion. First, we are able to follow a subset of workers across multiple employment spells. We find that workers who quit high-volatility jobs transition to new jobs with lower volatility. Second, workers quit firms with high pay volatility at about four times the rate at which they quit firms with low pay volatility. In a Burdett-Mortensen model, this estimated elasticity of quits with respect to volatility, combined with estimates of the elasticity of quits with respect to wages in the prior literature, together imply a willingness to pay to eliminate volatility equal to 12-25% of wages. However, we caution that the quit elasticities we measure are not causal and so this willingness-to-pay estimate is only suggestive.[1]

Our findings contribute to several active literatures in economics. To our knowledge, these are the first estimates of monthly earnings volatility based on administrative U.S. data, but they are similar to other existing estimates based on surveys in the U.S. and administrative data from other countries. Hannagan and Morduch (2015) surveys 235 low- and middle-income U.S. households on a monthly basis for one year and finds substantial amounts of income instability for this sample. Druedahl, Graber and Jørgensen (2023) estimates the volatility of monthly earnings in Denmark. That paper also finds substantial volatility, though earnings changes in Denmark are both less frequent and smaller than the earnings changes we document in the U.S. We complement these papers by examining earnings volatility for a representative U.S. population, exploring heterogeneity in the sources and characteristics of volatility across employment contract types, and investigating the extent to which this volatility reflects risk.

Our finding of widespread within-year earnings instability is at odds with the within-year implications of theoretical models based on annual data. Estimates of within-year earnings risk are central to modern heterogeneous agent macro models for two main reasons. First, these models are designed to make predictions about within-year consumption behavior. Second, the frequency of earnings shocks determines the household's optimal level of liquid savings, 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

[1]Other work has found quasi-experimental evidence that pay volatility does indeed make workers more likely to quit in the retail (Kesavan and Kuhnen 2017) and home health (Bergman, David and Song 2023) sectors.

forced to infer this risk by calibrating high-frequency earnings processes to low-frequency annual data.

We show that this calibration method often inadvertently generates estimates of within-year earnings risk which are smaller than what workers actually face. We compare our data on monthly earnings volatility to the monthly implications of benchmark models used in Kaplan, Moll and Violante (2018), Laibson, Maxted and Moll (2021), Kaplan and Violante (2022), and Crawley, Holm and Tretvoll (2022). This exercise shows that sizeable monthly earnings changes are much more common in the data than what is implied by these models. For example, the 75th percentile of the absolute percent change in earnings is between 1% and 9% in the models (as compared to 17% in the data), and the 90th percentile is between 10% and 13% in the models (as compared to 39% in the data).

Our estimates for the willingness to pay to eliminate pay volatility contribute to an active literature which uses a range of methods to estimate workers' valuation of job amenities. Maestas et al. (2023) use workers' stated preferences over job vignettes to compute the willingness to pay for several job attributes. They find that workers will accept wage reductions of 23-26% to move from a job with no paid time off to one with 20 days of paid time off. Because jobs with generous paid leave also generally have low earnings volatility, such a high willingness to pay for this kind of job using a stated-preference approach is consistent with the willingness to pay for eliminating volatility that we estimate using a revealed-preference approach.

Methodologically, our use of job transitions to estimate workers' valuation of amenities builds on several recent studies in labor economics.[2] Jarosch (2023) studies a model where firms differ by layoff risk and workers climb a job ladder with respect to moving to firms with lower layoff risk over time and documents empirical evidence consistent with this model in Germany. Lachowska et al. (2023) use job transitions and a discrete-choice model to document that many workers prefer jobs with more hours than they are offered. Other recent studies using job transitions to value amenities include Hoffman and Tadelis (2021), Lehmann (2023), Manning (2003), and Sorkin (2018). We complement this prior evidence on worker preferences for income risk and average hours with evidence on a conceptually distinct but complementary channel: worker preferences for hours volatility.

Overall, the analysis in this paper shows that workers face substantial monthly earnings volatility, that this volatility primarily reflects unpredictable risk rather than predictable changes or endogenous choices, that workers dislike this risk, and that the degree of this risk varies substantially across different types of employment contracts. This analysis suggests

[2]Job transitions are only one of several methods which have recently been used to measure valuation of job amenities. Bloemen (2008) measures workers' preferences for hours by combining survey response over desired hours with actual hours in a job search model. Rose and Shem-Tov (2023) documents persistent earnings losses after displacement for low-wage workers at high hours jobs and interprets the estimates as informative about low-wage workers' preferences for hours through a job search model. Mas and Pallais (2017) uses discrete-choice experiments to measure worker valuation of scheduling flexibility. Dube, Naidu and Reich (2022) uses hypothetical quit elasticities to estimate the willingness to pay for several different low-wage job amenities.

that high-frequency labor market shocks are an important source of risk and fragility which has been masked by past studies of annual labor market earnings.

2 Data and Benchmarking

Our primary data source consists of de-identified administrative earnings records from an anonymous U.S. payroll processor. Hereafter we refer to this processor as "PayrollCompany".

We work with PayrollCompany data spanning the years 2010 to 2023. During this time, PayrollCompany provided services to 2 million client firms. For any month that a firm is in the dataset, the data includes paycheck information for the universe of workers at that firm, covering between 2 million and 4 million workers each month.

Most PayrollCompany clients are small firms, though we use additional data (described below) to assess the external validity of our results for workers at larger firms.

Most of our study relies on the table containing information on paychecks. In this table, the fundamental unit of observation is the "pay item". The eight most common pay items are, in descending frequency order: base pay, bonuses, overtime, paid holiday, paid vacation, tips, commissions, and paid sick leave. For salaried workers we observe contracted pay per pay period while for hourly workers we observe a contracted hourly wage. We refer to both these variables as the "Base Wage".

We take two steps to clean the data. First, we focus on workers who can reliably be classified as hourly or salaried. To accomplish this, we drop workers without wage or salary data (11% of worker-months). We also drop workers who have an hourly wage which changes in more than half of their months worked (4% of worker-months). Of the worker-months we can reliably classify, 60% are hourly and 40% are salaried. Second, we drop the 0.2% of worker-months where the worker is paid for more than 400 hours out of a concern that this reflects the worker receiving back pay for more than one month of work.

In some of the analysis, we disaggregate salaried and hourly workers further. Among salaried workers, we separate between those that do and those that do not receive pay items tagged as bonuses. We 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.[3] This approach categorizes 44% of salaried workers as bonus recipients.

Among hourly workers, we stratify the sample into quartiles by average hours worked per week. The bottom quartile approximately corresponds to part-time workers while the top three quartiles correspond to workers working an average of at least 30 hours per week.[4]

[3] 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.

[4] While the U.S. Department of Labor does not list an exact definition of full-time employment, the Affordable Care Act defines a full-time job as one that requires employees to work a minimum of 30 hours per week.

We also observe some information about job spells, workers, and firms. We observe dates when workers are temporarily inactivated or permanently terminated in the payroll software. When a worker is inactivated or terminated in most cases the employer records the reason for the status change. We observe worker age, gender, and number of dependent children (as measured from IRS Form W-4). On the firm side we observe the firm's industry.

We conduct most of our analysis at the monthly level because this enables us to combine data on workers who are paid at different frequencies. If we were to literally aggregate paychecks to the monthly level and analyze total monthly pay, this would introduce an extra source of volatility for workers who are paid weekly or fortnightly. Because each year has 52.14 weeks, workers who are paid weekly get four checks in roughly eight months and five checks in the remaining four months. Similarly, workers who are paid fortnightly get three checks in roughly two months and two checks in the remaining ten months. To put workers at all pay frequencies on a common footing, we define "Total Earnings" as total earnings divided by the number of paychecks in each month. Therefore, we will count a worker who receives the same pay every paycheck as having no change in "Total Earnings" from one month to the next even if she receives five paychecks one month and four the next.

We are also able to follow workers who transition between multiple firms that are clients of PayrollCompany from August 2017 onward. We follow workers using an encrypted variable; no personally identifiable information (PII) is used in this research. In the transition sample we focus on workers who are employed by exactly two firms using PayrollCompany from 2017-2023, have tenure of at least four months at both firms, and have a gap of employment of more than one day but less than one year between the origin firm and the destination firm. We observe 69,900 such transitions.[5] For most of the transitions analysis, we restrict the sample to workers who are aged 25-54 at separation.

The PayrollCompany data appears to be representative of the prices of labor, the quantities of labor, and contract types for U.S. workers overall. In terms of the price of labor, Figure A-1 shows that the hourly wage distribution for PayrollCompany workers is similar to the wage distribution for all workers (panel a), with some evidence of modest positive selection among hourly workers relative to the population of hourly workers at small firms (panel b). In terms of the quantity of labor, Figure A-2 shows that the distribution of hours worked per quarter for PayrollCompany workers is similar to the population of workers in Washington state, and Figure A-3 shows that the monthly separation and hire rates are also similar. Furthermore, the relationship between quarterly fluctuations in hours and quarterly fluctuations in earnings at both the worker and firm level are similar in PayrollCompany and in Washington state earnings records (Table A-1).

[5]The sample of workers who change jobs are not representative of the labor force as a whole. As has been documented in prior work—e.g. Farber (1993)—lower-earning workers tend to have higher turnover and thus appear more frequently in our sample of job-switchers. Workers we observe separating from one firm and joining another are on average younger, lower income, and have shorter tenures (as measured at the origin firm). About 44% of workers in the transition sample at the origin firm are in the bottom quartile of hours in the main sample.

Workers at PayrollCompany have a similar distribution of contract types to U.S. workers overall. Table A-2 shows that the share of hourly versus salaried and the share that get bonuses are similar. Figure A-4 shows that the distribution of pay frequency (weekly, fortnightly, semimonthly or monthly) is similar as well.

Our analysis uses three samples. The broadest sample uses all workers with two complete calendar months of employment such that we can compute a change in pay from the prior month. This sample includes 25 million workers and 511 million worker-months. A second sample drops part-time hourly workers where part-time is defined as working less than an average of 30 hours per week. This removes 21% of the sample, leaving 17 million workers and 405 million worker-months. Finally, a third sample uses firm pay history to predict changes in workers' pay; this sample restricts to firms persisting in the data for at least three years with an average of at least eight workers over all months in the data and drops workers who experience a monthly change in log income within the largest or smallest 0.1% of months. This sample includes 1.1 million workers and 30 million worker-months. In most of the analysis, we work with a 1% or 5% sample of firms for computational tractability; the exception is that when we analyze workers who switch firms we use every single worker who switches between two companies that use the PayrollCompany software.

We supplement the PayrollCompany data with data on Chase bank customers from the JPMorganChase Institute (JPMCI) and data on U.S. workers from the Current Population Survey (CPS). Neither of these datasets is linked with PayrollCompany data; we simply compare distributions of earnings characteristics between different datasets. The JPMCI data are useful because we can compare pay volatility at both small firms and large firms and because we can understand the extent to which job-level earnings volatility translates into household-level earnings volatility. 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. The CPS is useful because it asks questions about unpaid vacation, which are not captured in the PayrollCompany data.

3 How High is Monthly Earnings Volatility?

This section describes the empirical patterns of monthly earnings volatility. Our main finding is that monthly earnings volatility is substantial in payroll data. We then explore the sources of this volatility and find significant heterogeneity between hourly and salaried workers. Finally, we show that volatility is just as high in bank account data as in payroll data, and we explore additional dimensions of heterogeneity in the bank data which we cannot observe in the payroll data.

3.1 Main Finding: Workers Face Substantial Earnings Volatility From Month to Month

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 histogram. Summary statistics from this distribution are shown in Table 1. These exhibits show that earnings changes are ubiquitous. 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 one-quarter of months, the change in pay is at least 17%.

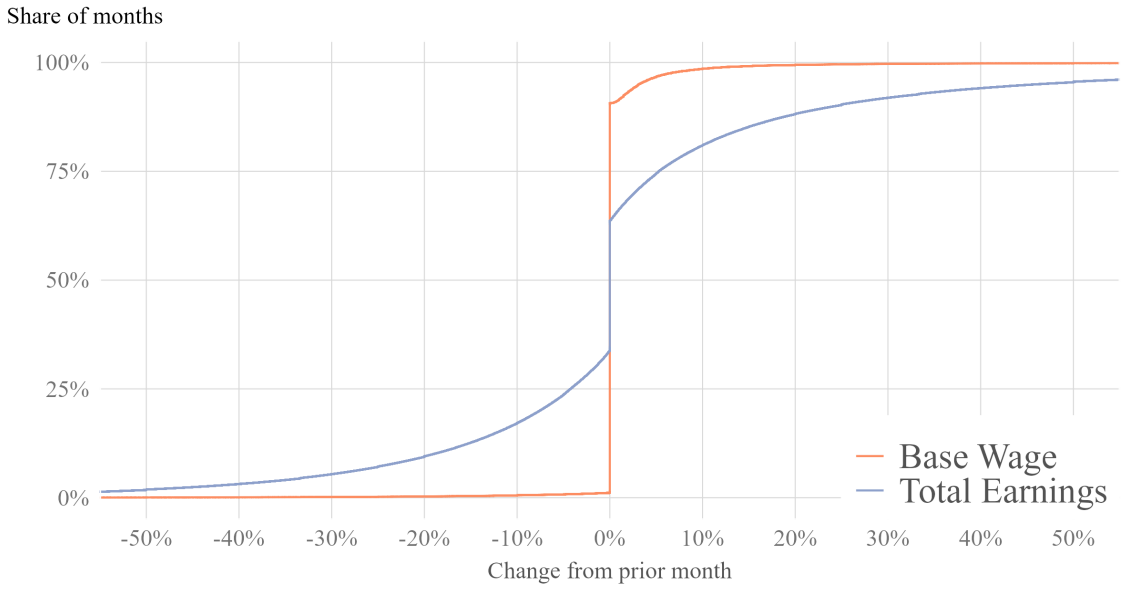
Figure 1 also shows that within-job changes in total earnings are much larger than within-job changes in base wages. This is true both on the extensive and intensive margins. In only 10% of months a worker will see a change in their base wage (hourly wage or per-pay-period base salary) from the prior month, but in 69% of months a worker will see a change in their earnings for the prior month.[6] On the intensive margin, Figure 1b shows that the magnitude of earnings changes is much larger than the magnitude of changes in base wages. Summarizing across intensive and extensive margins, the standard deviation of earnings changes is 28% while the standard deviation of wage changes is only 2%. 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.

One limitation of Figure 1 is that it combines data across 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. We therefore compute each worker's median absolute change and standard deviation. Figure A-5 shows the distribution of these worker-level statistics. There is substantial heterogeneity across workers by both measures. In addition, Figure A-6 shows substantial heterogeneity across industries. This motivates additional analysis of the data to uncover the drivers of systematic heterogeneity, beginning with stratifying by hourly versus salaried workers.

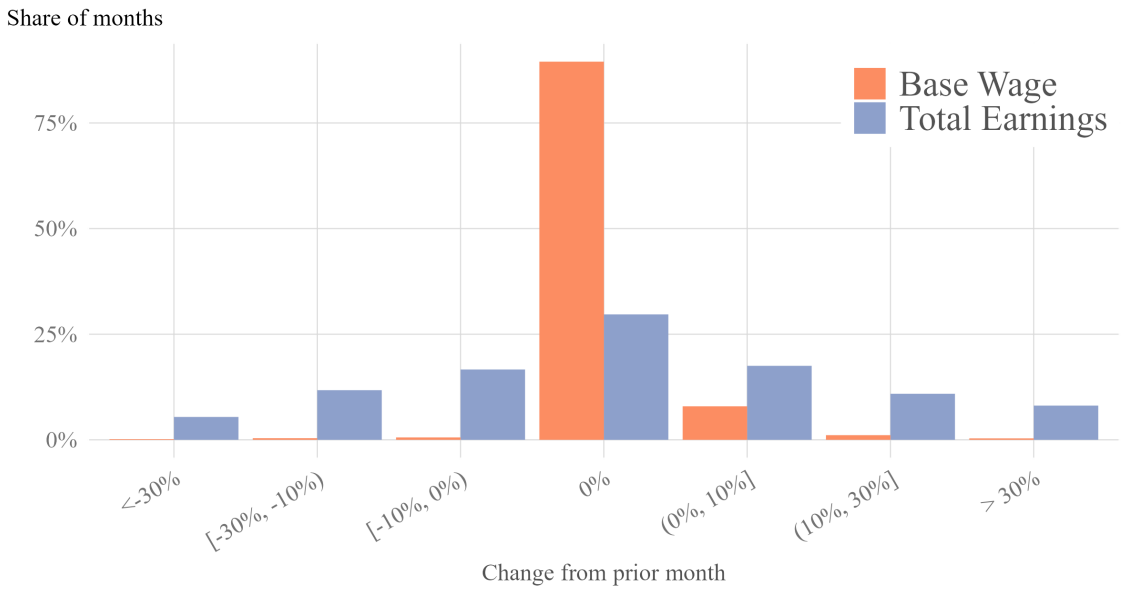
[6]Our finding that workers see base wage changes in only 10% of months is in line with the estimates on the frequency of wage adjustment from Grigsby, Hurst and Yildirmaz (2021). The 69% estimate of the share of months with changes in total earnings abstracts from pay schedule-induced fluctuations, as discussed in Section 2. If we were to include those the statistic would be even higher.

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 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.

Table 1: Summary Statistics of Earnings Changes

Sample	Variable	Share $\Delta \neq 0$	Median $\lvert\Delta\rvert$	75th percentile $\lvert\Delta\rvert$	Standard Deviation	Share $\Delta > 0$	Share $\Delta < 0$	Skewness	Kurtosis
Δ: First Difference									
All	Total Earnings	0.69	0.04	0.17	0.28	0.36	0.33	2.67	12.6
All	Base Wage	0.10	0	0	0.02	0.09	0.01	5.11	27.79
Hourly	Total Earnings	0.91	0.09	0.21	0.29	0.47	0.44	2.25	9.95
Hourly	Base Wage per Hour	0.12	0	0	0.02	0.11	0.01	4.76	24.47
Salaried	Total Earnings	0.35	0	0.05	0.28	0.19	0.15	3.32	16.74
Salaried	Base Salary	0.07	0	0	0.02	0.06	0.01	5.71	34.06
Δ: Lagged Median									
Hourly Full-time	Total Earnings	0.90	0.06	0.14	0.18	0.49	0.42	1.62	6.27
Hourly Full-time	Base Wage per Hour	0.18	0	0	0.03	0.17	0.01	3.66	14.22
Hourly Full-time	Hours	0.87	0.05	0.11	0.14	0.43	0.44	0.97	4.94
Salaried	Total Earnings	0.34	0	0.04	0.32	0.22	0.12	5.03	30.06
Salaried	Base Salary	0.11	0	0	0.04	0.1	0.01	4.74	28.03
Δ: Difference from Base Wage									
Salaried	Total Earnings	0.30	0	0.04	0.67	0.25	0.05	5.42	32.3
Δ: First Difference (Net Pay)									
All	Total Earnings	0.76	0.05	0.17	0.28	0.39	0.37	2.64	12.4
All: Chase Bank	Total Earnings	0.74	0.04	0.17	0.31	0.41	0.39	5.70	40.09

Notes: This table reports distributional statistics of first differences, differences relative to lagged median, and differences relative to base wage. The percent change of first differences for variable y is defined as $\frac{y_t - y_{t-1}}{y_{t-1}}$. The percent change relative to the lagged median is defined as $\frac{y_t - Median(\{y_s\}_{s\in t-1,t-2,t-3})}{Median(\{y_s\}_{s\in t-1,t-2,t-3})}$. The percent change relative to the base wage is defined as $\frac{TotalEarnings_t - BaseSalary_t}{BaseSalary_t}$. All rows report pre-tax earnings except the last two rows which report pay net of withholding and deductions. All data is from PayrollCompany except the last row which is data on Chase bank customers from JPMorganChase Institute (JPMCI). Hourly full-time defined as those who work an average of at least 30 hours per week. "Base Wage" in row 2 is defined as base wage per hour for hourly workers and per-pay-period base salary for salaried workers. Before computing higher-order moments (standard deviation, skew, kurtosis), the measures of change are winsorized at the 1st and 99th percentiles of non-zero changes. Table A-3 shows the standard deviation using alternative thresholds. PayrollCompany and JPMCI data are analyzed separately and were not merged as part of this analysis.

3.2 Heterogeneity in Earnings Volatility Between Hourly and Salaried Workers

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

Levels: Earnings volatility is larger for hourly than for salaried workers. Figure 2 shows that earnings fluctuations are the norm for hourly workers, while they are comparatively rare for salaried workers. Although both hourly and salaried workers have similarly infrequent changes in base wages, earnings vary much more often for hourly workers. For hourly workers, earnings change in 91% of months. So in 11 out of 12 months a year an hourly worker will have different monthly earnings relative to the prior month. These changes are often substantial in magnitude: pooling monthly changes across all hourly workers the median absolute change is 9% and the 75th percentile absolute change is 21%.

Although still volatile, salaried workers' pay is much more stable than hourly workers. They receive exactly their base pay in 70% of months. Because base pay changes infrequently, the median absolute change in total earnings is zero and the 75th percentile absolute change is 5%. However, conditional on a change occurring, their pay is more volatile. Thus the unconditional standard deviation is very similar for hourly and salaried workers and the unconditional kurtosis of pay changes is higher for salaried workers than hourly workers.

Sources: Hours instability drives earnings volatility for hourly workers. Figures 3 and 4 assess the sources of earnings changes and the prevalence of asymmetry. To make the possibility of asymmetry legible and to eliminate the "double counting" of temporary income changes, we measure the change in each variable relative to the median of the last three months (rather than the change from the prior month). The problem motivating this change is that a temporary increase in earnings in month t will show up as both an increase from $t-1$ to t and a decrease from t to $t+1$. If we instead use the change relative to the median of three lags, this eliminates the decrease from t to $t+1$. Figure 3 shows the cumulative distribution function and Figure 4 shows a binned histogram of the same series.

Figure 3a shows that monthly hours fluctuations are the key driver of monthly earnings fluctuations for hourly workers. The line for "Hours" is quite similar to the line for "Total Earnings." It is useful to emphasize that this finding is not mechanical. Hourly workers can have earnings volatility for three reasons other than changes in hours: from changes in base wage, from overtime, and from changes in non-hours-linked compensation (e.g. bonuses).

Our ability to assess the sources of earnings volatility for salaried workers is limited by data constraints. Figure 3b essentially recapitulates the results from Figure 2 that fluctuations other than base wage are responsible for most earnings fluctuations. Unfortunately, we do not have the ability to understand what exactly is shifting in the employment relationship from month to month because we do not observe hours for salaried workers and reporting of bonuses is incomplete.[7]

[7]When an employer wants to pay a bonus they can either enter a single larger number for the employee's

Figure 2: Within-Job Volatility for Hourly and Salaried Workers

(a) Cumulative Distribution Function



(b) Histogram



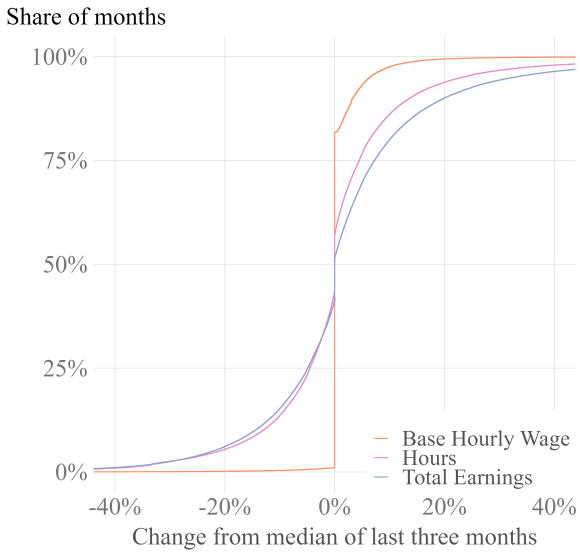
Notes: This figure shows the within-job distribution of the change in earnings and wages from the prior month. "Total Earnings" is average pay per paycheck to abstract from pay schedule-driven fluctuations. "Base Wage" is the hourly wage for hourly workers and the per-period base salary for salaried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

Asymmetry: Changes are more positively skewed for salaried workers. Hourly workers have earnings volatility which is mostly symmetric. Their changes are only slightly

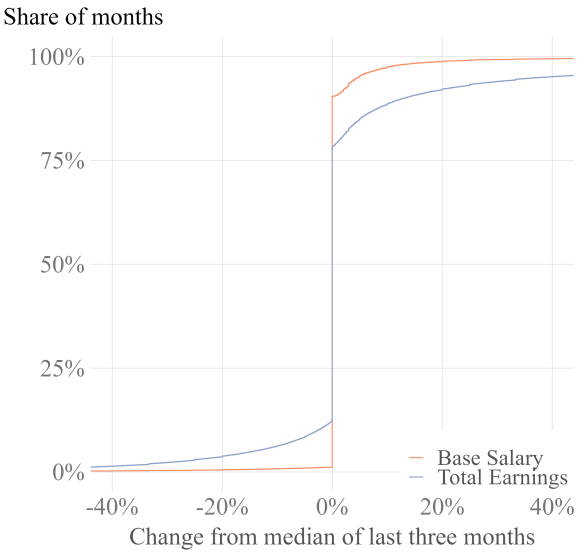
base pay or they can enter a separate pay item labeled "bonus". Because the tax treatment of bonuses is identical to that of base pay, the employer will only mark bonuses as such if it is useful for their internal record-keeping purposes. One piece of evidence we have found that payments which are effectively bonuses are often entered as base pay is that both base pay and bonuses for salaried workers surge by roughly equal amounts in December. This seasonal pattern suggests that half of bonus payments might not be labeled as such.

Figure 3: Sources of Within-Job Earnings Volatility

(a) Hourly Workers



(b) Salaried Workers



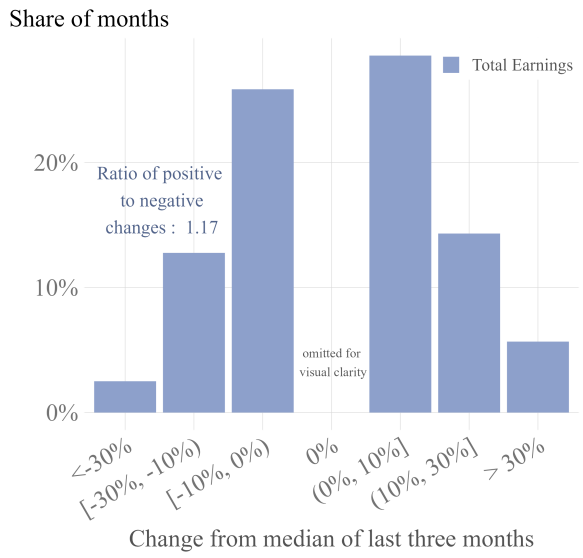
Notes: This figure shows the within-job distribution of $\frac{y_t - Median(\{y_s\}_{s\in t-1,t-2,t-3})}{Median(\{y_s\}_{s\in t-1,t-2,t-3})}$ for total earnings, base wages, and hours where t indexes months. If three months of history are not available the average of the past two months is taken. "Total Earnings" is average pay per paycheck to abstract from pay schedule-driven fluctuations. Sample in left panel is full-time workers, which is defined as those who work an average of at least 30 hours per week.

skewed to the upside. One useful summary statistic to assess asymmetry is the ratio of the share of positive changes to the share of negative changes. Figure 4a shows that positive *earnings* changes are slightly more common than negative changes (ratio 1.17). Meanwhile, Table 1 shows that positive and negative *hours* changes occur with roughly equal frequency (ratio: 0.98). The modest earnings asymmetry arises because base wage changes are almost always positive. The same pattern is also visible in Figure 3a where the distribution of negative hours and earnings changes is quite similar but the distribution of positive earnings changes is shifted out relative to the distribution of positive hours changes.

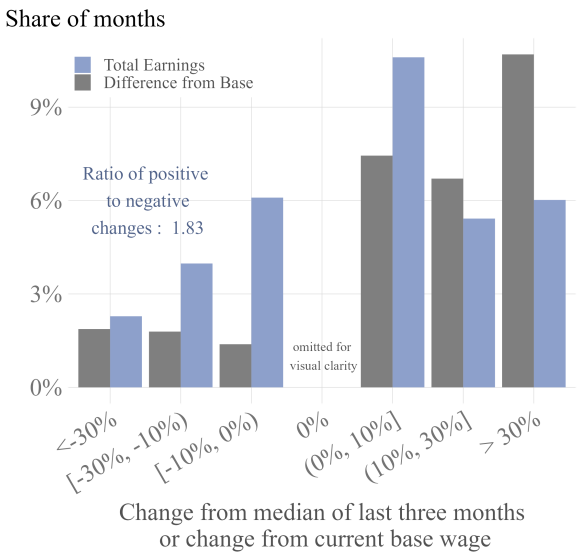
Salaried workers' earnings changes are much less common, but when they do occur they are much more strongly skewed to the upside. Figure 4b shows that the ratio of positive to negative changes is 1.8, so positive deviations are about 80% more frequent than negative deviations. The figure also shows the difference in total earnings as compared to the worker's base salary. A salaried worker might receive less than their base if they took unpaid leave or if the firm decreased the worker's true base salary but did not update the base salary variable (which is the modifiable default amount) in the payroll system. However, such events are rare empirically. Relative to the base wage, the ratio of positive to negative changes is 5.0. This indicates that most earnings volatility for salaried workers is additional income above-and-beyond their the base wage. The ratio statistics of 1.8 and 5.0 provide plausible bounds

Figure 4: Symmetry of Within-Job Earnings Volatility

(a) Hourly Workers



(b) Salaried Workers



Notes: This figure shows the within-job distribution of earnings changes. Both panels show $\frac{TotalEarnings_t - Median(\{TotalEarnings_s\}_{s\in t-1,t-2,t-3})}{Median(\{TotalEarnings_s\}_{s\in t-1,t-2,t-3})}$ where t indexes months and "Total Earnings" is average pay per paycheck to abstract from pay schedule-driven fluctuations. If three months of history are not available the average of the past two months is taken. Sample in left panel is full-time workers, which is defined as those who work an average of at least 30 hours per week. The right panel also shows $\frac{TotalEarnings_t - BaseSalary_t}{BaseSalary_t}$.

on the extent to which salaried workers perceive earnings changes as being skewed to the upside.[8] Unfortunately we do not have any data on base hours which would enable us to compute a comparable measure for hourly workers.

3.3 Earnings Volatility for Chase Bank Customers

We examine the distribution of the month-to-month change in PayrollCompany and in bank account data from JPMCI. This analysis of bank data builds on JPMCI reports examining links between income and expense volatility (Farrell and Greig 2015; Farrell, Greig and Yu 2019). The bank account and PayrollCompany data sources are not linked at the worker level, but we can compare the distribution of changes in the two datasets.

Using two different datasets to answer the same question is useful because they have distinct sources of measurement error. Volatility might be artificially inflated in PayrollCompany data if they include paychecks that are issued in error. Volatility might be artificially inflated in JPMCI data if workers split their paychecks between multiple accounts or shift payment modes.

In practice, Figure 5 shows that the distribution of changes is very similar in the two

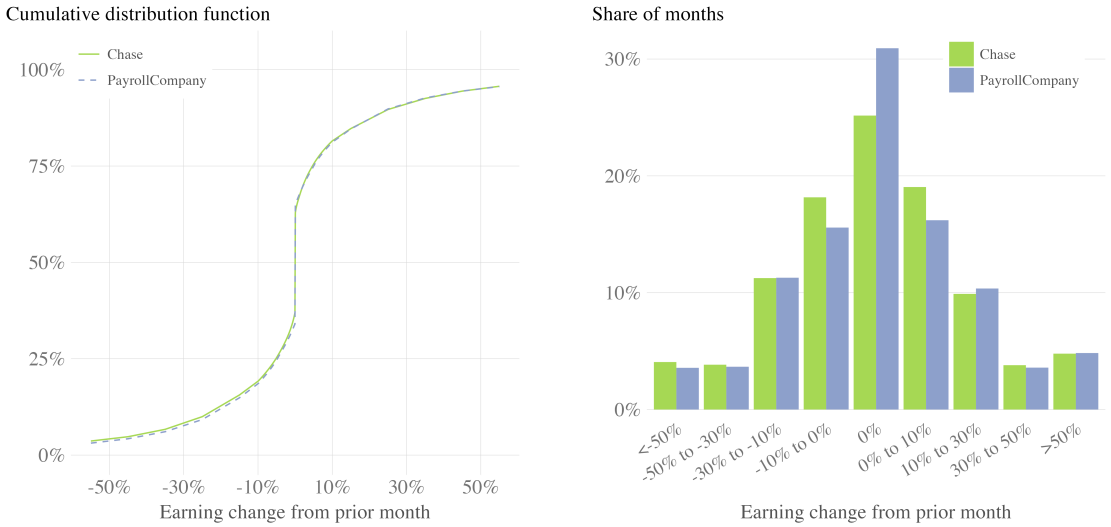
[8] We do not have any measures of salaried workers' expectations. If they form expectations using the median of the last three months then positive changes are 1.8 times as common as negative changes. If they form expectations using their base wage then positive changes are 5.0 times as common as negative changes.

datasets. This suggests either that these distinct sources of measurement error are not important or that they coincidentally happen to be the same size in both datasets. Table 1 reports statistics on the distribution of earnings changes in PayrollCompany using pre-tax earnings, in PayrollCompany using net pay, and in Chase using net pay. All three distributions are very similar.

The Chase data are also useful for addressing two limitations of the PayrollCompany data. First, PayrollCompany primarily serves small firms, whereas Chase serves employees at both small and large firms. Table A-4 shows that the standard deviation of earnings changes is similar across employees of small and large firms who have Chase bank accounts. This suggests that the lessons about pay volatility which are present for small firms using PayrollCompany will also apply to large firms.

Second, we find that job-level earnings volatility translates to household-level earnings volatility. Table A-4 shows that the standard deviation of earnings is the same at job level and for the sum of earnings across all jobs at the household level. This suggests that lessons about pay volatility at the job level will also apply to lessons about pay volatility at the household level.

Figure 5: Earnings Volatility in Payroll Company compared to Chase



Notes: This figure shows the distribution of earnings changes in two datasets. The PayrollCompany series shows changes in pre-tax earnings and is the same as in Figure 1. The "Chase" series is from JPMCI and shows changes in net pay. In both datasets, earnings is defined is average pay per paycheck to abstract from pay schedule-driven fluctuations. PayrollCompany and JPMCI data are analyzed separately and were not merged as part of this analysis.

4 To What Extent Does Earnings Volatility Reflect Risk?

In this section, we try to understand to what extent monthly earnings volatility reflects income risk from the workers' perspective. There are already many studies which measure annual

income risk.[9] That literature typically treats innovations to annual income as "shocks" or "risk" (see e.g. Guvenen et al. 2021). However, there are two reasons why *monthly* volatility might include fluctuations that would not be present in *annual* data and would not generally be considered "risk" from the worker's perspective. First, seasonal forces may induce earnings changes from one month to the next that recur every year and that workers can, in principle, easily plan for. Second, some monthly fluctuations may be driven by worker choices rather than exogenous factors outside of their control. For example, workers may occasionally choose to take time off work to enjoy leisure. If workers do not receive vacation pay as a benefit, then these choices to take vacation time will induce volatility in earnings but will not reflect risk. In the rest of this section we quantify the importance of each of these two potential sources of monthly earnings fluctuations, and then describe the distribution of earnings fluctuations purged of these sources of volatility.

4.1 How Much of Volatility Reflects Predictable Seasonal Patterns?

This section assesses the degree to which monthly earnings volatility reflects predictable seasonal patterns. Many seasonal or regularly recurring sources of earnings volatility, such as regular performance bonuses and seasonal trends in labor demand, can be plausibly identified in the PayrollCompany data by applying a straightforward regression framework to historical information about firm and worker pay. To reduce the potential for overfitting, we focus on a sample of larger firms with an average of at least eight employees per month (see Section 2 for details). We also restrict the sample to salaried or full-time hourly workers who are paid biweekly. We estimate regressions of the following form:

$$\log y_{i,j,t} - \log y_{i,j,t-1} = \boldsymbol{\beta} X_{i,j,t} + \epsilon_{i,j,t}. \tag{1}$$

Here, $y_{i,j,t}$ represents worker i's total earnings per paycheck at firm j in month t, $X_{i,j,t}$ represents a vector of covariates, including a constant term, and $\epsilon_{i,j,t}$ is an error term. We interpret the R^2 in these regressions as a measure of the power of our covariates to predict income changes. We report results for four samples—all hourly, all salaried, salaried without bonus receipt, and salaried with bonus receipt.

We study 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 $X_{i,j,t}$ seeks to capture firm-specific seasonality in pay by estimating firm by month fixed effects $\alpha_{j,m(t)}$ where $m(t)$ in an integer from 1 to 12 (e.g. January 2011 and January 2012 both have $m(t) = 1$). The second definition of $X_{i,j,t}$ seeks to capture 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

[9] See Gottschalk et al. (1994), Moffitt and Gottschalk (2002), Meghir and Pistaferri (2004), Blundell, Pistaferri and Preston (2008), and Altonji, Smith Jr. and Vidangos (2013).

from a year ago. However, because it requires information on the change in pay from thirteen months ago to twelve months ago, it can only be constructed for workers who have been in the sample for at least thirteen months.

Table 2: Predictability of Earnings Changes

	R^2 from regression		
	Hourly	Salaried	
Covariates	All	No bonus	Bonus
Firm x month FEs ($\alpha_{j,m(t)}$), including workers w/$\leq$ 12 month tenure	0.11	0.09	0.41
Firm x month FEs ($\alpha_{j,m(t)}$),	0.16	0.11	0.46
Month FEs ($\alpha_{m(t)}$) + 12-month lags ($\beta_{m(t)}(\log y_{i,j,t-12} - \log y_{i,j,t-13})$)	0.04	0.02	0.36
Number of firms	270	98	61
Number of workers	8,282	1,745	1,146
Number of worker-months	200,009	51,169	35,354

Notes: This table reports R^2s from regressions using equation (1). 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 2 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^2s ranging from 0.36 to 0.46. Salaried workers without bonuses and hourly workers have R^2s ranging from 0.02 to 0.16. 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.[10]

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 thirteen month tenure requirement because the R^2 is higher with firm-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. The firm x month FE specification includes twelve coefficients for each firm, while the month + 12-month lag specification includes only 24 coefficients in all. Thus, the former specification is likely to have a substantially larger R^2 simply because there the model is allowed to fit far more coefficients.

Although income changes are somewhat seasonal for salaried workers with bonuses, the main conclusion from Table 2 is that most income changes are not driven by predictable seasonal patterns. For the hourly workers and salaried workers without bonuses whose employment comprises 83% of months in the PayrollCompany data, the highest R^2 is 0.16.Put

[10]The analysis in this section defines the prediction target as $\log y_{i,j,t} - \log y_{i,j,t-1}$ for consistency with the analysis in Figure 1. However, this prediction target is poorly suited for studying the predictability of bonuses, which are usually zero. When 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.

otherwise, 84% of pay fluctuations are unrelated to predictable seasonal patterns using the regression framework in equation (1).

One important caveat of the analysis in this section is that although purging recurring seasonal variation aligns our measure of "risk" with previous analyses of annual data (where seasonal patterns wash out), neither our analysis nor previous analyses of annual data are intended to fully purge all variation that is potentially predictable to the worker. Instead, we answer the more narrow question of how much of monthly volatility is explained by one particular source of predictable variation (seasonality) which is not present in annual data. We conclude that predictable seasonal variation accounts for only a small share of monthly volatility. Fully purging all variation that is predictable to the worker requires more data on the worker's information set and actual expectations. We think this is an important and interesting question which is beyond the scope of this paper.

4.2 How Much of Volatility is Driven by Unpaid Leisure?

In this section, we attempt to quantify the contribution of unpaid leisure to earnings volatility. We limit our focus to hours volatility because salaried workers are much less likely to have earnings decreases and hours are the key source of earnings volatility for hourly workers. We focus on full-time workers because unpaid vacation is an ill-defined concept for part-time workers. This calculation requires two inputs: an estimate of how much unpaid vacation hourly workers use and a set of predictions for when workers choose to take unpaid vacation.

We estimate that unpaid vacation for full-time hourly workers is 1.58% of all hours worked based on public survey and administrative data. The CPS asks workers if they were physically present at work this week, the reason for their absence, and whether their absence was paid or unpaid. Each week, 0.37% of workers say that they were employed but absent from work for the entire week, the reason was vacation, and the vacation was unpaid. This statistic understates the prevalence of unpaid vacation because it excludes partial-week absences and because vacation is underreported in the CPS. After accounting for these issues, we estimate that hours of unpaid vacation are equal to 1.58% of hours worked for full-time hourly workers in the U.S. Section A.2.1 describes this calculation in detail.

Although our specific estimate of unpaid vacation relies on several assumptions, it is useful to note that some other evidence on paid vacation supports the view that unpaid vacation is rare. Maestas et al. (2023) uses stated-preference experiments to estimate workers' willingness to pay for paid vacation (along with several other job attributes). They find that workers will accept wage reductions of 16-18% for 10 days of paid vacation. They further note that if there are 250 work days a year then 10 days off is a 4% reduction in labor supply. Thus, they conclude that workers

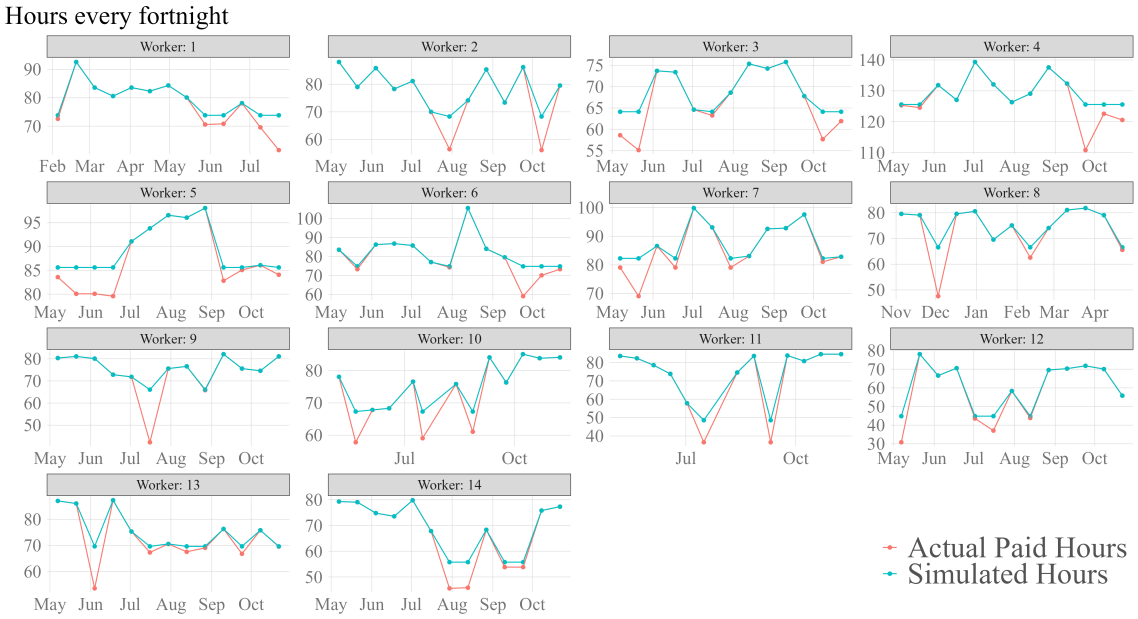
> are willing to sacrifice substantially more than 4 percent of their wages to work at a job with this amount of paid time off... one explanation for the higher valuation is that paid time off represents more than just a reduction in labor effort. Paid

time off also provides job protection, enabling a worker to take time off when desired and without threat of job loss.

Such a high worker valuation of *paid* time off suggests that *unpaid* time off must be rare or entail significant risks for workers.

We propose a simple algorithm to predict when workers take unpaid vacation. We assume that unpaid vacation occurs in the pay period with the lowest amount of hours paid over a six month window. If the unpaid vacation 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 vacation is also taken in the second lowest pay period. This algorithm continues until the unpaid vacation budget is exhausted or hours are equalized across all pay period within the six-month window. Figure 6 provides an illustration of the algorithm's predictions for workers at one firm.

Figure 6: Predictions of Unpaid Vacation Algorithm at One Firm

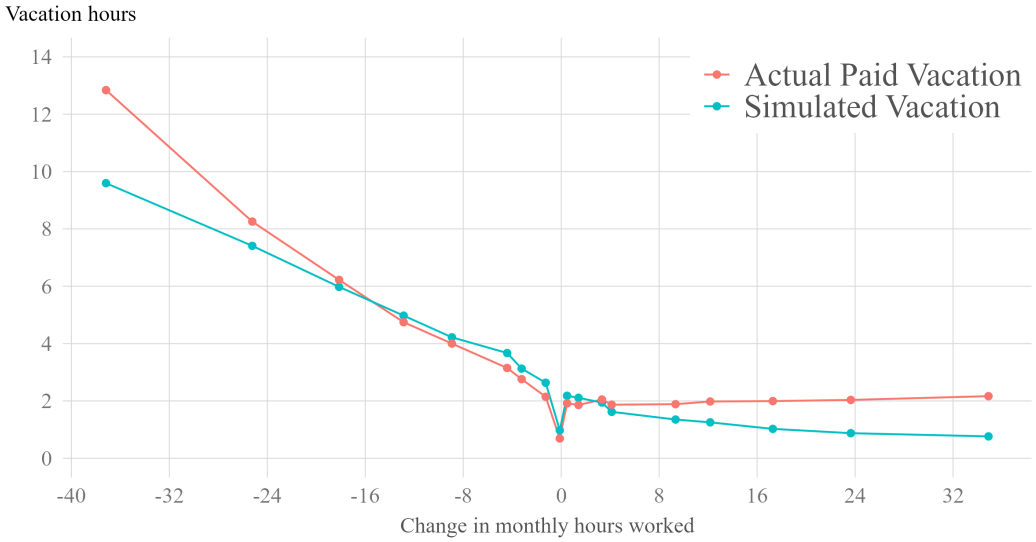


Notes: This figure illustrates the unpaid vacation algorithm for fourteen workers at one firm. We assume that workers have 24 hours of unpaid vacation to allocate across a six month time horizon.

We validate the *unpaid* vacation prediction algorithm by showing that it matches the joint distribution of the change in hours worked and the change in *paid* vacation. In Figure 7 we show deciles of the distribution for the change in hours worked together with average paid vacation hours for every decile. The red line in the figure shows that, on average, some paid vacation is taken in every month and that workers are much more likely to take paid vacation when hours worked are low. On the left side of the plot, in months where hours worked decrease, paid vacation hours increase with a slope of approximately 0.4. The blue line in the figure shows the joint distribution of the algorithm's prediction with the change

in hours worked. It closely follows the red line.

Figure 7: Validation of Unpaid Vacation Algorithm Using Paid Vacation



Notes: This figure reports average amounts of actual paid vacation, simulated vacation, 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% of worker-months, but the bin at zero corresponds to 9.7% of worker months. Paid vacation accounts for 2.26% of compensated hours in the PayrollCompany data. We therefore simulate the timing of unpaid vacation assuming that workers have a budget equal to 2.26% of hours.

Finally, Table 3 shows that estimated pay volatility changes little when we take into account unpaid vacation. This simply involves running the unpaid leave algorithm described in this section with two small tweaks relative to the analysis in Figure 7. First, we assume a budget for unpaid leave equal to 1.58% of hours based on the CPS. Second, we apply the algorithm to hours paid instead of hours worked. Finally, we convert hours changes into pay changes. Table 3 shows that the 75th percentile of the absolute change in pay falls from 13% to 12% and that the standard deviation falls from 18% to 17%.

Table 3: Summary Statistics of Earnings Changes Accounting for Unpaid Leave

Variable	Share $\Delta \neq 0$	Median $\lvert\Delta\rvert$	75th percentile $\lvert\Delta\rvert$	Standard Deviation	Share $\Delta > 0$	Share $\Delta < 0$
Δ: Change from Lagged Median						
Paid Hours	0.89	0.06	0.13	0.18	0.49	0.41
Paid + Simulated Unpaid Hours	0.92	0.05	0.12	0.17	0.5	0.42

Notes: The "Paid Hours" row shows actual earnings volatility for full-time hourly workers. The "Paid + Simulated Unpaid Hours" row shows what earnings volatility would be using the sum of actual paid hours plus simulated unpaid vacation hours. We assume that unpaid vacation is equal to 1.58% of hours paid. The simulation assumes that unpaid vacation is taken in the pay periods with the lowest paid hours (see text for details). Row 1 in this table is extremely similar to row 7 in Table 1 but it is not identical because of small technical differences. Before computing standard deviations, the measures of change are winsorized at the 1st and 99th percentiles of non-zero changes.

We stress that the estimates in this section are uncertain for at least two major reasons. First, actual vacation—paid or unpaid—is lumpy in ways that diminish its ability to reduce volatility relative to the algorithm's predictions of continous amounts of unpaid vacation which specifically target the months with the lowest number of hours. This means the algorithm might overstate the extent to which unpaid vacation is volatility reducing. Second, the estimates are entirely contingent on knowing the total amount of unpaid vacation which workers take. If our estimates of the unpaid leave budget are too low then our estimates of volatility including unpaid vacation will be too high.

4.3 Earnings Volatility Purged of Seasonality and Unpaid Vacation

We construct estimates of earnings risk by purging volatility arising from seasonality and unpaid vacation. We remove seasonality using regressions for three separate groups of workers (hourly, salaried without bonus, salaried with bonus) as in Section 4.1. For removing seasonality, we adjust only months which have a nonzero change in income. We further remove unpaid vacation for full-time hourly workers by applying the procedure from Section 4.2 to the sample from Section 4.1. Finally, we combine the distributions for each contract type to produce a single set of estimates.

Table 4 shows that accounting for seasonality and unpaid vacation leads to modest reductions in some measures of pay volatility. The 75th percentile of the absolute change in pay falls from 11% to 9% and the standard deviation falls from 23% to 17%.

Table 4: Summary Statistics of Earnings Changes Purged of Seasonality and Unpaid Vacation

Sample	Worker	Share $\Delta \neq 0$	Median $\lvert\Delta\rvert$	75th percentile $\lvert\Delta\rvert$	Standard Deviation	Share $\Delta > 0$	Share $\Delta < 0$	Kurtosis
Baseline	Hourly full-time	0.86	0.05	0.12	0.16	0.44	0.41	4.67
Remove seasonality	Hourly full-time	0.86	0.05	0.12	0.15	0.43	0.43	5
Remove unpaid vacation	Hourly full-time	0.86	0.04	0.10	0.14	0.44	0.41	5.02
Remove both	Hourly full-time	0.85	0.04	0.10	0.13	0.42	0.43	5.35
Baseline	Salaried	0.38	0	0.06	0.31	0.21	0.17	15.35
Remove seasonality	Salaried	0.38	0	0.05	0.21	0.20	0.18	21.66
Baseline	All	0.67	0.03	0.11	0.23	0.35	0.32	21.63
Remove both	All	0.66	0.02	0.09	0.17	0.34	0.33	25.14

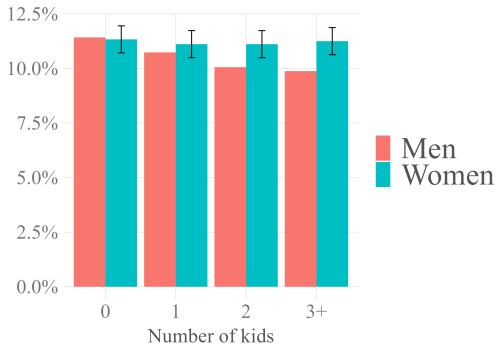
Notes: This table shows the impact of adjusting earnings volatility for seasonality and unpaid vacation. It uses the same definition of pay change (log first difference) and sample as Table 2. See Section 4.1 for details. The values are in log change with each group winsorized at 1% and 99% of non-zero changes. The "All" rows are constructed by taking the distributions for hourly full-time (weighted by 60%), salaried no bonus (weighted by 22%), and salaried bonus (weighted by 18%).

4.4 How Much of Volatility is Driven by Child Care Needs?

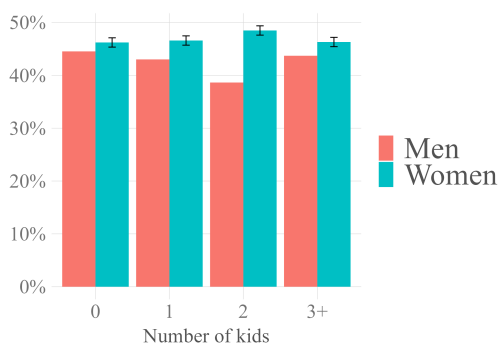
One possible reason that hours volatility could be a job *amenity* is because of child care needs, but in practice, such needs appear to explain only a small share of pay volatility. One way to get a sense of the role that child care needs play in driving pay volatility is to compare volatility for men and women by the number of dependent children they have. Figure 8 shows two measures of volatility by worker gender and number of children. Men and women without children have volatility which is statistically indistinguishable. Men with children have *lower* volatility than men without children, presumably because of lifecycle effects whereby older men tend to have more stable jobs. Women with children consistently have higher volatility than men with children and the differences between men and women are statistically significant. However, the magnitude of the differences are small. This leads us to conclude that child care needs are unlikely to be a major contributor to earnings volatility.

Figure 8: Earnings Volatility by Gender and Dependent Children

(a) Median of changes



(b) Standard deviation of changes



Notes: This figure shows earnings volatility stratified by gender and the number of dependent children reported on the workers' Form W-4. The vertical whiskers show 95% confidence intervals for the difference between men and women within each bin.

4.5 Where Does Remaining Volatility Come From?

So far we have shown that only a small subset of volatility is driven by predictable seasonal variation, leisure-related labor supply choices, or childcare needs. So where does the remaining non-seasonal, non-leisure, non-childcare-driven volatility come from? It could arise for worker-side reasons such as sick leave, medical leave, and family leave. It could also arise because of firm-side reasons such as fluctuations in the firm's labor demand, indivisibility of shifts, lack of managerial capacity to aggregate worker preferences, and downstream effects of worker-side absences. In this subsection we summarize lessons from two imperfect sources of information on this question, which provide suggestive evidence that roughly half of earnings volatility is attributable to firm-driven reasons and half to worker-driven reasons.

The first source leverages information reported by workers in the General Social Survey

(GSS), a nationally representative survey of U.S. residents. This survey includes a number of questions about how much control workers have over their total hours of work per week. Analyzing these questions, Lambert, Henly and Kim (2019) finds that 52% of hourly employees report have little or no input into their schedules, 31% have some input, and 17% decide freely or within limits. This suggests that at least half of hours volatility is due to schedules set by employers.

The second way we have for distinguishing between candidate sources of earnings volatility leverages information reported by firms in the PayrollCompany data. In particular, firms have the option to list workers as temporarily inactive. When they do so, they can choose from a comprehensive list of "reason codes" to explain why the worker is inactive. We analyze the reasons for inactivity spells *within* employment relationships (i.e. the worker is active at some point before and after the inactivity spell).

Table 5 shows that about half of inactivity periods within employment spells are attributed to firm-driven reasons (such as slow work) and half are due to worker-driven reasons (such as family leave). While we view this evidence as suggestive, it is imperfect because it only captures periods where workers are fully inactive. It does not capture all sources of both upside and downside volatility during periods when the worker is active, which is our primary focus in the rest of the paper.

We hope that future research with more detailed data is able to shed further light on the underlying reasons for within-job earnings volatility. Nevertheless, we note that one thing all the remaining explanations have in common is that they all look like income risk from the worker's perspective.

Table 5: Reasons for Unpaid Temporary Leave

Broad category	Reason	Share of spells	Share of spell-weeks
Worker-driven	Family and parental leave	0.181	0.145
Worker-driven	Medical leave	0.155	0.120
Worker-driven	Education leave	0.131	0.155
Worker-driven	Total	0.467	0.420
Firm-driven	Work is slow	0.211	0.225
Firm-driven	Seasonal employment	0.125	0.145
Firm-driven	Contract on hold	0.050	0.065
Firm-driven	Total	0.386	0.435
Other	Miscellaneous	0.147.	0.145

Notes: This table shows the six most common reason codes for temporary unpaid leave within employment relationships in the PayrollCompany data. Within "Miscellaneous", COVID-related inactivity accounts for 8% of spells and another 326 codes account for the remaining inactivity spells.

5 Workers Quit Volatile Jobs and Ascend a Job Ladder - Implying High Willingness to Pay to Eliminate Volatility

In this section, we provide evidence that earnings volatility is, on average, a job disamenity. In Section 5.1 we show that lower-wage workers are more likely to experience earnings volatility. In Section 5.2, we show that workers are more likely to quit jobs when those jobs have high earnings volatility (holding wages constant). In Section 5.3, we follow workers across multiple employment spells and show that after a separation, workers tend to find a new job with lower earnings volatility. Finally in Section 5.4, we use the quit elasticity to price workers' willingness to pay to eliminate pay volatility through the lens of a standard frictional labor market model.

Our estimates imply that workers have a high willingness to pay to eliminate volatility. We caution that the quit elasticities we measure are not causal and so this willingness-to-pay estimate is only suggestive. That said, other work has found quasi-experimental evidence that pay volatility does indeed make workers more likely to quit in the retail (Kesavan and Kuhnen 2017) and home health (Bergman, David and Song 2023) sectors.

5.1 Higher Wage Jobs Have Less Downward Earnings Volatility

We connect our new facts about earnings volatility with other well-known facts about hierarchy in the labor market. We divide workers into six groups of roughly equal size. For salaried workers, we divide sample into the 43% of workers who get bonuses and the 57% of workers who do not. For hourly workers, we divide the sample into quartiles by average hours worked.

The first two panels of Figure 9 show two previously well-known facts about the labor market. The left panel shows that salaried workers earn more per hour than hourly workers do.[11] Within salaried workers, workers who get bonuses get higher average compensation (Lemieux, MacLeod and Parent 2009). Full-time jobs pay more per hour than part-time jobs. The middle panel shows that separation rates move inversely with average wages.

The right panel of Figure 9 shows that downward pay volatility moves nearly in tandem with wages and separation rates. We plot the 10th, 25th, 50th, 75th and 90th percentile of the change in earnings. Earnings are almost completely stable for salaried workers without bonuses, indicating that much of the volatility for salaried workers documented in the prior figures may come from bonuses.[12] Among hourly workers, fewer hours per week is associated with more earnings risk.[13]

[11] We assume that all salaried workers work 40 hours per week. Their implied hourly wages would be even higher if we accounted for the fact that 14% of salaried workers are in fact part-time.

[12] There is some ambiguity about the level of volatility remaining after removing bonuses, because bonus reporting is incomplete. See footnote 7.

[13] This is partly mechanical. If all jobs have the exact same volatility of hours in levels then the jobs with lowest average hours will have the most volatility in % terms. However, part-time jobs also have more volatility in levels. The median absolute change in hours for each group is 3.5 for Q1, 3.53 for Q2, 1.5 for Q3, and 2.45

One interpretation of this common ranking across wages and downward pay volatility is as a disamenity: workers with higher productivity will "spend" some of it on higher wages and some of it in "buying" a contract with no downside risk as an amenity of their job. These contracts are likely more expensive to provide from the employer perspective (Lachowska et al. 2023). This interpretation relies on the well-known idea that in equilibrium workers who are more productive will choose jobs with both higher wages and higher amenities because amenities are a normal good (Lavetti 2023). Another interpretation of this common ranking across wages and downward pay volatility is as an amenity priced through compensating differentials: some workers accept lower wages in return for flexibility in choosing when they work.

Job separations, however, have the potential to differentiate between the disamenity view and the amenity view. The fact that job separations are higher for groups with higher pay volatility suggests that pay volatility is a disamenity. To probe this relationship further, we study the connection between pay volatility and job separations at the firm level in Section 5.2.

Figure 9: Volatility by Job Type



Notes: This plot shows wages, separations, and earnings volatility for six groups of workers which are roughly equal in size: salaried workers with bonuses, salaried workers without bonuses, and four quartiles of average hours worked. The plot conditions on six months of tenure at the firm because we found in exploratory data analysis that for shorter tenures it was hard to reliably classify workers into quartiles of hours worked. Earnings volatility is measured as quantiles of $\frac{TotalEarnings_t - Median(\{TotalEarnings_s\}_{s \in t-1,t-2,t-3})}{Median(\{TotalEarnings_s\}_{s \in t-1,t-2,t-3})}$ where t indexes months. The third panel shows volatility in hours for which the worker receives their base wage (thereby omitting overtime, which is relevant for Q4 hour/week workers). If we include overtime, the distribution of changes for Q4 hour/week workers is -0.15 for p10, -0.05 for p25, 0.07 for p75, 0.19 for p90).

for Q4.

5.2 Quits are Higher at Firms with Volatile Pay

Economists often use quits as a way to understand workers' preferences over job amenities (e.g. Sorkin 2018). Quits are particularly informative in labor markets with search frictions, imperfect competition, and/or heterogeneity in ability such that amenities cannot be priced using differences in wages (Bonhomme and Jolivet 2009).

In the spirit of this approach, we examine the firm-level relationship between quits and pay volatility. To construct a firm-level measure of pay volatility, we take the absolute value of the percent change in pay from the prior month for each worker-month and then aggregate by computing the median across all worker-months at a firm.[14] We use a parallel approach to construct a firm-level quit rate. We take a binary variable for quit in each worker-month and then aggregate to compute the mean monthly quit rate at the firm.[15]

Figure 10 shows that firms with volatile pay see workers separate at higher rates than firms with stable pay. At the 30% of firms with the most pay stability, the median worker-month has the same pay as each worker's previous month and these firms average 5 quits out of every 1000 workers employed each month. By contrast, the 30% of employers with the least pay stability have a quit rate four times larger. These firms average a median percent change in month-to-month pay of more than 15% and average nearly 20 quits for every 1000 workers employed each month. Figure A-7 shows that a similar relationship holds between individual workers' volatility and quit rates.

While the correlation between pay volatility and the quit rate is positive and strong, there is, at most, a weak relationship between pay volatility and separations initiated by the employer. As is visible in Figure 10, there is no relationship between the firing rate and pay volatility and only a weak relationship between the layoff rate and pay volatility. This pattern is consistent with an understanding of pay volatility as, on average, a job disamenity.

Table 6 converts the bivariate correlation between the quit rate and earnings volatility as illustrated in Figure 10 into a regression in column (1). Our preferred specification in column (2) is

$$P(\text{Quit})_j = \beta_0 + \beta_1 \text{Median}|\Delta|_j + \beta_2 \text{Wage}_j + e_j \tag{2}$$

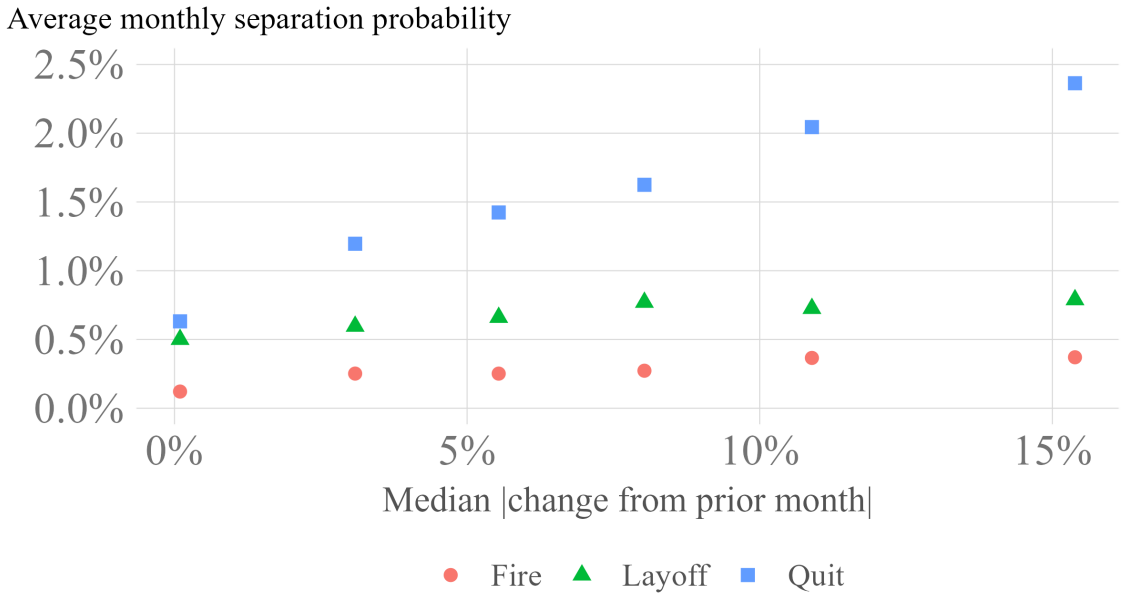
where j indexes firms. We control for wages because, as documented in Section 5.1, pay volatility is correlated with wages. We continue to find a strong relationship between volatility and quits after adding this control.

One difficulty in interpreting these quitting patterns as only reflecting the disamenity of volatile pay is that jobs with large pay volatility are potentially different from jobs with low

[14] We also explored specifications which instead use the standard deviation instead of the median. We find a positive, but weaker relationship between quits and volatility. We believe that the relationship is weaker because salaried workers tend to have a low quit rate and a lot of volatility as measured by the standard deviation.

[15] Although our baseline analysis weights each worker-month equally, we also conduct a robustness check where we instead weight each worker equally.

Figure 10: Firm-level Volatility and Separation Rate



Notes: This figure shows the relationship between firm-level pay volatility (measured using typical changes in worker earnings from prior month) and separation rates. See text for details on how pay volatility and separation rates are measured. Firms are stratified into deciles of volatility. The top decile is dropped for readability, but a version that includes the top decile is shown in Figure A-8. Because about 30% of observations have median earnings change of zero, there are are only 6 data points plotted here. Data is from PayrollCompany. We can only observe the separation type from May 2020 onward; during this period, 65% of separations are quits, 23% are layoffs, and 13% are fires.

pay volatility along many dimensions. For example, most salaried workers have a median absolute change in pay of zero while most hourly workers are categorized as having nonzero volatility in pay. Salaried and hourly jobs usually differ on dimensions beyond pay volatility and so we cannot attribute differing quit rates simply to the difference in volatility. To address some omitted variable concerns (such as the differences between hourly and salaried jobs), we study alternative specifications to the correlation illustrated in Figure 10.

Table 6 shows that the relationship between quits and volatility is robust to several different specification choices. To address concerns that pay volatility is correlated with unobserved characteristics of hourly jobs, we re-run the regression on the sub-sample of hourly workers and find similar results in Table 6, column (3). The fourth column runs the same regression as column (2), but weights firms by the product of the monthly employment and the number of months the firm is observed at PayrollCompany. The fifth column takes the specification from column (2) and adds industry fixed effects. The sixth column adds a linear control for average hours worked. This control is motivated by prior research which demonstrates that workers prefer jobs that offer more hours (Lachowska et al. 2023; Rose and Shem-Tov 2023) combined with the empirical fact that jobs with more hours per week tend to also provide more stable hours. The seventh column regresses the quit rate on

Table 6: Firm-level Volatility and Quits: Cross-Section

	Bivariate	Baseline	Paid Hourly	Equal Person Weights	Control Industry	Control Hours	X: Person
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Median $\|\Delta\|$	0.121***	0.109***	0.086***	0.107***	0.100***	0.112***	0.063***
	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.001)
Log wage		−0.005***	−0.007***	−0.007***	−0.004***	−0.006***	−0.010***
		(0.0004)	(0.001)	(0.0004)	(0.0004)	(0.0004)	(0.0002)
Average hours						0.00003***	
						(0.00001)	
Observations	17,202	17,202	16,154	17,202	17,202	17,202	472,309
Adjusted R^2	0.079	0.090	0.042	0.121	0.107	0.091	0.020

*p<0.1; **p<0.05; ***p<0.01

Notes: This table shows the relationship between pay volatility and quit rates. In columns (1)-(6) the unit of observation is a firm. See text for details on how we aggregate from worker volatility and quit rates to firm-level variables. In column (7) the unit of observation is a worker. Log wage refers to the hourly wage, assuming that salaried workers work 40 hours per week.

earnings volatility where the unit of observation is a worker (rather than a firm). Finally, in Table A-5, we show that the relationship between pay volatility and the quit rate remains stronger than the relationship between pay volatility and firm-initiated separations, even after controlling for wages. Across all specifications, the relationship between volatility and quits is economically and statistically significant.

Two concerns with interpreting the results in Table 6 as capturing worker distaste for pay volatility are reverse causality and selection. We first address reverse causality before discussing selection in Section 5.3.

One possibility is that when workers quit this causes an increase in pay volatility for coworkers who remain at the firm. If this were the case, it would be that high quit rates are the cause of earnings fluctuations for workers instead of pay volatility inducing workers to quit. In Table A-6 we address this concern in two ways. First, we use lagged volatility as a regressor in column (2) and find a similar coefficient to contemporaneous volatility in column (1). Second, in column (3) we examine whether volatility in a quarter where a firm has no workers quit predicts the likelihood that a worker quits in the following quarter. Because column (3) uses contemporaneous volatility as the key regressor, it remains susceptible to reverse causality. Column (4) uses the same restricted sample as column (3), but uses lagged volatility as the key regressor and is therefore not subject to concerns about reverse causality. We interpret the similarity of the coefficients on volatility in column (3) and column (4) as suggestive evidence that volatility causes quits rather than the reverse.

5.3 Workers Climb a Job Ladder to Reduce Volatility

The regressions above may also be biased by selection. For example, some firms may hire workers who value flexibility both on the intensive margin (more pay volatility) and the extensive margin (making them more likely to quit). Another possibility is that some firms

attract workers who are "flaky", both showing up at work inconsistently (generating pay volatility) and being more likely to quit. In both of these examples, an omitted time-invariant worker characteristic would explain a positive correlation between volatility and quits but the correlation would be uninformative about whether pay volatility is, on average, a disamenity.

One way to address these concerns about whether workers actually perceive volatility to be a disamenity is to analyze the types of jobs workers choose to take when they separate from one job in favor of another one. This approach holds worker characteristics fixed, and uses revealed preference to infer how workers value the bundle of characteristics at one job relative to another.

Presumably if workers quit one job and accept an alternative one this means that they prefer the bundle of characteristics offered by the new job. We follow workers who separate from a firm that uses PayrollCompany and are hired at another firm also using PayrollCompany. We can therefore check to see if indeed workers leave high volatility jobs for low volatility jobs.

We find that, on average, workers switch to jobs that are higher paying and lower volatility. Table 7 shows that the median wage for workers who *quit* jobs is \$22.34 per hour. The mean change in wage from origin to destination is \$3.97. We measure each worker i's volatility as $Vol_i = \text{Median}|\Delta_{it}|$. Across all quitters, the median volatility at the origin is 9.62%. The mean reduction in Vol_i is 0.97 percentage points which is a 10% reduction in volatility.

The fact that workers depart high-volatility jobs in favor of low-volatility jobs is consistent with an interpretation of pay volatility as a disamenity for the average worker. This interpretation relies on the idea discussed in Section 5.1 that workers will quit 'worse jobs' (either in terms of wage or volatility) for jobs that are higher wage, lower volatility, better on some unobserved dimension, or some combination of all of the above.

Table 7: Worker-level Job Ladder: Averages

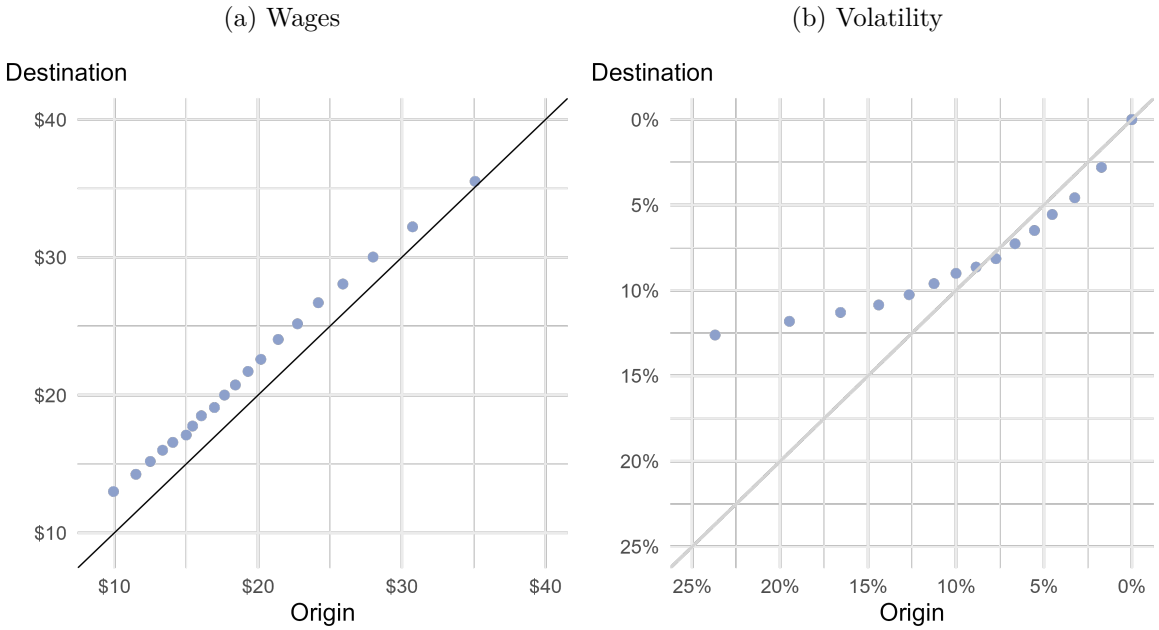
	Hourly wage	Volatility (Median $\|\Delta\|$)
Origin (median)	\$22.34	9.62%
Change from origin to destination	\$3.97	-0.97%
Row (2) / Row (1)	18%	-10%

Notes: This table documents the differences in wages and pay volatility between the origin and destination firms for 45,241 workers who quit the origin firm and join a destination firm both using PayrollCompany. We restrict our sample to workers aged 25-54 who are out of sample for 12 months or less and transition exactly once between May 2020 and January 2023. See text for details.

Next we move from studying the average change in job attributes from a transition to explore heterogeneity in job switches. We find the strongest evidence of a job ladder (i.e. workers leaving low-utility jobs for higher-utility jobs both in terms of higher wage and lower volatility) at the bottom end of the wage distribution and the high end of the volatility distribution. Figure 11 stratifies the sample of switchers by wages and volatility in the origin job. Panel (a) shows that there is wage growth throughout the entire origin wage

distribution. The largest wage growth upon job switch occurs for workers at the bottom of the wage distribution, which is consistent with the aggregate time trends for this period in the U.S. economy.

Figure 11: Worker-level Job Ladder: Heterogeneity



Notes: Panel (a) bins workers into vigintiles according to wage at origin. The x-axis is the wage at origin (the median within the bin). The y-axis is the wage at destination firm (the median within the bin). Workers across the wage distribution see wage gains at the new job. Panel (b) bins workers into vigintiles according to pay volatility at origin. On the x-axis is pay volatility at origin (the median within bin). On the y-axis is pay volatility at destination (the median within bin). Both axes are flipped so that higher values indicate a more positive amenity. Approximately 30% of workers have a zero volatility job at origin. Sample is workers who separate from an origin firm that uses PayrollCompany and is hired at a destination firm that also uses Payroll Company. See text for more details.

Panel (b) similarly shows that the largest decrease in volatility upon job switch occurs for workers at the high end of the volatility distribution. However, the patterns for volatility changes differ from wage changes in two ways. First, about 30% of workers have zero volatility at their origin job. The median destination volatility for these workers is also zero; this reflects workers changing from one salaried position to another. Second, workers with positive but low volatility see small *increases* in volatility. We believe that the most natural interpretation of this pattern is mean reversion. Recall from Figure 1 that earnings change frequently and wages change infrequently. If worker-level volatility is measured with any error then we will see workers switching from apparently low volatility jobs to apparently high volatility jobs (and vice versa).[16] Because wages are easier to measure, the same mean reversion pattern is not apparent there.

In unreported analysis, we also explore the extent to which workers trade off increased

[16]By measurement error, we do not mean that there are mistakes in the data. Rather, we mean that it is challenging to measure the true prevalence of volatility in pay from only a few months of data at each job.

wages or decreased volatility as part of job transitions. When workers leave low-wage jobs, they tend to move to jobs which are *both* higher wage and lower volatility. The fact that workers are moving from lower-wage, lower-stability work to higher-wage, higher-stability work suggests that jobs are vertically differentiated and jobs which are higher wage tend to be better on other dimensions as well. This is consistent with the vertical differentiation of jobs by wages, separation rates, and downward pay volatility documented in Figure 9.

The evidence in this subsection shows that the estimates in Table 6 are not driven entirely by selection and instead suggests a positive causal relationship between volatility and quits. However, the ideal strategy for eliminating potential bias from selection in the magnitude of the coefficients would be to find an instrument for volatility. We hope to add such an instrument in a future draft.

5.4 Willingness to Pay to Eliminate Volatility

We use the elasticity of quits to volatility to price workers' willingness to pay to eliminate frequent within-year pay volatility. In the Burdett-Mortensen model, the elasticity of quits with respect to randomly-assigned job amenity values, combined with the elasticity of quits with respect to randomly-assigned wages, can be used to measure how much a worker would give up in wages for a different value of the amenity. In Appendix A.2.2, we prove that in the Burdett-Mortensen model, the willingness to pay for a marginal unit of a job amenity is equal to the ratio of the quit elasticity to the amenity divided by the quit elasticity to wages.

We estimate an elasticity of quits with respect to volatility using essentially the same procedure from Section 5.2. As discussed above, our preferred measure of volatility—the median absolute change in pay—captures the magnitude of frequently-occurring pay fluctuations and usually designates salaried workers as having no pay fluctuations. We note that while salaried workers do in fact have meaningful pay fluctuations from month-to-month, these fluctuations are sufficiently infrequent that almost all salaried workers have a median pay change of zero. On a similar note, both groups of workers would still face annual income risk even if their pay changed only once per year. Thus, our exercise can be thought of as calculating the willingness to pay to move from the volatility at a typical hourly job to the amount of volatility which is typical at a salaried job.

Table 8 estimates the elasticity of quits with respect to volatility. Because we ideally want to estimate a log change in quits and we need to allow for frequent zeros as part of the dependent variable, we use a Poisson regression (Chen and Roth 2024). We find a semi-elasticity of quits with respect to average firm volatility of 5.8.[17] The interpretation of this coefficient is that moving from median pay fluctuations of 10% at the firm level to median pay fluctuations of 0 reduces the quit rate by 58%.

[17]This is a semi-elasticity because it captures the percent change in Y from a marginal (not percent) change in X. Because of the rest of the literature usually discusses quit elasticities, we refer to it informally elsewhere as a quit elasticity.

Table 8: Quit Elasticity

	Number of Quits	
	Firm	Person
	(1)	(2)
Median $\|\Delta\|$	5.759***	2.695***
	(0.069)	(0.025)
Implied WTP for zero volatility (using $d\log Q/dw$ from BDN)	−25%	−12%
Observations	17,202	472,475

Notes: This table reports the coefficients from a Poisson regression of quits on pay volatility. The dependent variable is the number of quits but because the regressions control for the number of workers at the firm, the coefficients can be interpreted as the proportional effect of each variable on the quit rate. A coefficient β in a Poisson regression should be interpreted as the change in relative risk so $E(y|x) = exp(\beta x)$. Column (1) examines this relationship at the firm level and column (2) examines this relationship at the worker level.

We estimate that a typical hourly worker would accept a pay cut of 12-25% to bring median pay volatility to zero. Concretely, using the willingness-to-pay formula of the ratio of the quit elasticity to a marginal unit of the amenity (in this case volatility, using the estimates from Table 8) divided by the quit elasticity to wages (using the estimate of 1.8 from Bassier, Dube and Naidu 2022, henceforth BDN), and multiplying by the median hourly worker's pay volatility of 8% gives

$$\textbf{Willingness to Pay} = \frac{\frac{d\log Q}{d\text{Median}|\Delta|}}{\frac{d\log Q}{d\log w}} \times E(\text{Median}|\Delta|) = \frac{5.8}{1.8} \times 0.08 = 25\%.$$

We prefer to use BDN's causal estimate of the quit elasticity to wages which relies on variation in firm wage policies in Oregon. We believe that such an estimate in the PayrollCompany data would likely be unreliable.[18] However, it is still useful to note that the elasticity of quits with respect to firm average wages using simple OLS is -0.717 (Table 8, column 2) in the PayrollCompany data and -0.753 in BDN's Table 1 column (3), suggesting a high degree of similarity in the quit elasticities with respect to wages between the two datasets. We also slightly prefer using our estimates of the quit elasticity with respect to firm-level volatility in column (1) of Table 8 rather than the elasticity with respect to individual-level volatility in column (2) because we think the coefficient estimated using individual-level volatility may suffer from attenuation bias due to measurement error.

[18]BDN use data on the universe of workers in Oregon and the procedure from AKM to measure firm wage policies. Identification and inference in AKM models is most reliable when a large number of workers are switching among a connected set of firms. Because PayrollCompany serves small firms and because most job switches are to firms outside the PayrollCompany universe, we are pessimistic about the reliability of AKM-based estimates in the PayrollCompany data.

The quit elasticity to pay volatility therefore implies a substantial willingness to pay to eliminate frequent within-year pay volatility, but has two major limitations. First, the biggest limitation is that we do not currently have a natural experiment which varies pay while holding all other job attributes constant. Second, even if we did have a natural experiment, it would not identify the fundamental reasons why pay volatility is a disamenity. In particular, it is not clear to what extent the quit elasticity reflects the disamenity of fluctuating hours (holding income constant) versus the disamenity of fluctuating income (holding hours constant). Survey evidence from workers in Schneider and Harknett (2019) suggests that fluctuating hours might be more important than fluctuating income.[19]

6 Implications for Estimates of Income Risk

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.[20] 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.[21]

Figure 12 compares the distribution of monthly earnings changes in the data to that inferred from several models using this approach. In particular, we simulate the continuous time model from KMV and compute its implications for monthly earnings changes. 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 Laibson, Maxted and Moll (2021, hereafter LMM) and Crawley, Holm and Tretvoll (2022, hereafter CHT). [22] Table 9 provides

[19] Schneider and Harknett (2019) finds that several measures of worker well-being are correlated with schedule instability. They then conduct a mediation analysis to understand the relative roles of economic insecurity and what they call "work-life conflict". The latter is measured by responses to four questions about whether a worker's role at work is in conflict with their roles in their family, which is one consequence of fluctuating hours. They find that work-life conflict is more important than economic insecurity.

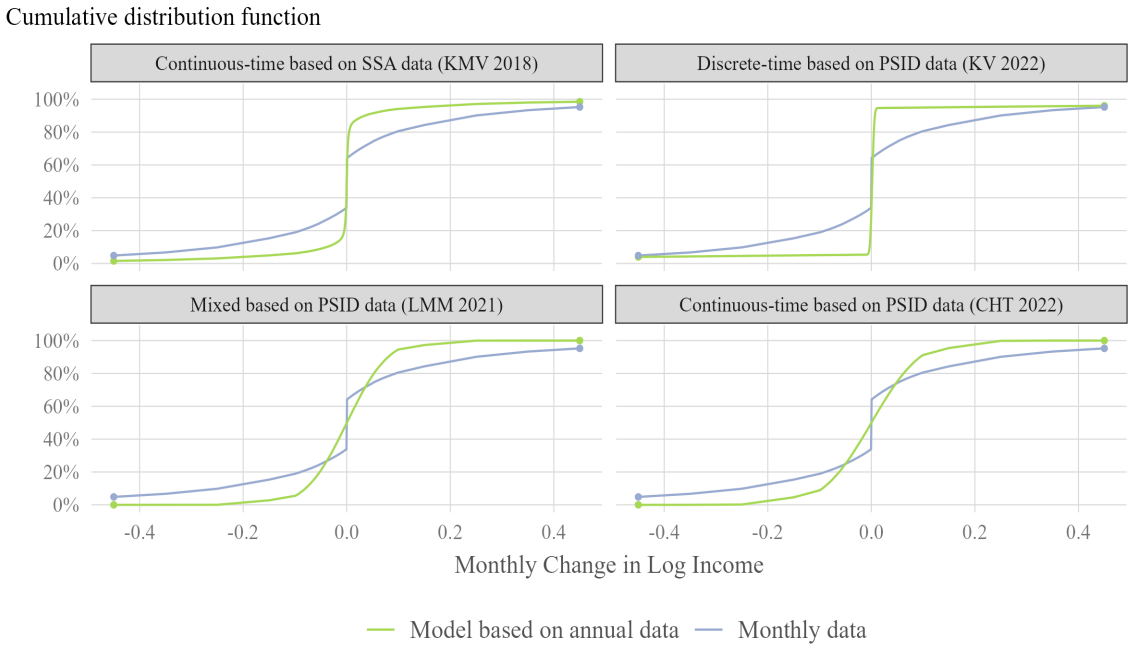
[20] This issue is particularly salient for these types of applications that focus on higher frequency phenomenon, but the level of risk at all time horizons could also be relevant even for lower frequency choices like retirement savings.

[21] KMV 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 generate by an earnings process that is dominated by large infrequent shocks."

[22] The KMV and KV models include a two-shock process that arrives with Poisson probability. Since KMV is a continuous time model, calculating monthly moments is straightforward. 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

summary statistics from these distributions.

Figure 12: 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), LMM is Laibson, Maxted and Moll (2021), and CHT is Crawley, Holm, and Tretvoll (2022). See footnote 22 for additional detail.

This exercise shows that sizeable 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% and 9% in the models (as compared to 17% in the data), and the 90th percentile is between 10% and 13% in the models (as compared to 39% in the data). Another interesting pattern shown in Table 9 is that deviations from the prior month are much more persistent in most 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). In the data, the shocks revert much more quickly.

In ongoing work, we are exploring alternative income processes that can jointly match earnings dynamics at both monthly and annual frequencies. Preliminary evidence suggests that this requires a model with three shocks. Persistent "career" shocks and more transitory "unemployment" shocks like those used in the past literature can help match annual earnings changes but miss most sub-annual earnings changes. The addition of a third "hours" shock

transitory shocks, and we re-estimate the size of shocks to match the same annual moments. The LMM 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 9: Earnings Risk in Monthly Data versus Models Calibrated to Annual Data

	Monthly data	Model based on annual data			
		KMV (2018)	KV (2022)	LMM (2021)	CHT (2022)
P90 - P10 Δ	0.43	0.07	0.01	0.16	0.19
Share $\|\Delta\| > 1\%$	0.64	0.30	0.11	0.87	0.89
Share $\|\Delta\| > 20\%$	0.22	0.07	0.09	0.00	0.01
50th percentile $\|\Delta\|$	0.04	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	0.25	0.17	0.30	0.06	0.07
Kurtosis	6	31	14	3	3
Crow-Siddiqui kurtosis	5	104	421	3	3
Positive persistence	0.35	0.95	0.92	0.58	0.37
Negative persistence	0.32	0.95	0.83	0.58	0.37

Notes: This table shows summary statistics of monthly log earnings changes ($\Delta = \log y_t - \log y_{t-1}$) in PayrollCompany data and in several benchmark models of earnings processes which are calibrated to annual data. KMV is Kaplan, Moll and Violante (2018), KV is Kaplan and Violante (2022), LMM is Laibson, Maxted and Moll (2021), 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 $\frac{\text{P97.5}-\text{P2.5}}{\text{P75}-\text{P25}}$. 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.

which is very frequent and also very transitory can then help to match these sub-annual earnings changes. After developing this alternative income process that better fits earnings dynamics at all time horizons, we plan to explore the implications of this high-frequency risk for consumption-savings problems.

7 Conclusion

In this paper, we document that within-year earnings fluctuations are substantial, that the fluctuations appear to reflect economic risk, and that on average the fluctuations appear to be a disamenity of economically meaningful magnitude. This suggests several questions which might be answered by future research. First, we know from prior work that firms already protect workers from—or put alternately, do not share the rents from—most of their productivity fluctuations (Guiso, Pistaferri and Schivardi 2005, Card et al. 2018, Balke and Lamadon 2022). What is the labor demand problem which firms are solving for which substantial within-year earnings risk is an equilibrium outcome? Second, what are the consequences of pay volatility for consumption and financial distress? Third, to what extent can markets where workers trade shifts improve welfare by allowing workers who want stable hours to receive stable hours and workers who want flexible hours to achieve that flexibility?

References

Altonji, Joseph G., Anthony A. Smith Jr., and Ivan Vidangos. 2013. "Modeling Earnings Dynamics." Econometrica, 81(4): 1395–1454.

Balke, Neele, and Thibaut Lamadon. 2022. "Productivity Shocks, Long-Term Contracts, and Earnings Dynamics." American Economic Review, 112(7): 2139–2177.

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.

Bloemen, Hans G. 2008. "Job Search, Hours Restrictions, and Desired Hours of Work." Journal of Labor Economics, 26(1): 137–179.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." American Economic Review, 98(5): 1887–1921.

Bonhomme, Stéphane, and Grégory Jolivet. 2009. "The Pervasive Absence of Compensating Differentials." Journal of Applied Econometrics, 24(5): 763–795.

Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." International Economic Review, 39(2): 257–273.

Bureau of Labor Statistics. 2024a. "Employer Costs Per Hour Worked for Employee Benefits." Accessed: Feb 22, 2024.

Bureau of Labor Statistics. 2024b. "Employer Costs Per Hour Worked for Employee Compensation." Accessed: Feb 22, 2024.

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.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with zeros? Some problems and solutions." Working Paper, 1–58.

Crawley, Edmund, Martin Blomhoff Holm, and Håkon Tretvoll. 2022. "A Parsimonious Model of Idiosyncratic Income." Washington: Board of Governors of the Federal Reserve System Working Paper 2022-026.

Druedahl, Jeppe, Michael Graber, and Thomas H. Jørgensen. 2023. "High Frequency Income Dynamics." Working Paper, 1–55.

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." National Bureau of Economic Research Working Paper 30441.

Farber, Henry S. 1993. "The Incidence and Costs of Job Loss: 1982-91." Brookings Papers on Economic Activity, 24(1): 73–132.

Farrell, Diana, and Fiona Greig. 2015. "Weathering Volatility." JPMorgan Chase Institute.

Farrell, Diana, Fiona Greig, and Chenxi Yu. 2019. "Weathering Volatility 2.0." JPMorgan Chase Institute.

Ganong, Peter, Damon Jones, Pascal J. Noel, Fiona E. Greig, Diana Farrell, and Chris Wheat. 2023. "Wealth, Race, and Consumption Smoothing of Typical Income Shocks." National Bureau of Economic Research Working Paper 27552.

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.

Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi. 2005. "Insurance within the Firm." Journal of Political Economy, 113(5): 1054–1087.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" Econometrica, 89(5): 2303–2339.

Hannagan, Anthony, and Jonathan Morduch. 2015. "Income Gains and Month-to-Month Income Volatility: Household Evidence from the US Financial Diaries." NYU Wagner Research Paper 2659883.

Hoffman, Mitchell, and Steven Tadelis. 2021. "People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis." Journal of Political Economy, 129(1): 243–285.

Jarosch, Gregor. 2023. "Searching for Job Security and the Consequences of Job Loss." Econometrica, 91(3): 903–942.

Kaplan, Greg, and Giovanni L. Violante. 2022. "The Marginal Propensity to Consume in Heterogeneous Agent Models." Annual Review of Economics, 14(1): 747–775.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." American Economic Review, 108(3): 697–743.

Kesavan, Saravanan, and Camelia M. Kuhnen. 2017. "Demand fluctuations, precarious incomes, and employee turnover." Working Paper, 1416–1435.

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." National Bureau of Economic Research Working Paper 31205.

Laibson, David, Peter Maxted, and Benjamin Moll. 2021. "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy." National Bureau of Economic Research Working Paper 29094.

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.

Lavetti, Kurt. 2023. "Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications." Journal of Economic Perspectives, 37(3): 189–212.

Lehmann, Tobias. 2023. "Non-Wage Job Values and Implications for Inequality." Working Paper, 1–45.

Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent. 2009. "Performance Pay and Wage Inequality." The Quarterly Journal of Economics, 124(1): 1–49.

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.

Manning, Alan. 2003. Monopsony in Motion: Imperfect Competition in Labor Markets. Princeton University Press.

Mas, Alexandre, and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." American Economic Review, 107(12): 3722–3759.

Meghir, Costas, and Luigi Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity." Econometrica, 72(1): 1–32.

Moffitt, Robert A., and Peter Gottschalk. 2002. "Trends in the Transitory Variance of Earnings in the United States." The Economic Journal, 112(478): C68–C73.

Rose, Evan K., and Yotam Shem-Tov. 2023. "How Replaceable Is a Low-Wage Job?" National Bureau of Economic Research Working Paper 31447.

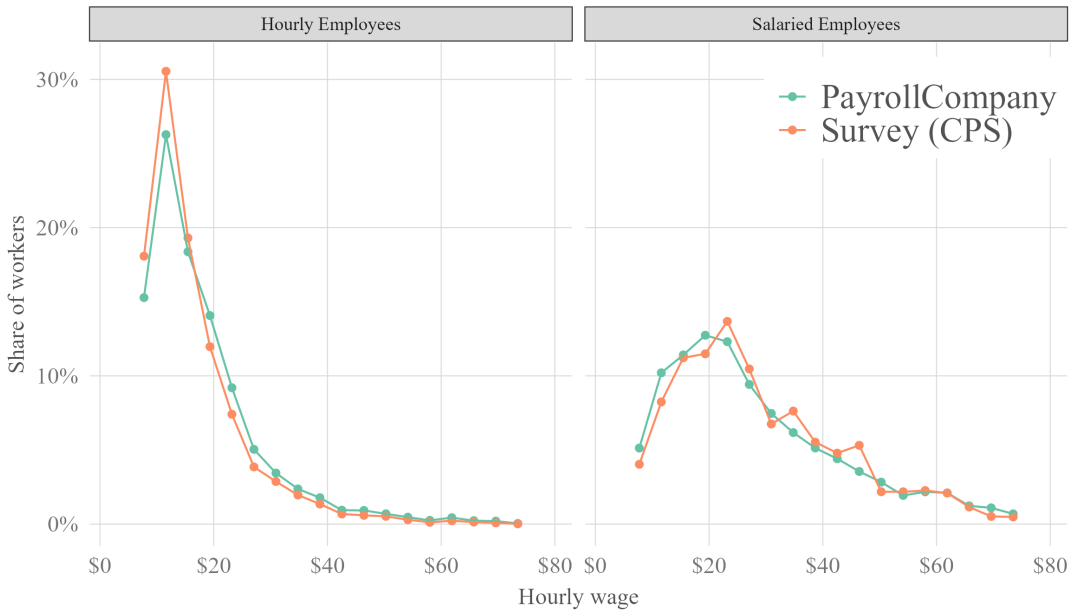
Schneider, Daniel, and Kristen Harknett. 2019. "Consequences of Routine Work-Schedule Instability for Worker Health and Well-Being." American Sociological Review, 84(1): 82–114.

Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." The Quarterly Journal of Economics, 133(3): 1331–1393.

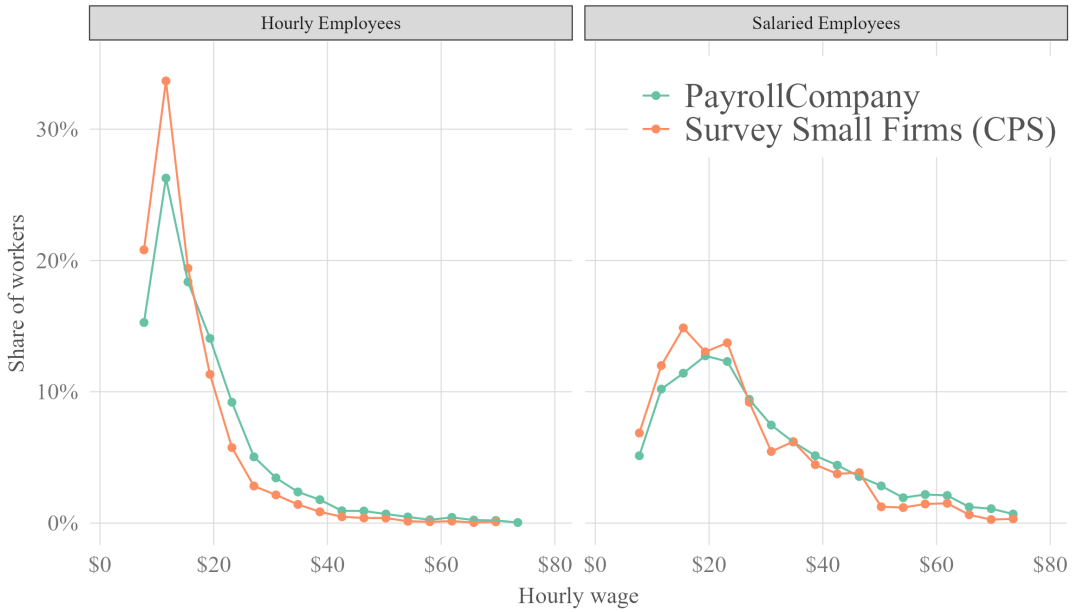
A Appendix Figures and Tables

Figure A-1: Wage Distribution

(a) CPS: All Workers

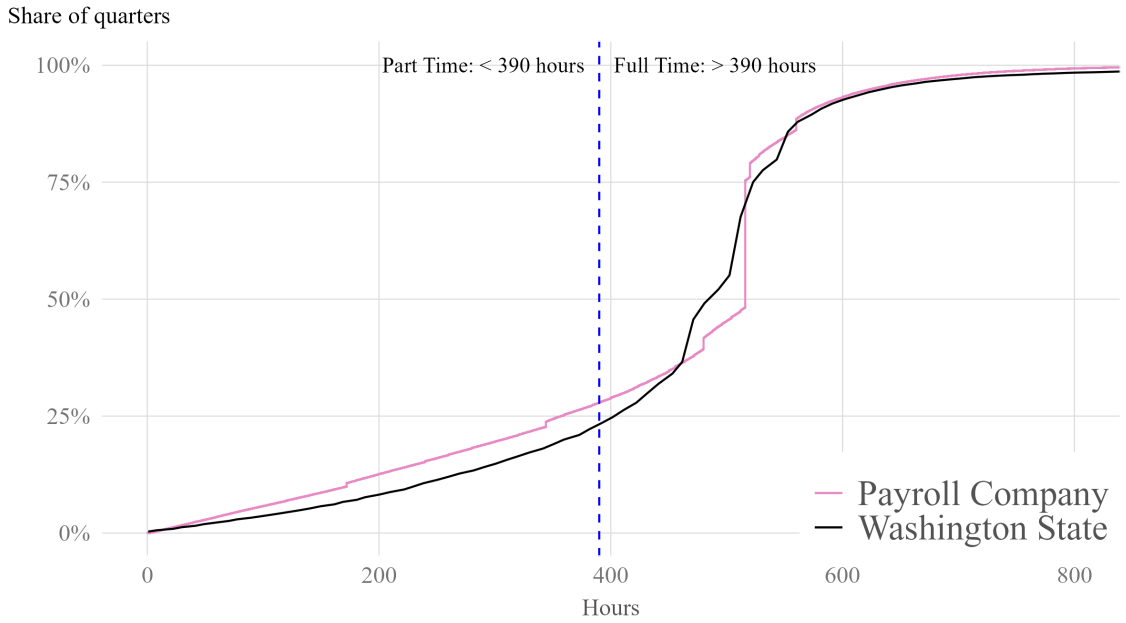


(b) CPS: Workers at Small Firms



Notes: This figure shows the hourly wage distribution in the PayrollCompany data (green) and in the Current Population Survey (red). 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

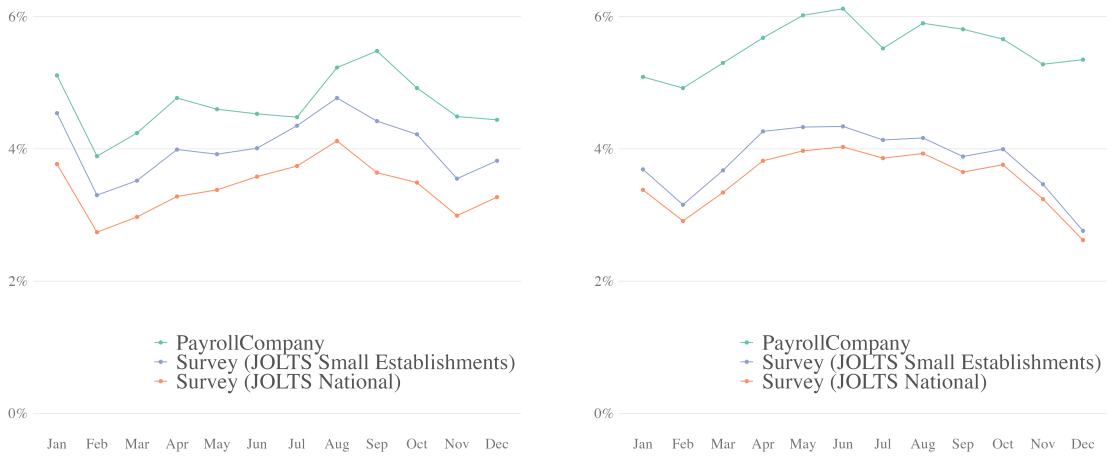


Notes: This figure shows the distribution of quarterly hours worked in the PayrollCompany data (green) and in the Washington state tax data (black). The Washington state tax data series is from Figure 2.B of Lachowska, Mas and Woodbury (2022). We take two steps to make the data series as comparable as 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 green distribution.

Figure A-3: Employee Turnover

(a) Separations

(b) Hires



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

Figure A-4: Pay Frequency

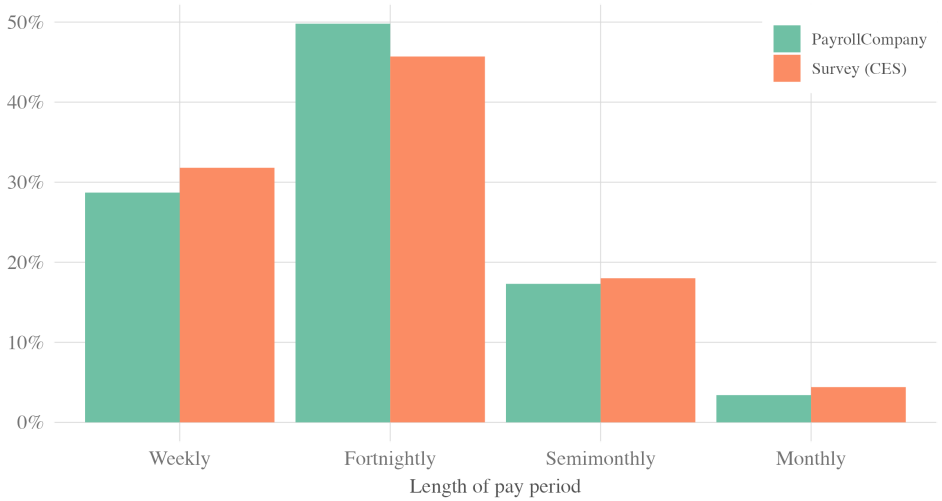
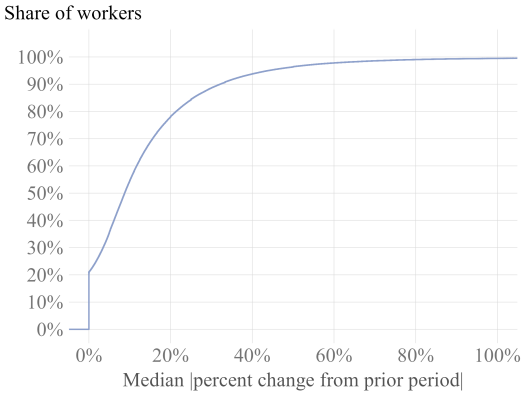
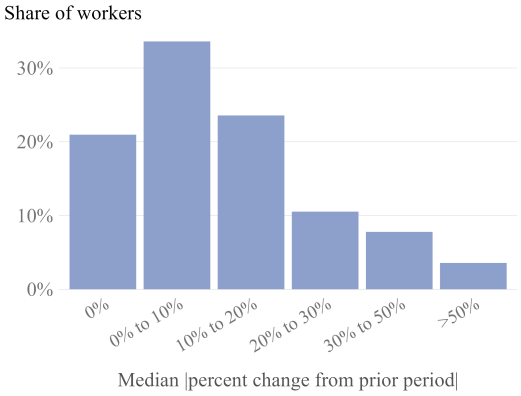


Figure A-5: Worker-level Volatility

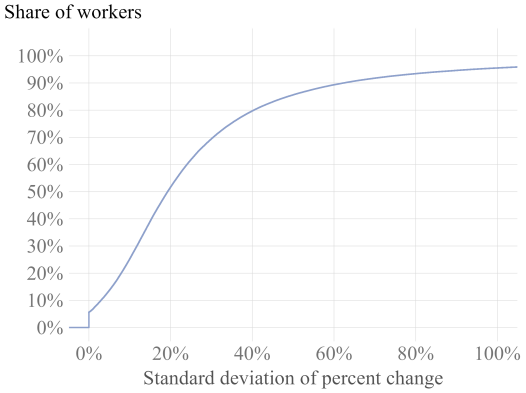
(a) CDF of Worker-level Median Absolute Change



(b) Histogram of Worker-level Median Absolute Change



(c) CDF of Worker-level Standard Deviation



(d) Histogram of Worker-level Standard Deviation

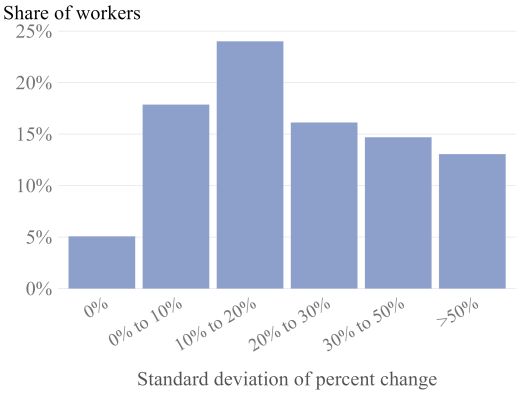


Figure A-6: Industry Heterogeneity

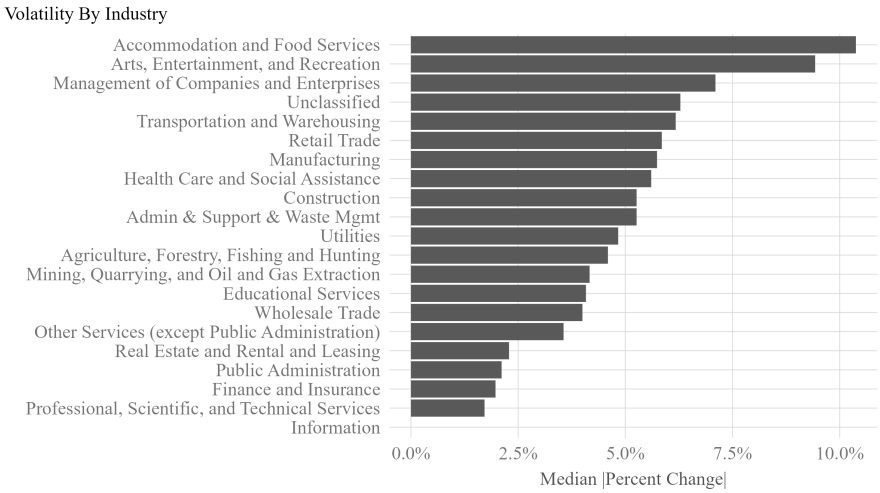


Table A-1: Correlation of Earnings and Hours Changes

	R^2 from $\Delta \log QuarterlyEarn = \alpha + \beta \Delta \log QuarterlyHours$	
Unit of aggregation	All workers	Firms using PayrollCompany
Worker	0.50	0.54
Firm	0.75	0.70

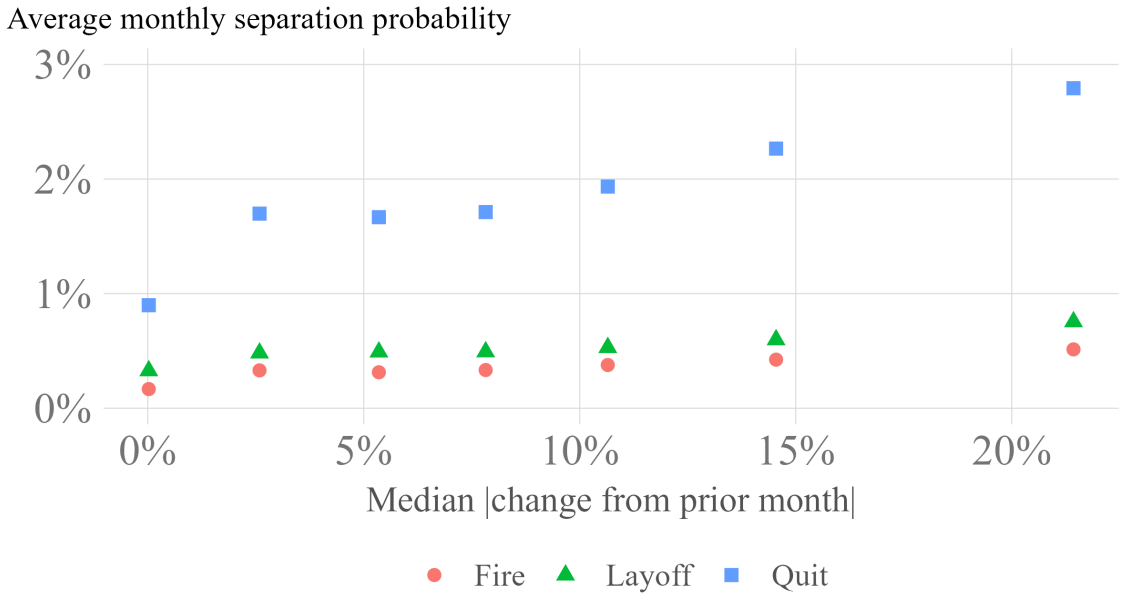
Notes: "All Workers" comes from Lachowska et al. (2023) analysis of Washington state, which asks employers to report quarterly earnings and hours. Sample in both datasets limited to full quarter employment.

Table A-2: Contract Type in PayrollCompany versus Representative Benchmarks

		Representative benchmarks	
	PayrollCompany	All workers	Workers at small firms
Hourly	60%	58%	60%
Salaried	40%	42%	40%
All: bonus	37%	40%	36%
All: no bonus	63%	60%	64%

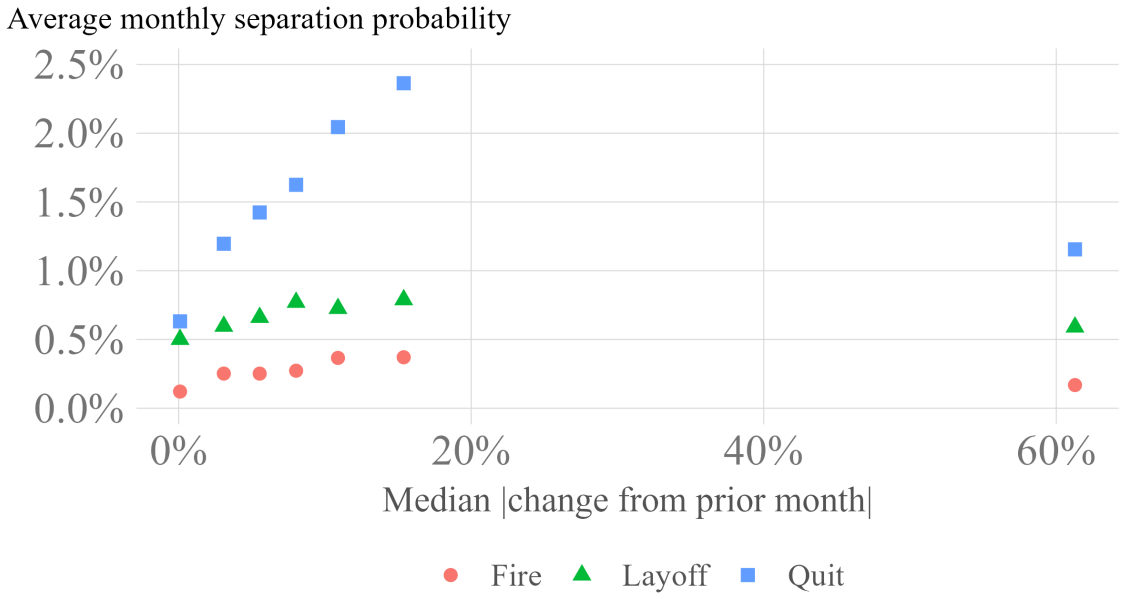
Notes: The representative share of workers that are hourly versus salaried comes from the Current Population Survey (CPS) using a sample of workers who respond to both the Earner Study and the ASEC. Small firms are defined for the CPS as less than 100 workers. In the PayrollCompany data, we classify 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.

Figure A-7: Worker-level Volatility and Separation Rate



Notes: This figure shows the relationship between worker-level pay volatility (measured using typical changes in earnings from prior month) and separation rates. See text for details on how pay volatility and separation rates are measured. Workers are stratified into deciles of volatility. This figure is identical to Figure 10 except it is at the worker level.

Figure A-8: Firm-level Volatility and Separation Rate – Including Top Decile Volatility



Notes: This figure shows the relationship between firm-level pay volatility (measured using typical changes in worker earnings from prior month) and separation rates. See text for details on how pay volatility and separation rates are measured. Firms are stratified into deciles of volatility. This figure is identical to Figure 10 except it includes the top decile of earnings volatility.

Table A-3: Earnings Volatility Under Different Winsorization Choices

Specification	Lower bound	Upper bound	SD
No winsorization			32%
Winsorize top/bottom 0.1% of nonzero changes	-2.65	2.60	31%
Winsorize top/bottom 0.5% of nonzero changes	-1.61	1.57	29%
Winsorize top/bottom 1% of nonzero changes	-1.22	1.18	27%
Winsorize top/bottom 5% of nonzero changes	-0.52	0.51	20%
Winsorize top/bottom 1% of all changes	-1.03	1.00	26%
Winsorize changes larger than 50%	Bottom 2.1% of data	Top 5% of data	20%

Notes: The variable is the first difference of log "Total Earnings" in every row except the last row where it is percent changes in pay.

Table A-4: Earnings Volatility in Chase

Aggregation Level	Condition	σ	Share of $\Delta \neq 0$	Median $\lvert\Delta\rvert$	75th perc $\lvert\Delta\rvert$
Job	None	0.31	0.74	0.04	0.17
	Less than 10 employees	0.27	0.69	0.03	0.15
	11-100 employees	0.27	0.74	0.04	0.16
	Greater than 100 employees	0.32	0.74	0.05	0.18
Household	None	0.33	0.78	0.05	0.18

Notes: Firm size is measured as the number of Chase customers who are employees of the firm. All statistics for Chase customers, including medians, reflect cells with multiple observations.

Table A-5: Quits, Layoffs, and Fires by Volatility

	Baseline	Layoffs	Fires
	(1)	(2)	(3)
Median $\lvert\Delta\rvert$	0.120^{***}	0.010^{***}	0.020^{***}
	(0.003)	(0.003)	(0.001)
Log wage	-0.005^{***}	-0.002^{***}	-0.0005^{***}
	(0.0004)	(0.0003)	(0.0001)
Observations	17,202	17,202	17,202
Adjusted R^2	0.096	0.007	0.015

Notes: This table shows the relationship between firm-level pay volatility (measured using typical changes in worker earnings from prior month) and separation rates. See text for details on how pay volatility and separation rates are measured. Column (1), labeled 'baseline' is the same regressions as A-6 Column (1) regresses the quit rate on pay volatility, column (2) regresses the layoff rate on pay volatility, and column (3) regresses the firing rat on pay volatility. Data is from PayrollCompany. We can only observe the separation type from May 2020 onward; during this period, 65% of separations are quits, 23% are layoffs, and 13% are fires.

Table A-6: Firm-level Volatility and Quits: Panel

	Quarterly quit rate			
	Baseline	Lagged X	No quits last qtr	No quits; Lagged X
	(1)	(2)	(3)	(4)
Median $\|\Delta\|$	0.114^{***}		0.057^{***}	
	(0.002)		(0.002)	
Lag median $\|\Delta\|$		0.108^{***}		0.051^{***}
		(0.002)		(0.002)
Log wage	-0.004^{***}	-0.004^{***}	-0.001^{***}	-0.001^{***}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Mean Quit Rate	0.011	0.011	0.007	0.007
Observations	160,710	138,943	109,437	108,908
Adjusted R^2	0.034	0.031	0.010	0.008

*p<0.1; **p<0.05; ***p<0.01

Notes: This table shows the positive relationship between pay volatility and quit rates. Column (1) in this table repeats column (1) from Table 6. Column (2) uses lagged firm pay volatility as the regressor. Column (3) restricts the sample to firms who did not have a worker quit in the previous quarter. Column (4) uses the same sample as column (3) and uses lagged firm pay volatility as the regressor. All regressions have a linear control for hourly wages, assuming that salaried workers work 40 hours per week. See text for details on how we aggregate from worker volatility and quit rates to firm-level variables.

B Appendix

A.2.1 Prevalence of Unpaid Vacation

The Current Population Survey asks workers if they were physically present at work this week, the reason for their absence, and whether their absence was paid or unpaid. Each week, 0.37% of workers say that they were employed but absent from work for the entire week, the reason was vacation, and the vacation was unpaid. We adjust this statistic for two reasons why it will understate the extent of unpaid vacation.

The first reason is partial-week absences. The CPS also measures whether full-time workers worked part-time this week and the reason for their partial-week absence. However, we cannot observe whether workers were paid for these partial-week absences. Thus, in order to include unpaid partial-week vacations in our calculation of the share of hours worked taken as unpaid vacation, we carry out the following procedure: first, we calculate the share of vacations that are taken over a full-week (33.08%) versus a partial-week (66.92%). Next, we calculate the median hours reduction associated with a full-week absence (40) and a partial-week absence (10). Taken together, these statistics allow us to estimate the share of hours reductions attributable to full-week (66.41%) and partial-week (33.59%) absences. We use this fact to rescale the full-week unpaid vacation rate (0.37%) by 1.51 (1/.6641) which gives us that unpaid vacation is equal to 0.56% of all hours worked. This rescaling procedure assumes that workers are equally likely to be paid for vacations regardless of duration.

The second reason is the underreporting of paid vacation in the CPS. The National Compensation Survey (NCS) asks employers detailed questions about their employee compensation costs, which can be used to calculate the share of wage and salary income associated with specific categories of benefits such as vacation leave. According to the NCS, paid vacation is equal to 5.24% of wage and salary income (Bureau of Labor Statistics 2024a, Bureau of Labor Statistics 2024b). Assuming that paid vacation compensation is distributed similarly to wage and salary income, then paid vacation is equal to 5.24% of hours worked in the NCS. According to the CPS (following the same procedure as above for unpaid leave), paid vacation is equal to 1.86% of hours worked. Note that we include all full-time workers in this calculation to more closely approximate the NCS sample. We follow the same procedure used above for unpaid leave: namely, by multiplying the full-week paid vacation rate (1.23%) by a rescaling factor (1.51) to account for the fact that hours reductions from full week absences are 66.4% of total hours reductions, which gives us that paid vacation is equal to 1.86% of hours worked. If we believe that the employer survey (NCS) is a more accurate source for the reporting of leave than the household survey (CPS), then the CPS underreports paid vacation by a factor of 2.82. If we assume that unpaid vacation is similarly underreported in the CPS and rescale the unpaid vacation rate by 2.82, then unpaid vacation is equal to 1.58% of hours worked. Note that the NCS captures leave for all workers whereas in the CPS, we derive an unpaid vacation estimate for full-time, hourly workers. Applying the NCS

rescaling factor assumes that vacation rates are similar between these two groups. Thus, our best estimate is that unpaid vacation for full-time, hourly workers is equal to between 0.56% (not rescaled to NCS) and 1.58% (rescaled to NCS) of all hours worked.

A.2.2 Burdett-Mortensen with a Job Amenity

We are interested in estimating a worker's willingness to pay for a job attribute (specifically, volatility σ). We are interested in characterizing the set of wages that compensates each worker for their respective volatility . The fundamental identification problem is that we do not observe pairs of jobs or vectors of jobs over which workers are in fact indifferent. This is the classic problem of estimating compensating differentials in the presence of worker-level heterogeneity and/or market-level frictions. Our solution to this problem is to use the separation rate as a measure of the utility provided by various job bundles. The model closely follows Chapters 2 and 8.2 of Manning (2003) which is itself a simplification of Burdett and Mortensen (1998).

Environment There are M_w workers who are equally productive and have identical preferences $U(w, \sigma)$. Not working delivers $U(b, 0)$.

There are M_f employers who are infinitesimally small relative to total employment each with constant returns-to-scale who can hire workers with productivity of p.

Each firm is endowed with a wage volatility $\sigma \in [0, \bar{\sigma}]$ and chooses w once and for all to maximize profits. Let the cumulative distribution function of firm-offered wages follow $F(w)$ and the joint distribution of wages and volatility be $G(w, \sigma)$.

Workers engage in random search. Offers arrive at rate λ drawn at random from the set of firms, which therefore have distribution $G(w, \sigma)$. Employed workers see their jobs destroyed at rate δ.

Worker Behavior Employed workers only accept offers to work at better jobs, each of which are evaluated using $U(w, \sigma)$. It is useful to note that for every offer (w, σ) there is a wage $\tilde{w}$ such that $U(w, \sigma) = U(\tilde{w}(\sigma), 0)$. This wage $\tilde{w}(\sigma)$ is a money metric for the value of an offer (w, σ). A non-employed worker will take any job that pays $\tilde{w}(\sigma) > b$ (since being unemployed involves no volatility). This specification of utility also gives rise to a distribution of monetary value of wage offers $\tilde{F}(\tilde{w})$.

Employer Behavior Employers solve

$$\max_{w} \pi = (p - w)N(w, \sigma; \tilde{F}) \tag{A1}$$

where $N(w, \sigma; \tilde{F})$ is the steady-state level of employment in a firm that pays w when the monetary distribution of offers $\tilde{F}$. For a firm's size to be constant it must be that

$$s(w, \sigma; \tilde{F})N(w, \sigma; \tilde{F}) = R(w, \sigma; \tilde{F}) \tag{A2}$$

where s is the separation rate and R is the recruitment rate. Manning (2003) proves that

the wage distribution $F(w)$ is continuous. The same argument applies to $\tilde{F}$.

Indifference Condition For any job offer $\{w, \sigma\}$ there is a set of job offers $\{w', \sigma'\}$ which deliver the same utility

$$U(w', \sigma') = U(w, \sigma) \forall \sigma' \in [0, \bar{\sigma}] \tag{A3}$$

Separation Rate The separation rate is a function of three forces

$$s(w, \sigma; \tilde{F}) = \delta + \lambda[1 - G(w, \sigma)] \tag{A4}$$

(1) the job destruction rate, (2) the offer arrival rate, and (3) the fraction of offers that are accepted. We will focus on just the quit rate $q = \lambda[1 - G(w, \sigma)]$.

Willingness to Pay to Reduce Volatility For a job offer $(w, \bar{\sigma})$, we are interested in finding the money metric equivalent offer $\tilde{w}(\bar{\sigma})$ which is defined as $\{\tilde{w}(\bar{\sigma}) : U(\tilde{w}, 0) = U(w, \bar{\sigma})\}$. Note that this is equivalent to characterizing the derivative $\frac{\partial w'}{\partial \sigma'}$ over the range from $\sigma' \in [0, \bar{\sigma}]$. We can then use the chain rule to express $\frac{\partial w'}{\partial \sigma'}$ as a function of the derivative of each variable with respect to quits:

$$\frac{\partial w'}{\partial \sigma'} = \frac{\partial w'}{\partial q} \times \frac{\partial q}{\partial \sigma'} = \frac{\frac{\partial q}{\partial \sigma'}}{\frac{\partial q}{\partial w'}}.$$

The willingness to pay to for volatility is therefore the ratio of $\frac{\partial q}{\partial \sigma'}$ to $\frac{\partial q}{\partial w'}$. We estimate $\frac{\partial q}{\partial \sigma'}$ in the paper and use an estimate of $\frac{\partial q}{\partial w'}$ from the prior literature.