

NBER WORKING PAPER SERIES

MONTHLY EARNINGS VOLATILITY AND HOUSEHOLD POOLING

Martin E. Andresen
Andreas R. Kostøl
Ross T. Milton
Corina Mommaerts
Luisa Wallossek

Working Paper 34563
<http://www.nber.org/papers/w34563>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2025

We thank Jeremy Lise and seminar audiences at the Celebratory Conference in Honor of Costas Meghir and the University of Wisconsin - Madison for helpful comments and suggestions. Andresen, Kostøl and Wallossek gratefully acknowledge support from the Norwegian Research Council (grant no. 345156). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Martin E. Andresen, Andreas R. Kostøl, Ross T. Milton, Corina Mommaerts, and Luisa Wallossek. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Monthly Earnings Volatility and Household Pooling

Martin E. Andresen, Andreas R. Kostøl, Ross T. Milton, Corina Mommaerts, and Luisa Wallossek
NBER Working Paper No. 34563

December 2025

JEL No. D13, D31, J12, J31

ABSTRACT

This paper examines monthly earnings volatility and its transmission to household earnings volatility using Norwegian data on the universe of monthly pay histories. We document substantial month-to-month earnings changes: within a job, while over one-quarter of months have no earnings changes, another quarter have at least a 23% change. Accounting for multiple jobs and non-employment increases volatility, while aggregating to households reduces volatility by 12-35%. Event studies around job loss and couple formation, along with decomposition and bounding exercises, show that most of this decline reflects pooling effects rather than sorting or responses to shocks.

Martin E. Andresen
University of Oslo
Department of Economics
martin.eckhoff.andresen@gmail.com

Andreas R. Kostøl
BI Norwegian Business School
Department of Economics
andreas.r.kostol@gmail.com

Ross T. Milton
University of Wisconsin - Madison
Department of Economics
rtmilton@wisc.edu

Corina Mommaerts
University of Wisconsin - Madison
Department of Economics
and NBER
cmommaerts@wisc.edu

Luisa Wallossek
University of Oslo
Department of Economics
luisa.wallossek@econ.uio.no

1 Introduction

Labor income volatility, and how it translates into consumption, is a key input to welfare and the design of safety net programs (Meghir and Pistaferri, 2011). A large literature documents individual earnings dynamics and volatility (e.g., Gottschalk and Moffitt, 1994; Meghir and Pistaferri, 2004; Guvenen et al., 2021) and to a lesser extent household income dynamics (e.g., Pruitt and Turner, 2020; Altonji et al., 2024), but this literature predominantly analyzes volatility at the annual level. However, most households earn income at much higher frequencies, and thus any sub-annual fluctuations in earnings would be missed by annual measures. If households struggle to smooth these higher frequency fluctuations (as much evidence suggests, see e.g., Larrimore et al., 2025), this could have important welfare implications as well as policy implications for the timing of safety net benefits. The recent rise of gig work, precarious work schedules, and short-term contract work makes within-year volatility all the more important to document (Schneider and Harknett, 2019; Garin et al., 2025).

In this paper, we examine monthly earnings volatility and the role of partners in shaping household earnings volatility using administrative earnings records from the universe of individuals in Norway from 2015 to 2023. This data has several advantages relative to other work in this space. First, we observe *monthly* labor earnings of individuals, a frequency that is seldom found in large datasets used to study earnings volatility.¹ Second, we observe all employment spells, which allows us to account for the roles of non-employment, job changes, and multiple job-holdings that hamper estimates based on data from banks or payroll processors. Finally, we are able to link individuals to their household members, which allows us to construct measures of both individual and household earnings volatility, which we believe we are the first to do at the monthly level.

We begin the paper by documenting the extent of month-to-month earnings changes in our sample, which consists of individuals aged 25 to 66 with labor earnings in at least parts of our nine-year sample period. To motivate the analysis at the monthly level, we first decompose the total variance of earnings within and across years, and find substantial within-year variation: on average, around 70% of the within-individual variance is within-year rather than across years, suggesting that annual earnings measures are insufficient for capturing the full extent of earnings volatility.

We then provide summary statistics on month-to-month earnings changes. Within

¹Notable exceptions are Druedahl et al. (2025), Ganong et al. (2025), and Brewer et al. (2025).

a job spell, we document that 29% of months have no change in earnings from one month to the next, while 25% have changes of 23% or more (the remaining 46% of months have non-zero but smaller changes). Changes are relatively symmetric, with earnings increasing in 36% of months and decreasing in 35%. These results are remarkably similar to the findings in Ganong et al. (2025) using payroll data from the US and Brewer et al. (2025) using tax data from the UK. Aggregating over job spells to allow for multiple job-holdings, job changes, and periods of non-employment slightly increases average month-to-month fluctuations, while aggregating to the household level decreases the extent of fluctuations.

Next, we summarize monthly earnings volatility using three complementary measures that highlight the strengths of our data and capture different features of volatility. These measures include two variants of arc percentage changes in monthly earnings – the mean absolute value and the within-individual (or household) standard deviation – and the within-individual (or household) coefficient of variation (CV) of monthly earnings. We focus on these measures for two main reasons (see Brewer et al. (2025) for more discussion). First, we want to capture the extensive margin of work. While much of the literature that estimates models of earnings processes has focused on measures involving log (annual) earnings (e.g., Meghir and Pistaferri, 2004; Blundell et al., 2008), such measures cannot handle periods of zero earnings, an issue that is exacerbated with higher frequency data. Second, our focus is on within-individual (or within-household) volatility, rather than across-individual inequality. Studies that measure volatility using standard deviations or variances (whether in logs or levels) often conflate within- and across-individual variation (see Shin and Solon (2011) for a discussion of this point). We thus avoid capturing across-unit inequality by calculating unit-specific standard deviations, which our data is well-suited to do since we have many monthly observations per unit. In the appendix, we show illustrative examples of the type of earnings fluctuations the three measures capture, and derive conditions under which household volatility is lower than average individual volatility.

We find substantial variation in monthly earnings. Even within a job spell, the average change in monthly earnings is 18%, as measured by the mean absolute arc percentage change, and the average within-job spell standard deviation of arc percentage changes in earnings is 29%. This shows that month-to-month earnings changes are both large and vary within individual job spells. Both of these measures focus on month-to-month changes, while the CV captures the overall spread of an individual’s earnings over time relative to average earnings. We find an average CV of 0.29, which implies that, on average, the standard deviation in earnings within

an individual is 29% of the individual’s mean earnings. Aggregating to allow for multiple job-holdings and periods of non-employment increases all three measures of volatility, but especially the CV and the standard deviation of the arc percentage change, likely because these measures place relatively more weight on large changes (e.g., extensive margin changes) than the mean absolute arc percentage change. In contrast, aggregating further to the household level leads to a decrease in volatility as measured by household-average earnings: for the subset of new couples in our sample, mean absolute changes decrease from an average change of 24% to 21% (a 12% decline), while the standard deviation of the change and the CV decrease by 24% and 35%, respectively.

In the final part of the paper, we examine why average household earnings are less volatile than individual earnings. We explore three main mechanisms: (1) pooling: averaging monthly earnings across individuals may mechanically lower volatility; (2) assortative matching: individuals may choose partners with positively (negatively) correlated earnings processes, which would attenuate (amplify) the pooling channel; and (3) partner insurance (or added worker effects): individuals may endogenously change their labor supply in response to a partner’s earnings shock. We first assess the role of partners as insurance and show that when an individual experiences an involuntary job loss, there is virtually zero response in partners’ earnings, immediately or up to 24 months afterwards. This result is remarkably similar regardless of gender or whether the outcome is extensive margin versus intensive margin partner responses.

We then conduct three exercises to differentiate between pooling and sorting. First, we test whether marital sorting *could* play a role by conducting a bounding exercise that estimates the minimum and maximum household earnings volatility that could arise if individuals were sorted to minimize and maximize household volatility. We find that maximized household volatility is over three times higher than minimized household volatility, and similarly, individual volatility is over twice as high as minimized household volatility. These patterns imply that if couples strategically sorted based on volatility it would significantly decrease household volatility.

Second, we explore whether marital sorting *does* play a role. In an event study analysis around the year couples form, we trace out year-by-year measures of volatility at the individual and household level. We show that true household volatility is remarkably similar to household volatility when individuals are randomly matched, suggesting that marital sorting does not play an important role in the reduced volatility of couples relative to individuals.² In contrast, we find that even volatility of ran-

²Despite the lack of marital sorting on earnings volatility, we show that there is substantial

domly matched couples is substantially lower than individual volatility, implying that the mechanical effect stemming from pooling earnings accounts for the gap between individual and household volatility. Moreover, volatility does not change meaningfully when couples form, again suggesting the limited role of partner insurance.

Finally, we quantify the role of pooling and sorting on household volatility directly. To do this, we decompose the differences between average individual volatility and household volatility using our CV measure into a set of pooling and sorting components. We again show that virtually all of the decrease in volatility when aggregating to the household level is due to mechanical pooling effects rather than assortative matching.

This paper contributes first and foremost to a large literature on the evolution of income risk and its role in explaining growing income inequality (e.g., Gottschalk and Moffitt, 1994, 2009; Dynan et al., 2012), and on income process estimation (e.g., Abowd and Card, 1989; Meghir and Pistaferri, 2004; Blundell et al., 2008; Guvenen et al., 2021).³ This literature has largely focused on volatility of male earnings, though some studies have examined household volatility (Ostrovsky, 2012; Altonji and Vidangos, 2013; Blundell et al., 2015; De Nardi et al., 2020; Halvorsen et al., 2024; Shiu et al., 2025), finding evidence that household volatility is typically lower than individual volatility. The vast majority of this literature, however, examines annual volatility. A few very recent papers leverage monthly earnings data, including data from one large payroll processor in the United States (Ganong et al., 2025), administrative data for men in Denmark (Drue Dahl et al., 2025), and a 1% sample of taxpayers in the UK (Brewer et al., 2025), and typically find that monthly volatility is more widespread than annual measures imply. Our paper adds to this nascent literature by capturing volatility arising not only within a job, but across jobs and non-employment spells and across household members, which these papers do not.

We also contribute to a literature on the role of partners in shielding individuals from volatility, either through marital sorting or added worker effects. While a body of work studies assortative matching in marriage along traits such as education or income and its effects on inequality (e.g., Greenwood et al., 2014; Hryshko et al., 2017; Eika et al., 2019; Chiappori et al., 2025), there has been much less focus on matching along volatility (one exception is Shore, 2015), though some work shows that sorting on other characteristics – such as sector or occupation – leads to sorting

positive sorting on earnings *levels*, especially on the permanent component of earnings (in line with the literature) but also on the annual and monthly components.

³See Moffitt et al. (2022) for a recent review and Guvenen et al. (2022) for a recent project documenting trends in income inequality and income dynamics around the world.

along earnings volatility (e.g., Hryshko et al., 2017; Busch et al., 2023). A separate literature studies the role of spousal labor supply as insurance against job loss (see, e.g., Lundberg, 1985; Cullen and Gruber, 2000; Stephens, 2002; Halla et al., 2020). Consistent with our findings, Pruitt and Turner (2020) and De Nardi et al. (2020) also find that the family is a relevant source of insurance due to pooling rather than partner labor supply responses, though they do not isolate causal responses to involuntary job loss.

Our findings of a lack of effects of marital sorting or partner insurance on monthly volatility begs the question of whether monthly volatility is welfare-relevant. If such higher frequency earnings fluctuations are easily smoothed through mechanisms such as savings or high-frequency transfers (e.g., safety net programs or informal insurance from family and friends), then it is no surprise that individuals do not necessarily choose partners or labor supply with monthly volatility in mind. While an examination of the consequences of monthly volatility on consumption and welfare is beyond the scope of this paper, other work suggests that monthly volatility can have important welfare implications. For example, survey data from the US suggest that variability in monthly income causes financial hardship and other adverse outcomes for some families, particularly low-income families (Larrimore et al., 2025; Gennetian et al., 2015), while Ganong et al. (2025) shows that accounting for monthly fluctuations increases individuals’ willingness to pay to eliminate volatility substantially relative to annual measures of volatility.

The paper proceeds as follows. In Section 2, we describe our data and sample restrictions, document the distribution of monthly earnings changes, and quantify the importance of within-year volatility relative to across-year volatility. Section 3 discusses our three volatility measures and quantifies the volatility of individual and household monthly earnings using each measure. We explore mechanisms underlying the difference between individual and household volatility in Section 4, including mechanical pooling, marital sorting, and partner insurance. Section 5 concludes.

2 Monthly earnings data and descriptives

The key to our empirical contribution is access to nine years of monthly labor earnings data in combination with population data, allowing us to provide a complete characterization of volatility at the individual and household level. In this section, we describe the data sources and provide some descriptive facts about the Norwegian setting. Norway is a small, high-income economy with a total labor force of around 3

million people, where employees account for 95 percent of workers, unemployment is low, and female labor force participation is high.⁴ Similar to the standard paycheck frequency in other European countries, “it is common [for workers] in Norway to receive payment once a month”, underscoring the central role of monthly wage earnings in Norwegian households.⁵

2.1 The data

Our main data source is monthly earnings records covering the universe of workers in Norway from 2015 to 2023. Employers file these records, typically through automated payroll systems, for all workers each month. The records contain total earnings and their components, including fixed salary, hourly pay, and overtime pay, and a range of other variables, including establishment and firm information. Note that (a) employers are required to file earnings statements for each worker every month and (b) months are the typical paycheck frequency, so we largely avoid issues related to sub-monthly pay periods that are a concern in other settings like the US and the UK (Brewer et al., 2025; Ganong et al., 2025).⁶ These earnings data are pre-tax, and thus will not capture differences in after-tax pay coming from the tax system or tax withholding, but it will capture additional payments such as end-of-the-year bonuses or June holiday pay. We keep earnings in nominal terms so that our measures of volatility do not pick up discrete changes in real earnings arising from inflation adjustments.⁷

We link these earnings records to annual population registers containing information on demographic characteristics, such as residence status and cohabitation status. We base our measure of partnership on cohabitation rather than marriage as

⁴See <https://ilostat.ilo.org/data/country-profiles/nor/>.

⁵See <https://www.arbeidstilsynet.no/en/pay-and-engagement-of-employees/pay-and-minimum-rates-of-pay>.

⁶The fact that months have slightly different numbers of days means that there may still be some variation in earnings across months for hourly workers. However, fewer than 20% of the job spells in our data are hourly workers, so we do not think this has a major impact on our results.

⁷Whether or not to adjust for inflation in the context of monthly volatility is not clear. While it can account for some of the monthly variation that is only nominal, it can also introduce excess volatility if households pay more attention to month-to-month changes in take home pay rather than month-to-month changes in purchasing power. For example, monthly-adjusted earnings are very rarely the same in two consecutive months despite a significant fraction of workers experiencing zero change in nominal earnings. Other work on monthly volatility has not coalesced on this issue: Ganong et al. (2025) do not mention inflation adjustments (but likely do not adjust, given the high percentage of zero month-to-month changes in earnings), Druedahl et al. (2025) use nominal pay, and Brewer et al. (2025) use monthly inflation adjustments, but find similar results when using nominal values.

cohabitation without marriage is common in Norway. The register includes a partner identification number for individuals who, as of January 1st of a given year, cohabit with a partner, whether married or unmarried. We define a partnership as having a reported partner identification number and civil status of married or single, excluding separated, divorced, and widowed individuals to minimize measurement error in partnership status. Together, the linked data cover all forms of formal employment and earnings for all individuals in Norway and allow us to link partners, making them well-suited for measuring household earnings volatility.

Sample restrictions From this universe of individuals in Norway, we restrict the sample to individuals with labor earnings. Specifically, we restrict the sample to individuals who are alive, a resident of Norway, and of working age, defined as age 25 to 66, for the entirety of the nine year period from 2015 to 2023, to avoid volatility stemming from initial labor force entry, retirement, and migration.⁸ We further restrict the sample to individuals with wage earnings over 1,000 NOK in at least one month of the sample period (1 USD is roughly 10 NOK). For months with no employment spells or with spells with monthly earnings below 1,000 NOK, we assign monthly earnings to be zero.⁹ We also drop a small number of individuals who ever (i) have reported negative earnings, (ii) have recorded employment spells without an employer ID, or (iii) are self-employed.¹⁰ We trim the remaining sample at the 1st and 99th percentiles of average individual earnings over the full period. Lastly, we drop individuals whose partners (at any point during the sample period) are dropped from the sample to keep the sample consistent when moving from individual to household volatility measures.

Our final sample consists of 767,219 individuals. See Table A.3 for details on how these restrictions impact our sample size. Much of the decline in sample size comes from dropping individuals who were not Norwegian residents for the full sample period, self-employed individuals, and individuals whose partners were dropped.

To analyze earnings volatility at the household (i.e., an individual and their part-

⁸Residency data is available only at the annual level, so it is still possible that we are picking up volatility arising from within-year moves abroad.

⁹This latter restriction is similar to Druedahl et al. (2025), which classifies individuals as unemployed if their earnings are below 1,000 DKK (roughly 1,500 NOK). We do this mainly because months with negligible earnings are similar to zero but can significantly skew mean-relative volatility measures.

¹⁰We drop self-employed individuals because it is only reported at the annual level, combines labor earnings and capital income, and is self-reported, which may contain more measurement error than employer-reported earnings. This is a common restriction in the earnings volatility literature (Gottschalk and Moffitt, 1994, 2009; Dynan et al., 2012).

ner) level as well as behavioral responses to partners' labor market shocks, we use three different couple subsamples. First, we use an *all couples subsample* that includes everyone who has a reported partner in a given year for the descriptive statistics in this section. Second, for some results in Sections 3 and 4 we use a *new couples subsample* that includes only couples for whom we observe couple formation during our sample period. Specifically, we restrict this sample to different-sex couples that we observe entering cohabitation (where event year $t = 0$ is the year prior to the year that the partnership is recorded, since the year partnership is recorded is based on cohabitation on January 1st so it is highly likely that they were cohabiting prior to that year), whom we do not observe in a partnership before, who cohabit for at least three years, and who are both observed at least two years prior to cohabitation.¹¹ This yields a balanced panel for event years $-2 \leq t \leq 2$. We further restrict the sample to couples where both partners' total earnings over the five years is positive, and only consider the first observed couple for both partners. In Section 4.2 we compare actual couples to random couples, which we create by assigning each female from the new couples subsample a randomly drawn male partner among the subset of men who began cohabiting in the same year.

Finally, for the job loss analysis in Section 4.1, we use a *job loss subsample* that restricts the sample to individuals who experienced an involuntary job loss between 2021-2022 and had a partner in the year they lost their job.¹² To concentrate on job losses that are economically meaningful shocks, we impose further restrictions on this sample. We restrict to individuals who have been employed at the firm for all 12 months prior to job loss and who have received wage income $\geq 1,000$ NOK at least once during the three months prior to job loss. To avoid capturing temporary layoffs, we further require that individuals do not return to that establishment (or any other establishment of the same company) within 12 months following the job loss. In addition, to better isolate shocks rather than anticipated job losses, we exclude job losses that occur in establishments that experienced substantial decreases in aggregate employment in the 12 months prior to the individual's job loss, defined as either at least one month-to-month decrease in the number of employees of over 10% in relative terms and 5 workers in absolute terms, or a decrease of 10% and 5 workers in average


¹¹We exclude same-sex couples here to be able to split couples into female and male partners, which we use for the random partner matching exercise in Section 4.2.

¹²To identify *involuntary* job losses, we use a unique feature of the Norwegian administrative data introduced in 2021 that records the reason why an employment spell ended. This allows us to identify displaced workers and exclude voluntary separations. We categorize individuals as experiencing an involuntary job loss from an establishment in month $t = 0$ if their employment spell with the establishment ends in that month and their employer initiated the termination.

employment between $-24 \leq m \leq -13$ and $-12 \leq m \leq -1$.

Summary statistics Table 1 reports summary statistics for our overall sample of individuals. The mean age in our sample is 46, and 50% of individuals are women. Most have a partner for at least parts of the sample period (64%), highlighting the importance of studying the role of partners in household earnings volatility. Monthly earnings are on average 42,000 NOK ($\approx 4,200$ USD). While the majority of individuals work most months in the sample period (on average 93 out of 108 months in total), zero earnings are not uncommon. Figure 1 plots the distribution of months with positive earnings and shows, for example, that over 10% of individuals have positive earnings in fewer than half of the months over the nine year sample period, highlighting the relevance of accounting for periods with zero earnings. Moreover, individuals on average have 3 employers over the nine-year sample period, with only 34% having the same employer for all spells, highlighting the importance of accounting for multiple employers.¹³

Figure 1: Distribution of months with positive earnings



Notes: Figure plots the cumulative distribution function of months with positive earnings among individuals in our sample. Individuals are observed for a total of 108 months.

¹³Most workers receive earnings from a fixed salary at least once (91%) and conditional on receiving salary at least once, they do so for the majority months (76%). At the same time, most workers also receive hourly pay at least once (71%), but for a smaller share of months (32%). Note that employees can receive both salary and hourly pay from the same employer in the same month.

Table 1: Individual summary statistics

Mean age	46
Share female	0.50
Share with a partner ever	0.64
Share of months with partner, conditional on ever	0.90
Individual-average monthly earnings: mean	42,108
Individual-average monthly earnings: standard deviation	23,372
Individual-average monthly earnings: 5th percentile	2,717
Individual-average monthly earnings: median	41,547
Individual-average monthly earnings: 95th percentile	84,392
Mean number of employers over sample period	3
Median number employers over sample period	2
Share with only one employer over sample period	0.34
Mean share of months with non-zero earnings	0.86
Share with earnings from salaried work ever	0.91
Share of months with salaried job, conditional on ever	0.76
Share with earnings from hourly work ever	0.71
Share of months with hourly job, conditional on ever	0.32
Number of individuals	767,219

Notes: Table reports summary statistics of individuals in our sample. Earnings in NOK. Individual-average monthly earnings are calculated as mean monthly earnings at the individual level over all months in the sample period (2015-2023), and then statistics (mean, standard deviation, percentiles) are calculated across individuals.

2.2 Distribution of monthly earnings changes

We next present descriptive statistics on the distribution of month-to-month changes in earnings. To do this, we measure monthly earnings changes as arc percentage changes, defined as $a_{i,m} = (y_{i,m} - y_{i,m-1}) / ((y_{i,m} + y_{i,m-1}) / 2)$ for individual or household i in month m .¹⁴ This is a common measure used to document volatility (see Brewer et al. (2025) for a comprehensive review), which we discuss in more detail, along with two other measures of volatility, in Section 3.1.

We find that individuals face frequent changes in their monthly earnings. Figure 2 plots histograms of the distributions of month-to-month earnings changes for different levels of aggregation. Panel (a) aggregates individual earnings across all jobs within a month, and shows histograms of total monthly earnings changes including and excluding months with zero earnings (dark and light pink, respectively). Excluding zero-earning months, 65% of month-to-month changes across all jobs are non-zero, and monthly earnings changes consist of both increases and decreases in earnings, but increases are slightly more common. Including months with zero earnings (darker pink bars) changes the distribution in two ways. First, the share of months in which earnings are the same in both months increases. Zero-earnings periods are often longer than one month, with a mean length of 7 months, and – by definition – earnings do not change during those periods. Second, very large increases and decreases become slightly more common, likely due to large earnings changes at the extensive margin when entering or exiting employment. In particular, the median absolute earnings difference at the extensive margin is 26,000 NOK, compared to 3,000 NOK at the intensive margin for periods with positive earnings.

Panel (b) of Figure 2 shows histograms of month-to-month changes in individual earnings (darker red bars) and average household earnings (lighter red bars) among couples, including all jobs and zero-earning months. Household earnings are less likely to experience zero change than individual earnings, and also less likely to experience very large (over $\pm 50\%$) changes, as shown by the lower light red bar than the darker red bar at exactly zero and the leftmost/rightmost bars, respectively.


Overall, Figure 2 shows a significant amount of variation in earnings at the monthly level.¹⁵ These numbers are remarkably similar to Ganong et al. (2025),

¹⁴By convention, we assign $a_{i,m} = 0$ when $y_{i,m} = y_{i,m-1} = 0$.


¹⁵Appendix Table A.4 reports further breakdowns of the distributions and Appendix Figure A.5 plots cumulative distribution functions. Appendix Figure A.6 shows histograms disaggregated by job (Panel a) and separately for salaried and hourly jobs (Panel b). Aggregation across jobs within a month does not change the distribution very much, while monthly earnings changes are much more common in hourly jobs than salaried jobs.

Figure 2: Distributions of month-to-month changes in earnings

(a) Individuals across all jobs: with vs without zero-earning months



(b) Individual vs household earnings among couples (all jobs, with zero-earning months)



Notes: Figure plots binned probability density functions of arc percentage changes in monthly earnings ($a_{i,m}$). Panel (a) plots earnings changes aggregated over all jobs, including months with zero total monthly earnings (darker pink bars) and excluding months with zero total monthly earnings (light pink bars). Panel (b) plots individual earnings (darker red bars) and the mean household earnings (lighter red bars) for the *all couples subsample*. For the specifications that include zero-earning months, in cases in which both months are zero we assign $a_{i,m} = 0$.

which uses US payroll data from 2010-2023, and Brewer et al. (2025), which uses UK monthly tax data from 2014-2019.¹⁶ Ganong et al. (2025) finds that within-job, 69% of months have a non-zero earnings change (compared to 71% in Norway), with 36% experiencing a positive change and 33% experiencing a negative change (compared to 36% and 35% in Norway, respectively), and the median absolute change is 4% (compared to our 7%). Their data, however, do not allow them to construct comprehensive measures across all jobs or for both partners in a household. Similarly, Brewer et al. (2025) find a mean absolute arc percentage change of 14% (compared to our 18%), but are also unable to calculate household earnings volatility.

2.3 Monthly versus annual variance decomposition

We next decompose the within-unit variance of monthly earnings over our sample period into within- and across-year components (where a unit is either an individual-job, individual, or household). This exercise quantifies the extent to which annual data captures earnings variation over time. To do this, we compute the within-unit variance of monthly earnings across all months, and decompose this variance into the variance of mean annual earnings and the within-year variance.¹⁷

Table 2 reports the ratios of these components to the total variance for several specifications that vary whether months of zero earnings are included and whether the unit of observation is the individual-job, individual (summing across jobs), or household (summing across jobs and across partners). Within a job, on average, 85% of the total variance is within-year. Including earnings from all jobs reduces this to 72%. Including months with zero earnings reduces the within-year share of the total variance further, to 66% for both individual and household earnings, but still well over half of the variance of total monthly earnings is within-year. Overall, this variance decomposition shows that across specifications, there is more variance within years than across, highlighting the relevance of higher frequency data for understanding earnings volatility.¹⁸

¹⁶Ganong et al. (2025) reports summary statistics using the percent change of first differences within an individual-job spell, which is similar to our arc percentage measure but divided by y_{m-1} rather than $0.5(y_m + y_{m-1})$.

¹⁷This requires estimating the variance of the annual means within person, which will be inflated by statistical noise because the annual means are estimated from at most 12 monthly observations. We bias-correct for this by subtracting off the average sampling variance of yearly means, assuming observations are i.i.d. More conservative methods of inference that allow for serial dependence would decrease precision, and, if anything, increase the within-year share of variation.

¹⁸The naive, unadjusted estimates produce a very similar pattern of results.

Table 2: Variance decomposition of total monthly earnings

	Excluding zeros		Including zeros	
	Individual-job (1)	Individual (2)	Individual (3)	Household (4)
Total variance (1,000s NOK)	355,461	429,343	445,001	323,452
Share within-year	0.85	0.72	0.66	0.66
Share across years	0.15	0.28	0.34	0.34

Notes: Table reports total variance of monthly earnings across the sample period for the unit of observation (as denoted in the second heading). Columns (1) and (2) exclude months with zero earnings, while columns (3) and (4) include months with zero earnings. The share of the variance across years is calculated as the variance of mean annual earnings divided by the total variance, and the share of the variance within-year is the remaining variance. Column (4) uses the *all couples subsample*. The across-year variance is bias-corrected by subtracting off the average sampling variance of the yearly means.

3 Volatility of individual and household earnings

In this section, we first discuss three measures of volatility that we use to quantify monthly volatility in our data. We then report estimates of monthly earnings volatility and discuss how our estimates that exclude non-employment spells, multiple job-holdings, and job changes compare to estimates that include those features, as well as how estimates at the individual level compare to estimates at the household level.

3.1 Measuring volatility

Prior research has measured earnings volatility in a variety of ways, each with strengths and limitations. Standard approaches often focus on logged earnings – such as the variance or standard deviation of log earnings or their first differences – which summarize relative fluctuations but cannot accommodate periods of zero earnings (e.g., Gottschalk and Moffitt, 1994; Meghir and Pistaferri, 2004). This is a key limitation when working with high-frequency data, for which periods of zero earnings are likely to be more prevalent. For this reason, studies using monthly or quarterly data more commonly use other measures. One such measure is the coefficient of variation (the standard deviation of earnings divided by mean earnings), which can incorporate zeros and facilitate comparisons across individuals or groups, but are less informative about the timing of fluctuations (e.g., Keys, 2008; Moffitt and Ribar, 2008; Gennetian et al., 2015). Another measure that is common in studies using monthly or quarterly data is percentage or arc percentage changes in earnings between periods, which capture the magnitude and direction of earnings movements

and can accommodate transitions into and out of work (e.g., Dynan et al., 2012).

Guided by the necessity to handle zero-earning months, we use three primary measures of earnings volatility. The first is the within-unit normalized standard deviation, captured by the coefficient of variation (CV_i), where a “unit” i in our context is an individual-job, individual, or household, depending on the specification. The CV of unit i is then simply the ratio of the standard deviation of monthly earnings to the mean. This measure is well-suited for our analysis because it can handle zero-earnings months (conditional on the unit having at least one month in which $y_{i,m} > 0$, which our sample satisfies),¹⁹ it is scale-invariant, it treats positive and negative deviations identically (i.e., it is symmetric), and it allows for derivable comparisons across individuals and households within our framework, as discussed below. Moreover, in contrast to much of the literature that constructs measures of volatility using standard deviations or variances, this measure captures variation *within* a unit and not cross-sectional variation across units, which separates within-unit volatility from across-unit inequality (we show across-unit measures in Appendix Table A.5).

Our two other measures of earnings volatility are based on month-to-month arc percentage changes in earnings, $a_{i,m}$, as first introduced in Section 2.2. This is defined as the first difference in monthly earnings divided by the average earnings over those two months, with the convention that $a_{i,m} = 0$ if earnings are zero in both months. We focus on the absolute value of the arc percentage change ($|a_{i,m}|$) and the within-unit standard deviation of the arc percentage change ($SD_i(a_{i,m})$). These measures are also symmetric and scale-invariant, and are easily interpretable. One notable departure from much of the literature, however, is our focus again on within-unit standard deviations to capture volatility as opposed to inequality.

Our three measures perform slightly differently for different earnings patterns. For one, CV_i captures a longer-run (“multi-period”) notion of volatility, while the measures based on arc percentage changes capture month-to-month notions of volatility. There are two differences between our two arc percentage change measures. First, the $SD_i(a_{i,m})$ measure places more weight on larger earnings changes than the $|a_{i,m}|$ measure. This is also true about the CV_i measure and thus makes these measures more susceptible to outliers.²⁰ Second, since standard deviations are calculated within-unit, they represent deviations from units’ long-term trends in percent changes while the

¹⁹Note that when estimating the CV in subperiods (as we do in Section 4), mean earnings are not always non-zero. In those cases, we define $CV_i = 0$ if earnings are zero in all periods.

²⁰To evaluate the role of outliers, we also show robustness to median CV_i and median $SD_i(a_{i,m})$ across the population rather than means.

mean absolute deviation represents the average total magnitude of percent changes. Appendix A works through illustrative examples of these differences.

3.2 Individual monthly earnings volatility

Table 3 reports means of our three measures of monthly earnings volatility across different sample specifications.²¹ To account for the fact that units can differ in the number of months they appear in the sample in some specifications (e.g., for specifications that exclude zeros), we first calculate the within-unit mean of $|a_{i,m}|$ in column (2) to mimic the fact that $CV_i(y_{i,m})$ and $SD_i(a_{i,m})$ are within-unit, and then weight each within-unit value by the number of months that unit is observed.²² Moreover, to facilitate comparison across samples (i.e., down a column) for our CV_i measure, we assign the denominator for each sample to be the mean of raw earnings over all months, including zero-earning months even if the sample does not include zero-earning months. This ensures that differences in CV_i across samples is driven by changes to the standard deviation of earnings (the numerator) rather than mean earnings (the denominator).

Panel (A) of Table 3 reports estimates at the individual-job spell level. Within a job, the mean absolute monthly change in earnings is 18%, the standard deviation of monthly changes is 0.286, and the coefficient of variation for earnings across the sample period is 0.290. One way to interpret the latter estimate is that, for someone with earnings of 40,000 NOK (roughly the average monthly earnings in our sample) the standard deviation of their earnings over the sample period is approximately 29% of their income, or 11,600 NOK. Overall, all three measures show a substantial amount of volatility across months even within a job spell, suggesting that not only do job-to-job transitions and periods of non-employment generate volatility, but also earnings changes within a job.

Residualized volatility Some of this volatility could, of course, be capturing predictable changes rather than earnings uncertainty. For example, monthly earnings could follow seasonal patterns, including seasonal employment changes within indus-

²¹Appendix Table A.5 report analogous estimates for medians across units as well as calculating the $CV(y_{i,m})$ and $SD(a_{i,m})$ within *and* across units rather than within unit.

²²In other words, $|a_{i,m}|$ is inherently *unit-by-month*-specific while $CV_i(y_{i,m})$ and $SD_i(a_{i,m})$ are *unit*-specific, so by aggregating $|a_{i,m}|$ to the unit level, each volatility measure contains the same number of observations. This ensures that all volatility measures weights all individuals the same way.

Table 3: Monthly volatility measures

	$CV_i(y_{i,m})$ (1)	$ a_{i,m} $ (2)	$SD_i(a_{i,m})$ (3)
<i>Panel A: Individual-job</i>			
Excluding zeros	0.290	0.178	0.286
Excluding zeros, residualized	0.280	0.186	0.293
<i>Panel B: Individual</i>			
Excluding zeros	0.377	0.184	0.296
Including zeros	0.676	0.229	0.427
Including zeros, new couples	0.515	0.240	0.435
<i>Panel C: Household</i>			
Including zeros	0.552	0.218	0.377
Including zeros, new couples	0.337	0.210	0.332

Notes: Table reports estimates of means of our three volatility measures (columns 1–3) over different units of observation (panels A–C) and different subsamples (rows within panels). Column (1) reports the weighted mean of within-unit coefficients of variation, in which each coefficient of variation is scaled by the mean of raw earnings over all months (including zero-earning months) to maintain comparability across rows, column (2) reports the weighted mean of within-unit mean absolute arc percentage changes, and column (3) reports the weighted mean of the within-unit standard deviation of arc percentage changes, where all three measures are weighted by the number of months the unit is observed (which may differ across units for rows that exclude zeros and/or restrict to couple subsamples). The unit of observation is individual earnings from a particular job in panel (A), individual earnings aggregated across all jobs in panel (B), and per capita household (individual plus partner, if applicable) earnings, aggregated across all jobs and defined at the individual level (i.e., for couples there is an observation for each partner) in panel (C). The first row in panel (C) includes singles and couples, while the second row in panel (C) and the third row in panel (B) restricts the sample to the *new couples subsample*. For the “excluding zeros” rows, we exclude months with zero earnings, and for the “residualized” row, we use residualized rather than raw earnings (column 1) or earnings changes (columns 2 and 3), that remove firm-by-month-of-the-year effects (all three columns) and individual-firm fixed effects (column 1 only).

try, holiday pay (which most Norwegian employees are entitled to), and bonuses.²³ To account for such predictable volatility, we residualize earnings from firm-by-month-of-the-year fixed effects (to capture within-firm seasonal variation) and individual-firm fixed effects (to capture time-invariant individual-firm characteristics). Similarly, we residualize earnings changes from firm-by-month-of-the-year fixed effects, because individual-by-firm effects are differenced out in the first-differences equation. Appendix Table A.6 reports that the (partial) R^2 for these regressions are 0.2041 and 0.2488 for levels and changes, respectively. It also shows that these fixed effects can explain more volatility for employees who are paid hourly compared to salaried employees. Limiting the sample to individuals that are employed for at least 12 months at a given firm over the sample period has almost no impact on the partial R^2 estimates.²⁴

The second row of Panel (A) in Table 3 reports volatility measures based on these residualized earnings and earnings differences. Interestingly, using residualized earnings results in only a slight change in volatility across all three volatility measures. Interestingly, for the arc percentage change measures it results in a slight *increase* in volatility, likely due to the fact that residualization drastically reduces the likelihood of zero-change months.

Multiple jobs and the extensive margin Panel (B) of Table 3 reports volatility measures based on individual monthly earnings summed across all jobs. This allows us to not only provide a more accurate characterization of individual-level volatility but also to examine two potentially important components of volatility. The first is multiple job-holdings and job switching, which could potentially be sources of insurance and thus decrease volatility relative to within-job volatility.²⁵ The first row in Panel (B) shows, in contrast to the insurance channel, that accounting for multiple job-holdings and job switching increases volatility slightly for the arc percentage change measures and more meaningfully for the coefficient of variation measure.

²³Holiday pay is typically paid out in June based on earnings the previous calendar year. For details on holiday pay in Norway, see <https://www.arbeidstilsynet.no/en/working-hours-and-organisation-of-work/holiday/holiday-pay/>.

²⁴It may seem puzzling that our partial R^2 results imply a drop in the standard deviation of residualized earnings by $1 - \sqrt{1 - 0.2488} = 0.13$, while the drop in the average CV is only around 3 percent, which reflects the change in the average SD of residualized earnings as we keep the denominator fixed. This difference arises because the partial R^2 is a global measure, weighted by individuals residual variance and number of observations, whereas the average CV gives equal weight to each worker and averages square roots of variances rather than variances themselves.

²⁵Alternatively, individuals may treat a wage increase in one job as an outside option during wage bargaining in another job, and potentially reinforce transitory or persistent shocks to earnings. Lachowska et al. (2022) finds evidence for such effects at the upper end of the wage distribution.

The second component of volatility that we capture by aggregating across all jobs is the extensive margin. Because we observe the universe of jobs, we can also account for periods of zero earnings. On average, individuals have positive earnings for 81% of months, which highlights the potential relevance of the extensive margin. Including months with zero earnings further increases the average absolute month-to-month change in earnings from 18% to 23% and increases the average standard deviation of month-to-month changes in earnings from 30% to 43%. The coefficient of variation increases even more, from 0.377 excluding months with zero earnings to 0.676 including zero-earning months. This is in line with the theoretical prediction that the CV should increase when we include periods with zero income, for which we provide a proof in Appendix B.

3.3 Household monthly earnings volatility

Given that most individuals in our sample have a partner for at least part of the sample period, a more complete picture of earnings volatility is one that accounts for volatility of partners as well. We define household volatility as the volatility of mean per capita household earnings. If volatility was measured using the standard deviation (or equivalently, the variance) of earnings, average volatility would indeed unambiguously decrease upon pooling. However, the standard deviation is not our preferred volatility measure, and when moving to our preferred volatility measures, it is ambiguous how household volatility compares to individual volatility and thus an empirical question.²⁶

Panel (C) of Table 3 reports measures of household volatility, where households are defined as an individual and their partner (if they have one). Overall, earnings volatility decreases slightly across all measures (first row of Panel C). However, this may understate the role of partners, as more than 40% of the monthly observations are for single individuals. The third row in Panel (B) and the second row in Panel (C) thus restrict the sample to only couples, and in particular newly formed couples to be consistent with our analyses in Section 4.2 for which they are the focus. These rows show that household volatility declines by 12-35% depending on the measure. This is a substantial decline, particularly for volatility measures that place more weight on large changes in earnings (CV_i and $SD_i(a_{i,m})$). The next section investigates the

²⁶Appendix C provides a proof for the standard deviation result as well as conditions under which the standard deviation of household volatility is lower than the volatility of the lower-volatility partner. It also provides a derivation that relates the change in the CV measure of volatility to the relative means, standard deviations and correlation between spouses' earnings.

reasons behind this decline.

4 Understanding household volatility

This section explores three potential mechanisms behind our finding that household volatility is lower than individual volatility. First, how partners react to each other’s earnings shocks (i.e., partner insurance, or added worker effects) could affect household volatility. If partners insure each other’s labor market shocks (e.g., by working more hours if their partner loses their job, or changing jobs to make earnings less correlated), then household volatility will be lower than individual volatility; conversely, if partner earnings are complements (e.g., if partners take time off from work at the same time), then household volatility will be more similar to individual volatility.

Second, there could be a mechanical pooling effect. As shown in Section 3, simply pooling earnings at the household level typically lowers the volatility of average earnings within a couple relative to the average volatility of individual earnings. Finally, the extent to which pooling reduces volatility depends crucially on marital sorting (i.e., who marries whom): negatively correlated earnings processes between spouses will reduce household volatility relative to individual volatility to a much larger extent than positively correlated earnings processes.

This section examines the significance of these mechanisms by focusing on couples. We begin by testing for added worker effects using unique data on involuntary job loss, showing that there is little evidence of such effects in our setting. Given this finding, we then turn to two exercises that quantify the relative importance of household pooling and assortative matching.

4.1 The role of partner labor supply responses to job loss

We first explore the household insurance mechanism whereby a partner may offset lost earnings due to involuntary job loss. Our sample for this analysis is the *job loss subsample*, which includes individuals who experience an involuntary job loss from a meaningful worker-firm relationship and who have a partner in the year of job loss, as described in more detail in Section 2.

Figure 3 shows monthly event study estimates of employment and unemployment insurance benefit receipt (Panel a), and earnings of individuals and their partners

(Panel b) around the time of an involuntary job loss.²⁷ We define employment as having a spell with an employer, with no restrictions on earnings (e.g., an individual is still employed if they are taking unpaid leave from an employer). Panel (a) shows that after a job loss, individuals experience an immediate 50% decline in employment in the month after job loss. Employment gradually but only partially recovers to around 80% by 12 months after the shock, persisting at least 24 months out. The receipt of unemployment insurance benefits increases, reaching a peak of around 20% of individuals three months following job loss, but by 24 months following job loss, fewer than 10% receive unemployment insurance benefits despite only 80% being employed at that point. About 15% of displaced workers have neither found a new job nor received unemployment benefits 24 months after their job loss.

Panel (b) of Figure 3 reports monthly earnings of individuals experiencing job loss (solid circles) and their partners (hollow diamonds). In line with the employment effects, there is an immediate 60% drop in monthly earnings of individuals who experience job loss relative to the month prior to job loss.²⁸ 24 months after job loss, earnings recover somewhat, but average monthly earnings is still around 20% lower than it was before the job loss.

In contrast, while job losses lead to a permanent decrease in earnings for displaced individuals, we observe no visible employment or earnings response from their partners, either immediately following the displacement or in anticipation (a concern that other studies have pointed out, e.g., Stephens, 2002; Hendren, 2017). We also don't see any responses when we restrict the sample to displaced males (see Appendix Figure A.9), as is typically the focus of the literature on spousal labor supply responses.²⁹

One possible reason that partners do not respond is that displaced individuals can cushion their earnings loss with unemployment benefits, which up to 20% claim in the months after their job loss. To explore this, Appendix Figure A.8 splits the sample into displaced individuals who found a new job and/or received unemployment insurance benefits in at least one month during $1 \leq m \leq 12$ and those that had neither a new job nor unemployment benefits during that time. Intuitively, the latter group


²⁷Figure A.7 expands the sample to all displaced individuals (not restricted to individuals with partners) and shows similar results.

²⁸The sharp *increase* in earnings during the month of job loss can be attributed to severance pay, outstanding wages or holiday pay for the following year, which are often paid out immediately upon termination.


²⁹Note that we restrict our sample to couples in which both have positive labor income at some point during our nine-year sample period. As a result, the labor force participation of women is higher in our sample than in the total Norwegian population. However, Norway in general has a high female labor force participation of 81% for women aged 25 to 64 (OECD, 2025).

Figure 3: Labor market outcomes of individuals and partners after job loss

(a) Employment and unemployment insurance benefits



(b) Earnings



Notes: Figure plots the means and 95% confidence intervals of monthly labor market outcomes for individuals and their partners over time using the *job loss subsample*. Panel (a) plots the share of individuals in month m that are employed (mid-blue), receive unemployment insurance benefits (light blue), or either (darker blue), as well as the share of their partners who are employed (pink). The share of individuals employed is 100% for $-12 \leq m \leq -1$ by construction. Panel (b) plots monthly earnings for the individuals who experience a job loss (dark red) and their partners (light red) relative to $m = -1$. Vertical red line is the month they lost their job involuntarily ($m = 0$).

experiences a larger shock to their available earnings than the former, so one might expect a larger partner response for the latter group. While these results must be considered suggestive because the sample split is endogenous, we still do not find a response in partner earnings even among this group.

Another potential reason for the absence of observed partner responses is the potential presence of positively correlated shocks between partners, which could occur if partners work in similar occupations, industries, or local labor markets (Cullen and Gruber, 2000). Such correlation would attenuate any potential partner response. If this was the case, however, we would expect to see a decrease in employment among partners as well, which we do not observe.

Overall, our results suggest that partners, on average, do not insure one another’s labor supply shocks (as measured by job loss) by increasing – or even changing at all – their labor supply. This is perhaps unsurprising given Norway’s very high rate of female labor force participation, but also surprising given that existing evidence is mixed. While some European studies find evidence of spousal insurance through increased female employment (Halla et al., 2020), studies from other northern European countries, such as Denmark and the Netherlands, find small (or zero) responses (Andersen et al., 2023; De Nardi et al., 2020).³⁰

4.2 The role of marital sorting

We next examine the role of assortative matching in shaping earnings volatility. To do this, we conduct event study analyses of volatility around couple formation using the *new couples subsample*, as described in Section 2. For comparison, we also construct a measure of *random* household volatility as well as measures of the lower and upper bounds of household volatility based on hypothetical matching patterns that minimize and maximize household volatility, respectively.

This exercise provides several useful comparisons. First, the bounds of household volatility provide a sense of how large of a role marital sorting *could* play. Second, the difference between random household volatility and true household volatility provides a sense of how large of a role marital sorting *does* play. Third, the difference between individual volatility and random household volatility provides a sense of the mechanical earnings pooling effect, as this comparison is purged of marital sorting effects.

³⁰Spousal insurance could still exist for other types of shocks in these countries. For example, Autor et al. (2019) shows that spousal labor supply provides insurance against disability insurance denials, Persson (2020) finds that spouses increase labor supply after the elimination of survivors benefits, and Fadlon and Nielsen (2021) shows that spouses increase labor supply following the death of their spouse.

Finally, tracing volatility from several years prior to couple formation to several years after couple formation shows how volatility evolves around couple formation. Such event time effects around couple formation could arise, for example, from life-cycle changes (e.g., labor supply changes around childbirth, see Kleven et al., 2024) or the inclusion of partner insurance effects after couple formation.

Figure 4 shows measures of individual volatility (purple triangles) and true household volatility (pink with solid diamonds) from five years prior to couple formation to five years following couple formation, computed separately for each event year using within-year monthly earnings variation.³¹ In addition, the dashed pink lines with hollow diamonds show household volatility when couples are randomly matched, and the darker and lighter blue dashed lines with hollow circles show household volatility when couples are matched to maximize or minimize household volatility, respectively.³² Panel (a) reports our coefficient of variation measure, panel (b) reports the absolute arc percentage change measure, and panel (c) reports the standard deviation of the arc percentage change measure.³³ Overall, measures of household volatility are always lower than measures of individual volatility, and are roughly the magnitudes of those reported in Table 3.³⁴

Our first result is that marital sorting *could* play an important role in determining how individual volatility translates to household volatility. In particular, a comparison of the matching patterns that minimize and maximize household volatility show that maximum possible household volatility is over three times higher than minimum possible household volatility, and similarly, individual volatility is over twice as high as minimum possible household volatility. Moreover, the fact that the maximum possible household volatility only slightly exceeds individual volatility suggests that most possible matching patterns would reduce household volatility. These patterns imply that strategic marital sorting based on volatility could significantly decrease household volatility. The extent to which it does, however, depends on true matching patterns.


³¹Note that we still calculate household volatility prior to couple formation, but it should only be interpreted as what household volatility would have looked like if they had begun cohabiting earlier in the absence of behavioral responses affecting their earnings processes.

³²To find the matching pattern that minimizes household volatility, we solve an optimal assignment problem and consider the set of N females indexed by j and N males indexed by i who all (truly) match in the same calendar year, and solve $\min_{\{x_{ij}\}_{i,j=1}^N} \sum_{i=1}^N \sum_{j=1}^N x_{ij} V_{ij}$ subject to $\sum_i x_{ij} = \sum_j x_{ij} = 1$ and $x_{ij} \in 0, 1$ for household volatility measure V_{ij} separately for each event time year. We also find the analogous matching pattern that maximizes household volatility.

³³See Appendix Figure A.10 for individual volatility measures broken down by male and female.

³⁴The magnitudes are not exactly the same because of sample restriction differences and because measures here are computed over a calendar year span rather than the nine-year span.

Figure 4: Volatility around couple formation



Notes: Figure plots yearly means of our three main volatility measures for the *new couples subsample* over time relative to the event of cohabitation in year 0. Panel (a) plots the within-unit coefficient of variation, panel (b) plots the mean absolute arc percentage change in monthly earnings, and panel (c) plots the within-unit standard deviation of the arc percentage change in monthly earnings. We compute each measure separately for each event year using within-year monthly earnings variation. Individual-level measures of volatility are plotted in purple triangles, per-capita household-level measures in solid pink diamonds, and random household measures in dashed hollow pink diamonds. Dashed lighter and darker blue lines with hollow circles indicate the bounds for volatility that could be reached if couples were re-matched every period to achieve the lowest and highest possible household volatility, respectively.

Turning to the true matching patterns, the comparison of true household volatility to random household volatility suggests that marital sorting does *not* play an important role in the reduced volatility of couples relative to individuals. Interestingly, random and true household volatility are remarkably similar; the fact that true household volatility is slightly higher than random household volatility prior to couple formation suggests partners' earnings processes are positively correlated prior to couple formation, but this difference is quite small. Thus, the observed differences between individual volatility and household volatility do not appear to be driven by marital sorting.

Instead, we find that the mechanical effect stemming from pooling earnings is an important driver of the reduced volatility for couples. One way to see this is from the comparison of household volatility for random couples versus individual volatility. Random household volatility is substantially lower than individual volatility (e.g., a coefficient of variation of roughly 0.25 compared to 0.4, respectively), implying that pooling mechanically lowers volatility. Another way to see this is from the comparison of household volatility of true couples prior to their formation versus individual volatility. Under the assumption that they are not yet a couple 4-5 years prior to cohabitation, the fact that household volatility is lower than individual volatility also suggests a mechanical effect of pooling on volatility.

A final result from Figure 4 is that volatility does not change meaningfully when couples form. For all three volatility measures, there are no visible discontinuities around the time couples start cohabiting ($t = 0$). Moreover, the gap between individual and household volatility remains virtually unchanged over time. These patterns suggest that changes to earnings processes do not seem to occur around the time of household formation, again suggesting limited role for added worker effects or other endogenous responses to couple formation in shaping volatility in our setting.

4.2.1 Marital sorting on earnings levels

It is perhaps surprising that couples do not sort on earnings *volatility* in our setting given that a large literature finds strong assortative matching on earnings *levels* (Greenwood et al., 2014; Eika et al., 2019; Chiappori et al., 2025). Here, we show evidence that these two empirical facts are not in conflict: individuals in our setting do sort on earnings levels. In the next section, we show that this does not translate to an overall effect of sorting on volatility because sorting on earnings levels has a very minor influence on volatility, and individuals also sort on other dimensions that negate this minor volatility effect. To investigate sorting on earnings levels, we

extend the standard analysis to examine not only sorting on permanent and annual earnings, but also monthly earnings.

We model individual earnings at the monthly level for partner $J \in (M, F)$ as the sum of three uncorrelated components: a permanent term, an annual term, and a monthly term:

$$y_{J_i,m} = \alpha_{J_i} + \pi_{J_i,t(m)} + \theta_{J_i,m} \quad (1)$$

where α is the permanent component, which captures long-run earnings differences between individuals, π captures annual fluctuations to earnings in year t , and θ captures monthly fluctuations. This structure implies a simple three-way covariance decomposition:

$$\text{Cov}(y_{M_i,m}, y_{F_i,m}) = \text{Cov}(\alpha_{M_i}, \alpha_{F_i}) + \text{Cov}(\pi_{M_i}, \pi_{F_i}) + \text{Cov}(\theta_{M_i}, \theta_{F_i}) \quad (2)$$

and an analogous decompositions for the variance of male and female earnings. Because the data contains repeated observations at the monthly level over multiple years, the three components of the variance and covariance decompositions are identified from (1) mean earnings over the sample period, (2) mean earnings by year t , and (3) monthly earnings, at the individual and couple level respectively. Intuitively, the (co)variance of the long-run means identify the permanent component, comparing this to the (co)variance of yearly means pins down the annual component, and the residual variation identifies the monthly component.

Table 4 presents the covariance decomposition for the new couples sample, and reports results for true couples in column (1) and random couples in column (2). Overall, the total covariance is much higher among true couples than random couples (row 1), which indicates the presence of positive assortative matching on earnings levels.³⁵ Two-thirds of this covariance stems from the permanent earnings component, while annual and monthly components split the remainder.³⁶ This suggests that couples mostly sort on permanent earnings, but with some role for sorting on yearly and monthly earnings.

In column (3) of Table 4, we report results for stable couples, defined as couples

³⁵The total covariance is non-zero for random couples because we randomize partners within their true year of initial cohabitation, making the earnings of our random partners somewhat more correlated than a completely random individual.

³⁶Note that when aggregating the model to the annual level and estimating a permanent-annual model (as is more typically done in the literature), we find a permanent share around 77%, considerably larger than the 67% we estimate using monthly data. The comparable permanent share in Hyslop (2001) is 90%.

Table 4: Covariance decomposition of partner earnings

	New couples		Stable couples
	True (1)	Random (2)	True (3)
Total covariance	187	25	180
Share permanent component α	0.67	-0.19	0.68
Share yearly component π	0.19	0.83	0.18
Share monthly component θ	0.14	0.36	0.14
N couples	6,546	6,546	192,794
N monthly observations (per spouse)	705,702	705,702	20,821,752

Notes: Table reports the covariance (row 1) between partners' earnings for the new couples sample (columns 1 and 2) and the stable couples sample (column 3). Covariance shares (rows 2-4) are the share of the covariance attributable to the permanent, annual, and monthly components using the three-way earnings model described in the text. Covariances in millions of NOK².

whom we observe cohabiting together for the full sample period 2015-2023. These couples are arguably more similar to the samples used in other papers in the literature that estimates assortative matching. The results are remarkably similar to new couples, both in the total covariance and in the shares attributable to the different components.

The covariances reported in Table 4, along with estimated variance terms (not shown), imply that the correlation of partners' permanent components is 0.20 for new couples and a remarkably similar 0.19 for stable couples. This is lower than estimates from the literature. Hyslop (2001), for instance, reports a correlation of permanent components of 0.57, suggesting that assortative matching on earnings levels are less pronounced in Norway.

The correlation between partners' monthly components is 0.08, significantly higher than for random couples, implying positive assortative matching also on high-frequency earnings.³⁷

Table 4 suggests that there is marital sorting on earnings levels, but because it uses earnings data both before and after couple formation for the new couples sample, some

³⁷Unfortunately, we cannot identify the correlation of the annual components in practice due to a weak identification issue that arises when annual shocks are small relative to permanent heterogeneity (a common issue in the income dynamics literature, see e.g., Baker and Solon, 2003; Arellano et al., 2017). Note that this identification problem does not affect the estimation of covariances or covariance shares reported in Table 4.

of the positive correlation in earnings could be driven by behavioral responses after couple formation rather than sorting. To examine this further, we plot event studies of the correlations of the annual and monthly components. Because of the weak identification issue discussed above, which would be amplified by smaller samples for event-time specific estimates, we use a residual-based approach instead of the structural three-way decomposition. This approach constructs annual and monthly earnings shocks directly from the data by removing long-run and couple-year means, and allows us to estimate empirical correlations at each event time. These reduced-form shocks mix the structural yearly and monthly components, but they avoid the weak-identification problem and provide event-time profiles of how partners earnings co-move in the short run.³⁸

Figure 5 reports the correlation of annual earnings (Panel a) and the correlation of monthly earnings (Panel b) for true couples (solid darker blue lines) and randomly-matched couples (lighter dashed blue lines). For both components, the correlation for true couples is typically positive, while the correlation for random couples is closer to zero.³⁹ Moreover, both panels also show virtually no change in correlation around the time when couples form, suggesting that these correlations are more likely due to sorting than a behavioral response once cohabiting. This, together with the permanent correlation discussed above, suggests that couples are positively sorted on all three components of earnings *levels*. We next return to the role of sorting on earnings *volatility*.


4.3 A decomposition of marital sorting on volatility and pooling

In a final exercise, we propose a statistical decomposition to quantify the role of assortative matching and mechanical pooling in understanding the difference between individual and household volatility. To do this, we decompose the difference between the volatility of average household earnings and the average volatility of individual

³⁸The residual-based annual shocks equal the structural yearly shock π plus the average monthly shock. As a result, their covariance equals the structural yearly covariance plus a contamination term coming from the covariance of the average of the monthly shocks, typically with the same sign so that the covariance is overestimated. The variances of the residual-based annual shock also include an additional term that mechanically dampens correlations. The residual-based monthly shocks contains no contamination term but may be attenuated. Our pooled estimate is a variance weighted average of the event-time specific estimates.

³⁹The correlation for random partners is slightly positive instead of exactly zero because we draw random partners from the pool of individuals that (truly) matched within the same calendar year the individual started cohabiting with their true partner. As a result, a random partner is somewhat more similar than any random person (e.g., they are typically at a more similar career stage than a random person who started cohabiting in a different year).

Figure 5: Partner correlations of yearly and monthly components of earnings



Notes: Figure plots correlations between partners of the yearly component π (Panel a) and the monthly component θ (Panel b) of earnings from Equation (1) over time relative to the year (for Panel a) or month (for Panel b) of couple formation.

earnings, or ΔV_i for volatility measure V for a household i . To fix ideas, we first describe a simple decomposition of the change in the variance in earnings upon pooling, but next move to using the coefficient of variation as our volatility measure (with derivations and proofs in the appendix). As in the rest of the paper, we focus on *within-unit* variation because the focus of the paper is on volatility as opposed to cross-sectional inequality.⁴⁰

Let monthly earnings of males be y_{M_i} and females y_{F_i} , as above but suppressing the monthly subscript. Denote the variances of their monthly earnings by $\sigma_{M_i}^2$ and $\sigma_{F_i}^2$, respectively, means by μ_{M_i} and μ_{F_i} , and the correlation between them by ρ_i . The difference between the variance of average household earnings and the average variance is then:

$$\Delta \text{Var}_i = \text{Var} \left(\frac{y_{M_i} + y_{F_i}}{2} \right) - \frac{1}{2} [\text{Var}(y_{M_i}) + \text{Var}(y_{F_i})] \quad (3)$$

$$= \underbrace{-0.25 (\sigma_{M_i}^2 + \sigma_{F_i}^2)}_{\text{Pooling effect}} + \underbrace{0.5 \rho_i \sigma_{M_i} \sigma_{F_i}}_{\text{Assortative matching effect}} \quad (4)$$

The first term in this expression is the pooling effect: by pooling earnings, volatility mechanically decreases. The second term is the effect of assortative matching on earnings changes: positively correlated earnings processes between partners can

⁴⁰Hyslop (2001) conducts a similar variance decomposition exercise on household earnings, but also includes cross-sectional variation, which incorporates assortative matching on permanent levels.

dampen volatility decreases from pooling (and in the extreme case where $\rho_i = 1$, can fully undo volatility decreases from pooling), while negatively correlated earnings processes can magnify volatility decreases.

While the change in variance allows for a simple analytical decomposition, it is not one of our preferred measures of volatility. To maintain consistency with the rest of our analyses, we derive an analogous decomposition for the change in the coefficient of variation, ΔCV_i . The difference between the CV of household average earnings and the average CV of individual earnings can be expressed as:

$$\Delta CV_i = \frac{\sqrt{\sigma_{M_i}^2 + \sigma_{F_i}^2 + 2\rho_i\sigma_{M_i}\sigma_{F_i}}}{\mu_{M_i} + \mu_{F_i}} - \frac{1}{2} \left(\frac{\sigma_{M_i}}{\mu_{M_i}} + \frac{\sigma_{F_i}}{\mu_{F_i}} \right) \quad (5)$$

Because this expression is non-linear in the parameters of interest, we rely on a second-order Taylor approximation to further decompose it into interpretable terms (see Appendix C for details). Another implication of this non-linearity is that the decomposition generates many additional terms beyond those in the variance decomposition of Equation (4). These terms can be broadly grouped into (1) a homogeneous benchmark, which captures the change in volatility if all couples had the same (average) parameters $\bar{\mu}_M, \bar{\mu}_F, \bar{\sigma}_M, \bar{\sigma}_F$ and there was no sorting, (i.e., $\rho_i = 0$); (2) heterogeneity in earnings process parameters across individuals $\mu_{M_i}, \mu_{F_i}, \sigma_{M_i}, \sigma_{F_i}$ (evaluated at $\rho_i = 0$), which does not depend on the matching pattern; (3) how the first two terms change when evaluated using the average correlation $\bar{\rho}$; and (4) assortative matching on mean individual earnings, the variance of individual earnings as well as cross-moments that depend on the matching pattern. The first two groups of terms capture pooling effects while the second two groups of terms capture assortative matching effects. In the results that follow we group the terms in (3) and (4) into a single broad category that capture all types of sorting that depend on the matching pattern, but show disaggregated effects in Appendix C.

To quantify the magnitude of each of these terms, we use the *new couples sample*, and estimate the standard deviations and means for both partners and their correlation separately for the two years pre-cohabitation and the two years post-cohabitation. The sample variance-covariance matrix for these parameters across couples is a biased estimate of the true population level variance-covariance matrix due to estimation noise (we have 24 monthly observations pre-cohabitation and 24 monthly observations post-cohabitation for each couple). To account for this, we de-bias the estimates by subtracting the average of the sampling variance (see, e.g., Becker, 2000), which we estimate by bootstrap using 500 repetitions. Combining the


resulting variance-covariance estimates with the means of the parameters allows us to quantify the contribution of each term to the change in volatility upon pooling earnings, as shown in Appendix C.

Figure 6 reports the magnitudes of each decomposition term in $\mathbb{E}(\Delta CV_i)$, where the contribution of each component is indicated by the horizontal lighter blue bars for the pre-cohabitation period and the horizontal darker blue bars for the post-cohabitation period. There are four main takeaways from these results. The first is that the contributions from the various components are relatively stable from before to after cohabitation. If individuals engaged in endogenous responses to partnership formation, for instance through changes in jobs or labor supply that would make earnings less correlated or more stable, we would expect these relationships to change. As this is not the case, we take this as another piece of evidence that added worker effects or endogenous responses to short-term partner earnings shocks are not an important feature in our setting.

Second, our decomposition allows us to understand why the reduction we see when true couples pool earnings is relatively similar to what we would see from random matching (as shown in Figure 4). In Figure 6, we see that the mechanical pooling effect (the homogeneous benchmark) is substantial, and accounts for around 60% of the total decline in volatility. Third, individual heterogeneity – the variance across individuals in earnings and risk as well as within-individual covariances between earnings and risk – on net also accounts for a sizable share of the decline in volatility of couples relative to individuals.

The final takeaway from Figure 6 is that the contribution of assortative matching, as summarized by the final component, has very little effect on household volatility compared to individual volatility. Unlike the pooling terms, these remaining terms capture the influence of the particular way in which couples form. This include assortative matching on levels – i.e., the fact that higher earning males tend to partner with higher earning females – and assortative matching on risk – i.e., the fact that males with more variable earnings tend to partner with females with more variable earnings, as well as cross-moment covariances. These components are illustrated separately in Figure A.11, but while some of them slightly contribute to increasing the benefits of pooling, others decrease it. Taken together, the terms that are affected by the particular way couples match sum to approximately zero, and thus much of the decline in household volatility relative to individual volatility is driven by mechanical pooling rather than assortative matching. This explains why the current matching pattern produces a drop in volatility that is remarkably similar to the

Figure 6: $\mathbb{E}(\Delta CV_i)$ decomposition, before and after couple formation



Notes: Figure reports the estimated components of ΔCV_i for the 24 months prior to couple formation (lighter blue bars) and the 24 months post-couple formation (darker blue bars) for the *new couples subsample*. The homogeneous benchmark is the mechanical pooling effect if all couples had the average earnings processes and partner earnings were uncorrelated. The within-gender heterogeneity bars aggregate the three components that are gender-specific: the variance across individuals of average earnings, the variance across individuals of earnings risk, and the covariance between average earnings and earnings risk, again if partner earnings were uncorrelated. The sorting bars aggregate the components that are specific to the observed matching behavior: how the homogeneous benchmark changes if partner earnings are correlated as in the data, how the within-gender heterogeneity changes if partner earnings are correlated as in the data, assortative matching on earnings levels and on earnings risk, the covariance of male average earnings and female earnings risk and vice versa, the variance of the correlation between male and female earnings, the covariances between ρ and the average earnings of males and females, and the covariances between ρ and the earnings risk of males and females. Variance-covariance estimates are de-biased by subtracting off the average of the within-couple sampling variance-covariance matrix, which is estimated via bootstrap.

pattern suggested by random couples.

5 Conclusion

This paper documents the extent of monthly earnings volatility in Norway at the individual and household level. Using administrative data covering the universe of workers from 2015 to 2023, we show that the bulk of individual earnings variation occurs within years, rather than across them – highlighting the limitations of annual data in capturing the earnings dynamics that households face. Monthly changes in earnings are both frequent and sizable, with volatility further amplified by multiple job-holdings and non-employment spells.

Despite this high degree of individual volatility, household-level earnings are noticeably less volatile. We conduct event study analyses as well as decomposition and bounding exercises that reveal that this reduction arises primarily from mechanical pooling: combining the earnings of two partners smooths earnings streams in a way that random matching replicates almost exactly. In contrast, we find little evidence that marital sorting on volatility or partner labor supply adjustments following job loss play an important role in reducing household volatility relative to individual volatility. Taken together, these results suggest that household-level insurance operates primarily through pooling, rather than through who marries whom or through behavioral responses to shocks.

The absence of net sorting on volatility raises questions for future research: why do individuals match along traits such as education and permanent income, but not along volatility? Understanding this puzzle may shed light on how risk considerations enter household formation decisions. More broadly, our evidence emphasizes that earnings volatility is not just an annual phenomenon but a monthly reality for workers and families, and suggests that recognizing the high-frequency nature of labor income risk is crucial for evaluating both the need for and the effectiveness of social insurance.

References

- Abowd, John M and David Card**, “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 1989, *57* (2), 411–445.
- Altonji, Joseph G. and Ivan Vidangos**, “Modeling Earnings Dynamics,” *Econometrica*, 2013, *81* (4), 1395–1454.
- , **Daniel Giraldo Páez, Disa M. Hynsjo, and Ivan Vidangos**, “Marriage Dynamics, Earnings Dynamics, and Lifetime Family Income,” 2024.
- Andersen, Asger Lau, Amalie Sofie Jensen, Niels Johannesen, Claus Thustrup Kreiner, Søren Leth-Petersen, and Adam Sheridan**, “How Do Households Respond to Job Loss? Lessons from Multiple High-Frequency Datasets,” *American Economic Journal: Applied Economics*, October 2023, *15* (4), 1–29.
- Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme**, “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework,” *Econometrica*, 2017, *85* (3), 693–734. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA13795](https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA13795).
- Autor, David, Andreas Kostøl, Magne Mogstad, and Bradley Setzler**, “Disability benefits, consumption insurance, and household labor supply,” *American Economic Review*, 2019, *109* (7), 2613–2654.
- Baker, Michael and Gary Solon**, “Earnings Dynamics and Inequality among Canadian Men, 1976–1992: Evidence from Longitudinal Income Tax Records,” *Journal of Labor Economics*, April 2003, *21* (2), 289–321. Publisher: The University of Chicago Press.
- Becker, Betsy J.**, “Multivariate meta-analysis,” in “Handbook of applied multivariate statistics and mathematical modeling,” San Diego, CA, US: Academic Press, 2000, pp. 499–525.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, “Consumption inequality and partial insurance,” *American Economic Review*, 2008, *98* (5), 1887–1921.
- , **Michael Graber, and Magne Mogstad**, “Labor income dynamics and the insurance from taxes, transfers, and the family,” *Journal of Public Economics*, July 2015, *127*, 58–73.

- Brewer, Mike, Nye Cominetti, and Stephen P. Jenkins**, “What Do We Know About Income and Earnings Volatility?,” *Review of Income and Wealth*, May 2025, 71 (2), e70013.
- Busch, Christopher, Rocio Madera, and Fane Groes**, “Income Dynamics of Couples: Correlated Risks and Heterogeneous Within-Household Insurance,” 2023.
- Chiappori, Pierre-André, Mónica Costa-Dias, Costas Meghir, and Hanzhe Zhang**, “Changes in Marital Sorting: Theory and Evidence from the US,” *Journal of Political Economy*, 2025, *Forthcoming*.
- Cullen, Julie Berry and Jonathan Gruber**, “Does unemployment insurance crowd out spousal labor supply?,” *Journal of Labor Economics*, 2000, 18 (3), 546–572.
- Druehl, Jeppe, Michael Graber, and Thomas H. Jørgensen**, “High Frequency Income Dynamics,” October 2025.
- Dynan, Karen, Douglas Elmendorf, and Daniel Sichel**, “The Evolution of Household Income Volatility,” *The B.E. Journal of Economic Analysis & Policy*, December 2012, 12 (2).
- Eika, Lasse, Magne Mogstad, and Basit Zafar**, “Educational Assortative Mating and Household Income Inequality,” *Journal of Political Economy*, December 2019, 127 (6), 2795–2835. Publisher: The University of Chicago Press.
- Fadlon, Itzik and Torben Heien Nielsen**, “Family labor supply responses to severe health shocks: Evidence from Danish administrative records,” *American Economic Journal: Applied Economics*, 2021, 13 (3), 1–30.
- Ganong, Peter, Pascal Noel, Christina Patterson, Joseph Vavra, and Alexander Weinberg**, “Earnings Instability,” *NBER Working Paper 34227*, 2025.
- Garin, Andrew, Emilie Jackson, and Dmitri Koustas**, “New gig work or changes in reporting? Understanding self-employment trends in tax data,” *American Economic Journal: Applied Economics*, 2025, 17 (3), 236–270.
- Gennetian, Lisa A, Sharon Wolf, Heather D Hill, and Pamela A Morris**, “Intrayear household income dynamics and adolescent school behavior,” *Demography*, 2015, 52 (2), 455–483.
- Gottschalk, Peter and Robert Moffitt**, “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, 1994, 1994 (2), 217.

- **and** —, “The Rising Instability of U.S. Earnings,” *Journal of Economic Perspectives*, December 2009, *23* (4), 3–24.
- Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos**, “Marry your like: Assortative mating and income inequality,” *American Economic Review*, 2014, *104* (5), 348–353.
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song**, “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?,” *Econometrica*, 2021, *89* (5), 2303–2339.
- , **Luigi Pistaferri, and Giovanni L. Violante**, “Global trends in income inequality and income dynamics: New insights from GRID,” *Quantitative Economics*, 2022, *13* (4), 1321–1360.
- Halla, Martin, Julia Schmieder, and Andrea Weber**, “Job Displacement, Family Dynamics, and Spousal Labor Supply,” *American Economic Journal: Applied Economics*, October 2020, *12* (4), 253–287.
- Halvorsen, Elin, Hans A Holter, Serdar Ozkan, and Kjetil Storesletten**, “Dissecting idiosyncratic earnings risk,” *Journal of the European Economic Association*, 2024, *22* (2), 617–668.
- Hendren, Nathaniel**, “Knowledge of future job loss and implications for unemployment insurance,” *American Economic Review*, 2017, *107* (7), 1778–1823.
- Hryshko, Dmytro, Chinhui Juhn, and Kristin McCue**, “Trends in earnings inequality and earnings instability among US couples: How important is assortative matching?,” *Labour Economics*, 2017, *48*, 168–182.
- Hyslop, Dean R.**, “Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings,” *American Economic Review*, September 2001, *91* (4), 755–777.
- Keys, Benjamin**, “Trends in Income and Consumption Volatility, 1970–2000,” in D. Joliffe and J. Ziliak, eds., *Income Volatility and Food Assistance in the United States*, Upjohn Institute for Employment Research, 2008.
- Kleven, Henrik, Camille Landais, and Gabriel Leite-Mariante**, “The child penalty atlas,” *Review of Economic Studies*, 2024, p. rdae104.

- Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A. Woodbury**, “Wage Posting or Wage Bargaining? A Test Using Dual Jobholders,” *Journal of Labor Economics*, April 2022, *40* (S1), S469–S493. Publisher: The University of Chicago Press.
- Larrimore, Jeff, Alicia Lloro, Zofsha Merchant, Ellen A Merry, Fatimah Shaalan, Julie Siwicki, and Mike Zabek**, “Economic Well-Being of US Households in 2024,” Technical Report, Board of Governors of the Federal Reserve System (US) 2025.
- Lundberg, Shelly**, “The added worker effect,” *Journal of Labor Economics*, 1985, *3* (1, Part 1), 11–37.
- Meghir, Costas and Luigi Pistaferri**, “Income Variance Dynamics and Heterogeneity,” *Econometrica*, January 2004, *72* (1), 1–32.
- and —, “Earnings, consumption and life cycle choices,” in “Handbook of labor economics,” Vol. 4, Elsevier, 2011, pp. 773–854.
- Moffitt, Robert and David C. Ribar**, “Variable Effects of Earnings Volatility on Food Stamp Participation,” in D. Joliffe and J. Ziliak, eds., *Income Volatility and Food Assistance in the United States*, Upjohn Institute for Employment Research, 2008.
- , **John Abowd, Christopher Bollinger, Michael Carr, Charles Hokayem, Kevin McKinney, Emily Wiemers, Sisi Zhang, and James Ziliak**, “Reconciling trends in us male earnings volatility: Results from survey and administrative data,” *Journal of Business & Economic Statistics*, 2022, *41* (1), 1–11.
- Nardi, Mariacristina De, Giulio Fella, and Gonzalo Paz-Pardo**, “Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare,” *Journal of the European Economic Association*, April 2020, *18* (2), 890–926.
- OECD**, “Employment and unemployment by five-year age group and sex - indicators,” 2025.
- Ostrovsky, Yuri**, “The correlation of spouses’ permanent and transitory earnings and family earnings inequality in Canada,” *Labour Economics*, 2012, *19* (5), 756–768.
- Persson, Petra**, “Social insurance and the marriage market,” *Journal of Political Economy*, 2020, *128* (1), 252–300.

- Pruitt, Seth and Nicholas Turner**, “Earnings Risk in the Household: Evidence from Millions of US Tax Returns,” *American Economic Review: Insights*, June 2020, *2* (2), 237–254.
- Schneider, Daniel and Kristen Harknett**, “Consequences of routine work-schedule instability for worker health and well-being,” *American sociological review*, 2019, *84* (1), 82–114.
- Shin, Donggyun and Gary Solon**, “Trends in men’s earnings volatility: What does the Panel Study of Income Dynamics show?,” *Journal of Public Economics*, August 2011, *95* (7-8), 973–982.
- Shiu, Ji-Liang, Sisi Zhang, and Peter Gottschalk**, “Family Income Dynamics 1970–2018: Putting the Pieces Together,” *Journal of Labor Economics*, 2025, *43* (S1), S123–S151.
- Shore, Stephen H**, “The co-movement of couples incomes,” *Review of Economics of the Household*, 2015, *13* (3), 569–588.
- Stephens, Melvin Jr.**, “Worker Displacement and the Added Worker Effect,” *Journal of Labor Economics*, July 2002, *20* (3), 504–537.

Appendices

A Volatility measures for different earnings processes

A.1 Volatility measures and illustrative earnings processes


To illustrate how different volatility measures deal with different types of earnings processes we compare four exemplary earnings processes y_1 , y_2 , y_3 , and y_4 , which are plotted in Figure A.1. All four earnings paths have the same mean over time $1 \leq t \leq 20$. For y_1 , earnings increase once and remain at the higher level afterwards. For y_2 , earnings increase and immediately decrease again by the same amount from period to period. For y_3 , earnings increase and decrease by a twice as large amount but only half as often as for y_2 . We define y_3 such that the mean absolute arc-percent change is the same as for y_2 . Lastly, for y_4 , earnings grow at a constant growth rate every period.

Table A.1 reports point estimates for our main measures of volatility for y_1 to y_4 . The table shows that which earnings process is more or less volatile varies depending on the volatility measure. For the coefficient of variation, we see that $CV(y_3) > CV(y_4) > CV(y_1) = CV(y_2)$. Two properties of the CV explain this order. First, deviations from mean earnings enter the variance quadratically, which is why $CV(y)$ put larger weights on larger changes. Having many small changes (y_2) yields a smaller $CV(y)$ compared fewer but larger changes (y_3). As a result, $CV(y_3) > CV(y_2)$. Second, the order is irrelevant for the CV. One permanent change is the same as many transitory changes, as long as there is the same number of periods for each earnings level. As a result, $CV(y_1) = CV(y_2)$.

Let $\bar{a}(y) \equiv \frac{1}{T} \sum_{t=1}^T |a_t(y)|$ denote the mean absolute arc-percent change of income process y . We see that $\bar{a}(y_2) = \bar{a}(y_3) > \bar{a}(y_4) > \bar{a}(y_1)$. Two properties of arc percentage changes explain this order. First, it is a linear measure. Larger changes (y_3) do not receive more weight than smaller changes (y_2). Because of that, we can construct y_3 such that the mean absolute arc-percent change satisfies $\bar{a}(y_2) = \bar{a}(y_3)$, with half as many but twice as large increases. Second, the order is crucial. One permanent increase (y_1) implies $\bar{a}(y) = 0$ for all t except one, whereas alternating increases and decreases (y_2) imply $\bar{a}(y) > 0$ for all t . As a result, $\bar{a}(y_2) > \bar{a}(y_1)$. Note that this is only true for the *absolute* value of arc-percent changes, since otherwise increases and decreases of the same relative size cancel out.

Our third measure of volatility is the standard deviation of arc percentage changes,

Appendix Figure A.1: Illustrative earnings processes



Notes: Figure shows four illustrative earnings processes: y_1 , y_2 , y_3 , and y_4 , which all have the same mean. We report our three volatility measures for each process in Table A.1.

Appendix Table A.1: Selected volatility measures for illustrative earnings processes

	y_t	CV	$ a_t $	$SD(a_t)$
	(1)	(2)	(3)	(4)
y_1	100.00	0.21	0.02	0.09
y_2	100.00	0.21	0.40	0.41
y_3	100.00	0.39	0.40	0.57
y_4	100.00	0.29	0.05	0.00

Notes: Table reports mean values for (1) earnings, (2) the coefficient of variation of earnings, (3) the absolute value of the arc percentage change, and (4) the standard deviation of the arc percentage change for the four earnings processes y_1 to y_4 that are plotted in Figure A.1.

which measures the volatility of *earnings changes*. Here, $SD(a_t(y_3)) > SD(a_t(y_2)) > SD(a_t(y_1)) > SD(a_t(y_4)) = 0$. A key difference to $CV(y)$ and $|a_t(y)|$, which measure volatility of *earnings*, is that $SD(a_t) = 0$ for constant earnings growth rates (y_4), whereas $CV(y)$ and $|a_t(y)|$ are always > 0 if earnings are not the same across all periods.

A.2 Volatility measures and log earnings processes

In this section, we simulate a simple log earnings process in order to show how the parameters of that process affect the volatility measures used in the paper. The process has individual heterogeneous earnings profiles (Güvenen, 2009) and both permanent and transitory shocks.

Earnings for individual i are assumed to follow the process

$$\log y_{i,t} = t\beta_i + u_{i,t} + \epsilon_{i,t} \quad (6)$$

where $y_{i,t}$ is the earnings of individual i at time t , β_i is an individual-specific time trend, $u_{i,t}$ is a permanent component, and $\epsilon_{i,t}$ is a transitory shock. The permanent component follows a random walk according to:

$$u_{i,t} = u_{i,t-1} + \eta_{i,t} \quad (7)$$

The shocks, the time trends, and the initial conditions are all assumed to be normally distributed with variances σ_η^2 , σ_ϵ^2 , σ_β^2 , and $\sigma_{u_0}^2$ respectively. Each of these may also be correlated within couples with correlations ρ_{u_0} , ρ_β , ρ_η , and ρ_ϵ respectively, but are assumed to be independent across couples and across time. All individuals have a partner.

Appendix Table A.2: Baseline simulation parameters


Parameter	Description	Value
σ_{u_0}	Initial standard deviation of the permanent component	$\sqrt{0.3}$
σ_η	Standard deviation of the permanent shock	$\sqrt{0.125}$
σ_ϵ	Standard deviation of the transitory shock	$\sqrt{0.06}$
σ_β	Standard deviation of the time trend	$\sqrt{0.01}$
ρ_{u_0}	Correlation within couple of initial permanent components	0.4
ρ_η	Correlation within couple of permanent shocks	0.15
ρ_ϵ	Correlation within couple of transitory shocks	0.2
ρ_β	Correlation within couple of time trends	0.4

The baseline parameters used for the simulation are shown in Appendix Table A.2. We simulate the process for 5000 couples over 12 periods and calculate our three volatility measures at the individual and household level. For each of the eight parameters, we vary the value and re-run the simulation holding other parameters constant at the baseline values. In order to facilitate comparison across the different volatility measures, we standardize each measure to have a value of 1 at the baseline parameter values by dividing by the value of the measure at those parameter values.


The results are shown in Appendix Figure A.2. All three measures respond in the same direction to changes in each of the σ parameters. However, the within-couple $SD_i(a_{i,m})$ is nearly unresponsive to changes in the variance of time trends. The CV reflects deviations from the unit-specific mean, so it is especially sensitive to components that push earnings away from that mean in a systematic way, such as larger variance and higher correlation in the trend terms $(\sigma_\beta, \rho_\beta)$, which steepen earnings profiles and thereby raise the dispersion of levels over time. By contrast, $|a_{i,m}|$ and $SD_i(a_{i,m})$ are based on month-to-month arc percentage changes, so they respond primarily to high-frequency innovations $(\varepsilon_{i,t})$, or their correlation across spouses which acts to amplify them.

Panel (e) illustrates that increasing ρ_β raises the CV of household earnings but slightly lowers $SD_i(a_{i,m})$. A higher correlation between partners' trends makes household earnings levels drift further from their average, boosting dispersion in levels (and thus the CV), while at the same time smoothing the household time series so that transitory shocks are a smaller share of month-to-month changes, which reduces $SD_i(a_{i,m})$. The three measures respond in the same direction to changes in the correlation between partners of u_0 , η , and ϵ .


Appendix Figure A.2: Effect of Earnings Process Parameters on Volatility Measures




(a) σ_β




(b) σ_ϵ




(c) σ_η




(d) σ_{u_0}




(e) ρ_β



(f) ρ_ϵ



(g) ρ_η



(h) ρ_{u_0}

B Proof: Including zero-earning periods increases CV

Lemma 1. *The coefficient of variation of earnings that includes periods of zero-earnings is weakly larger than the coefficient of variation of earnings that excludes periods of zero-earnings, under a model in which earnings $y_{i,m}$ of unit i in month m are equal to zero with probability π and otherwise distributed according to $y_{i,m}^*$ with mean μ and variance σ^2 .*

Proof.

$$\mathbb{E}(y_{i,m}) = (1 - \pi)\mu$$

and

$$\begin{aligned}\text{Var}(y_{i,m}) &= \mathbb{E}(y_{i,m}^2) - (\mathbb{E}(y_{i,m}))^2 \\ &= (1 - \pi)(\mu^2 + \sigma^2) - (1 - \pi)^2\mu^2\end{aligned}$$

because $\mathbb{E}(y_{i,m}^2) = (1 - \pi)\mathbb{E}((y_{i,m}^*)^2) = (1 - \pi)(\mu^2 + \sigma^2)$, so

$$CV(y_{i,m}) = \frac{\sqrt{(1 - \pi)(\mu^2 + \sigma^2) - (1 - \pi)^2\mu^2}}{(1 - \pi)\mu}$$

This is weakly greater than $CV(y_{i,m}^*) = \frac{\sigma}{\mu}$ because:

$$\begin{aligned}\frac{\sqrt{(1 - \pi)(\mu^2 + \sigma^2) - (1 - \pi)^2\mu^2}}{(1 - \pi)\mu} &\stackrel{?}{\geq} \frac{\sigma}{\mu} \\ (1 - \pi)(\mu^2 + \sigma^2) - (1 - \pi)^2\mu^2 &\stackrel{?}{\geq} \sigma^2(1 - \pi)^2 \\ (\mu^2 + \sigma^2) - (1 - \pi)\mu^2 &\stackrel{?}{\geq} \sigma^2(1 - \pi) \\ (\mu^2 + \sigma^2) &\stackrel{?}{\geq} (\mu^2 + \sigma^2)(1 - \pi) \\ 1 &\geq (1 - \pi)\end{aligned}$$

where $0 \leq \pi \leq 1$. □

C A framework for understanding volatility changes upon pooling

In the first part of this appendix, we provide simple proofs for why the standard deviation of average earnings are weakly smaller than the average of the standard deviations of average earnings, and a condition for when the standard deviation of average earnings is smaller than the standard deviation of the lower-volatility spouse.

In the second part of this appendix, we supplement the discussion in Section 4.3 by providing details of a framework for understanding the sources behind the difference between individual volatility and household volatility, focusing on the coefficient of variation (CV).

C.1 Standard deviation

Lemma 2. *The standard deviation of average earnings is weakly lower than the average standard deviation of individual earnings.*

Proof.

$$\begin{aligned} \frac{1}{2} \sqrt{\sigma_{M_i}^2 + \sigma_{F_i}^2 + 2\rho\sigma_{M_i}\sigma_{F_i}} &\stackrel{?}{\leq} \frac{1}{2} (\sigma_{M_i} + \sigma_{F_i}) \\ \sigma_{M_i}^2 + \sigma_{F_i}^2 + 2\rho\sigma_{M_i}\sigma_{F_i} &\stackrel{?}{\leq} (\sigma_{M_i} + \sigma_{F_i})^2 \\ 2\sigma_{M_i}\sigma_{F_i}(\rho - 1) &\leq 0 \end{aligned}$$

because $\rho \leq 1$. □


Lemma 3. *The standard deviation of average earnings is not necessarily lower than the standard deviation of the low-SD partner:*

Proof.

$$\frac{1}{2} \sqrt{\sigma_{H(i)}^2 + \sigma_{L(i)}^2 + 2\rho\sigma_{H(i)}\sigma_{L(i)}} - \sigma_{L(i)}$$

where $H(i)$ and $L(i)$ denote the high- and low-SD partner in couple i . This implies that the standard deviation of average earnings is only less than the standard deviation of the low-SD partner if $\rho < \frac{3k^2-1}{2k}$, where $k = \frac{\sigma_{L(i)}}{\sigma_{H(i)}} \in [0, 1]$. This condition is illustrated in Figure A.3. □

Appendix Figure A.3: Parameters when household SD equals low-SD spouse



C.2 Coefficient of variation

Let there be a population households indexed by i , each consisting of women F_i and man M_i with monthly earnings processes y_{F_i} and y_{M_i} (with subscript m suppressed), characterized by means μ_{F_i} and μ_{M_i} and standard deviations σ_{F_i} and σ_{M_i} (all positive and finite). Let the correlation between the earnings of man and woman in couple i be ρ_i . The CV of the mean earnings of couple i is then:

$$CV_i\left(\frac{y_{F_i} + y_{M_i}}{2}\right) = \frac{\sqrt{\sigma_{F_i}^2 + \sigma_{M_i}^2 + 2\rho_i\sigma_{F_i}\sigma_{M_i}}}{\mu_{F_i} + \mu_{M_i}}.$$

Let the relative means and standard deviations in couple i be denoted by $m_i = \frac{\mu_{F_i}}{\mu_{M_i}}$ and $k_i = \frac{\sigma_{F_i}}{\sigma_{M_i}}$. Then the change in volatility upon pooling is:

$$\begin{aligned} \Delta CV_i &= \frac{\sqrt{\sigma_{M_i}^2 + \sigma_{F_i}^2 + 2\rho_i\sigma_{M_i}\sigma_{F_i}}}{\mu_{M_i} + \mu_{F_i}} - \frac{1}{2} \left(\frac{\sigma_{M_i}}{\mu_{M_i}} + \frac{\sigma_{F_i}}{\mu_{F_i}} \right) \\ &= \frac{\sigma_{M_i}}{\mu_{M_i}} \left[\frac{\sqrt{1 + k_i^2 + 2\rho_i k_i}}{1 + m_i} - \frac{1}{2} \left(1 + \frac{k_i}{m_i} \right) \right] \end{aligned}$$

Solving $\Delta CV_i < 0$ yields a condition for when a couple will see reduced volatility

upon pooling:


$$\rho_i < \frac{\frac{(1+m_i)^2}{4} \left(1 + \frac{k_i}{m_i}\right)^2 - (1+k_i^2)}{2k_i} := \rho_i^* \quad (8)$$

This is always satisfied in the special cases of $m_i = 1$ (equal means) or $m_i = k_i$ (equal CVs). One might also be interested in when pooling reduces volatility compared to one or the other's partner's volatility, for which we have:

$$\begin{aligned} \Delta CV_{M_i} &= \frac{\sqrt{\sigma_{M_i}^2 + \sigma_{F_i}^2 + 2\rho_i}}{\mu_{M_i} + \mu_{F_i}} - \frac{\sigma_{M_i}}{\mu_{M_i}} \\ \Delta CV_{M_i} < 0 &\Rightarrow \rho_i < \frac{m_i^2 + 2m_i - 3 - 4k_i^2}{8k_i} := \rho_i^M \\ \Delta CV_{F_i} &= \frac{\sqrt{\sigma_{M_i}^2 + \sigma_{F_i}^2 + 2\rho_i}}{\mu_{M_i} + \mu_{F_i}} - \frac{\sigma_{F_i}}{\mu_{F_i}} \\ \Delta CV_{F_i} < 0 &\Rightarrow \rho_i < \frac{k_i^2(1 + 2m_i - 3m_i^2) - 4m_i^2}{8k_i m_i^2} := \rho_i^F \end{aligned}$$

These conditions are plotted in Appendix Figure A.4.

Appendix Figure A.4: Parameters when household volatility equals individual or mean volatility



Next, we are interested in the population level expectations of this change:

$$\mathbb{E}(\Delta CV_i) = \mathbb{E} \left[\frac{\sqrt{\sigma_{M_i}^2 + \sigma_{F_i}^2 + 2\rho_i \sigma_{M_i} \sigma_{F_i}}}{\mu_{M_i} + \mu_{F_i}} - \frac{1}{2} \left(\frac{\sigma_{M_i}}{\mu_{M_i}} + \frac{\sigma_{F_i}}{\mu_{F_i}} \right) \right]$$

Let bars over parameters denote means over the population. Let $\theta = (\mu_M, \mu_F, \sigma_M, \sigma_F, \rho)$ be the collection of random parameters in the population of couples, and $g(\theta) = \frac{\sqrt{\sigma_M^2 + \sigma_F^2 + 2\rho\sigma_F\sigma_M}}{\mu_M + \mu_F} - \frac{1}{2} \left[\frac{\sigma_M}{\mu_M} + \frac{\sigma_F}{\mu_F} \right]$ be the function that computes the difference between the CV of the mean and the mean of the CV's. Doing a second order Taylor approximation around the means of the parameters yields

$$g(\theta) \approx g(\bar{\theta}) + q(\bar{\theta})'(\theta - \bar{\theta}) + \frac{1}{2}(\theta - \bar{\theta})'H(\bar{\theta})(\theta - \bar{\theta})$$

where q is the vector of first derivatives of g and H the Hessian of second derivatives. Taking expectations, the second term vanishes because $\mathbb{E}(\theta - \bar{\theta}) = 0$. The quadratic term is the covariances matrix of the parameters, and we are left with

$$\begin{aligned} \mathbb{E}(g(\theta)) &\approx g(\bar{\theta}) + \frac{1}{2}\text{tr}(H(\bar{\theta})\text{Cov}(\theta)) \\ &= g(\bar{\theta}) + \frac{1}{2} \sum_k \sum_l H_{kl}(\bar{\theta}) \text{Cov}(\theta_k, \theta_l) \end{aligned}$$

where $\text{tr}(A) = \sum_i A_{ii}$ is the trace function and the second line just writes it out in component-wise form. The effect of income pooling thus depend crucially on the population level variance-covariance matrix of the five parameters across couples.

Writing this out explicitly, using the expression for g , yields:

$$\begin{aligned}
\mathbb{E}(\Delta_i) \approx & \underbrace{\frac{\sqrt{\bar{\sigma}_M^2 + \bar{\sigma}_F^2 + 2\bar{\rho}\bar{\sigma}_F\bar{\sigma}_M}}{\bar{\mu}_M + \bar{\mu}_F} - \frac{1}{2}\left(\frac{\bar{\sigma}_M}{\bar{\mu}_M} + \frac{\bar{\sigma}_F}{\bar{\mu}_F}\right)}_{\text{homogeneous benchmark}} \\
& + \underbrace{\frac{1}{2}(c_{11}\text{Var}(\mu_M) + c_{22}\text{Var}(\mu_F))}_{\text{variance in earnings levels}} + \underbrace{\frac{1}{2}(c_{33}\text{Var}(\sigma_M) + c_{44}\text{Var}(\sigma_F))}_{\text{risk dispersion}} \left. \vphantom{\frac{1}{2}(c_{33}\text{Var}(\sigma_M) + c_{44}\text{Var}(\sigma_F))} \right\} \begin{array}{l} \text{additional} \\ \text{terms} \\ \text{for random} \\ \text{couples} \\ \text{benchmark} \end{array} \\
& + \underbrace{c_{13}\text{Cov}(\mu_M, \sigma_M) + c_{24}\text{Cov}(\mu_F, \sigma_F)}_{\text{within-partner correlation of risk and levels}} \\
& + \underbrace{c_{12}\text{Cov}(\mu_M, \mu_F)}_{\text{assortative matching on earnings levels}} + \underbrace{c_{34}\text{Cov}(\sigma_M, \sigma_F)}_{\text{assortative matching on risk}} \left. \vphantom{c_{34}\text{Cov}(\sigma_M, \sigma_F)} \right\} \begin{array}{l} \text{sorting} \\ \text{components} \end{array} \\
& + \underbrace{c_{14}\text{Cov}(\mu_M, \sigma_F) + c_{23}\text{Cov}(\mu_F, \sigma_M)}_{\text{cross-partner terms between levels and risk}} + \underbrace{\frac{1}{2}c_{55}\text{Var}(\rho)}_{\text{Correlation dispersion}} \\
& + \underbrace{c_{15}\text{Cov}(\mu_M, \rho) + c_{25}\text{Cov}(\mu_F, \rho) + c_{35}\text{Cov}(\sigma_M, \rho) + c_{45}\text{Cov}(\sigma_F, \rho)}_{\text{cross-moment terms with } \rho}
\end{aligned}$$

, where $S = \sqrt{\bar{\sigma}_F^2 + \bar{\sigma}_M^2 + 2\bar{\rho}\bar{\sigma}_F\bar{\sigma}_M}$, $D = \bar{\mu}_F + \bar{\mu}_M$, $c_{xy} = \frac{\partial^2 g}{\partial x \partial y}$

$$\begin{aligned}
c_{11} = c_{22} &= \frac{2S}{D^3} - \frac{\sigma_g}{\mu_g^3} \begin{array}{l} \leq 0 \\ \geq 0 \end{array} \\
c_{12} &= \frac{2S}{D^3} > 0, \quad c_{33} = \frac{\bar{\sigma}_F^2(1 - \bar{\rho}^2)}{D S^3} > 0, \quad c_{55} = -\frac{(\bar{\sigma}_M\bar{\sigma}_F)^2}{D S^3} < 0 \\
c_{44} &= \frac{\bar{\sigma}_M^2(1 - \bar{\rho}^2)}{D S^3} > 0, \quad c_{34} = -\frac{\bar{\sigma}_M\bar{\sigma}_F(1 - \bar{\rho}^2)}{D S^3} < 0, \\
c_{13} = c_{24} &= -\frac{1}{2\bar{\mu}_g^2} - \frac{\bar{\sigma}_g + \bar{\rho}\bar{\sigma}_{-g}}{D^2 S} < 0 \\
c_{14} &= -\frac{\bar{\sigma}_F + \bar{\rho}\bar{\sigma}_M}{D^2 S} < 0, \quad c_{23} = -\frac{\bar{\sigma}_M + \bar{\rho}\bar{\sigma}_F}{D^2 S} < 0 \\
c_{15} = c_{25} &= -\frac{\bar{\sigma}_M\bar{\sigma}_F}{D^2 S} < 0 \\
c_{35} &= \frac{\bar{\sigma}_F^2}{D S^3}(\bar{\sigma}_F + \bar{\rho}\bar{\sigma}_M) > 0, \quad c_{45} = \frac{\bar{\sigma}_M^2}{D S^3}(\bar{\sigma}_M + \bar{\rho}\bar{\sigma}_F) > 0
\end{aligned}$$


where the sign of the cross term coefficients requires that the mean within couple correlation $\bar{\rho}$ is not too negative: $\bar{\sigma}_F + \bar{\rho}\bar{\sigma}_M > 0$ and $\bar{\sigma}_M + \bar{\rho}\bar{\sigma}_F > 0$.

For interpretation purposes, it makes sense to further break up the first compo-

nents (the homogeneous benchmark and the within-gender variances) into components evaluated at $\bar{\rho} = 0$ and the difference between the effect at the true $\bar{\rho}$ and this benchmark. Because the homogeneous benchmark and the within-gender (co)variances are independent of the matching pattern, these terms only depend on the matching pattern through the average correlation $\bar{\rho}$. To separate the terms that depend on the matching pattern from those that do not, we evaluate the first four bracketed terms separately for $\rho = 0$ and $\rho = \bar{\rho}$. In practice, we separate these into $c_{xy}(0)\text{Cov}(x, y)$ and $[c_{xy}(\bar{\rho}) - c_{xy}(0)]\text{Cov}(x, y)$ to evaluate them separately.

D Additional Figures


Appendix Figure A.5: CDFs of month-to-month changes in earnings




Notes: Figure plots cumulative distribution functions of arc percentage changes in monthly earnings ($a_{i,m}$). Panel (a) plots earnings changes within a job separately for salaried (solid line) and hourly (dashed line) job spells. Since some jobs can have both salaried earnings and hourly earnings, we assign a job as salaried in month m if salaried earnings are greater than hourly earnings, and vice versa. Panel (b) plots earnings changes within a job (dashed line) and aggregated earnings over all jobs (solid line), excluding months with zero earnings. Panel (c) plots earnings changes aggregated over all jobs, including zero-earnings months (solid line) and excluding zero-earnings months (dashed line). Panel (d) plots individual earnings changes across all jobs and including zero-earnings months (solid line) and mean household earnings changes across all jobs and including zero-earnings months (dashed line), both for the *all couples subsample*.

Appendix Figure A.6: Distributions of month-to-month changes in earnings, other breakdowns

(a) Individuals: by job vs across all jobs, excl. zero-earning months




(b) Individuals by job: salaried vs hourly jobs




Notes: Figure plots binned probability density functions of arc percentage changes in monthly earnings ($a_{i,m}$). Panel (a) plots earnings changes within a job (lighter blue bars) and aggregating earnings over all jobs (darker blue bars), excluding months with zero earnings. Panel (b) plots earnings changes within a job separately for salaried (darker orange bars) and hourly jobs (lighter orange bars). Since some jobs can have both salaried earnings and hourly earnings, we assign a job as salaried in month m if salaried earnings are greater than hourly earnings, and vice versa.

Appendix Figure A.7: Labor market outcomes of individuals after job loss (full sample)

(a) Employment and unemployment insurance benefits (full sample)



(b) Earnings (full sample)




Notes: Figure plots the means and 95% confidence intervals of monthly labor market outcomes for individuals over time for all individuals who experience an involuntary job loss. This is similar to Figure 3, but for the full sample of individuals who lose their job, including those without a partner. Panel (a) plots the share of individuals in month m that are employed (mid-blue), receive unemployment insurance benefits (light blue), or either (darker blue). The share of individuals employed is 100% for $-12 \leq m \leq -1$ by construction. Panel (b) plots monthly earnings relative to $m = -1$. Vertical red line is the month they lost their job involuntarily ($m = 0$).

Appendix Figure A.8: Labor market outcomes of individuals and partners after job loss, subsample analysis




Notes: Figure plots the means (relative to $m = -1$) and 95% confidence intervals of monthly earnings for individuals (solid lines) and their partners (dashed lines) over time using the *job loss subsample*. The sample is further split into individuals who find a new job and/or receive unemployment insurance benefits within 12 months of job loss (red) and those who do not (orange). Vertical red line is the month they lost their job involuntarily ($m = 0$).

Appendix Figure A.9: Labor market outcomes of males and their female partners after male job loss




Notes: Figure plots the means and 95% confidence intervals of monthly labor market outcomes for males who lose their job and their female partners over time using the *job loss subsample*, restricted to males who lose their job. Panel (a) plots the share of males in month m that are employed (mid-blue), receive unemployment insurance benefits (light blue), or either (darker blue), as well as the share of their female partners who are employed (pink). The share if males employed is 100% for $-12 \leq m \leq -1$ by construction. Panel (b) plots monthly earnings for the males who experience a job loss (dark red) and their female partners (light red) relative to $m = -1$. Vertical red line is the month they lost their job involuntarily ($m = 0$).

Appendix Figure A.10: Female and male volatility around couple formation



Notes: Figure plots yearly means of our three main volatility measures at the individual level for the *new couples subsample* over time relative to the event of cohabitation in year 0, separately for males (hollow circles) and females (solid circles). Panel (a) plots the within-unit coefficient of variation, panel (b) plots the mean absolute arc percentage change in monthly earnings, and panel (c) plots the within-unit standard deviation of the arc percentage change in monthly earnings. We compute each measure separately for each event year using within-year monthly earnings variation.

Appendix Figure A.11: $\mathbb{E}(\Delta CV_i)$ decomposition, before and after couple formation, expanded



Notes: Figure reports the estimated components of ΔCV_i for the 24 months prior to couple formation (lighter blue bars) and the 24 months post-couple formation (darker blue bars) for the *new couples subsample*. The homogeneous benchmark is the mechanical pooling effect if all couples had the average earnings processes and partner earnings were uncorrelated. The within-gender heterogeneity bars include the three components that are gender-specific: the variance across individuals of average earnings, the variance across individuals of earnings risk, and the covariance between average earnings and earnings risk, again if partner earnings were uncorrelated. The sorting bars include the components that are specific to the observed matching behavior: how the homogeneous benchmark changes if partner earnings are correlated as in the data, how the within-gender heterogeneity changes if partner earnings are correlated as in the data, assortative matching on earnings levels and on earnings risk, the covariance of male average earnings and female earnings risk and vice versa, the variance of the correlation between male and female earnings, the covariances between ρ and the average earnings of males and females, and the covariances between ρ and the earnings risk of males and females. Variance-covariance estimates are de-biased by subtracting off the average of the within-couple sampling variance-covariance matrix, which is estimated via bootstrap.

E Additional Tables

Appendix Table A.3: Sample sizes

Sample	N individuals	Share
Alive & working age full sample period		
Total wage income > 0	2,523,126	1.00
Norwegian resident full sample period	2,000,762	0.79
Never missing employer information	1,994,547	0.79
Never negative wage income	1,709,871	0.68
Never income from self-employment	1,278,785	0.51
Trim top / bottom 1% total earnings	1,243,332	0.49
Drop if partner has been dropped	767,219	0.30

Notes: Table shows how the sample size changes following the sample restrictions applied to the data.

Appendix Table A.4: Summary statistics: arc percentage changes

	Individual-job Excl. zeros			Individual Incl. zeros		Household Incl. zeros	
	All	Salaried	Hourly	All	New Couples	All	New Couples
	(1a)	(1b)	(1c)	(2a)	(2b)	(3a)	(3b)
Share $a_{i,m} \neq 0$	0.71	0.66	0.97	0.65	0.66	0.76	0.86
Share $a_{i,m} > 0$	0.36	0.34	0.50	0.33	0.34	0.39	0.44
Share $a_{i,m} < 0$	0.35	0.32	0.47	0.32	0.32	0.37	0.41
Distribution of arc-percentage changes							
Mean $a_{i,m}$	0.01	0.00	0.01	0.01	0.01	0.00	0.01
P25 $a_{i,m}$	-0.06	-0.04	-0.24	-0.05	-0.05	-0.07	-0.08
P50 $a_{i,m}$	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P75 $a_{i,m}$	0.08	0.05	0.29	0.06	0.07	0.08	0.10
Distribution of absolute arc-percentage changes							
Mean $ a_{i,m} $	0.19	0.14	0.40	0.23	0.24	0.22	0.21
P25 $ a_{i,m} $	0.00	0.00	0.09	0.00	0.00	0.00	0.02
P50 $ a_{i,m} $	0.07	0.04	0.26	0.05	0.06	0.07	0.09
P75 $ a_{i,m} $	0.23	0.17	0.62	0.22	0.23	0.24	0.26
N individuals	766,780	643,480	315,841	767,219	13,092	767,219	13,092

Notes: Table reports estimates for the distribution of arc percentage changes in monthly earnings across the population in our sample of individuals (columns 1a, 1b, 1c, 2a, and 3a) and the *new couples subsample* (columns 2b and 3b). For columns 1a – 1c, monthly earnings are measured at the individual-by-job level so if an individual has several employers in one month, they enter several times in that month. Column 1a includes all jobs, column 1b includes jobs with predominantly salary earnings, and column 1c includes jobs with predominantly hourly earnings. For (2a)– (3b), earnings are aggregated across jobs including zero-earnings periods. Finally, (3a) and (3b) report the measures for average household earnings, while all other columns report individual earnings. (2a) and (3a) include single and couple households, while (2b) and (3b) only includes the subsample of new couples.

Appendix Table A.5: Alternative monthly volatility measures

	Medians			Cross-sectional	
	$CV_i(y_{i,m})$ (1)	$ a_{i,m} $ (2)	$SD_i(a_{i,m})$ (3)	$CV(y_{i,m})$ (4)	$SD(a_{i,m})$ (5)
<i>(A) Individual-job</i>					
Excluding zeros	0.239	0.067	0.240	0.631	0.333
Excluding zeros, residualized	0.221	0.075	0.229	0.626	0.332
<i>(B) Individual</i>					
Excluding zeros	0.275	0.074	0.266	0.580	0.338
Including zeros	0.371	0.052	0.388	0.746	0.495
Including zeros, new couples	0.316	0.057	0.401	0.694	0.512
<i>(C) Household</i>					
Including zeros	0.316	0.074	0.320	0.639	0.441
Including zeros, new couples	0.265	0.092	0.284	0.524	0.385

Notes: Table reports analogous estimates to Table 3, but medians instead of means in columns (1)–(3) and the cross-sectional population-level $CV(y_{i,m})$ and $SD(a_{i,m})$ without aggregating to the individual-level in columns (4) and (5), respectively.

Appendix Table A.6: Variation explained by residualizing earnings

Tenure at firm	Job type	(1)	(2)
		$y_{i,m}$	$\Delta y_{i,m}$
All	All	0.2041	0.2488
	Salary	0.1975	0.2388
	Hourly	0.3624	0.4336
≥ 12 months	All	0.2093	0.2542
	Salary	0.2006	0.2425
	Hourly	0.3765	0.4450

Notes: Table reports the R^2 from regressing earnings on firm-by-month-of-the-year and individual-firm fixed effects (column 1) and regressing earnings changes on firm-by-month-of-the-year fixed effects (column 2) for different samples. Row 1 represents the full sample, with no restrictions on tenure of individual i with firm f , and for all job types. Row 2 restricts to months with salaried earnings, and row 3 restricts to months with hourly earnings. We also report the R^2 for the subsample of individuals who work at least 12 months for f during the sample period in rows 4 – 6.