

DISCUSSION PAPER SERIES

DP15395

DISSECTING IDIOSYNCRATIC EARNINGS RISK

Elin Halvorsen, Hans Holter, Serdar Ozkan and Kjetil
Storesletten

MACROECONOMICS AND GROWTH



DISSECTING IDIOSYNCRATIC EARNINGS RISK

Elin Halvorsen, Hans Holter, Serdar Ozkan and Kjetil Storesletten

Discussion Paper DP15395
Published 21 October 2020
Submitted 20 October 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Macroeconomics and Growth

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Elin Halvorsen, Hans Holter, Serdar Ozkan and Kjetil Storesletten

DISSECTING IDIOSYNCRATIC EARNINGS RISK

Abstract

This paper examines whether nonlinear and non-Gaussian features of earnings dynamics are caused by hours or hourly wages. Our findings from the Norwegian administrative and survey data are as follows: (i) Nonlinear mean reversion in earnings is driven by the dynamics of hours worked rather than wages since wage dynamics are close to linear while negative changes to hours are transitory and positive changes are persistent. (ii) Large earnings changes are driven equally by hours and wages, whereas small changes are associated mainly with wage shocks. (iii) Both wages and hours contribute to negative skewness and high kurtosis for earnings changes, although hour-wage interactions are quantitatively more important. (iv) When considering household earnings and disposable household income, the deviations from normality are mitigated relative to individual labor earnings: changes in disposable household income are close to symmetric and less leptokurtic.

JEL Classification: E24, H24, J24, J31

Keywords: Earnings Dynamics, Income Shocks, Insurance, wages, Hours, Higher-order earnings risk, Skewness, Kurtosis

Elin Halvorsen - elin.halvorsen@ssb.no
Statistics Norway

Hans Holter - hans.holter@econ.uio.no
University of Oslo

Serdar Ozkan - serdar.ozkan@utoronto.ca
University of Toronto

Kjetil Storesletten - kjetil.storesletten@econ.uio.no
University of Oslo and CEPR

Acknowledgements

We thank seminar participants at the 2017 NBER Summer Institute; Working Group on Consumption—Micro to Macro (Cambridge), the 2017 Meeting of the Society for Economic Dynamics (Edinburgh), the 2018 North American Winter Meeting of the Econometric Society (Philadelphia), Barcelona GSE, 2018 Annual Research Conference at Bank of Spain, Institute for Fiscal Studies (London), Toulouse School of Economics, University of Cologne, University of Bonn, University of Southampton, UNC Charlotte, Universitat Autònoma Barcelona, Insper, Fundação Getúlio Vargas, University of Florida, the Federal Reserve Bank of Cleveland, University of Essex, Center for Retirement Research at Boston College, Statistics Norway, and the University of Oslo for helpful comments and suggestions. We are grateful for support from the Research Council of Norway, grant numbers: 302661 (VAM), 261655 and 309393 (INTPART), 219616 and 267428 (the Oslo Fiscal Studies program).

Dissecting Idiosyncratic Earnings Risk*

Elin Halvorsen[†]

Hans A. Holter[‡]

Serdar Ozkan[§]

Kjetil Storesletten[¶]

August 7, 2020

Abstract

This paper examines whether nonlinear and non-Gaussian features of earnings dynamics are caused by hours or hourly wages. Our findings from the Norwegian administrative and survey data are as follows: (i) Nonlinear mean reversion in earnings is driven by the dynamics of hours worked rather than wages since wage dynamics are close to linear while negative changes to hours are transitory and positive changes are persistent. (ii) Large earnings changes are driven equally by hours and wages, whereas small changes are associated mainly with wage shocks. (iii) Both wages and hours contribute to negative skewness and high kurtosis for earnings changes, although hour-wage interactions are quantitatively more important. (iv) When considering household earnings and disposable household income, the deviations from normality are mitigated relative to individual labor earnings: changes in disposable household income are close to symmetric and less leptokurtic.

JEL Codes: E24, H24, J24, J31.

Keywords: Earnings dynamics, Income shocks, Insurance, Wages, Hours, Higher-order earnings risk, Skewness, Kurtosis.

*We thank seminar participants at the 2017 NBER Summer Institute; Working Group on Consumption—Micro to Macro (Cambridge), the 2017 Meeting of the Society for Economic Dynamics (Edinburgh), the 2018 North American Winter Meeting of the Econometric Society (Philadelphia), Barcelona GSE, 2018 Annual Research Conference at Bank of Spain, Institute for Fiscal Studies (London), Toulouse School of Economics, University of Cologne, University of Bonn, University of Southampton, UNC Charlotte, Universitat Autònoma de Barcelona, Insper, Fundação Getúlio Vargas, University of Florida, the Federal Reserve Bank of Cleveland, University of Essex, Center for Retirement Research at Boston College, Statistics Norway, and the University of Oslo for helpful comments and suggestions. We are grateful for support from the Research Council of Norway, grant numbers: 302661 (VAM), 261655 and 309393 (INTPART), 219616 and 267428 (the Oslo Fiscal Studies program).

[†]Statistics Norway; elin.halvorsen@ssb.no; www.ssb.no/en/forskning/ansatte/elin-halvorsen

[‡]University of Oslo; hans.holter@econ.uio.no; <https://sites.google.com/site/hansaholter/>

[§]University of Toronto; serdar.ozkan@utoronto.ca; www.serdarozkan.me

[¶]University of Oslo; kjetil.storesletten@econ.uio.no; <http://folk.uio.no/kjstore/>

1 Introduction

The nature of income dynamics and the distribution of idiosyncratic shocks are crucial for behavioral choices of consumption, savings, and leisure, and influence the design of optimal social insurance and taxation. While the early literature studying idiosyncratic income fluctuations focused mostly on linear and symmetric models of risk, recent contributions have explored *nonlinearities* and *nonnormalities* (e.g., [Arellano *et al.* \(2017\)](#); [Guvenen *et al.* \(2014, 2019\)](#)). In particular, this literature has documented that the persistence of innovations is not uniform but exhibits systematic asymmetries—for example, that large negative earnings shocks are less persistent than positive changes—and that the distribution of innovations to income displays strong left skewness and excess kurtosis than normally distributed shocks. Much of this literature has focused on fluctuations in *individual annual earnings*. However, for many questions in economics, such as optimal taxation and consumption-savings choices, it is important to understand not only the dynamics of earnings but also the dynamics of each of the components that labor income comprises: hours worked and hourly wages.¹ Moreover, the relevant risk for the household is in disposable household earnings and not individual labor earnings before taxes.

This paper decomposes earnings shocks into changes in hours and changes in wage rates and studies the extent to which the nonlinear and non-Gaussian aspects of the fluctuations in male earnings are driven by hours, wages, or their interactions. We also examine the role of specific sources of large earnings shocks such as job changes. Finally, we examine the extent to which the non-Gaussian aspects of male earnings changes are passed through to household earnings and household disposable income. We study these questions using nonparametric methods building on [Guvenen *et al.* \(2014, 2019\)](#), which enables us to detect the sources behind the nonnormalities and nonlinearities in a descriptive and intuitive way. To this end, we use administrative panel data from Norway. This data set covers the entire population and is derived from a combination of administrative registers such as annual tax records and employment registers. To derive a high-quality measure of hours worked for the entire population, we develop an imputation procedure for hours worked based on merging data from the Norwegian Labor

¹For example, if the asymmetric persistence, negative skewness, and high variance of earnings changes were due entirely to changes in hours worked, then the policy prescriptions for optimal taxation and social insurance might differ from what they would be if these features instead were driven by changes in hourly wage rates.

Force Survey with the administrative registry data. Specifically, using the Labor Force Survey data, we estimate actual annual hours as a function of observables available in the registry data. We then impute a measure of hours in the registry data by applying the estimated model. This imputation procedure is an independent contribution of our paper.²

We start by decomposing contemporary earnings growth into hours and (hourly) wage components conditional on workers' age and past earnings. For a majority of workers with recent earnings around the median, hours and wage growth are about equally important in accounting for large changes in earnings, whereas small earnings changes are mainly driven by wage growth. Low and high earners exhibit different patterns, however. For individuals with low past earnings, hours changes account for a larger fraction of earnings growth than does wage growth. For high earners, this pattern is reversed, with wage growth accounting for most of their earnings fluctuations. The main events associated with large negative or positive earnings shocks are transitions in and out of long-term sickness, transitions between full-time and part-time work, and job changes.

We next document that the persistence of male earnings changes in Norway is highly asymmetric, a finding broadly consistent with other studies for U.S. and Norwegian workers (c.f., [Arellano *et al.* \(2017\)](#); [Guvonen *et al.* \(2019\)](#)). Small shocks and large positive changes are essentially permanent, whereas large negative shocks are transitory for most workers. The exception is the high earners, for whom negative changes are highly persistent and positive ones are more transitory. Exploiting the administrative nature of the registry data—which includes even those who drop out of the workforce—our methodology allows us to capture effects working through both the intensive and the extensive margins of labor supply.

We investigate the dynamics and mean reversion patterns of hours worked versus hourly wages in order to understand the drivers of the nonlinear persistence of earnings. We uncover a sharp dichotomy between hours and wages. Changes in wage rates are highly persistent. This holds true for both positive and negative changes, and for small and large changes alike.³ In contrast, the persistence of changes in hours worked turns

²The administrative data contain a measure of contractual hours worked, reported by employers. However, this variable does not include overtime and is prone to measurement error as employers often fail to accurately update hours changes.

³The only exceptions are the workers at the lower and higher ends of the recent earnings distribution. For low-income workers, large negative wage changes tend to be less persistent, whereas for high-income earners, large positive wage growth tends to be less persistent.

out to be highly nonlinear. In particular, moderate and large reductions in hours worked tend to be transitory and have mostly disappeared five years after an initial fall, whereas increases in hours worked are permanent. This holds true for all workers except for those with the highest recent earnings. We conclude that the nonlinear persistence of individual earnings changes in Norway is mainly driven by hours worked and not hourly wages. Namely, earnings declines for the majority of workers are transitory because they are to a larger extent driven by hours declines, which are transitory. Again, the exception is the high earners, for whom declines have a somewhat larger persistence because earnings reductions for these workers are primarily driven by wage declines, which are persistent.

We then turn to the higher-order moments of the distribution of individual earnings changes in Norway. A first observation is that the higher-order moments are remarkably similar to the corresponding moments reported by [Guvenen *et al.* \(2019\)](#) for U.S. workers. The variance of earnings growth is falling in age and in recent earnings. Earnings growth is not symmetric but negatively skewed, and the left skewness becomes more pronounced as individuals get older or their earnings increase.⁴ Moreover, this distribution is highly leptokurtic; that is, most individuals experience very small earnings changes in a given year, and a small but non-negligible fraction sees extreme changes. The variance of earnings growth is larger in the U.S. than in Norway. However, the changes in the variance over the life cycle and between income groups are very similar between these two economies. The same goes for the levels of the higher-order moments and their variation over the life cycle and between income groups. We conclude that, despite the differences between Norway and the U.S. in their welfare state and labor market institutions, the nonnormalities and nonlinearities in earnings dynamics are very similar. This might reflect that these dynamics are driven by similar underlying economic mechanisms.

We study the distributions of hours and wage growth and their role in driving the distribution of earnings growth and find that both hours growth and wage growth display non-Gaussian features. Hours growth (wage rate growth) is more (less) negatively skewed and has higher (lower) kurtosis than earnings growth. To quantify the importance of hours and wages in the higher-order moments of log earnings growth, we apply an exact statistical decomposition where the skewness (kurtosis) of earnings growth is a weighted sum of the skewness (kurtosis) of hours and wage growth plus a residual term

⁴Women instead have close to zero skewness, that is, an almost symmetric distribution of earnings changes.

that broadly captures whether large hours and wage changes coincide. We label the residual term as co-skewness (co-kurtosis). All terms contribute to the left skewness of earnings growth, with co-skewness being the main component (i.e., workers experience large declines in hours and wages simultaneously). Similarly, the co-kurtosis term is the largest contributor to the excess kurtosis.

We also investigate the role of job changes—one of the main events associated with large earnings changes—in driving the skewness and kurtosis of earnings changes. We find that the earnings growth for job switchers exhibits less negative skewness and less excess kurtosis than for job stayers. Moreover, earnings changes for job switchers is close to being normally distributed (especially for low earners). Therefore, the skewness and kurtosis of earnings growth tend to be driven by job stayers.

Finally, we examine how the dynamics of household earnings and household disposable income differ from male earnings dynamics. We find that the higher-order moments for disposable income growth and, to a smaller extent, household earnings growth differ sharply from those of male earnings growth: changes in disposable income have substantially lower variance than changes in male earnings. Moreover, growth in household earnings and disposable income is less negatively skewed, with disposable income growth being close to symmetrically distributed. Thus, the Norwegian system of taxes and transfers provides substantial insurance against fluctuations in income.

Our paper is structured as follows. Section 2 reviews some of the recent literature on earnings dynamics. Section 3 describes the data and empirical methodology. Section 4 decomposes earnings risk into changes in hours and changes in wage rates and studies their dynamics and contributions to earnings growth. Section 5 studies the higher-order moments of wage and hours growth and their contributions to the higher order moments of earnings growth. It also examines how the higher-order moments are affected when going from male labor earnings to household earnings and household disposable income. We then conclude in Section 6.

2 Related Literature

A number of empirical studies have focused on the distributional properties of income shocks, including classic contributions by Lillard and Willis (1978), Lillard and Weiss (1979), MaCurdy (1982), and Abowd and Card (1989). Most of this work has been based on survey data on individual labor earnings. Because of limited sample sizes, the approaches in these studies have been parametric. Guvenen *et al.* (2014, 2018) rely instead

on administrative data from the U.S. Social Security Administration (SSA) to study the distribution of individual annual earnings changes. They document that male earnings changes are strongly negatively skewed (negative shocks have a thicker tail than positive shocks) and more leptokurtic than a normal distribution (thicker tails for both positive and negative shocks). However, several important questions cannot be answered because of SSA data limitations. One shortcoming is the lack of detailed information on components of annual individual earnings, such as hours worked, hourly wages, unemployment spells, and other nonworking spells. Another shortcoming is missing information on other household members and other types of income such as government taxes and transfers and capital income. The Norwegian data used in this paper contain all this information and are the basis of our investigation.

Some recent papers consider higher-order moments of household labor income risk and household disposable income risk but focus mainly on the variation over the business-cycle. [Pruitt and Turner \(2018\)](#) use IRS data to study the distribution of household labor income risk over the business cycle. In line with our findings for Norway, they find that household labor income growth in the U.S. is significantly less negatively skewed than individual earnings growth. [Busch *et al.* \(2018\)](#) use data from the U.S., Germany, and Sweden to study variations in higher-order earnings risk over the business cycle. They argue that the family and the welfare state are only moderately efficient in terms of mitigating business cycle risk.

The paper closest to our study of higher-order moments is [De Nardi *et al.* \(2019\)](#), which is developed contemporaneously with our paper. They study the distribution of household income before and after taxes and transfers as well as the distribution of hours and wage growth in the Netherlands. In line with our findings for Norway, they find that wage growth is negatively skewed and that there is significant insurance from the family and the welfare state. They do not study asymmetric mean reversion of earnings, as we do. Another difference between our paper and [De Nardi *et al.* \(2019\)](#) is that they rely on an employer-reported hours measure in the Dutch register data (mainly contracted hours). We are instead able to apply a high-quality imputation of hours worked based on the Norwegian Labor Force Survey, exploiting detailed government and employer records. We also have access to contracted hours, reported by the employer, in the Norwegian register data and include this as part of our imputation procedure. However, we document that contracted hours data contain large and systematic measurement error relative to survey data on actual hours. In particular, changes in contracted hours are

biased toward zero relative to changes in actual hours.

While our paper and the papers discussed above pursue a nonparametric approach to studying income shocks, [Arellano *et al.* \(2017\)](#) put more structure on the income process. They develop a quantile-based panel data framework to study the nature of income persistence and the transmission of income shocks to consumption where log earnings is modeled as the sum of a general Markovian persistent component and a transitory innovation. Using both the Panel Study of Income Dynamics (PSID) and Norwegian register data, they document that the persistence of earnings shocks is nonlinear (asymmetric). Our results for earnings are in line with their paper, and our contribution is to extend the study of nonlinearities to hours and wage rates.

Our paper also contributes to the broader literature on idiosyncratic risk and risk sharing, including [Cutler and Katz \(1992\)](#), [Deaton and Paxson \(1994\)](#), [Attanasio and Davis \(1996\)](#), [Blundell and Preston \(1998\)](#), [Krueger and Perri \(2005\)](#), [Blundell *et al.* \(2008\)](#), and [Heathcote *et al.* \(2014\)](#). [Blundell *et al.* \(2014\)](#) study insurance in Norway through two specific channels: the welfare state and the family. They find that the welfare state provides a large reduction in the variance of persistent and transitory shocks relative to individual labor income and that this reduction is larger for low- and medium-skilled workers relative to high-skilled workers. Our findings are in line with these results.

3 Data and Methodology

3.1 Income and Labor Data

Our analysis uses data from four different data sources between 1993 and 2014. The first data set is Administrative Tax and Income Records, which contains a set of detailed information on income and taxes for the entire Norwegian population from 1993 onward. In addition, this register contains age, gender, household composition, country of origin, and education variables. Our measure of labor earnings is comprehensive and includes wages and salaries from all employment, including bonuses and other irregular payments. In addition, we have information on business income from self-employment.⁵ Tax records are of high quality because most information is third-party reported to the tax authorities, and very little is self-reported. For example, employers are obliged to

⁵Among Norwegians, 5% have just business income but no labor earnings. An additional 5% have both labor earnings and business income, although for this group, business income tends to be small relative to labor income.

report information on earnings payments. All values are deflated using the (Laspeyres) Consumer Price Index.

The basic tax unit is an individual. However, by using family identifiers from the population register, we pool individual incomes of spouses for both married and cohabiting couples to calculate household income.

The second data set is administered by the Norwegian Labor and Welfare Administration Register, which contains the start and end dates of spells for unemployment, parental leave, sickness, and disability benefits at the daily level. We use this information to impute hours worked, which we discuss below. The data set also contains accurate information about public insurance through transfers and taxes. Transfers include unemployment benefits, sickness benefits, paid parental leave, remuneration for participation in various government activity programs, disability benefits, public pensions, and other social welfare payments.

Third, the Employment Register is a matched employer-employee data set. All employers are required to report contractual hours, employment duration, sector, and industry to the government. Information about contractual hours of work is limited to the period from 2003 to 2014, since prior to 2003 only full-time and part-time hours was reported.⁶ The Employment Register covers the entire labor force, except for self-employed workers and freelancers. This amounts to 90% of the labor force and 77% of the prime-age population (25 to 60 years old). For individuals with multiple jobs during the year, we define main employment as the job that accounts for the largest share of annual earnings and measure annual contractual hours as the sum of contractual hours worked in all jobs.

3.2 Imputation of Hours

Contractual hours in the the Employment Register has some significant weaknesses as a measure of hours worked. First, the register was originally administered by the Social Security Administration and used for calculating work-related benefits. Therefore, contractual hours do not include overtime. Second, the register does not cover income from employment that amounts to less than four hours per week (on average) or seven days per month. A large fraction of work paid by the hour may therefore remain unreported. Lastly, the register contains substantial measurement error. It is well known

⁶We drop data after 2014, after which a new system for reporting hours worked was introduced, leading to another major break in the data series.

that employers often forget to update changes in employment spells or hours worked. Because of these shortcomings, we turn to the Norwegian Labor Force Survey (AKU) to obtain a better measure of hours worked. This data set contains high-quality survey data on actual hours worked but has a limited sample size. The purpose of this survey is to measure employment and hours worked. As it turns out, all individuals present in the Labor Force Survey are also present in the register data. We merge the two data sources using individual identification numbers and design a novel imputation approach to infer actual hours worked for the entire population.

The Labor Force Survey records weekly hours worked. Each individual is surveyed up to eight consecutive quarters. We use only those who are present in all eight quarters and impute actual annual hours in year t as $h_t^{LFS} = 13 \cdot \sum_{q=1}^4 h_{t,q}^{LFS}$, where $h_{t,d}^{LFS}$ is weekly hours in quarter q of year t . We then regress actual annual hours h_{it}^{LFS} from the Labor Force Survey on information in the register data:

$$h_{it}^{LFS} = \alpha h_{it}^{REG} + \beta X_{it} + \epsilon_{it}, \quad (1)$$

where h_{it}^{REG} is contractual hours reported according to the Employment Register and X_{it} contains a rich set of observables from the register data: sickness days, parental leave days, unemployment days, part-time, sector, labor earnings, country of origin, and education. We estimate the model separately for men and women and for each recent earnings quintile (see Section 3.3 for our measure of recent earnings). We use the estimated model in (1) to impute actual work hours for the individuals that are not present in the Labor Force Survey.⁷ We add bootstrapped errors to the imputed hours, using the approach of resampling residuals from the original regression. Residuals are clustered by gender and recent earnings and then drawn randomly within these bins to match the imputed hours based on the whole register population.

The two most important covariates are contractual hours and labor earnings. The number of days receiving benefits for sickness, parental leave, and unemployment is also an important predictive variable. This is not surprising since the number of days on benefits is very accurately measured—it is based on actual benefit payments—whereas the number of employer-reported contractual hours in the Employment Register often misses such benefits spells. The estimations show that our model has greater explanatory power for women than for men, partly because of contractual hours and earnings having

⁷See Appendix B for a full set of results.

a higher correlation with actual hours worked for women than for men.

How good is our imputation of hours worked? The explanatory power of our imputation is relatively high, measured in terms of overall R -squared: about 0.19 for men and 0.41 for women. This is comparable to the explanatory power of Mincer-type linear regressions on data from the PSID. Regressions on these data with annual hours worked as the dependent variable and standard covariates as explanatory variables (gender, education, a quartic in age, and, most importantly, annual earnings) yield an overall R -squared of 0.45 for women and 0.16 for men (see Online Appendix D for details).⁸ We also evaluate the quality of our imputation by considering out-of-sample predictions. More precisely, we first estimate the model on a random half of the sample and use the estimates to predict hours for the second half of the sample. We find that adjusted root mean square errors (RMSE) are similar for the two samples, only slightly higher out of sample, with a difference that is not statistically significant. Figure 1 documents the success of our imputation approach. The left panel shows the kernel distribution of the actual annual hours as measured in the Labor Force Survey together with the kernel distribution of the imputed hours from the register data.⁹

Given that our focus is to quantify the importance of hours changes in earnings risk, we investigate the average hours growth conditional on earnings changes from our imputation approach compared to the Labor Force Survey and Employment Register data. In particular, in the right panel of Figure 1, we rank individuals present in the Labor Force Survey into 10 bins based on their one-year earnings changes, where earnings data are from the register data. For each bin, we plot the log change of average earnings between t and $t + 1$ on the x -axis against the corresponding change in average annual hours on the y -axis for three hours measures: actual hours from the Labor Force Survey, contracted hours from the Employment Register, and our imputed hours measure.


The figure shows that our imputed hours changes are remarkably close to the actual changes in the Labor Force Survey data. For example, changes in imputed hours account for 91% of the changes in actual hours for men experiencing large earnings changes (cal-

⁸If earnings is dropped as an explanatory variable, the R -squared falls substantially for the PSID, whereas it remains high for our Norwegian register data, as a result of detailed information on days receiving unemployment and sickness benefits.

⁹We have experimented with alternative methodologies and models, such as quantile regression, regression in growth instead of levels, allowing for more interaction between the explanatory variables, other forms of more flexible parameterization, or some machine learning algorithms. In the end, none of these alternatives outperformed the simple linear OLS in terms of explanatory power and out-of-sample prediction.

culated as the average log change for the group as a whole). In contrast, changes in contracted hours (from the administrative data) are significantly smaller than those of the survey data, especially for large earnings changes. For example, changes in contracted hours are less than half of the changes in actual hours worked. We conclude that contracted hours data contain large and systematic measurement error relative to survey data on actual hours. In particular, both small and large changes in measured hours are biased toward zero for contracted hours relative to the changes in actual hours according to the survey data. In the rest of the paper, we use the imputed hours as our measure of annual hours worked.

FIGURE 1 – Actual Hours, Contracted Hours and Imputed Hours



Note: The left panel displays the kernel distribution of the actual annual hours in the Labor Force Survey and the kernel distribution of the imputed hours from the register data. The right panel displays the average one-year changes in contracted annual hours, imputed annual hours, and actual annual hours (as measured in the Labor Force Survey) for individuals with different one-year changes in annual earnings.

3.3 Sample Selection and Empirical Methodology

We follow a nonparametric empirical methodology building on [Güvenen *et al.* \(2019, 2014\)](#). The fundamental idea is to group workers with similar observables at a sufficiently fine level so that they can be thought of as approximately ex ante identical. Then, for each such group, we investigate the properties of income changes as a proxy for the nature of idiosyncratic risk that individuals within that group are facing. This methodology allows us to uncover the heterogeneity in the nonnormalities and nonlinearities in earnings dynamics that different groups of workers face.

Base Sample

Our *base sample* is a revolving panel of 25- to 60-year-old workers with a reasonably strong labor market attachment. We first define an individual-year earnings observation as being *admissible* for that year if the individual (i) is between 25 and 60 years old and (ii) has wage earnings above $Y_{\min,t}$, which is equal to 5% of median earnings, and (iii) works more than 200 (imputed) hours per year.¹⁰ Then, for each year t between 1998 and 2013, we select individuals that are admissible in $t-1$ and in at least two more years between $t-5$ and $t-2$. This condition ensures that the individual has a reasonably strong labor market attachment. Given these restrictions, our sample consists of 28.9 million individual-year observations in total, which is roughly 900,000 males and 800,000 females per year.

Worker Groupings

One of the key observables we group workers with is their recent earnings (RE) between $t-1$ and $t-5$, \bar{Y}_{t-1}^i . By requiring individuals to have at least three years of admissible income in the last five years, we ensure that we can compute a reasonable measure of this average past income. We compute each individual's RE \bar{Y}_{t-1}^i by summing his or her annual wages normalized by age effects between $t-1$ and $t-5$:

$$\bar{Y}_{t-1}^i \equiv \sum_{s=1}^5 \frac{\tilde{Y}_{t-s,h-s}^i}{\exp(d_{h-s})},$$

where $\tilde{Y}_{t-s,h-s}^i$ denote the annual wage earnings of individual i who is h years old in year t . The constants d_{h-s} are age dummies from regressing log individual earnings on a full set of age, gender, and cohort dummies. Next, we group workers by their gender and age in $t-1$. Within each of these groups, we rank workers into 10 deciles with respect to their recent earnings \bar{Y}_{t-1}^i .¹¹

¹⁰Heathcote *et al.* (2010a) also restrict attention to individuals who work at least 200 hours per year. Guvenen *et al.* (2019) choose $Y_{\min,t}$ as the income from one quarter of full-time work at half of the legal minimum wage, which corresponds to 5% of median earnings in the U.S. but they have no information on hours worked. The constraint (iii) on hours curtails the sample by only 1% relative to constraints (i) and (ii).

¹¹The implication of this analysis is that RE percentiles are age group dependent. The advantage of this approach is that it ensures each RE group contains a similar number of observations, whereas grouping workers based on the RE distribution in the overall sample will result in too many younger workers appearing in lower RE percentiles and vice versa for middle-age workers. As a robustness check, we first group workers based on the RE distribution in the overall sample, and then within each RE group, we classify workers by age. We find that our main conclusions are robust to this change.

Growth Rate Measures

We use two types of measures of growth in the variables of interest Z (e.g., earnings, income, hours worked, and hourly wages). The first measure focuses on growth in average allocation \bar{Z} for a group of similar workers. In particular, for a group j of workers who have the similar observable characteristics, V_{t+1}^j (e.g., they are in the same age group, have similar recent earnings in $t - 1$, and have experienced a similar earnings change between t and $t + 1$), we define the average growth between t and $t + k$ as follows:

$$\Delta RA_{t,t+k}^Z = \log(\bar{Z}_{t+k,h+k}^j | V_{t+1}^j) - \log(\bar{Z}_{t,h}^j | V_{t+1}^j),$$

where $\bar{Z}_{t,h}^j \equiv \sum_{i=1} Z_{t,h,i}^j$ and $Z_{t,h,i}^j$ is variable Z for individual i in group j .

We refer to this as the *representative agent (RA)* change. One major advantage of this approach is that it incorporates the extensive margin when going forward. For example, even though a person drops out of the labor market, the impact of his or her (zero) earnings will be included through the average change for the group. Another advantage of the representative agent change is that the persistence of the original shock to earnings is crystallized in a nonparametric way. We assume that future idiosyncratic shocks to individuals in the group are independent. These future innovations will therefore wash out across group members. Thus, the evolution of the group mean captures by construction the expected evolution after the initial shock for the group. We study both short-term and long-term RA changes.

When investigating the distribution of shocks (in Section 5), we focus instead on an individual-oriented measure because we are interested in the higher-order moments of the individual dynamics of various variables. We work with the log growth rate in individual-specific variables between t and $t + k$:

$$\text{log change: } \Delta_{\log}^k z_t^i \equiv z_{t+k,h+k}^i - z_{t,h}^i,$$

where $z_t^i \equiv \log Z_{t,h}^i - d_h^z$ denotes log of variable z net of age effects of the same variable. This is a widely used growth rate measure, and its higher-order moments for a log-normal distribution are familiar to most readers (zero skewness and a kurtosis coefficient of 3). But it is also well-known that observations close to zero need to be dropped or winsorized at an arbitrary value. Thus, when we use $\Delta_{\log}^k z_t^i$, we drop individuals from

the base sample whose data are not admissible in t or $t + k$.¹²

4 Dissecting Earnings Dynamics

Changes in earnings are due to changes in hours worked or the hourly wage rate or both. This section documents the extent to which earnings dynamics are driven by hours versus hourly wages. The answer to this question matters for many economic questions, including risk sharing and social insurance arrangements (see, e.g., [Heathcote *et al.*, 2010b](#) and [Conesa *et al.*, 2009](#)). A large literature, dating back to seminal papers by [Abowd and Card \(1989\)](#), [MaCurdy \(1981\)](#), [Altonji \(1986\)](#), and [Abowd and Card \(1989\)](#), has studied the covariance structure of changes in wages and hours. Most of the focus has been on uniform relations between movements in wages and movements in hours. In particular, data restrictions have made it difficult to examine possible heterogeneity in the covariance structure of wage and hours growth. In this paper, we exploit the sheer size of our administrative data and our novel imputation of hours to document the heterogeneity in the co-movement of hours and wage growth between small and big and negative and positive earnings changes of workers with different earnings histories.

We first quantify the importance of hours and wage changes in the impact of earnings changes, and then we turn to their persistence and document how hours versus hourly wages varies in accounting for the asymmetries and nonlinearities in earnings dynamics.

4.1 Decomposing Earnings Changes to Hours and Wage Growth


To document the heterogeneity in the covariance structure of wage and hours growth, we plot changes in hours and hourly wages against changes in earnings for different groups of workers. For this purpose, on top of conditioning workers with respect to their age and recent earnings \bar{Y}_{t-1}^i , we also group them with respect to their earnings growth. In particular, as described previously, we first group workers into “young” (ages 25 to 35) and “prime age” (ages 36 to 55), and then within each age group, we rank them into 10 deciles with respect to their recent earnings (\bar{Y}_{t-1}^i) in $t - 1$. Next, within each age and RE group, we further sort workers into 20 quantiles according to their earnings growth

¹²The individual log change $\Delta_{\log}^k z_t^i$ ignores some potentially valuable information on the extensive margin. For example, the long-term unemployed must be dropped (note that this caveat does not apply to the representative agent measure). Thus, for robustness we also conduct most of our individual-based analysis with an arc-percent change measure, $\Delta_{\text{arc}}^k z_t^i = 2(Z_{t+k}^i - Z_t^i) / (Z_{t+k}^i + Z_t^i)$, which is not prone to this caveat and is commonly used in the firm dynamics literature. Our results are qualitatively robust to the choice of the individual-based growth rate measure. Results are available upon request.

from t to $t + 1$. We then treat each such finely defined group as homogeneous and plot the growth of their average hours and hourly wages on the y -axis conditional on their earnings growth between t and $t + 1$ on the x -axis.¹³ To control for age effects and differences in mean reversion between different groups of workers, we normalize changes on both the x - and y -axes such that their values at the median quantile cross at zero.

We start with the 40% of prime-age males (36-55 years old) who are in the middle of the recent earnings distribution (the 4th to 7th deciles with little differences between them). The results are plotted in Figure 2. Note first that for large negative earnings changes, hours growth is roughly as large as wage rate growth. For example, the group of workers whose earnings decline around 60 log points on average experience a decline of about 30 log points in hours and a decline of 30 log points in wage rates. However, for large positive changes, the split is slightly more skewed toward hourly wage changes. Second, smaller earnings changes (both gains and losses) are mainly driven by wage changes. For example, for men who experience a loss of 10 log points in earnings, more than 70% of this loss is from a decline in wage rates. These results illustrate the heterogeneity in the covariance structure of wage and hours growth over the earnings change distribution.

FIGURE 2 – Contribution of Hours and Wages to Earnings Shocks, 4th-7th RE Deciles

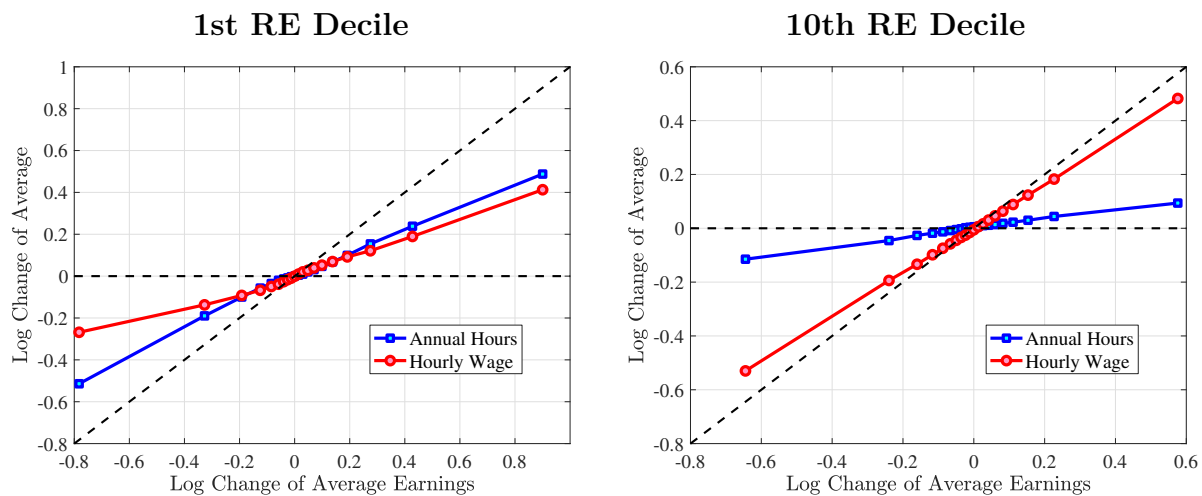


Note: The figure displays the one-year representative agent change (log change of averages) for imputed hours and imputed wage rates for 20 different groups of prime-age males (ages 36 to 55) in the 4th-7th RE deciles, plotted against their contemporaneous one-year log change in average annual earnings.

¹³The results when using the alternative measure of changes—the *average of log earnings change* within each group, $\Delta_{\log}^k z_t^i$ —are qualitatively similar and are available upon request.

We next consider the role of hours and wage changes for the top and bottom of the recent earnings groups. Figure 3 plots the changes in hours worked and wage rates against changes in earnings for the bottom decile (left panel) and top decile (right panel) of recent earnings (see Figure A.6 in the Appendix C.2 for the 2nd and 8th RE deciles). For the bottom decile of recent earnings, changes in hours worked are more important than changes in wage rates in accounting for earnings changes, especially for large earnings declines. This result is flipped for the top earners: high earners experience only minor changes in hours worked, and thus most of their earnings changes are from changes in wage rates. These findings suggest that different economic mechanisms are behind the earnings dynamics of high and low earners. These results are consistent with previous research showing that unemployment risk is an important component of idiosyncratic risk for low-income workers, whereas wage fluctuations are the main drivers of income risk of workers at the higher end of the income distribution who have more stable jobs (Karahan *et al.* (2019)).

FIGURE 3 – Contribution of Hours and Wage Rates to Earnings Shocks



Note: The figure displays the one-year representative agent change (log change of averages) for imputed hours and imputed wage rates for 20 different groups of prime-age males (ages 36 to 55) in the 1st RE decile (left panel) and 10th RE decile (right panel), plotted against their contemporaneous one-year log change in average annual earnings.

4.2 Asymmetric Mean Reversion

We now document how the persistence and mean reversion of earnings changes can be attributed to the dynamics of hours versus wage rates. We start by studying the persistence of earnings. We first illustrate how the persistence of earnings changes varies


by the magnitudes of the initial change. To this end, we plot the change in average earnings after five years against the initial change for prime-age males. The x -axis has the initial (average) change $y_{t+1}^i - y_t^i$ for each quantile of workers, sorted by the size of their earnings shock. The y -axis plots the representative agent change in earnings, that is, the change in the log of average earnings for each such quantile from t to $t + 5$, $\log \mathbb{E} [Y_{t+5}^i] - \log \mathbb{E} [Y_t^i]$, where Y_t^i is the income level net of age and time effects of individual i . If the initial change is permanent, then $\mathbb{E} [Y_{t+5}^i] = \mathbb{E} [Y_{t+1}^i]$ because individual changes after $t + 1$ wash out across people in the quantile. In this case, the observations will line up along the 45-degree line. Conversely, if the change between t and $t + 1$ is transient, then $\mathbb{E} [Y_{t+5}^i] = \mathbb{E} [Y_t^i]$, and the observations will line up on the x -axis. Note also that our representative agent approach incorporates changes in the extensive margin of labor supply after the initial change. While all people in the sample by construction satisfy the sample restrictions in periods t and $t + 1$, we do not impose any restriction for period $t + 5$. Therefore, the entire quantile, even those with zero earnings and hours, is included in $t + 5$.

Consider first the workers around the median of recent earnings (4th to 7th deciles). Figure 4 reveals a striking pattern: both small changes and large positive changes are close to permanent. However, for the 10% of workers who experience the largest negative changes (i.e., reductions in earnings of more than 15%), the earnings changes are more transitory. For example, workers who experience an initial 35% decline (45 log points) in earnings relative to time t have on average a reduction in earnings of just 15% five years later.

For the bottom decile of RE workers, these patterns are even more pronounced: earnings losses are transitory, whereas earnings gains are, for all practical purposes, permanent. For example, workers in this group who experienced around a 55% decline (80 log points) in their earnings between t and $t + 1$ see their earnings in 5 only 10% less than their t values, whereas we do not see any mean reversion for either small changes or large positive changes.

Consider now the top decile of RE workers. For this group, small changes are also permanent. However, different from workers with lower recent earnings, large positive changes are quite transient and large negative changes are highly persistent. For example, for the workers with initial earnings increases of 75% (55 log points), only half of the initial change remains after five years. Conversely, the workers who experience an initial 50% decline (70 log points) see a sustained 40% decline (50 log points) five years later.

FIGURE 4 – Persistence of Earnings Changes, Prime Age Males



Note: The figure displays the five-year representative agent change (log change of averages) in earnings for 20 different groups of prime-age males (ages 36 to 55) in the 1st RE decile, 4th-7th RE decile, and 10th RE decile, plotted against their respective one-year log change in average annual earnings.

These findings are consistent with [Arellano *et al.* \(2017\)](#) and very similar to what [Guvenen *et al.* \(2019\)](#) documented for the U.S.

4.2.1 Mean Reversion of Hours and Wage Changes

What explains this asymmetric mean reversion for earnings? To answer this question, we investigate the persistence of hours and wage rate changes separately. In line with the strategy above, we group workers with respect to age and recent earnings as well as their hours or wage rate growth between t and $t + 1$. For each quantile of change in hours or wage rate, we then plot the log of the five-year growth of averages within the quantile (t to $t + 5$) on the y -axis against the log of the one-year growth of the average variable between t and $t + 1$.

Starting with the persistence of hours change, the left panel of [Figure 5](#) shows the five-year log change in average hours for different sizes of impulses for workers around the median of recent earnings, as well as those in the bottom and top deciles (similar to [Figure 4](#)). We show this graph for only the prime-age group, but young workers display a similar pattern (see [Figure A.16](#) in the Online Appendix). Note first that hours declines

are very transitory, whereas hours increases are close to permanent. This finding is in line with the evidence in [Krusell *et al.* \(2011\)](#) that the duration of employment spells is much larger than the duration of unemployment spells. However, some differences can be seen in how persistent hours changes are between different income groups. In particular, for the bottom earners, declines are fully transitory and increases are fully permanent. As we move to higher recent earnings deciles, increases in hours become slightly less persistent and declines become somewhat more persistent. These differences are likely due to different events leading to hours changes for different income groups. For example, for bottom earners the main drivers of hours changes are likely to involve transitions between non-employment and employment or from part-time to full-time work, whereas for higher-income workers, changes in hours may be due to more flexible work conditions, such as the possibility to work overtime or have multiple employments. We will revisit this issue in [Section 4.3](#) when we associate real-life events to earnings changes.

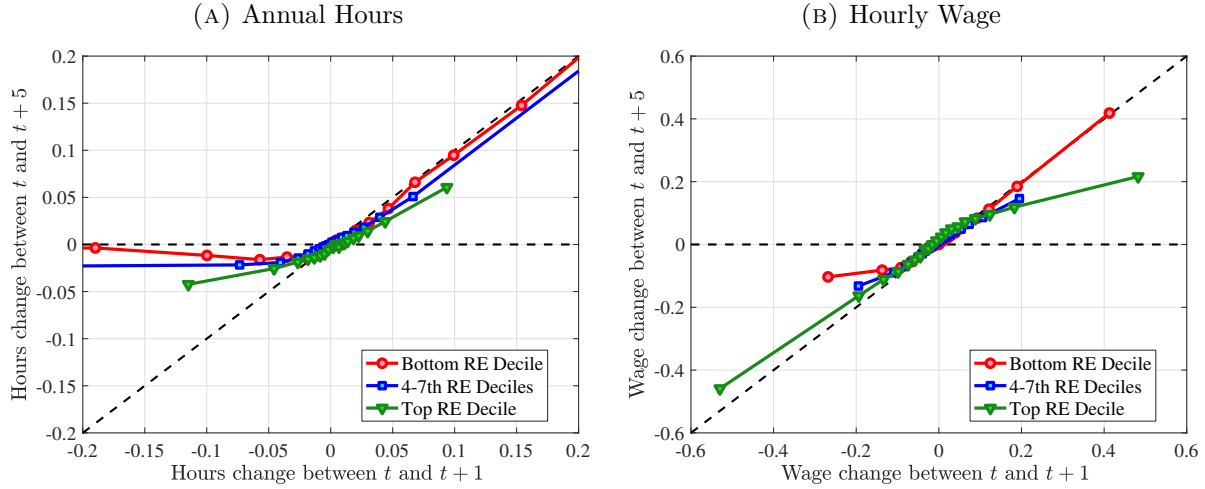
The right panel of [Figure 5](#) shows the persistence of wage rate changes. Unlike hours changes, wage rate changes are very persistent regardless of whether they are positive or negative. Differences can be noted across recent earnings groups: large wage rate increases for top earners and large wage rate losses for bottom earners tend to be somewhat less persistent. One factor contributing to the low persistence of earnings gains for the top earners could be the fact that high earners receive some of their income in the form of bonuses and stock options, and this income is highly cyclical (c.f. [Parker and Vissing-Jørgensen, 2010](#)).

We conclude that the nonlinear persistence of earnings documented in [Figure 4](#) is largely a result of nonlinear persistence in hours worked. In particular, hours changes exhibit nonlinear persistence patterns that are very similar to those of earnings. In contrast, wage rate dynamics are more linear, especially for workers around the median. Earnings declines are very persistent for high earners because hours do not move much for these workers—the fall in their earnings is due to a decline in wage rates.

4.3 The Life Events Associated with Large Earnings Shocks

A natural question for the purpose of dissecting idiosyncratic earnings changes is, what are the major events in an individual’s life that lead to small versus large earnings changes (e.g., see [Cochrane, 1991](#))? Our data set allows us to link individuals’ earnings to information available in other administrative data sets. In particular, we focus on five

FIGURE 5 – Persistence of Hours and Wage Changes by RE Decile



Note: The left panel displays the five-year representative agent change (log change of averages) in imputed annual hours for 20 different groups of prime-age males (ages 36 to 55) in the 1st RE decile (red line), 4th-7th RE decile (blue line) and 10th RE decile (green line), plotted against their respective one-year log change in imputed annual average hours. The right panel displays the corresponding figure for imputed hourly wage rates.

important events workers experience that are known to have significant effects on their earnings: transitioning into and out of (i) unemployment, (ii) long-term sickness, (iii) part-time work, (iv) parental leave, and (v) job change to a different firm. We investigate the likelihood of these events in six groups of workers, sorted by the size of their earnings change between t and $t + 1$. Rows (1)-(5) of Table I show the fraction of workers within each earnings change group who experience the above listed five life-cycle events. In rows (6)-(9), we report the average hours and wage rate changes for each group of workers both in the same year (between t and $t + 1$) and five years later (between t and $t + 5$). The upper panel shows the entire sample of males, and the second and third panels show the results for the bottom and top recent earnings deciles, respectively.

The events corresponding to the largest earnings changes (gains or losses of more than 50 log points, columns (1) and (6)) and the intermediate changes (changes of 25 to 50 log points, columns (2) and (5)) are quite similar. The only exception is that parental leave is more likely in columns (2) and (5) than it is in columns (1) and (6). As for the minor changes—the ones smaller than 25 log points—they are less likely to be associated with these events.

The most frequent cause of large losses (>50 log points) is long-term sickness, 25%, followed by change of employer, 19%, and going from full-time to part-time, 15%. Only 8% of those suffering large losses have experienced unemployment. However, an unem-

ployment spell is on average longer than a sickness spell. The average number of weeks with sickness benefits (for males in our base sample) is around 8 weeks, whereas an average unemployment spell is 20 weeks. For those experiencing the largest earnings losses, the average decline in log hours is slightly smaller than the average log hourly wage loss: -0.40 versus -0.44. However, as discussed above, wage rate declines are substantially more persistent than hours declines. After five years, wages are down by 19 log points, whereas hours is only 3 log points lower.

The events behind the large positive earnings changes are relatively symmetric to the events associated with large earnings losses. The events most frequently associated with large positive changes are change of employer, 23%; going from part-time to full-time, also 23%; and returning from long-term sickness, 25%. The average change in log hours for workers with the largest positive earnings shocks is 0.40, which is somewhat smaller than the average change in the log hourly wage rate, 0.45. Increases in both hours and wages are quite persistent as well.

Overall, the events patterns are similar even if we study the bottom and top deciles separately, with the exception that for the top earners, large earnings losses are to a smaller extent caused by unemployment or sickness. For the top group, large earnings losses are more closely associated with firm change than anything else. Conversely, for low income earners, large earnings gains are chiefly associated with changes of employer and changes from part-time to full-time. Results in rows (6)-(9) of Table I confirms the findings from Figure 5 that negative shocks are transitory and positive shocks are permanent for the bottom decile, whereas for top earners, large negative changes in hourly wage rates are more persistent than negative changes in hours.

TABLE I – Important Life Cycle Events Associated with Earnings Changes

All Life-cycle event into/out of		Annual Earnings Change, $\Delta y \in$					
		One-Year Earnings Loss			One-Year Earnings Gain		
		< -0.5	$[-0.5, -0.25)$	$[-0.25, 0.0)$	$[0.0, 0.25)$	$[0.25, 0.5)$	≥ 0.5
	(1)	(2)	(3)	(4)	(5)	(6)	
(1)	Unemployment	0.08	0.06	0.02	0.02	0.08	0.10
(2)	Long-term sickness	0.23	0.23	0.09	0.10	0.22	0.25
(3)	Part-time	0.15	0.11	0.05	0.07	0.16	0.23
(4)	Parental leave	0.06	0.08	0.04	0.05	0.08	0.05
(5)	Firm change	0.19	0.22	0.12	0.13	0.21	0.23
(6)	$\mathbb{E} \left[\Delta_{\log}^1 h_t^i \right]$	-0.40	-0.16	-0.03	0.03	0.17	0.40
(7)	$\mathbb{E} \left[\Delta_{\log}^5 h_t^i \right]$	-0.03	-0.03	-0.03	0.00	0.13	0.37
(8)	$\mathbb{E} \left[\Delta_{\log}^1 w_t^i \right]$	-0.44	-0.18	-0.04	0.05	0.17	0.45
(9)	$\mathbb{E} \left[\Delta_{\log}^5 w_t^i \right]$	-0.19	-0.15	-0.04	0.04	0.13	0.40
(10)	# of Obs.	104,727	298,777	2,219,654	1,973,893	320,891	111,353
Lowest decile (RE=1)		(1)	(2)	(3)	(4)	(5)	(6)
(1)	Unemployment	0.10	0.09	0.04	0.05	0.10	0.10
(2)	Long-term sickness	0.20	0.19	0.10	0.10	0.14	0.13
(3)	Part-time	0.16	0.12	0.06	0.13	0.25	0.32
(4)	Parental leave	0.03	0.04	0.03	0.02	0.03	0.02
(5)	Firm change	0.19	0.22	0.14	0.17	0.24	0.28
(6)	$\mathbb{E} \left[\Delta_{\log}^1 h_t^i \right]$	-0.38	-0.16	-0.03	0.05	0.17	0.37
(7)	$\mathbb{E} \left[\Delta_{\log}^5 h_t^i \right]$	-0.00	-0.01	-0.02	0.03	0.13	0.37
(8)	$\mathbb{E} \left[\Delta_{\log}^1 w_t^i \right]$	-0.50	-0.19	-0.05	0.05	0.17	0.51
(9)	$\mathbb{E} \left[\Delta_{\log}^5 w_t^i \right]$	-0.05	-0.06	-0.01	0.06	0.13	0.58
(10)	# of Obs.	17,813	22,993	124,680	145,885	39,418	37,547
Top decile (RE=10)		(1)	(2)	(3)	(4)	(5)	(6)
(1)	Unemployment	0.05	0.02	0.00	0.00	0.02	0.10
(2)	Long-term sickness	0.08	0.09	0.05	0.05	0.08	0.25
(3)	Part-time	0.15	0.11	0.06	0.06	0.11	0.23
(4)	Parental leave	0.05	0.08	0.05	0.06	0.08	0.05
(5)	Firm change	0.31	0.24	0.13	0.12	0.21	0.23
(6)	$\mathbb{E} \left[\Delta_{\log}^1 h_t^i \right]$	-0.18	-0.09	-0.02	0.03	0.08	0.16
(7)	$\mathbb{E} \left[\Delta_{\log}^5 h_t^i \right]$	-0.03	-0.03	-0.02	0.00	0.04	0.10
(8)	$\mathbb{E} \left[\Delta_{\log}^1 w_t^i \right]$	-0.67	-0.25	-0.06	0.05	0.26	0.64
(9)	$\mathbb{E} \left[\Delta_{\log}^5 w_t^i \right]$	-0.50	-0.23	-0.05	0.04	0.14	0.30
(10)	# of Obs.	15,274	34,208	310,628	298,130	32,251	11,015

The table sorts individuals into six groups according to the size of their earnings change, defined as the percentage change in earnings from t to $t + 1$. Rows (1)-(5) in the table display the fraction in each earnings change group who experienced each of these events (not mutually exclusive). Rows (6)-(9) show the corresponding percentage change in imputed hours and hourly wage in each group in the same period (from t to $t + 1$) and five years later (from t to $t + 5$). Average over all years 1993-2014, males only.

5 Higher-Order Earnings Risk

We now turn to the higher-order moments of individual earnings changes. In order to investigate “transitory” and “persistent” innovations separately, it is useful to distinguish between income growth over short (one-year) and long (five-year) horizons (i.e., between t and $t+1$ and from t to $t+5$). The persistent component of changes becomes more salient the longer the horizon (Güvenen *et al.* (2019)).¹⁴ We document statistics from both one-year and five-year growth distributions, but in the main text we mainly focus on five-year changes since persistent changes are economically more important for consumption and savings behavior.¹⁵ The results for one-year changes are qualitatively similar (see Appendix C). We consider growth in earnings, hours, and hourly wages. In constructing these figures, we calculate the average of the moment of interest for each age/RE group over the years between 2003 and 2014- k .

Figure 6 displays the distribution of one-year (left panel) and five-year (right panel) individual earnings growth for male workers in the base sample defined in Section 3.3, along with Gaussian densities chosen to have the same standard deviation as in the data. Note that a normal distribution would feature zero skewness (symmetric) and a kurtosis of 3. The earnings growth distribution displays left (negative) skewness and excess kurtosis relative to a Gaussian density. In other words, workers face an earnings change distribution with a longer left tail relative to the right, and there are far more people with very small and very large changes and fewer people with intermediate changes. These qualitative properties are in line with findings for many other countries.

5.1 Higher-Order Moments for Male Earnings Growth


We now document the higher-order moments of earnings growth in Norway. Appendix C.1 contains analogous results for the U.S. as well as results for one-yearly earnings growth in Norway.¹⁶ For comparability with earlier work we focus on men. The results

¹⁴Güvenen *et al.* (2019) argue that in the commonly used random-walk permanent/transitory model, skewness is mainly driven by permanent changes. Moreover, as k increases, the variance and kurtosis of k -year log change $\Delta_{\log}^k y_t^i$ reflects more of the distribution of permanent innovations than that of transitory ones.

¹⁵A weakness with this approach is that transitory shocks will be present—albeit less pronounced—even in the changes at the five-year horizon. An alternative approach would have been to model transitory and permanent changes following the methodology of Arellano *et al.* (2017). We prefer the current descriptive approach because we find it more transparent.

¹⁶For comparability, Figures A.1 and A.2 (made with data from Güvenen *et al.* (2019)) and Figures A.3 and A.4 use identical sample selection criteria. In particular, the minimum imputed hours threshold is not imposed for this Norwegian sample.

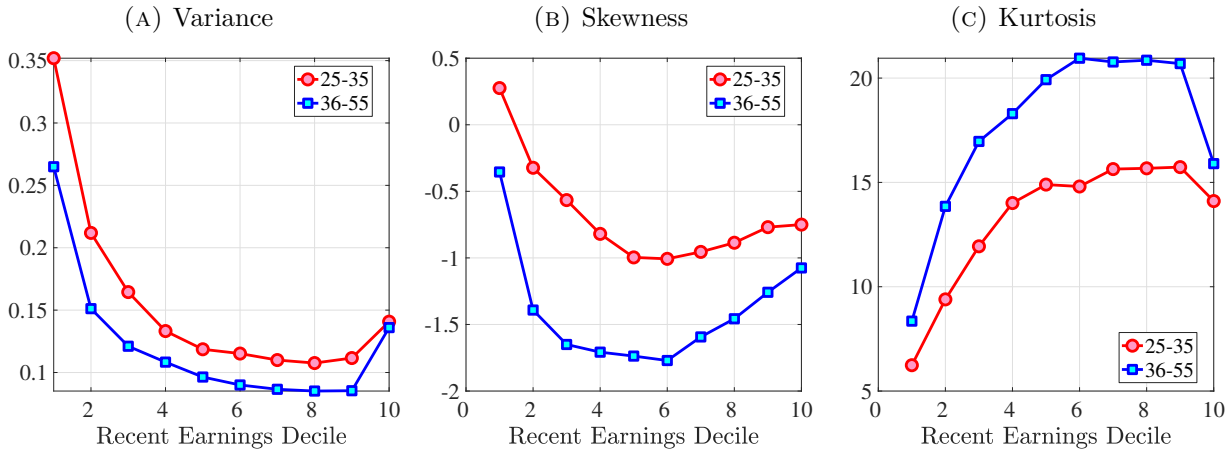
FIGURE 6 – Histograms of One- and Five-Year Log Earnings Changes



Note: The figure plots the empirical densities of one- and five-year earnings changes superimposed on Gaussian densities with the same variance. The data are for male workers in the base sample defined in Section 3 and $t = 2011$.

for women are reported in Appendix F.

FIGURE 7 – Cross-Sectional Moments for Five-Year Earnings Growth in Norway



Note: The figure displays the higher-order moments of five-year log earnings changes ($y_{t+5} - y_t$) for young males (red line) and prime-aged males (blue line) for each decile of RE.

Starting with the second moment, Figure 7 shows the variance of five-year earnings growth between t and $t + 5$ conditional on workers' age and RE in $t - 1$. Workers differ significantly in the dispersion of earnings growth they face with respect to their RE. In particular, for prime-aged workers, the variance declines monotonically from around 0.29 for the bottom decile of RE to roughly 0.09 for the 90th percentile, after which it increases to 0.14. The life-cycle variation is smaller than the differences across RE groups, with the variance of shocks being largest for the young workers (ages 25-35).

These patterns are qualitatively similar to those found for the U.S., albeit that overall earnings growth is less volatile in Norway (Figures A.1 and A.2).

Next, Figure 7 shows that almost all workers face a left-skewed distribution of five-year log earnings growth regardless of their RE and age, meaning that experiencing very large persistent declines in earnings is more likely than seeing very large increases. However, skewness—measured here as the third standardized moment—is more negative for prime-age workers.^{17,18} Thus, it seems that the older an individual gets or the higher his current earnings, the more gradual will be the upward movements and the more drastic will be the fall in earnings. This could be driven by a smaller scope for increases and a larger scope for declines.

Figure 7 plots the fourth standardized moment of five-year earnings growth by age and RE. This kurtosis measure increases monotonically from around 10 for the bottom earners up to around 35 for the 8th decline of RE in the middle-age group. That is, most high earners see even smaller earnings changes, and few experience very large ones. Moreover, kurtosis tends to increase with age for all RE groups, especially in the first 10 years of their careers. The RE and age variations in kurtosis of annual earnings growth in Norway data are similar to those documented for the U.S.

5.2 The Distributions of Wage and Hours Growth

In this section, we study the extent to which the distributions of hours and wage growth display non-Gaussian features and investigate the extent to which the higher-order moments for earnings are driven by changes in hours versus changes in hourly wages. We follow a similar graphical methodology as in Section 5.1 and study the cross-sectional moments of hours and wage growth conditional on 3 age groups and 10 deciles of recent earnings.

Higher-order moment decomposition We decompose higher-order moments of earnings growth into hours and wages using a statistical decomposition, which we lay out in the following lemma.

¹⁷Figure A.5 in Appendix C plots the percentile based skewness measure, Kelly’s skewness, by RE and age, which displays similar patterns.

¹⁸Earnings changes for women are significantly less negatively skewed. For example, for women younger than 45, earnings changes are symmetric regardless of recent earnings.

Lemma 1. *If x and y are two random variables, then*

$$\begin{aligned}
skew(x + y) &= \left(\frac{std(x)}{std(x + y)} \right)^3 \cdot skew(x) + \left(\frac{std(y)}{std(x + y)} \right)^3 \cdot skew(y) \\
&+ \underbrace{\frac{3}{(std(x + y))^3} (cov(x^2, y) + cov(x, y^2) - 2(E\{y\} + E\{x\}) \cdot cov(x, y))}_{co-skewness} \\
\\
kurt(x + y) &= \left(\frac{var(x)}{var(x + y)} \right)^2 kurt(x) + \left(\frac{var(y)}{var(x + y)} \right)^2 kurt(y) \\
&+ \frac{4}{(var(x+y))^2} \left[E \left\{ [x - E(x)]^3 [y - E(y)] \right\} + E \left\{ [x - E(x)] [y - E(y)]^3 \right\} \right] \\
&+ \underbrace{\frac{6}{(var(x+y))^2} E \left\{ [x - E(x)]^2 [y - E(y)]^2 \right\}}_{co-kurtosis}
\end{aligned}$$

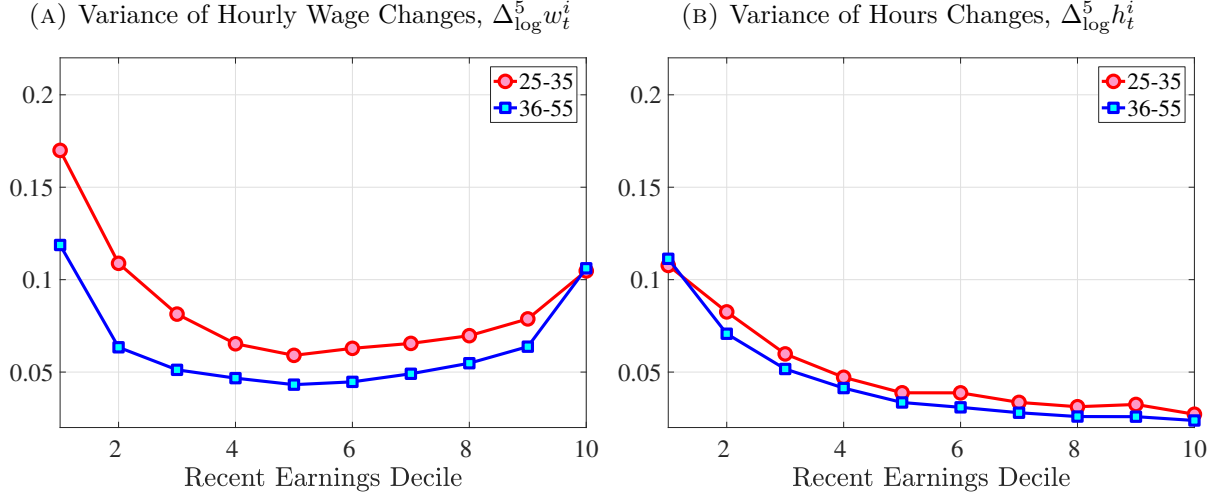
For the proof, see Appendix A.

The lemma shows that the skewness (kurtosis) of the sum of two random variables is equal to the weighted sum of skewness (kurtosis) of individual variables plus a co-skewness (co-kurtosis) term. The weights are determined by the ratio of the variance of individual variables to the variance of the sum. Thus, the more volatile variable will account for a larger share of the moments for the sum of the variables. A negative (positive) co-skewness indicates that both variables tend to undergo extreme negative (positive) deviations at the same time. Similarly, if two random variables exhibit a high level of co-kurtosis, they tend to undergo extreme deviations concurrently.

Second Moment: Variance

Figure 8 shows the variance of changes in wage rates (left panel) and hours (right panel), including how these variances change over the life cycle and across the recent earnings distribution. The age and income variations in the variance of both the wage and hours changes are qualitatively similar to those of annual earnings growth in Figure 7. In particular, the volatility of both hours and wage changes is falling with age. Moreover, lower-income workers tend to face a more dispersed distribution of wage rate and hours growth than higher-income workers do. The one exception is that top earners face a more volatile wage rate growth than workers with slightly lower recent earnings, whereas the volatility of hours worked is monotone decreasing in recent earnings RE. This suggests

FIGURE 8 – Variance of Five-Year Log Hourly Wage and Hours Growth



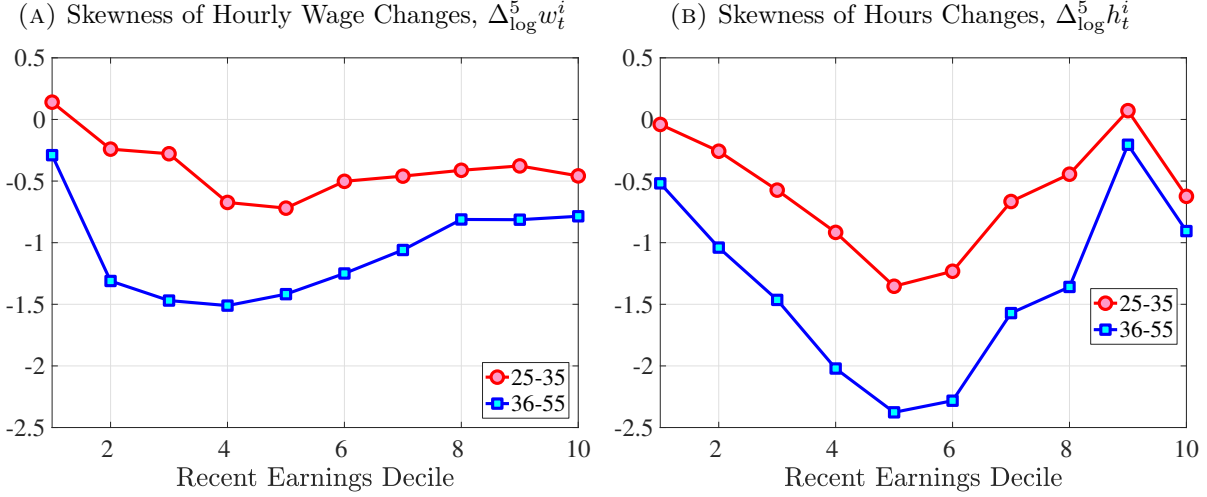
Note: The figure plots the variances of five-year wage rate changes (left panel) and hours changes (right panel) by RE decile for young men (red line) and prime-aged men (blue line).

that the increase in the U-shaped profile of the variance of earnings growth at the top end of the RE distribution in Figure 7 is due to the more volatile wage rate growth for the top earners. This is consistent with the view of [Parker and Vissing-Jørgensen \(2010\)](#), for example, that the earnings volatility of high earners is affected by a performance-based compensation structure such as bonuses and stock options.

For most groups of workers, wage rates are more volatile than hours. For example, the variance of wage changes varies from around 0.2 for the bottom young earners to 0.05 for upper-middle-income old workers, whereas the corresponding figures for hours are only around 0.10 for the bottom young earners and 0.02 for upper-middle-income old workers. This is consistent with the findings from the PSID (c.f., [Heathcote *et al.* \(2014\)](#)). It is also consistent with the evidence from Section 4.1 that earnings changes for high earners tend to be driven by changes in wage rates rather than hours.

The variance of earnings growth is higher than the variance of both earnings growth and wage growth. In Appendix C.3, we plot the decomposition of variance of five-year earnings growth to its wage and hours change components as well as their covariance (see Figure A.7). The variance of wage rate changes dominates that of hours changes. The small but positive covariance term also contributes to the variance of earnings growth.

FIGURE 9 – Skewness of Five-Year Log Hourly Wage and Hours Growth



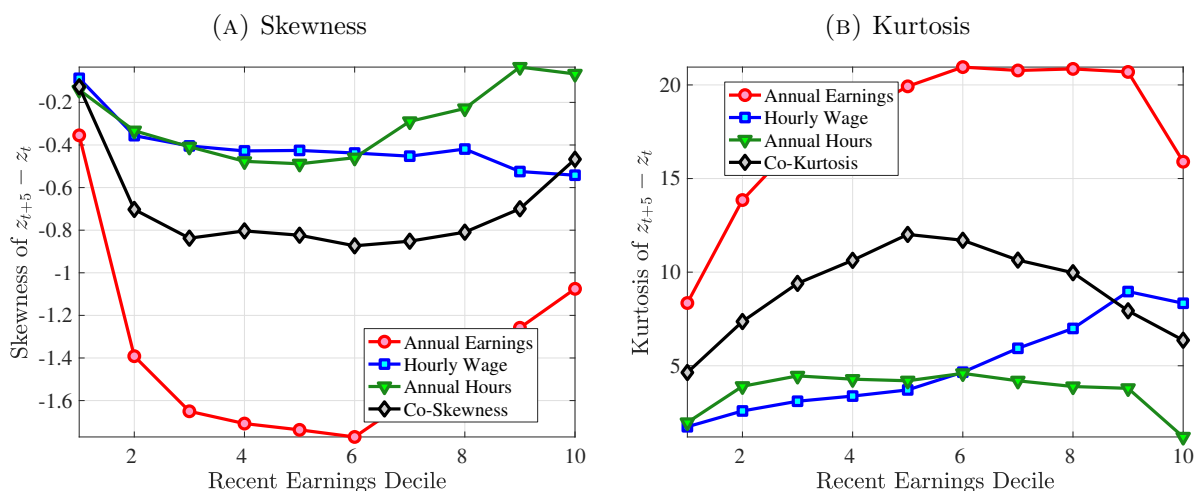
Note: The figure plots the skewness of five-year wage rate changes (left panel) and hours changes (right panel) by RE decile for young men (red line) and prime-aged men (blue line).

Third Moment: Skewness

Figure 9 documents the third moment of hours and wage rate growth. Again, the age and income variations in the skewness of these variables are qualitatively similar to those of annual earnings growth in Figure 7. First, skewness follows a U-shaped pattern over the RE distribution, with middle RE workers facing a more left-skewed distribution of five-year wage and hours changes. Second, the distributions of five-year growth in both wage rates and hours are more left (negatively) skewed for prime-age workers relative to younger workers. Note that hours changes are more negatively skewed than the wage rate changes.

We next decompose the skewness of the earnings growth distribution into hours and wage growth components as well as the co-skewness term as defined in Lemma 1 (the left panel of Figure 10). Recall that hours and wage growth contribute to the skewness of earnings growth according to the ratio of their variance to the variance of earnings growth. Thus, even though hours growth is more left skewed, the more volatile wage growth accounts for a larger share of the left skewness of earnings growth, especially above the median RE. More importantly, the main driver of the left skewness of the earnings changes is the co-skewness term. This term captures that hours and wages tend to simultaneously undergo large negative deviations and simultaneously experience gradual and small increases. Thus, when workers experience large hours cuts, they also see their wages decline sharply. This finding is consistent with the literature studying

FIGURE 10 – Decomposition of Skewness and Kurtosis of Five-Year Log Earnings Growth



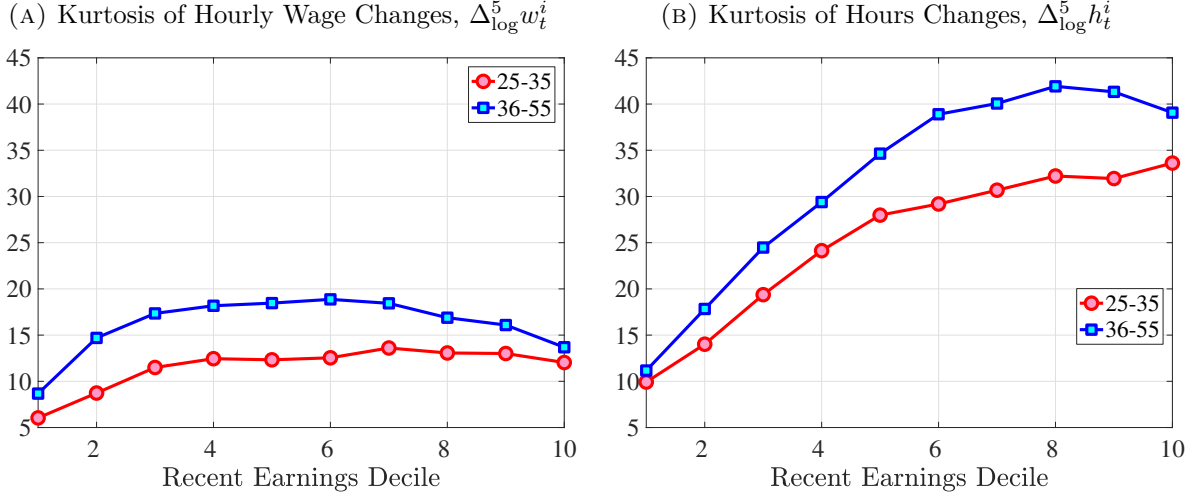
Note: The figure plots a decomposition of skewness (left panel) and kurtosis (right panel) of five-year log earnings changes (red line) for prime-aged men into the skewness/kurtosis of log wage changes (blue line), the skewness/kurtosis of log hours changes (green line), and the co-skewness/co-kurtosis between log wage and log hours changes (black line). Each dot represents a decile of RE. The decomposition is based on Lemma 1.

the labor market dynamics associated with unemployment, where large initial declines in hours and wage rates are followed by gradual recoveries involving long-lasting reductions (scarring effects) in employment and wages (e.g., [Jacobson *et al.* \(1993\)](#); [Von Wachter *et al.* \(2009\)](#), and [Huttunen *et al.* \(2011\)](#) for Norway).

Fourth Moment: Kurtosis

Finally, Figure 11 shows the income and age profiles of the fourth standardized moment of five-year wage (left panel) and hours growth (right panel). Both variables display very large excess kurtosis. The kurtosis of wage rate changes and how it varies with age and recent earnings groups are very similar to those of earnings growth (Figure 7). The kurtosis of hours growth is significantly higher than the kurtosis of wage and earnings growth. Thus, hours changes are less frequent but more extreme, so when they change, the changes are large. Furthermore, the kurtosis of hours changes also displays an increasing profile over the RE distribution, compared to a hump-shaped profile for the kurtosis of wage changes. As for the age variation, older workers are facing more leptokurtic hours, and wage growth distributions similar to the distribution of earnings changes. These features of hours changes are broadly consistent with transitions into or out of unemployment or part-time work for prime-age male workers, especially for those at the higher end of the RE distribution. Such flows feature infrequent but large changes in hours.

FIGURE 11 – Kurtosis of Five-Year Log Hourly Wage and Hours Growth



Note: The figure plots the kurtosis of five-year wage rate changes (left panel) and hours changes (right panel) by RE decile for young men (red line) and prime-aged men (blue line).

The right panel of Figure 10 decomposes the kurtosis of five-year earnings growth into the contributions from changes in work hours, changes in hourly wages, and the co-kurtosis term using Lemma 1. Up to the median RE, both the wage and the hours changes contribute equally to the kurtosis of earnings growth. For higher RE deciles, the wage rate growth becomes more important in accounting for the kurtosis of earnings changes, partly because wage changes are larger for these workers. However, the dominant driver of earnings kurtosis tends to be the co-kurtosis term (except for the top RE groups). Recall that this term is high if hours and wages tend to undergo large changes concurrently.


Taking Stock In this section, we have shown that both hours and wage changes are slightly negatively skewed but highly leptokurtic. Both hours and wage rates contribute to the higher-order moments of log earnings growth. However, the main culprit for the large negative skewness and kurtosis of earnings growth turns out to be the co-skewness and co-kurtosis terms, respectively. Thus, earnings growth is negatively skewed and leptokurtic, mainly because of the interaction between hours changes and wage rate changes.¹⁹

¹⁹Note that this conclusion may depend on the severity of measurement error. Appendix A.2 shows how the presence of classical measurement error would cause a downward bias in the estimates of skewness and kurtosis of a variable observed in the data.

5.3 Job Stayers and Switchers

Section 4.3 documented that a change of employer is a key event causing large earnings changes. A change of employer can happen either via an unemployment spell or through a direct job-to-job movement. In this section, we study the earnings growth distributions of job stayers and job switchers separately and quantify their role in driving the higher-order moments of earnings growth. To this end, we first identify job stayers and job switchers. We define a job switcher as an individual whose main employer is different between years t and $t + 1$, where the main employment is the job that accounts for the largest share of annual earnings. The rest of the population is composed of job stayers. In other words, a job stayer is a worker whose main employer has remained unchanged for two years in a row, and the job spell is contiguous. This classification of stayer workers is similar to the previous literature (Card *et al.* (2013)).²⁰

FIGURE 12 – Moments of One-Year Log Earnings Growth: Stayers vs. Switchers



Note: The figure displays the higher-order cross-sectional moments of one-year log earnings growth ($y_{t+1} - y_t$) for job switchers (blue line), job stayers (green line), and all prime-aged males (red line) for each RE decile.

Figure 12 shows the cross-sectional moments of one-year earnings growth for stayers and switchers separately. Annual earnings changes for switchers tend to be substantially more dispersed, more symmetric (less left skewed), and significantly less leptokurtic than those for stayers. Figure A.27 in the Appendix confirms that a similar pattern holds for women. The differences between stayers and switchers are qualitatively different from the findings in Guvenen *et al.* (2019) for the U.S. in that they find that annual earnings


²⁰Guvenen *et al.* (2019) impose a substantially tighter definition of a job stayer, namely, that the worker must have had some income in $t - 1$ and $t + 2$ from the same firm that was his main employer in periods t and $t + 1$.

growth for males is more symmetric for stayers than it is for switchers. One important contributor to strong negative skewness for stayers in Norway is sickness leave, as workers who receive some sickness benefits have more negative skewness than the stayers who do not experience sickness (details are available upon request). Note that by regulation these workers remain employed by the same firm during their sickness leave.

The Role of Stayers and Switchers in Higher Order Moments of Earnings Growth

How important are the switchers in driving the nonnormal features of annual earnings growth relative to the stayers? It turns out that the contributions of two mutually exclusive groups, such as job stayers and job switchers, to the cross-sectional skewness and kurtosis of earnings changes can be decomposed using two simple formulas that we state in the following lemma.

FIGURE 13 – Moment Decomposition of Earnings Growth: Stayers vs. Switchers



Note: The figure displays the contributions of job switchers (blue line) and job stayers (green line) to the population skewness and kurtosis of one-year earnings growth for prime-aged males (red line) for each RE decile.

Lemma 2. Suppose the sample S can be split into two mutually exclusive groups, $S = S_1 \cup S_2$ where $S_1 \cap S_2 = \emptyset$. The third central moment (skewness) can then be decomposed into components stemming from S_1 and S_2 ,

$$skew(y) = \underbrace{\frac{1}{(std(y))^3} \int_{\{i \in S_1\}} (y_i - E(y))^3 dF(y)}_{skewness\ due\ to\ S_1} + \underbrace{\frac{1}{(std(y))^3} \int_{\{i \in S_2\}} (y_i - E(y))^3 dF(y)}_{skewness\ due\ to\ S_2}.$$

The fourth central moment (kurtosis) can be decomposed into components stemming from S_1 and S_2 ,

$$kurt(y) = \underbrace{\frac{1}{(std(y))^4} \int_{\{i \in S_1\}} (y_i - E(y))^4 dF(y)}_{kurtosis \text{ due to } S_1} + \underbrace{\frac{1}{(std(y))^4} \int_{\{i \in S_2\}} (y_i - E(y))^4 dF(y)}_{kurtosis \text{ due to } S_2}.$$

In Figure 13, we use Lemma 2 to decompose the population skewness and kurtosis of one-year earnings growth into the contribution of job stayers and the contribution of job switchers.²¹ The negative skewness and high kurtosis of earnings growth is overwhelmingly driven by job stayers. The reason is twofold. First, the skewness is more negative and the kurtosis larger for stayers than for switchers. Second, there are very few switchers relative to the number of stayers. The expressions in Lemma 2 show that the contributions from each of the mutually exclusive subsets to the third and fourth centralized moments of their union are proportional to their population size.


5.4 Cross-Sectional Moments of Household Labor and Disposable Income Growth

Thus far, we have focused on the distribution of fluctuations in male labor earnings growth. However, for many economic questions, one is more interested in household disposable income after taxes and transfers, rather than individual male or female labor income before taxes. For example, for consumption and risk sharing, the relevant risk is that for disposable income. In this section, we investigate the cross-sectional moments of shocks to household labor income and household disposable income for married couples. Our measure of household disposable income is after-tax, after-transfer household labor income excluding capital income. Transfers include unemployment benefits, sickness benefits, paid parental leave, remuneration for participation in various government activity programs, disability benefits, public pensions, and other social welfare payments.

In Figure 14a we plot the variance of five-year income growth for individual earnings, household earnings, and household disposable income for males living as part of a couple (married and cohabiting men) by recent earnings decile (see Figure A.13 in the Appendix for the moments for one-year income growth measures). Consider first the difference between the variance of changes in individual male earnings (red line in Figure 14a) and

²¹We focus on one-year earnings growth when analyzing stayers versus switchers since this approach allows an unambiguous classification of switchers versus stayers.

FIGURE 14 – Moments of Five-Year Log Household Earnings and Disposable Income Growth



Note: The figure displays the higher-order cross-sectional moments of five-year log earnings changes (red line), log household earnings changes (blue line), and log disposable income changes (green line) for prime-aged males, plotted for each decile of RE. The sample comprises married and cohabiting men and their households.

changes in household earnings (blue line in Figure 14a). The fact that the variance of changes is somewhat lower for household earnings than for individual earnings for men below the median RE, and is approximately the same for men above the median, suggests that couples face slightly lower earnings risk than individual men do, although only for men below the median.

Consider next the variance of changes in disposable household income (green line in Figure 14a). These variances are substantially lower than for individual and household labor earnings. This suggests that taxes and transfers provide substantial insurance against individual earnings risk and that this holds true across the income distribution. For men below the median, the effect is especially large: the variance of changes to disposable income is reduced by more than half relative to household earnings, and for men in the lowest decile of RE who are part of a couple, the variance of income growth after taxes and transfers is less than a third of the variance of the growth of individual earnings. For low-income men, the main culprit for reducing the variance of changes is the generous Norwegian welfare benefits (unemployment insurance, disability insurance, social aid, and cash benefits to families with children).²² Our findings are in line with [Blundell *et al.* \(2014\)](#), who study how risk changes over the life cycle using the same

²²A more detailed analysis shows that the progressive tax system also plays a role, especially for men above the median. However, the generous transfers are the main driver of the big reduction in the variance of income growth relative to earnings growth.

administrative data from Norway. They estimate models with transitory and persistent income shocks and show that the variances of both transitory and persistent shocks fall substantially when going from individual earnings to household disposable income.²³

Consider now skewness. Figure 14b plots the skewness for five-year income growth for individual male earnings, household earnings, and household disposable income by RE decile. As can be seen from the figure, changes in household earnings are substantially less negatively skewed than male earnings changes. Note that this reduction in left skewness is not due to behavioral changes in spousal earnings. In Figure A.22 in the appendix, we show that spouses' earnings are remarkably unresponsive to changes in male earnings, for both positive and negative changes. Thus, we interpret the reduction in negative skewness as a mechanical second-earner effect. Figure 14b reveals that the effect of going to disposable income is even starker: five-year changes in disposable income have close to zero skewness.²⁴ In conclusion, both the tax and transfer system and spousal income contribute to removing the negative skewness of wage earnings growth, making five-year household disposable income growth close to symmetric.

In Figure 14c, we plot the kurtosis for five-year income growth for individual earnings, household earnings, and household disposable income for males living as part of a couple by RE decile. The kurtosis of household earnings growth is significantly lower than for individual labor income. For men above the median, the kurtosis is reduced by almost one half when going from individual labor income to household labor income. Note that the kurtosis of disposable income growth is higher than for household labor income growth. Disposable income growth is more leptokurtic than household earnings because of the nature of government transfers. The Norwegian social insurance system is designed to replace earnings shortfalls for full-time workers. This lowers the variance of disposable income changes. However, some infrequent changes, such as transitions between full-time and part-time work, are not insured and will trigger large changes in both individual earnings and disposable income (see also Table I). Thus, the welfare state magnifies the tendency to see frequent small changes and infrequent large changes, inducing higher kurtosis.

Overall, these results show that the higher-order moments for disposable income

²³They also calculate *individual* disposable income and find that the variances of shocks to individual disposable income are just slightly larger than those of household disposable income.

²⁴Figure A.13 in the appendix shows that the skewness of one-year disposable income changes is even closer to zero and is actually positive for men in the eighth decile.

growth and, to a smaller extent, household earnings growth differ sharply from those of male earnings growth. In particular, disposable income is essentially symmetrically distributed. Moreover, changes in disposable income have substantially lower variance than changes in male earnings.²⁵ However, the kurtosis for disposable income growth is comparable to the kurtosis of individual earnings growth: slightly larger for poorer households and lower for the richest ones.

6 Conclusion

This paper studies the drivers behind two important deviations from standard linear and symmetric models of labor income risk, namely, asymmetric mean reversion of earnings changes—large negative changes are less persistent than positive changes—and more negative skewness and higher kurtosis relative to a Gaussian distribution. Using Norwegian administrative register data and labor survey data, we decompose earnings into hours and hourly wages and examine the extent to which the nonlinearities in mean reversion and non-Gaussian higher-order moments are caused by hours or hourly wages.

We find that, first, individual labor income dynamics in Norway are remarkably similar to their counterparts in U.S. data, in terms of both nonlinear persistence and higher-order moments. Second, we find that the nonlinear mean reversion in earnings is, to a large extent, driven by the dynamics of hours worked, whereas wage rate dynamics play a less important role. Indeed, hours dynamics are very nonlinear: negative changes are relatively transient, whereas positive changes are close to permanent. In contrast, wage rate dynamics are close to being linear, with both positive and negative changes being highly persistent. These findings are based on following groups of people who experience a similar initial change, and this approach captures the effects working through both the intensive and extensive margins of labor supply.

Third, we find that both wage rates and hours worked contribute to the negative skewness and high kurtosis of individual earnings changes. However, the interaction between hours and wages—captured by the co-skewness and co-kurtosis terms—is quantitatively most important. Finally, the Norwegian register data allows us to identify individuals in households, something that is not possible using U.S. administrative data for example. We show that when considering household earnings and disposable household income,

²⁵Our findings for individual versus household earnings are in line with [Pruitt and Turner \(forthcoming\)](#). Using U.S. administrative tax records, they find that total household earnings growth has lower variance and less negative skewness compared to individual males' earnings growth.

the deviations from normality are mitigated relative to the higher-order moments for male labor earnings. In fact, changes in disposable household income are close to being symmetric.

While our study is based on data from Norway, we believe the findings also have general validity for other countries, including the U.S. The fact that earnings dynamics for Norway and the U.S. are quantitatively and qualitatively similar despite the differences in labor market institutions across these countries suggests that the Norwegian and U.S. earnings dynamics may be driven by similar economic mechanisms.

References

- ABOWD, J. M. and CARD, D. (1989). On the covariance structure of earnings and hours changes. *Econometrica*, **57** (2), 411–45.
- ALTONJI, J. G. (1986). Intertemporal substitution in labor supply: Evidence from micro data. *Journal of Political Economy*, **94** (3), S176–S215.
- ARELLANO, M., BLUNDELL, R. and BONHOMME, S. (2017). Earnings and consumption dynamics: a nonlinear panel data framework. *Econometrica*, **85** (3), 693–734.
- ATTANASIO, O. and DAVIS, S. J. (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy*, **104** (6), 1227–62.
- BLUNDELL, R., GRABER, M. and MOGSTAD, M. (2014). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics*.
- , PISTAFERRI, L. and PRESTON, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, **98** (5), 1887–1921.
- and PRESTON, I. (1998). Consumption inequality and income uncertainty. *The Quarterly Journal of Economics*, **113** (2), 603–640.
- BUSCH, C., DOMEIJ, D., GUVENEN, F. and MADERA, R. (2018). Asymmetric business-cycle risk and social insurance. *Working paper*.
- CARD, D., HEINING, J. and KLINE, P. (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics*, **128** (3), 967–1015.
- COCHRANE, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy*, **99** (5), 957–76.
- CONESA, J. C., KITAO, S. and KRUEGER, D. (2009). Taxing capital? not a bad idea after all! *American Economic Review*, **99** (1), 25–48.
- CUTLER, D. M. and KATZ, L. F. (1992). Rising inequality? changes in the distribution of income and consumption in the 1980’s. *American Economic Review*, **82** (2), 546–51.
- DE NARDI, M., FELLA, G., KNOEF, M., OOIJEN, R. V. and PARDO, G. P. (2019). *Family and Government Insurance: Wage, Earnings, and Income Risks in the Netherlands and the U.S.* Working paper.
- DEATON, A. and PAXSON, C. (1994). Intertemporal choice and inequality. *Journal of Political Economy*, **102** (3), 437–67.
- GUVENEN, F., KARAHAN, F. and OZKAN, S. (2018). *Inequality over the Life Cycle in the Presence of Higher-Order Risk*. Working paper, University of Minnesota.
- , —, — and SONG, J. (2019). *What Do Data on Millions of U.S. Workers Say About Labor Income Risk?* Working Paper 20913, National Bureau of Economic Research.
- , OZKAN, S. and SONG, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, **122** (3), 621–660.

- HEATHCOTE, J., PERRI, F. and VIOLANTE, G. L. (2010a). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967-2006. *Review of Economic Dynamics*, **13** (1), 15–51.
- , STORESLETTEN, K. and VIOLANTE, G. L. (2010b). The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy*, **118** (4), 681–722.
- , — and — (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*.
- HUTTUNEN, K., MØEN, J. and SALVANES, K. G. (2011). How destructive is creative destruction? effects of job loss on job mobility, withdrawal and income. *Journal of the European Economic Association*, **9** (5), 840–870.
- JACOBSON, L. S., LALONDE, R. J. and SULLIVAN, D. G. (1993). Earnings losses of displaced workers. *American Economic Review*, **83** (4), 685–709.
- KARAHAN, F., OZKAN, S. and SONG, J. (2019). Anatomy of lifetime earnings inequality: Heterogeneity in job ladder risk vs. human capital. *FRB of New York Staff Report*, (908).
- KRUEGER, D. and PERRI, F. (2005). Does income inequality lead to consumption inequality? evidence and theory. *Review of Economic Studies*, **73** (1), 163–193.
- KRUSELL, P., MUKOYAMA, T., ROGERSON, R. and ŞAHİN, A. (2011). A three state model of worker flows in general equilibrium. *Journal of Economic Theory*, **146** (3), 1107–1133.
- LILLARD, L. A. and WEISS, Y. (1979). Components of variation in panel earnings data: American scientists 1960-70. *Econometrica*, **47** (2), 437–454.
- and WILLIS, R. J. (1978). Dynamic aspects of earnings mobility. *Econometrica*, **46** (5), 985–1012.
- MACURDY, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of political Economy*, **89** (6), 1059–1085.
- (1982). The Use of Time Series Process to Model the Error Structure of Earnings in a Longitudinal Data Analysis. *Journal of Econometrics*, **18**, 83–114.
- PARKER, J. A. and VISSING-JØRGENSEN, A. (2010). The increase in income cyclicality of high-income households and its relation to the rise in top income shares. *Brookings Papers on Economic Activity*, **41** (2), 1–70.
- PRUITT, S. and TURNER, N. (2018). The nature of household labor income risk. *Working paper*.
- and — (forthcoming). Earnings risk in the household: Evidence from millions of us tax returns. *American Economic Review: Insights*.
- VON WACHTER, T., SONG, J. and MANCHESTER, J. (2009). *Long-Term Earnings Losses due to Mass Layoffs During the 1982 Recession: An Analysis Using U.S. Administrative Data from 1974 to 2004*. Tech. rep., Columbia University.

Appendices

A Derivations

A.1 Proof of Lemma 1

The formula for skewness is:

$$skew(x) = \frac{1}{(std(x))^3} E(x - E(x))^3.$$

Consider the skewness of a sum of stochastic variables:

$$\begin{aligned} skew(x+y) &= \frac{1}{(std(x+y))^3} E(x+y - E(x+y))^3 \\ &= \frac{1}{(std(x+y))^3} E(x - E(x) + y - E(y))^3 \\ &= \frac{1}{(std(x+y))^3} E\left\{[x - E(x)]^3 + [y - E(y)]^3 + 3[x - E(x)]^2[y - E(y)] + 3[x - E(x)][y - E(y)]^2\right\} \\ &= \left(\frac{std(x)}{std(x+y)}\right)^3 \frac{E\left\{[x - E(x)]^3\right\}}{(std(x))^3} + \left(\frac{std(y)}{std(x+y)}\right)^3 \frac{E\left\{[y - E(y)]^3\right\}}{(std(y))^3} \\ &\quad + \frac{3}{(std(x+y))^3} (cov(x^2, y) - 2E\{x\} \cdot cov(x, y) + cov(x, y^2) - 2E\{y\} \cdot cov(x, y)) \\ &= \left(\frac{std(x)}{std(x+y)}\right)^3 \cdot skew(x) + \left(\frac{std(y)}{std(x+y)}\right)^3 \cdot skew(y) \\ &\quad + \frac{3}{(std(x+y))^3} (cov(x^2, y) + cov(x, y^2) - 2(E\{y\} + E\{x\}) \cdot cov(x, y)) \end{aligned}$$

where the second to last equality is due to

$$\begin{aligned} &E\left((x - E\{x\})^2 (y - E\{y\})\right) \\ &= E(x^2 y) - E\{y\}E(x^2) - 2E\{x\}(E(xy) - E\{y\}E\{x\}) + E(E\{x\}^2 (y - E\{y\})) \\ &= E(x^2 y) - E\{y\}E(x^2) - 2E\{x\} \cdot cov(x, y) \\ &= cov(x^2, y) - 2E\{x\} \cdot cov(x, y) \end{aligned}$$

The formula for the fourth central moment, kurtosis, is:

$$kurt(x) = \frac{1}{(std(x))^4} E(x - E(x))^4.$$

Consider now the kurtosis of a sum of stochastic variables and keeping in mind that: $(x+y)^4 =$

$x^4 + 4x^3y + 6x^2y^2 + 4xy^3 + y^4$ we have:

$$\begin{aligned}
\text{kurt}(x+y) &= \frac{1}{(\text{std}(x+y))^4} E(x+y - E(x+y))^4 \\
&= \frac{1}{(\text{var}(x+y))^2} E(x - E(x) + y - E(y))^4 \\
&= \frac{1}{(\text{var}(x+y))^2} E \left\{ [x - E(x)]^4 + [y - E(y)]^4 + 4[x - E(x)]^3 [y - E(y)] + 4[x - E(x)] [y - E(y)]^3 \right\} \\
&\quad + \frac{1}{(\text{var}(x+y))^2} E \left\{ 6[x - E(x)]^2 [y - E(y)]^2 \right\} \\
&= \left(\frac{\text{var}(x)}{\text{var}(x+y)} \right)^2 \text{kurt}(x) + \left(\frac{\text{var}(y)}{\text{var}(x+y)} \right)^2 \text{kurt}(y) + X,
\end{aligned}$$

where X captures a set of co-variance and co-skewness terms,

$$X \equiv \frac{2}{(\text{var}(x+y))^2} \left[2E \left\{ [x - E(x)]^3 [y - E(y)] \right\} + 2E \left\{ [x - E(x)] [y - E(y)]^3 \right\} + 3E \left\{ [x - E(x)]^2 [y - E(y)]^2 \right\} \right].$$

A.2 The Effect of Classical Measurement Error in Hours on Estimates of Higher Order Moments

How does a possible measurement error affect the estimates of higher order moments?

Even though our imputation methodology does a fairly good job in producing a better measure of hours than the one available in the register data, it is inevitable that there may still be some measurement error. How would measurement error affect our empirical findings? In particular, how are the estimates of the higher order moments of hours and wage changes biased due to the measurement error? To address this issue, consider a simple model of measurement error, $\hat{z} = z + \hat{\varepsilon}$, where \hat{z} denotes the measure of a variable observed in the data, $\hat{\varepsilon}$ denotes measurement error, and z is the true value of the variable. The following lemma establishes how measurement error influences the measured moments.

Lemma 3.

Assume that measurement error $\hat{\varepsilon}$ is independent of the true variable z . Then, the estimates for the skewness and kurtosis of the measured variable \hat{z} are given by:

$$\begin{aligned}
\text{skew}(\hat{z}) &= \left(\frac{\text{var}(z)}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^{\frac{3}{2}} \text{skew}(z) + \left(\frac{\text{var}(\hat{\varepsilon})}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^{\frac{3}{2}} \text{skew}(\hat{\varepsilon}) \\
\text{kurt}(\hat{z}) &= \left(\frac{\text{var}(z)}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^2 \text{kurt}(z) + \left(\frac{\text{var}(\hat{\varepsilon})}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^2 \text{kurt}(\hat{\varepsilon}) + \frac{6 \cdot \text{var}(\hat{\varepsilon}) \cdot \text{var}(z)}{(\text{var}(z) + \text{var}(\hat{\varepsilon}))^2}
\end{aligned}$$

where \hat{z} *skew* and *kurt* are the third and fourth centralized moments for skewness and excess kurtosis, respectively.

Proof. See sections [A.2.1](#) and [A.2.2](#) below. □

Suppose that measurement error is classical, i.e., that it has a symmetric distribution with $skew(\hat{\varepsilon}) = 0$ and $kurt(\hat{\varepsilon}) > 0$. In this case, the lemma shows that the skewness of the measured hours and wage growth would be biased towards zero relative to the true skewness (see Section A.2). For kurtosis the bias could go either way. Thus, we believe that our findings in this section are upper bounds for the true skewness of hours and wage changes.

Let z and \hat{z} denote the true value of the random variable z and the measured value of z , respectively. Assume that the measurement error in z is classical,

$$\hat{z} = z + \hat{\varepsilon}.$$

and this measurement error is inherited in the empirical moments:

$$var(\hat{z}) = var(z) + var(\hat{\varepsilon})$$

A.2.1 Skewness

Consider now the third moment (skewness),

$$\begin{aligned} skew(\hat{z}) &= \frac{1}{(\sigma(\hat{z}))^3} E(\hat{z} - E(\hat{z}))^3 = \frac{1}{(var(z) + var(\hat{\varepsilon}))^{\frac{3}{2}}} E(z + \hat{\varepsilon} - E(z))^3 \\ &= \frac{1}{(var(z) + var(\hat{\varepsilon}))^{\frac{3}{2}}} E(\hat{\varepsilon}^3 + 3\hat{\varepsilon}^2(z - \bar{z}) + 3\hat{\varepsilon}(z - \bar{z})^2 + (z - \bar{z})^3) \\ &= \frac{E\hat{\varepsilon}^3 + E(z - \bar{z})^3}{(var(z) + var(\hat{\varepsilon}))^{\frac{3}{2}}} \\ &= \left(\frac{var(z)}{var(z) + var(\hat{\varepsilon})}\right)^{\frac{3}{2}} \frac{E(z - \bar{z})^3}{(var(z))^{\frac{3}{2}}} + \left(\frac{var(\hat{\varepsilon})}{var(z) + var(\hat{\varepsilon})}\right)^{\frac{3}{2}} \frac{E\hat{\varepsilon}^3}{(var(\hat{\varepsilon}))^{\frac{3}{2}}} \\ &= \left(\frac{var(z)}{var(z) + var(\hat{\varepsilon})}\right)^{\frac{3}{2}} \cdot S(z) + \left(\frac{var(\hat{\varepsilon})}{var(z) + var(\hat{\varepsilon})}\right)^{\frac{3}{2}} \cdot S(\hat{\varepsilon}). \end{aligned}$$

Thus, the measured skewness is a weighted sum of skewness of the true z , ($S(z)$), and skewness of measurement error, ($S(\hat{\varepsilon})$), where the weights do not sum to unity. It follows that the true skewness of z is given by:

$$S(z) = \left(\frac{var(z) + var(\hat{\varepsilon})}{var(z)}\right)^{\frac{3}{2}} S(\hat{z}) - \left(\frac{var(\hat{\varepsilon})}{var(z)}\right)^{\frac{3}{2}} \cdot S(\hat{\varepsilon})$$

Thus if we assume that the measurement error is normally distributed, then the true skewness is more pronounced than the the measured skewness:

$$S(z) = \left(\frac{var(z) + var(\hat{\varepsilon})}{var(z)}\right)^{\frac{3}{2}} S(\hat{z})$$

A.2.2 Kurtosis

Finally, consider the fourth moment,

$$\text{kurt}(z) \equiv \frac{1}{(\sigma(z))^4} E(z - E(z))^4$$

This implies that measured kurtosis of z can be expressed as

$$\begin{aligned} \text{kurt}(\hat{z}) &= \frac{1}{(\sigma(\hat{z}))^4} E(\hat{z} - E(\hat{z}))^4 = \frac{1}{(\text{var}(z) + \text{var}(\hat{\varepsilon}))^2} E(z + \hat{\varepsilon} - E(z))^4 \\ &= \frac{1}{(\text{var}(z) + \text{var}(\hat{\varepsilon}))^2} E(\hat{\varepsilon}^4 + 4(\hat{\varepsilon})^3(z - \bar{z}) + 6(\hat{\varepsilon})^2(z - \bar{z})^2 + 4\hat{\varepsilon}(z - \bar{z})^3 + (z - \bar{z})^4) \\ &= \frac{(E\hat{\varepsilon}^4 + 6\text{var}(\hat{\varepsilon}) \cdot \text{var}(z) + E(z - \bar{z})^4)}{(\text{var}(z) + \text{var}(\hat{\varepsilon}))^2} \\ &= \left(\frac{\text{var}(\hat{\varepsilon})}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^2 \frac{E\hat{\varepsilon}^4}{(\text{var}(\hat{\varepsilon}))^2} + \frac{6\text{var}(\hat{\varepsilon}) \cdot \text{var}(z)}{(\text{var}(z) + \text{var}(\hat{\varepsilon}))^2} + \left(\frac{\text{var}(z)}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^2 \frac{E(z - \bar{z})^4}{(\text{var}(z))^2} \\ &= \left(\frac{\text{var}(z)}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^2 \cdot \text{kurt}(z) + \left(\frac{\text{var}(\hat{\varepsilon})}{\text{var}(z) + \text{var}(\hat{\varepsilon})} \right)^2 \cdot \text{kurt}(\hat{\varepsilon}) + \frac{6 \cdot \text{var}(\hat{\varepsilon}) \cdot \text{var}(z)}{(\text{var}(z) + \text{var}(\hat{\varepsilon}))^2} \end{aligned}$$

It follows that the true kurtosis of z is given by

$$\text{kurt}(z) = \left(1 + \frac{\text{var}(\hat{\varepsilon})}{\text{var}(z)} \right)^2 \text{kurt}(\hat{z}) - \left(\frac{\text{var}(\hat{\varepsilon})}{\text{var}(z)} \right)^2 \cdot \text{kurt}(\hat{\varepsilon}) - 6 \cdot \frac{\text{var}(\hat{\varepsilon})}{\text{var}(z)}$$

Then the true excess kurtosis of z is given by

$$\text{excess kurt}(z) = \left(1 + \frac{\text{var}(\hat{\varepsilon})}{\text{var}(z)} \right)^2 \text{excess kurt}(\hat{z}) - \left(\frac{\text{var}(\hat{\varepsilon})}{\text{var}(z)} \right)^2 \cdot \text{excess kurt}(\hat{\varepsilon})$$

Again if we assume that measurement error is normally distributed then the true excess kurtosis is more pronounced than the measured excess kurtosis:

$$\text{excess kurt}(z) = \left(1 + \frac{\text{var}(\hat{\varepsilon})}{\text{var}(z)} \right)^2 \text{excess kurt}(\hat{z})$$

B Imputing Hours in the Register Data

Due to the shortcomings of the contracted hours measure in the register data, we design an imputation approach based on the Norwegian Labor Force Survey to obtain a better measure of hours worked. This data set contains high quality survey data on actual hours worked but has a limited sample size. All individuals present in the Labor Force Survey are also present in the register data. We merge the two data sources using individual identification numbers and design a novel imputation approach to infer actual hours worked for the entire population.

The Labor Force Survey records weekly hours worked. Each individual is surveyed up to eight consecutive quarters. We use only those who are present in all eight quarters and impute actual annual hours in year t as: $h_t^{LFS} = 13 \cdot \sum_{q=1}^4 h_{t,q}^{LFS}$, where $h_{t,d}^{LFS}$ is weekly hours in quarter q of year t . We then regress actual annual hours h_{it}^{LFS} from the Labor Force Survey on information in the register data:

$$h_{it}^{LFS} = \alpha h_{it}^{REG} + \beta X_{it} + \epsilon_{it},$$

where h_{it}^{REG} is contractual hours reported according to the Employment Register and X_{it} contains a rich set of observables from the register data: sickness days, parental leave days, unemployment days, part time, sector, labor earnings, country of origin, and education. We estimate the model separately for men and women and for each recent earnings quintile. Tables A.1 and A.2 below contain the regression results for males and females. The estimated coefficients are used to impute actual work hours for the individuals that are not present in the Labor Force Survey. We add bootstrapped errors to the imputed hours, using the approach of re-sampling residuals from the original regression. Residuals are clustered by gender and recent earnings, and then drawn randomly within these bins to match the imputed hours based on the whole register population.

Using Lasso to Determine the Variables that Contribute Most in the OLS

To get a feeling for which of the covariates in tables A.1 and A.2 that have more explanatory power we use Lasso. The lasso minimizes the residual sum of squares (RSS) subject to a constraint on the absolute size of coefficient estimates. The lasso shrinks each (standardized) coefficient by a constant factor, truncating at zero. This means that the stronger the restriction, the more coefficients are set to zero. Thus, by varying the constraint, we can obtain a list of predictors that contribute the most. For ease of exposition we report results for the most important variables for all men and all women, respectively, without conditioning on recent earnings group.

Males			
	Added to constant:	R-sq	L1-Norm
1	Log earnings	0.023	32.2
2	Register hours/Days on sickness leave	0.103	164.6
3	Days unemployed	0.143	188.2
4	Days on labor market program	0.158	196.4
5	Public sector/College or university	0.190	219.1
Females			
	Added to constant:	R-sq	L1-Norm
1	Log earnings	0.061	47.1
2	Register hours	0.245	222.6
3	Days parental leave	0.365	328.2
4	Days on sick leave	0.372	332.0
5	Days unemployed	0.410	328.7

The two most important covariates are contractual hours and labor earnings. The number of days receiving benefits for sickness, parental leave, and unemployment are also important predictive variables. This is not surprising since days on benefits are very accurately measured—it is based on actual benefit payments—while the employer-reported contractual hours in the Employment register often miss such benefits spells. The estimations show that our model has a larger explanatory power for women than for men. This is partly due to contractual hours and earnings having a higher correlation with actual hours worked for women than for men.

The explanatory power of our imputation is relatively high, as measured in terms of overall R-squared: about 0.19 for men and 0.41 for women. This is comparable to the explanatory power of Mincer-type linear regressions on data from the PSID. Regressions on these data with annual hours worked as the dependent variable and standard covariates as explanatory variables (gender, education, a quartic in age, and, most importantly, annual earnings), yield an overall R-squared of 0.45 for women and 0.16 for men (see the Online Appendix D for details).²⁶ We also evaluate the quality of our imputation by considering out-of-sample predictions. More precisely, we first estimate the model on a random half of the sample and use the estimates to predict hours for the second half of the sample. We find that adjusted root mean square errors (RMSE) are similar for the two samples, only slightly higher out-of-sample, with a difference that is not statistically significant.

²⁶If earnings is dropped as an explanatory variable, the R-squared falls substantially for the PSID, while it remains high for our Norwegian register data, due to detailed information on days receiving unemployment and sickness benefits.

TABLE A.1 – Regression Results for Actual Hours (Males)

Variable	Recent Earnings Quintile				
	1	2	3	4	5
Register hours	0.152**	0.049	0.096**	0.087*	0.296**
Log earnings	340.4**	329.6**	412.6**	249.5**	205.4**
Days parental leave	-2.86*	-2.37*	-1.80*	-4.48**	-4.90**
Days on sick leave	-3.40**	-3.38**	-2.53**	-4.06**	-5.34**
Days unemployed	-2.56**	-2.98**	-2.52**	-3.39**	-2.71**
Days on labor market program	-1.76**	-3.13**	-3.05**	-2.14	-4.41*
Part time dummy	-9.93	17.1	-6.36	-36.4	-35.4
Public sector dummy	-86.5*	-93.3**	-49.2**	-64.7**	70.9*
Age	13.8	-18.8*	7.65	-10.3	-1.79
(Age) ²	-0.143	0.198	-0.093	0.101	0.022
Educational level (ref = primary)	0	0	0	0	0
Upper secondary	6.94	-27.9	-40.8	-22.6	-114.5
College/university	-21.1	-107.3	-101.0	-76.2	-166.8
Year (ref = 2003)	0	0	0	0	0
2004	-55.7	-21.9	-14.9	-39.0	24.4
2005	10.6	-26.3	12.8	-12.5	53.1
2006	13.7	-34.3	-22.0	-32.3	15.4
2007	-42.6	-71.7*	-35.9	-43.1	40.2
2008	-22.2	-75.3*	-40.1	-73.9*	-18.0
2009	-121.8*	-90.6*	-138.3**	-107.2**	-51.2
2010	-78.2*	-81.8*	-54.5*	-74.4*	-47.7
2011	-59.8	-82.5*	-43.1	-56.9	-33.8
2012	-23.6	-67.1	-25.2	-12.9	17.3
2013	-79.5*	-77.3*	-66.2	-28.9	21.3
2014	-115.3*	-99.1*	-74.4*	-39.5	14.9
Country of origin (ref = Norway)	0	0	0	0	0
Nordic countries	118.6	-62.9	24.8	-25.6	50.2
Poland	-113.5	-5.2	71.7	-1041.0**	-167.9
Western Europe, except Nordic	112.1	-54.1	-39.1	-35.4	30.1
EU countries in Eastern Europe	-247.2	-66.3	67.1	-102.3	-259.4
Russia and non-EU Eastern Europe	-126.8	-52.5	-44.6	-275.8*	92.8
North America and Oceania	-55.6	-20.4	94.1	-85.7	15.1
Africa	-65.9	-25.0	175.8	-281.7*	74.8
Asia	-57.7	-39.3	-152.9*	-40.6	-4.2
South and Central America	15.1	73.4	-114.5	37.9	-162.9
Constant	-3019	-2003	-3775	-1283	-1420
R - squared	0.227	0.163	0.113	0.104	0.083
No. of observations	4 621	5 269	5 813	5 751	5 508

* p < 0.05, ** p < 0.01

Note: The table displays the coefficients from the regression of male hours from the Norwegian Labor Force Survey on observables in the register data, equation 1 in the main text. We split the sample in five recent earnings quintiles and report the results by quintile.

TABLE A.2 – Regression Results for Actual Hours (Females)

Variable	Recent Earnings Quintile				
	1	2	3	4	5
Register hours	0.243**	0.224**	0.293**	0.221**	0.148**
Log earnings	406.9**	509.0**	401.4**	354.2**	214.8**
Days parental leave	-2.14**	-1.40**	-2.28**	-3.25**	-4.66**
Days on sick leave	-1.53**	-1.03**	-2.27**	-2.52**	-3.68**
Days unemployed	-1.17**	-1.07**	-1.58**	-1.54**	-2.76**
Days on labor program	-0.867**	-0.848**	-0.963**	-0.224	-3.33**
Part time dummy	-14.2	43.5*	38.3	23.3	1.50
Public sector dummy	-45.3**	-59.0**	-36.1**	-34.7*	-21.9
Age	-21.3	-6.33	5.79	-11.4	-24.3*
(Age) ²	0.238	0.638	-0.077	0.131	0.281
Educational level (ref = primary)	0	0	0	0	0
Upper secondary	-7.01	-14.2	-24.2	1.58	-57.3
College/university	-17.3	-121.8	-131.9	-72.16	-88.0
Year (ref = 2003)	0	0	0	0	0
2004	-30.7	-18.3	-60.5	-31.4	-20.9
2005	-25.1	8.97	-16.3	20.3	5.15
2006	-33.3	-19.5	-52.2*	-41.2	-4.51
2007	-52.6*	-31.4	-64.1	-71.2*	-21.8
2008	-11.8	-27.2	-28.2	-52.2*	-23.9
2009	-61.5*	-64.5*	-100.6*	-95.3*	-46.9
2010	-3.63	-52.7*	-54.7*	-61.8*	-34.7
2011	12.4	-34.9	-52.3*	-59.9*	-29.4
2012	-17.3	-35.1	-43.6	-67.3*	-13.8
2013	-20.5	-48.0	-58.2*	-77.2*	12.3
2014	-38.8	-53.4*	-77.3*	-90.8*	-6.7
Country of origin (ref = Norway)	0	0	0	0	0
Nordic countries	-53.8	99.9	-14.6	-1.34	-51.6
Poland	-29.7	-64.6	48.1	-117.8	-528.9*
Western Europe, except Nordic	312.8**	18.4	30.8	-39.3	37.1
EU countries in Eastern Europe	-69.5	43.4	8.09	-8.87	-334.3
Russia and non-EU Eastern Europe	147.7*	41.3	7.22	-71.4	-157.2
North America and Oceania	-183.8	-136.9	-4.28	-86.5	45.9
Africa	-39.1	11.3	37.3	-394.4*	93.6
Asia	19.5	32.5	7.10	9.23	-2.29
South and Central America	75.5	116.3	14.8	201.5	163.8
Constant	-3640	-5106	-4078	-2848	-844.2
R - squared	0.366	0.354	0.392	0.336	0.297
No- of observations	4 875	4 936	5 322	5 491	5 355

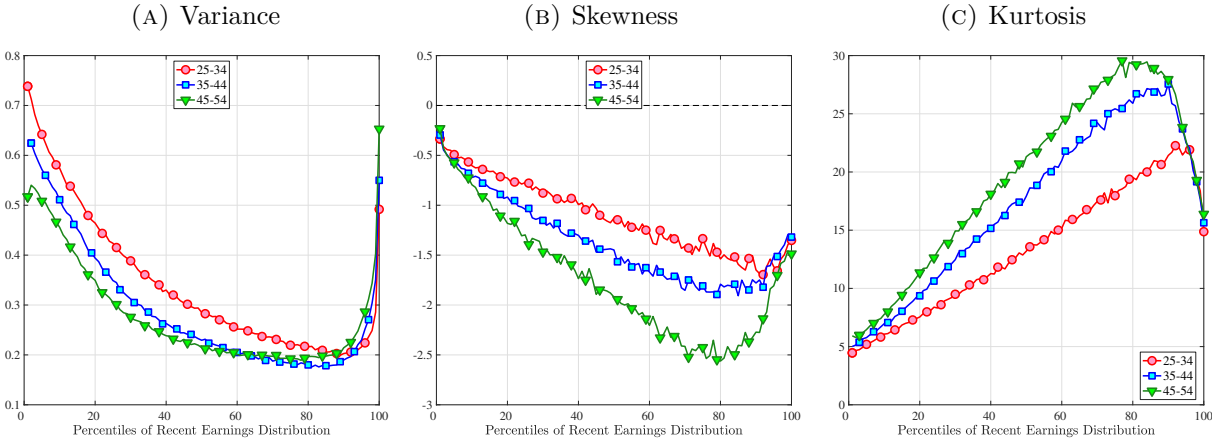
* p < 0.05, ** p < 0.01

Note: The table displays the coefficients from the regression of female hours from the Norwegian Labor Force Survey on observables in the register data, equation 1 in the main text. We split the sample in five recent earnings quintiles and report the results by quintile.

C Additional Figures and Tables

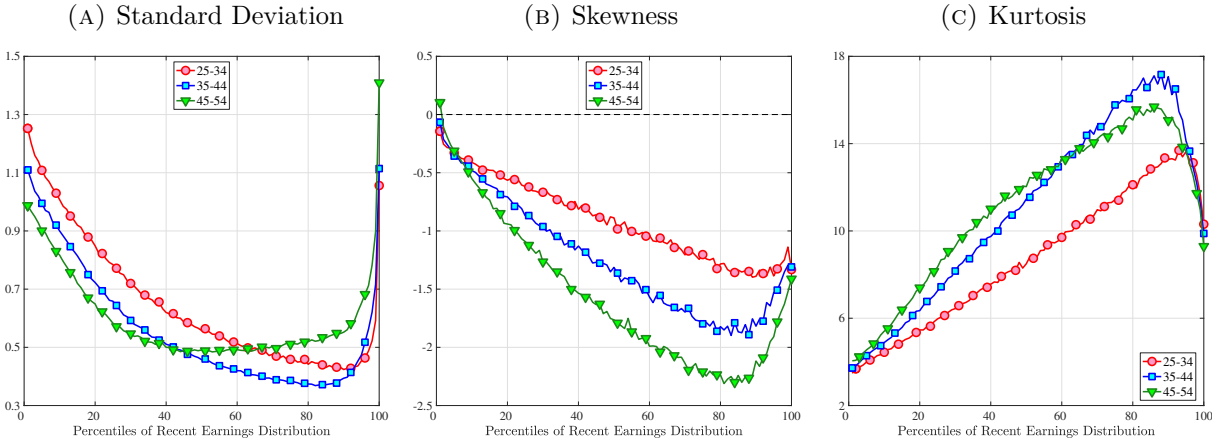
C.1 Cross-Sectional Moments of Earnings Growth in Norway and in the US

FIGURE A.1 – Cross-Sectional Moments of One-Year Log Earnings Growth in the US




Note: The figure displays the higher order moments of one-year earnings growth ($y_{t+1} - y_t$) for three age groups in the U.S. Source: [Guvenen et al. \(2019\)](#). Each point on these figures represents one percentile of RE.

FIGURE A.2 – Cross-Sectional Moments of Five-Year Log Earnings Growth in the US




Note: The figure displays the higher order moments of five-year earnings growth ($y_{t+5} - y_t$) for three age groups in the U.S. Source: [Guvenen et al. \(2019\)](#). Each point on these figures represents one percentile.

FIGURE A.3 – Cross-Sectional Moments of One-Year Log Earnings Growth in Norway




Note: The figure displays higher-order moments of one-year earnings growth ($y_{t+1} - y_t$) for three age groups in Norway.

FIGURE A.4 – Cross-Sectional Moments of Five-Year Log Earnings Growth in Norway



Note: The figure displays higher-order moments of five-year earnings growth ($y_{t+5} - y_t$) for three age groups in Norway.


FIGURE A.5 – Kelly Skewness of Earnings Growth for Males



Note: The figure displays Kelley Skewness for one-year earnings growth (left panel) and five-year earnings growth (right panel) for three age groups in Norway.

C.2 Decomposing Earnings Growth to Hours and Wages


FIGURE A.6 – The Contributions of Hours and Wage Rates to Earnings Shocks



Note: The figure displays the one-year Representative Agent changes (log change of averages) for imputed hours and imputed wage rates for 20 different groups of prime-age males (ages 36 to 55) in the 3rd RE decile (left panel) and 8th RE decile (right panel), plotted against their contemporaneous one-year log change in average annual earnings.

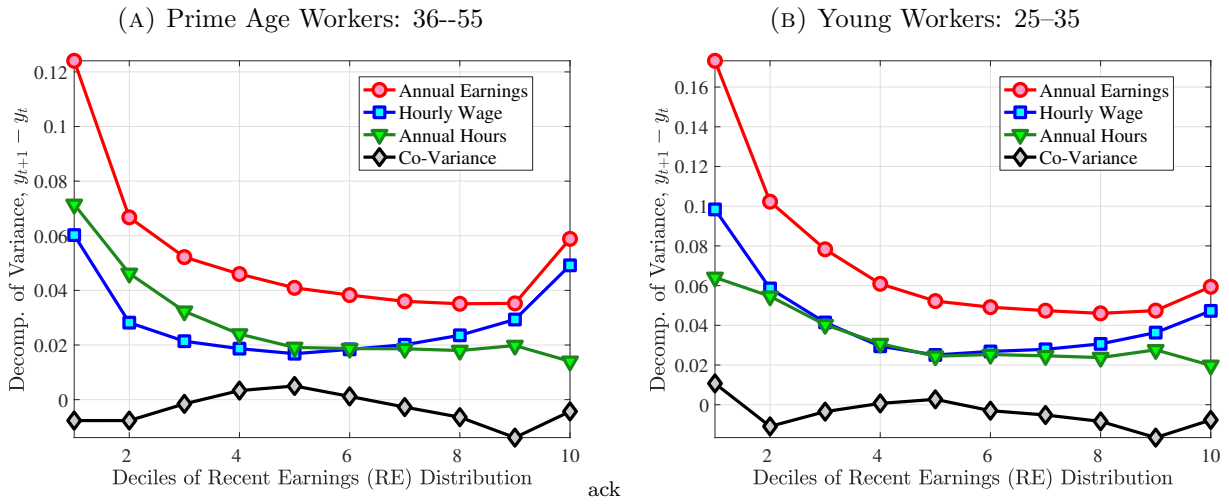
C.3 Additional Figures for Decomposition of Moments

FIGURE A.7 – Decomposition of Variance of Five-Year Growth




Note: The figure decomposes the variance of five-year log earnings changes (red line) into the variance of log wage changes (blue line), the variance of log hours changes (green line) and the co-variance between log wage and log hours changes (black line). Each dot represents a decile of RE. The decomposition is based on Lemma 1.

FIGURE A.8 – Decomposition of Variance of One-Year Growth




Note: The figure decomposes the variance of one-year log earnings changes (red line) into the variance of log wage changes (blue line), the variance of log hours changes (green line) and the co-variance between log wage and log hours changes (black line). Each dot represents a decile of RE. The decomposition is based on Lemma 1.

FIGURE A.9 – Decomposition of Skewness of Five-Year Growth, Young Workers



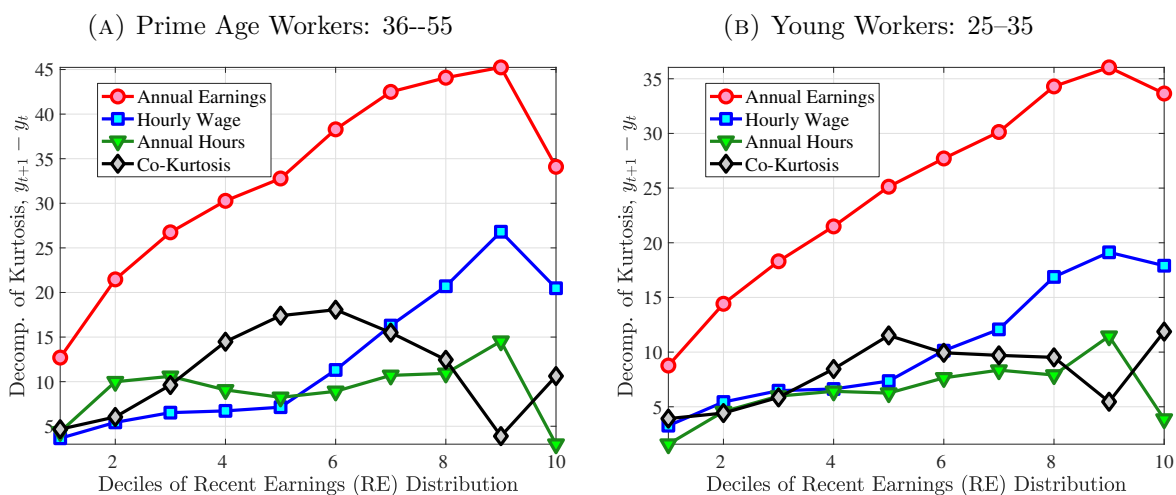
Note: The figure decomposes the skewness of five-year log earnings changes (red line) for young workers (age 25-35) into the skewness of log wage changes (blue line), the skewness of log hours changes (green line) and the co-skewness between log wage and log hours changes (black line). Each dot represents a decile of RE. The decomposition is based on Lemma 1.

FIGURE A.10 – Decomposition of Skewness of One-Year Growth




Note: The figure decomposes the skewness of one-year log earnings changes (red line) into the skewness of log wage changes (blue line), the skewness of log hours changes (green line) and the co-skewness between log wage and log hours changes (black line). Each dot represents a decile of RE. The decomposition is based on Lemma 1.

FIGURE A.11 – Decomposition of Kurtosis of One-Year Growth



Note: The figure decomposes the kurtosis of one-year log earnings changes (red line) into the kurtosis of log wage changes (blue line), the kurtosis of log hours changes (green line) and the co-kurtosis between log wage and log hours changes (black line). Each dot represents a decile of RE. The decomposition is based on Lemma 1.


FIGURE A.12 – Decomposition of Kurtosis of Five-Year Growth, Young Workers: 25–35



Note: The figure decomposes the kurtosis of five-year log earnings changes (red line) into the kurtosis of log wage changes (blue line), the kurtosis of log hours changes (green line) and the co-kurtosis between log wage and log hours changes (black line). Each dot represents a decile of RE. The decomposition is based on Lemma 1.

C.4 Cross-Sectional Moments of One-Year Growth in Household Earnings and Disposable Income

FIGURE A.13 – Cross-Sectional Moments of One-Year Log Household Earnings and Disposable Income Growth



Note: The figure displays higher-order moments of one-year log earnings changes (red line), log household earnings changes (blue line) and log disposable income changes (green line) for prime aged males, plotted for each decile of RE. The sample comprises married and co-habiting men and their households.

D Imputing Hours in the PSID

For comparison with our imputation of hours in Norwegian data, we obtain data from the Panel Study of Income Dynamics, 1968-1997. We regress yearly hours worked on a number of demographic variables as well as annual earnings:

$$h_{it} = \beta X_{it} + \epsilon_{it},$$

where X_{it} in addition to the variables in Table A.3 below contains cohort dummies.

TABLE A.3 – Predicting Hours Worked in the PSID


Variable	Men		Women	
log(earnings)		0.1847**		0.4121**
age	0.1918**	0.1381*	-0.0673	-0.1270
age ²	-0.0067**	-0.0051*	0.0035	0.0042
age ³	0.0001**	0.0001*	-0.00006	-0.0001
age ⁴	-6.34e-07**	-5.10e-07*	3.44e-07	2.54e-07
years of education	0.0614**	0.0495**	0.1468**	0.1723**
(years of education) ²	-0.0042**	-0.0032*	-0.0097**	-0.0124**
(years of education) ³	0.0001**	0.00007	0.0002*	0.0002**
Age * education	-0.00004	-0.0002**	-0.0002	-0.0005**
Marital status	0.0582**	0.0240**	-0.1632**	-0.0908**
Number of children	0.0119**	0.0089**	-0.0763**	-0.0203**
R - squared	0.0345	0.1559	0.0771	0.4505
No. of observations	52996	52996	40910	40910

* $p < 0.05$, ** $p < 0.01$

We observe that the explanatory power (as measured by R-squared) of our variables in PSID, like in the Norwegian register data, is much greater for women than for men. It also matters significantly for R-squared whether we include annual earnings as an explanatory variable. For men R-squared increases from 0.034 to 0.156 when annual earnings is included, whereas for women R-squared increases from 0.077 to 0.451.


E Additional Results for Men

FIGURE A.14 – The Contributions of Hours and Wage Rates to Earnings Shocks for 4th-7th RE Deciles (Young Men)




Note: The figure displays the one-year Representative Agent change, i.e., log change of averages, for imputed hours and imputed wage rates for 20 different groups of young males (ages 25 to 35), in the 4th-7th RE deciles, plotted against their contemporaneous one-year log change in average annual earnings.

FIGURE A.15 – The Contributions of Hours and Wage Rates to Earnings Shocks (Young Men)



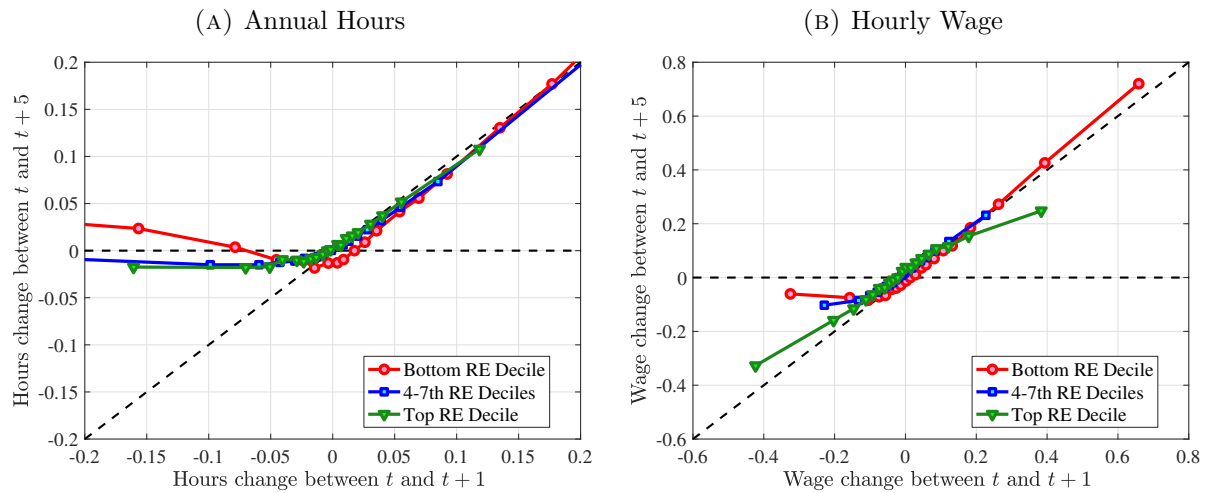
Note: The figure displays the one-year Representative Agent change, i.e., log change of averages, for imputed hours and imputed wage rates for 20 different groups of young males (ages 25 to 35) in the 1st and 10th RE deciles, plotted against their contemporaneous one-year log change in average annual earnings.

FIGURE A.16 – Persistence of Earnings Changes, Young Males



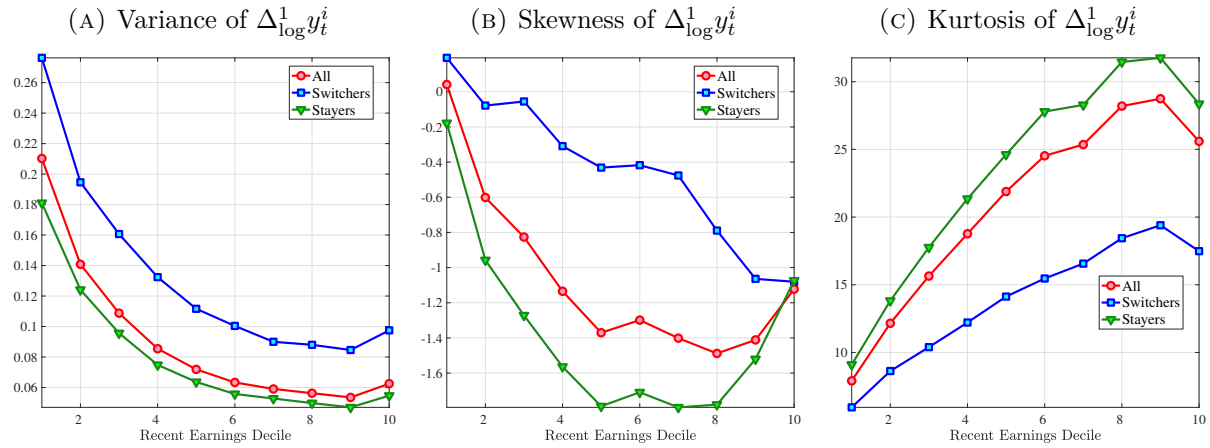
Note: The figure displays the five-year Representative Agent change (log change of averages) in earnings for 20 different groups of young males (ages 25 to 35) in the 1st RE decile (red line), 4th-7th RE decile (blue line) and 10th RE decile (green line), plotted against their respective one-year log change in average annual earnings.

FIGURE A.17 – The Persistence of Hours and Wage Changes by RE Decile, Young Males




Note: The left panel displays the five-year Representative Agent change (log change of averages) in imputed annual hours for 20 different groups of young males (ages 25 to 35) in the 1st RE decile (red line), 4th-7th RE decile (blue line) and 10th RE decile (green line), plotted against their respective one-year log change in imputed annual average hours. The right panel displays the corresponding figure for imputed hourly wage rates.

FIGURE A.18 – Cross-Sectional Moments of One-Year Log Earnings Growth (Young Males): Stayers vs. Switchers ($y_{t+1} - y_t$)



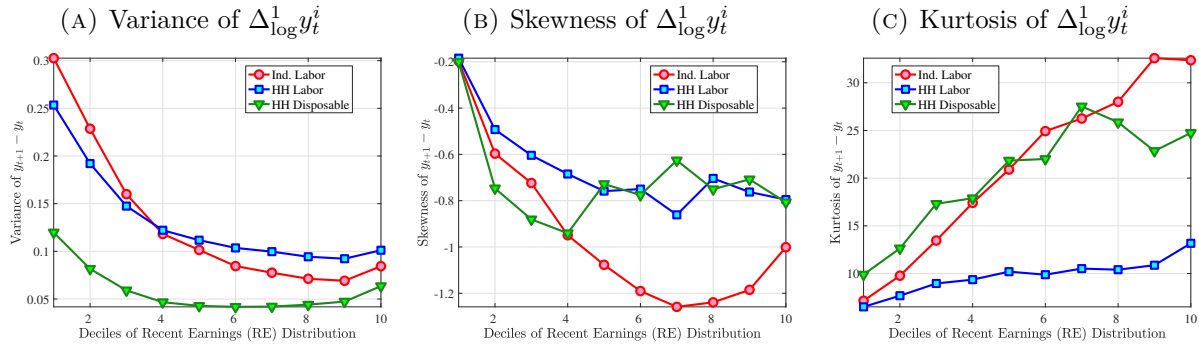
Note: The figure displays the higher-order moments of log earnings growth for job switchers (blue line), job stayers (green line), and all young males (red line) for each RE decile.

FIGURE A.19 – Contribution of Job Stayers and Job Switchers to Skewness and Kurtosis of Earnings Growth (Young Males)



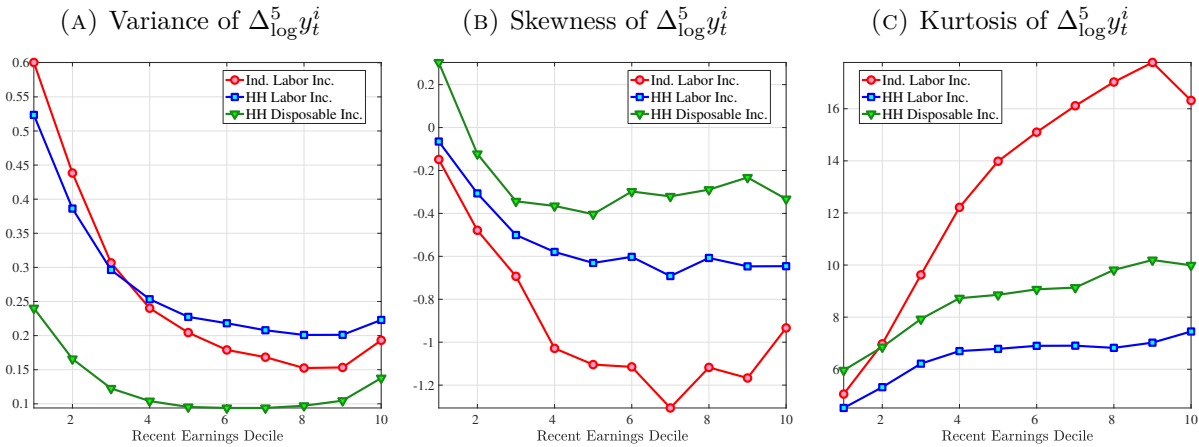
Note: The figure displays the contributions of job switchers (blue line) and job stayers (green line) to the population skewness and kurtosis of one-year earnings growth for young males (red line) for each RE decile.

FIGURE A.20 – Cross-Sectional Moments of One-Year Log Household Earnings and Disposable Income Growth (Young Males)




Note: The figure displays the higher order moments of one-year log earnings changes (red line), log household earnings changes (blue line) and log disposable income changes (green line) for young males, plotted for each decile of RE. The sample comprises married and co-habiting men and their households.

FIGURE A.21 – Cross-Sectional Moments of Five-Year Log Household Earnings and Disposable Income Growth (Young Males)



Note: The figure displays the higher order moments of five-year log earnings changes (red line), log household earnings changes (blue line) and log disposable income changes (green line) for young males, plotted for each decile of RE. The sample comprises married and co-habiting men and their households.

FIGURE A.22 – Spousal Responses to Male Earnings Shocks



Note: The figure displays the one-year Representative Agent change (log change of averages) of spouse annual earnings and imputed hours for 20 different groups of prime-age males (ages 36 to 55), plotted against their contemporaneous one-year log change in average annual earnings. The sample comprises married and co-habiting men and women.

F Results for Females

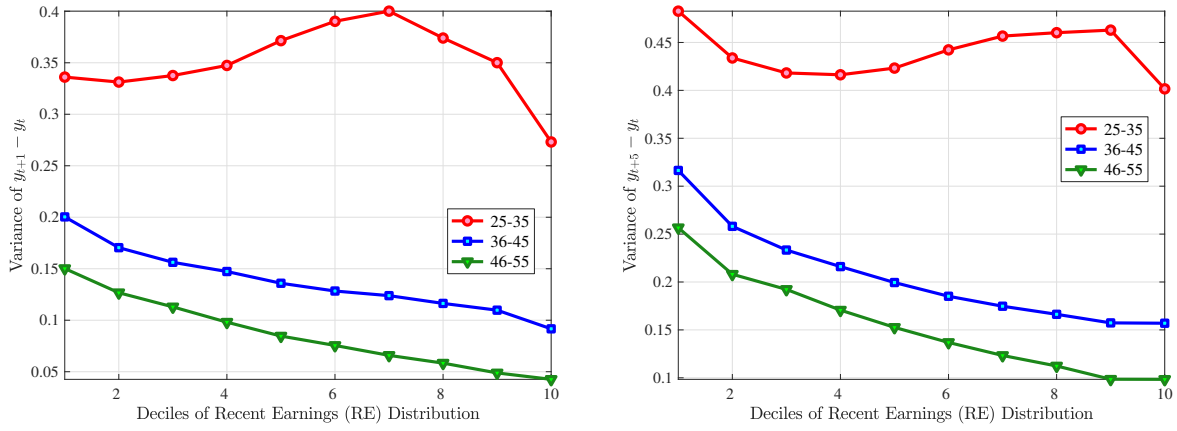
F.1 Higher Order Moments of Earnings Growth

Figure A.23 displays the variance of earnings growth for women. The results for women above 35 are both qualitatively and quantitatively quite similar to those for men. One difference is that the U-shape with higher variance for the highest earners is less pronounced for women. Women between 25 and 35 have more volatile earnings shocks than males across all RE groups and there is no decline in volatility by RE. In table A.4 below we show that parental leave is a key driver of large female earnings shocks and this is likely to explain the high volatility of earnings growth for young women.

Figure A.24 displays the skewness of earnings growth for females. The negative skewness is less pronounced for younger women than it is for men and varies less by earnings group. For older women (46-55) the negative skewness is similar to that of men in the same age group. However, we do not observe the U-shape for high top earners.


In Figure A.26 we plot the kurtosis of one-year earnings growth (left) and five-year earnings growth (right) for women. For young women the kurtosis of earnings growth is generally much lower than for young men and the increase by earnings group is less pronounced. For the group, 25-35, the kurtosis of one-year earnings growth is close to 3 (the value for a normal distribution). For older women (45-55) the kurtosis of earnings growth is high and sharply increasing in recent earnings, with an inverted U-shape for top earners.

FIGURE A.23 – Variance of Female Earnings Growth




Note: The figure displays the variance of one-year (left) and five-year (right) earnings growth for females by age and RE decile.

FIGURE A.24 – Skewness of Female Earnings Growth



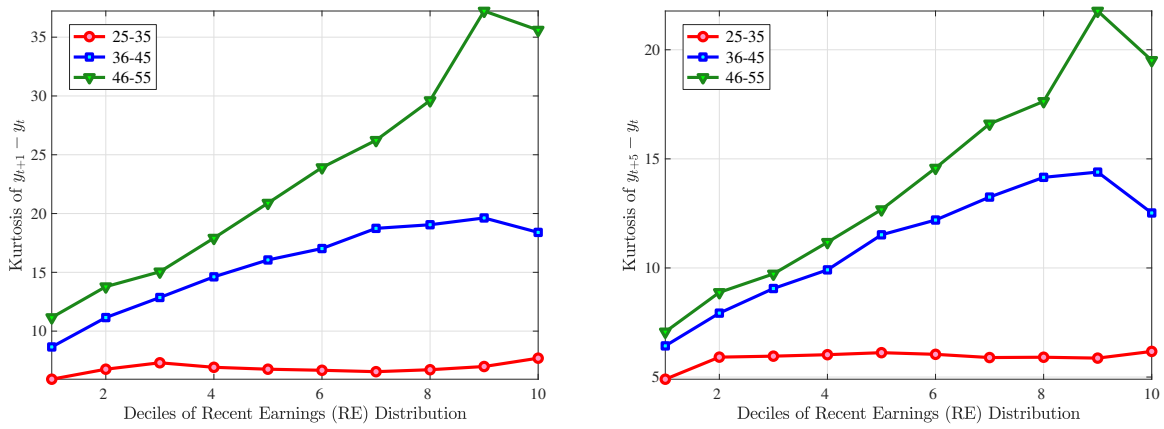
Note: The figure displays the skewness of one-year (left) and five-year (right) earnings growth for females by age and RE decile.

FIGURE A.25 – Kelly Skewness of Female Earnings Growth



Note: The figure displays the Kelly skewness of one-year (left) and five-year (right) earnings growth for females by age and RE decile.


FIGURE A.26 – Kurtosis of Female Earnings Growth



Note: The figure displays the kurtosis of one-year (left) and five-year (right) earnings growth for females by age and RE decile.


F.2 Job Stayers and Job Switchers

FIGURE A.27 – Variance of Female Earnings Growth




Note: The figure displays the variance of log earnings growth for job switchers (blue line) and job stayers (green line) for each RE decile for all females (age 25-55).

FIGURE A.28 – Skewness of Female Earnings Growth



Note: The figure displays the skewness of log earnings growth for job switchers (blue line) and job stayers (green line) for each RE decile for all females (age 25-55).

FIGURE A.29 – Kurtosis of Female Earnings Growth



Note: The figure displays the kurtosis of log earnings growth for job switchers (blue line) and job stayers (green line) for each RE decile for all females (age 25-55).


F.3 Cross Sectional Moments of Household Labor and Disposable Income Growth

Figure A.30 displays the variance of one-year and five-year income growth for the three income measures by income decile for females living as part of a couple. For females the income of the spouse is an even more important insurance mechanic than for males. The insurance from the spouse is also more important than the tax and transfer system for one-year income growth.

In Figure A.31 we plot the skewness of one-year and five-year income growth of females living as part of a couple for the three income types by recent earnings decile. The pattern is the same as for men with spousal income and the government both contributing to reducing the negative skewness.


Figure A.32 displays the kurtosis of one-year and five-year income growth of females living as part of a couple for the three income types by recent earnings decile. Adding the income from the spouse increases kurtosis for low earning females and decreases it for high earnings females.

FIGURE A.30 – Variance of the Growth of Different Income Measures (females in couple)




Note: The figure displays the variance of one-year (left) and five-year (right) log earnings changes (red line), log household earnings changes (blue line) and log disposable income changes (green line) for females of all ages (25-55), plotted for each decile of RE. The sample married and co-habiting females and their households.

FIGURE A.31 – Skewness of the Growth of Different Income Measures (females in couple)



Note: The figure displays the skewness of one-year (left) and five-year (right) log earnings changes (red line), log household earnings changes (blue line) and log disposable income changes (green line) for females of all ages (25-55), plotted for each decile of RE. The sample married and co-habiting females and their households.

FIGURE A.32 – Kurtosis of the Growth of Different Income Measures (females in couple)



Note: The figure displays the kurtosis of one-year (left) and five-year (right) log earnings changes (red line), log household earnings changes (blue line) and log disposable income changes (green line) for females of all ages (25-55), plotted for each decile of RE. The sample married and co-habiting females and their households.

F.4 Life Events Associated with Large Female Earnings Shocks

Table A.4 displays the fraction of female workers with large and small positive and negative earnings changes that experience certain events. A key difference between women and men is that one of the most frequent events experienced by women with large earnings changes is parental leave, 23% of negative changes and 25% of positive changes. Also for women long term sickness, 28% of negative changes and 27% of positive changes, and change of employer, 17% of negative changes and 25% of positive changes, is associated with large negative income shocks. Another key difference between men and women is that men with large earnings shocks on average experience a larger change in their hourly wage and smaller change in their hours. Whereas the log of the hourly wage rate for women with large income shocks on average only declines by 0.36 for negative shocks and increases by only 0.32 for positive shocks. The average change in log hours is, however, -0.57 and 0.55 for women with large negative and positive shocks respectively. For men with large earnings shocks the average change in log wage rates is -0.44 / 0.45 and the average change in log hours is -0.40 / 0.40 for declines / increases in earnings.

TABLE A.4 – Important Life Cycle Events Associated with Earnings Changes (Females)

All Life-cycle event into/out of		Annual Earnings Change, $\Delta y \in$					
		One-Year Earnings Loss			One-Year Earnings Gain		
		< -0.5	$[-0.5, -0.25)$	$[-0.25, 0.0)$	$[0.0, 0.25)$	$[0.25, 0.5)$	≥ 0.5
		(1)	(2)	(3)	(4)	(5)	(6)
(1)	Unemployment	0.04	0.04	0.01	0.02	0.05	0.05
(2)	Long-term sick	0.28	0.29	0.13	0.16	0.27	0.27
(3)	Part time	0.09	0.09	0.04	0.06	0.13	0.15
(4)	Parental leave	0.23	0.06	0.00	0.01	0.12	0.25
(5)	Firm change	0.17	0.20	0.12	0.14	0.23	0.25
(6)	$\mathbb{E} \left[\Delta_{\log}^1 h_t^i \right]$	-0.57	-0.23	-0.04	0.034	0.24	0.55
(7)	$\mathbb{E} \left[\Delta_{\log}^5 h_t^i \right]$	-0.10	-0.07	-0.07	0.00	0.22	0.56
(8)	$\mathbb{E} \left[\Delta_{\log}^1 w_t^i \right]$	-0.36	-0.13	-0.03	0.04	0.11	0.32
(9)	$\mathbb{E} \left[\Delta_{\log}^5 w_t^i \right]$	-0.08	-0.09	-0.04	0.02	0.07	0.27
(10)	# of Obs.	270,101	279,867	2,484,702	2,310,440	320,050	271,395
Lowest decile (RE=1)		(1)	(2)	(3)	(4)	(5)	(6)
(1)	Unemployment	0.07	0.07	0.03	0.03	0.06	0.06
(2)	Long-term sick	0.25	0.24	0.14	0.11	0.15	0.14
(3)	Part time	0.12	0.10	0.05	0.09	0.18	0.25
(4)	Parental leave	0.12	0.04	0.00	0.01	0.03	0.08
(5)	Firm change	0.23	0.22	0.15	0.17	0.27	0.35
(6)	$\mathbb{E} \left[\Delta_{\log}^1 h_t^i \right]$	-0.51	-0.21	-0.05	0.05	0.22	0.48
(7)	$\mathbb{E} \left[\Delta_{\log}^5 h_t^i \right]$	-0.02	0.01	0.02	0.09	0.25	0.50
(8)	$\mathbb{E} \left[\Delta_{\log}^1 w_t^i \right]$	-0.37	-0.14	-0.03	0.04	0.13	0.35
(9)	$\mathbb{E} \left[\Delta_{\log}^5 w_t^i \right]$	-0.04	-0.04	-0.02	0.03	0.12	0.35
(10)	# of Obs.	17,033	21,290	119,260	151,221	44,480	37,994
Top decile (RE=10)		(1)	(2)	(3)	(4)	(5)	(6)
(1)	Unemployment	0.03	0.03	0.00	0.01	0.02	0.02
(2)	Long-term sick	0.24	0.23	0.10	0.12	0.23	0.27
(3)	Part time	0.09	0.10	0.05	0.06	0.10	0.10
(4)	Parental leave	0.35	0.13	0.01	0.02	0.18	0.36
(5)	Firm change	0.19	0.22	0.13	0.13	0.22	0.21
(6)	$\mathbb{E} \left[\Delta_{\log}^1 h_t^i \right]$	-0.55	-0.21	-0.03	0.02	0.23	0.57
(7)	$\mathbb{E} \left[\Delta_{\log}^5 h_t^i \right]$	-0.12	-0.09	-0.09	-0.05	0.19	0.57
(8)	$\mathbb{E} \left[\Delta_{\log}^1 w_t^i \right]$	-0.35	-0.14	-0.04	0.05	0.12	0.28
(9)	$\mathbb{E} \left[\Delta_{\log}^5 w_t^i \right]$	-0.17	-0.16	-0.04	0.04	0.04	0.20
(10)	# of Obs.	29,724	30,992	299,107	250,623	25,216	19,040

The table sorts individuals into six groups according to the size of their earnings change, defined as the percentage change in earnings from t to $t+1$. Rows (1)-(5) in the table display the fraction in each earnings change group who experienced each of these events (not mutually exclusive). Rows (6)-(9) show the corresponding percentage change in imputed hours and hourly wage in each group in the same period (from t to $t+1$) and five years later (from t to $t+5$). Average over all years 1993-2014, females only.