Asymmetric Business-Cycle Risk and Social Insurance*

Christopher Busch[†] David Domeij[‡] Fatih Guvenen[§] Rocio Madera[¶]

August 18, 2018

Abstract

This paper studies the business-cycle variation in higher-order (labor) income risk—that is, risks that are captured by moments higher than the variance of income changes. We examine the extent to which such risks can be smoothed within households or with government social insurance and tax policies. We use panel data from three countries that differ in many aspects relevant for our analysis: the United States, Germany, and Sweden. Our analysis has three main results. First, analyzing individual gross labor income, we document that skewness is procyclical and dispersion (variance) is flat and acyclical in Germany and Sweden, as was previously documented for the United States. The same patterns hold true for groups defined by education, gender, occupation, and public- versus private-sector jobs. Second, within-household smoothing appears to be not effective at mitigating business cycle fluctuations in skewness, and household-level income displays cyclical patterns that are very similar to individual income. Third, government tax and transfer programs blunt some of the largest declines in incomes, reducing procyclical fluctuations in skewness. The resulting welfare gain—through the lens of a structural model—amounts to 1.3% in consumption-equivalent terms for Sweden (for which we are able to perform this calculation). However, the remaining risk (in household disposable income) is still substantial: households are willing to pay 4.6% of their consumption to completely eliminate procyclical fluctuations in skewness.

JEL Codes: D31, E24, E32, H31

Keywords: Idiosyncratic income risk, countercyclical risk, business cycles, procyclical skewness, social insurance policy.

---

*For helpful critiques and comments, we thank Nick Bloom, Alexander Ludwig, Kurt Mitman, Luigi Pistaferri, Kjetil Storesletten, and Anthony Smith, as well as seminar and conference participants at various institutions. Busch acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563). Domeij acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation. Guvenen acknowledges financial support from the National Science Foundation.

[†]MOVE-UAB and Barcelona GSE; chris.busch@movebarcelona.eu; www.chrisbusch.eu

[‡]Stockholm School of Economics; david.domeij@hhs.se; https://sites.google.com/view/daviddomeij/

[§]University of Minnesota, FRB Minneapolis, and NBER; guvenen@umn.edu; www.fguvenen.com

[¶]Southern Methodist University; romadera@smu.edu; www.rocio-madera.com

1 Introduction

There is broad consensus among economists that idiosyncratic earnings risk rises in recessions. While this much is agreed upon, there is still an active debate on two sets of key questions. One, what is the precise nature of idiosyncratic income risk, and how does it change in recessions?[1] And two, how successful are various ways by which individual income fluctuations are mitigated in an economy, which prevents these fluctuations from affecting an individual's consumption? The answers to these questions have far-reaching implications for a wide set of both positive and normative macroeconomic issues in the context of, e.g., asset pricing and the effectiveness of monetary and fiscal policy.

A simple way to approach the first set of questions is to measure risk by the variance of income changes and measure how much it rises in recessions.[2] While this is certainly a useful first step, it cannot be the last: in principle, income risk depends on the entire distribution, so higher-order moments, such as skewness and kurtosis could matter just as much—and perhaps more—for income risk. Two distributions of income shocks with the same variance can imply very different amounts of risk if they differ in higher-order moments. To emphasize this point, we will refer to income risk captured by the latter components as *higher-order* income risk and study how it changes over the business cycle.

Turning to the second question, two prominent sources of income smoothing are the household, where spouses can act together to mitigate fluctuations in their individual incomes, and the government, which operates a rich tax and transfer system, parts of which are specifically designed to insure against income losses, and more so during recessions. Thus, one goal of this paper is to understand the extent to which the rise in individual income risk in recessions is mitigated by households and government policies.

In order to provide a broad perspective on both individual income dynamics and insurance mechanisms, we study panel data on individuals and households from three

[1]Measuring income *risk* faces well-understood challenges because researchers observe income fluctuations in the data, but (often) do not have access to other information the worker may have to distinguish what is the anticipated component and what is a surprise. This is a difficult issue that has been addressed in only a few papers in the literature (see, Pistaferri (2001); Cunha *et al.* (2005); Guvenen (2007); Guvenen and Smith (2014)). In this paper, we will follow the bulk of the literature that treats income fluctuations as unanticipated.

[2]We use earnings and income interchangeably throughout the paper.

countries—the United States, Germany, and Sweden. These countries differ in important aspects relevant for our analysis, such as household structures, the extent of their social safety nets, labor market institutions and compositions (unions, employment in public sector), among others. The data sets we use are based on social security records (SIAB for Germany), tax register data (LINDA for Sweden), and household surveys (PSID for the United States and SOEP for Germany), and cover more than three decades in each country.

The focus of our analysis is on the business-cycle variation in the skewness (or *asymmetry*—hence the title) of the distribution of income changes in addition to the variance. Consistent with some recent work (reviewed below), we find that skewness fluctuations are procyclical and are a critical component of changing idiosyncratic risk over the business cycle. We have also examined the business-cycle variation in the fourth moment—the kurtosis—but did not find large and robust cyclical patterns. That said, the *average level* of the kurtosis of income changes is very high, meaning that the distribution has very high concentration at the center as well as long and thick tails. These long tails interact with, and amplify, the effects of skewness fluctuations to generate a large rise in idiosyncratic risk in recessions.

Our analysis yields four sets of results. First, in all three countries, the variance of *individual gross income growth* is relatively stable over time and is acyclical, whereas the skewness is robustly procyclical.[3] This finding both confirms the empirical evidence in Guvenen *et al.* (2014) from US administrative data and shows that it holds more broadly—in two other developed economies, as well as in survey data (the PSID and SOEP). In addition, we show that this result is robust across demographic groups defined by gender, skill/education, occupation, and private/public sector employment. Furthermore, the administrative data set we employ from Germany contains information on workers' establishment and their full-time status. Using this information we examine the extent to which the cyclicality of skewness is driven by variation in work hours or employer changes. We find that even for full-time workers who are continuously employed at the same establishment, the variance of income growth is flat and acyclical and the skewness is procyclical, with magnitudes that are similar to the overall population. This indicates that the results are not exclusively driven by changes in annual hours worked (i.e., periods of unemployment).

[3]In most of our empirical analysis, we focus on a time-series regression of the moments of individual income growth on an indicator of business cycles (contemporaneous GDP growth) in that country. This is a simple but useful way to summarize the cyclicality of each moment.

Second, moving to *household-level gross income*, we find evidence that households are not very effective at mitigating business cycle fluctuations in skewness of individual level incomes. In order to assess smoothing within households, we form "synthetic" (or random) couples by randomly matching men and women in a way that matches the overall distribution of couples by age and education level in the economy. By construction, the individual incomes within the random couples are not coordinated. Thus, if there was successful insurance within households against business-cycle variation in individual incomes, the joint income of actual couples would be less cyclical than the joint income of the artificial couples.[4] We do not find this to be the case. Actual couples have similar or higher cyclicalities in most dimensions we measure, suggesting limited smoothing of business cycle variation in higher-order risk.[5] Part of the explanation for this result could be positively correlated incomes within the household, stemming from the assortative nature of marital matching or from shocks that are semi-aggregate (regional, or industry-level where spouses work in similar industries, etc.).

Third, we move to *post-government household income*—that is, income after all government tax and transfer programs are accounted for—and find that the procyclical fluctuations in skewness are smaller in all three economies. When zooming in on the upper and lower half of the distribution of earnings changes, we find interesting heterogeneity across countries. In both the United States and Germany, the lower half of the distribution becomes non-cyclical when moving from pre- to post-government earnings. In Sweden, the upper half becomes non-cyclical, while the lower half also becomes less cyclical (as measured by changes in the point estimates), however it still systematically varies with the business cycle.

Fourth, and finally, we investigate how effective government policy is in insuring households against business-cycle fluctuations of earnings risk, and how much households value this insurance. To do that, we need a structural model where households can self-insure through borrowing and saving, wherein we can translate the statistics on income changes into distributions of underlying income shocks. To generate non-zero skewness and excess kurtosis consistent with the data, we specify the econometric model for (log) income as the sum of a persistent process and a transitory component, where all innovations are drawn from mixtures of normals. We estimate the

[4]Two ways to think about this are that actual households could be formed by partners chosen to have negatively correlated income shocks, or that partners could react with their labor supply to income shocks of their spouse.

[5]Notice that there could be smoothing of the level of risk or other aspects of it. Here we focus on the business cycle variation only.

parameters of this process separately for pre- and post-government household labor income by matching moments on the cyclical properties of the higher-order income risk documented in the empirical part.

We then feed these parameters into a variant of the partial insurance model of Heathcote *et al.* (2014) to quantify the welfare gains. We conduct this analysis for Sweden only, because even this fairly rich income process does not fit the dynamics of the higher-order moments we target sufficiently well for the United States and Germany, which makes a comparison of pre- versus post-government income less reliable. For Sweden, we find that the degree of overall insurance provided by the existing tax and transfer system amounts to a welfare gain of 1.3% in consumption equivalent terms (CEV). However, the remaining risk (in post-government household-level income) is still substantial: households are willing to pay 4.6% of their consumption to completely eliminate procyclical fluctuations in skewness.

The paper is organized as follows. The next section discusses the data sources, and Section 3 describes the empirical approach. Section 4 presents the results for gross individual income for various groups in the population. Section 5 expands the analysis to households and post-tax-transfer income. Section 6 uses a structural consumption-savings model with partial insurance to quantify the welfare benefits of governments' social insurance policies in the three countries under study. Section 7 concludes.

Related Literature

Earlier empirical work in the literature was limited by the small sample size and time span of the available survey-based panel data sets, such as the PSID, leading researchers to make parametric assumptions to obtain identification. One common assumption is that shocks to earnings are Gaussian, which implies zero skewness. Restricting attention to the changes in the mean and variance of income shocks, Storesletten *et al.* (2004) concluded that the variance of income shocks in the US data is countercyclical.[6]

More recently, Guvenen *et al.* (2014) studied the earnings histories of a 10% representative panel of US males from SSA records. The large sample size allowed them to relax parametric assumptions as well as to examine variations in skewness. They found that the variance of income shocks is stable over the business cycle and is robustly acyclical, whereas the skewness of shocks varies significantly over time in a procyclical fashion. The current paper goes substantially beyond their analysis by studying

[6]Using a similar approach, Bayer and Juessen (2012) studied the cyclicality of the variance in Germany, the UK, and the US, and different patterns in Germany and the UK relative to the US and attributed it to differences in institutions.

two new countries and four data sets, using income measures that include significant sources of smoothing, and employing a structural consumption-savings model to quantify the welfare benefits of government insurance and the cost of remaining (uninsured) fluctuations in higher-order income risk.

Taking a different methodological approach, Busch and Ludwig (2018) adapt the parametric approach of Storesletten *et al.* (2004) to allow for skewness fluctuations and analyze the cyclicality of labor income risk in Germany and the United States. They come to the same substantial conclusion as we do, namely, that variation of income risk over the business cycle is asymmetric. In ongoing work, Angelopoulos *et al.* (2018) follow the approach in the present paper to study the cyclicality of higher-order risk in the United Kingdom using panel data from the British Household Panel Survey. They confirm the same finding of strongly procyclical skewness for the UK since the early 1990s. Similarly, Harmenberg and Sievertsen (2018) document procyclical skewness of individual earnings changes in administrative Danish data. In a recent paper, Pruitt and Turner (2018) analyze individual and household-level income dynamics using United States tax records from the IRS. They also document procyclical skewness of income changes for both male and household incomes. Unlike in our four data sets, they find countercyclical dispersion of male (not household) earnings growth.

A couple of recent papers aim at exploring the role played by hours versus wages for the observed cyclical dynamics of earnings changes. In an analysis of administrative unemployment insurance data from Washington State, Kurman and McEntarfer (2017) document procyclical skewness of hourly wage changes. They also explicitly show that the share of workers realizing a wage cut increases substantially in recessions. Pora and Wilner (2017) document in French administrative data that the distribution of earnings changes was more negatively skewed in the 2008 recession than in the directly preceeding period. Conditioning on income, they find that for high-income workers, hourly wages account for this change of the distribution, while for low-income workers hours worked are more important. In data from Italian social security records, Blass-Hoffman and Malacrino (2016) find a larger role for hours changes in driving fluctuations in skewness of earnings changes than what we find for Germany.

A growing number of theoretical and quantitative studies emphasizes the importance of the higher-order moments of income shocks for various economic questions. In asset pricing, several papers have found that the procyclical skewness of consumption (and income) growth helps explain various puzzling features of asset prices (Mankiw (1986), Constantinides and Ghosh (2014), Schmidt (2016)). Recent research on monetary and fiscal policy also emphasizes the role of higher-order income risk in shaping optimal policy or in modifying the standard channels through which policy works. Examples include Kaplan *et al.* (2016) who examine the monetary transmission mechanism in the presence of leptokurtic shocks, and Golosov *et al.* (2016) who find that, in a Mirleesian setting, the optimal tax schedule is greatly affected by whether or not one accounts for higher order moments of income shocks.

2 The Data

This section provides an overview of the data sets we use in our empirical analysis, the sample selection criteria, as well as the variables used in the subsequent empirical analyses. Given the diversity of our data sources, we relegate the details to Appendix A. Briefly, we employ four longitudinal data sets corresponding to three different countries: the Panel Study of Income Dynamics (PSID) for the United States, covering 1976 to 2010;[7] the Sample of Integrated Labour Market Biographies (SIAB[8]) and the German Socio-Economic Panel (SOEP) for Germany, covering 1976 to 2010 and 1984 to 2011, respectively; and the Longitudinal Individual Data Base (LINDA) for Sweden, covering 1979 to 2010. The PSID and the SOEP are survey-based data sets. The PSID has a yearly sample of approximately 2,000 households in the core sample, which is representative of the US population; the SOEP started with about 10,000 individuals (or 5,000 households) in 1984 and, after several refreshments, covers about 18,000 individuals (10,500 households) in 2011.[9]

The SIAB is based on administrative social security records and our initial sample covers on average 370,000 individuals per year. It excludes civil servants, students, and self-employed workers, which make up about 20% of the workforce. From the perspective of our analysis, the SIAB has two caveats: (i) income is top-coded at the limit of income subject to social security contributions, and (ii) individuals cannot be linked to each other, which prohibits identification of households. We deal with (i) by fitting a Pareto distribution to the upper tail of the wage distribution[10] and with (ii)

[7]The PSID contains information since 1967. We choose our benchmark sample to start in 1976 because of the poor coverage of income transfers before the 1977 wave. We complement our results using a longer period whenever possible and pertinent.

[8]We use the factually anonymous scientific use file SIAB-R7510, which is a 2% draw from the Integrated Employment Biographies data of the Institute for Employment Research (IAB).

[9]These numbers refer to observations after cleaning but before sample selection. Only the representative SRC sample is considered in the PSID. The immigrant sample and high-income sample of the SOEP are not used, because they cover only subperiods.

[10]The imputation is done separately for each year by subgroups defined by age and gender. For

by using data from SOEP for all household-level analyses. Throughout the analysis, we focus on West Germany, which for simplicity we refer to as Germany. LINDA is compiled from administrative sources (the Income Register) and tracks a representative sample with approximately 300,000 individuals per year.

For each country, we consider three samples: two at the individual level—one for males and one for females—and one at the household level. The samples are constructed as revolving panels: for a given statistic computed based on the time difference between years $t-s$ and t, the panel contains individuals who are ages 25 to 59 in periods $t-s$ and t ($s=1$ in the case of Sweden and Germany, and $s=2$ in the case of the United States) and have yearly labor earnings above a minimum threshold in both years. This threshold is defined as the earnings level that corresponds to 520 hours of employment at half the legal minimum wage, which is about $1,885 US dollars for the United States in 2010.[11] To avoid possible outliers, we exclude the top 1% of earnings observations in the PSID and SOEP, but not in LINDA (which is from administrative sources). For each individual, we record age, gender, education, and gross labor earnings. By gross earnings we mean a worker's compensation from his/her employer before any kind of government intervention in the form of taxes or transfers.

The household sample is constructed by imposing the same criteria on the household head and adding specific requirements at the household level. More specifically, a household is included in our sample if it has at least two adult members, one of them being the household head,[12] that satisfy the age criterion and household income that satisfies the income criteria. At the household level, we analyze pre- and post government earnings. Pre-government earnings is defined as the sum of gross labor earnings earned by the adults in the household. Post-government earnings is constructed by adding taxes and transfers.

workers with imputed wages, across years, we preserve the relative ranking within the age-specific cross-sectional wage distribution. The procedure follows Daly *et al.* (2014); see Appendix A.3 for details.

[11] For the United States, we use the federal minimum wage. There is no official minimum wage in Sweden or Germany during this period. For Germany, we follow Fuchs-Schündeln *et al.* (2010) and take a minimum threshold of 3 euros (in year 2000 euros) for the hourly wage. For Sweden, the effective hourly minimum wage via labor market agreements was around SEK 75 in 2004 (Skedinger, 2007). For other years, we adjust the minimum wage by calculating the mean real earnings for each year, estimating a linear time trend for these means and removing that time trend from the SEK 75 minimum wage.

[12] In PSID and SOEP, the head of a household is defined within the data set. In LINDA, the head of a household is defined as the sampled male.

3 Empirical Approach

Measuring Income Volatility over the Business Cycle

For each year, we calculate robust statistics of log s-year changes in income. We consider different choices of s in order to distinguish between earnings growth over short and long horizons, and interpret these as corresponding to “transitory” and “persistent” earnings shocks.

More specifically, we compute moments $m[\Delta_s y_t]$, where $y_t \equiv \ln Y_t$, $\Delta_s y_t \equiv y_t - y_{t-s}$, and Y_t denote income in period t. The moments m we consider are: the log differential between the 90th and 10th percentiles (L9010), the Kelley measure of skewness, and the upper (L9050) and lower (L5010) tails. For Germany and Sweden, s refers to 1-year changes. Due to the biennial structure of the PSID from the 1997 wave on, our analyses for the United States refer to 2-year changes instead.[13]

We do not impose any parametric assumption on the dynamics of income but instead analyze the behavior of the tails of the distribution of earnings changes. We think this is important since interpretations when using the variance as a summary statistic of the distribution alone can be misleading. To see this point, consider a widening of both the upper and lower tails of a normally distributed variable. That is, P90 is shifted to the right and P10 is shifted to the left. This certainly implies an increase in the variance; the opposite, however, is not necessarily true. Think of the case in which only the lower tail shifts to the left. Then the overall dispersion of the distribution increases, but if we were to interpret this increase in isolation, we would wrongfully conclude that not only one tail expands, but both of them expand.

Similarly, unchanged overall dispersion does not imply an unchanged distribution, but can be observed when both tails move together (i.e., one tail shrinks while the other expands). Both of these last two scenarios imply a change of the relative size of the tails–a feature summarized by the skewness of the distribution. In our empirical analysis, these are the two scenarios we observe when considering cyclicality: either overall dispersion does not change while skewness does, or dispersion is cyclical, caused by one tail expanding and the other shrinking.

A Continuous Measure of Business Cycles

Some important macroeconomic variables do not perfectly synchronize with expansions and contractions, as classified by the NBER dating committee for example,

[13] We calculate overlapping s-year differences up to $\Delta_s y_{1996}$, and non-overlapping s-year differences from then and up to $\Delta_s y_{2010}$, for $s = 2, 4$.

but their fluctuations might have an impact on earnings. For example, the US stock market experienced a significant drop in 1987, officially classified as an expansion year, and indeed the skewness of household income growth dips in that year (Figure 2a). Similarly, the US economy displayed an overall weakness in 1993 and 1994, which is evident in a range of economic variables, but these years are technically classified as expansion years. Other examples (e.g., 1996) are easy to find for Germany and Sweden. Therefore, the main focus of our analysis will be on the comovement of higher-order moments of earnings changes with a continuous measure of business cycles. We use the (natural) log growth rate of GDP—that is $\Delta_s GDP_t \equiv \ln(GDP_t) - \ln(GDP_{t-s})$—as our measure of aggregate fluctuations. More specifically, we regress each moment m of the log income change between $t-s$ and t on a constant, a linear time trend, and the log growth rate of GDP between year $t-s$ and t :

$$m\left(\Delta_s y_t\right) = \alpha + \gamma t + \beta^m \times \Delta_s(GDP_t) + u_t. \tag{1}$$

For a quantitative interpretation of the results reported in the next sections, Figure 1 reports the short-run volatility of GDP growth for each country and displays the cyclical component of log GDP as a reference.[14]

4 Empirical Results: Gross Individual Income

In this section, we address four questions concerning higher-order risk for individual earnings. First, we ask whether the countercyclical skewness and the acyclical dispersion are US-only phenomena or robust features of business cycles that can be observed in other countries whose labor markets differ greatly from that in the US. For example, according to the OECD (2016), 10.7% of US workers are unionized and 11.9% are covered by trade union agreements. In Germany, the equivalent shares are 18.1% and 57.6%, respectively. In Sweden, 67.3% are unionized and the overwhelming majority (89%) of workers is covered by trade union agreements.[15]

Second, we ask whether business-cycle variation in higher-order income risk differs across observationally distinct groups, defined by gender, education, private/public

[14]Throughout the paper, shaded areas indicate recessionary episodes. For the US, we classify recession episodes based on the NBER peak and trough dates, with the exception that we classify 1980-1983 as a single "double-dip" recession. For Germany and Sweden, we classify recessions based on Economic Cycle Research Institute (ECRI) dates, with the exception that we classify 2001-2003 as a recession in Sweden, since Swedish GDP fell by a similar magnitude to that in the US and Germany during these years, as seen in Figure 1.

[15]The numbers refer to 2013.

Figure 1: Cyclical Component of Quarterly GPD Growth: US, Germany, and Sweden

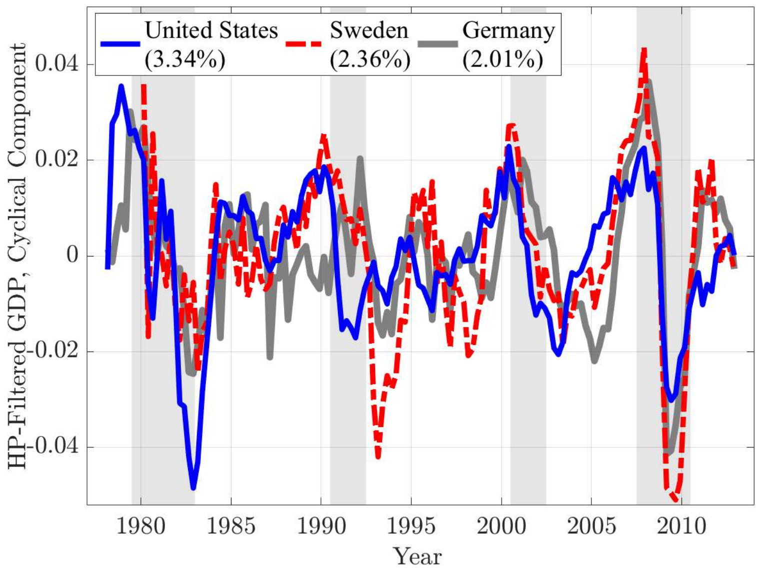

Note: The shaded areas indicate US recessions. The series for Germany corresponds to West Germany up to and including 1990Q4, and to (Unified) Germany from 1991Q1 on. The cyclical components are obtained by HP-filtering the series for GDP per capita from 1970Q1 to 2014Q1. The numbers in parentheses next to each country indicate the standard deviation of the (unfiltered) short-run GDP growth series over the period 1976-2010, where short-run is one-year difference for Germany and Sweden, and two-year difference for the United States, to be consistent with the micro data used in our analysis.

sector employment and occupation. Third, we ask whether the cyclicality of earnings changes can be attributed mainly to changes in hours worked or to changes in wages, or both. Fourth, we ask whether the countercyclicality of skewness and the acyclicality of dispersion found in US administrative earnings data are also borne out in US survey data (e.g., the PSID). This question is important because earlier papers that used the PSID and adopted parametric methods found a strongly countercyclical variance of shocks. This raises the question: is it the data set or the methodology that accounts for these different conclusions?

Cyclicality of Dispersion

In Table I, we report the cyclicality of four key statistics computed from the distribution of earnings changes of individual workers. To provide a comparative discussion, we report the results for all three countries in the same table. For now, we focus on the first row of each panel, corresponding to the sample of male workers in each country. In the United States, L9010 for males is acyclical, as seen from the statistically insignificant coefficient (–0.54 with a *t-stat* of –1.38). Turning to Sweden and Germany,

Table I: Cyclicality of Income Growth Moments: Gross Individual Income

	L9010	Kelley	L9050	L5010
	United States			
Males	–0.54	2.25***	0.68**	–1.23***
	(–1.38)	(4.79)	(2.49)	(–4.27)
Females	0.40	1.17***	0.86**	–0.47**
	(1.39)	(3.01)	(2.57)	(–2.38)
	Sweden			
Males	–0.11	3.74***	0.91***	–1.01***
	(–1.22)	(4.00)	(3.80)	(–3.74)
Females	0.43**	1.64***	0.67***	–0.24**
	(2.24)	(3.33)	(3.09)	(–2.67)
	Germany (SIAB)			
Males	0.15	5.48***	0.95***	–0.80***
	(0.36)	(5.80)	(3.14)	(–4.11)
Females	0.34	2.55**	0.80	–0.46*
	(0.48)	(2.05)	(1.25)	(–1.80)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3 for SIAB and LINDA, 2 for PSID). Asterisks ($*, **, ***$) denote significance at the 10%, 5%, and 1% levels.

the L9010 for male earnings is also acyclical.[16]

Overall, we conclude that in all three countries, the dispersion of earnings changes does not display any robust pattern of cyclicality, judging from these regressions. In addition to being acyclical, the dispersion of earnings changes is quite flat over time (see the left panels of Figure 2). These figures should be compared with typical calibrations in the literature that assume that the volatility of earnings shocks doubles or triples during recessions. Here, the largest movements are on the order of 10% to 15%, and they show no signs of cyclicality.

Cyclicality of Skewness

We next turn to the cyclical behavior of skewness. Column 2 in Table I reports one measure of asymmetry, Kelley skewness, defined as follows:

$$\mathcal{S}_k = \frac{(P90 - P50) - (P50 - P10)}{(P90 - P10)}.$$

[16]All regression results based on SIAB data are robust to various robustness checks that address issues of top-coding and a structural break in the wage variable. See Appendix C for details.

Figure 2: Standard Deviation, Skewness, and Tails of Short-Run Income Growth; United States, Sweden, and Germany (SIAB): All Males

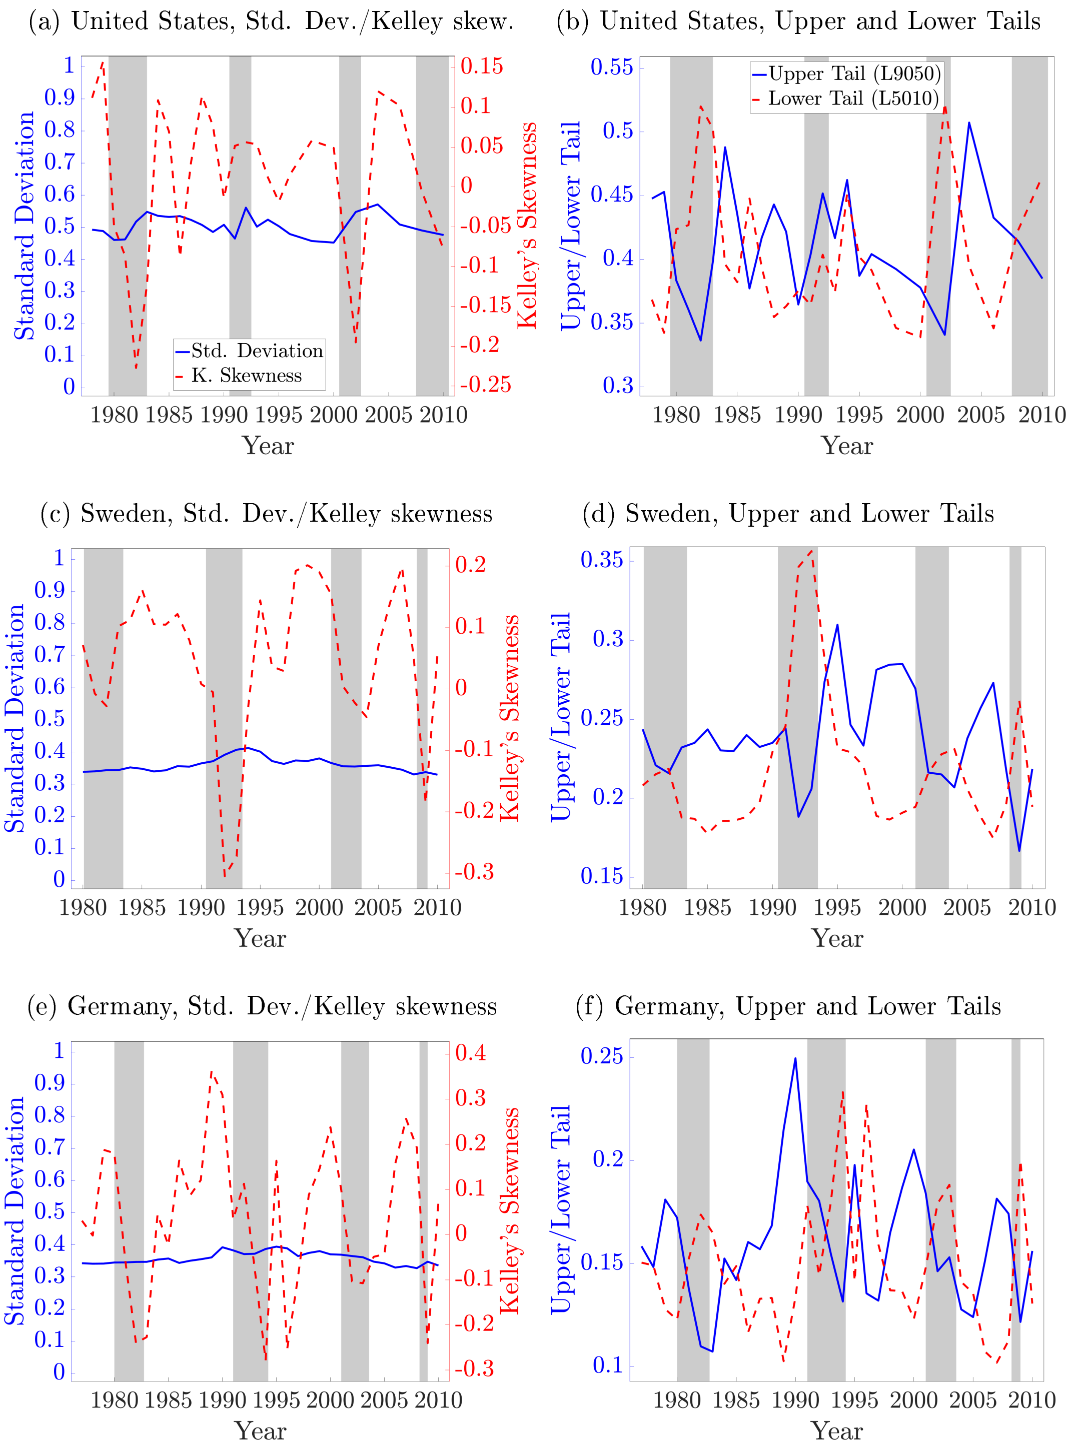

Note: Linear trend removed, centered at sample average. Shaded areas indicate recessionary periods (see footnote 14).

This measure has several attractive features compared with the standardized third moment. First, it is not sensitive to extreme observations, since it does not depend on observations beyond the 90th and 10th percentiles of the distribution. It is therefore our preferred measure of skewness, especially when considering the survey data from the US and Germany (GSOEP) where potential outliers and measurement issues could be more important.[17] Second, the particular value of Kelley skewness has a simple interpretation, in terms of the relative lengths of the upper and lower tails. In particular,

$$\frac{P90 - P50}{P90 - P10} = 0.5 + \frac{\mathcal{S}_k}{2}, \tag{2}$$

which can be used to compute the fraction of overall dispersion (P90-P10) that is accounted for by the upper tail (P90-50) and consequently by the lower tail (P50-P10).

In all three countries, Kelley skewness is procyclical (see the left panels of Figure 2). This pattern is particularly striking in Sweden and Germany, where movements in Kelley skewness are almost perfectly synchronized with the business cycle as defined by ECRI. The notable exception is the fall in Kelley skewness in 1996, but note that the cyclical component of GDP did indeed fall in 1996, as displayed in Figure 1. Furthermore, Table I shows that the procyclicality of Kelley skewness is (statistically) significant at the 1% level in all three countries. The coefficient is 1.67 for the US, 3.74 for Sweden, and 5.48 for Germany, showing more cyclicality when moving from the US to Sweden and the most cyclicality for Germany.

In order to interpret these coefficients, we need to take into account the volatility of GDP growth itself in each country, as measured by the standard deviations reported in Figure 1. Thus, for example, if a typical recession in Sweden entails a drop in GDP growth of two standard deviations (from $+1$ to -1 sigmas, for a swing of $2 \times 0.0236 = 0.0472$), Kelley skewness will fall by $0.0472 \times 3.74 = 0.177$. For the sake of discussion, suppose $\mathcal{S}_k^{\text{exp.}} = 0$ in an expansion, then $\mathcal{S}_k^{\text{rec.}} = -0.18$, which in turn implies from equation (2) that the upper tail to lower tail ratio, $(P90 - P50)/(P50 - P10)$, goes from 50/50 to 41/59 from an expansion to a recession. This is a large change in the relative size of each tail, especially for a country like Sweden, which might be thought of as displaying lower business-cycle risk (because of the high unionization rate, among other factors).[18]

[17]We have also analyzed the third standardized moment and found very similar results.

[18]The corresponding changes in $\mathcal{S}_k$ for the US and Germany are 0.15 and 0.22, respectively.

Inspecting the Tails

At the expense of some oversimplification, it might be useful to think about a shift toward more negative skewness as arising from either a compression of the right tail or an expansion of the left tail or both. Thus, a follow-up question is: which one of these changes is driving the cyclical changes in skewness for each country? Again, the pattern is particularly striking in Sweden (see panel d of Figure 2). It shows that the top tail is procyclical, whereas the bottom tail is countercyclical. The last two columns of Table I show that this pattern is present and (statistically) significant in all three countries. This means that, in a recession, the positive half of the shock distribution compresses relative to the median, whereas the negative half expands. Thus, the shift toward negative skewness happens through the process of both tails moving in unison during recessions.

Furthermore, notice that for all three countries, it turns out that the magnitude of movement of each tail is similar to each other. For example, for Sweden, the coefficient for L9050 is 0.91 and for L5010 it is –1.01. The corresponding coefficients are 0.68 and –1.23 for the US, and 0.95 and –0.80 for Germany. Therefore, as log GDP growth fluctuates over the business cycle, the shrinking of one tail is matched closely by the expansion of the other tail, making total dispersion, the L9010, move very little over the cycle. As a result, skewness becomes more negative in recessions without any significant change in the variance.

This analysis shows that the behavior of higher-order risk is best understood by separately studying the top and bottom tails over the cycle, which can move together or independently. Focusing simply on a directionless moment, such as the L9010 or the variance, can miss important asymmetries that might matter for the nature of earnings risk. As we will see in a moment, whenever we observe cyclical dispersion, it is driven by *asymmetric* movements of the tails and should not be thought of as a pure change in L9010 or the variance (which would imply an expansion/compression of *both* tails).

Survey versus Administrative Data

As noted earlier, it is not possible to link individual data from the SIAB data set to obtain household-level information. This is why we use survey data (PSID for the US and SOEP for Germany) to answer questions regarding insurance provided within households and by the government. These data sets, however, suffer from having fairly few observations, which may imply that higher moments are imprecisely estimated.

Specifically, we have rerun the regression in equation (1) using moments from the SSA data (reported in Guvenen *et al.* (2014)), and from SOEP data. The resulting coefficients for US males using SSA data for each of the four moments are –0.07 , 2.31*** , 1.02*** , and –1.09**, respectively. These numbers are strikingly similar to those in the first row of the top panel in Table I. The equivalent numbers using SOEP data are –1.33**, 1.76***, –0.21, and –1.12***. While these numbers differ somewhat from those in the first row of the bottom panel in Table I, they tell the same story. In particular, male earnings changes in both SOEP and SIAB are characterized by asymmetric movements of the tails rather than uniform expansions and contractions of both tails.[19] The main difference is that, in the SOEP, there is evidence of countercyclical dispersion, which was not observed in the SIAB.

However, this is best understood by looking directly at the tails. The lower tail is countercyclical in both data sets, whereas the upper tail is procyclical in SIAB but acyclical in SOEP. As a result, the L9010 is acyclical in SIAB and countercyclical in SOEP. This is yet another example in which limiting the analysis to the overall measure of dispersion gives an incomplete picture: the L9010 is countercyclical, but due to an expansion of the lower tail in contractions while the upper tail is unchanged, not to a symmetric expansion of both tails. This evidence of asymmetric risk is reflected in procyclical skewness.

4.1 Differences by Gender

In examining the cyclicality of higher-order risk for female workers (the second row of each panel in Table I), we see two main patterns. First, Kelley's measure of skewness is always procyclical as indicated by the positive coefficient on log GDP growth, which is highly significant for Sweden and the US (1% level), and significant for Germany (5% level). Second, inspecting the top and bottom tails separately (last two columns), we observe the expected pattern of cyclicality whenever the coefficient is significant. In particular, the L9050 is procyclical and significant for the US and Sweden, whereas the L5010 is countercyclical and significant for all three economies. Thus, just as in the case of male workers, the behavior of the variance is driven by an asymmetric movement of the two tails rather than a uniform expansion of both tails.

In our view, this finding reiterates our earlier point that the L9010 or the variance are not ideal statistics to focus on when it comes to measuring higher-order earnings risk over the business cycle. In comparing patterns in higher order risk across gender,

[19] We have also run regression 1 using the standard deviation of earnings changes as our measure for overall dispersion instead, and the coefficients are small (0.07 (SIAB), –0.12 (SOEP)) and insignificant (t-stat of 0.42 (SIAB), –0.54 (SOEP)) in both data sets.

the magnitudes of the fluctuations in both Kelley skewness and the upper and lower tails separately are somewhat attenuated for women compared with men.

4.2 Differences across Groups of Workers

To shed light on the possible sources of cyclicality of higher-order income risk, we now examine whether it differs across observationally distinct groups. First we divide male and female workers into groups by education (college versus non-college graduates) or by private and public employment. These are two dimensions by which the three countries differ greatly. In Germany, 12% of men and 8% of women are college educated. In Sweden and the US, the equivalent numbers are 16 and 25 for males and 17 and 25 for females, respectively. Differences in the size of public sector employment are even larger. Defining public sector employment as employment in public administration, health care, and education (sectors which in Germany and Sweden are dominated by public sector jobs or by jobs funded by the public),[20] the share of public sector employment in Sweden is more than twice as large as in Germany or the US.[21] Moreover, public sector jobs are often thought of as less risky, offering generous employment protection and less volatile compensation, so it is interesting to ask if this is borne out in the data.

For each of these groups, we analyze higher-order income risk by first computing average (standardized) moments across years and countries by quartiles of (standardized) log GDP change as shown in Figures 3 and 4. The standardization of moments and log GDP changes is performed independently for each country before pooling across countries, which implies that a deviation from zero indicates a standardized deviation from the country-specific mean of the moment. For each quartile, the bars correspond to the average moment for (ordered from the left) the full sample, college graduates, non-college graduates, private employment, and public employment, respectively.

Figure 3 shows that the nature of income risk is qualitatively similar across all male subgroups: overall dispersion is acyclical (panel a), Kelley skewness is procyclical (panel b), the upper tail is procyclical (panel c), and the lower tail is countercyclical

[20] Formally, we classify a worker as working in the public sector, if he/she works in these sectors in both years t and $t-s$. Historically, most workers in these sectors were employed by the public; this is less true today.

[21] In Sweden, about 23% of men and 63% of women work in the public sector (these figures have been relatively stable over the considered time period). In Germany, a stable 10% of men work in the public sector, while the share of women steadily increased from about 23% to about 36% over the considered time period. In the US, 13% of males and 18% of females are employed in the public sector.

Figure 3: Higher-Order Moments by Quartiles of Log GDP Change: Males

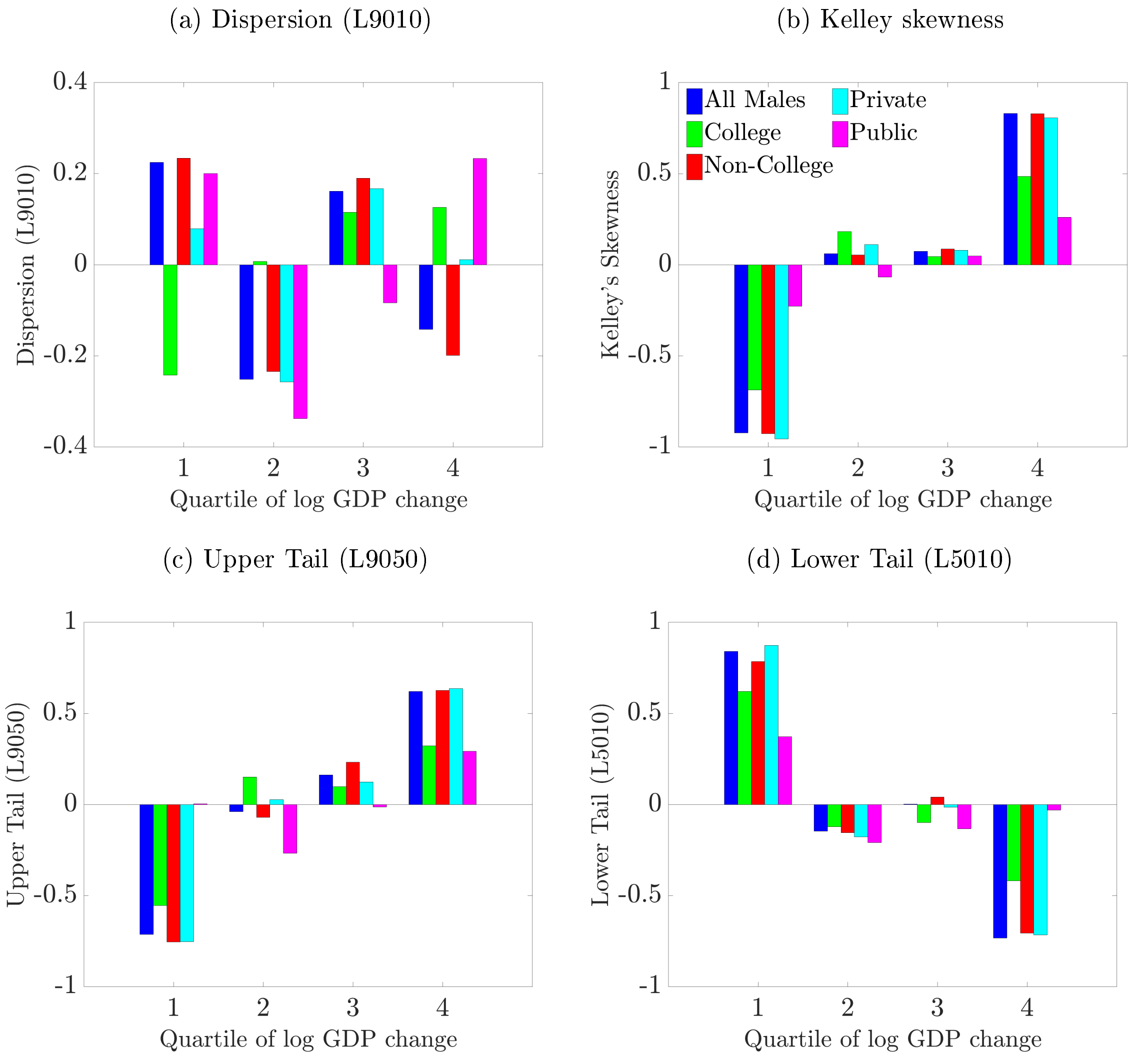

Note: For different samples, each bar shows the average moment across years and countries by quartiles of log GDP change. Both log GDP changes and moments are standardized by country.

(panel d). Figure 4 shows a similar picture for women and, as noted above, shows that fluctuations in earnings risk is somewhat attenuated relative to men. For both males and females, we see a strong asymmetric cyclical change of the distribution of earnings changes across groups.

For each group and country, we estimated our baseline regression (equation 1). The estimated sensitivity coefficients are displayed in Figure 5. (Further details are in in Appendix B; see tables B.1 through B.4). Each panel in the figure shows, starting

Figure 4: Higher-Order Moments by Quartiles of Log GDP Change: Females

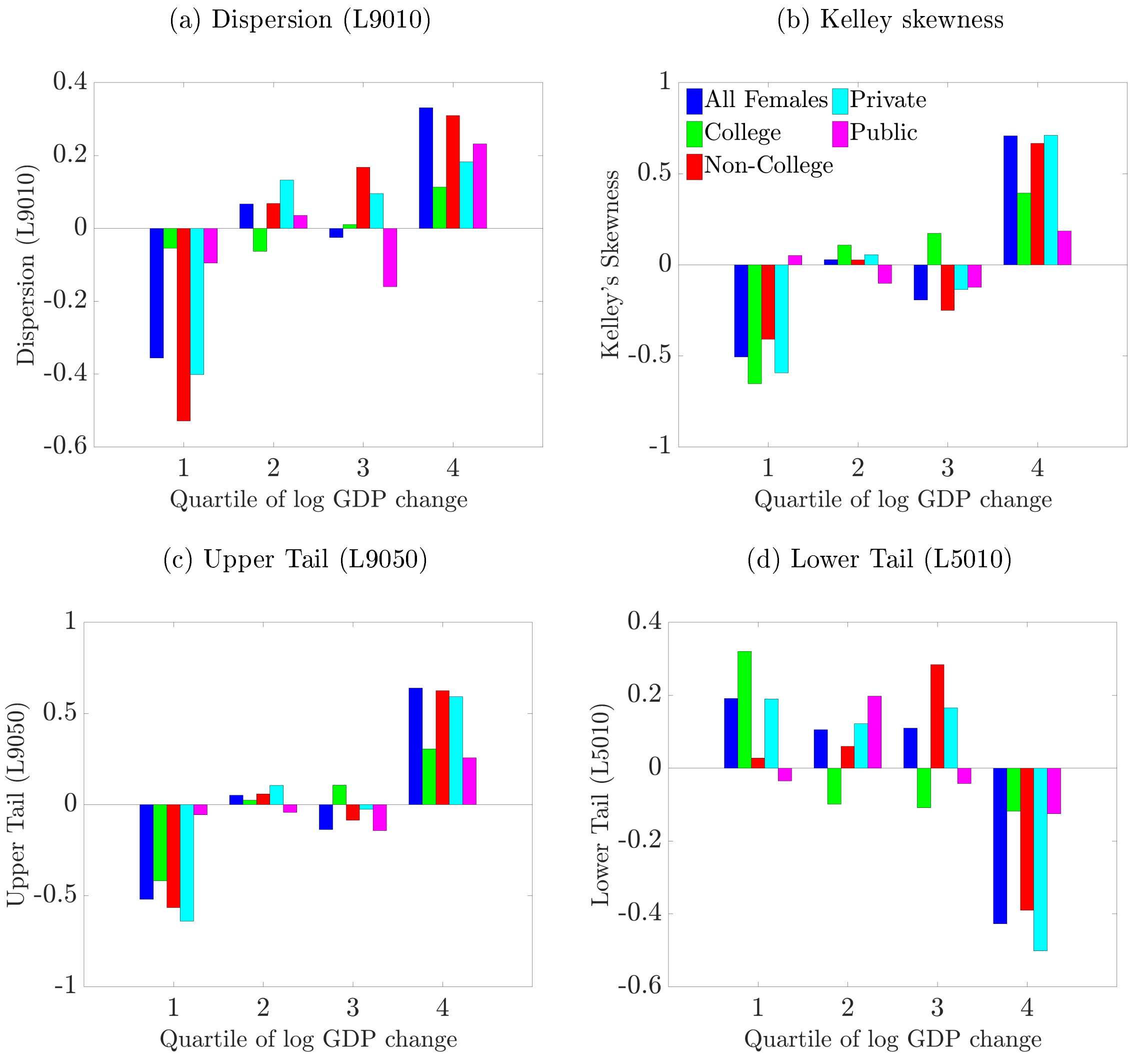

Note: See notes to figure 3.

from the left, the regression coefficients along with 95% confidence intervals for males (solid) in Sweden (red, triangles), Germany (green, squares), and the US (blue, bullets), followed by the equivalent regression coefficients for females (dotted). Within each country-gender grouping, the coefficients are (ordered from the left) those from the full sample, college graduates, non-college graduates, private employment, and public employment, respectively.

Figure 5 confirms the picture that emerged in Figures 3 and 4: higher-order earnings risk is similar across groups. However, we see some noteworthy differences. The magnitude of cyclicality is stronger for non-college graduates as compared to college

Figure 5: Cyclicality of Higher-Order Moments: Income vs Wages (Sweden, Germany (SIAB), and the United States)

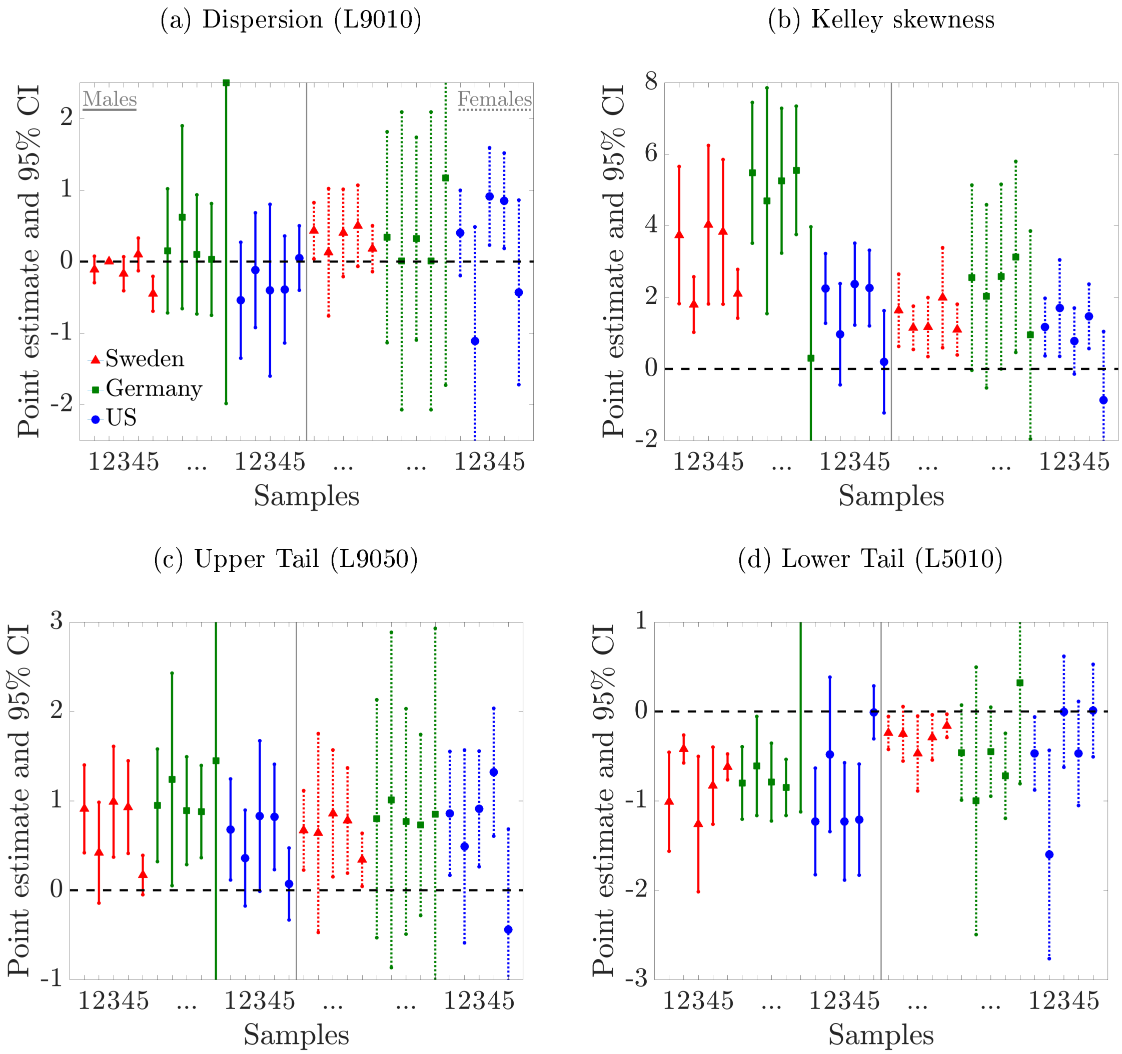

Note: The samples are (1) earnings: full sample , (2) earnings: college graduates, (3) earnings: non-college graduates, (4) earnings: private sector, (5) earnings: public sector. In each figure, the left (right) half shows the results for males (females). For details of samples, see text. For the regressions, see note to Table I. Each marker reports the coefficient on log GDP change.

graduates. The difference is particularly large for males in the US and Sweden, where the regression coefficient for Kelley skewness is about two to three times larger for non-college graduates (insignificant 0.97 vs. 2.37*** for the US and 1.80*** vs. 4.03*** for Sweden). Moreover, the magnitude of cyclicality for public sector workers is weaker in all countries—and insignificant in the cases of Germany and the US.

In Sweden, the procyclicality of Kelley's measure of earnings is lower for the public sector (2.10*** for males and 1.10*** for females) compared with the private sector (3.83*** for males and 1.99*** for females). For males, this is due to differences in the top tail; it compresses strongly for private sector employees, whereas it is acyclical in the public sector. The L5010 gap, on the other hand, fluctuates by comparable magnitudes for both groups. For women, the reduced cyclicality is due to both tails fluctuating slightly less.

Overall, it is somewhat surprising that for workers in the public sector in a country like Sweden with a reputation for high levels of public insurance, there is robust evidence of higher downside risk in recessions—compression of the top and expansion of the bottom—even if the magnitudes are somewhat smaller than in the private sector. This finding further strengthens the conclusion of this section that increasing downside earnings risk appears to be a robust feature of business cycles in developed countries.

Differences across Occupations

We now turn to occupations and explore the heterogeneity of cyclical earnings changes along this dimension. We are able to conduct this analysis for Germany; the SIAB provides time-consistent occupational codes based on the KldB-88, the 1988 version of the classification of occupations by the German Federal Employment Agency. We run the cyclicality regressions separately for each occupation, where a worker contributes to the earnings changes of occupation j from $t-1$ to t if in year $t-1$ he or she works in that occupation.

We first consider the most broad categories in the KldB-88, which defines five *occupational areas*: (1) farming, gardening, animal breeding, fishing, and similar occupations; (2) mining and mineral extraction; (3) manufacturing and fabrication; (4) technical occupations like engineering or laboratory work; and (5) service occupations. For each occupation, we ran the baseline regression (equation 1) of each moment on GDP growth. Detailed results can be found in Appendix B. The results are quite similar to those for the full sample. In each occupational category, the variance of income growth is acyclical. And for male workers in manufacturing occupations, technical occupations, and service occupations, skewness is procyclical, resulting from a procyclical upper tail and counter-cyclical lower tail. These same trends are present for female works but with less statistical significance, as in the full sample (see table B.5 in the appendix).

We then conduct a more disaggregated analysis—at the expense of relatively small sample sizes for some occupations—and rerun the regressions for 30 *occupational segments*. The same general patterns in variance, skewness, and the top and bottom tails of the distribution remain, but with significant variation. For example, the coefficient on Kelley skewness ranges from 2.36 to 17.87, which implies large differences in the asymmetric cyclical dynamics of income risk across occupations. Tables B.6 and B.7 report summary statistics of the distribution of coefficients across the 30 occupational segments.

Summing up, we find that broad occupational groups experience similar cyclicality, particularly in manufacturing, technical, and service occupations. Regressions at a finer level of disaggregation point towards heterogeneity of earnings cyclicality across occupations.

4.3 Cyclicality: Earnings versus Wages

A natural question that is raised by these results is whether the observed cyclicality of earnings changes can be attributed mainly to changes in wages or to increased risk of unemployment in economic downturns. We take advantage of the rich information on labor market attachment in the SIAB, in particular we exploit information on the duration of each employment spell and on whether it is a part-time or full-time job. Focusing on full-time workers, we analyze the cyclicality of the distribution of wage changes and compare the results to the ones on earnings changes. We define a worker as full time if his or her full-time spells add up to at least 50 weeks of employment in a given year. (A less strict definition of full-time workers as 45 weeks of employment does not change the results.) The wage variable is the average daily wage rate, where the average is taken over all full-time spells. The same measure has also been used in Dustmann *et al.* (2009) and Card *et al.* (2013).[22]

In Table II, rows 1 and 4 reproduce the results from Table I for completeness. The first set of new results are in rows 2 and 5: these report the cyclicality regressions using average daily wages instead of annual earnings. The main finding for both males and females is that the cyclicality of wages for full-time workers is remarkably similar to the cyclicality of earnings. Specifically, both measures of dispersion of wages are acyclical, as was the case for earnings, and the point estimates for both skewness measures are very close for wages and earnings.[23] Naturally, the dispersion of earnings changes is

[22]In Germany, a full-time worker is entitled to an annual vacation time of 4 to 6 weeks, which is counted as part of the employment spell.

[23]The sample of full-time female workers contains about 73% of women (who make up only 54% of the observations) that contribute to the measures of earnings changes for women. The corresponding figures are 88% of individuals and 82% of observations for males. This implies that part-time employment plays a more important role for the female sample.

Table II: Cyclicality of Individual Earnings vs. Wages; Germany (SIAB)

	L9010	Kelley	L9050	L5010
		Males		
Earnings	0.15	5.48***	0.95***	–0.80***
	(0.36)	(5.80)	(3.14)	(–4.11)
Full-Time Wages	–0.09	4.73***	0.30***	–0.39***
	(–0.54)	(6.31)	(3.77)	(–3.20)
Full-Time Wages	–0.12	4.98***	0.28***	–0.40***
(Firm Stayers)	(–0.81)	(5.78)	(3.29)	(–3.20)
		Females		
Earnings	0.34	2.55**	0.80	–0.46*
	(0.48)	(2.05)	(1.25)	(–1.80)
Full-Time Wages	0.03	2.12***	0.17**	–0.14
	(0.18)	(5.11)	(2.61)	(–1.58)
Full-Time Wages	0.02	2.28***	0.16***	–0.14
(Firm Stayers)	(0.13)	(4.84)	(3.17)	(–1.61)

Note: See notes for Table I.

wider than the distribution of wage changes, which is reflected by the point estimates on the tails (last two columns), which are about half as big for wage changes.

What remains is the question of what happens to the wages of workers that stay at the same firm. We therefore further restrict the sample to those workers that work at least 50 weeks for the same employer in two consecutive years.[24] The second set of new results is in rows 3 and 6: the cyclicality regressions for average daily wages for those workers who work at the same firm. The remarkable result is that even for those, we observe the same qualitative pattern of cyclicality of wage changes. By and large, these results strongly indicate that the cyclicality results are driven by changes in wages, and not by hours, even for full-time workers.

5 Introducing Insurance

We now turn to various sources of insurance available and gauge the extent to which they are able to mitigate downside risk over the business cycle.

[24]The sample of full-time female workers that do not switch firms contains about 61% of women (who make up about 40% of the observations) that contribute to the measures of earnings changes for women. The corresponding figures are 80% of individuals and 65% of observations for males.

5.1 Within-Family Insurance

In the previous section, we have shown that individual earnings risk over the business cycle is captured by higher-order moments. While it is important to understand the underlying nature of labor income risk and the systematic differences across groups, most of our samples are composed of individuals in cohabitation.[25] Assuming pooling of resources within the household, the relevant income measure for many economic decisions is the joint labor income in the household, not individual incomes. We therefore shift our attention to joint labor earnings at the household level in order to shed light on the role of informal insurance mechanisms within the household. As mentioned earlier, it is not possible to link individuals in SIAB, so for the household-level analysis of Germany we rely on SOEP data instead.

Several mechanisms that are potentially relevant for household income dynamics are at work simultaneously. An active insurance channel against income losses of one earner is spousal labor supply adjustments, both along the intensive and extensive margins—sometimes referred to as "added worker effect" (e.g., Blundell *et al.* (2016), Attanasio *et al.* (2005), or Pruitt and Turner (2018)). A passive insurance channel is simply the existence of two income streams as opposed to one. Male labor income on average constitutes 71%, 60%, and 62% of household earnings in the United States, Sweden, and Germany, respectively. Thus, if, for example, male income dropped in recessions while female income stayed constant, this would translate into household income react less to aggregate changes than male earnings. However, to the degree that spouses work in same regions, industries, or firms, and are exposed to similar cyclical income shocks, this channel is unlikely to provide a significant amount of insurance.

In order to assess the active insurance provided within households, we consider the cyclicality of income for actual households in comparison to income changes for randomly formed couples. Any endogenous response of spousal earnings is by construction not existent for synthetic couples. To the extent that each spouse in actual households endogenously responded to a shock to the other's income (i.e., a strong "added worker effect"), we would see household income fluctuate less compared with that of synthetic households.

We consider three sets of randomly formed couples. First, we randomly pair heads and spouses for each $t-1$ to t change. To each synthetic couple we apply the same

[25] Only 12% of our benchmark individual sample in the United States lives in a single-person household.

Table III: Cyclicality of Earnings Growth Moments: Actual vs. Synthetic Households

	L9010	Kelley	L9050	L5010
	United States			
Actual households	0.04	1.91***	0.81***	–0.78***
Synthetic households:				
Fully random match	0.09	1.33***	0.64***	–0.55***
Random | spouses' ages	–0.04	1.61***	0.71***	–0.75***
Random | age and educ.	–0.01	1.59***	0.72***	–0.73***
	Sweden			
Actual Households	–0.02	2.24***	0.50***	–0.52*
Synthetic households:				
Fully random match	–0.21***	1.72***	0.31***	–0.52***
Random | spouses' ages	–0.20***	1.76***	0.32***	–0.53***
Random | age and educ.	0.02	1.82***	0.46***	–0.43***
	Germany (SOEP)			
Actual households	–1.31***	1.88**	–0.05	–1.26***
Synthetic households:				
Fully random match	–0.99***	1.28***	–0.12	–0.87***
Random | spouses' ages	–1.15***	1.02**	–0.25	–0.89***
Random | age and educ.	–1.19***	1.01**	–0.28	–0.91***

Note: See notes for Table I. Maximum lag-length considered for Newey-West standard errors in case of SOEP is 3. The parameter for the synthetic couples is the mean over 250 bootstrap repetitions. The regression for Sweden with education starts in 1991.

selection criteria as for the actual households. We next control for some characteristics of the actual household formation, as sorting along those dimensions might in part explain higher cyclicality of actual household incomes. Thus, we make the synthetic couples more similar to actual households, while still isolating potential added worker effects. Specifically, we control for age (seven age groups) or for both age and education. Table III shows the regression coefficients for each country and each randomization. Overall, in all three economies, earnings changes of actual households are more cyclical than earnings changes of randomly paired couples: there is no clear sign of insurance against the cyclicality of individual incomes provided at the household level. This is in line with evidence in Pruitt and Turner (2018), who document in tax data from the United States that recessions are times in which female earnings growth is only weakly correlated with male earnings growth, and female labor supply adjustments along the extensive margin in reaction to male earnings losses are less pronounced (i.e., the "added worker effect" is weaker).

5.2 Government: Taxes and Social Insurance Policy

Focusing on the household as the relevant unit, we analyze the effectiveness of social policy in mitigating business-cycle risk in addition to any insurance arrangements made within households. We evaluate the total insurance effect of the tax and transfer system by analyzing the cyclicality of post-government earnings as compared to household gross earnings. In order to gain insights into the effectiveness of different policies, we then evaluate the relative importance of several subcomponents of transfers using the empirical tools employed in the previous analysis on artificial income measures that in turn add different transfers to household gross earnings.

The Overall Effect of the Tax and Transfer System

We begin with a brief discussion on the overall effect of the government, comparing the cyclicality of pre- and post-government measures of household earnings listed in rows 1 and 2 of Table IV. Figures 6 and 7 visualize the findings. We find that the tax schedule and social insurance policy are important sources of insurance against aggregate fluctuations in all three economies, with very similar overall effects.

Motivated by the considerations from the above sections, we directly consider the reactions of the upper and lower tails of income changes. In all three economies, downside risk is mitigated successfully by the tax and transfer system. In both the United States and Germany, the lower tail of post-government earnings changes is unresponsive to the business cycle–while significantly countercyclical for pre-government earnings. In Sweden, considering the point estimates, lower tail countercyclicality appears to be dampened but is still statistically significant (from a point estimate of –0.52 to –0.38).

Considering the cyclicality of the upper tail reveals differences between the countries. In Germany, it is unresponsive to the cycle for both pre- and post-government earnings. While both the US and Sweden display procyclicality in the L9050 of pre-government earnings changes, the L9050 of post-government earnings changes is acyclical in Sweden, but still procyclical in the United States. The different reactions of the tails translate into procyclical overall dispersion of post-government earnings changes in the US (with a t-statistic of 1.57) and countercyclical dispersion in Sweden. Taken together, the reaction of overall dispersion and tails results in procyclicality of Kelley's skewness measure for both countries, although the procyclicality is much smaller for post- than for pre-government earnings.

Figure 6: Standard Deviation and Skewness of Short-Run Earnings Growth: United States, Germany (SOEP), and Sweden

(a) United States, Std. Dev.

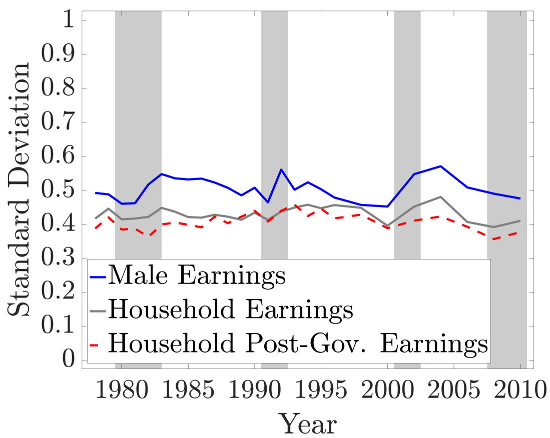

(b) United States, Kelley skewness

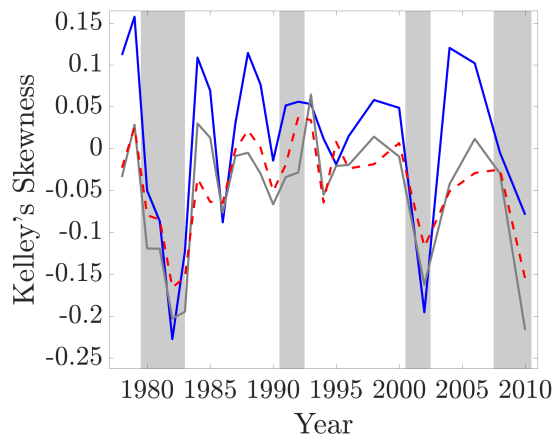

(c) Sweden, Std. Dev.

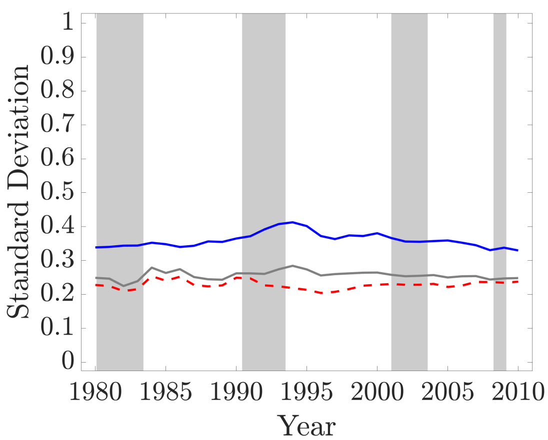

(d) Sweden, Kelley skewness

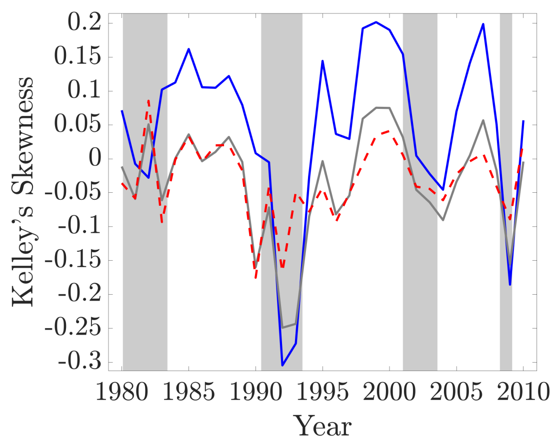

(e) Germany, Std. Dev.

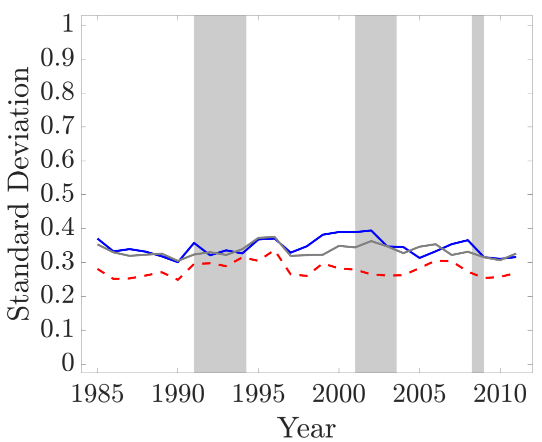

(f) Germany, Kelley skewness

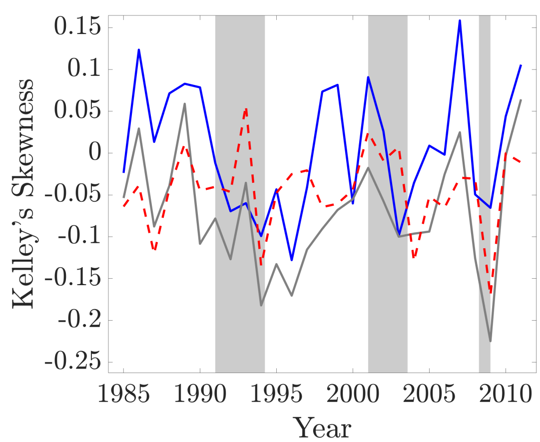

Note: See notes to Figure 2

Figure 7: Tails of Short-Run Earnings Growth: United States, Germany (SOEP), and Sweden

(a) United States, Upper Tail (L9050)

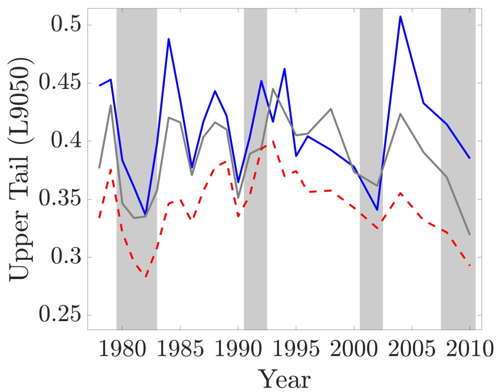

(b) United States, Lower Tail (L5010)

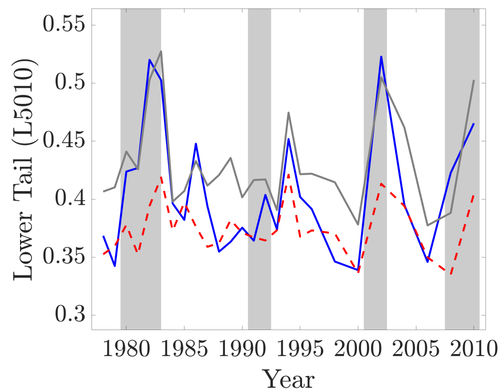

(c) Sweden, Upper Tail (L9050)

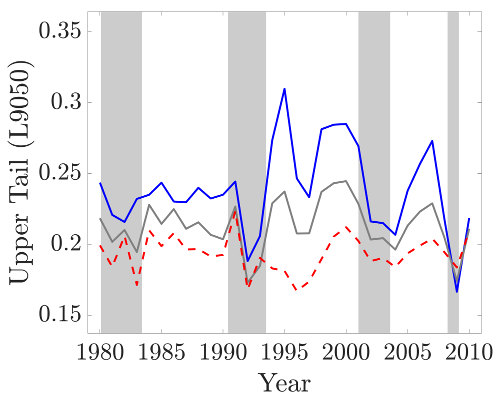

(d) Sweden, Lower Tail (L5010)

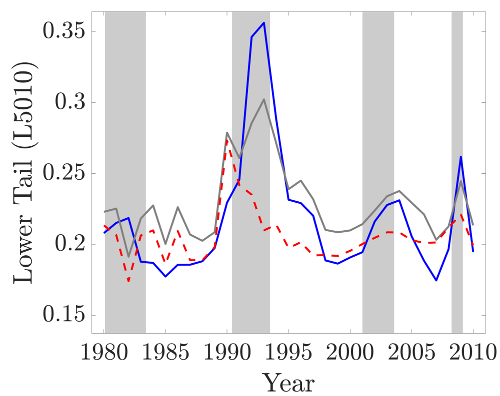

(e) Germany, Upper Tail (L9050)

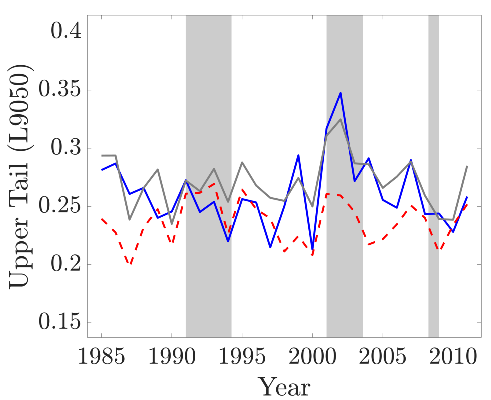

(f) Germany, Lower Tail (L5010)

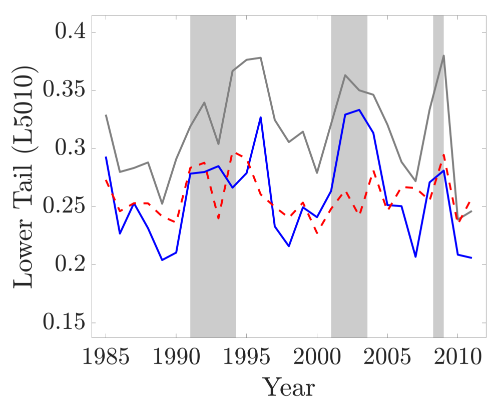

Note: See notes to Figure 2

Table IV: Business Cycle Variation in Household Income Change Distribution

	L9010	Kelley	L9050	L5010
	United States			
Pre-Gov	0.04	1.91***	0.81***	–0.78***
	(0.15)	(6.57)	(5.93)	(–3.78)
Post-Gov	0.34	1.09***	0.55***	–0.21
	(1.57)	(3.40)	(3.20)	(–1.43)
	Sweden			
Pre-Gov	–0.02	2.24***	0.50***	–0.52*
	(–0.08)	(3.33)	(4.94)	(–2.00)
Post-Gov	–0.41*	0.94**	–0.03	–0.38**
	(–2.00)	(2.38)	(–0.44)	(–2.33)
	Germany (SOEP)			
Pre-Gov	–1.31***	1.88**	–0.05	–1.26***
	(–3.60)	(2.68)	(–0.18)	(–4.26)
Post-Gov	–0.18	0.66	0.07	–0.25
	(–1.09)	(0.85)	(0.32)	(–1.28)

Note: See notes for Table I. Maximum lag length considered for Newey-West standard errors in case of SOEP is 3.

To sum up, the analysis suggests that asymmetric risk in recessions is mitigated, though not completely eliminated, by taxes and transfers. In all countries, the tax and transfer system plays an important role at insuring away downside risk. In Sweden, an additional effect is lowered upside risk in expansions.

Components of Government Social Insurance

The measure of post-government earnings used so far lumps many different transfers received and taxes paid by households. While this measure is appropriate for assessing the overall effect of the tax and transfer system, it is not as well suited for understanding the success of different social policies that specifically aim at mitigating downside risk or aiding low-income families, who can be expected to be especially vulnerable in recessionary periods. Therefore, we now consider different types of transfers separately.

We consider three main groups of transfers that are comparable across countries and for each country are consistently measured over time. The groups are (1) labor-market-related policies, (2) aid to low-income families, and (3) pension payments.[26]

[26]The components are measured as follows.

Labor-market-related policies mainly consist of unemployment benefit payments. This component of social insurance policy is of particular importance for the mitigation of increased downside household earnings risk in recessions, if the nature of downside risk is (temporary) job loss of household head or spouse.

The second component considered, aid to low-income families, consists of several measures of social insurance policies specifically aimed at at-risk households. The relevance of this type of transfer can therefore be expected to matter most for low-income households who have a higher likelihood of satisfying at-risk criteria in the course of a recession. The third component, pension payments and disability insurance, is not by construction related to the business cycle. The business cycle can, however, still play a relevant role for household members, who may take up pension payments or disability insurance instead of unemployment payments if they decide to leave the labor market upon job loss.

The results of the cyclicality analysis are listed in Table V. As for the estimates of total taxes and transfers, we compare the coefficients to the ones from the household gross earnings analysis in the first row of each panel in Table IV.[27] The estimates suggest that out of the three groups of transfers, only labor-market-related transfers (which have unemployment benefits as the main component) play a role in the reduction of downside risk. However, the lower tail remains significantly cyclical in all three countries, which is indicative of a major role played by the tax system (in combination with the transfers). The other two components of transfers do not have any impact on cyclicality as measured by our cyclicality regressions: for all three economies, the point estimates when adding aid to low-income families or pensions are almost identical to the ones for gross earnings.

In Germany, we additionally look at individual-level incomes in our larger sample based on the SIAB database. Besides individual earnings, SIAB also contains infor-

"Labor-market-related policies" in all three data sets are unemployment benefits; in LINDA additionally labor market programs; in PSID additionally workers' compensation.

"Aid to low-income families": LINDA: family support, housing support, cash transfers from the public (no private transfers); SOEP: subsistence allowance, unemployment assistance (before 2005), unemployment benefits II (since 2005); PSID: Supplemental Security Income; Aid to Families with Dependent Children (AFDC); Food Stamps; Other Welfare.

"pension payments: LINDA: (old-age) pensions; SOEP: combined old-age, disability, civil service, and company pensions; PSID: combined (old-age) social security and disability (OASI).

[27]Recall that in order to be in the year t base sample for the analysis, the lowest considered income measure of a household needs to be above the income threshold for that year. This way, we ensure that the sample is stable at the lower end of the distribution and that the results are not driven by low-income households entering the sample for one type of transfer but not for another type.

Table V: Cyclicality of Household Earnings: Transfers Added Separately

Household Earnings	L9010	Kelley	L9050	L5010
	United States			
+ Labor transfers	0.23	1.56***	0.73***	-0.50***
	(1.12)	(5.73)	(4.84)	(-3.33)
+ Aid to low-income	0.04	1.86***	0.80***	-0.75***
	(0.19)	(6.09)	(6.25)	(-3.66)
+ Pensions	0.04	1.69***	0.73***	-0.68***
	(0.20)	(5.52)	(5.60)	(-3.28)
	Sweden			
+ Labor transfers	–0.22	1.14***	0.13*	–0.35**
	(–1.23)	(4.23)	(2.04)	(–2.58)
+ Aid to low-income	–0.07	2.11***	0.42***	–0.49**
	(–0.38)	(3.72)	(4.51)	(–2.47)
+ Pensions	–0.07	2.34***	0.48***	–0.55**
	(–0.43)	(3.55)	(4.50)	(–2.68)
	Germany (SOEP)			
+ Labor transfers	–1.09***	1.34**	–0.13	–0.96***
	(–2.96)	(2.50)	(–0.60)	(–3.65)
+ Aid to low-income	–1.32***	1.66**	–0.11	–1.21***
	(–3.82)	(2.40)	(–0.47)	(–4.08)
+ Pensions	–1.21***	1.80***	–0.04	–1.17***
	(–3.30)	(3.10)	(–0.18)	(–4.58)

Note: See notes for Tables I and IV.

mation on unemployment benefits at the individual level. Table VI shows results for individual-level regressions for male and female earnings separately, when unemployment benefits are excluded (rows 1 and 3) and included (2 and 4). These individual level results line up well with the household level analysis conducted using SOEP data; labor market transfers has some, but limited, effect in mitigating the cyclicality. Together with the household-level analysis, this suggests that the German tax system (or interaction between taxes and transfers) is a primary reason for post-government earnings being acyclical.

Table VI: Cyclicality of Individual Earnings Including Unemployment Benefits in Germany (SIAB)

	L9010	Kelley	L9050	L5010
Male earnings	0.11	5.71***	0.97***	–0.86***
	(0.26)	(5.32)	(2.93)	(–4.40)
+Unempl. benefits	0.15	5.12***	0.84**	–0.70***
	(0.34)	(5.24)	(2.61)	(–4.01)
Female earnings	0.46	2.69*	0.89	–0.44*
	(0.60)	(1.92)	(1.26)	(–1.74)
+ Unempl. benefits	0.50	2.43*	0.82	–0.32
	(0.67)	(1.82)	(1.22)	(–1.43)

Note: See notes for Table I. Differences between estimates in Table I are due to regressions starting in 1981 instead of 1976.

6 Welfare Analysis

In this section, we quantitatively investigate how successful government policy is in insuring households against business-cycle fluctuations of earnings risk—and how much households value this insurance. For that purpose, we estimate income processes for pre-government household labor income and, separately, for post-government household income. The process is specified flexibly as a mixture of normals with time-varying moments to allow for cyclical higher-order risk.

Our estimation is based on earnings data, so we use a model to simulate the consumption profiles of households facing either pre- or post-government income streams. Specifically, we use a variant of the partial insurance model by Heathcote *et al.* (2014) to quantify the welfare gains. In the model, there is full insurance against transitory shocks and partial insurance against permanent shocks to income. We assume that there is no additional insurance against permanent shocks to post-government household income, and then calculate the corresponding degree of partial insurance against pre-government household income shocks. We also reinterpret this measure in terms of consumption-equivalent variation.

We pursue this analysis for Sweden only, because we need both pre- and post-government income to be captured well by the specific process we use in the estimation and which we feed into the model. Although the income process we choose is quite rich and flexible, the estimated models for Germany and the US fail to match some of the higher-order moments we target. Without a very good fit, the welfare effects we measure will be affected by the relative differences in the fit of the models to pre- and post-government household income moments. Of course, the same point applies to Sweden as well, but the fit is significantly better than the other two countries giving us some confidence that the measured welfare figures are meaningful. Before going to the results of the model-based analysis in Section 6.2, the next subsection discusses the estimation of the income process.

6.1 Estimation of Pre- and Post-Government Income

Let Y_t denote household earnings in period t, and define $y_t \equiv \log Y_t$. We assume y_t evolves according to the following process (for expositional reasons, we do not indicate pre- and post-government):

$$\begin{aligned} y_t &= z_t + \xi_t \\ z_t &= z_{t-1} + \zeta_t \end{aligned} \tag{3}$$

where ξ_t is an *iid* transitory shock, drawn from a mixture of two normals $\mathcal{N}(\bar{\mu}_\xi, \sigma^2_{\xi,i})$, i=1,2, with probabilities $p_{\xi,i}$ and $1-p_{\xi,i}$, respectively, $\bar{\mu}_\xi$ is chosen such that $\mathbb{E}\left(e^{\xi}\right) = 1$, and ζ_t denotes a permanent shock with time-varying and business-cycle-dependent distribution, modeled as in McKay (2014). This specification allows the process to match excess kurtosis found in the data.

In particular, ζ_t follows a mixture of three normals $\mathcal{N}(\bar{\mu}_{\zeta,t} + \mu_{\zeta,i} - \phi_i x_t, \sigma^2_{\zeta,i})$, with respective probability $p_{\zeta,i}$, $i = 1, 2, 3$, where $\sum_{i=1}^{3} p_{\zeta,i} = 1$, x_t is standardized log GDP growth, and $\bar{\mu}_{\zeta,t}$ is chosen such that $\mathbb{E}(e^{\zeta}) = 1$. We use GDP growth as the empirical measure of aggregate fluctuations in order to make the quantitative results easily interpretable in relation to the empirical estimates shown in Section 4. The parameters ϕ_i determine how much of aggregate risk is transmitted to idiosyncratic earnings risk and are estimated alongside the other parameters that characterize the distributions of the shocks.

We estimate the set of parameters

$$\chi = (\sigma_{\xi,1}, \sigma_{\xi,2}, p_{\xi,1}, \mu_{\zeta,2}, \mu_{\zeta,3}, \sigma_{\zeta,1}, \sigma_{\zeta,2}, p_{\zeta,1}, p_{\zeta,2}, \phi_2, \phi_3)$$

by simulated method of moments (SMM).[28] We target the time series of L9050 and L5010 of the 1-, 3-, and 5-year earnings change distributions, the average of the Crow-

[28]For identification purposes, we impose $\mu_{\zeta,2} \geq 0$, $\mu_{\zeta,3} \leq 0$, and $\phi_1 = 0$. With this assumption,

Siddiqui (1967) measure of kurtosis of 1-, 3-, and 5-year changes,[29] as well as the age profile of the cross-sectional variance from ages 25 to 60. Table E.1 in Appendix E shows the parameter estimates for pre- and post-government income. Appendix E includes the comparison between the simulated moments at these parameters and the empirical moments as well as further details of the estimation.

6.2 Quantitative Model

Shocks differ in their nature: some shocks are insurable, while others are not. In order to assess the welfare implications of the existing tax and transfer system, we map the estimated earnings process into a quantitative framework. Heathcote *et al.* (2014) set up a model populated by a continuum of islands, each of which is in turn populated by a continuum of agents. Two types of shocks exist in their economy: one common to all members of an island and the other specific to an individual. An island refers to a group of agents that are described by the same history of uninsurable shocks. Islands can be thought of as a network of family members, who perfectly share the risks faced by each individual. If, for example, all family members work in the same industry and live in the same region, there will be shocks that hit every member equally and hence cannot be insured within the family network.

Importantly for the quantitative analysis, there is no need to define empirical counterparts to the model islands. Given some assumptions on the market structure, outlined below, Heathcote *et al.* (2014) show the existence of a non-trade equilibrium in the spirit of Constantinides and Duffie (1996). In this equilibrium, there is no asset trade across islands, while agents within an island insure themselves perfectly against the individual-specific shocks. This reflects insurable (within-island) and uninsurable (island-level) shocks. A major advantage of their framework is that it allows for an analytical solution of an incomplete markets model. Crucially for us, this result does not depend on any distributional assumption of the shocks.

Model Structure

We employ a version of their model in which we abstract from endogenous labor supply. We also stay agnostic about the specific functional form of the tax and transfer system. Instead, we confront the model agents with the estimated pre-government earnings process and derive the implied consumption profile faced by expected utility

the time-varying means of the three mixtures will control the center, right tail, and left tail of the distribution of ζ, respectively. For practical purposes, we further assume $p_{\zeta,2} = p_{\zeta,3}, \sigma_{\zeta,2} = \sigma_{\zeta,3}$.

[29] The Crow-Siddiqui measure is a robust percentile-based measure of kurtosis (see appendix E).

maximizers, whose only choice (on top of engaging in asset trade) is consumption. We then consider the alternative world in which agents face the estimated post-government earnings process and derive their consumption profiles. Because some of the permanent shocks are insured while others are not, we need to slightly adjust notation of the income process.

Specifically, income is assumed to follow

$$
\begin{aligned}
y_t &= \alpha_t + \varepsilon_t \\
\alpha_t &= \alpha_{t-1} + \omega_t \\
\varepsilon_t &= \kappa_t + \theta_t \\
\kappa_t &= \kappa_{t-1} + \eta_t \\
\theta_t &\sim F_{\theta,t} \\
\eta_t &\sim F_{\eta,t} \\
\omega_t &\sim F_{\omega,t}
\end{aligned}
\tag{4}
$$

where α_t is the "island-specific" component, which is common to a continuum of agents, and ε_t is the "individual" component, which in turn has a permanent part κ_t and a transitory part θ_t.

Agents live finite lives. Each period a mass $(1-\delta)$ of newborns enters the economy with age 0. The probability of survival from age a to age $a+1$ is constant at δ. Newborn agents maximize discounted lifetime utility. For the per period utility function, we use log utility: $u(c_t) = \log(c_t)$.

Age 0 agents entering in year τ hold zero financial wealth and are allocated to an island of agents that then share the same sequence of uninsurable shocks $\{\omega_t\}_{t=\tau}^{\infty}$. Within islands, agents can trade a full set of state-contingent claims for individual-specific shocks. Between-island trading contracts are non-existent.

In equilibrium, log consumption and consumption change are given by[30]

$$
\log c_t(\mathbf{x}_t, \varepsilon_t) = \alpha_t + \log \int \exp(\varepsilon_t)\, dF^a_{\varepsilon,t}, \tag{5}
$$

$$
\Delta \log c_t = \omega_t + \left(\log \frac{\int \exp(\eta_t)\, dF_{\eta,t} \int \exp(\theta_t)\, dF_{\theta,t}}{\int \exp(\theta_{t-1})\, dF_{\theta,t-1}} \right). \tag{6}
$$

[30]The derivation of consumption outlined in Heathcote *et al.* (2014) carries over one to one to our simplified version of their model, which abstracts from the tax function and endogenous labor supply.

Note that the uninsurable shock ω_t translates one for one to consumption. The individual realizations of the two insurable shocks, however, do not affect consumption: given perfect risk-sharing, all members of an island consume the mean realization of these shocks.

Distribution of Shocks

In order to use the model framework for our analysis, we need to translate the estimated earnings process (3) into the process specified in (4). The transitory component ξ_t in (3) directly translates into θ_t from (4). Permanent earnings changes, ζ_t, in (3) are drawn from a mixture of three normal distributions. In specification (4), the overall permanent earnings change is given by $(\omega_t + \eta_t)$, the insurable (individual-level) and the uninsurable (island-level) part. We assume that both types of permanent shocks are drawn from time-varying mixture distributions. We scale the estimated parameters of the permanent shocks such that the variance of η_t (ω_t) is equal to a fraction λ (1-λ) of the overall variance of the permanent shock ζ. Specifically, $\eta_t \sim \mathcal{N}(\lambda^{1/2}\mu_{\zeta,i,t}, \lambda\sigma^2_{\zeta,i})$ with probability $p_{\zeta,i}$ for i=1,2,3 and $\omega_t \sim \mathcal{N}((1-\lambda)^{1/2}\mu_{\zeta,i,t}, (1-\lambda)\,\sigma^2_{\zeta,i})$ with probability $p_{\zeta,i}$.

This scaling implies that the first three moments of η_t and ω_t are given by $E[\eta_t] = \lambda^{1/2}E[\zeta_t]$, $var[\eta_t] = \lambda var[\zeta_t]$, and $skew[\eta_t] = skew[\zeta_t]$ (for ω replace λ with $1-\lambda$); see Appendix F for further details.[31] In this setup, λ is a measure of the degree of partial insurance against permanent shocks: it measures the share of the total variance of permanent shocks that is accounted for by the insurable component.

6.3 Quantitative Results

For a given degree of partial insurance, we simulate income and consumption profiles for a large number of agents. Agent i's expected discounted lifetime utility when facing the pre- or post-government income streams is given by

$$U_i^j\left(\left\{c_{i,a}^j\left(\lambda^j\right)\right\}_a\right) = \sum_a (\beta\delta)^{a-1}\, u\left(c_{i,a}^j\left(\lambda^j\right)\right),$$

for $j = pre, post$, and $c_{i,a}^j$ is consumption of agent i at age a when facing income stream j given that the degree of partial insurance is λ^j. Utilitarian welfare is given by $W^j\left(\lambda^j\right) = \sum_i U_i^j\left(\left\{c_{i,a}^j\left(\lambda^j\right)\right\}_a\right)$.

Now consider the following experiment. Agents face either the pre- or the post-government income process. We assume that permanent shocks to post-government

[31] We then adjust the cross-sectional mean, such that $E[exp(\eta)] = E[exp(\omega)] = 1$.

income translate one for one into consumption changes; that is there is no insurance against permanent shocks: $\lambda^{post} = 0$. This assumption is motivated by empirical results in Blundell *et al.* (2016), who find that the degree of partial insurance on top of government and family transfers is very close to zero. Now, we search for the degree of partial insurance against permanent shocks to the pre-government income process such that the agents are indifferent ex ante between the post-government income process with $\lambda^{post} = 0$ and the pre-government process with $\lambda^{pre} > 0$. The term λ^{pre} can thus be interpreted as the degree of partial insurance provided by the government under the assumption that there is no additional partial insurance. For the per-period utility function, we choose log utility.

For reasons discussed above we only perform this analysis for Sweden. We find $\lambda^{pre} = 0.43$, which means that the existing tax and transfer schedule in Sweden corresponds to insuring households against 43% of permanent shocks to household labor income, as shown in Table VII.

In order to assess the magnitude of this degree of partial insurance in terms of welfare, we use the model to calculate the consumption equivalent variation (CEV) that makes agents in the world with the pre-government income stream and no partial insurance indifferent to the world with the pre-government income stream and partial insurance of the size given by λ^{pre}. The 43% partial insurance translates into a CEV of 14.3% when assuming log utility. Hence, the existing tax and transfer system provides sizable insurance. Note that this calculation abstracts from any first-order effects: both the level effect of the tax and transfer system on average income and the variation in average income changes are taken out of the equation.

When interpreting these results, it is important to notice that government policy not only reduces the cyclicality of shocks but also reduces the overall level of cross-sectional dispersion. In order to filter out the effect of this change of "initial variance," we impose in a second run of the same experiment that the cross-sectional variance at age 25 (when agents are born in the model) is the same as for the pre-government process. As expected, this step takes away some of what is measured as insurance before. After adjusting the level of income variance at age 25, we get $\lambda^{pre} = 0.06$. The overall welfare gain of moving from the pre- to the post-government income stream adjusted to the same initial variance is 6%, which translates into a CEV of about 1.3%.

Given this already sizable insurance, what is the scope of additional government policy as a means of insurance against cyclical risk? In order to approach this question, we consider the same experiment for a counterfactual income process. Assume that on top of what the government already does, cyclicality is completely shut down for the post-government income stream. For this experiment, we set $\phi_2 = \phi_3 = 0$, thus imposing the distribution of idiosyncratic income changes that corresponds to periods of average GDP growth. This implies an even stronger degree of insurance of about 64% (or 27% when adjusting for initial variance). Considering the CEV connected to those insurance parameters, the scope of additional insurance is sizable: through the lens of the model, when adjusting for initial variance effects, an additional welfare gain of about 4.6% is possible, abstracting from first-order effects.

Table VII: Welfare Gains of the Tax and Transfer System

Scenario	λ^{pre}	CEV	λ^{pre} (adjustment)	CEV (adjustment)
Pre to Post	43%	14.26%	6%	1.27%
Pre to Post*	64%	17.53%	27%	5.91%

Note: The term λ^{pre} denotes the degree of partial insurance against permanent shocks; the per-period utility function is $\log(c)$. * indicates that the cyclicality of the permanent shocks is shut down. See text for details on the scenarios. The CEV columns denote the corresponding consumption equivalent variation associated with the change from the world with the pre-government income stream and no partial insurance to a world with the pre-government income stream and partial insurance of the size given by λ^{pre}.

Digging deeper into the structure of the model economy, Figure 8 shows moments of the cross-sectional consumption distribution for the cohort that lives through the Swedish macroeconomic history. Considering the variance that accumulates over time, it is apparent that in the two recessions of the early 1990s and the late 2000s the distribution becomes more dispersed. If no government insurance had been available for Swedish households, this increase would have been very strong. During the first three years of the 1990s, the distribution widens by more than in the whole preceding decade. Considering the accumulated consumption variance that is realized given the actual post-government income stream, this increase is mitigated, as can be seen even more directly by considering the cross-sectional variance when the post-government income stream is adjusted for the initial variance difference.

In panel b of Figure 8, we see how the asymmetric cyclical dynamics of income changes translates into asymmetry of the cross-sectional consumption distribution. The clear pattern is that the steep increases of dispersion do not come in a symmetric way. Instead we see that the distribution clearly becomes more left-skewed as measured by Kelley skewness. Zooming in on the tails, L9050 and L5010, establishes that the steep

Figure 8: Moments of Cross-Sectional Consumption Distribution

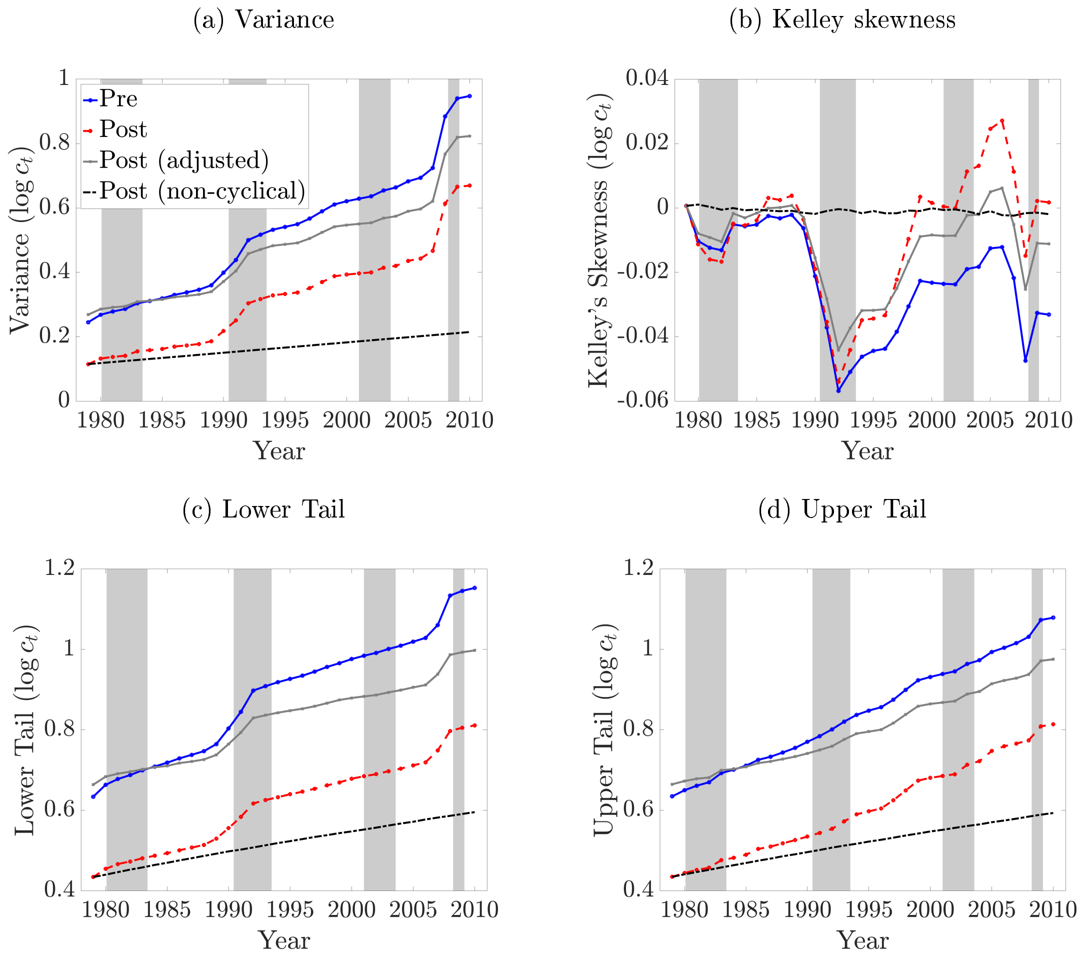

Note: Each figure shows a moment of the simulated cross-sectional consumption distribution for a cohort that lives through the Swedish macroeconomic history and faces, (i), the estimated pre-government income process; (ii), the estimated post-government income process; (iii), the post-government income process adjusted for initial variance; or, (iv), the post-government income process that eliminates all cyclicality of the distribution of shocks.

increase of dispersion comes from the lower tail widening more in the two recessionary periods as compared to the other time periods.

When the cyclicality of the post-government income process is shut down ($\phi_2 = \phi_3 = 0$), the variance increases only mildly and does not change toward a negatively skewed distribution in recessionary periods. These elements together explain the potential welfare gain of further reducing cyclicality.

7 Conclusion

This paper has characterized how higher-order income risk varies over the business cycle, as well as the extent to which such risk can be smoothed within households or with government social insurance policies. We have studied panel data on individuals and households from the United States, Germany, and Sweden, covering more than three decades for each country. This allowed us to take a broad perspective when approaching the two sets of questions raised in the introduction. One, what is the precise nature of idiosyncratic income risk, and how does it change in recessions? And two, how successful are various ways by which individual income fluctuations are mitigated in an economy, which prevents these fluctuations from affecting an individual's consumption?

We documented first that the underlying variation in higher-order risk is similar across these countries that differ in many details of their labor markets. In particular, in all three countries, the variance of earnings changes is almost entirely constant over the business cycle, whereas the skewness becomes much more negative in recessions. We further showed that these general patterns hold true for different groups defined by education, gender, public- versus private-sector jobs, and occupation.

Second, within-household smoothing appears to be not effective at mitigating individual-level business cycle fluctuations in skewness. It is worth emphasizing that this does not contradict the existence of family insurance in general. Instead, it points towards family insurance reaching its limits in particularly hard times. It is consistent with a lower ability of each spouse to respond to the other's income change in recessions. Also, to the extent that spouses work in, e.g., the same regional labor market, or industry they can be expected to be exposed to similar semi-aggregate shocks. The detailed evaluation of this channel is on our agenda and left for future research.

Third, government-provided insurance—unemployment insurance, aid to low-income households, social security benefits, among other transfers and taxes—plays an important role in reducing the cyclicality of downside risk in all three countries. An interesting assessment that is out of the scope of this paper would be to quantify the relative roles of automatic stabilizers, active expansions of the social safety net, and tax progressivity (which could be an important driver of changes in the upper tail of income changes). Through the lens of a structural model with partial insurance against permanent income shocks, the degree of overall insurance provided by the existing tax and transfer system amounts to 14% in consumption-equivalent terms in

Sweden. After isolating the gains from a lower initial variance at age 25, the insurance gain amounts to about 1.3%. However, the remaining risk in post-government household-level income is still substantial: individuals are willing to pay 4.6% of their consumption to completely eliminate procyclical fluctuations in skewness. Of course, this exercise leaves open the question, which policy could achieve such a smoothing. An analysis of explicit policies is out of the scope of this paper and is left for further research.

References

Angelopoulos, K., Lazarakis, S. and Malley, J. (2018). *Time variation in labour income risk in Great Britain*. Working paper, University of Glasgow.

Attanasio, O., Low, H. and Sanchez-Marcos, V. (2005). Female labor supply as insurance against idiosyncratic risk. **3**.

Bayer, C. and Juessen, F. (2012). The life-cycle and the business-cycle of wage risk - cross-country comparisons. *Economics Letters*, **117** (3), 831 – 833.

Blass-Hoffman, E. and Malacrino, D. (2016). *Employment Time and the Cyclicality of Earnings Growth*. Tech. rep., Department of Economics, Stanford University.

Blundell, R., Pistaferri, L. and Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *American Economic Review*, **106** (2), 387–435.

Busch, C. and Ludwig, A. (2018). Higher-order income risk over the business cycle: A parametric approach.

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.

Constantinides, G. M. and Duffie, D. (1996). Asset pricing with heterogeneous consumers. *The Journal of Political Economy*, **104** (2), 219–240.

— and Ghosh, A. (2014). *Asset Pricing with Countercyclical Household Consumption Risk*. Working paper, University of Chicago.

Crow, E. L. and Siddiqui, M. M. (1967). Robust estimation of location. *Journal of the American Statistical Association*, **62** (318), 353–389.

Cunha, F., Heckman, J. and Navarro, S. (2005). Separating uncertainty from heterogeneity in life cycle earnings. *Oxford Economic Papers*, **57** (2), 191–261.

Daly, M., Hryshko, D. and Manovskii, I. (2014). Reconciling estimates of earnings processes in growth rates and levels, manuscript.

Dustmann, C., Ludsteck, J. and Schönberg, U. (2009). Revisiting the german wage structure. *Quarterly Journal of Economics*, **124** (2), 843–881.

Fitzenberger, B., Osikominu, A. and Völter, R. (2006). Imputation rules to improve the education variable in the IAB employment subsample. *Schmollers Jahrbuch. Zeitschrift für Wirtschafts- und Sozialwissenschaften*, **Jg. 126**.

Fuchs-Schündeln, N., Krueger, D. and Sommer, M. (2010). Inequality trends for germany in the last two decades: A tale of two countries. *Review of Economic Dynamics*, **13** (1), 103–132.

Golosov, M., Troshkin, M. and Tsyvinski, A. (2016). Redistribution and social insurance. *American Economic Review*, **106** (2), 359–386.

Guvenen, F. (2007). Learning your earning: Are labor income shocks really very persistent? *American Economic Review*, **97** (3), 687–712.

— (2011). Macroeconomics with heterogeneity: A practical guide. *Federal Reserve Bank of Richmond Economic Quarterly*, **97** (3), 255–326.

—, Ozkan, S. and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, **122** (3), 621–660.

— and Smith, A. A. (2014). Inferring Labor Income Risk and Partial Insurance from Economic Choices. *Econometrica*, **82** (6), 2085–2129.

Harmenberg, K. and Sievertsen, H. H. (2018). *The Labor-Market Origins of Cyclical Income Risk*. Working paper, Institute for International Economic Studies.

Heathcote, J., Storesletten, K. and Violante, G. L. (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*.

Kaplan, G., Moll, B. and Violante, G. L. (2016). *Monetary Policy According to HANK*. Working paper, Princeton University.

Kurman, A. and McEntarfer, E. (2017). *Downward Wage Rigidity in the United States: New Evidence from Administrative Data*. Working paper.

Mankiw, N. G. (1986). The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics*, **17** (1), 211–219.

McKay, A. (2014). *Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics*. Working paper, Boston University.

OECD (2016). Economic policy reforms 2016: Going for growth interim report. OECD Publishing.

Pistaferri, L. (2001). Superior information, income shocks and the permanent income hypothesis. *Review of Economics and Statistics*, **83** (3), 465–476.

Pora, P. and Wilner, L. (2017). The individual dynamics of wage income in france during the crisis. *Economie et Statistique / Economics and Statistics*, (494-496), 179–199.

Pruitt, S. and Turner, N. (2018). *The Nature of Household Labor Income Risk*. Working paper.

Schmidt, L. D. W. (2016). *Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk*. Working paper, University of Chicago.

Shin, D. and Solon, G. (2011). Trends in men's earnings volatility: What does the panel study of income dynamics show?,. *Journal of Public Economics*, **95** (7-8), 973–982.

Skedinger, P. (2007). *The Design and Effects of Collectively Agreed Minimum Wages: Evidence from Sweden*. Working Paper Series 700, Research Institute of Industrial Economics.

Storesletten, K., Telmer, C. I. and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, **112** (3), 695–717.

vom Berge, P., Burghardt, A. and Trenkle, S. (2013). *Sample-of-Integrated-Labour-Market-Biographies Regional-File 1975-2010 (SIAB-R 7510)*. Tech. rep., Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB).

Appendices

A Data Appendix

This appendix briefly describes the variables used for each of the data sets and lists the numbers of observations after the sample selection steps.

A.1 PSID

Variables

Demographic and Socioeconomic

Head and Relationship to Head. We identify *current* heads and spouses as those individuals within the family unit with `Sequence Number` equal to 1 and 2, respectively. In the PSID, the man is labelled as the household head and the woman as his spouse. Only when the household is headed by a woman alone is she considered the head. If the family is a split-off family from a sampled family, then a new head is selected.

Age. The age variable recorded in the PSID survey does not necessarily increase by 1 from one year to the next. This may be perfectly correct, since the survey date changes every year. For example, an individual can report being 20 years old in 1990, 20 in 1991, and 22 in 1992. We thus create a consistent age variable by taking the age reported in the first year that the individual appears in the survey and add 1 to this variable in each subsequent year.

Education Level. In the PSID, the education variable is not reported every year and it is sometimes inconsistent. To deal with this problem, we use the highest education level that an individual ever reports as the education variable for each year. Since our sample contains only individuals that are at least 25 years old, this procedure does not affect our education variable in a major way.

Income

Individual Male Wages and Salaries. This is the variable used for individual income in the benchmark case. It is the answer to the question: *How much did (Head) earn altogether from wages or salaries in year t-1, that is, before anything was deducted for taxes or other things?* This is the most consistent earnings variable over time reported in the PSID, as it has not suffered any redefinitions or change in subcomponents.[32]

Individual Male Labor Earnings. `Annual Total Labor Income` includes all income from wages and salaries, commissions, bonuses, overtime and the labor part of self-employment (farm and business income). Self-employment in PSID is split into asset and labor parts using a 50-50 rule in most cases. Because this last component has been inconsistent over time,[33] we subtract the labor part of business and farm income before 1993.

Individual Female Labor Earnings. There is no corresponding `Wages and Salaries` variable for spouses. We use `Wife Total Labor Income` and follow a similar procedure as in the case of heads.

Annual Hours. For heads and wives, annual hours is defined as the sum of annual hours worked on main job, extra jobs, and overtime. It is computed using usual hours of work per week times the number of actual weeks worked in the last year.

Pre-Government Household Labor Earnings. Head and wife labor earnings.

Post-Government Household Labor Earnings. Pre-government household earnings *minus* taxes *plus* public transfers, as defined below.

Taxes. The PSID reports own estimates for total taxes until 1991. For the remaining years, we estimate taxes using TAXSIM.

Public Transfers. Transfers are considered at the family unit level whenever possible. We group social and welfare programs into three broad categories. Due to changes in the PSID design, the specific definition of each program is different every year. We give an overview below and leave the specific replication details for the online Data Appendix.

Transfers

We refer to Table **??** in the main text for a description of the three groups of programs considered, as well as their subcomponents. In the PSID, obtaining an annual amount

[32]See Shin and Solon (2011) for a comparison of PSID male earnings variables in inequality analyses.
[33]In particular, total labor earnings included the labor parts of farm and business income up to the 1993 survey but not in subsequent waves.

of each type of benefits is almost wave-specific. Every few survey years, the level of aggregation within the family unit and across welfare programs is different for at least one of our groups. To impose some common structure, we establish the following rules.

For survey years 1970-1993[34] and 2005-2011, the total annual amount of each program is reported for the head, spouse, and others in the family unit. Occasionally, the amount appears combined for several or all members.[35] Because in those cases it is impossible to identify separate recipiency of each member, we consider the benefit amount of the whole family. That is, we add up all available information for all family members, whether combined or separately reported.

In survey years 1994-2003, most benefits (except Food Stamps and OASDI) are reported separately for the head and the spouse only. The way amounts are reported changes as well. First, the reported amount ($\$X$) received is asked. Second, the frequency of that amount ($\$X$ per year, per month, per week, etc.) is specified. We convert all amounts to a common frequency by constructing a monthly amount $\$x$ using these time values. Finally, the head and spouse are asked during which months the benefit was received. The final annual recipiency of transfers is then obtained by multiplying $\$x$ by the number of months this benefit was received. For Food Stamps and OASDI, we follow the rules described for the other waves.

Detailed Sample Selection

We start with an initial sample of 584,392 SRC individuals interviewed between 1976 and 2011. We then impose the next criteria every year. The number of individuals kept at each stage in the sample selection is listed in Table A.1. Previous to this selection process, we have cleaned the raw data and corrected duplicates and inconsistencies (for example, zero working hours with positive labor income). We also require that the individuals have non top-coded observations in income.

1. The individual must be from the original main PSID sample (not from the Survey of Economic Opportunities or Latino subsamples).

2. In the benchmark individual sample, we select male heads of family. In the reference household sample, we require at least two adult members in the unit and

[34]Our main sample refers to survey years 1977-2011, but complementary results are provided for the annual subsample of the PSID, that is, for 1970-1997. We drop the first two waves in all cases, since benefits such as OASDI, UI, and WC are only reported for the family head and benefits such as SSI are not reported at all.

[35]This is always the case for Food Stamps.

that individuals had no significant changes in family composition. More specifically, we require that they responded either "no change" or "change in family members other than the head or wife" to the question about family composition changes.

3. The household must not have missing variables for the head or wife labor income, or for education of the head. The individuals must not have missing income or education themselves.
4. The individual must not have income observations that are outliers. An outlier is defined as being in the top 1% of the corresponding year.
5. We require the income variable of analysis to be positive.
6. Household heads must be between 25 and 65 years old.

Table A.1: Number of Observations Kept in Each Step: PSID

	Male Heads	Households	All Females
SRC	586,187	586,187	586,187
Family Composition	90,106	75,202	110,711
Non-Missing y or College	83,039	69,443	97,990
Positive Income	63,875	58,551	54,214
Outliers	63,065	57,262	53,257
Age Selection	54,593	50,102	45,330
Final #Obs for transitory changes	42,623	38,171	33,687
Final #Obs for persistent changes	34,985	30,985	27,269

Note: Table lists number of person-year, or household-year, observations in the three panels for the sample from PSID.

A.2 LINDA

Variables

<u>Demographic and Socioeconomic</u>

Head and Relationship to Head. LINDA is compiled from the Income Register based on filed tax reports and other registers. Statistics Sweden samples individuals and then adds information for all family members, where family is defined for tax purposes. This implies that there is no information about "head of households." We therefore define the head of a household as the sampled male.

Age. As defined by Statistics Sweden.

Education Level. LINDA contains information about education from 1991 and onward. An individual is assigned "college" education if he/she has at least three years of university education.

Private/Public employment An individual is defined as as working in the public sector, if he/she works in public administration, health care or education. LINDA contains consistent comparable information for the years 1991 and onward. For the years 1991-1992, we use SNI90 codes 72000-72003, 90000-93999, and >=96000 to define public sector employment. For 1993-2006, we use SNI92 codes 64110-64202, 73000-74110, 75000-92000, 92500-92530, and >=96000. For 2007-2010, we use SNI2007 codes 64110-64202, 73000-74110, 75000-92000, 92500-92530, and >=96000.

Income

For the years 1985-2010, we use the measures suggested by Statistics Sweden to be comparable between years in LINDA. We construct comparable measures for the years 1979-1984.

Individual Labor Earnings. Labor earnings consist of wages and salaries, the part of business income reported as labor income, and taxable compensation for sick leave and parental leave.

Pre-Government Household Labor Earnings. Defined as the sum of individual labor income within the family.

Post-Government Household Labor Earnings. Post-government earnings is calculated as pre-government earnings *minus* taxes *plus* public transfers.

Taxes. LINDA provides observations of total taxes paid by the individual. Since taxes paid on capital income constitute a small part of total tax payments, and since we cannot separate taxes on capital income from those on labor income, we assume that all taxes are labor income taxes.

Public Transfers. LINDA provides observations of total public transfers at the individual level (Statistics Sweden has individualized transfers given to families) and at the household level. We also consider three subcategories of transfer as listed below.

Transfers

Transfers in subcategories 1 and 3 are individual-level transfers. Transfers in subcategory 2 are family level transfers but have been individualized by Statistics Sweden.

For each subcategory, we take all transfers received by all members of the households.

- *HH-level transfers subcategory 1 (labor market transfers)*: sum of unemployment benefits received by all members of household.
- *HH-level transfers subcategory 2 (family aid)*: sum of transfers to support families received by all members of household.
- *HH-level transfers subcategory 3 (pensions)*: sum of old-age pensions received by all members of household.

Detailed Sample Selection

To be included in the individual sample, the individual has to be sampled and between 25 and 60 years old. A family is included in the household sample if the sampled individual is a man between 25 and 60 years old and there are at least two members ages 25-60 in the family.

A.3 SIAB

We use the scientific use file SIAB-R7510 provided by the Institute for Employment Research (IAB). The SIAB data from which the scientific use file is constructed are a 2% random sample of all individuals covered by a data set called IEB. This data set is from four different sources, which can be identified in the data. For construction of our sample, we use earnings data stemming from BeH (employee history) and transfer data from LeH (benefit recipient history). Records in BeH are based on mandatory social security notifications from employers and hence cover individuals working in employment subject to social security, which excludes civil servants, students, and self-employed individuals. A new spell starts whenever there is a new notification, which happens when a new employment relationship changes, an ongoing contract is changed, or a new calendar year starts. BeH covers all workers subject to social security contributions, which excludes civil servants, self-employed individuals and students. For details on the data set, see vom Berge *et al.* (2013).

Variables

Demographic and Socioeconomic

Head and Relationship to Head. SIAB does not contain information on households. We use only individual-level data.

Age. Birth year is reported consistently in SIAB data.

Education Level. Each individual spell in SIAB contains information on the highest degree of formal education as reported by the employer. In order to construct a consistent measure of education we apply imputation rules proposed by Fitzenberger *et al.* (2006).

Private/Public Employment. An individual is defined as working in the public sector, if he/she works in public administration, health care or education. SIAB contains consistent comparable information for all years of the sample. We use the classification WZ93 as provided in the data, which aggregates 3-digit codes of the original WZ93 classification into 14 categories. The industry of an employer is registered once a year and assigned to the worker spells of that year. This implies that for some individual spells, there is no information on the industry. For each year, a worker is assigned the industry from the longest spell in that year. We classify as public employment those in sectors 13 (3-digit WZ93 801-804, 851-853: Education, social, and health-care facilities) and 14 (751-753, 990: public administration, social security).

Income

Individual Labor Earnings. We calculate annual earnings as the sum of total earnings from all valid spells for each individual. As marginal employment spells were not reported before 1999, we drop marginal employment in the years where they are reported in order to obtain a time consistent measure. For the same reason, we drop spells with a reported average daily wage rate below the highest marginal employment threshold in the sample period, which is 14.15 euros (in 2003 euros). The available data have two drawbacks: the structural break of the wage measure in 1984 and top-coding.

Structural Break in Wage Measure. Since 1984 the reported average daily wage rate from an employment spell includes one-time payments. We correct for this structural break following a procedure based on Dustmann *et al.* (2009): we rank individuals from 1976 to 1983 into 50 quintiles of the annual full-time wage distributions. Then we fit locally weighted regressions of the wage growth rate from 1982 to 1983 on the quintiles in 1983 and the same for 1983 to 1984. We then define as the correction factor the difference between the quintile-specific smoothed value of wage growth between 1984 and 1983. The underlying assumption is that wage growth should be higher from 1983 to 1984 because the wage measure includes one-time payments. In order to control for overall wage growth differences, we subtract the average of the correction factor of the second to 20th quintiles. The resulting percentile-specific correction factor is then applied to wages in 1976-1983.

Imputation of Top-Coded Wages. Before aggregating earnings from all spells, we correct full-time wage spells for the top-coding. We therefore follow Daly *et al.* (2014) and fit a Pareto tail to the cross-sectional wage distribution. The Pareto distribution is estimated separately for each year by age group and sex. We define seven age groups: 25-29, 30-34,...,55-60. As a starting point for the Pareto distribution, we choose the 60th percentile of the subgroup-specific distribution. As in Daly *et al.* (2014), we draw one random number by individual, which we then apply to the annual specific distributions when assigning a wage to the top-coded workers. We apply the imputation method to the annual distribution of average full-time wages, and hence an individual can be below the cutoff limit if, for example, from two full-time spells in a year only one is top-coded. We therefore define as the top-coding limit the annual specific limit minus 3 DM (1995 DM) as in Dustmann *et al.* (2009).

Transfers

In SIAB we observe consistently over time unemployment benefits at the individual level.

Detailed Sample Selection

To be included in the sample, the individual has to be between 25 and 60 years old and earn a gross income above 520*0.5*minimum wage. We drop all workers that have at least one spell reported in East Germany.

A.4 SOEP

Variables

Demographic and Socioeconomic

Head and Relationship to Head. For each individual in the sample, SOEP reports the relationship to the head of household in any given wave. Whenever there is a non-couple household, (that is no spouse is reported), the reported head is classified as head. Whenever we observe a couple household and the reported head is a male, we keep this; when the reported head is a female and the reported spouse is a male, we reclassify the male to be head and the female to be spouse.

Age. The age is measured by subtracting the year of birth from the current year.

Education Level. The education variable used categorizes the obtained maximum education level by ISCED 1997. An individual with category 6 is assigned "college" education; an individual with categories 1-5 is assigned "non-college." Category 6 includes a degree obtained from a university, from technical college, from a university abroad, and a PhD. An individual still in school (category 0) is assigned a missing value. For a small number of individuals, the described procedure yields inconsistencies in the sense that for some year t, the assignment is "college" and some later year $t+s$ the assignment is "non-college"; in these cases, we assign "college" to the later year.

Income and Hours

Individual Labor Income. Labor earnings are calculated from individual labor income components and include income from first job, secondary job, 13th and 14th salary, Christmas bonus, holiday bonus, and profit sharing. For consistency with the PSID measure, we assign 50% of income from self-employment to labor income.

Household-Level Labor Income. Defined as the sum of individual labor income of head and spouse.

Annual Hours. SOEP measures the average actual weekly hours worked and the numbers of months an individual worked. From these measures SOEP, provides a constructed measure of annual hours worked of an individual.

Pre-Government Household Labor Earnings. Head and spouse labor earnings.

Post-Government Household Labor Earnings. Pre-government household earnings *minus* taxes *plus* public transfers, as defined below.

Taxes. SOEP provides estimates of total taxes at the household level.

Public Transfers. Transfers are considered at the family unit level and at the individual level. We group social and welfare programs into three broad categories as listed below.

Transfers

Transfers are partly observed at the individual level and partly at the household level. For each subcategory, we take all transfers received by all members of the households.

- *HH-level transfers*: we use transfers received by all individual household members in order to calculate measures that are consistent over time. For each individual, total transfers are the sum of the following components: old-age pensions, widow's pensions, maternity benefit, student grants, unemployment benefits, subsistence allowance, unemployment assistance (up to 2004); at the hh-level we measure received child allowances and the total unemployment benefits II received by all household members (since 2005 replacing unemployment assistance).
- *HH-level transfers subcategory 1 (labor market transfers)*: sum of unemployment benefits received by all members of household.
- *HH-level transfers subcategory 2 (family aid)*: sum of subsistence allowance of all members, + sum of unemployment assistance received by all members (up to 2004), + hh-level measure of unemployment benefits II (since 2005).
- *HH-level transfers subcategory 3 (pensions)*: sum of old-age pensions received by all members of household.

Sample Selection

In order to be in the initial sample for a year, the individual or household head must be between ages 25 and 60 and live in West Germany. In order to have a consistent sample, we drop the immigrant subsample and the high-income subsample. This gives initial sample sizes of 87,582 individual-year observations for the male sample, 76,249 individual-year observations for the female sample, and 76,051 household-year observations for the household sample (see Table A.2). The sample selection then follows the steps listed below for each sample. All cross-sectional statistics are calculated using appropriate cross-sectional individual or household weights, respectively.

1. drop if no info on education or if no degree obtained yet
2. drop if currently working in military
3. drop if no info on income
4. drop if no info on hours worked
5. keep if income >0 and hours >520

Table A.2: Number of Observations Kept in Each Step: SOEP

selection step	Male Heads	Households	All Females
initial	87,582	76,051	76,249
drop if no coll. info	86,737	75,310	75,270
drop if in military	86,712	75,293	75,268
drop if no obs on ymin	79,547	75,070	50,374
drop if no obs on hours	79,547	75,070	50,374
keep if >=520 hrs and ymin>0	77,265	71,389	42,245
drop top 1% of ymin per year	76,404	70,627	41,830
drop if ymin<.5*520*min wage	76,268	70,097	41,434
Final #Obs for transitory changes	64,572	59,209	31,612
Final #Obs for persistent changes	38,399	34,792	16,792

Note: Table lists number of person-year, or household-year, observations in the three panels for the sample from SOEP.

6. drop if in highest percentile (sample outliers)
7. drop if below 520*0.5**minimum wage*, where *minimum wage* is set to be 6€ in year 2000 euros
8. for transitory change measure: keep if in sample in t and t-1
9. for permanent change measure: keep if in sample in t and t-5

B Cyclicality of Individual Earnings by Groups

Tables B.1 to B.7 show results of the individual level earnings regressions discussed in Section 4 by subgroups.

Table B.1: Cyclicality of Male Earnings, by Education Groups

	L9010	Kelley	L9050	L5010
	United States			
College Graduates	–0.12	0.97	0.36	–0.48
	(–0.31)	(1.42)	(1.39)	(–1.15)
Non-College	–0.40	2.37***	0.83*	–1.23***
	(–0.69)	(4.29)	(2.04)	(–3.88)
	Sweden			
College Graduates	–0.00	1.80***	0.42	–0.42***
	(–0.01)	(4.93)	(1.58)	(–5.72)
Non-College	–0.17	4.03***	0.99***	–1.26***
	(–1.52)	(3.86)	(3.39)	(–3.53)
	Germany (SIAB)			
College Graduates	0.62	4.70***	1.24**	–0.61**
	(1.01)	(3.10)	(2.17)	(–2.29)
Non-College	0.10	5.26***	0.89***	–0.79***
	(0.25)	(5.41)	(3.07)	(–3.78)

Note: See Table I for explanations.

Table B.2: Cyclicality of Female Earnings, by Education Groups

	L9010	Kelley	L9050	L5010
	United States			
College graduates	–1.11	1.70**	0.49	–1.60***
	(–1.44)	(2.61)	(0.94)	(–2.84)
Non-college	0.91**	0.78*	0.91***	–0.00
	(2.77)	(1.75)	(2.91)	(–0.01)
	Sweden			
College graduates	0.13	1.15***	0.64	–0.25
	(0.31)	(4.03)	(1.22)	(–1.74)
Non-college	0.50*	1.81***	0.75***	–0.25**
	(1.96)	(3.40)	(2.78)	(–2.71)
	Germany (SIAB)			
College graduates	0.01	2.03	1.01	–1.00
	(0.01)	(1.65)	(1.12)	(–1.39)
Non-college	0.32	2.58**	0.77	–0.45*
	(0.47)	(2.08)	(1.27)	(–1.88)

Note: See Table I for explanations.

Table B.3: Cyclicality of Individual Earnings, Public vs. Private Sector Employment, Males

	L9010	Kelley	L9050	L5010
	United States			
Private	–0.39	2.26***	0.82***	–1.21***
	(–1.08)	(4.43)	(2.88)	(–4.03)
Public	0.05	0.20	0.07	–0.01
	(0.23)	(0.29)	(0.36)	(–0.07)
	Sweden			
Private	0.10	3.83***	0.93***	–0.83***
	(0.93)	(4.02)	(3.81)	(–4.08)
Public	–0.45***	2.10***	0.17	–0.62***
	(–3.93)	(6.55)	(1.64)	(–9.11)
	Germany			
Private	0.03	5.55***	0.88***	–0.85***
	(0.08)	(6.44)	(3.55)	(–5.64)
Public	2.50	0.30	1.45	1.06
	(1.16)	(0.17)	(1.08)	(1.01)

Note: See Table I for explanations.

C Robustness of the Empirical Results

We perform a number of robustness checks for the analyses based on SIAB data, which deal with (i) top-coding of incomes and (ii) a structural break in the income measure in 1984. In addition to Kelley skewness, we consider two alternatives: two versions of Hinkley's measure of skewness. Instead of L9050 and L5010, these measures relate L8550 and L5015 or L8050 and L5020, respectively.

The first four rows of table C.1 show the results of the regressions for male and female earnings wages, respectively. The results are the ones from the main text and serve as a comparison to the robustness analyses. Columns 7-12 show the results for the two versions of Hinkley's skewness measures and the corresponding tails. Compared to Kelley skewness and L9050 and L5010, the estimates show that the substantive conclusion is also robust for these smaller log percentile differentials. Rows 5 and 6 show the results for the wage regressions when applying a less strict criterion of working full-time for only 45 weeks in two consecutive years. Again, the results are very similar to those reported for 50 weeks.

Table B.4: Cyclicality of Individual Earnings, by Sector of Employment, Females

	L9010	Kelley	L9050	L5010
	United States			
Private	0.85**	1.47***	1.32***	–0.47
	(2.64)	(3.38)	(3.82)	(–1.67)
Public	–0.43	–0.87	–0.44	0.01
	(–0.69)	(–0.94)	(–0.81)	(0.04)
	Sweden			
Private	0.50*	1.99***	0.78**	–0.29**
	(1.87)	(3.02)	(2.81)	(–2.43)
Public	0.18	1.10***	0.34**	–0.16**
	(1.19)	(3.29)	(2.43)	(–2.61)
	Germany			
Private	0.01	3.13**	0.73	–0.72***
	(0.01)	(2.44)	(1.50)	(–3.15)
Public	1.17	0.95	0.85	0.32
	(0.84)	(0.68)	(0.85)	(0.59)

Note: See Table I for explanations.

In order to ensure that top-coding does not drive our results, we redo the analysis using reduced samples in which an individual is considered in the distribution of income changes from t to $t+1$ only if income is below the top-coding thresholds in both t and $t+1$. About 11% and 2% of all observations are top-coded in the male and female base samples, respectively. Table C.2 shows the results of the respective regressions for earnings, wages, and wages of firm stayers for both males and females. Second, we rerun the regressions completely ignoring top-coding, that is, all individuals from the base sample are in the sample, but with their reported incomes again for earnings, wages, and wages of stayers. Results are in Table C.3.

A rerun of the regression analysis using only observations after 1983, thereby dropping all years for which the reported income measure does not include one-time payments such as bonuses, does not change the results (see the lower panel of Table C.3).

Table B.5: Cyclicality of Earnings by Occupational Area: Germany (SIAB)

	L9010	Kelley	L9050	L5010
		Males		
Farming and related	4.56	5.64	3.80	0.76
	(1.23)	(1.51)	(1.52)	(0.45)
Mining, Mineral Extraction	2.62	3.23	1.32**	1.30
	(1.25)	(1.39)	(2.43)	(0.72)
Manufacturing, Fabrication	0.17	11.39***	2.00***	-1.83***
	(0.20)	(5.53)	(3.21)	(-3.99)
Technical Occupations	0.13	12.36***	1.51**	-1.38***
	(0.19)	(4.04)	(2.72)	(-3.64)
Service Occupations	0.59	8.89***	1.76**	-1.17***
	(0.68)	(3.92)	(2.41)	(-3.09)
		Females		
Farming and related	2.90	0.96	2.06	0.84
	(0.73)	(0.31)	(0.71)	(0.61)
Mining, Mineral Extraction	-5.59	12.26	1.61	-7.20**
	(-1.02)	(1.54)	(0.34)	(-2.59)
Manufacturing, Fabrication	-0.72	10.59***	2.48*	-3.20***
	(-0.48)	(4.95)	(2.00)	(-6.01)
Technical Occupations	-0.75	8.44**	1.41	-2.16***
	(-0.83)	(2.70)	(1.56)	(-2.82)
Service Occupations	0.85	4.09	1.45	-0.60
	(0.59)	(1.63)	(1.13)	(-1.15)

Note: See notes for Table I.

Table B.6: Cyclicality of Earnings by Occupational Segments: Males, Germany (SIAB)

	L9010	Kelley	L9050	L5010
	Distribution of Beta Coefficients			
Mean	0.71	8.78	2.06	–1.35
P10	–0.83	3.68	0.90	–2.91
Median	0.46	8.35	1.75	–1.29
P90	2.62	13.09	3.29	–0.18
Min	–1.55	2.36	0.66	–7.09
Max	4.56	17.87	7.89	1.30
	Distribution of t-Statistics			
Mean	0.34	4.11	2.59	–2.21
P10	–0.90	1.39	1.43	–4.93
Median	0.46	3.23	2.46	–2.22
P90	1.19	7.73	3.73	–0.18
Min	–2.30	1.01	0.99	–6.62
Max	2.01	11.46	6.36	0.72

Note: The table displays moments of the distribution of beta coefficients (upper panel) and t-statistics (lower panel) from separate regressions for each of the 30 occupational segments. See notes for Table I.

Table B.7: Cyclicality of Earnings by Occupational Segments: Females, Germany (SIAB)

	L9010	Kelley	L9050	L5010
	Distribution of β^m across Occupations			
Mean	–0.10	6.87	1.98	–2.08
P10	–3.33	0.96	0.89	–4.45
Median	0.12	6.30	2.06	–2.00
P90	1.97	12.26	3.10	0.69
Min	–6.40	–0.56	–4.13	–7.20
Max	3.61	13.50	3.80	0.84
	Distribution of t-Statistics			
Mean	–0.11	2.35	1.43	–1.84
P10	–1.53	0.31	0.53	–4.28
Median	0.05	1.81	1.38	–1.35
P90	0.71	5.41	2.27	0.52
Min	–2.06	–0.07	–0.72	–6.30
Max	0.95	6.15	4.04	0.61

Note: The table displays moments of the distribution of beta coefficients (upper panel) and t-statistics (lower panel) from separate regressions for each of the 30 occupational segments. See notes for Table I.

Table C.1: Sensitivity of Regression Results - SIAB I

	Std Dev	L9010	Skew	Kelley	L9050	L5010	Hinkley 1	Hinkley 2	L8550	L8050	L5015	L5020
Male Earnings	0.07	0.15	14.42***	5.48***	0.95***	-0.80***	5.84***	5.85***	0.51***	0.32***	-0.54***	-0.36***
	(0.42)	(0.36)	(4.28)	(5.80)	(3.14)	(-4.11)	(9.85)	(7.51)	(4.10)	(3.57)	(-4.77)	(-3.43)
Female Earnings	0.10	0.34	4.34*	2.55**	0.80	-0.46*	2.75**	2.71***	0.43	0.25	-0.24**	-0.14*
	(0.47)	(0.48)	(1.77)	(2.05)	(1.25)	(-1.80)	(2.62)	(3.85)	(1.40)	(1.65)	(-2.56)	(-1.87)
Male Wages	0.01	-0.09	14.55***	4.73***	0.30***	-0.39***	4.94***	4.88***	0.22**	0.18**	-0.28**	-0.20**
	(0.23)	(-0.54)	(4.58)	(6.31)	(3.77)	(-3.20)	(4.35)	(3.37)	(2.59)	(2.66)	(-2.55)	(-2.07)
Female Wages	0.04	0.03	8.98*	2.12***	0.17**	-0.14	2.20***	2.09***	0.14**	0.11**	-0.09	-0.04
	(0.66)	(0.18)	(2.02)	(5.11)	(2.61)	(-1.58)	(4.79)	(4.67)	(2.68)	(2.65)	(-1.24)	(-0.83)
Male Wages	0.01	-0.08	13.20***	4.65***	0.31***	-0.39***	4.88***	4.85***	0.23**	0.18***	-0.29**	-0.20**
(45 weeks)	(0.27)	(-0.54)	(4.55)	(6.60)	(3.90)	(-3.30)	(4.50)	(3.48)	(2.70)	(2.78)	(-2.61)	(-2.09)
Female Wages	0.04	0.04	8.80*	2.07***	0.17**	-0.14	2.20***	2.10***	0.14**	0.12**	-0.09	-0.05
(45 weeks)	(0.72)	(0.25)	(2.02)	(5.21)	(2.72)	(-1.57)	(4.85)	(4.72)	(2.73)	(2.66)	(-1.23)	(-0.84)

Note: See notes for Table I.

Table C.2: Sensitivity of Regression Results - SIAB II

	Std Dev	L9010	Skew	Kelley	L9050	L5010	Hinkley 1	Hinkley 2	L8550	L8050	L5015	L5020
	Not top-coded workers only:											
Male Earnings	0.08	0.26	14.49***	4.98***	0.96**	-0.70***	4.83***	4.65***	0.48***	0.31***	-0.44***	-0.28***
	(0.41)	(0.53)	(4.26)	(4.28)	(2.53)	(-3.07)	(6.66)	(8.86)	(3.13)	(3.40)	(-4.06)	(-3.08)
Male Wages	-0.01	-0.05	8.76***	3.39***	0.23***	-0.28***	3.49***	3.36***	0.19***	0.14***	-0.20**	-0.14*
	(-0.14)	(-0.29)	(6.07)	(10.76)	(3.52)	(-2.91)	(8.43)	(8.09)	(3.74)	(3.49)	(-2.34)	(-2.00)
Male Wages	-0.03	-0.08	11.41***	3.66***	0.22***	-0.30***	3.67***	3.48***	0.17***	0.13***	-0.21**	-0.14**
(stayers)	(-0.75)	(-0.52)	(5.77)	(9.09)	(4.12)	(-2.96)	(7.52)	(7.65)	(4.14)	(3.63)	(-2.42)	(-2.10)
Female Earnings	0.09	0.33	4.67*	2.54*	0.80	-0.46*	2.72**	2.67***	0.43	0.25	-0.23**	-0.13*
	(0.45)	(0.47)	(1.90)	(2.03)	(1.24)	(-1.83)	(2.57)	(3.76)	(1.40)	(1.67)	(-2.46)	(-1.71)
Female Wages	0.04	0.05	2.04	2.05***	0.17**	-0.12	2.11***	2.12***	0.13***	0.11**	-0.08	-0.05
	(0.71)	(0.31)	(0.66)	(4.42)	(2.64)	(-1.34)	(4.10)	(4.56)	(2.77)	(2.74)	(-1.08)	(-0.92)
Female Wages	0.02	0.03	3.87	2.17***	0.16***	-0.12	2.25***	2.18***	0.13***	0.10***	-0.08	-0.05
(stayers)	(0.56)	(0.25)	(0.78)	(4.11)	(3.16)	(-1.38)	(4.04)	(4.46)	(3.24)	(2.99)	(-1.15)	(-0.98)

Note: See notes for Table I.

Table C.3: Sensitivity of Regression Results - SIAB III

	Std Dev	L9010	Skew	Kelley	L9050	L5010	Hinkley 1	Hinkley 2	L8550	L8050	L5015	L5020
						Ignore Top-Coding:						
Male Earnings	0.07	0.15	14.72***	5.68***	0.91***	-0.76***	5.97***	6.17***	0.48***	0.30***	-0.52***	-0.37***
	(0.40)	(0.40)	(4.30)	(5.70)	(3.14)	(-4.68)	(10.66)	(8.02)	(4.27)	(3.46)	(-5.60)	(-4.51)
Male Wages	-0.01	-0.09	13.36***	4.93***	0.29***	-0.38***	5.40***	5.38***	0.21**	0.15*	-0.30***	-0.22***
	(-0.27)	(-0.69)	(4.24)	(7.59)	(4.59)	(-3.86)	(5.11)	(3.76)	(2.57)	(1.95)	(-3.59)	(-3.22)
Male Wages	-0.04	-0.13	16.06***	5.19***	0.27***	-0.39***	5.55***	5.36***	0.19**	0.14*	-0.31***	-0.22***
(stayers)	(-1.15)	(-1.06)	(3.47)	(6.67)	(4.03)	(-3.85)	(4.63)	(3.51)	(2.11)	(1.85)	(-3.58)	(-3.17)
Female Earnings	0.10	0.34	4.36*	2.51*	0.79	-0.45*	2.70**	2.63***	0.41	0.24	-0.23**	-0.13*
	(0.48)	(0.48)	(1.76)	(2.02)	(1.24)	(-1.78)	(2.61)	(3.98)	(1.39)	(1.65)	(-2.55)	(-1.84)
Female Wages	0.03	0.02	1.91	2.03***	0.15**	-0.13	2.16***	2.17***	0.13**	0.11**	-0.09	-0.05
	(0.65)	(0.16)	(0.60)	(4.79)	(2.64)	(-1.51)	(4.68)	(5.10)	(2.74)	(2.68)	(-1.29)	(-1.11)
Female Wages	0.02	0.01	3.52	2.20***	0.14***	-0.13	2.35***	2.27***	0.12***	0.10***	-0.09	-0.05
(stayers)	(0.48)	(0.09)	(0.75)	(4.45)	(3.31)	(-1.58)	(4.70)	(4.90)	(3.30)	(3.01)	(-1.40)	(-1.18)
						1984-2010:						
Male Earnings	-0.04	-0.07	13.25***	5.10***	0.81***	-0.88***	5.82***	6.22***	0.46***	0.31***	-0.61***	-0.42***
	(-0.26)	(-0.18)	(3.84)	(5.85)	(2.96)	(-3.99)	(11.04)	(9.21)	(3.83)	(3.12)	(-5.48)	(-4.33)
Female Earnings	0.04	0.30	3.86	2.46*	0.75	-0.46	2.79**	2.88***	0.42	0.26	-0.26**	-0.16*
	(0.21)	(0.39)	(1.51)	(1.84)	(1.10)	(-1.65)	(2.46)	(3.65)	(1.24)	(1.49)	(-2.29)	(-1.87)

Note: See notes for Table I.

D Long-Run Earnings Growth

Figures D.1 to D.3 characterize the distribution of long-run earnings growth, that is, five-year changes for Germany and Sweden, and four-year changes for the United States.

Figure D.1: Standard Deviation of Long-Run Earnings Growth: United States, Sweden, and Germany

(a) United States

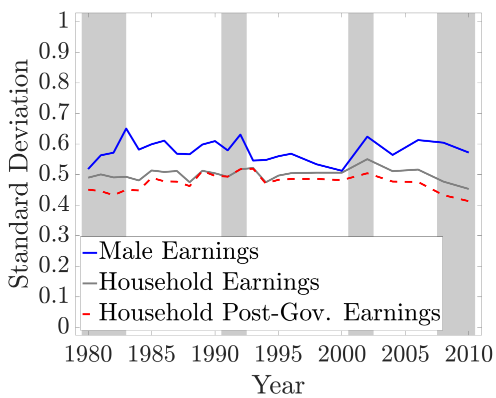

(b) Sweden

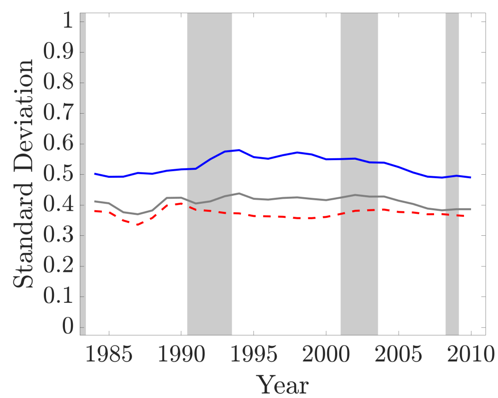

(c) Germany

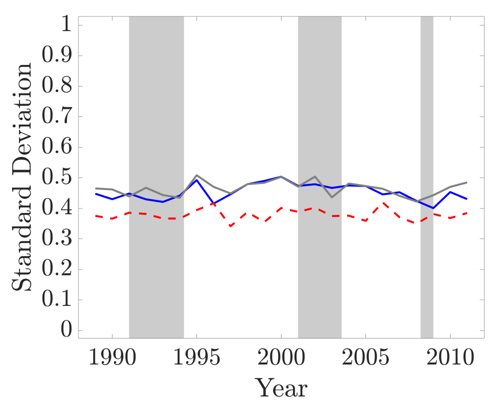

Note: See notes to Figure 2

Figure D.2: Kelley skewness of Long-Run Earnings Growth: United States, Sweden, and Germany (SOEP)

(a) United States

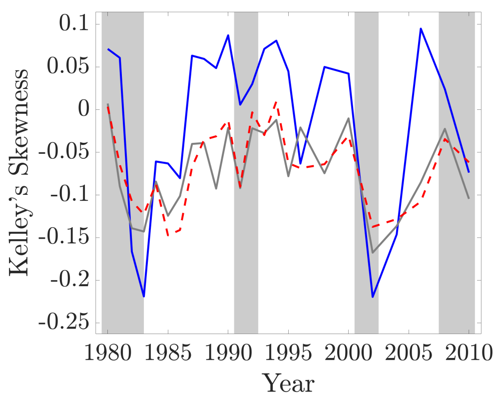

(b) Sweden

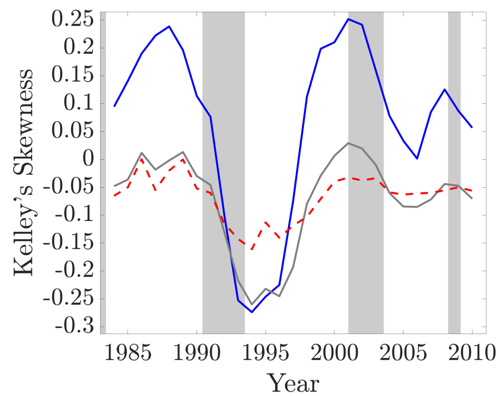

(c) Germany

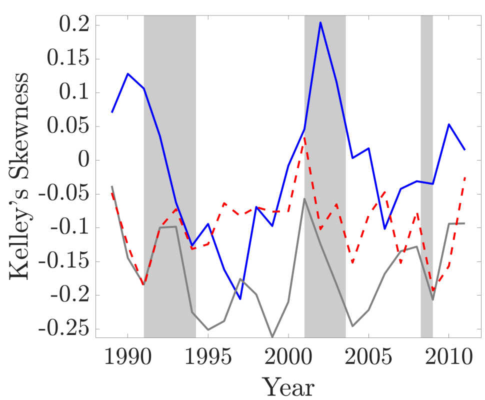

Note: See notes to Figure 2

Figure D.3: Standard Deviation, Kelley skewness, and Tails of Long-Run Earnings Growth: Germany, IAB Sample

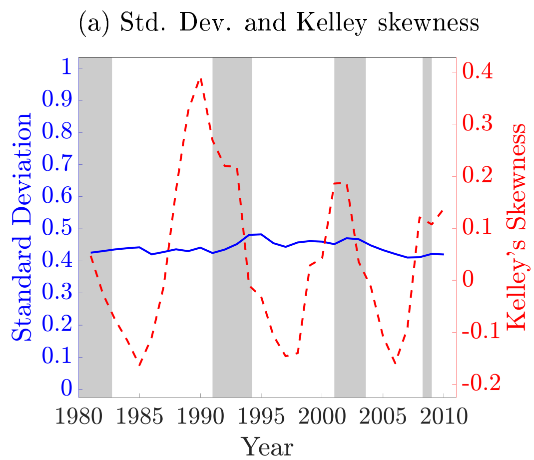

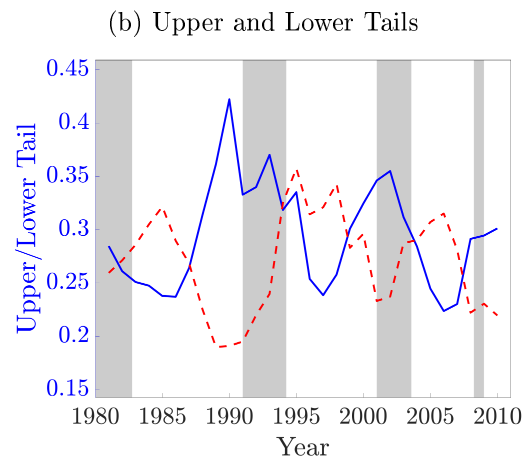

Note: Linear trend removed, centered at sample average.

E Details on the Estimation and Simulation

The vector of parameters to be estimated is

$$\chi = (\sigma_{\xi,1}, \sigma_{\xi,2}, p_{\xi,1}, \mu_{\zeta,2}, \mu_{\zeta,3}, \sigma_{\zeta,1}, \sigma_{\zeta,2}, p_{\zeta,1}, p_{\zeta,2}, \phi_2, \phi_3),$$

which we estimate by simulated method of moments (SMM). We target the time series of L9050 and L5010 of the 1, 3, and 5-year earnings changes distribution, the average of the Crow-Siddiqui measure of kurtosis of 1-, 3-, and 5-year changes, as well as the age profile of the cross-sectional variance from ages 25 to 60. The Crow-Siddiqui measure of kurtosis (Crow and Siddiqui, 1967) is defined as

$$\mathcal{CS} = \frac{(P97.5 - P2.5)}{(P75 - P25)}.$$

This gives 213 moments for our estimation of the income process for Sweden.

To construct the simulated income profiles over time, we write earnings growth as a function of the shocks, using equation (3):

$$y_t - y_{t-s} \;\; = \;\; \xi_t - \xi_{t-s} + \sum_{j=0}^{s-1} \zeta_{t-j}, \tag{7}$$

for the different horizons $s = 1, 3, 5$.

The simulated series of the life-cycle variance profile of log earnings is computed as follows. We assume a time-invariant distribution of shocks by imposing $x_t \equiv 0$. Notice that this assumes that the variance accumulates linearly over the life cycle. We then normalize the series so that the variance at age 25 in the simulation is 0. Finally, we rescale the resulting simulated profile to exhibit the same mean as its empirical counterpart.

We simulate these profiles $R = 10$ times for $I = 10,000$ individuals and compute the moments corresponding to the aforementioned targets. To find $\hat{\chi}$, we minimize the average scaled distance between the simulated and empirical moments. A weighting matrix is used to scale the life-cycle profile. In particular, we weight the variance profile with 0.2 and the remaining moments with 0.8. For the optimization part, we use a global version of the Nelder-Mead algorithm with several quasi-random restarts, as described in Guvenen (2011).

Simulation of Income Profiles

Let c_n^m denote the empirical moment n ($n = 1, \cdots, N$) that corresponds to cross-sectional target m
($m \in \{L5010(\Delta^1 y_t), L5010(\Delta^3 y_t), L5010(\Delta^5 y_t), \cdots, var(y_{age=25}), ..., var(y_{age=60})\}$). In each simulation, we draw a matrix of random variables
$X_r = \{\xi_1^i, \xi_2^i, ..., \xi_{T-1}^i, \xi_T^i, \zeta_1^i, \zeta_2^i, ..., \zeta_{T-1}^i, \zeta_T^i\}_{i=1}^I$, where T denotes the last year available in the data. We simulate these profiles $R = 10$ times. For each simulation, we then calculate the respective simulated moments $d_n^m(\chi, X_r)$ given the parameter vector χ.

Optimization

We minimize the scaled deviation $F(\chi)$ between each data and simulated moment

$$min_\chi F(\chi)' W F(\chi)$$

where F is defined as

$$\begin{aligned} F_n(\chi) &= \frac{d_n^m(\chi) - c_n^m}{|c_n^m|} \\ d_n^m(\chi) &= \frac{1}{R}\sum_{r=1}^{R} d_n^m(\chi, X_r) \end{aligned}$$

Because our goal is to capture as closely as possible the business-cycle fluctuations of idiosyncratic income risk, we impose the mean of the medium-run income changes to be as in the data. We adjust the weighting matrix such that the cross-sectional moments get a weight of 80% and the life-cycle moments get a weight of 20%.

Parameter Estimates and Model Fit

Parameter Estimates

As noted at the beginning of this section, we estimate income processes for pre-government household labor income and, separately, for post-government household income. Table E.1 shows the parameter estimates. Figures E.1 and E.2 compare the simulated moments at these parameters and the empirical moments.

Table E.1: Estimated Parameter Values

Parameter	Description	Pre-Gov.	Post-Gov.
$p_{\xi,1}$	Mixture prob. of ξ distribution	0.892	0.877
$\sigma_{\xi,1}$	St. dev. of ξ distribution mix. comp. 1	0.055	0.047
$\sigma_{\xi,2}$	St. dev. of ξ distribution mix. comp. 2	0.628	0.401
$p_{\zeta,1}$	Weight of center of ζ distribution	0.981	0.981
$p_{\zeta,2}$	Weight of right tail of ζ distribution	0.010	0.009
$p_{\zeta,3}$	Weight of left tail of ζ distribution	0.010	0.009
$\sigma_{\zeta,1}$	St. dev. of center of ζ distribution	0.086	0.057
$\sigma_{\zeta,2}$	St. dev. of right tail of ζ distribution	0.020	0.009
$\sigma_{\zeta,3}$	St. dev. of left tail of ζ distribution	0.020	0.009
$\mu_{\zeta,2}$	Mean of right tail of ζ distribution	0.002	0.008
$\mu_{\zeta,3}$	Mean of left tail of ζ distribution	-0.158	-0.065
ϕ_2	Aggregate risk transmission upper tail	1.186	1.240
ϕ_3	Aggregate risk transmission lower tail	0.467	0.229
M	# moments targeted in estimation	213	213

Note: Estimated parameters for gross household labor income (Pre-Gov.) and household income after taxes and transfers (Post-Gov.) in Sweden.

Figure E.1: Pre-Government Income Fit: Sweden

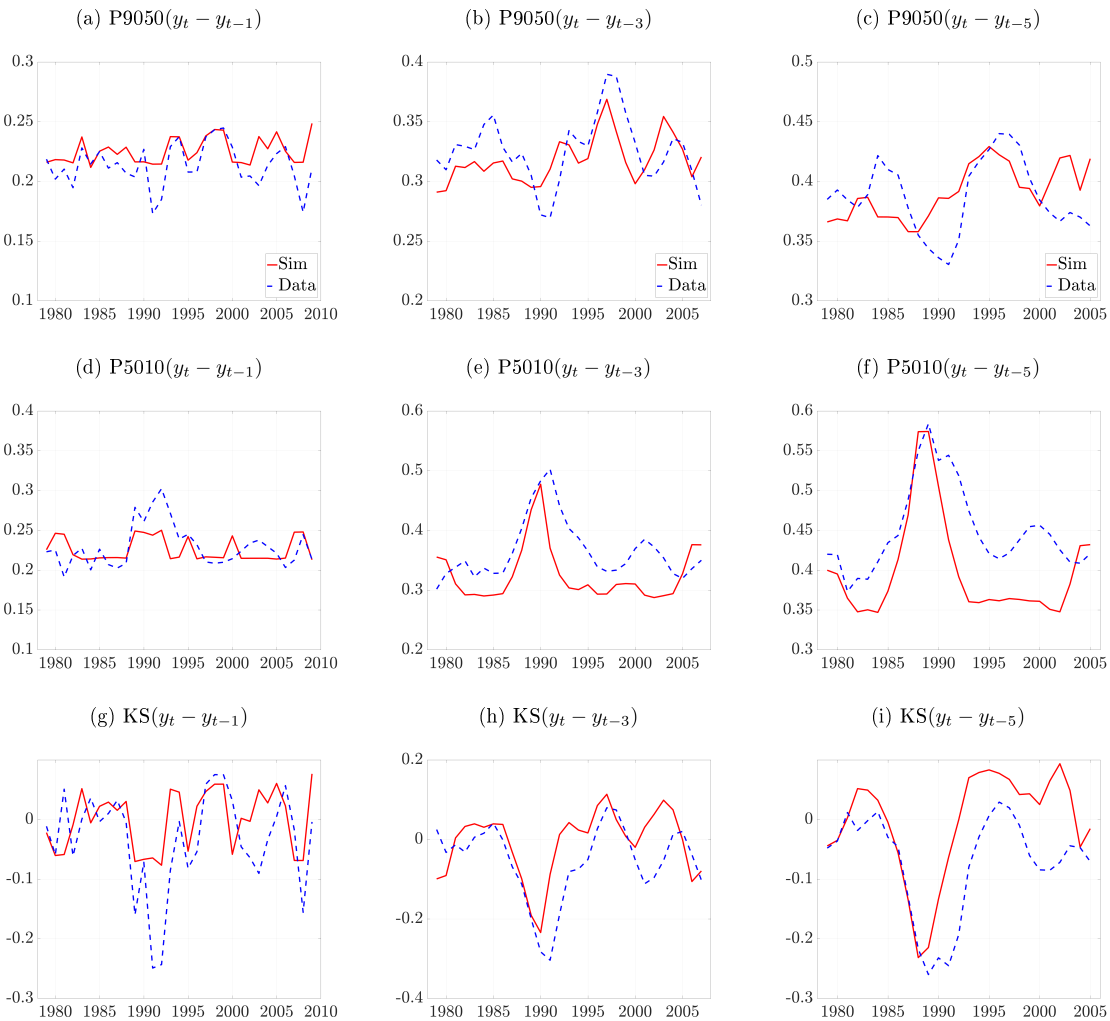

Note: Each panel shows the time series of a moment of short-run, medium-run, or long-run income changes together with the corresponding moment implied by the estimated income process.

Figure E.2: Post-Government Income Fit: Sweden

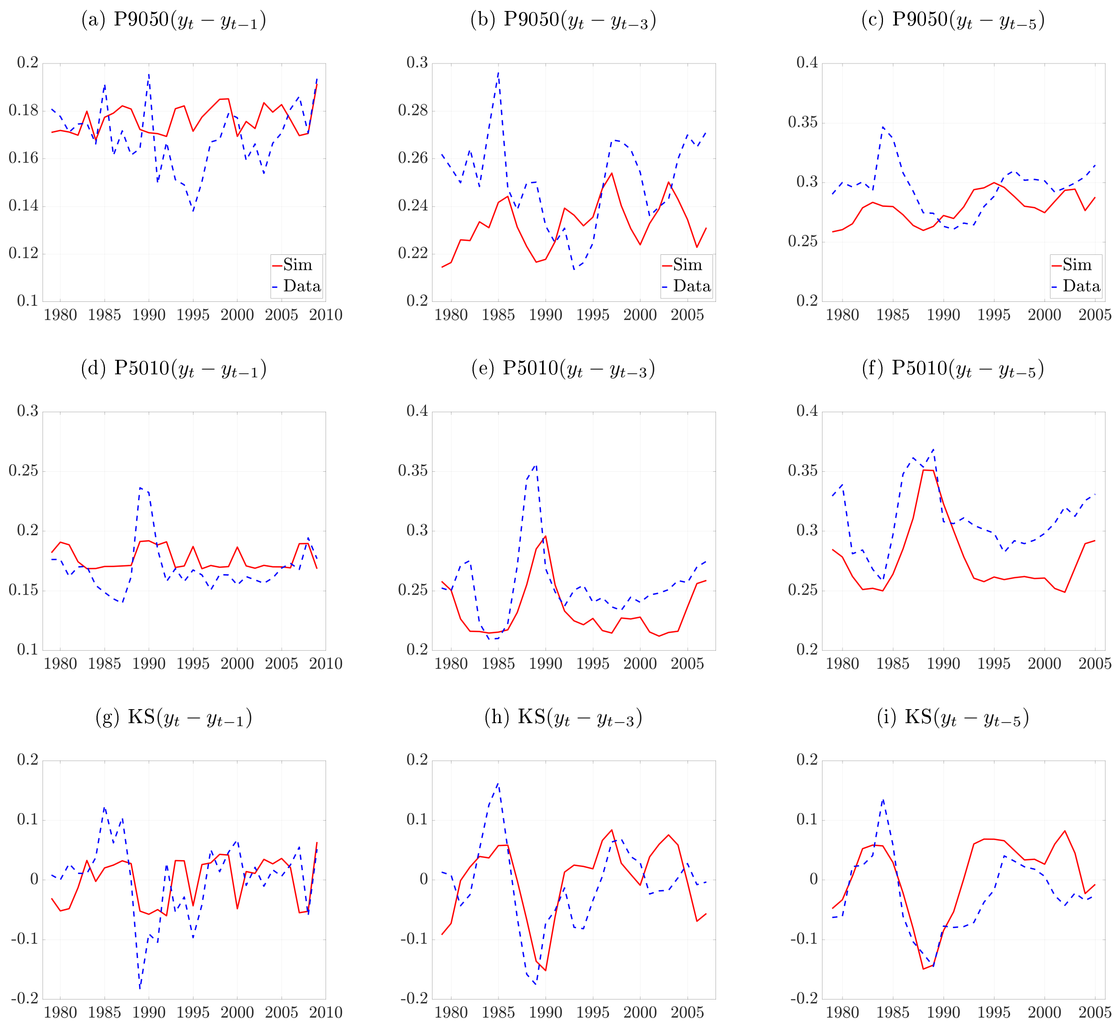

Note: See notes to figure E.1.

F Quantitative Model

Given estimates of the income process, we scale the parameters of the permanent shocks ζ to feed them into the model; fraction λ is insurable and the rest is uninsurable. This scaling implies that the first three standardized moments of the distribution of insurable shocks are given as below: for the first three moments of the uninsurable shocks, simply replace λ with $1-\lambda$.

$$
\begin{aligned}
E\left[\eta_t\right] &= \sum_{i=1}^{3} p_{\zeta,i}\mu_{\eta,i,t} = \sum_{i=1}^{3} p_{\zeta,i}\lambda^{1/2}\mu_{\zeta,i,t} = \lambda^{1/2}\sum_{i=1}^{3} p_{\zeta,i}\mu_{\zeta,i,t} = \lambda^{1/2}E\left[\zeta_t\right] \equiv \lambda^{1/2}\mu_{\zeta,t} \\
var\left[\eta_t\right] &= \sum_{i=1}^{3} p_{\zeta,i}\left(\sigma_{\eta,i}^2 + \mu_{\eta,i,t}^2\right) - \left(E\left[\eta_t\right]\right)^2 = \sum_{i=1}^{3} p_{\zeta,i}\left(\lambda\sigma_{\zeta,i}^2 + \lambda\mu_{\zeta,i,t}^2\right) - \left(\lambda^{1/2}E\left[\zeta_t\right]\right)^2 \\
&= \lambda\left(\sum_{i=1}^{3} p_{\zeta,i}\left(\sigma_{\zeta,i}^2 + \mu_{\zeta,i,t}^2\right) - E\left[\zeta_t\right]^2\right) = \lambda var\left[\zeta_t\right] \\
skew\left[\eta_t\right] &= \frac{1}{var\left[\eta_t\right]^{3/2}}\sum_{i=1}^{3} p_{\zeta,i}\left(\mu_{\eta,i,t} - E\left[\eta_t\right]\right)\left[3\sigma_{\eta,i}^2 + \left(\mu_{\eta,i,t} - E\left[\eta_t\right]\right)^2\right] \\
&= \frac{1}{\lambda^{3/2}var\left[\zeta_t\right]^{3/2}}\sum_{i=1}^{3} p_{\zeta,i}\left(\lambda^{1/2}\mu_{\zeta,i,t} - \lambda^{1/2}E\left[\zeta_t\right]\right)\left[3\lambda\sigma_{\zeta,i}^2 + \left(\lambda^{1/2}\mu_{\zeta,i,t} - \lambda^{1/2}E\left[\zeta_t\right]\right)^2\right] \\
&= \frac{1}{\lambda^{3/2}var\left[\zeta_t\right]^{3/2}}\sum_{i=1}^{3} p_{\zeta,i}\lambda^{1/2}\left(\mu_{\zeta,i,t} - E\left[\zeta_t\right]\right)\left[\lambda\left(3\sigma_{\zeta,i}^2 + \left(\mu_{\zeta,i,t} - E\left[\zeta_t\right]\right)^2\right)\right] \\
&= \frac{1}{var\left[\zeta_t\right]^{3/2}}\sum_{i=1}^{3} p_{\zeta,i}\left(\mu_{\zeta,i,t} - E\left[\zeta_t\right]\right)\left[\left(3\sigma_{\zeta,i}^2 + \left(\mu_{\zeta,i,t} - E\left[\zeta_t\right]\right)^2\right)\right] \\
&= skew\left[\zeta_t\right]
\end{aligned}
$$