Asymmetric Business Cycle Risk and Government Insurance

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

March 25, 2016

Work in Progress. Comments Welcome.

Abstract

This paper studies the business-cycle variation in higher-order income risk—i.e., risks that are captured by moments higher than the variance. A key focus of our analysis is the extent to which such risks can be smoothed within households or with government social insurance policies. To provide a broad perspective on these questions, we study panel data on individuals and households from the United States, Germany, and Sweden, covering more than three decades of data for each country. We find that the underlying variation in higher-order risk is remarkably 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 of these shocks becomes much more negative in recessions. Government provided insurance, in the form of unemployment insurance, welfare benefits, aid to low income households, and the like, plays a more important role reducing downside risk in all three countries; the effectiveness is weakest in the United States, and most pronounced in Germany. Using a model with partial insurance against income shocks, we calculate that the welfare benefits of stabilizing higher-order income risk over the business cycle through the existing tax and transfer system range from 4% of annual consumption for Germany to 7% for Sweden, when assuming log utility and a modest degree of partial insurance beyond taxes and transfers.

JEL Codes: D31, E24, E32, H31

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

[*]CMR, University of Cologne; busch@wiso.uni-koeln.de

[†]Stockholm School of Economics; david.domeij@hhs.se

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

[§]University of Minnesota; made0088@umn.edu

1 Introduction

This paper studies how higher-order income risk varies over the business cycle as well as the extent to which such risks can be smoothed within households or with government social insurance policies. By higher-order income risk, we refer to risks that are captured by not only the variance of income shocks, but also their skewness and kurtosis. These higher order moments of the data can be a major source of risk for individuals as we show in this paper.

To provide a broad perspective on these questions, we study panel data on individuals and households from the United States, Germany, and Sweden, covering more than three decades of data for each country. It is useful to begin by putting our analysis in context. A broad range of empirical evidence indicates that idiosyncratic income risk rises in recessions. Earlier work in the literature was based on small survey-based panel datasets, such as the Panel Study of Income Dynamics (PSID), which required researchers to make parametric assumptions to obtain identification. The earlier studies in the literature have restricted attention to the changes in the mean and variance of income shocks and concluded that the variance of income shocks is countercyclical (e.g., Storesletten *et al.* (2004)). In recent work, Guvenen *et al.* (2014) used a very large panel dataset on earnings histories from the U.S. Social Security Administration (SSA) records. Using non-parametric techniques, they found that the variance of income shocks is very stable over time and is robustly acyclical, whereas the left-skewness of shocks varies significantly over time in a countercyclical fashion.

Despite important advantages, the SSA data also have three shortcomings: (i) earnings data are available only for individuals, and it is not possible to link household members to each other, (ii) no information is available on taxes and transfers (unemployment insurance, welfare payments, gifts, etc.), and (iii) no information is available on skills/education. Furthermore, Guvenen *et al.* (2014) focus on males with no corresponding information on women.

This paper makes three contributions. First, applying non-parametric techniques and using robust statistics, we document that the variance of individual labor earnings growth is flat and acyclical in all three countries, whereas the left-skewness of shocks is strongly countercyclical. Therefore, we conclude that applying the same method to survey and administrative data yields the same substantive conclusions.

Second, we find that the underlying variation in higher-order risk is remarkably similar across these countries that differ in many details of their labor markets. In

particular, in all three countries, the variance of earnings shocks is almost entirely constant over the business cycle, whereas the skewness of these shocks becomes significantly more negative in recessions.

Third, we find that insurance provided within households or by the government plays an important role in reducing downside risk, but that how and to what extent differs between the countries. Within-household provided insurance reduces the countercyclicality in the skewness of earnings in Sweden, but evidence of within-household insurance is much weaker in United States and in Germany. Government provided insurance, in the form of unemployment insurance, welfare benefits, aid to low income households, and the like, plays a more important role in all three countries; the effectiveness is weakest in the United States, and strongest in Germany.

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 (before-government) individual earnings and examines how the patterns of cyclicality vary by gender, education, and type of employment. Section 5 expands the analysis to households and includes various types of government social insurance policies to examine their impact on the cyclicality of higher-order risk. Section 6 uses a structural life-cycle 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.

1.1 Related Literature

[To be added]

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;[1] the Sample of Integrated Labour Market Biographies (SIAB[2]) 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 2000 households in the core sample, which is representative of the U.S. 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.[3]

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, which make 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[4] and with (ii) 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 and $t+k$, the panel contains individuals who are aged 25 to 59 in periods t and $t+k$ ($k = 1$ or 5) 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 $1885 US

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

[2]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).

[3]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 sub-periods.

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

dollars for the United States in 2010.[5] 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 labor earnings.

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,[6] that satisfy the age criterion and household income that satisfies the income criteria. At the household level, we will analyze several income measures. We start with labor earnings and then add various transfers, taxes, and capital income. To ensure that the sample is consistent across our analyses, the condition that earnings exceed our minimum threshold is imposed on the minimum earnings across all household income measures. Household earnings are converted into adult-equivalent units by dividing by the square root of household size.

Classifying Expansions and Recessions

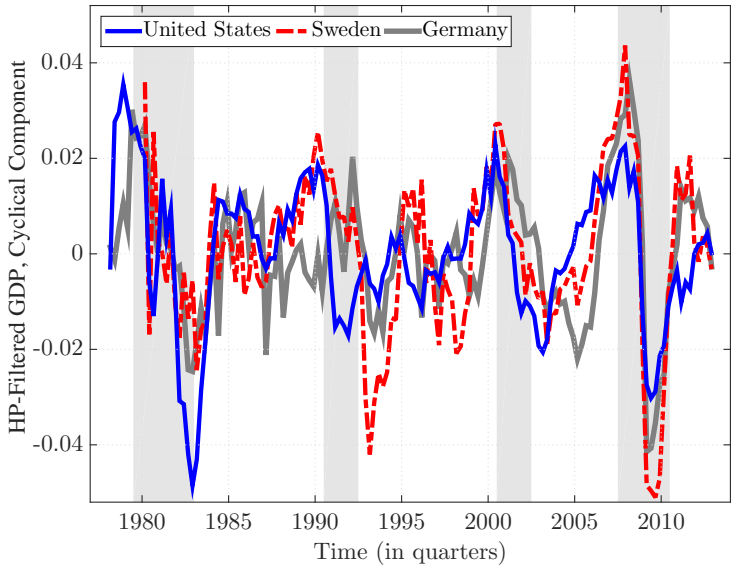
For the United States, the classification of expansionary and recessionary episodes is based on the NBER peak and trough dates, with small timing variations. Given the time span covered by our sample, we classify the following years as recessions: 1980–1983, 1991–92, 2001–2002, and 2008–2010. The main difference compared to the NBER list is that we treat the 1980–1983 period as a single "double-dip" recession because of the short duration of the intervening expansion and the lack of recovery in the unemployment rate. Based on this classification, there are four expansions and four recessions during our sample period.

For both Germany and Sweden, we base the dating of expansions and recessions on data from the Economic Cycle Research Institute (ECRI), which applies the NBER methodology to OECD countries since 1948. The classification is consistent with various aggregate measures of the German and Swedish economies, respectively. In the

[5]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 a 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.

[6]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.

Figure 1: Cyclical Component of Quarterly GPD Growth: U.S., Germany, and Sweden



Note: The shaded areas in the figure indicate U.S. recessions according to our classification described in the text. The series for Germany corresponds to West Germany up to and including 1990Q4, and to (Unified) Germany from 1991Q1 on. The cyclical component is obtained by HP-filtering the series for GDP per capita from 1970Q1 to 2014Q1.

time period covered by the panel data, recession periods for Germany (peak to trough) are from January 1980 to October 1982, January 1991 to April 1994, January 2001 to August 2003, and April 2008 to January 2009. Our sample period hence covers four recessions and four expansions. For Sweden, ECRI recession periods are from February 1980 to June 1983, June 1990 to July 1993, and April 2008 to March 2009. This leaves us with three recessions and three expansions during our sample period.

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 and plot moments $m\left[\Delta_s y_t\right]$, where $y_t \equiv \ln Y_t$ (natural logarithm) and $\Delta_s y_t \equiv y_t - y_{t-s}$. The moments m we consider are: the standard

deviation, the (Kelly) skewness, and the top (L90-50) and bottom (L50-10) tails.[7] For Germany and Sweden, s refers to 1- and 5-year changes. Due to the biennial structure of the PSID from the 1997 wave, our analyses of earnings for the United States refer to 2- and 4-year changes instead.[8]

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. This 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. Notice how the overall dispersion of the distribution increases here as well, but if we were to interpret this increase in isolation we would wrongfully conclude that not only one tail, 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.

We conclude that, when measuring income volatility, the tails should be explicitly analyzed. Furthermore, when relying on summary statistics of the distribution, limiting the analysis to the variance cannot possibly identify the nature of the change, yielding misleading results. Higher-order moments, like skewness, should be then considered. Note how any assumption on the distribution of income shocks would drive our results: a (log-) normal distribution cannot capture changes in skewness, for example. This is why, and in light of recent evidence on male earnings growth using administrative data for the United States (Guvenen *et al.*, 2014), we take a skeptical—non parametric—point of view.

[7] $L9050 \equiv P90 - P50$ in logs, and analogously for $L5010$.

[8] 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 $s = 2, 4$.

Broadening the Definition of Business Cycles

Some of the important macroeconomic variables do not perfectly synchronize with expansions and recessions, but their fluctuations might have an impact on earnings. For example, the U.S. stock market experienced a significant drop in 1987, during an expansion, and we can see in the time series analysis how the third moment falls sharply in that year. Similarly, the U.S. economy displayed an overall weakness in 1993–1994, which is evident in a range of economic variables, but these years are technically classified as part of an expansion by the NBER dating committee. Other examples are easy to find for Germany and Sweden (e.g., 1996). Therefore, the main focus of our analysis will be on the co-movement of higher-order moments of earnings changes with a continuous measure of business cycles.

For this part, we consider the four moments m defined above for the graphical analysis, and add two more. In particular, we compute correlations between GDP growth and (i) the standard deviation, (ii) the log differential between the 90th and 10th percentiles (L90-10), (iii) the skewness, measured as the third standardized moment, (iv) the Kelly's measure of skewness, and (v) the upper (L90-50) and (vi) lower (L50-10) tails. 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. Therefore, we consider the following regression of 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 :

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

For a quantitative interpretation of the results reported in the next sections, Table 1 reports the short- and long-run volatility of GDP growth for each country and year sample considered along the paper.

4 Empirical Results: Gross Individual Earnings

In this section, we examine the cyclical behavior of the dispersion and the skewness of earnings changes in gross labor earnings for individuals. By gross earnings we mean a worker's compensation from his/her employer before any kind of government intervention in the form of taxes, benefits, welfare, unemployment insurance, and so on. In the next section, we will turn to household earnings and construct various measures

Table 1: Short- and Long-Run GDP Growth Volatility: United States, Germany, and Sweden

		Std. Dev. of GDP Growth	
	Data period	short-run	long-run
United States	1969-2010	3.37%	4.28%
	1976-2010	3.34%	4.44%
Germany	1976-2010	2.01%	3.95%
	1984-2011	2.08%	3.83%
Sweden	1976-2010	2.36%	5.42%

Note: Short-run is 1-year difference for Germany and Sweden, and 2-year difference for the United States. Long-run is 5-year difference for Germany and Sweden and 4-year difference for the United States.

of household earnings that lead up to disposable income. We will then compare how the cyclical behavior of these higher-order moments is affected by insurance provided within the household and by various government social insurance policies.

We address three questions about higher-order risk in individual incomes. First, we ask if the counter-cyclicality of skewness and the acyclicality of dispersion found in U.S. administrative earnings data is also borne out in U.S. survey data, e.g., the PSID. This question is important because earlier papers that used the PSID *and* adopted parametric methods found strongly countercyclical variance of shocks. So the question is: is it the data set or is it the methodology that accounts for these different conclusions?

Second, we ask if the countercyclical skewness and the acyclical dispersion is a US-only phenomenon or a robust feature of business cycles that can be seen in other countries, such as Sweden and Germany, that differ greatly from the U.S. in many dimensions of their labor market structure. For example, Sweden has a very high unionization rate and a share of public employment that is nearly three times that in the United States.

Finally, we examine how the business cycle variation in higher-order income risk differs between (i) men and women, (ii) those with high- and low-education, and (iii) those employed in the private versus public sector. Examining cyclicality for these observationally distinct groups can shed light on the sources of the cyclicality, providing a deeper understanding of systematic components in the response to business cycle

Table 2: Cyclicality of Male Earnings, by Education Groups

	Std Dev	L9010	Skew	Kelly	L9050	L5010
			United States (PSID)			
All Males	0.20	–0.11	5.93**	1.67***	0.57***	–0.68***
	(1.02)	(–0.51)	(2.63)	(5.00)	(3.71)	(–3.96)
College Graduates	0.36*	0.25	2.90	0.58*	0.35*	–0.10
	(2.04)	(0.90)	(0.59)	(1.97)	(1.98)	(–0.64)
Non-College	0.10	–0.38	6.06	1.84***	0.52*	–0.90***
	(0.35)	(–0.84)	(1.58)	(4.17)	(1.83)	(–3.19)
			Sweden (LINDA)			
All Males	0.05	–0.11	9.91***	3.74***	0.91***	–1.01***
	(1.18)	(–1.22)	(4.08)	(4.00)	(3.80)	(–3.74)
College Graduates	0.04	–0.00	3.94**	1.80***	0.42	–0.42***
	(0.29)	(–0.01)	(2.16)	(4.93)	(1.58)	(–5.72)
Non-College	0.06	–0.17	11.21***	4.03***	0.99***	–1.26***
	1.07	(–1.52)	(4.09)	(3.86)	(3.39)	(–3.53)
			Germany (SIAB)			
All Males	0.07	0.15	14.42***	5.48***	0.95***	–0.80***
	(0.42)	(0.36)	(4.28)	(5.80)	(3.14)	(–4.11)
College Graduates	0.16	0.62	6.96**	4.70***	1.24**	–0.61**
	(1.06)	(1.01)	(2.11)	(3.10)	(2.17)	(–2.29)
Non-College	0.06	0.10	15.63***	5.26***	0.89***	–0.79***
	(0.31)	(0.25)	(4.73)	(5.41)	(3.07)	(–3.78)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in a 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%, 1%)-level.

fluctuations.

To answer these questions, we start by computing correlations between earnings innovations and GDP growth. Next, we plot some of the different moments over time and inspect their fluctuations over the business cycle.

Cyclicality of Dispersion

In Table 2, we report the cyclicality of six key statistics computed from the distribution of earnings changes of male 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. The

first column reports the cyclicality of the standard deviation. In the United States, the standard deviation is acyclical, as seen from the small (0.20) and statistically insignificant (*t-stat* of 1.02) coefficient.[9] In the next column, we report another measure of dispersion—the log 90-10 differential—which is just as acyclical. Therefore, the PSID data is consistent with the findings of Guvenen *et al.* (2014) regarding the acyclicality of dispersion from the much larger SSA administrative data.

A natural follow up question is whether this acyclicality is specific to the United States, or whether it also holds in Sweden and/or Germany, which in many ways have very different labor markets. As seen in the first column of the middle panel, both measures of dispersion are acyclical in Sweden, with very small and insignificant coefficients. Turning to Germany (bottom panel), standard deviation and L90-10 are again acyclical.[10]

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 (left panel of Figure 2). These figures should be compared with typical calibrations in the literature that assume the volatility of earning 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

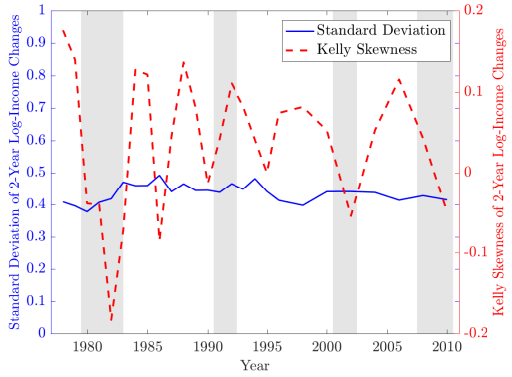
We next turn to the cyclical behavior of skewness. Column 3 reports the third standardized moment of earnings changes. While this measure is well known, it has a tendency of being sensitive to outliers, which can be a concern for the U.S. and German (SOEP) data, since these are surveys, with possibly large measurement error and modest sample sizes. This is less of an issue for Sweden and the German SIAB data, given the larger sample size and higher data quality. Nevertheless, to alleviate such concerns, in column 4 we also report another measure of asymmetry, called Kelly's

[9] We have repeated this calculation using moments from the SSA data, as reported in Guvenen *et al.* (2014). The results for all males turn out to be surprisingly similar, but even stronger. In particular, the coefficients for each of the 6 moments are -0.18^{*}, -0.07, 4.73^{***}, 2.31^{***}, 2.31^{***}, 1.02^{***}, -1.09^{**}, respectively.

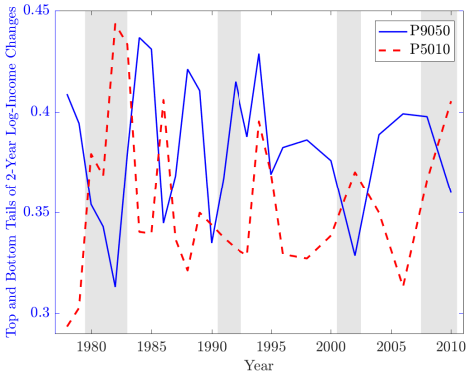
[10] All regression results based on SIAB data are robust to various robustness checks that address issues of top-coding and a structural break in the wage variable. See appendix B for details.

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

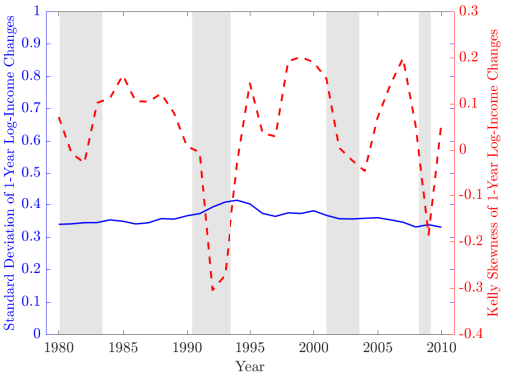
(a) United States, SD (left) and KS (right)



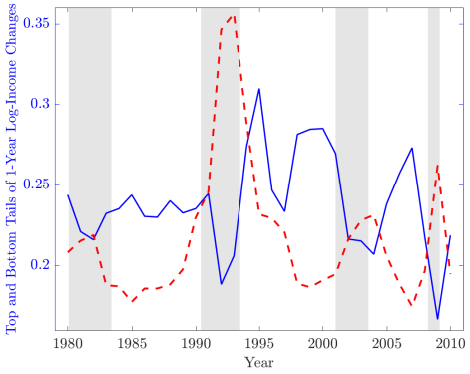
(b) United States, Upper and Lower Tail



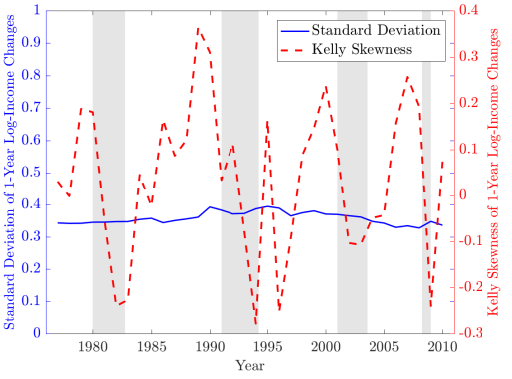
(c) Sweden, SD (left) and KS (right)



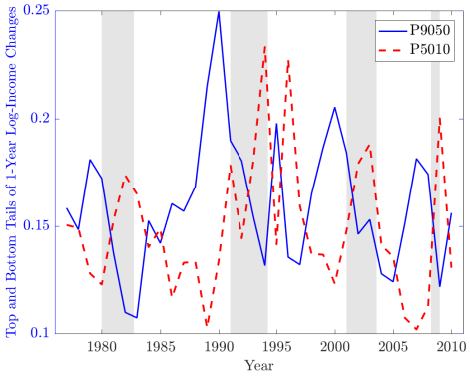
(d) Sweden, Upper and Lower Tail



(e) Germany, SD (left) and KS (right)



(f) Germany, Upper and Lower Tail



Note: Linear trend removed, centered at sample average.

skewness, defined as:

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

This measure has several attractive features. First, it is much less sensitive to extreme observations, since it does not depend on observations beyond the 90th and 10th percentiles of the distribution. This deals with the concern about potential outliers. Because of this advantage, it is our preferred measure of skewness, especially for the U.S. and Germany where measurement issues could be more important. Second, the particular value of Kelly's skewness has a simple interpretation, in terms of the relative lengths of the top and bottom tails. In particular,

$$\frac{P90 - P50}{P90 - P10} = 0.5 + \frac{\mathcal{S}_k}{2}, \tag{1}$$

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

Armed with these definitions, we turn to Table 2. In all three countries, Kelly's skewness is procyclical and (statistically) significant at the 1% level. The coefficient is about 1.7 for the U.S., double (3.7) for Sweden, and about 5.5 for Germany, showing more cyclicality when moving from the U.S. to Sweden and most for Germany. 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 $2 \times 0.0236 = 0.0472$), Kelly's skewness will fall by $0.0472 \times 3.7 = 0.18$. 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 (1) 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 (due to the high unionization rate, among others). Finally, the coefficient on the third moment measure is also positive in all three countries, consistent with Kelly's skewness, and is significant at the 1% level in both Sweden and Germany, and at the 5% level in the United States.[11]

[11] The corresponding changes in $\mathcal{S}_k$ for the U.S. and Germany are: 0.11 and 0.22 respectively.

Inspecting the Tails

At the expense of some oversimplification, it might be useful to think about a shift towards 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? The last two columns of Table 2 report the cyclicality of the L9050 and L5010. Notice that in all three countries the top tail is procyclical, whereas the bottom tail is countercyclical. 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 towards negative skewness happens through 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 the U.S., the coefficient for L9050 is 0.57 and for L5010 is –0.68. The corresponding coefficients are 0.91 and –1.01 for Sweden, 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 the 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 variance, can miss important asymmetries that can 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 variance (which would imply an expansion/compression of *both* tails).

It is useful to dig a bit deeper to see if the patterns regarding higher-order risk documented so far are concentrated to certain subgroups of the economy or whether they are pervasive across the economy. For this purpose, we examine the same set of statistics separately by (i) gender groups, (ii) skill groups, and (iii) private- versus public-sector workers.

Table 3: Cyclicality of Female Earnings, by Education Groups

	Std Dev	L9010	Skew	Kelly	L9050	L5010
	United States (PSID)					
All females	0.45***	0.40*	–1.42	0.62*	0.48**	–0.08
	(3.56)	(1.85)	(–0.54)	(1.97)	(2.61)	(–0.52)
College graduates	0.25	–0.60	–5.00	1.08*	0.18	–0.78**
	(1.07)	(–1.70)	(–1.09)	(1.77)	(0.57)	(–2.50)
Non-college	0.54***	0.79***	0.49	0.59	0.67***	0.12
	(4.59)	(3.59)	(0.25)	(1.46)	(3.14)	(0.53)
	Sweden (LINDA)					
All females	0.27***	0.43**	4.39**	1.64***	0.67***	–0.24**
	(3.26)	(2.24)	(2.76)	(3.33)	(3.09)	(–2.67)
College graduates	0.05	0.13	3.06*	1.15***	0.64	–0.25
	(0.28)	(0.31)	(1.82)	(4.03)	(1.22)	(–1.74)
Non-college	0.31***	0.50*	5.72***	1.81***	0.75***	–0.25**
	(3.04)	(1.96)	(3.52)	(3.40)	(2.78)	(–2.71)
	Germany (SIAB)					
All females	0.10	0.34	4.34*	2.55**	0.80	–0.46*
	(0.47)	(0.48)	(1.77)	(2.05)	(1.25)	(–1.80)
College graduates	0.04	0.01	1.55	2.03	1.01	–1.00
	(0.14)	(0.01)	(0.93)	(1.65)	(1.12)	(–1.39)
Non-college	0.10	0.32	4.58	2.58**	0.77	–0.45*
	(0.47)	(0.47)	(1.69)	(2.08)	(1.27)	(–1.88)

Note: See Table 2 for explanations.

4.1 Differences by Gender

We now turn to the cyclicality of higher-order risk for female workers (Table 3) and examine how they compare to the patterns for males. Focusing on the first row of each panel, we see three main patterns. First, the standard deviation of earnings changes is *procyclical* for U.S. and Swedish women but acyclical for German women. This is different from men, who displayed acyclical dispersion in all countries. Second, Kelly's measure of skewness is always procyclical—left-skewness is countercyclical—as indicated by the positive coefficient on log GDP growth, which is highly significant for Sweden (1% level), significant for Germany (5% level), and only slightly significant for the U.S. (10% level). The third central moment is only significantly procyclical in Sweden and Germany. As noted before, this might be due to the smaller sample size for the U.S.

Third, inspecting the top and bottom tails separately (last two columns), we observe the expected pattern of cyclicality, whenever the coefficient is significant. In particular, L9050 is procyclical and significant for the U.S. and Sweden, whereas the L5010 is countercyclical and significant for Sweden and Germany.[12] Thus, just as for 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 variance is not an ideal statistic to focus on when it comes to measuring higher-order earnings risk over the business cycle. The case of U.S. women is an excellent illustration of this point: the highly significant procyclicality of variance is entirely driven by the upper tail. Finally, it is worth noting that the magnitudes of the fluctuations in both Kelly's skewness and in the upper and lower tails separately are somewhat attenuated for women compared with men.

4.2 Differences By Education Level

Economists have studied extensively how the average earnings and employment of different skill groups vary over the business cycle. We divide workers into two groups—college graduates and non-college workers—based on the highest education level they have acquired. Starting with males (Table 2), in all three countries, both education groups display procyclical Kelly's skewness that is statistically significant. In the United States and Sweden, the magnitude of cyclicality is quite a bit stronger for less educated workers—about three times stronger in both the U.S. and Sweden. In Germany, the difference between the groups goes in the same direction but is quantitatively much smaller. Again, turning to standard deviation, there is not a very clear pattern in any country: variance is either acyclical or when it is pro-cyclical, the magnitude is small.

Turning to women in Table 3, Sweden again emerges as the country with the clearest patterns. Skewness is strongly procyclical for both education groups; the lower tail is counter-cyclical for non-college and the upper tail is procyclical for non-college graduates. The variance for non-college graduates is also procyclical, although the magnitude is quite small. As discussed before, this happens because the top end of the shock distribution collapses more than the expansion of the bottom end during recessions.

[12] It is somewhat surprising that women in the U.S. seem to face less downside risk as measured by the L5010 differential compared with these two European countries.

Table 4: Cyclicality of Individual Earnings, by Sector of Employment, Males

	SD	L9010	Skew	Kelly	L9050	L5010
			Sweden (LINDA)			
Private	0.12	0.10	13.62***	3.83***	0.93***	–0.83***
	(1.45)	(0.93)	(3.31)	(4.02)	(3.81)	(–4.08)
Public	–0.08	–0.45***	9.24***	2.10***	0.17	–0.62***
	(–1.34)	(–3.93)	(6.40)	(6.55)	(1.64)	(–9.11)
			Germany (SIAB)			
Private	0.06	0.03	15.81***	5.55***	0.88***	-0.85***
	(0.38)	(0.08)	(4.18)	(6.44)	(3.55)	(–5.64)
Public	0.25	2.50	5.90**	0.30	1.45	1.06
	(0.71)	(1.16)	(2.04)	(0.17)	(1.08)	(1.01)

Note: See Table 2 for explanations.

In the U.S., the variance is more robustly procyclical for less educated women and acyclical for college graduates. In Germany, both non-college female workers and college graduates display acyclical dispersion of earnings changes. The overall pattern observed for non-college workers resembles closely the estimates for all females, where a countercyclical lower tail (L5010) appears to be driving the volatility of earnings changes, which displays significantly procyclical skewness. For more educated workers, the estimates are statistically insignificant. Again, the lack of significance for the U.S. estimates and the German college-educated might be due to the lower share of employment among women, especially in the earlier part of the sample, due to the relatively modest sample size (in the U.S.) and low female college share (in the early part of the sample in Germany).

4.3 Differences Between Private- vs. Public-Sector Workers

One of the most pronounced differences between the economies of the United States and European countries, and Sweden in particular, is the size of the public employment sector in the latter. 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. Because of data limitations for the U.S., we are able to conduct this analysis for Sweden and Germany only.

Our data set does not include a direct indicator of public sector employment. However, some sectors in Sweden and Germany are dominated by public sector jobs (or, more broadly, by jobs funded by the public). So, we define a worker as working in the public sector, if he/she works in public administration, health care, or education in both years t and $t+k$ (where $k = 1, 5$).[13] The split between private/public employment varies substantially by gender: in Sweden about 23% of men work in the public sector compared with a whopping 63% for women (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.

Table 4 reports the cyclicality regressions separately for workers in private sector versus public employment (by country). Interestingly, the results suggest stark differences for the two European countries. First, the standard deviation shows acyclicality for both countries in both sectors. Looking at the measures of skewness, the results display strong procyclicality for both sectors in Sweden, but only for the private sector in Germany–the third standardized moment displays cyclicality, which is not reflected by the more stable Kelly measure. Still, the magnitude of the procyclicality of Kelly's measure is lower for the Swedish public sector (2.10) compared with the private sector (3.83). A reasonable initial reaction is that this might be due to the lower tail risk being better insured in the public sector. However, this conjecture is wrong. Looking at the top and bottom tails separately it is evident that the L5010 gap fluctuates by comparable magnitudes for both groups and is in fact statistically significant at that 0.1% level or more. What is different is the top tail: it compresses strongly for private sector employees, whereas it is acyclical in the public sector. As a side note, because of this, the L9010 measure of dispersion is slightly countercyclical in the public sector, unlike for private sector workers. In the case of Germany, however, the strong cyclicality of both tails for private sector workers is not mirrored in the public sector, and it appears that earnings movements for publicly employed do not react to the business cycle.

Turning to female workers (table 5), we see the variance of earnings changes being slightly procyclical in both sectors in Sweden (as was the case for all women). Turning to skewness, it is again procyclical for both sectors, with a somewhat smaller coefficient in the public sector as was the case for males. However, the top tail is procyclical and the lower tail is countercyclical in both sectors. The magnitudes however are smaller

[13] Historically most workers in these sectors were employed by the public; this is less true today.

Table 5: Cyclicality of Individual Earnings, by Sector of Employment, Females

	SD	L9010	Skew	Kelly	L9050	L5010
			Sweden (LINDA)			
Private	0.34**	0.50*	8.46***	1.99***	0.78**	–0.29**
	(2.47)	(1.87)	(3.47)	(3.02)	(2.81)	(–2.43)
Public	0.20**	0.18	3.00**	1.10***	0.34**	–0.16**
	(2.38)	(1.19)	(2.63)	(3.29)	(2.43)	(–2.61)
			Germany (SIAB)			
Private	0.05	0.01	5.45*	3.13**	0.73	–0.72***
	(0.34)	(0.01)	(1.84)	(2.44)	(1.50)	(–3.15)
Public	0.16	1.17	1.70	0.95	0.85	0.32
	(0.49)	(0.84)	(0.76)	(0.68)	(0.85)	(0.59)

Note: See Table 2 for explanations.

in the public sector. In the case of Germany, the results for the private sector closely resemble those of all women–even the L9050 now reacts significantly– while, as for males, the distribution of earnings changes in the public sector appears to be acyclical.

Overall, it is somewhat surprising that, even for workers in the public sector and 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, even with very different labor market institutions. The results for Germany are in line with initial conjectures of public sector employment being less affected by general economic conditions.

4.4 Differences Across Occupations

From industries 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 now run the cyclicality regressions separately for each occupation,

where a worker contributes to the earnings changes of occupation j from t to $t+1$ if in year t he or she works in that occupation.

We first consider the highest level of disaggregation in the KldB-88, which defines five *occupational areas*. This allows us to see how different broad groups of occupations differ in terms of their cyclicality. The occupational areas are (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. First considering males, the upper panel of table 6 shows that the beta-coefficients for the regressions of the dispersion measures are statistically insignificant for all five occupational areas. The tail measures and both skewness measures display the same pattern as the aggregate regressions for manufacturing occupations, technical occupations, and service occupations. Farming and related occupations, display acyclicality for the tail measures, with only the central moment based skewness measure displaying pro-cyclicality. Mining and related occupations show pro-cyclical upside chances as measured by L9050, which drives overall pro-cyclial skewness.

Turning to females, the lower panel of table 6 shows similar results. Women working in an occupation related to mining, or mineral extraction; to manufacturing, or fabrication; or in technical occupations show the by now familiar pattern of lower tail countercyclicality: downside earnings changes are more pronounces in contractions. In the second occupation group, mining/mineral extraction, this downside risk drives overall dispersion measured by the standard deviation to be countercyclical. Women working in service occupations face an acyclical distribution of earnings changes, similarly women in farming and related occupations.

Next, we consider a more disaggregated analysis and rerun the regressions for 30 *occupational segments*. The upper panels of tables 7 and 8 summarize the distributions of the coefficient of interest across the 30 regressions for males and females, respectively.[14] The lower panels summarize the corresponding distributions of t-statistics. Again, first considering males, we see that there is some variation across occupations in terms of their earnings cyclicality. Still, the general pattern seen for aggregate male earnings is dominant across occupations. Consider the t-statistics for the lower tail cyclicality: the median clearly displays statistically significant countercyclicality; and while some occupational segments are characterized by acyclical earnings changes,

[14] E.g., the first entry gives the mean of the 30 beta coefficients from the occupation-specific regression, the second gives the 10^{th} percentile, and so on.

Table 6: Cyclicality of Earnings by occupational area; Germany (SIAB)

	SD	L9010	Skew	Kelly	L9050	L5010
				Males		
Farming and related	0.77	4.56	31.97***	5.64	3.80	0.76
	(1.13)	(1.23)	(3.73)	(1.51)	(1.52)	(0.45)
Mining, Mineral Extraction	0.60	2.62	11.87***	3.23	1.32**	1.30
	(1.28)	(1.25)	(3.17)	(1.39)	(2.43)	(0.72)
Manufacturing, Fabrication	0.17	0.17	37.28***	11.39***	2.00***	-1.83***
	(0.40)	(0.20)	(4.91)	(5.53)	(3.21)	(-3.99)
Technical Occupations	0.19	0.13	20.37*	12.36***	1.51**	-1.38***
	(0.55)	(0.19)	(1.75)	(4.04)	(2.72)	(-3.64)
Service Occupations	0.16	0.59	22.03***	8.89***	1.76**	-1.17***
	(0.54)	(0.68)	(4.42)	(3.92)	(2.41)	(-3.09)
				Females		
Farming and related	1.32	2.90	12.01*	0.96	2.06	0.84
	(1.10)	(0.73)	(1.94)	(0.31)	(0.71)	(0.61)
Mining, Mineral Extraction	-4.24*	-5.59	13.23	12.26	1.61	-7.20**
	(-1.74)	(-1.02)	(0.64)	(1.54)	(0.34)	(-2.59)
Manufacturing, Fabrication	0.12	-0.72	15.18**	10.59***	2.48*	-3.20***
	(0.27)	(-0.48)	(2.26)	(4.95)	(2.00)	(-6.01)
Technical Occupations	0.20	-0.75	7.15	8.44**	1.41	-2.16***
	(0.56)	(-0.83)	(0.87)	(2.70)	(1.56)	(-2.82)
Service Occupations	0.21	0.85	7.21	4.09	1.45	-0.60
	(0.51)	(0.59)	(1.61)	(1.63)	(1.13)	(-1.15)

Note: See notes for Table 2.

there is no occupation, for which the lower tail is procyclical (the maximum value of t-statistics is 0.72). The mirror image holds for the upper tail cyclicality, and while a small number of occupations displays cyclicality of the L9010, this vast majority is described by acyclical dispersion: the 10^{th} and 90^{th} percentiles of the t-statistics are small in absolute value.

Turning to females, again, the lower tail is countercyclical for most, and acyclical for some occupational segments. Similarly, the upper tail is procyclical for most, acyclical for some occupational segments. The dispersion measures also reflect what we see for males and the dispersion of the earnings change distribution is acyclical for most occupations. For both males and females the tail movements translate into the cyclicality procyclicality of Kelly's skewness in the by now familiar way.

Summing up, we find that broad occupational groups experience similar cyclicality

Table 7: Cyclicality of Earnings by occupational segments – Males; Germany (SIAB)

	SD	L9010	Skew	Kelly	L9050	L5010
	Distribution of Beta Coefficients					
Mean	0.18	0.71	28.89	8.78	2.06	-1.35
P10	-0.39	-0.83	8.40	3.68	0.90	-2.91
Median	0.20	0.46	27.18	8.35	1.75	-1.29
P90	0.60	2.62	42.92	13.09	3.29	-0.18
Standard Deviation	0.38	1.30	16.42	3.52	1.37	1.47
Min	-0.91	-1.55	-7.05	2.36	0.66	-7.09
Max	0.87	4.56	65.22	17.87	7.89	1.30
	Distribution of t-Statistics					
Mean	0.43	0.34	3.24	4.11	2.59	-2.21
P10	-0.66	-0.90	0.93	1.39	1.43	-4.93
Median	0.45	0.46	3.19	3.23	2.46	-2.22
P90	1.13	1.19	4.69	7.73	3.73	-0.18
Standard Deviation	0.76	0.87	1.45	2.71	1.07	1.79
Min	-1.31	-2.30	-0.53	1.01	0.99	-6.62
Max	2.04	2.01	6.72	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 2.

with farming and mining related occupations being less cyclical. Regressions at finer level of disaggregation point towards interesting heterogeneity of earnings cyclicality across occupations.

4.5 Cyclicality of Earnings vs. 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. The SIAB contains detailed 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

Table 8: Cyclicality of Earnings by occupational segments – Females; Germany (SIAB)

	SD	L9010	Skew	Kelly	L9050	L5010
	Distribution of Beta Coefficients					
Mean	0.04	-0.10	11.28	6.87	1.98	-2.08
P10	-2.06	-3.33	0.16	0.96	0.89	-4.45
Median	0.26	0.12	10.73	6.30	2.06	-2.00
P90	0.89	1.97	24.83	12.26	3.10	0.69
Standard Deviation	1.12	2.25	12.63	3.94	1.37	1.96
Min	-4.24	-6.40	-29.52	-0.56	-4.13	-7.20
Max	1.32	3.61	46.52	13.50	3.80	0.84
	Distribution of t-Statistics					
Mean	0.36	-0.11	1.39	2.35	1.43	-1.84
P10	-1.20	-1.53	0.01	0.31	0.53	-4.28
Median	0.52	0.05	1.23	1.81	1.38	-1.35
P90	1.11	0.71	2.69	5.41	2.27	0.52
Standard Deviation	0.84	0.80	1.10	1.81	0.89	1.73
Min	-1.74	-2.06	-0.94	-0.07	-0.72	-6.30
Max	1.83	0.95	3.99	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 2.

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); Card *et al.* (2013).[15]

In Table 9, rows 1 and 4 reproduce the results from Tables 2 and 3 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 are 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.[16] Naturally, the dispersion of earnings

[15]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.

[16]The sample of full-time female workers contains about 73% of women (who make for only 54% of the observations) that contribute to the measures of earnings change for women. The corresponding

Table 9: Cyclicality of Individual Earnings vs. Wages; Germany (SIAB)

	SD	L9010	Skew	Kelly	L9050	L5010
				Males		
Earnings	0.07	0.15	14.42***	5.48***	0.95***	–0.80***
	(0.42)	(0.36)	(4.28)	(5.80)	(3.14)	(–4.11)
Full-Time Wages	0.01	–0.09	14.55***	4.73***	0.30***	–0.39***
	(0.23)	(–0.54)	(4.58)	(6.31)	(3.77)	(–3.20)
Full-Time Wages	–0.01	–0.12	17.82***	4.98***	0.28***	–0.40***
(Firm Stayers)	(–0.26)	(–0.81)	(4.54)	(5.78)	(3.29)	(–3.20)
				Females		
Earnings	0.10	0.34	4.34*	2.55**	0.80	–0.46*
	(0.47)	(0.48)	(1.77)	(2.05)	(1.25)	(–1.80)
Full-Time Wages	0.04	0.03	8.98*	2.12***	0.17**	–0.14
	(0.66)	(0.18)	(2.02)	(5.11)	(2.61)	(–1.58)
Full-Time Wages	0.03	0.02	11.96*	2.28***	0.16***	–0.14
(Firm Stayers)	(0.53)	(0.13)	(1.70)	(4.84)	(3.17)	(–1.61)

Note: See notes for Table 2.

changes is 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.

A question that remains is 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 both year t and t+1.[17] 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 even for full time workers and not by hours.

figures were 88% of individuals and 82% of observations for males. This implies that part-time employment plays a more important role for the female sample.

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

4.6 Who suffers in recessions?

The results presented so far do not exploit the panel structure of our data beyond considering one-year (respectively two-year) earnings changes. When thinking of downside labor market risk in recessions, an important question is whether this risk is concentrated at a group of workers who experience bad luck repeatedly–or if income changes during the last recession are forgotten when it comes to the current one. We provide an answer to this question by analyzing the distribution of income changes during a given recession conditional on the income change during the last recession and a set of individual characteristics.

Formally, we estimate a set of quantile regressions of the following form

$$Q_\tau\left(\triangle^{rec}_{s(t_1)} y_{i,t_1}\right) = \alpha(\tau) + \beta_1(\tau)\triangle^{rec}_{s(t_0)} y_{i,t_0} + \beta_2(\tau)\,\text{std}\left(\{\triangle_k y_{i,t}\}^{t_1}_{t=t_0+s(t_0)+1}\right) + \mathbf{x_{i,t_1}}\gamma(\tau) + \varepsilon_{i,t_1}, \tag{2}$$

with the τ^{th} quantile of the income change during the recession starting in year t_1 on the left-hand-side. The two explanatory variables of main interest are $\triangle^{rec}_{s(t_j)} y_{i,t_j} \equiv \left(y_{i,t_j+s(t_j)} - y_{i,t_j}\right)/s(t_j)$ and $std\left(\{\triangle_k y_{i,t}\}^{t_1}_{t=t_0+s(t_0)+1}\right)$: the log income change during recession j normalized by the number of years of the recession, and the standard deviation of transitory income changes ($k = 1$ or 2 depending on the data set) between two recessions, respectively. $\alpha(\tau)$ is the quantile-specific constant and$\mathbf{x_{i,t_1}}$is a vector of observables at the beginning of the recession. It includes a dummy for public sector employment, a quadratic in age interacted with education dummies, as well as dummies for the recessionary period.

Recessions are defined based on NBER recession dates in the case of the US and on ECRI recession dates for both Sweden and Germany, as discussed in section 2. We calculate the income change during a recession as the annualized change of annual earnings from the last non-recessionary year to the last year of the recession.[18] For the US and Germany our data covers four recessionary periods, and in the case of Sweden we exploit three recessions.

Figure 3 plots the sets of coefficients $\beta_1(\tau)$ and $\beta_2(\tau)$for a wide range of τ. The patterns of the estimates are widely symmetrical for Germany and Sweden. In the case

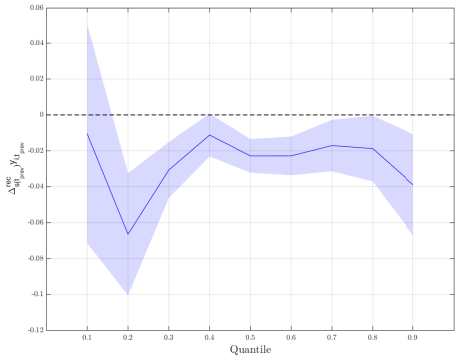
[18]For example, the 1990s recession in Germany is dated from 01/1991 to 04/1994 and the early 2000s recession is dated from 01/01 to 08/2003. This implies that we calculate the recessionary income changes as $\triangle^{rec}_{s(1990)} y_{i,1990} = (y_{1994} - y_{1990})/4$ and $\triangle^{rec}_{s(2000)} y_{i,2000} = (y_{2003} - y_{2000})/3$. The volatility between these two recessions is defined as the standard deviation of income changes $\{y_{1995} - y_{1994}, y_{1996} - y_{1995}, ..., y_{2000} - y_{1999}\}$

of PSID the results are less clear, which might in part be due to sample size. Consider first the point estimate on the income change during the previous recession in the case of Germany at the median of current income change, which is equivalent to what a regular OLS regression would tell us based on the mean. The estimation implies that, after conditioning on other observables, the median worker's income change during the current recession is the more negative, the bigger the income gain during the last recession was. This points towards mean reversion across recessions: those who gained a lot during the last recession, on average, loose during the current recession. The positive slope of the estimates from the quantile regressions allow us to see that the dispersion of conditional income changes increases with last recession's income gains: bigger past income gain shifts down the quantiles below the median more than the median, while the quantiles above the median are shifted up (or less down than the median). The point estimates imply that an increase of the lagged income gain by 1 log point leads to a widening of the L9010 log differential by 6 log points.. If the coefficients lined up on a straight line crossing the point estimate of the median, this increase in dispersion would be completely symmetric around the median. While this roughly holds for the quantile range 0.2–0.8, we see quite some asymmetry once going beyond that range: the point estimates at the median, 10^{th} and 90^{th} percentiles imply that about 57% of the widening of the L9010 comes from dispersion below the median.

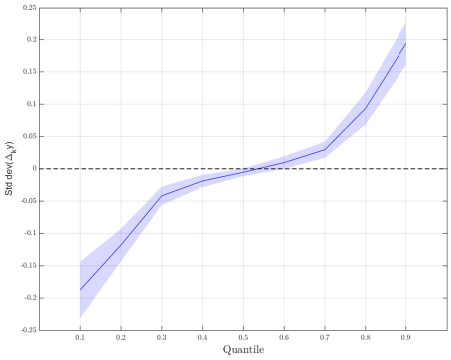
Summing up, we find that larger income gains (losses) in the previous recession imply a downward (upward) shift of the median worker's conditional income change, the distribution is more (less) dispersed and more skewed to the left (right) as measured by Kelly's skewness. Along the same lines, we can discuss the estimates for the second variable of interest, the volatility of in-between-recession income changes, which display the same pattern. We find that the median of the conditional distribution of income changes does react positive, while overall dispersion around the median increases symmetrically with the volatility measure. This is rather intuitive in that it simply states that those workers who are on a highly volatile income path between the recessions face a rather wide distribution of possible outcomes during the recession. While the result on previous recession income changes is less clear, the US estimates based on PSID display the same pattern when it comes to between-recession volatility.

Figure 3: Who suffers in recessions–quantile regression estimates: United States, Sweden, and Germany (SIAB); All Males.

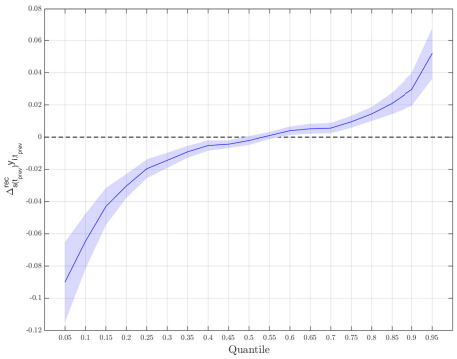
(a) United States, Change Previous Recession



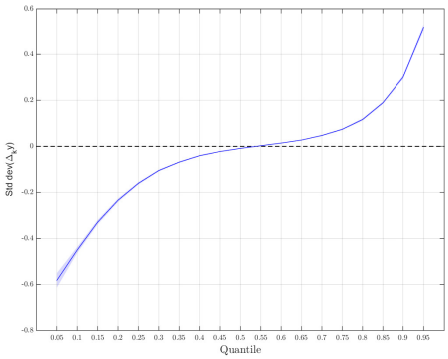
(b) United States, Volatility btw. Recessions



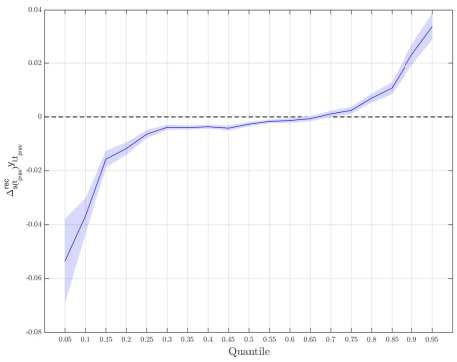
(c) Sweden, Change Previous Recession



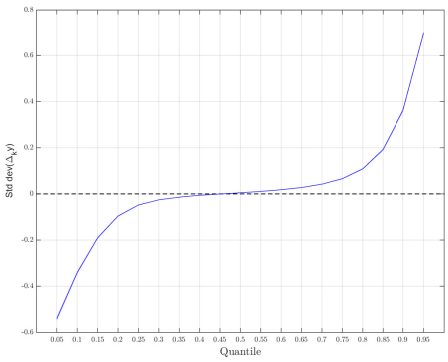
(d) Sweden, Volatility btw. Recessions



(e) Germany, Change Previous Recession



(f) Germany, Volatility btw. Recessions



Note: Figures plot the point estimates on lagged recession income change and on between-recession volatility from quantile regressions (2). For details see text. The 95% confidence bands are constructed using the standard errors from 200 bootstrap repetitions.

5 Introducing Insurance

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

5.1 Within-Family Insurance

In the previous section, we have shown that higher-order moments drive individual earnings risk over the business cycle. 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 by individuals in cohabitation.[19] 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 income. 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 we rely on SOEP data instead.

Mixed Evidence of Within-Family Insurance

The first row of each panel in Table 10 displays the cyclicality of each moment of household earnings changes. In order to get a feeling for the decrease (or increase) of exposure to business cycle fluctuations, we compare these results to the corresponding measures for male earnings from Table 2. Additional evidence comes from the graphical analysis of the dispersion, skewness, and the tails, of male earnings changes and household earnings changes in Figures 4 and 5, respectively.

Addressing the fact that we use different data, the panel for Germany also displays the regression results for male earnings in the SOEP panel. The qualitative results using SIAB and SOEP data broadly line up: standard deviation of earnings changes is acyclical, the lower tail is strongly countercyclical, and skewness—as measured by either the third central moment or using Kelly's measure—is procyclical. The upper tail is not sensitive to the cycle, whereas the lower tail responds strongly to aggregate fluctuations. This causes overall dispersion to be slightly countercyclical when

[19] Only 12% of our benchmark individual sample in the United States lives in a singe-person household, for example.

Table 10: Cyclicality of Household Earnings

	Std.Dev	L9010	Skew	Kelly	L9050	L5010
			United States (PSID)			
Earnings	0.40**	0.23	4.67***	1.97***	0.93***	–0.71***
	(2.10)	(0.74)	(3.01)	(6.17)	(4.96)	(–3.20)
Post-Gov	0.42***	0.59**	0.59	1.17***	0.72***	–0.14
	(3.25)	(2.44)	(0.19)	(3.13)	(3.42)	(–0.86)
Disposable	0.48**	0.63*	2.67	1.13***	0.74***	–0.12
	(2.24)	(1.90)	(1.05)	(4.83)	(3.75)	(–0.65)
			Sweden (LINDA)			
Earnings	0.10	–0.02	8.92***	2.24***	0.50***	–0.52*
	(1.20)	(–0.08)	(4.96)	(3.33)	(4.94)	(–2.00)
Post-Gov	–0.11	–0.41*	–1.06	0.94**	–0.03	–0.38**
	(–1.45)	(–2.00)	(–0.40)	(2.38)	(–0.44)	(–2.33)
Disposable	–0.13	–0.43	5.16**	1.50***	0.06	–0.49**
	(–1.29)	(–1.64)	(2.45)	(3.89)	(0.61)	(–2.67)
			Germany (SOEP)			
Males	–0.12	–1.33**	11.09*	1.76***	–0.21	–1.12***
	(–0.54)	(–2.33)	(2.05)	(5.95)	(–0.75)	(–3.56)
Earnings	–0.12	–1.31***	1.00	1.88**	–0.05	–1.26***
	(–0.70)	(–3.60)	(0.49)	(2.68)	(–0.18)	(–4.26)
Post Gov	0.13	–0.18	–8.89*	0.66	0.07	–0.25
	(0.61)	(–1.09)	(–1.78)	(0.85)	(0.32)	(–1.28)
Disposable	0.18	–0.16	–6.21	0.56	0.05	–0.22
	(1.05)	(–1.11)	(–1.52)	(0.67)	(0.21)	(–1.19)

Note: See notes for Table 2.

measured by L9010, again driven by one of the tails alone. Overall, the measure of downside risk is robust across data sets, while upside chances in booms seem to be slightly underestimated in SOEP.

Considering cyclicality of dispersion, the patterns and magnitudes for household earnings line up with the ones described for individual male earnings for all countries: household earnings changes display no cyclicality of dispersion. This is true especially for Sweden and Germany, while in the United States the positive coefficient (0.40) becomes significant at the 5% level. However, this behavior of the variance—and the L9010—is driven by asymmetric movements of the tails, as we will comment on in the next paragraph. The countercyclical measure of dispersion (as measured by L9010) for

Germany is again driven by the lower tail and thus the overall pattern here mirrors the one of male earnings dispersion.

The analysis of Kelly's skewness—and the inspection of the tails—yields very interesting results when comparing the three countries. In Sweden, intra-family insurance plays an important role in reducing downside risk over the business cycle as captured by a coefficient on Kelly's skewness of about 2.2 (compared to 3.7 for male earnings). This difference is mainly driven by the reaction of the lower tail being halved when moving from male earnings to household earnings. Repeating the illustrative calculation from above, this would imply a move from an upper tail to lower tail ratio of 50/50 in a typical expansion to 45/55 in a recession—much smaller compared to the change to 41/59 for male earnings.

Evidence of within-family insurance is weaker for the United States and Germany. In both economies, the results are slightly in favor of higher downside risk in recessions as measured by Kelly's skewness in the case of gross household earnings than in the case of individual male earnings. The differences are rather small, though. Considering the tails separately for the two countries, reveals important differences. While the slightly stronger reaction of Kelly's skewness is driven by higher procyclicality of upward movements in household earnings as compared to male earnings in the United States, the opposite is true in Germany. As for male earnings, the upper tail of earnings changes is not cyclical in Germany and the lower tail widens more for household earnings.

In order to shed further light on the insurance within households, we consider the cyclicality of income for actual households in comparison to income changes for randomly formed couples. This way we want to see if there is anything special about households visible in the data, or if the dynamics of household income just represent the dynamics of male and female income. We therefore randomly pair heads and spouses for each t to t+1 change. For each random couple, we make sure that artificial income is above the lower income. The first set of results in each country panel of table 11 shows the bootstrapped mean, standard deviation and 10-90 confidence band of the regression coefficients. In both the US and Germany, we find the random couples to experience lower downside risk than actual households as measured by the cyclicality of L5010. For Sweden, the random couples' L5010 shows the same cyclicality as actual households. The next rows show the same results when not randomly pooling all heads and spouses, but controlling for some observables on the side of the head. When we control for age, we group heads into 7 age groups and in the pool of spouses for each

age group are all spouses of heads in the actual data. Finally, we do the random coupling by age and education groups. As expected, the cyclicality experienced by random couples is more and more similar to actual households. Still, for the US and Germany we find actual households experiencing slightly higher cyclicality of earnings changes than their artificial counterparts. This suggests that the correlation between head's and spouse's labor market income is higher than for a random counterpart and uncontrolled characteristics play some role - like, e.g., most heads and spouses working in the same local labor markets.

We conclude that the responses of gross household earnings are heterogeneous across countries, with Sweden being the only economy where the family plays a clear insurance role against aggregate fluctuations. However, it is hard to extract further conclusions in disconnection to taxes and transfers payed and received by the household. In order to shed light on this issue, we move on to considering the role of social insurance policy over the business cycle.

5.2 Government 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 on 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 income measures that in turn add certain transfers to household gross earnings.

For the analysis of subcomponents, 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) "pensions," and are listed in Table 12. 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

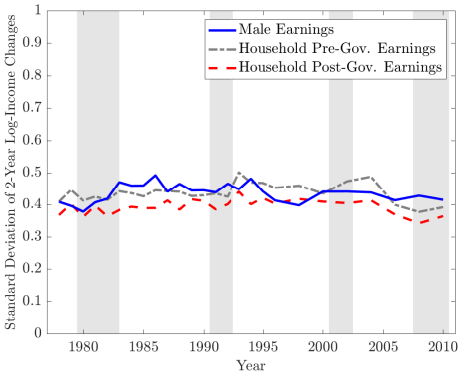
Table 11: Cyclicality of Earnings for Random Couples

	Std.Dev	L9010	Skew	Kelly	L9050	L5010
			United States (PSID)			
Actual Households	0.40**	0.23	4.67***	1.97***	0.93***	–0.71***
Random couples	0.39	0.21	8.12	1.27	0.63	-0.42
	(0.21)	(0.22)	(6.90)	(0.31)	(0.15)	(0.19)
	[0.09 - 0.67]	[–0.07 - 0.49]	[0.29 - 17.13]	[0.84 - 1.69]	[0.42 - 0.82]	[–0.66 - –0.19]
Random by age	0.37	0.03	7.57	1.51	0.65	–0.62
	(0.27)	(0.30)	(8.15)	(0.35)	(0.18)	(0.24)
	[0.00 - 0.71]	[–0.35 - 0.45]	[–2.71 - 17.57]	[1.05 - 1.94]	[0.42 - 0.89]	[–0.92 - –0.33]
Random by age & educ.	0.43	0.01	6.19	1.49	0.63	-0.62
	(0.26)	(0.29)	(7.25)	(0.34)	(0.18)	(0.23)
	[0.12 - 0.74]	[–0.40 - 0.34]	[–3.44 - 15.36]	[1.06 - 1.92]	[0.37 - 0.87]	[–0.92 - –0.32]
			Sweden (LINDA)			
Actual Households	0.10	–0.02	8.92***	2.24***	0.50***	–0.52*
Random couples	0.019	–0.21	7.09	1.72	0.31	–0.52
	(0.02)	(0.03)	(1.02)	(0.05)	(0.02)	(0.02)
	[–0.01 - 0.04]	[–0.25 - –0.16]	[5.94 - 8.57]	[1.66 - 1.79]	[0.28 - 0.33]	[–0.55 - –0.49]
Random by age	0.02	–0.20	7.41	1.76	0.32	–0.,53
	(0.03)	(0.03)	(1.19)	(0.06)	(0.02)	(0.02)
	[–0.01 - 0.05]	[–0.25 - –0.16]	[6.01 - 8.83]	[1.68 - 1.85]	[0.30 - 0.35]	[–0.56 - –0.50]
Random by age & educ.	0.07	0.02	8.10	1.82	0.46	–0.43
	(0.02)	(0.03)	(1.07)	(0.06)	(0.02)	(0.02)
	[0.04 - 0.10]	[–0.02 - 0.06]	[6.69 - 9.58]	[1.72 - 1.89]	[0.43 - 0.48]	[–0.47 - –0.40]
			Germany (SOEP)			
Actual Households	–0.12	–1.31***	1.00	1.88**	–0.05	–1.26***
Random couples	–0.21	–0.99	2.82	1.28	–0.12	–0.87
	(0.20)	(0.29)	(6.00)	(0.48)	(0.18)	(0.22)
	[–0.49 - 0.03]	[–1.35 - –0.62]	[–4.79 - 11.12]	[0.64 - 1.88]	[–0.35 - 0.11]	[–1.14 - –0.58]
Random by age	–0.35	–1.15	1.26	1.02	–0.25	–0.89
	(0.19)	(0.32)	(6.15)	(0.57)	(0.21)	(0.25)
	[–0.60 - –0.12]	[–1.57 - –0.73]	[–6.82 - 8.85]	[0.29 - 1.78]	[–0.53 - 0.01]	[–1.23 - –0.58]
Random by age & educ.	–0.35	–1.19	0.32	1.01	–0.28	–0.91
	(0.17)	(0.33)	(5.96)	(0.56)	(0.21)	(0.25)
	[–0.55 - –0.14]	[–1.65 - –0.80]	[–7.60 - 7.77]	[0.25 - 1.70]	[–0.54 - –0.01]	[–1.24 - –0.60]

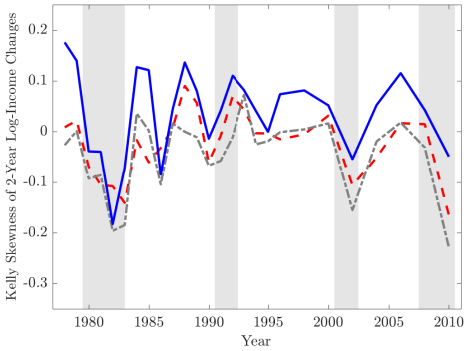
Note: See notes for Table 2. The parameter for the random couples is the mean over 250 bootstrap repetitions. In parentheses is the standard deviation, in brackets are the 10th and 90th percentiles. The regression for Sweden with education starts in 1991.

Figure 4: Standard Deviation (Left) and Skewness (Right) of Short-Run Earnings Growth: United States, Germany (SOEP), and Sweden

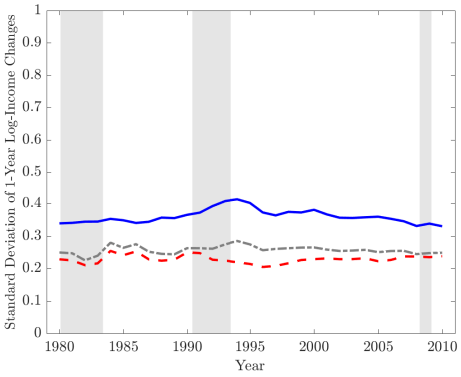
(a) United States, Std. Dev.



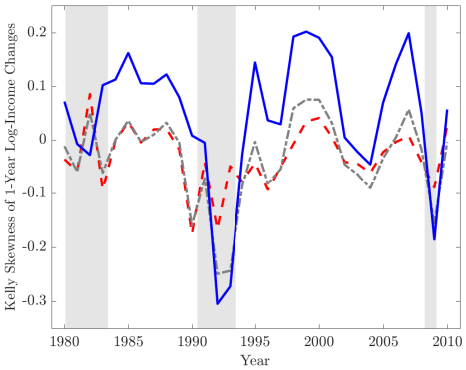
(b) United States, Kelly's Skewness



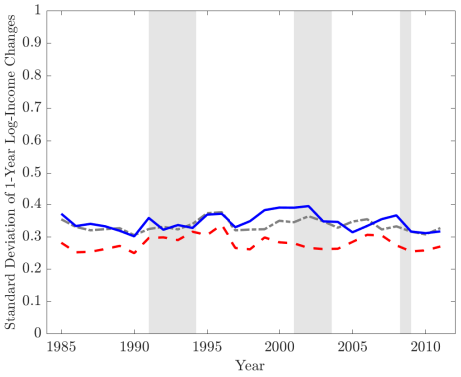
(c) Sweden, Std. Dev.



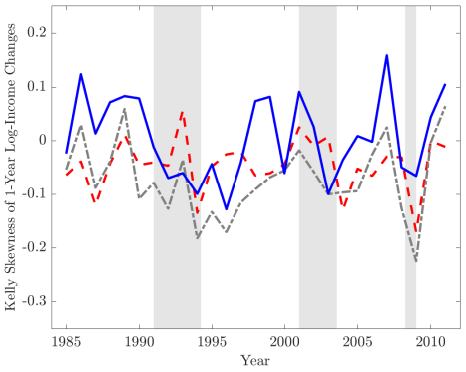
(d) Sweden, Kelly's Skewness



(e) Germany, Std. Dev.



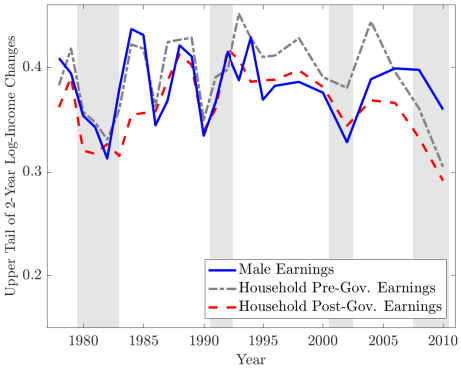
(f) Germany, Kelly's Skewness



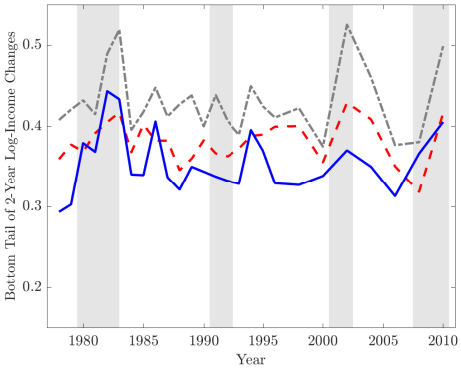
Note: Linear trend removed, centered at sample average.

Figure 5: Tails of Short-Run Earnings Growth: United States, Germany (SOEP), and Sweden

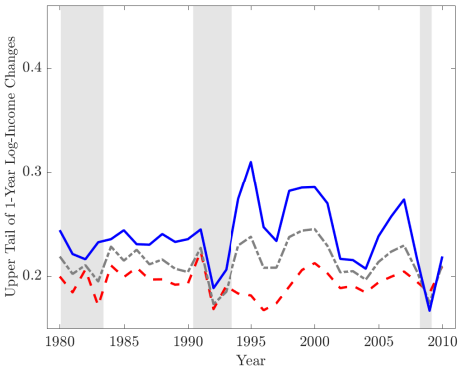
(a) United States, P90-50



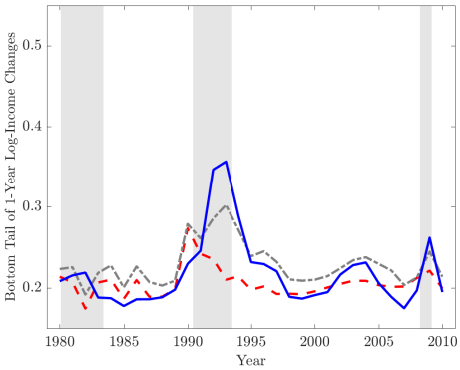
(b) United States, P50-10



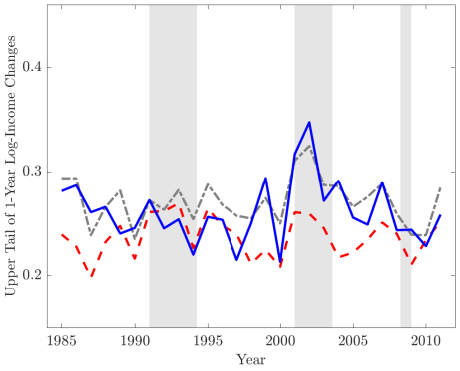
(c) Sweden, P90-50



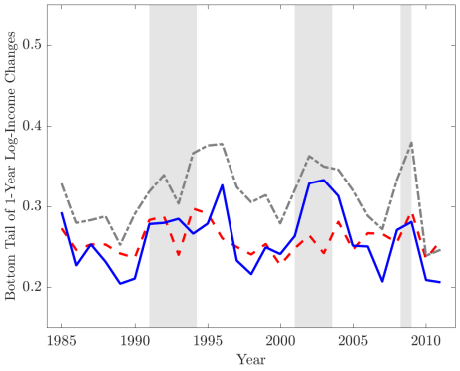
(d) Sweden, P50-10



(e) Germany, P90-50



(f) Germany, P50-10



Note: Linear trend removed, centered at sample average.

Table 12: Components of Social Policy

	LINDA	SOEP	PSID
1. Labor Market Transfers:	Unemployment benefits; Labor market programs	Unemployment benefits	Unemployment benefits; Workers' compensation
2. Aid to Low-Income Families:	Family support; Housing support; Cash transfers from the public; (no private transfers)	Subsistence allowance; Unemployment assistance (up to 2004); Unemployment benefits II (since 2005)	Supplemental Security Income; Aid to Families with Dependent Children (AFDC); Food Stamps; Other Welfare
3. Social Security and Pensions:	(Old Age) Pensions	Combined old-age, disability, civil service, and company pensions	Combined (Old Age) Social Security and Disability (OASI)

Note: Table lists the measures used in the three data sets to construct subcomponents of transfers.

relevance of this type of transfer can therefore be expected to matter most for low-income households who have a higher likelihood of falling down to fulfilling 'at-risk' criteria in the course of a recession. The third component, pension payments, is not directly connected to business cycle considerations. It can still play a relevant role for household members near or at retirement age, who may take up pension payments instead of unemployment payments if they decide to leave the labor market upon job loss.

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 10. Again, Figures 4 and 5 visualize the findings. We find that social policy is an important source of insurance against aggregate fluctuations in all three economies, with very similar overall effects. Motivated by the considerations from 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, lower tail counter-cyclicality is dampened but 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 U.S. and Sweden reveal procyclicality of 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 translates into procyclical overall dispersion of post-government earnings changes in the U.S., and countercyclical dispersion in Sweden. Summarizing the reaction of overall dispersion and tails results in procyclicality of Kelly's skewness measure for both countries. This analysis reveals the importance of considering the tails separately.

To sum up, the analysis suggests that downside risk in recessions is mitigated by taxes and transfers. In Sweden, an additional effect are lowered upside chances in expansions. This lines up with considerations of Sweden as a country with a high degree of redistribution.

The Role of Subcomponents of Social Policy

The measure of post-government earnings used so far lumps a lot of very different transfers received and taxes paid by households. While this measure is appropriate for assessing the overall effect of the tax and transfers system, it is not as well suited for understanding the success of different social policies that specifically aim at mitigating downside risk or that aims at aiding low-income families, who can be expected to be especially vulnerable in recessionary periods. Therefore, we now consider different types of transfers separately. The results of the cyclicality analysis are listed in Table 13. As for for the estimates of total taxes and transfers, we compare the coefficients to the ones from the household gross earnings analysis in row 1 of Table 10. 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 results are not driven by low-income households entering the sample for a certain type of transfer but are not in the sample when considering another.

The results in Table 13 show that labor market related transfers (which have unemployment benefits as the main component) are successful in mitigating downside risk in recessions in all three economies. This supports the consideration of an increased incidence of unemployment in recessions. In the United States the cyclicality

Table 13: Cyclicality of Household Earnings - transfers added separately

	Std.Dev	L9010	Skew	Kelly	L9050	L5010
			United States (PSID)			
+ Labor transfers	0.69**	0.60	–0.88	1.59***	0.92***	–0.33
	(2.29)	(1.54)	(–0.24)	(5.20)	(4.20)	(–1.34)
+ Aid to low-income	0.37**	0.21	5.25***	1.90***	0.89***	–0.69***
	(2.14)	(0.77)	(3.92)	(6.13)	(5.16)	(–3.33)
+ Pensions	0.33*	0.22	4.86**	1.82***	0.86***	–0.64***
	(1.97)	(0.80)	(2.74)	(5.61)	(4.79)	(–3.06)
			Sweden (LINDA)			
+ Labor transfers	–0.05	–0.22	4.26***	1.14***	0.13*	–0.35**
	(–0.66)	(–1.23)	(2.99)	(4.23)	(2.04)	(–2.58)
+ Aid to low-income	0.02	–0.07	7.32***	2.11***	0.42***	–0.49**
	(0.29)	(–0.38)	(4.79)	(3.72)	(4.51)	(–2.47)
+ Pensions	0.03	–0.07	11.2***	2.34***	0.48***	–0.55**
	(0.71)	(–0.43)	(4.61)	(3.55)	(4.50)	(–2.68)
			Germany (SOEP)			
+ Labor transfers	–0.07	–1.09***	–1.76	1.34**	–0.13	–0.96***
	(–0.40)	(–2.96)	(–0.73)	(2.50)	(–0.60)	(–3.65)
+ Aid to low-income	–0.15	–1.32***	2.54	1.66**	–0.11	–1.21***
	(–0.83)	(–3.82)	(1.05)	(2.40)	(–0.47)	(–4.08)
+ Pensions	–0.08	–1.21***	–0.96	1.80***	–0.04	–1.17***
	(-0.40)	(–3.30)	(–0.44)	(3.10)	(–0.18)	(–4.58)

Note: See notes for Table 2.

estimate of the lower tail of income changes is no longer statistically significant when considering these transfers (with the point estimate halved compared to household gross earnings). In the two European countries, the point estimates are cut by about a third. Considering the upper tail, in both the United States and Germany, there is no change when moving from gross earnings to earnings plus labor market transfers; in Sweden, L9050 widens less in expansions. The other to components of transfers do not have any impact on cyclicality as measured by our cyclicality regressions. For all three economies, the point estimates are almost identical to the ones for gross earnings.

In order to get a feeling for the importance of labor market related transfers relative to measures that are added when moving to post-government earnings, we compare the coefficients to the ones from the household post-government earnings analysis in row 2 of Table 10. It turns out that the results differ by country. In the case of Sweden, labor market transfers basically go the whole way from pre- to post-government earnings

cyclicality in the lower tail. Similarly, in the U.S. a big chunk of the way from pro-cyclical L5010 of gross earnings changes to acyclical L5010 of post-government earnings changes is accounted for by labor market transfers (the point estimate is statistically insignificant and thus supports the interpretation of acyclical income changes after considering labor market transfers).[20] For Germany, other components not captured separately matter a lot on top of labor market transfers for the mitigation of downside risk.

Moving to the upper tail, a big chunk on the way from procyclicality in gross earnings changes to acyclicality of L9050 of post-government earnings changes in the case of Sweden is already taken out from labor market transfers: the point estimate is only 0.13 (and significant at the 10% level only) as compared to a highly (statistically) significant estimate of 0.5. In the case of the U.S., the upper tail is unaffected by this component of social policy. Hence, the way to lower procyclicality of L9050 of post-government earnings changes is accounted for by other elements of the tax and transfer system.

While SIAB data set includes information on individuals and not on households, we do have information on unemployment benefits at the individual level. Table 14 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). As expected, including unemployment benefits reduces downside risk in recessions quite significantly for male workers. Workers who have a period of non-employment and hence zero earnings for a fraction of the year do not fall down to actually having zero income but receive unemployment benefits during that period. For females, the insurance effect of unemployment benefits goes in the expected direction as measured by reduced downside risk in recessions and reduced cyclicality of Kelly's skewness measure. Hence, these individual level results line up well with the household level analysis conducted using SOEP data.

5.3 Sensitivity of results to choice of lag length

All results reported in the text refer to the distribution of what we label transitory, i.e., one-year changes of several income measures.[21] Given the focus of Storesletten *et al.*

[20] Again, this might partly be due to modest sample size.

[21] Recall that for the US we define two-year changes as transitory in order to account for the biannual nature of the PSID since 1997.

Table 14: Cyclicality of Individual Earnings including unemployment benefits in Germany (SIAB)

	Std.Dev.	L9010	Skew	Kelly	L9050	L5010
Male Earnings	0.07	0.11	14.56***	5.71***	0.97***	–0.86***
	(0.44)	(0.26)	(3.92)	(5.32)	(2.93)	(–4.40)
+Unempl. benefits	0.28	0.15	8.63**	5.12***	0.84**	–0.70***
	(1.19)	(0.34)	(2.28)	(5.24)	(2.61)	(–4.01)
Female Earnings	0.14	0.46	4.45	2.69*	0.89	–0.44*
	(0.66)	(0.60)	(1.63)	(1.92)	(1.26)	(–1.74)
+ Unempl. benefits	0.20	0.50	2.66	2.43*	0.82	–0.32
	(0.82)	(0.67)	(1.06)	(1.82)	(1.22)	(–1.43)

Note: See notes for Table 2. Difference to estimates in 2 and 3 are due to the fact that regressions start in 1981 instead of 1976.

(2004) or Guvenen *et al.* (2014), to which we relate our results, on persistent income changes this choice needs to be discussed. The main reason for us to focus on one-year changes is that we choose a regression framework as our main tool of analysis. We make this choice, because we compare the cyclicality of income risk across countries. While for the US it is widely accepted to base the dating of business cycles on NBER recession dates, this dating is less clear cut for both Germany and Sweden. More generally, it is not clear that in a cross-country comparison the dating of business cycles is of the same quality in terms of capturing actual economic conditions. Our regression framework allows a very clear interpretation and comparison of cyclicality of income changes.

Moving to five-year changes–which are closer to capturing persistent changes–would imply problems with the regression analysis for two reasons. One option would be to use non-overlapping five-year changes of income and gdp, another would be to use overlapping changes. The first option would give too few data points for a regression analysis, while the second would open the door to usual problems of overlapping data.

The time-series of five-year changes is shown in figures A.1 to A.3 in appendix C. Comparison to the one-year changes suggest the same qualitative patterns.

6 Welfare Analysis

In this section, we quantitatively address the question of how successful government policy is in insuring households against business cycle fluctuations of earnings risk. 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 to allow for cyclical higher-order risk. Given estimated processes, we quantify the welfare gain of the existing tax and transfer system for the average household as the consumption equivalent variation (CEV) that is necessary to make households facing the pre-government income stream indifferent to facing the post-government income stream.

Our estimation is based on earnings data, so we use a model to simulate consumption profiles of households facing either pre- or post- government income streams. It is well documented that households are able to partially insure against income changes via various mechanisms (see, e.g., Blundell *et al.* (2008)). Reflecting this, we use a variant of the partial insurance model by Heathcote *et al.* (2014) to quantify the welfare gains. We do this exercise for all three countries. Before going to the results of the model-based welfare analysis in section 6.2, the next subsection discusses the estimation of the income process.

6.1 Estimation of Pre- and Post-Government Earnings

Details of the Parametric Specification

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:

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

where ε_t is an *iid* transitory shock with distribution $\mathcal{N}(\mu_\theta, \sigma_\theta)$,[22] and ζ_t denotes a permanent shock with time-varying and business-cycle dependent distribution. In par-

[22] μ_θ is chosen such that $\mathbb{E}\left(e^\theta\right) = 1$.

ticular, ζ_t follows a mixture of 3 normals:

$$\zeta_t \sim \begin{cases} \mathcal{N}(\mu_{1t,}, \sigma_1) & \text{with probability } p_1 \\ \mathcal{N}(\mu_{2t}, \sigma_2) & \text{with probability } p_2 \\ \mathcal{N}(\mu_{3t}, \sigma_3) & \text{with probability } p_3 \end{cases} \tag{4}$$

where $\sum_{j=1}^{3} p_j = 1$.

Notice that business-cycle risk is transmitted to household earnings only through the permanent shock. Blundell *et al.* (2016), among others, concluded that transitory shocks are insurable almost entirely. We thus concentrate in ζ_t as our main idiosyncratic component exposed to aggregate risk. We follow McKay (2014) and incorporate the dynamics of aggregate risk into household earnings, as opposed to imposing discrete events for recessions and expansions. Namely, the time-varying means are specified as:

$$\begin{aligned} \mu_{1t} &= \bar{\mu}_t \\ \mu_{2t} &= \bar{\mu}_t + \mu_2 - \phi x_t \\ \mu_{3t} &= \bar{\mu}_t + \mu_3 - \phi x_t \end{aligned}$$

where $\bar{\mu}_t$ is a normalization such that $E(e^{\zeta_t}) = 1$ for all t, and $x_t \equiv -(\log \frac{GDP_{t+1}}{GDP_t})$ captures the business cycle.[23] 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 parameter ϕ is will be estimated alongside the other parameter and will determine how aggregate risk is transmitted to household earnings risk.

Using 3, we can write earning growth as

$$y_{t+\tau} - y_t = \theta_{t+\tau} - \theta_t + \sum_{j=1}^{\tau} \zeta_{t+j}, \tag{5}$$

for different horizons τ.

We further impose $p_2 = p_3, \sigma_2 = \sigma_3$, as in McKay (2014). We are thus left with 7 parameters to estimate:

$$\theta = (\sigma_\theta, p_1, \mu_2, \mu_3, \sigma_1, \sigma_2, \phi)$$

[23] To be precise: in the estimation we standardize log gdp changes.

Estimation

We estimate θ by Method of Simulated Moments (MSM). We target the (i) p10, (ii) p50, and (iii) p90 of the 1, 3, and 5-year earnings changes distribution, as well as the age profile of the cross-sectional variance from 25 to 60. In the case of the United States, we target the 2, 4, and 6-year changes. We do this to be able to use the whole sample of the PSID. For ease of notation, we will describe out methodology assuming the changes horizons as in the case of Sweden and Germany, but the change to the US is trivial. To construct the simulated income profiles, we use equation 5. In order to accomodate the mean shift of the distribution of earnings changes over time, we shift the simulated earnings changes such that the mean of 3 year earnings changes fits the data.

The total number of moments N depends on the available years, as well as the frequency. Therefore, this number will vary by country. We have $N = 261$, $N = 2297$, and $N = 252$ moments in the case of Germany, Sweden, and the United States, respectively. We simulate these profiles $R = 10$ times, and minimize the average scaled distance between the simulated and empirical moments using a global version of the Nelder-Mead algorithm with several restarts, as described in Guvenen (2011). More details on the simulation and estimation are included in Appendix D.

Parameter Estimates

Tables 16a and 16b show the parameter estimates for the three economies.

Table 15: Estimated parameter values

Parameter	Description	Value		
		US	Sweden	Germany
p_1	Weight of center of ζ distribution	0..9252	0..8944	0..8162
p_2	Weight of right tail of ζ distribution	0..0374	0..0528	0..0919
p_3	Weight of left tail of ζ distribution	0..0374	0..0528	0..0919
σ_ε	St. dev. of transitory income shock	0..2167	0..1128	0..1522
$\sigma_{1,\zeta}$	St. dev. of center of ζ distribution	0..0870	0..0061	0..0053
$\sigma_{2,\zeta}$	St. dev. of right tail of ζ distribution	0..0161	0..0138	0..0138
$\sigma_{3,\zeta}$	St. dev. of left tail of ζ distribution	0..0161	0..0138	0..0138
μ_2	Mean of right tail of ζ distribution	0..2428	0..2183	0..0726
μ_3	Mean of left tail of ζ distribution	-0..3471	-0..4129	-0..3665
ϕ	Aggregate risk transmission parameter	0..2270	0..0135	0..0875

(a) Pre-Government Earnings

Parameter	Description	Value		
		US	Sweden	Germany
p_1	Weight of center of ζ distribution	0..9252	0..9106	0..9288
p_2	Weight of right tail of ζ distribution	0..0374	0..0447	0..0356
p_3	Weight of left tail of ζ distribution	0..0374	0..0447	0..0356
σ_ε	St. dev. of transitory income shock	0..1985	0..0916	0..1384
$\sigma_{1,\zeta}$	St. dev. of center of ζ distribution	0..0624	0..0065	0..0557
$\sigma_{2,\zeta}$	St. dev. of right tail of ζ distribution	0..0185	0..0123	0..0117
$\sigma_{3,\zeta}$	St. dev. of left tail of ζ distribution	0..0185	0..0123	0..0117
μ_2	Mean of right tail of ζ distribution	0..0462	0..0212	0..0101
μ_3	Mean of left tail of ζ distribution	-0..4355	-0..3338	-0..1346
ϕ	Aggregate risk transmission parameter	0..0122	0..0117	0..2989

(b) Post-Government Earnings.

6.2 Quantitative Model

The earnings changes we estimated reflect both actual "shocks" in the sense of unforeseen income changes, as well as expected changes (cf. Guvenen and Smith, 2014). Further, shocks differ in their nature, and some shocks are insurable, while others are not. Heathcote *et al.* (2014, HSV hereafter) set up a model populated by a continuum of islands, each of which in turn is 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 individual specific. Island refers to a group of workers 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, e.g., 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 model estimation, there is no need to define empirical counterparts to the model islands, and all parameters can be identified solely with aggregate moments. Given some assumptions on the market structure, outlined below, HSV 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.

Model structure

We employ a version of the HSV 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 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.

Specifically, household income is assumed to follow

$$\log y_t = \alpha_t + \varepsilon_t,$$

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

$$\begin{aligned} \alpha_t &= \alpha_{t-1} + \omega_t \\ \varepsilon_t &= \kappa_t + \theta_t \\ \kappa_t &= \kappa_{t-1} + \eta_t \\ \theta_t &\sim F_{\theta,t} \\ \eta_t &\sim F_{\eta,t} \\ \omega_t &\sim F_{\omega,t} \end{aligned} \tag{6}$$

Workers live finite lives. Each period a mass $(1-\delta)$ of new-borns enters the economy with age 0. The probability of survival from age a to age $a+1$ is constant at δ. New-born agents maximize discounted life-time utility. The felicity utility functions is CRRA and given by

$$u\left(c_t\right) = \frac{c_t^{1-\gamma} - 1}{1-\gamma}$$

Age 0 agents entering in year b hold zero financial wealth and are allocated to an island of agents that then share the same sequence of uninsurable shocks $\{\omega_t\}_{t=b}^{\infty}$. There exists a full set of state-contingent claims for individual-specific shocks. Across islands, individuals can trade contracts contingent on their individual level shocks, while inter-island contracts are non-existent.

In equilibrium, log consumption and consumption change is given by:[24]

$$\log c_t\left(\mathbf{x}_t, \varepsilon_t\right) = \alpha_t + \log \int \exp\left(\varepsilon_t\right) dF_{\varepsilon,t}^{a}, \tag{7}$$

$$\Delta \log c_{t+1} = \omega_{t+1} + \left(\log \frac{\int \exp\left(\eta_{t+1}\right) dF_{\eta,t+1} \int \exp\left(\theta_{t+1}\right) dF_{\theta,t+1}}{\int \exp\left(\theta_t\right) dF_{\theta,t}} \right) \tag{8}$$

Note that the uninsurable shock ω_{t+1} that realizes for an individual translates one for one to consumption. The individual realizations of the two insurable shocks, however, do not at all 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, we need to translate the estimated earnings process (3) into the process specified in (6). In the estimation the transitory component is assumed to be drawn from a Gaussian distribution, which directly translates into the distribution θ_t from (6). Permanent earnings changes are drawn from a mixture of three normal distributions. In specification (6), the overall permanent earnings change is given by $(\omega_t + \eta_t)$, the insurable (individual-level) and the uninsurable (island-level) part. We now 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 the fraction λ (1-λ) of the overall variance of the

[24] the derivation of consumption outlined in HSV carries over 1:1 to our simplified version of their model, which abstracts from the tax function and endogenous labor supply

permanent shock ζ:

$$\eta_t \sim \begin{cases} \mathcal{N}(\lambda^{1/2}\mu_{1,t}, \lambda\sigma_1^2) & \text{with probability } p_1 \\ \mathcal{N}(\lambda^{1/2}\mu_{2,t}, \lambda\sigma_2^2) & \text{with probability } p_2 \\ \mathcal{N}(\lambda^{1/2}\mu_{3,t}, \lambda\sigma_3^2) & \text{with probability } p_3 \end{cases} \tag{9}$$

$$\omega_t \sim \begin{cases} \mathcal{N}((1-\lambda)^{1/2}\mu_{1,t}, (1-\lambda)\sigma_1^2) & \text{with probability } p_1 \\ \mathcal{N}((1-\lambda)^{1/2}\mu_{2,t}, (1-\lambda)_2^2) & \text{with probability } p_2 \\ \mathcal{N}((1-\lambda)^{1/2}\mu_{3,t}, (1-\lambda)\sigma_3^2) & \text{with probability } p_3 \end{cases} \tag{10}$$

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 E.[25]

In this setup, λ is the measure of the degree of partial insurance against permanent shocks: it simply measures the share of the total permanent variance that is accounted for by the insurable component. In line with the estimation we draw changes to x_t from a Standard Normal distribution and for each t calculate the parameters of the transitory and permanent shock components. In order to break the permanent shocks into an insurable and an uninsurable component we exogenously pick a value for λ and then simulate income and consumption profiles for a number of model agents.

Household i's lifetime utility when facing the pre- or post-government income streams is given by:

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

for $x = pre, post$, and $c_{i,a}^x$ is consumption of household i at age a when facing income stream x.

Utilitarian welfare in the world of the pre-government income stream when agents receive per-period consumption equivalent variation CEV is given by $W^{pre}(CEV) \equiv \sum_i U_i^{pre}\left(\left\{(1+CEV)\, c_{i,a}^{pre}\right\}_a\right)$, similarly utilitarian welfare in the post-government world is $W^{post} \equiv \sum_i U_i^{p\acute{o}st}\left(\left\{c_{i,a}^{post}\right\}_a\right)$. We search for the CEV that ensures $W^{pre}(CEV) = W^{post\,\prime}$, which due to homogeneity of the per-period utility function can be calculated in closed-form as

$$CEV = \left(\frac{W^{post}}{W^{pre}}\right)^{1/(1-\gamma)} - 1. \tag{11}$$

[25] We then shift the η and ω shocks by d_η and $d_\omega$such that $E[exp(\eta + d_\eta)] = E[exp(\omega + d_\omega)] = 1$.

We evaluate the welfare gains for several values of the degree of partial insuranceλ and coefficients of relative risk aversion γ. For the remaining model parameters we follow HSV, i.e., we set the annual survival probability to $\delta = 0.996$, the preference discount factor to $\beta = 0.95$. We simulate the model economy for 10 cohorts, each consisting of 10,000 individuals that live for a maximum of 100 years and calculate their expected present value utility at model age 1. Each cohort lives through a different macroeconomic history of log gdp changes, which we draw independently from a Normal distribution. The initial distribution of earnings at model age 1 is drawn from a log-normal distribution that resembles the overall dispersion at age 25 from the data. To be precise, we assume that earnings at age 25 are composed by the sum of a transitory component and two permanent components. We draw the transitory component from a log-normal distribution with variance of transitory shocks and then draw the uninsurable and insurable parts of the permanent shocks from log-normal distributions with λ and $(1-\lambda)$ times the remaining variance. Thus, total variance at model age 1 corresponds to data variance at age 25. Then, agents accumulate shocks as described above. Using (11), we calculate the CEV for each cohort and then take the average over cohorts to calculate the welfare gains which are shown in table 16.

The degree of partial insurance λ is set to be the same in the two alternative worlds, i.e., pre- or post-government income. This reflects that we explicitly do not want to capture any partial insurance by the tax and transfer system. As expected, a higher value of partial insurance implies a lower gain from introducing the tax and transfer system: households are insured better against shocks and so the individual realization of earnings changes is less relevant for actual consumption. Consider the case of log preferences, i.e., $\gamma = 1$. If the channels of partial insurance available to households beyond taxes and transfers reduce earnings volatility by 50 percent, the introduction of the tax and transfer system amounts to a welfare gain of about 6% of annual consumption for the average American household facing the pre-government income stream. These numbers are 7% and 4% for Sweden and Germany, respectively. Varying the degree of relative risk aversion, again the results reflect the intuition that the gains of moving to a world with less cyclical earnings are higher when the aversion against cyclicality is higher.

Summarizing, in all three economies, the welfare gains resulting from the tax and transfer system are large and of similar magnitude. The Swedish tax and transfer system is estimated to yield a higher welfare gain than its German and American counterparts. However, comparing the US to Germany, it is surprising that our results

point towards slightly higher gains in the United States.

Table 16: Welfare Gains of the Tax and Transfer System

	$\lambda = 0.75$	$\lambda = 0.5$	$\lambda = 0.25$
		$\gamma = 1$	
United States	3.2	6.4	9.7
Sweden	3.3	6.7	10.1
Germany	1.9	3.7	5.6
		$\gamma = 2$	
United States	6.4	13.4	21.5
Sweden	6.9	14.6	23.1
Germany	4.1	8.6	13.6

Note: λ denotes the degree of partial insurance against permanent shocks. γ denotes the degree of relative risk aversion. The table entries show the consumption equivalent variation in percentages that is necessary to make households indifferent between the (compensated) pre-government income stream and the post-government income stream.

The described welfare analysis is based on an imposed structure with strong assumptions. We complement these results with a more reduced form estimate of the welfare gains. Consider Table 17: it shows moments of the distribution of income changes coming from pooled samples for each country for the income measures considered throughout the analysis in the preceding sections. Let us focus on the moments of pre- and post-government household income. In all three countries, the tax and transfer system overall closes the distribution of earnings changes: both upper and lower tails are smaller. Also a finding common to all economies is that the lower tail is affected more than the upper tail, which is reflected by a less negatively skewed distribution. For example, in Sweden the distribution of post-government earnings changes is even symmetric on average. For the United States, the described pattern is very weak: Kelly's skewness is only slightly affected and the measure based on central moments is unchanged to the second digit.

Not only is the overall dispersion smaller, and the distribution more symmetric, when comparing post-government household income to household gross earnings, also the share of households for which the annual income movements are very small is higher. This is what the higher kurtosis of post-government income suggests. Again, the effect is weakest for the United States and most pronounced for Sweden. Note that, in all three countries, the distribution of both pre- and post-government income changes is far away from a log-normal distribution (which has a kurtosis of 3).

Table 17: Moments of $\Delta_s y$

	Std.Dev	L9010	Skew	Kelly	Kurt	L9050	L5010	Nobs
	United States (PSID)							
Male Earnings	0.44	0.75	–0.42	0.03	13.69	0.39	0.36	42,698
HH Earnings	0.45	0.83	–0.24	–0.04	10.86	0.40	0.43	38,314
HH Post Gov	0.41	0.76	–0.24	–0.03	12.26	0.37	0.39	38,602
HH Disposable	0.39	0.79	–0.19	–0.02	10.50	0.39	0.40	38,632
+ Labor transfers	0.44	0.81	–0.24	–0.04	11.33	0.39	0.42	38,327
+ Aid to low-income	0.45	0.83	–0.25	–0.04	10.98	0.40	0.43	38,326
+ Pensions	0.45	0.83	–0.22	–0.04	10.85	0.40	0.43	38,354
	Sweden (LINDA)							
Male Earnings	0.36	0.45	–0.27	0.04	13.57	0.23	0.22	1,907,421
HH Earnings	0.26	0.45	–0.47	–0.04	14.22	0.21	0.23	1,113,760
HH Post Gov	0.21	0.35	–0.04	0.00	23.03	0.17	0.17	1,113,759
HH Disposable	0.20	0.35	0.03	0.01	17.64	0.18	0.17	1,113,759
+ Labor transfers	0.23	0.40	–0.44	–0.04	16.37	0.19	0.21	1,077,255
+ Aid to low-income	0.28	0.45	–0.34	–0.04	17.54	0.22	0.24	1,077,255
+ Pensions	0.25	0.44	–0.21	–0.01	15.00	0.21	0.22	1,077,255
	Germany (SOEP)							
Male Earnings	0.35	0.52	–0.23	0.00	15.89	0.26	0.26	64,572
HH Earnings	0.33	0.58	–0.78	–0.08	13.32	0.27	0.31	59,161
HH Post Gov	0.28	0.50	–0.11	–0.04	16.09	0.24	0.26	58,725
HH Disposable	0.27	0.50	–0.11	–0.03	14.99	0.24	0.26	58,853
+ Labor transfers	0.32	0.57	–0.70	–0.07	13.93	0.26	0.30	59,173
+ Aid to low-income	0.33	0.58	–0.83	–0.08	13.29	0.27	0.31	59,199
+ Pensions	0.32	0.57	–0.66	–0.06	13.55	0.27	0.30	59,166

Note: For Germany and Sweden, all moments refer to 1-year income differences, i.e. $s = 1$. For the United States, the reference sample is 1976–2010, with $s = 2$. Moments for $s = 1$ are reported in the appendix for the sample 1969-1996.

7 Conclusion

This paper has studied how higher-order income risk varies over the business cycle, as well as the extent to which such risks can be smoothed within households or with government social insurance policies. To provide a broad perspective on these questions, we studied panel data on individuals and households from the United States, Germany, and Sweden, covering more than three decades of data for each country. We find that the underlying variation in higher-order risk is remarkably similar across these countries that differ in many details of their labor markets. In particular, in all three countries, the variance of earnings shocks is almost entirely constant over the

business cycle, whereas the skewness of these shocks becomes much more negative in recessions. Government provided insurance, in the form of unemployment insurance, welfare benefits, aid to low income households, and the like, plays a more important role reducing downside risk in all three countries; the effectiveness is weakest in the United States, and most pronounced in Germany. For Sweden we find that insurance provided within households plays a similar role. We calculate that the welfare benefits of social insurance policies for stabilizing higher-order income risk over the business cycle range from 1% of annual consumption for the United States to 4.5% for Sweden.

References

Blundell, R., Pistaferri, L. and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, **98** (5), 1887–1921.

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

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.

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**.

Guvenen, F. (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.

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

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

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 unite 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, she is 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: *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[26].

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[27], 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, it 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.

[26] See Shin and Solon (2011) for a comparison of PSID male earnings variables in inequality analyses.

[27] In particular, total labor earnings included the labor parts of farm and business income up to the 1993 survey but not in subsequent waves.

Public Transfers. Transfers are considered at the family unit level, when possible. We group social and welfare programs in 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.

Household Disposable Income. We construct this variable from Household Taxable Income (Head's and wife's income from assets, earnings, and net profit from self-employment) *minus* taxes *plus* public transfers.

Transfers

We refer to Table 12 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 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[28] and 2005–2011, the total annual amount of each program is reported for the head, spouse and others in the family unit. In occasions, the amount appears combined for several or all members.[29] Because in those cases it is impossible to identify separate recipiency of each member, we consider the benefit amount of the whole family. This 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 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.

[28] Our main sample refers to survey years 1977-2011, but complementary results are provided for the annual subsample of the PSID. This 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.

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

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 I. 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 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.

A.2 LINDA

Variables

Demographic and Socioeconomic

Head and Relationship to Head. LINDA is compiled from the Income Register based on filed tax reports and other registers. Statistics Sweden samples individuals

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

	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.

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 eduction from 1991 and onwards. An individual is assigned "college" education if it has at least 3 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 onwards. For the years 1991-92 the public sector employment is defined as those we use SNI90 codes 72000-72003, 90000-93999 and >=96000. For 1993-2006 we use SNI92 codes 64110-64202, 73000-74110, 75000-92000, 92500-92530 and >=96000. For 2007 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 measure for the years 1979-1984.

Individual labor income. 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.

Household Disposable Income. Disposable income consists of the sum of factor income and minus taxes and plus public transfers.

Taxes. LINDA provides observations of total taxes paid by the individual. Since taxed 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 subcategory 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 aged 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 is a 2% random sample of all individuals covered by a dataset 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. A new spell starts whenever there is a new notification, which happens when either a new employment relationship changes, an ongoing contract is changed, or with the start of a calendar year. BeH covers all workers subject to social security contributions, which excludes civil servants, self-employed 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 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 income. We calculate annual earnings as the sum of total earning 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 to obtain a time consistent measure. For the same reason we drop spells with reported average daily wage rate below the highest marginal employment threshold in the sample period, which is 14.15 Euros (in 2003 Euros). There are two drawbacks in the available data: 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 1982-1983 on the quintiles in 1983 and the same for 1983-1984. 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 1983-1984 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 starting point for the Pareto 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, e.g., from two full-time spells in a year only one is top-coded. We therefore define as 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 which 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, i.e., 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 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 university, from technical college, from a university abroad, and a PhD. An individual still in school (category 0) is assigned a missing. 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 includes income from first job, secondary job, 13th and 14th salary, christmas bonus, holiday bonus, 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 in three broad categories as listed below.

Household Disposable Income. We construct this variable from Household Taxable Income (Head's and wife's income from assets, earnings, and net profit from self-employment) *minus* taxes *plus* public transfers. SOEP provides a measure of household asset flows, which is calculated as income from renting minus operating costs, plus dividend income.

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

Table A.2: Number of observations in the three panels after each selection step

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.

individual-year observations for the female sample, and 76,051 household-year observations for the household sample. 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
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 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 Kelly's skewness we consider two alternatives–2 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 A.3 show the results of the regressions for male and female earnings wages, respectively. The results are the ones from the main text and serve for 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 Kelly's skewness and L9050 and L5010, the estimates show that the substantive conclusion is robust also 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 as the reported ones for 50 weeks.

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 A.4 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, i.e., 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 table A.5.

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 (lower panel of table A.5).

C Additional Figures

Table A.3: Sensitivity of Regression Results – SIAB I

	Std Dev	L9010	Skew	Kelly	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 2.

Table A.4: Sensitivity of Regression Results – SIAB II

	Std Dev	L9010	Skew	Kelly	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 2.

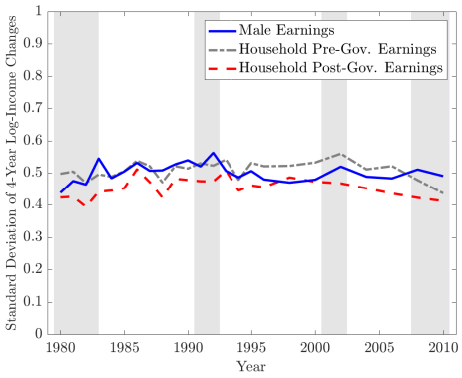
Table A.5: Sensitivity of Regression Results – SIAB III

	Std Dev	L9010	Skew	Kelly	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)

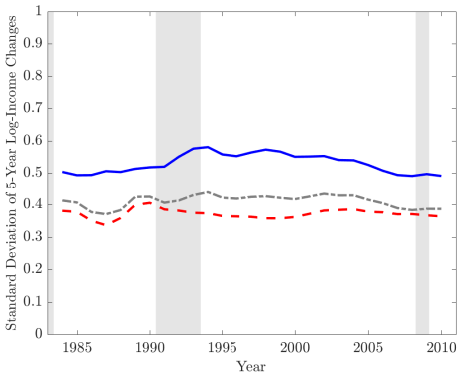
Note: See notes for Table 2.

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

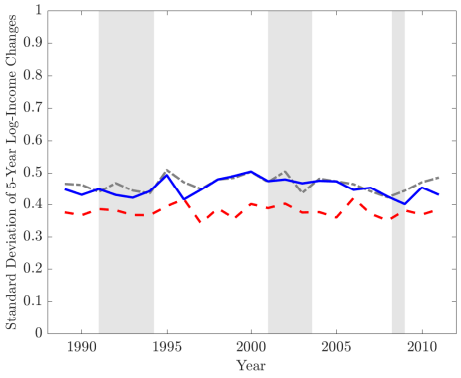
(a) United States



(b) Sweden



(c) Germany



Note: Linear trend removed, centered at sample average.

D Estimation and Simulation Details

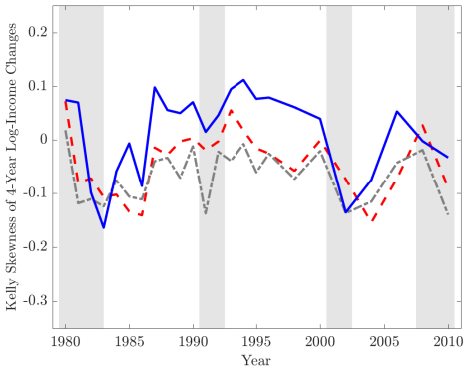
Simulation of Income Profiles

Let c_n^m denote the empirical moment n ($n = 1, \cdots, N$) that corresponds to cross-sectional target m ($m \in \{p10(\Delta^1 y_t), p10(\Delta^3 y_t), p10(\Delta^5 y_t), \cdots, var(y_{age=25}), ..., var(y_{age=60})\}$).

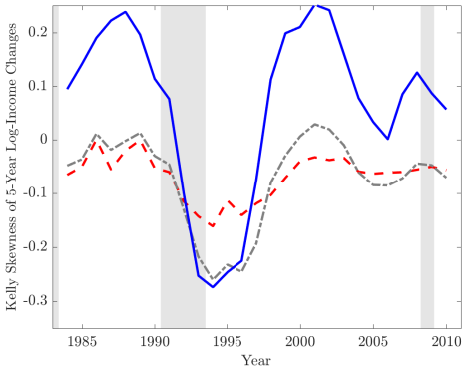
In each simulation, we draw a vector of random variables $X_r = \{\varepsilon_1, \varepsilon_2, ..., \varepsilon_{T-1}, \varepsilon_T, \zeta_1, \zeta_2, ..., \zeta_{T-1}, \zeta_T\}$, where T denotes the last year available in the data in levels. We simulate these profiles $R = 10$ times.

Figure A.2: Kelly Skewness of Long-Run Earnings Growth: United States, Germany, and Sweden

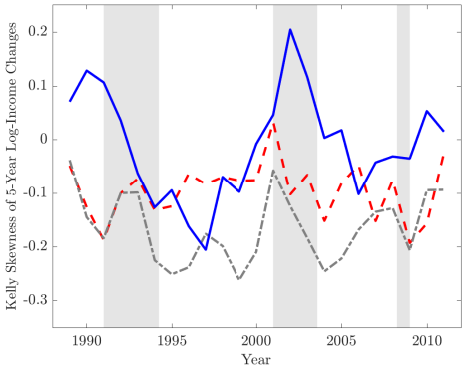
(a) United States



(b) Sweden



(c) Germany



Note: Linear trend removed, centered at sample average.

Optimization

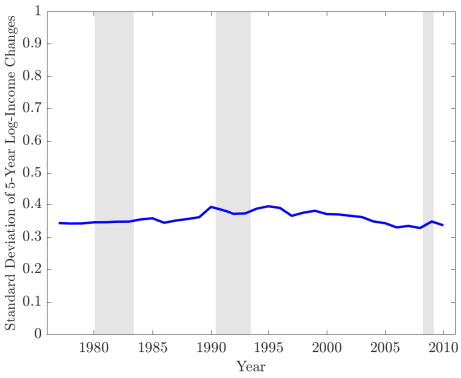
We minimize the scaled deviation $F(\theta)$ between each data and simulated moment.

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

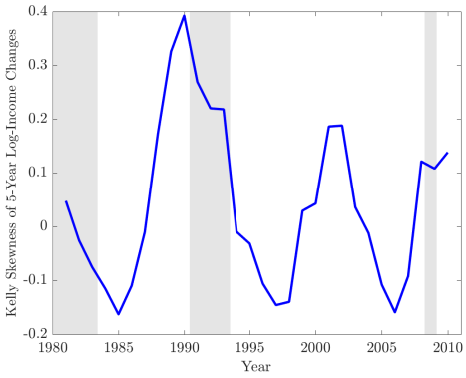
Where F is defined as:

Figure A.3: Standard Deviation, Skewness, and Tails of Long-Run Earnings Growth: Germany, IAB Sample

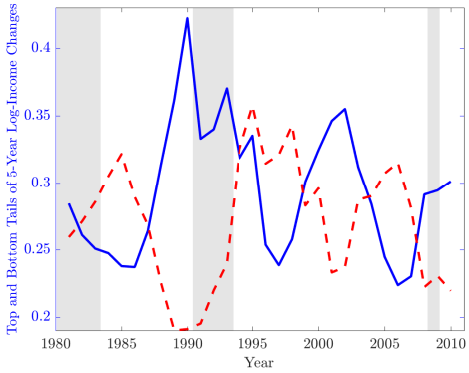
(a) Standard Deviation: Males



(b) Kelly Skewness: Males



(c) Upper and Lower Tails: Males



Note: Linear trend removed, centered at sample average.

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

Notice that the scaling can be problematic in the case of the P50 and the mean, which are very close to zero. γ^m is chosen such that, for those two moments, the denominator is, on average, equal to the sample average of the P90 moment. The average is calculated within the same group m. This is,

$$
\begin{aligned}
\gamma^m &= \frac{1}{T-1}\left[\sum_{i=1}^{T-1}|c_n^{p90,1}| - \sum_{i=1}^{T-1}|c_n^m|\right] \quad for \quad m \in \{p50(\Delta^1 y_t)\} \\
\gamma^m &= \frac{1}{T-3}\left[\sum_{i=1}^{T-3}|c_n^{p90,3}| - \sum_{i=1}^{T-3}|c_n^m|\right] \quad for \quad m \in \{p50(\Delta^3 y_t)\} \\
\gamma^m &= \frac{1}{T-5}\left[\sum_{i=1}^{T-5}|c_n^{p90,5}| - \sum_{i=1}^{T-5}|c_n^m|\right] \quad for \quad m \in \{p50(\Delta^5 y_t)\}
\end{aligned}
$$

and $\gamma^m = 0$ otherwise.

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%.

Data for Estimation

See Excel Data_for_estimation.xlsx.

Model Fit

Pre-Government Household Earnings[30]

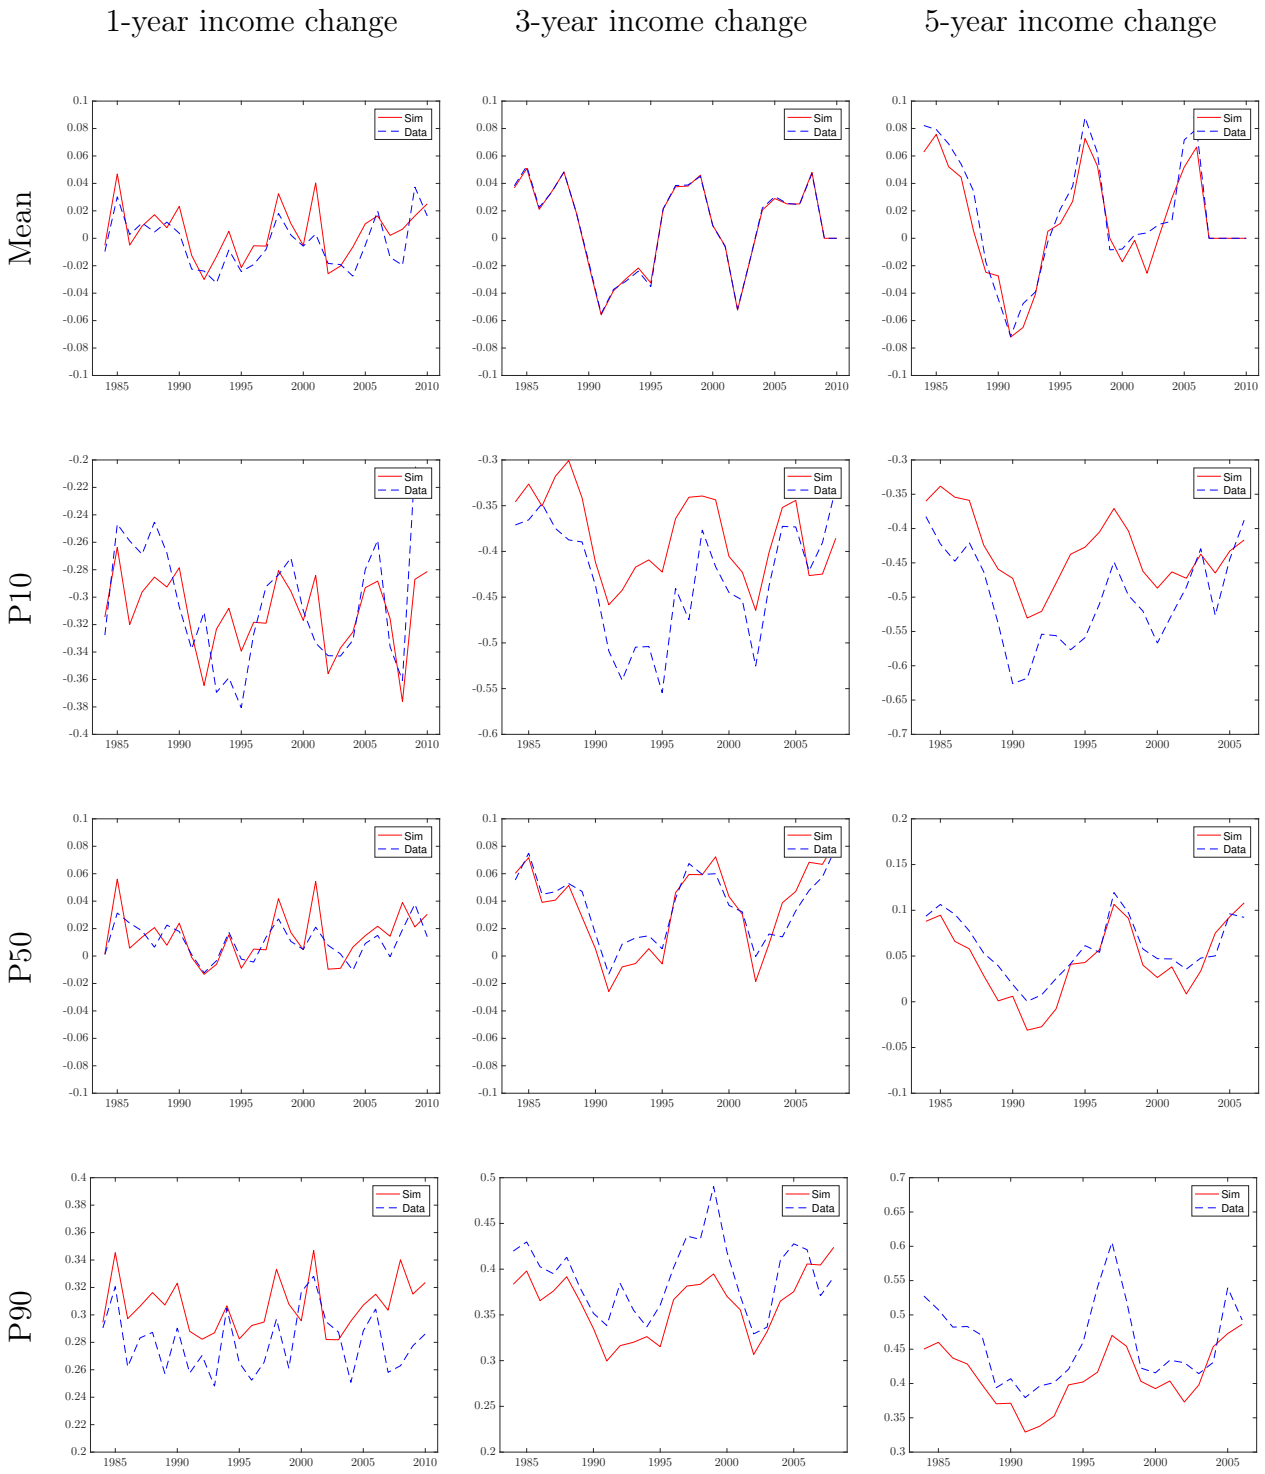


Figure A.4: Germany - Simulated (solid red line) and Data (dashed blue line) moments of earnings changes.

[30]Remember that the mean of the 3-year changes (4 in the US) is imposed to be as in the data.

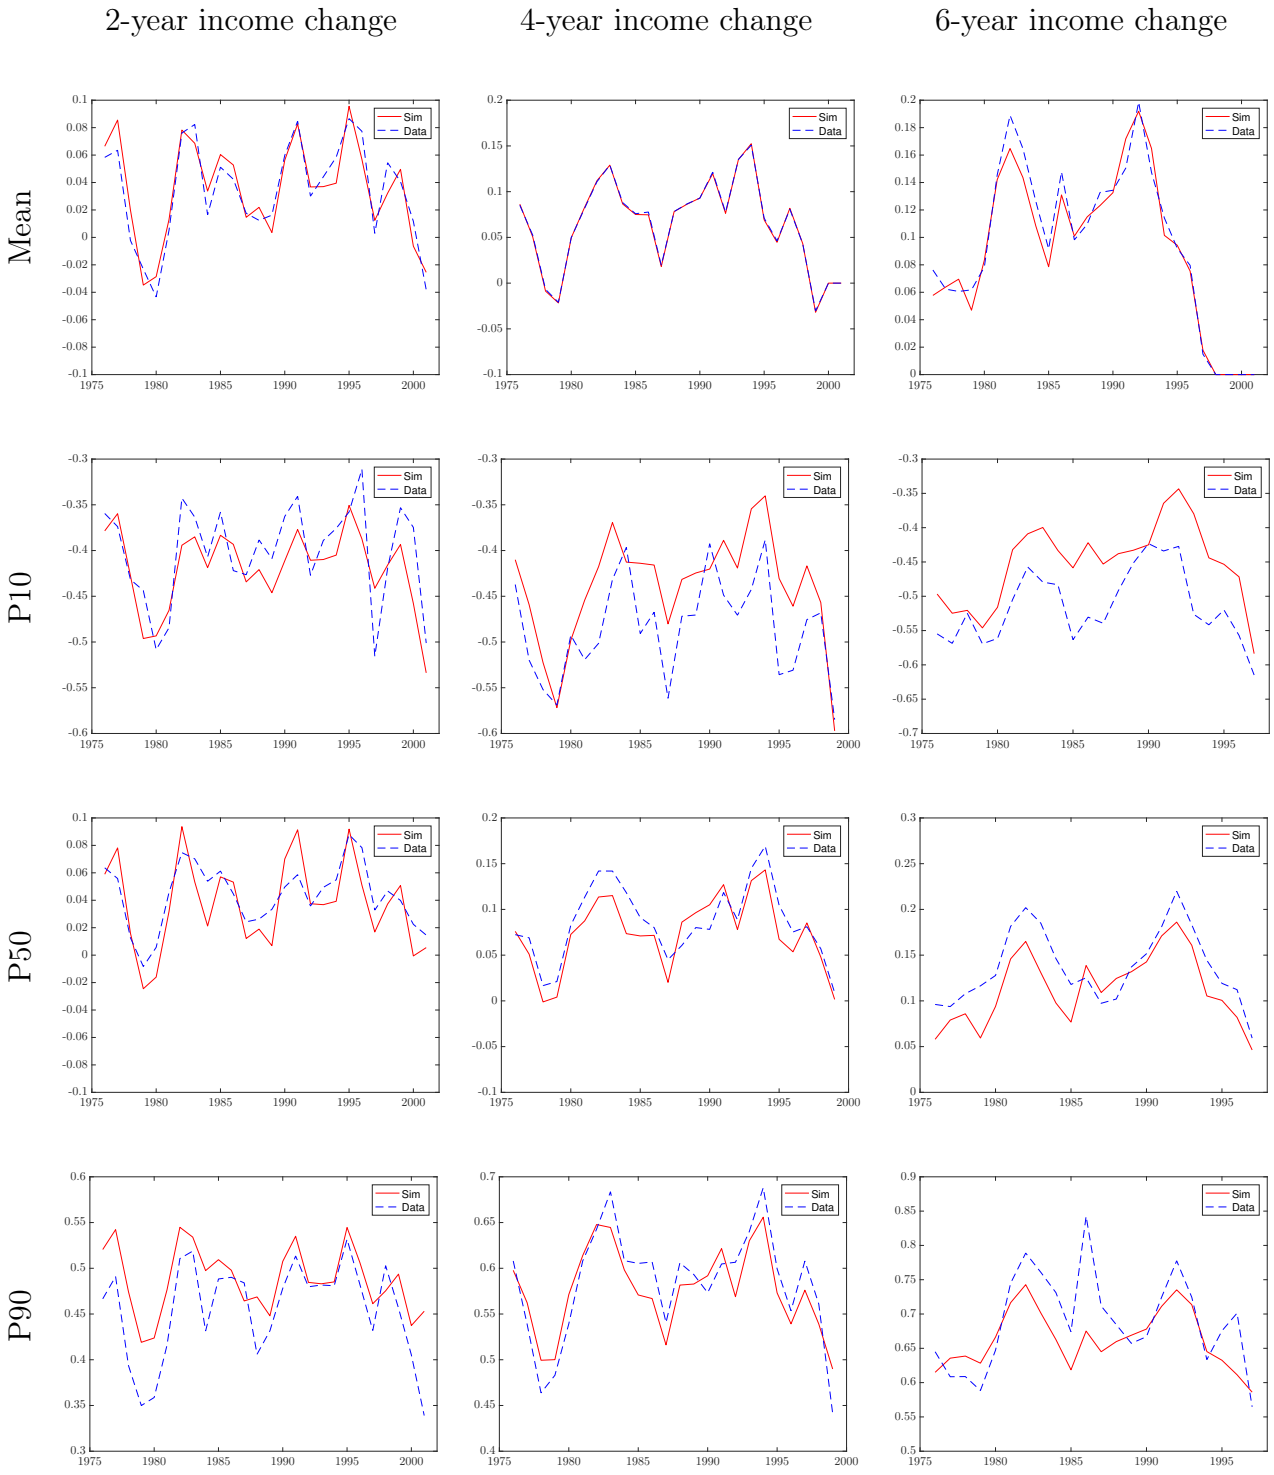


Figure A.5: United States - Simulated (solid red line) and Data (dashed blue line) moments of earnings changes.

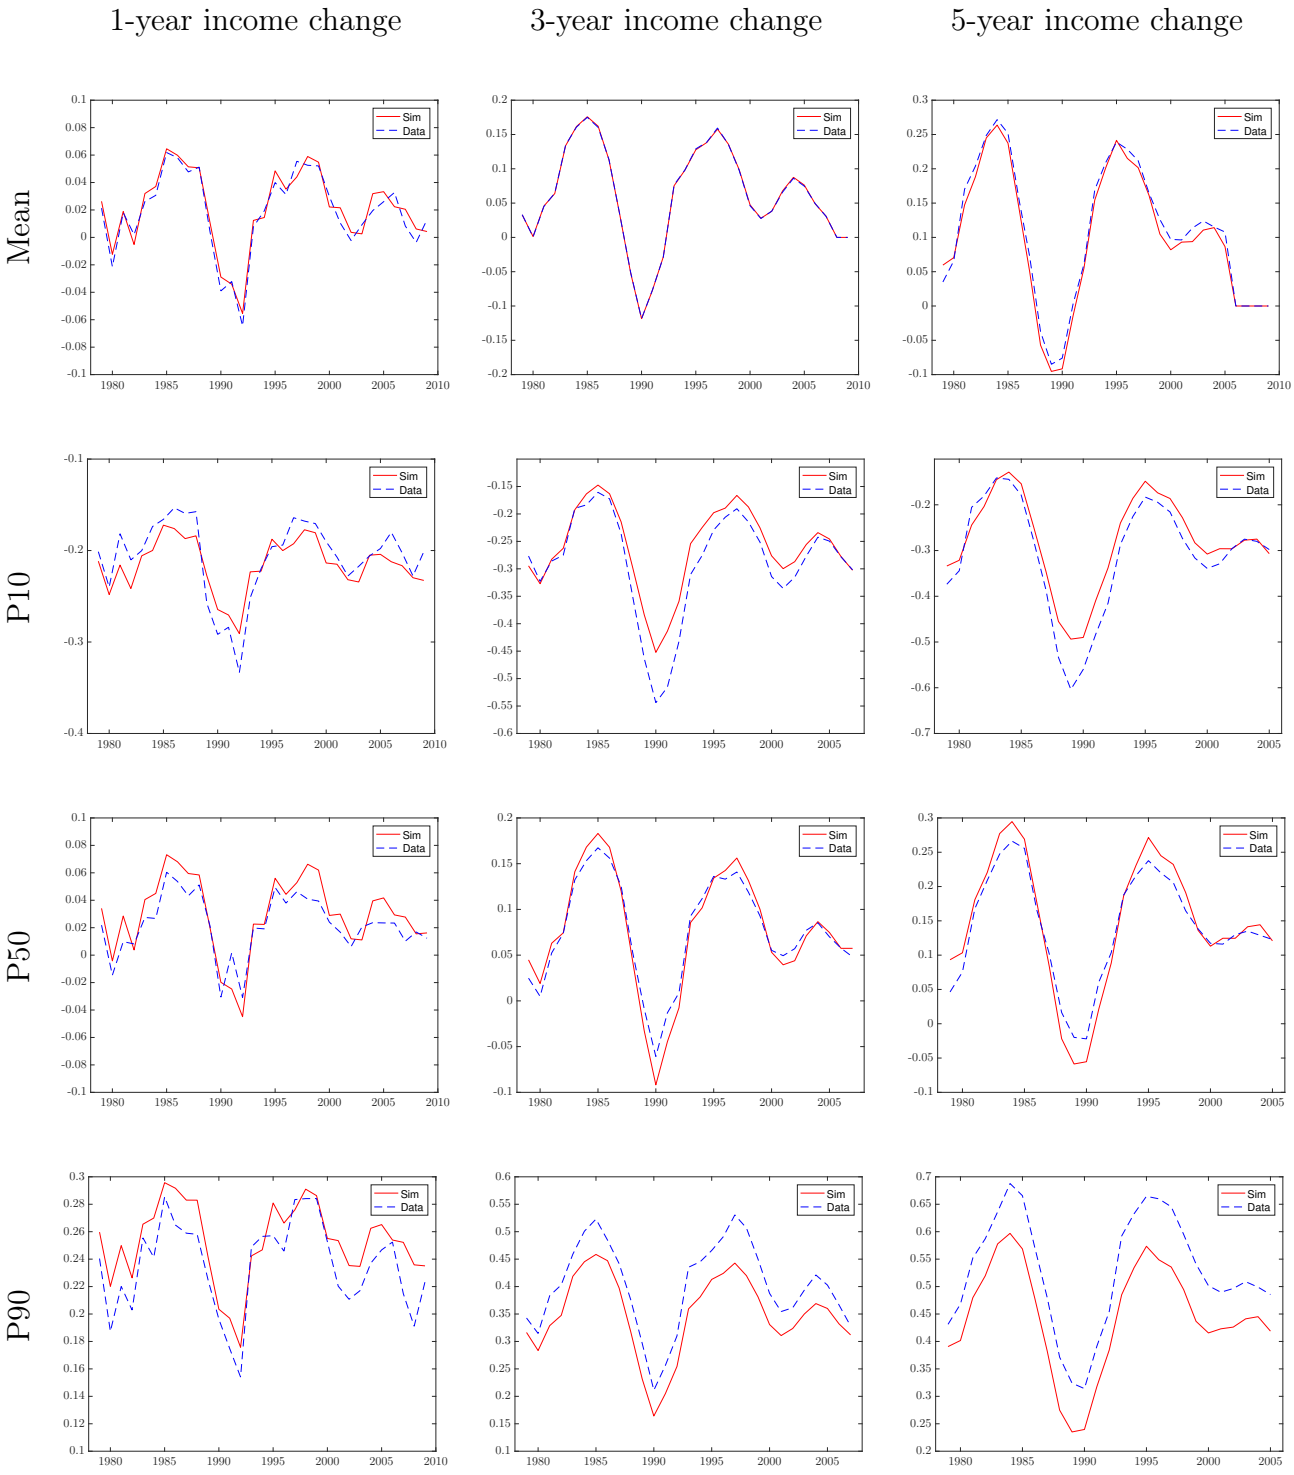


Figure A.6: Sweden - Simulated (solid red line) and Data (dashed blue line) moments of earnings changes.

Post-Government Household Earnings

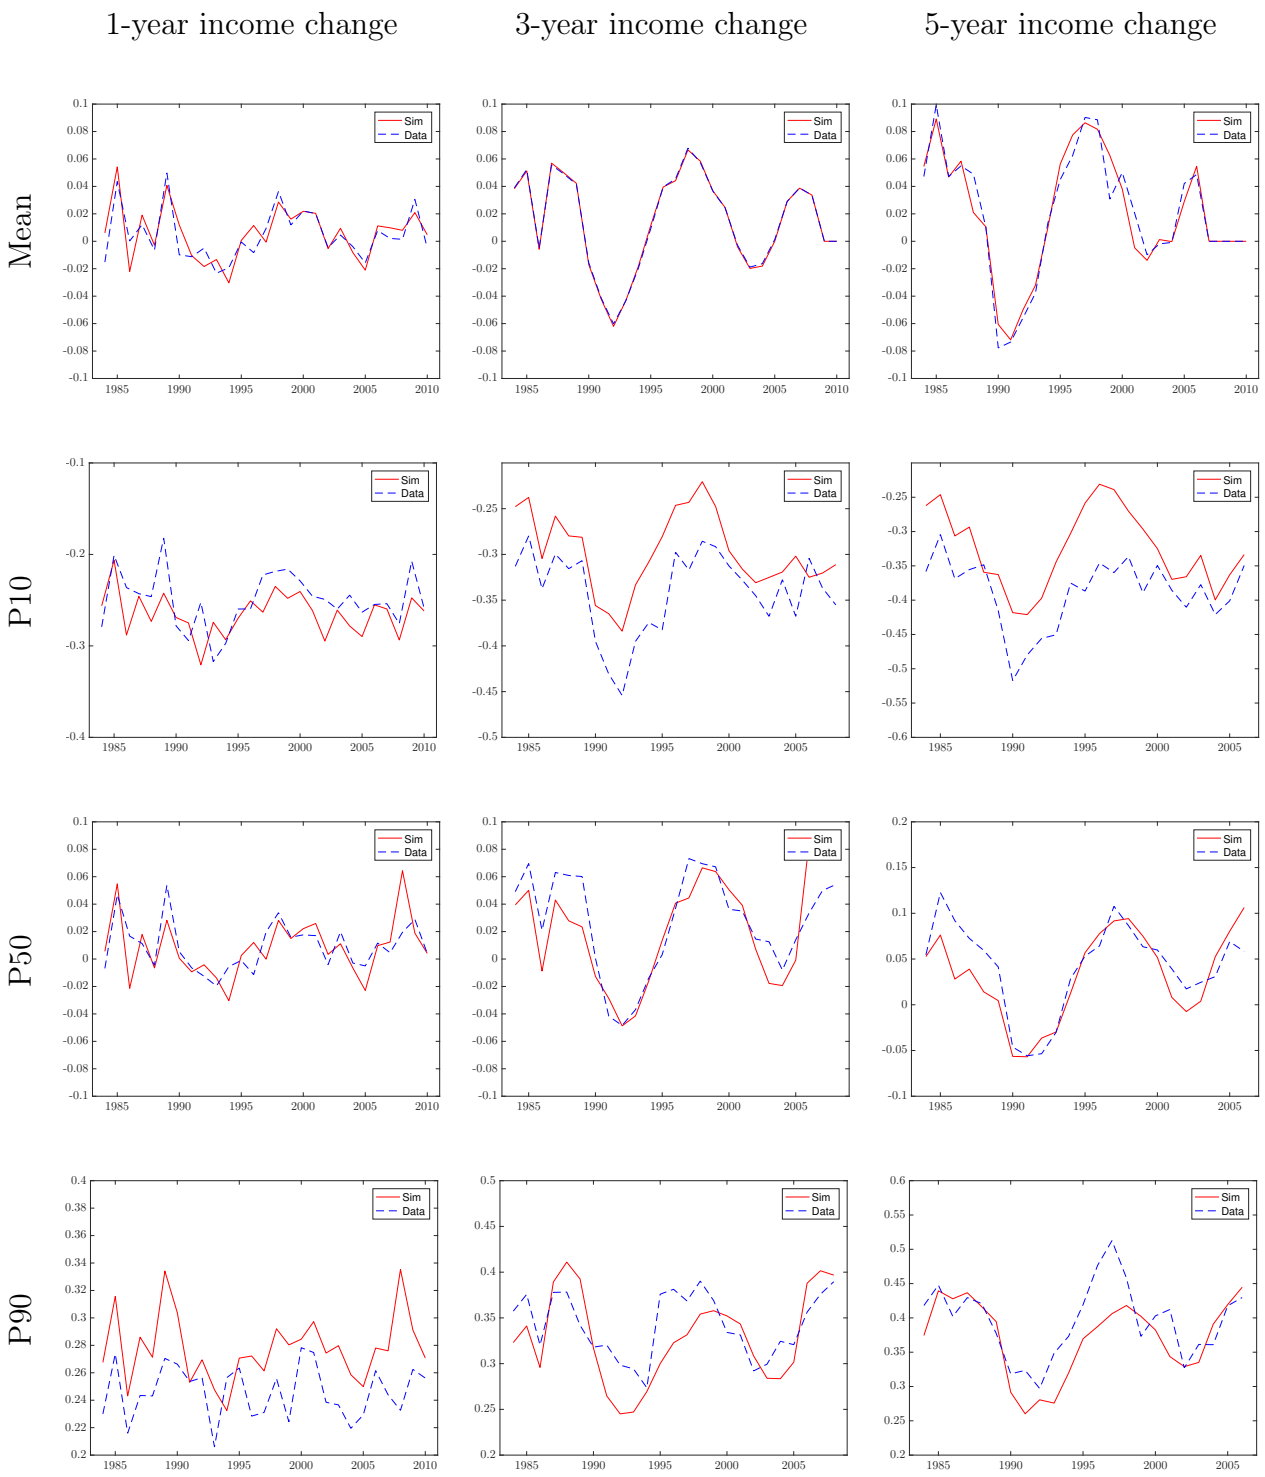


Figure A.7: Germany - Simulated (solid red line) and Data (dashed blue line) moments of earnings changes.

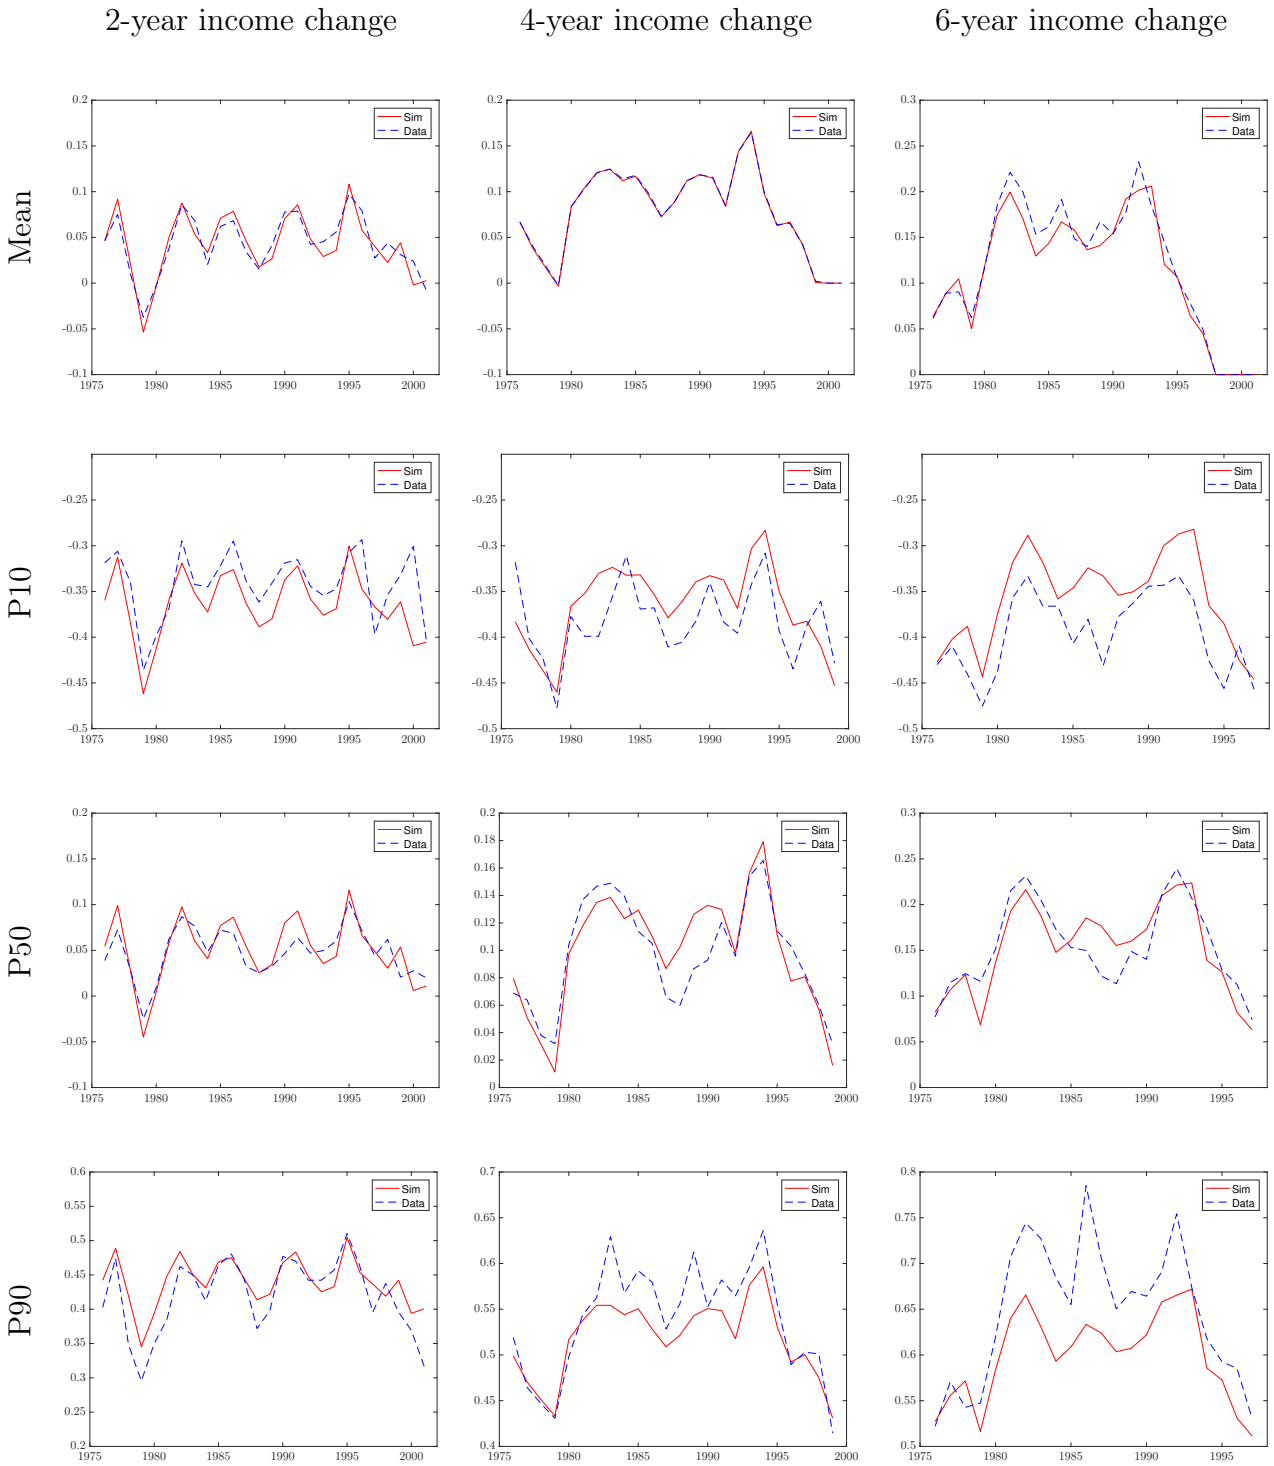


Figure A.8: United States - Simulated (solid red line) and Data (dashed blue line) moments of earnings changes.

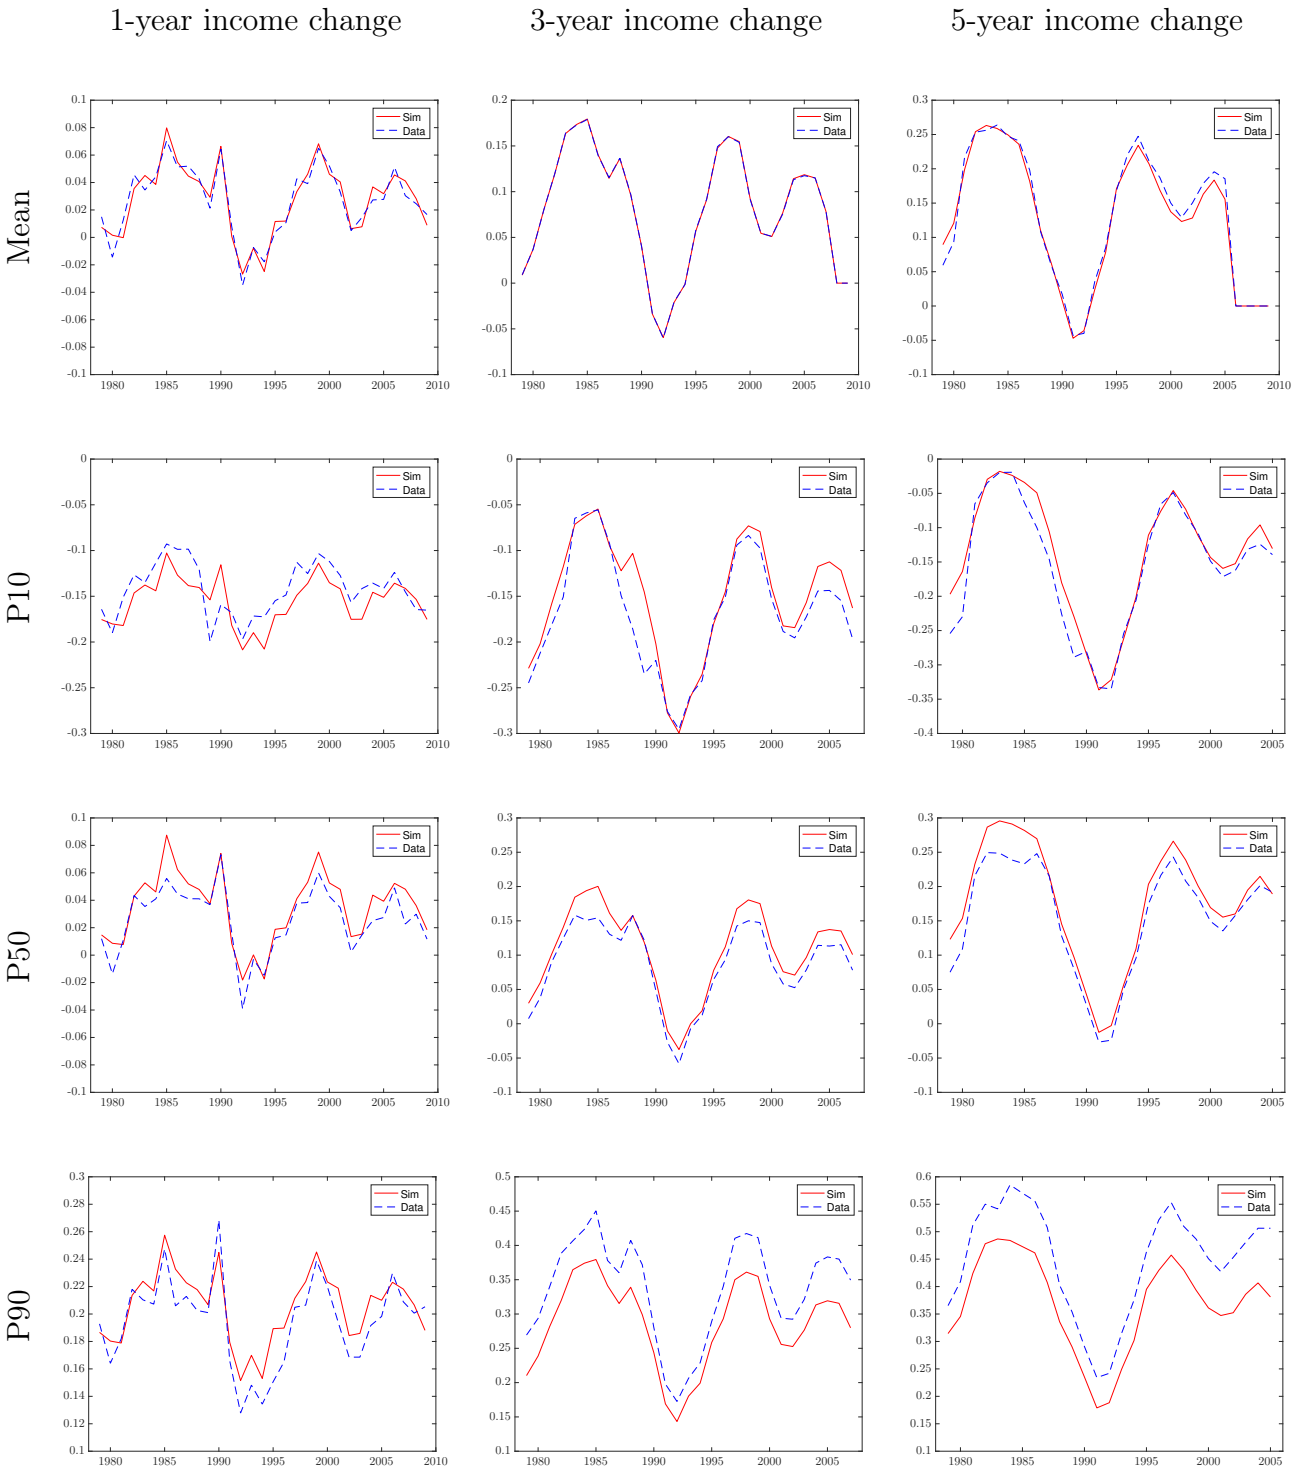
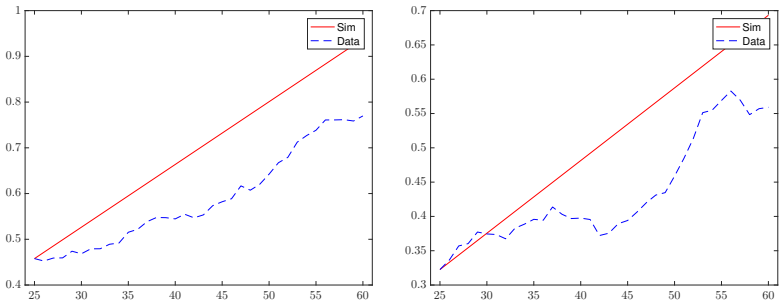


Figure A.9: Sweden - Simulated (solid red line) and Data (dashed blue line) moments of earnings changes.

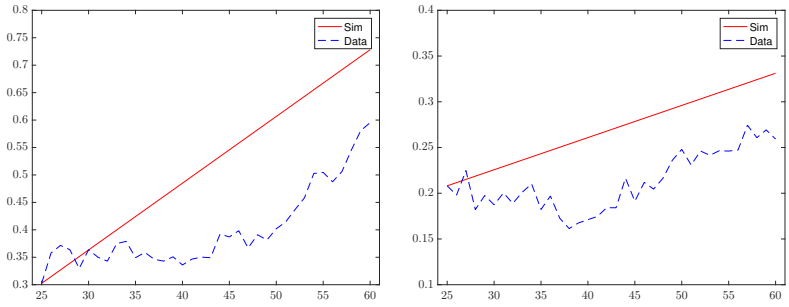
Variance of Pre- and Post- Government Household Earnings over the Life-Cycle

Figure A.10: Simulated (solid red line) and Data (dashed blue line) moments of life-cycle earnings variance.

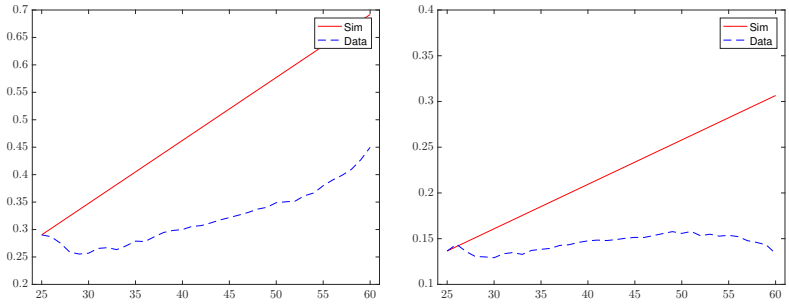
(a) United States, pre (left panel) and post (right panel)



(b) Germany, pre (left panel) and post (right panel)



(c) Sweden, pre (left panel) and post (right panel)



E Quantitative Model

Given estimates of the income process, we scale the parameters of the permanent shocks ζ to feed them into the model – fraction $\lambda$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 $1-\lambda$.

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