Higher-Order Income Risk and Social Insurance Policy Over the Business Cycle

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

February 13, 2015

Preliminary and Incomplete

Abstract

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. 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 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 5% for Germany.

JEL Codes: D31, E24, E32, H31

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

*CMR, University of Cologne
†Stockholm School of Economics
‡University of Minnesota, FRB of Minneapolis, and NBER; guvenen@umn.edu; www.fguvenen.com
§University of Minnesota

1 Introduction

In this paper, we study the effectiveness of social policies in insuring households against fluctuations in higher-order income over the business cycle. 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 being 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 survey-based panel datasets, such as the Panel Study of Income Dynamics (PSID), which required researchers to make parametric assumptions to obtain identification. Restricting attention to business cycle variation in the mean and variance of income shocks, these studies concluded that the variance of income shocks is countercyclical (e.g., Storesletten *et al.* (2004)). In recent work, Guvenen *et al.* (2014) use a very large panel dataset on earnings histories from the U.S. Social Security Administration (SSA) records. Using non-parametric techniques, they find 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 important 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.), (iii) no information is available on skills/education. Furthermore, the authors 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 countercyclical. The time span of all the databases we analyze covers more than 30 years. 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 are almost entirely constant over the business cycle, whereas the skewness of these shocks becomes much 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 significantly reduces the countercyclicality in the skewness of earnings in Sweden, but evidence of within-household insurances 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 most pronounced in Germany.

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 the Data Appendix. Briefly, we employ three 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 German Socio-Economic Panel (SOEP) for Germany, covering 1984 to 2011, and the Longitudinal Individual Data Base (LINDA) for Sweden, covering 1979 to 2010. Data for the PSID and the SOEP have been collected via surveys and questionnaires. 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.[2] LINDA is compiled from administrative sources (the Income Register) and tracks approximately 300,000 individuals per year.

[1] The PSID runs from 1967 to 2010. We select a subsample as our benchmark due to data quality, especially for income transfers. We complement our results using alternative year samples whenever possible and pertinent.

[2] 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.

For each country, we construct three samples: two at the individual level—one for males and one for females—and one at the household level. The individual samples contain those individuals aged 25 to 59 in periods t and $t+k$ ($k = 1$ or 5) and have yearly labor earnings above a minimum threshold. 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 dollars in 2010.[3] For each individual, we record age, gender, education, and labor earnings. 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).

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 if it has at least two members, one of them being the household head,[4] that satisfy the age criterion and household income that satisfies the income criteria. At the household level, we will analyze several different 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, we impose the condition that the minimum earnings across all measures we consider exceed our minimum threshold. 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 recession. The reasons are 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 contractions during our sample period.

[3]For the case of 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.

[4]In PSID and SOEP the head of a household is defined within the database. 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

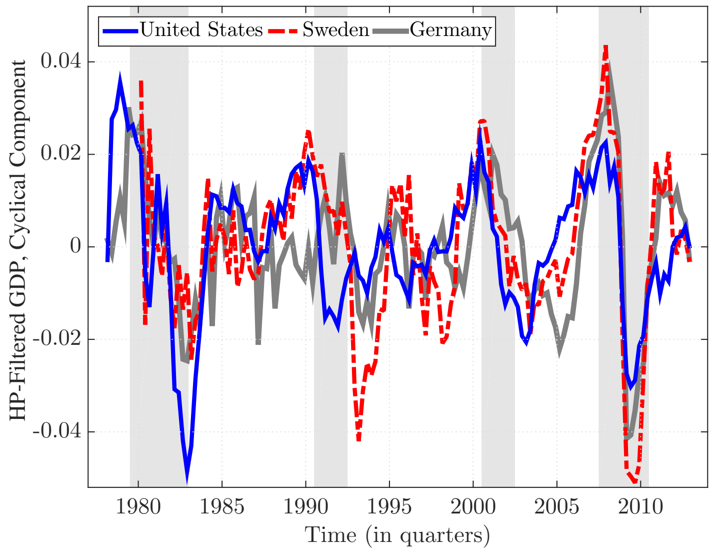

Note: The shaded areas in the figure indicate U.S. recessions according to our classification described in the text. The series for Germany corresponds to West Germany up to and including 1990Q4, and to (Unified) Germany from 1991Q1 on. The cyclical component is obtained by HP-filtering the series for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

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

3 Empirical Approach

Measuring Income Volatility over the Business Cycle

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

More specifically, we compute moments $m[\Delta_s y_t]$, where $y \equiv \log Y$ and $\Delta_s y_t \equiv y_t - y_{t-s}$. The measures moments m we consider are: the standard deviation, the (Kelly) skewness, and the top (L90-50) and bottom (L50-10) tails.[5] For the case of 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.[6]

We do not impose any parametric assumption on the dynamics of income but instead explicitly analyze the tails of the distribution of earnings changes. It will turn out that it is important to understand the behavior of both tails and that focussing on the variance as a summary statistic of the distribution alone can be misleading. To see this point, it is useful to think of the distribution in terms of its tails. An increase in the variance of a normally distributed variable implies a widening of both upper and lower tails: P90 is shifted to the right and P10 is shifted to the left. However, the overall dispersion of the distribution can be higher when only one of the two tails widens. 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 scenarios imply a change of the relative size of the tails—a feature summarized in 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 the tails should be explicitly analyzed, and when relying on summary statistics of the distribution, not only the variance should be taken into account, but more importantly, the skewness. Note how any assumption on the distribution of income shocks would drive our results: a (log-) normal distribution cannot capture changes in skewness. In light of recent evidence on male earnings growth using administrative data for the United States, we want to take a skeptical point of view.

[5] $L9050 = P90 - P50$ in logs, and equivalently for $L5010$.

[6] We calculate overlapping 2- and 4-year differences up to the 1994–1996 change (survey years 1995-1997), and non-overlapping 2- and 4-year differences from then until the 2008–2010 change (corresponding to survey years 2009-2011).

Table I: 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	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.

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.

Specifically, we consider six statistics of earnings growth: (i) standard deviation, (ii) log differential between the 90th and 10th percentiles, denoted as L90-10, (iii) Skewness, measured as the third standardized moment, (iv) Kelly's measure of skewness, (v) L90-50 and (vi) L50-10. We use the (natural) log growth rate of GDP—$\Delta_s GDP_t \equiv \ln(GDP_t) - \ln(GDP_{t-s})$—as our measure of aggregate fluctuations. Therefore we consider the following regression of 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 :

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

4 Empirical Results: Raw 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 compensation from one's 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 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.

In this section we examine three questions about higher-order risk in individual incomes. First, we ask if the counter-cyclicality of skewness and 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?

Then, we ask if the countercyclical skewness and acyclical dispersion is a US-only phenomenon, or is it a robust feature of business cycles more generally 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 varies 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 into the sources of the cyclicality, providing a deeper understanding of systematic components in the response to business cycle fluctuations.

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

Table II: Cyclicality of Male Earnings, by Education Groups

	Std Dev	L9010	Skew	Kelly	L9050	L5010
			United States			
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			
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	0.22***	0.40**	7.92**	2.66***	0.87***	–0.46***
	(3.81)	(2.46)	(2.48)	(4.38)	(4.06)	(–5.50)
Non-college	0.03	–0.31**	11.69***	4.14***	0.96***	–1.27***
	–0.50	(–2.29)	(4.93)	(4.14)	(3.47)	(–3.77)
			Germany			
All 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)
College graduates	0.20	–1.33	35.09***	1.94*	–0.22	–1.11**
	(0.44)	(–1.66)	(3.95)	(2.02)	(–0.57)	(–2.08)
Non-college	–0.28	–1.41**	3.79	1.34***	–0.34	–1.07***
	(–1.47)	(–2.64)	(0.54)	(2.91)	(–1.25)	(–3.42)

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. In parentheses is the Newey-West t-statistic (maximum lag length considered: 3 for SOEP and LINDA, 2 for PSID). Asterisks ($*, **, * * *$) denote significance at the (10%, 5%, 1%)-level. Household income measures adjusted by square root of household size.

Cyclicality of Dispersion

In Table II, we report the cyclicality of six statistics of the distribution of earnings changes for male workers. To provide a comparative discussion we report the results for all three countries in the same table. The first column reports the cyclicality of the standard deviation. In the United States, for male workers, the standard deviation is acyclical, as seen from the small coefficient (0.20), which is also insignificant. 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 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 two columns of the second panel, both measures of dispersion are acyclical in Sweden, with very small coefficients that are insignificant. Turning to Germany (bottom panel), the standard deviation is again acyclical, whereas L90-10 shows some signs of counter-cyclicality. As we comment on in a moment, a closer inspection will show that the countercyclical variation in L90-10 is driven by a change in skewness.

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 (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, being a third order moment, it has a tendency of being sensitive to outliers, which can be a concern for the U.S. and German data, since these are surveys (with possibly large measurement error) with modest sample sizes. This is less of an issue for Sweden, 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 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 raised above. 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

Figure 2: Standard Deviation (Left) and Skewness (Right) of Short-Run Earnings Growth: United States, Germany, and Sweden

(a) United States, Std. Dev.

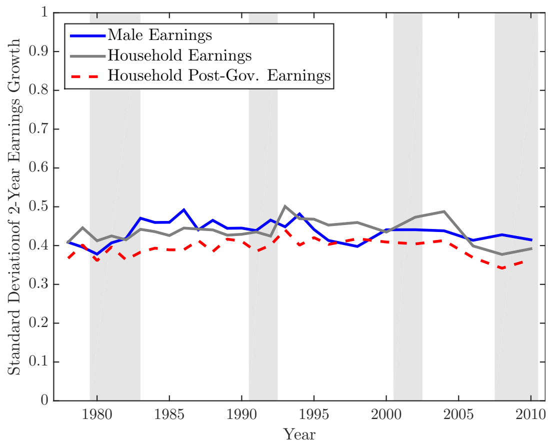

(b) United States, Kelly's Skewness

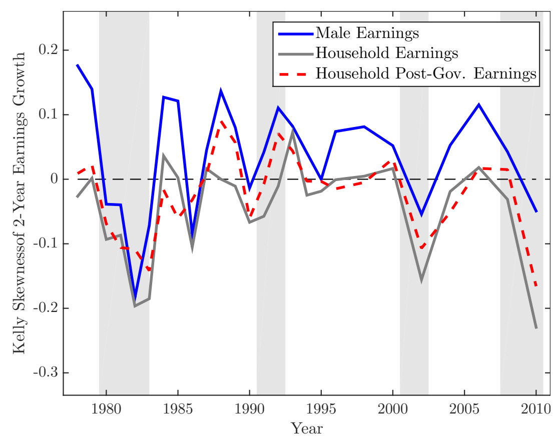

(c) Sweden, Std. Dev.

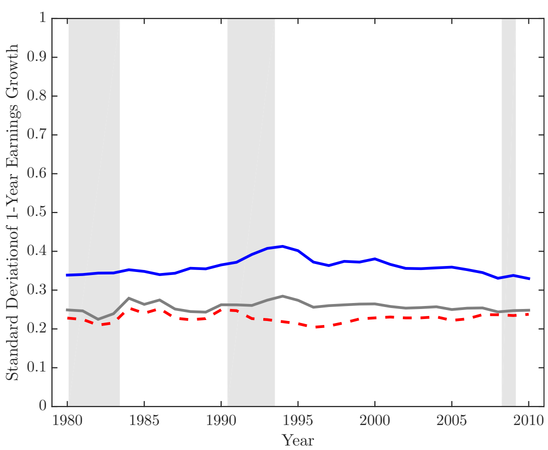

(d) Sweden, Kelly's Skewness

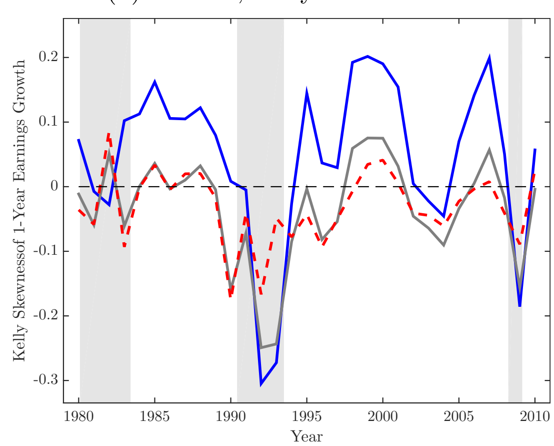

(e) Germany, Std. Dev.

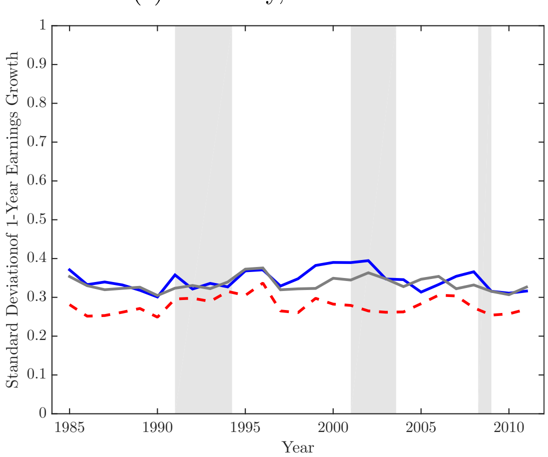

(f) Germany, Kelly's Skewness

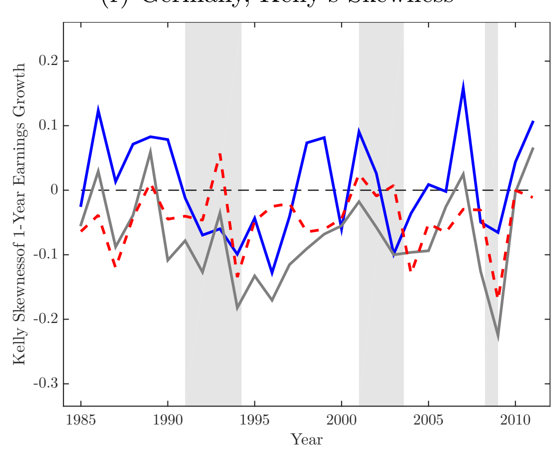

Note: Linear trend removed, centered at sample average.

Figure 3: Tails of Short-Run Earnings Growth: United States, Germany, and Sweden

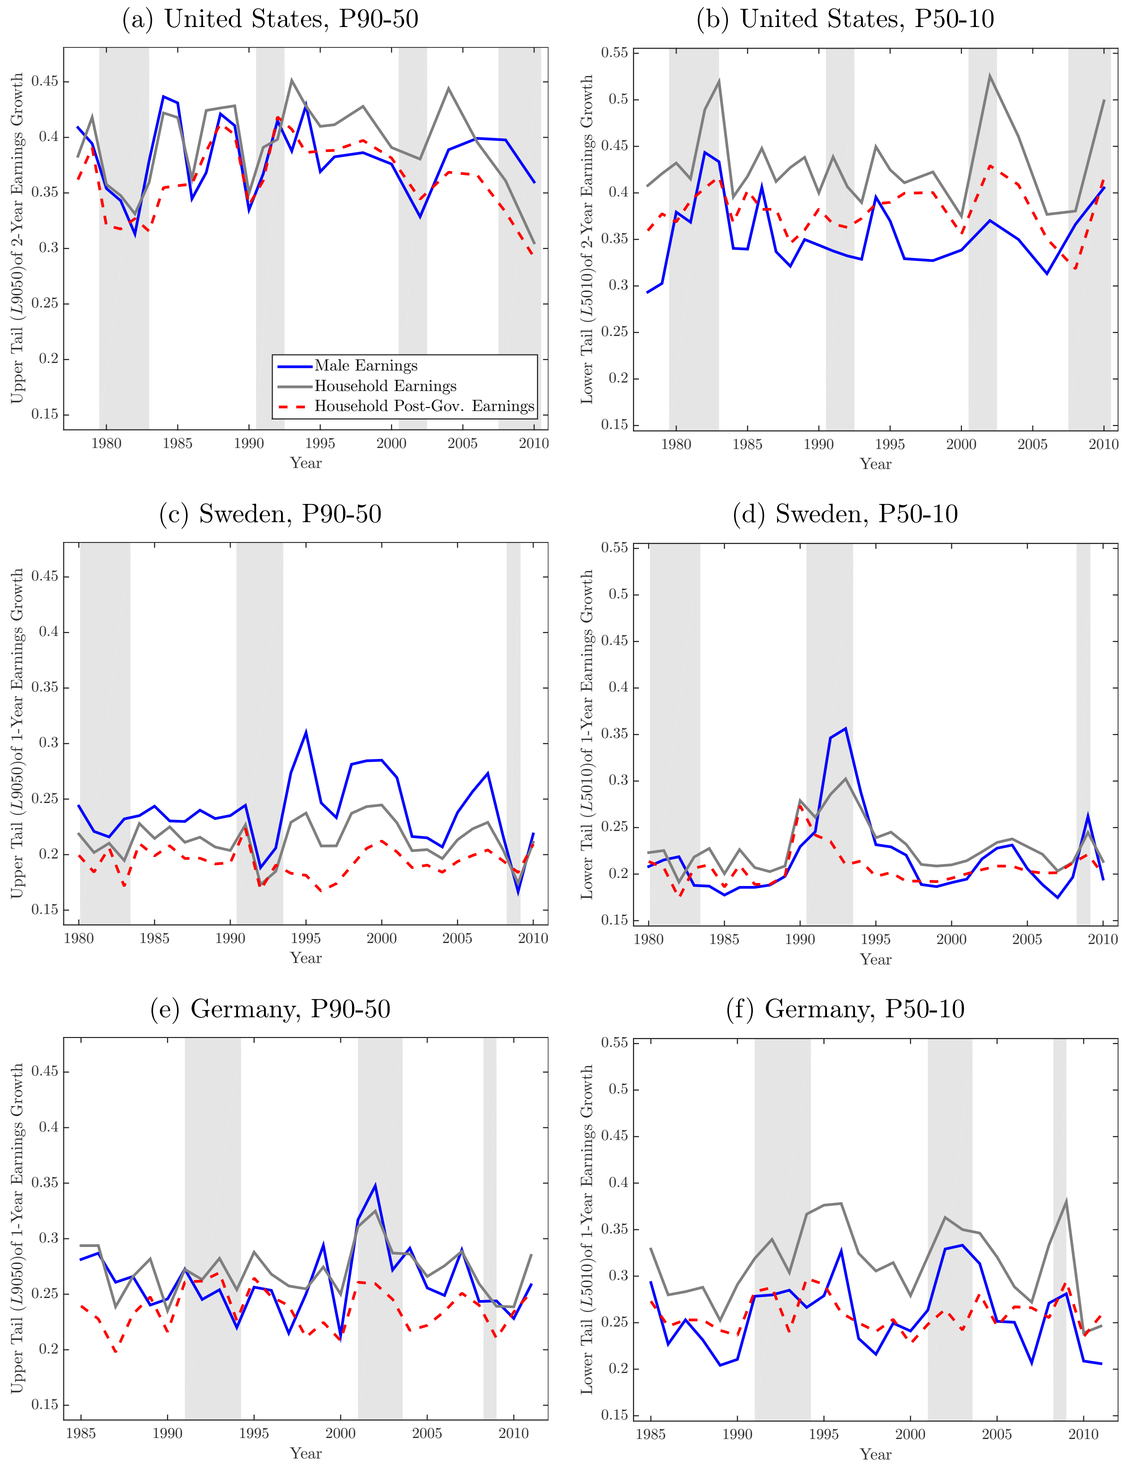

Note: Linear trend removed, centered at sample average.

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 II. In all three countries, Kelly's skewness is procyclical and (statistically) significant at 1%. Interestingly, the coefficient is about 1.7 for the U.S. and Germany, and double (3.7) for Sweden, showing more cyclicality for the latter country. 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 is often thought as displaying lower business cycle risk. Finally, the coefficient on the third moment measure is also positive in all three countries, consistent with Kelly's skewness, and is significant at 1% level in Sweden, and at 5% level in the U.S. and 10% level in Germany.[7]

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 shrinking of the right tail or an expansion of the left tail or both. Thus, a follow-up question is: which one of these changes are driving the cyclical changes in skewness for each country? The last two columns report the cyclicality of the L9050 and L5010. Notice that in the U.S. and Sweden, 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 concert during recessions.

Furthermore, notice that for both of these 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. Therefore, as log GDP growth fluctuates over the business cycle, the

[7]The corresponding changes in $\mathcal{S}_k$ for the U.S. and Germany are: 0.11 and 0.07 respectively.

Table III: Cyclicality of Female Earnings, by Education Groups

	Std Dev	L9010	Skew	Kelly	L9050	L5010
	United States					
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					
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	0.26***	0.43*	3.52**	1.60***	0.68***	–0.25**
	(3.16)	(1.85)	(2.12)	(3.98)	(3.76)	(–2.12)
Non-college	0.32**	0.53*	6.80***	1.89***	0.79***	–0.26**
	(2.62)	(1.88)	(3.87)	(3.39)	(2.67)	(–2.70)
	Germany					
All females	–0.55**	–1.26*	2.54	1.72	–0.08	–1.18***
	(–2.43)	(–1.86)	(1.10)	(1.71)	(–0.13)	(–3.15)
College graduates	–0.59	–2.30	20.78**	2.28	–0.16	–2.14**
	(–0.96)	(–1.55)	(2.39)	(1.63)	(–0.17)	(–2.21)
Non-college	–0.66**	–1.08	–2.46	1.08	–0.23	–0.85*
	(–2.49)	(–1.37)	(–0.91)	(0.86)	(–0.30)	(–1.99)

Note: See Table II for explanations.

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. Notice that the behavior is different in Germany, were the L9050 is acyclical (coefficient of –0.21 with a low t-statistic) whereas the coefficient for L5010 is –1.12 and significant.[8] Therefore, skewness is procyclical *and* dispersion (at least as measured by L9010) is countercyclical,

[8]The analysis for Germany overall suggests that male earnings changes are very cyclical at the lower end of the distribution. Inspection of the time-series of the tails of the distribution reveals that the 50-10 gap of transitory earnings changes widens in all three covered recessions. The upper tail appears to close in the course of the three recessions. The time-series of permanent earnings changes suggest a similar pattern for the two 1990s and early 2000s contractions. As expected from the tails, the implied measure of skewness shows that left-skewness of transitory earnings changes increases substantially in contractionary periods. Also the distribution of permanent earnings changes is more skewed to the left for the years directly after the two recessions for which enough subsequent years are observed in order to assess permanent earnings changes.

because the expansion of the lower tail is not offset by a shrinking of the top tail. However, it is clear that this rise in dispersion is driven by asymmetric movements in tails, and should not be thought of as a pure rise in variance (which would imply an expansion of *both* tails).

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 greatly for the nature of earnings risk.

We next dig deeper to see if the patterns we found regarding higher-order risk 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.

4.1 Differences in Cyclicality by Gender

We now turn to the cyclicality of higher-order risk for female workers and examine how they compare to the patterns discussed for males. Table III reports the analogous set of results for females. 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 countercyclical for German women. This is different from men who displayed acyclical dispersion. Second, Kelly's measure of skewness is always procyclical as indicated by the positive coefficient on log GDP growth, but the coefficient is only significant for Sweden, at 1% level, and for the United States, at 5% level. The third central moment is only significantly procyclical in Sweden. As noted before, this might be due to the smaller sample sizes for the U.S. and Germany.

Third, inspecting the top and bottom tails separately, it is evident that they always have the expected pattern of cyclicality when their coefficient is significant. In particular, the L9050 is procyclical and significant for the U.S. and Sweden, whereas the L5010 is countercyclical and significant for Sweden and Germany.[9] Thus, just as for men, the behavior of the variance is driven by the two tails not moving in concert and therefore is more of a reflection of skewness rather than dispersion increasing uniformly (with both tails expanding). Therefore, we conclude that the variance does not seem to be right statistic to focus on when it comes to measuring higher-order earnings risk over

[9]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.

the business cycle. Finally, it is worth noting that the magnitudes of fluctuation in both Kelly's skewness and in the upper and lower tails separately is somewhat muted for women compared with men.

4.2 Differences in Cyclicality by Skill Groups

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 II), in all three countries, both education groups display procyclical Kelly's skewness that is statistically significant. In United States and Sweden, the magnitude of cyclicality is quite a bit stronger for less educated workers—about three times stronger in the U.S. and two times stronger in Sweden. In Germany, the opposite is the case but the difference between groups seems small. Again, turning to standard deviation, there is not very clear pattern in any country: variance is either acyclical or when it is pro- or counter-cyclical, the magnitude is generally small.

Turning to women in Table III, Sweden again emerges as the country with the clearest patterns. Skewness is strongly procyclical for both education groups; the lower tail is robustly counter-cyclical and the upper tail is procyclical for both skill groups. The variance is also procyclical for both education groups, 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.

In the U.S., the variance is more robustly procyclical for less educated women and acyclical for college graduates. In Germany, the lower tail (L5010) appears to be the only statistic that is statistically significant in its cyclicality. Again, the lack of significance for these two countries might be due to the lower share of employment among women (especially in the earlier part of the sample) and relatively modest sample sizes.

4.3 Differences in Cyclicality by Type of Employment: Private- vs. Public Sector

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.[10] Public sector jobs are often thought of as less risky, offering

[10]An individual belongs to sector x if the individual works in sector x in both t and t+1 (or t+5). This of course implies that some observations are lost (those switching between sectors).

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. and Germany, we are able to conduct this analysis for Sweden only. This is also perhaps the most interesting country to consider in this regard given the large share of public employment and general impression regarding lower risk coming from the welfare state policies.

Table IV: Cyclicality of Individual Earnings, by Sector of Employment – Sweden

	SD	L9010	Skew	Kelly	L9050	L5010
			Males			
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)
			Females			
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)

Note: See Table II for explanations.

Our data set does not include a direct indicator of public sector employment. However, some sectors in Sweden 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 $k = 1, 5$).[11] The split between private/public employment varies substantially by gender: 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).

Table IV reports the cyclicality regressions separately for workers in private sector versus public employment (by gender). For males, the standard deviation is acyclical for both sectors and both measures of skewness are strongly procyclical. However, the magnitude of the procyclicality of Kelly's measure is lower for the 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

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

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 in fact is statistically significant at 0.1% level or better. 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.)

Turning to female workers, we see the variance of earnings changes being slightly procyclical in both sectors (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 in the public sector.

Overall, it is somewhat surprising that even for workers in the public sector and in a country like Sweden with a reputation for 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. 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 Introducing Insurance

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 between-group differences, most of our samples are composed by individuals in cohabitation.[12] 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.

[12] Only 12% of our benchmark individual sample in the United States lives in a singe person household. This number is 13% and XX for Germany and Sweden, respectively.

Mixed Evidence of Within-Family Insurance

The first row of each panel in Table V 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 II. Additional evidence comes from the graphical analysis of the dispersion, skewness, and the tails, of male earnings changes and household earnings changes in Figures 2 and 3, respectively.

Table V: Cyclicality of Household Earnings

	Std.Dev	L9010	Skew	Kelly	L9050	L5010
			United States			
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			
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			
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 II.

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 – the lower tail widens more for household earnings.

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

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

The second component considered, "aid to low-income families," consists of several measures of social insurance policies specifically aimed at at-risk households. The relevance of this type of transfer can therefore be expected to matter most for low-income households who have a higher likelihood of 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 V. Again, Figures 2 and 3 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 lowered (from a point

Table VI: 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 to2004 (incl.)) Unemployment benefits II (since2005 (incl.))	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.

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 the L9050 of pre-government earnings changes, the L9050 of post-government earnings changes is acyclical in Sweden, but still procyclical in the U.S.. 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. It is therefore well suited to get a feeling for the overall effect of the tax and transfers system. Its is not well suited to actually judge the success of social policy that specifically aims at mitigating downside risk or that aims low-income families, who can be expected to be especially vulnerable in recessionary periods, because access to private insurance arrangements is limited for those families. Therefore, we now consider the stabilizing effect of different policies by adding received transfers of different types to gross earnings and thereby generating artificial income measures at the household level. The results of the cyclicality analysis are listed in Table VII. 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 V. 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.

Our analyses suggest that in all three economies labor market related transfers (which have unemployment benefits as the main component) are successful in mitigating down-

Table VII: Cyclicality of Household Earnings - transfers added separately

	Std.Dev	L9010	Skew	Kelly	L9050	L5010
	United States					
+ 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					
+ 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					
+ 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 II.

side risk in recessions. This supports the consideration of an increased incidence of unemployment in recessions. In the United States the cyclicality 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 V. 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 procyclical 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[13]). 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.

6 Welfare Analysis

In this section, we perform a simple quantitative exercise in order to gain some insight on the welfare gains coming from the tax and transfer system. Suppose that households all live in autarky, so that their consumption equals their income. Next, we impose a CRRA per period utility function and, for each household, calculate the present value of lifetime utility. This calculation is done for pre-government income and for post-government income. One can now calculate the percentage compensation of per-period consumption, that is necessary to make households indifferent between the compensated pre-government income stream and the post-government income stream.

Household i's lifetime utility in autarky with pre-government income is given by:

$$U_i^{pre}\left(\{c_{ia}^{pre}\}_a\right) = \sum_{a=25}^{60} \beta^{a-25} u\left(c_{ia}^{pre}\right) \underset{(autarky)}{=} \sum_{a=25}^{60} \beta^{a-25} u\left(y_{ia}^{pre}\right),$$

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

where

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma},$$

with γ the degree of (constant) relative risk aversion. y_{ia}^{pre} is household i's pre-government income and c_{ia}^{pre} is household i's consumption at age i. We can calculate (utilitarian) welfare for a given compensation λ as

$$W^{pre}(\lambda) \equiv \sum_{i=1}^{H} U_i^{pre}\left(\left\{(1+\lambda)\, c_{ia}^{pre}\right\}_a\right),$$

where H denotes the total number in the sample.

Now consider the government that gives transfers of different types and levies taxes. Note that taxes and transfers do not only affect the volatility of the income stream, but also mean income. Given the focus of this paper on social policy as a form of insurance against business cycle fluctuations, we adjust post-government income for first-order effects in the following way:

$$\tilde{y}_i^{post} \equiv \frac{y_i^{post}}{\bar{y}_i^{post}} \bar{y}_i^{pre},$$

where $\bar{y}_i^{pre}$ and $\bar{y}_i^{post}$ are household i's mean lifetime pre-government and post-government income, respectively. For each household, adjusted post-government income matches pre-government mean income over the life cycle and preserves the relative changes of the post-government income stream. Given (adjusted) post-government income, we calculate welfare in the same way as for pre-government income. As a measure of the welfare effect of social policy we now calculate the percentage change λ of per-period consumption that is necessary to make individuals indifferent between the compensated pre-government income stream and the post-government income stream: $W^{pre}(\lambda) \equiv W^{post}$.

The homogeneity of the per-period utility function allows a closed-form calculation of λ as

$$\lambda = \left(\frac{W^{post}}{W^{pre}(0)}\right)^{\frac{1}{1-\gamma}} - 1.$$

For the calculation we impose a discount factor of $\beta = 0.95$ and a level of relative risk aversion of $\gamma = 2$. Starting with our baseline household sample, we keep only households with at least 10 years of observations in order to being able to calculate a meaningful average income over the life cycle. This leaves us with 2,410 households in the PSID sample and with 3,027 households in the GSOEP sample.

Table VIII shows the calculated welfare gains. Focusing on the second order effect of the tax and transfer system, we find positive welfare gains in all three economies. Gains are high in Germany and Sweden, where households would need a compensation of 5.14% and 4.54%, respectively, of annual pre-government income in order to be indifferent to the actual income stream after taxes and transfers. Given average yearly gross income of about \$43,100 and \$37,350[14] this amounts to an average compensation of about \$2,215 and \$1,700, respectively. In the United States, households need a compensation of only about 0.92%, which amounts to about \$390 given an annual pre-government income of about \$42,700 [15]

Table VIII: Welfare gains of the tax and transfer system

	λ
United States	0.92%
Sweden	4.54%
Germany	5.14%

Note: λ denotes the percentage increase of per-period consumption that is necessary to make households indifferent between the pre-government income stream and the (level-adjusted) 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 IX: 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.

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 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—in Sweden, the distribution of post-government earnings changes is even symmetric on average. For the United States the described pattern is only very weak:

[14]Dollar values take 2010 as the base year. The conversion from Euro (SEK) to U.S. dollars uses the annual average of daily exchange rates for 2010.

[15]Remember that the household income measures are adjusted for household size.

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 would imply a kurtosis of 3).

Table IX: Moments of $\Delta_s y$

	Std.Dev	L9010	Skew	Kelly	Kurt	L9050	L5010	Nobs
				United States				
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				
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				
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 5% for Germany.

References

Guvenen, F., Ozkan, S. and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, **122** (3), 621–660.

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.

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

Appendix A: Data

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

PSID

Variables Definition

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 and Hours

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[16].

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

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[17], 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.

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 VI 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

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

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[18] 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[19]. 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.

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).

[18] 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.

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

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.

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.

LINDA

[TBA]

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 individual component considered, we take all transfers received by all members of the households.

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

- *HH-level transfers subcategory 1 (labor market transfers)*: sum of unemployment benefits received by all members of household.

- *HH-level transfers subcategory 2 (family aid)*: sum of subsistence allowance of all members, + sum of unemployment assistance received by all members (up to 2004), + hh-level measure of unemployment benefits II (since 2005).

- *HH-level transfers subcategory 3 (pensions)*: sum of old-age pensions received by all members of household.

Sample Selection

In order to be in the initial sample for a year, the individual or household head must be between ages 25 and 60 and live in West Germany. In order to have a consistent sample, we drop the immigrant subsample and the high income subsample. This gives initial sample sizes of 87,582 individual-year observations for the male sample, 76,249 individual-year observations for the female sample, and 76,051 household-year observations for the household sample. 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

Appendix B: Additional Figures

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.

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

(a) United States

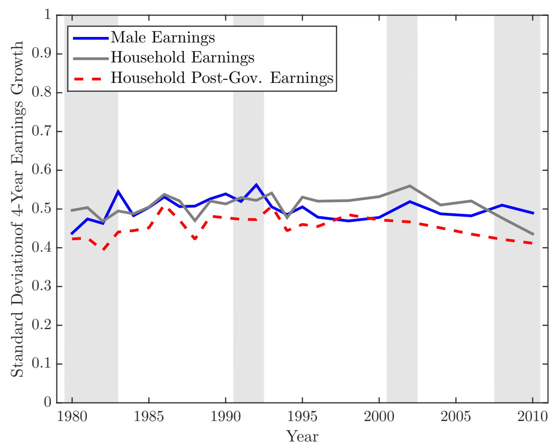

(b) Sweden

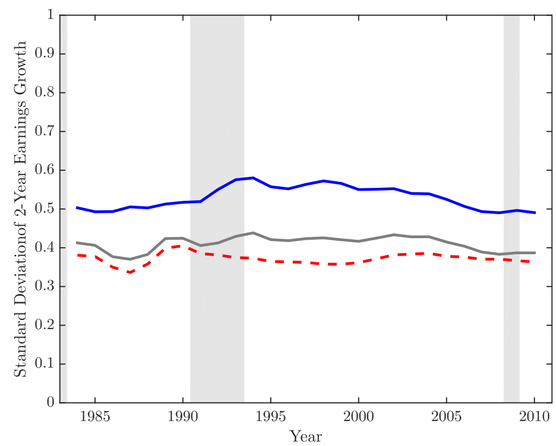

(c) Germany

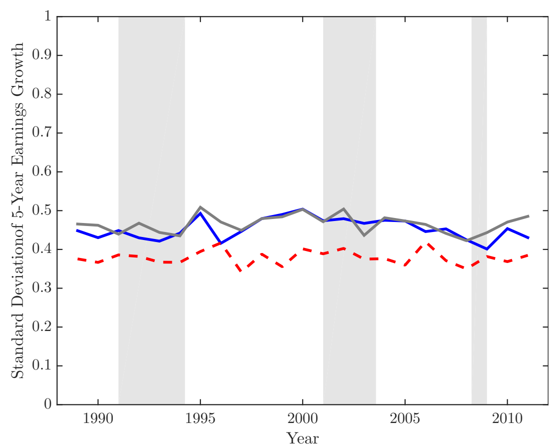

Note: Linear trend removed, centered at sample average.

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

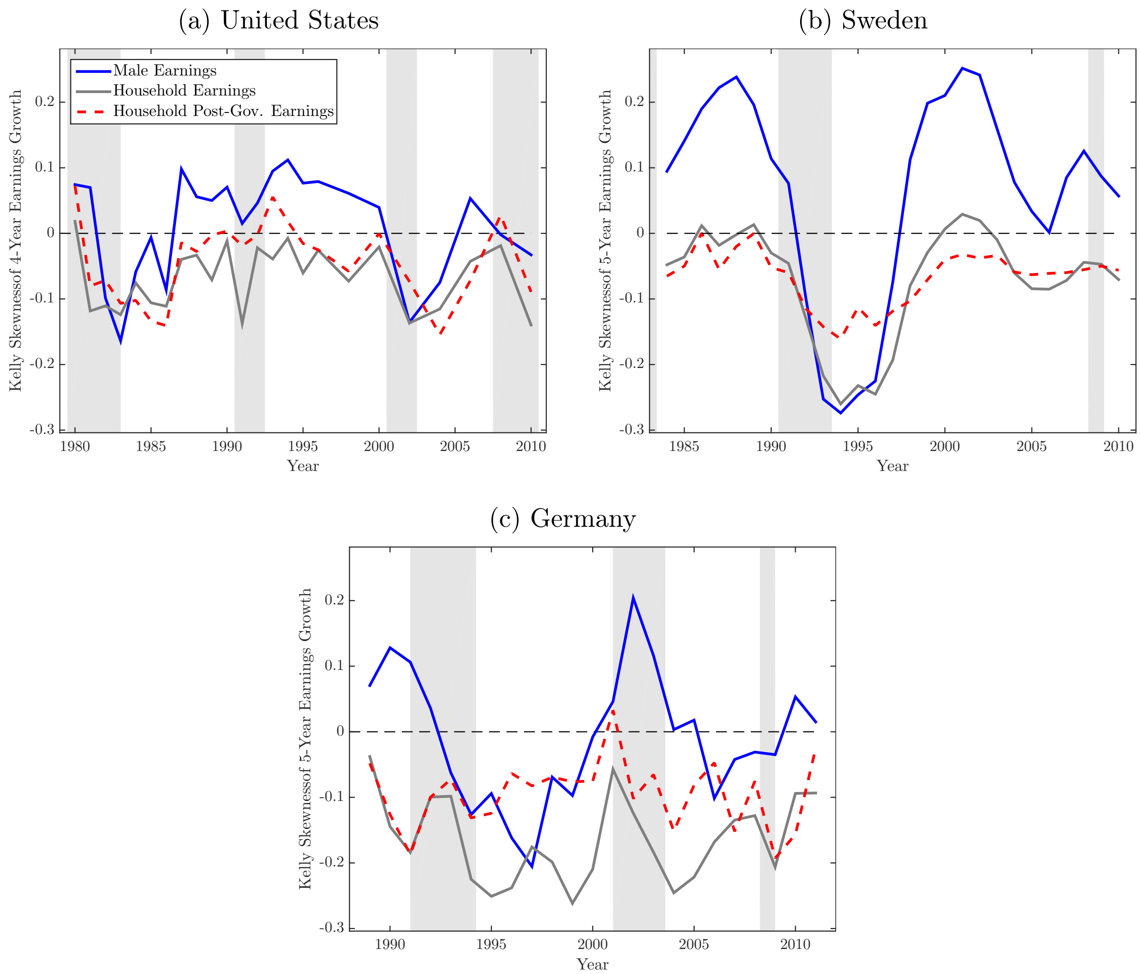

Note: Linear trend removed, centered at sample average.