

Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance*

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

May 20, 2020

Abstract

Recent studies have shown that idiosyncratic labor income risk becomes more left-skewed during recessions. This procyclical skewness arises from a combination of higher downside risk and lower chances of upward surprises during recessions. While this much is known, some important open questions remain. For example, how robust are these patterns across countries that differ in their institutions and policies, as well as across genders, education groups, and occupations, among others? What is the contribution of wages versus hours to procyclical skewness of earnings changes? To what extent can skewness fluctuations in individual earnings be smoothed within households or with government policies? Using panel data from the United States, Germany, Sweden, and France, we find four main results. First, the skewness of individual income growth (before-tax/transfer) is procyclical while its variance is flat and acyclical in all three countries. Second, this result holds even for full-time workers continuously employed in the same establishment, indicating that the hours margin is not the main driver; additional analyses of hours and wages confirm that both margins are important. Third, within-household smoothing does not seem effective at mitigating skewness fluctuations. Fourth, tax-and-transfer policies blunt some of the largest declines in incomes, reducing procyclical fluctuations in skewness.

Keywords: Idiosyncratic income risk, skewness, countercyclical risk.

*For helpful comments, we thank Nick Bloom, Alexander Ludwig, Kurt Mitman, Luigi Pistaferri, Anthony Smith, and Kjetil Storesletten, as well as seminar and conference participants at various institutions. Busch acknowledges financial support from ERC Advanced Grant (Agreement no. 324048 “APMPAL”) and from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D. Domeij acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation. Guvenen acknowledges financial support from the National Science Foundation. For the part of the analysis using DADS data, this work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the ‘Investissements d’Avenir’ programme (reference: ANR-10-EQPX-17 - Centre d’accès sécurisé aux données – CASD). Previous versions of this paper were circulated under the title “Asymmetric Business Cycle Risk and Social Insurance”.

[†]Universitat Autònoma de Barcelona, MOVE, and Barcelona GSE; chris.busch@movebarcelona.eu; www.chrisbusch.eu

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

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

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

1 Introduction

Recent empirical studies have shown that idiosyncratic labor income risk becomes more left-skewed during recessions. This rise in left-skewness arises from a combination of larger *downside risks* and smaller *upward surprises* during recessions. Put differently, the center of the earnings change distribution remains quite stable over the business cycle, whereas the upper tail compresses and lower tail expands in recessions and vice versa in expansions, resulting in procyclical skewness fluctuations. A striking example of this phenomenon could be seen during the Great Recession: between 2007 and 2009, the average decline in the labor earnings of US men was almost 7%—the largest two-year decline since the Great Depression—whereas the median change in labor earnings was +0.1%—slightly positive. The large mean decline was entirely driven by the upper and lower tails collapsing during those two years as opposed to a negative aggregate shock pulling down the entire earnings distribution ([Güvenen, Ozkan and Song \(2014\)](#)). Therefore, skewness fluctuations can potentially matter both at the micro level (i.e., the idiosyncratic risk faced by workers) and the macro level (for understanding the behavior of aggregates).

While the procyclical skewness of earnings changes has been well documented, some important questions naturally raised by these new facts remain open. In this paper, we aim to shed light on four of these related questions. First, how robust are these patterns across countries—which differ in their institutions and policies—as well as across genders, education groups, and occupations, among others? Second, what is the contribution of hourly wages versus hours worked to the procyclical skewness of earnings changes? Third, to what extent are households able to smooth the skewness fluctuations in the earnings growth of each spouse, thereby mitigating the effect on the household’s consumption and welfare? Fourth, and finally, how effective are government social insurance policies (i.e., the tax-and-transfer systems) in smoothing skewness fluctuations over the business cycle?

To address these questions, we use five panel datasets on earnings histories from four different countries, which collectively provide the information necessary for the empirical analysis. The bulk of our analysis focuses on three countries—the United States, Germany, and Sweden—which differ in important dimensions relevant for our analysis, such as household structures, the tax-and-transfer systems, and labor market institutions, among others. The datasets we use are based on Social Security records (the Sample of Integrated Labour Market Biographies, or SIAB, for Germany), tax

register data (the Longitudinal Individual Data Base, or LINDA, for Sweden), and household surveys (the Panel Study of Income Dynamics, PSID, for the United States and the German Socio-Economic Panel, SOEP for Germany), covering more than three decades in each country. We complement the main analysis with another administrative panel dataset from France (Declaration Annuelle des Donnees Sociales, DADS), which has more detailed information on hours and wages, allowing us to shed more light on the relative contribution of each to the skewness fluctuations in earnings changes.

Our analysis yields four results. First, starting with before-tax-and-transfer (“gross”) individual earnings growth, we find that skewness is robustly procyclical in all three countries, with substantial fluctuations from peak to trough. In fact, if anything, the fluctuations are larger in Sweden and Germany compared with the United States. To give one example, the skewness of individual earnings growth in Germany went from 0.31 in 1990, the peak year before the start of a deep recession, to -0.28 in 1994 (the trough), using the Kelley skewness measure which is a robust and convenient statistic (Figure 3). Put differently, these figures imply that, in 1990, the gap between the 90th percentile and the median (hereafter, P90–P50) of the earnings growth distribution was twice as large as the gap between the median and the 10th percentile (P50–P10), whereas this ratio had completely flipped by 1994, with the lower tail (P50–P10) growing to twice the size of the P90–P50 by 1994 (using eq. (2) below). The changes were just as large for Sweden. In contrast to these large swings in skewness, the variance of earnings growth is mostly flat and acyclical—not countercyclical as it was typically modeled in the earlier literature.

These findings both confirm the empirical evidence found by [Güvenen *et al.* \(2014\)](#) from US administrative data and show that they hold more broadly—in administrative data from two other developed economies, as well as in survey data (the PSID and SOEP). In addition, we show that this result is robust across sub-populations defined by gender, education, occupation, and private/public sector employment. Moreover, the cyclicity of skewness also holds for five-year income changes, which shows that the procyclical swings in skewness is present in the persistent component of earnings.

Second, we find that changes in hours and wages are both critical in generating the procyclical skewness in earning changes. We establish this result in several ways. Starting with Germany, while the SIAB dataset does not report work hours, daily wages can be calculated for full-time workers. Using this information, we can focus on full-time workers who are also continuously employed at the same establishment, a subsample where many potential sources of variation in hours and wages are either

absent or much more limited (e.g., unemployment, large drops in hours, changes in wages due to job changes, among others). Even for these strongly-attached workers, changes in daily wages is robustly procyclical, whereas the variance continues to be acyclical. We find the same result for Sweden by merging in extra information from LISA, a separate administrative database.

While these results clearly show that skewness fluctuations are not primarily driven by the hours margin, they pertain to full-time workers and there is no continuous measure of hours to study its cyclicity directly.¹ Thus, to bring more direct evidence we use the DADS, based on French Social Security records, which reports information on work hours and wages for all workers. We find that changes in both wages and hours are strongly procyclical with similar magnitudes to each other (Table III). Looking at subsamples, in the sample of strongly-attached workers (same as defined above), procyclical skewness is almost entirely driven by changes in wages, with the skewness of hours changes showing no cyclicity. The pattern is partially reversed for the rest of the baseline sample, for whom skewness is procyclical for changes in both wages and hours, but the latter is twice as volatile as the former. Collectively, these separate strands of evidence all point to procyclical skewness as a robust property of fluctuations in changes in individual earnings, wages, and hours.

Third, moving from individual earnings to household earnings, we did not find any evidence indicating that households are able to mitigate the higher in downside risk during recessions in each spouse’s individual earnings. For example, comparing the cyclicity of the earnings growth of actual households to synthetic households that are formed by randomly pairing unrelated men and women shows that the procyclicality of skewness is not any lower for actual households. We have also studied the response of spousal earnings to changes in the head’s earnings to see if there was any evidence that the larger downside risk and smaller upside surprises in recessions triggered a stronger spousal response. We have not found evidence of such a response despite examining households across the entire earnings distribution and earnings changes across the distribution. This is not completely surprising given that spouses are facing the same labor market conditions as the heads during recessions, so they are likely to have difficulty increasing their work hours or finding a second job.

Fourth, moving from gross to disposable (or post-government) household income,

¹The PSID and the SIAB are not suitable for the purposes of this analysis because teasing out the *cyclicity* of the *skewness of changes* in a variable essentially involves triple-differencing the data, which exacerbates the measurement error in reported hours in survey data (which is a more severe problem for hours than annual earnings—see Bound *et al.* (2001) for evidence from validation studies).

we find that the tax-and transfer system reduces the procyclicality of skewness in all three economies. In the United States and Sweden, the elasticity of Kelley skewness with respect to GDP growth is about half as large for the post-government household earnings measure compared with its pre-government counterpart. However, this similar effect on skewness in the two countries are driven by different sources: In the United States, the tax-and-transfer system mainly reduces the cyclicity of the lower tail, whereas the opposite is true in Sweden—the major effect is on the upper tail, which becomes acyclical, with a smaller effect on the lower tail.² We also unbundle the components of the tax-and-transfer system and find differences in the effectiveness of each component in each country. Overall, we conclude that the tax-and-transfer system plays an important role in reducing the magnitude of procyclical fluctuations in the skewness for households. Our analysis does not address the costs of the tax-and-transfer system, which should clearly be weighed against any potential benefit. Furthermore, the reduced procyclicality of skewness in some cases comes from the reduced procyclicality of the upper tail, partly achieved through progressive taxation.

We have also examined the extent of business cycle fluctuations in the fourth moment—the kurtosis—of earnings changes but did not find large and robust cyclical patterns. That said, one aspect of kurtosis matters greatly for evaluating the effects of skewness fluctuations. Basically, earnings changes are highly leptokurtic—they have long and fat tails—which interact with, and amplify, the effects of skewness fluctuations to generate a large rise in idiosyncratic risk in recessions.

The paper is organized as follows. The next section discusses the data sources, and Section 3 describes the empirical approach. Section 4 presents the results for gross individual income for the three countries. Section 5 zooms in on various groups in the population, and presents the results on the wages versus hours margin. Section 6 expands the analysis to households and post-tax-transfer income. Section 7 concludes.

Related Literature

Earlier empirical work in the literature was limited by the small sample size and time span of the available survey-based panel datasets, such as the PSID, leading researchers to make parametric assumptions to obtain identification. One common

²As we discuss further in Section 6.2, the results for Germany were mixed. On the one hand, in the SOEP data, the skewness of post-government household earnings changes is essentially acyclical. On the other hand, the SOEP data also shows some important differences from SIAB data in cyclicity patterns for individuals, which raises some uncertainty about the reliability of this result for Germany.

assumption is that shocks to earnings are Gaussian, which implies zero skewness. Restricting attention to the changes in the mean and variance of income shocks, [Storesletten, Telmer and Yaron \(2004\)](#) concluded that the variance of income shocks in the US data is countercyclical.³

[Guvenen *et al.* \(2014\)](#) revisit this question using a large panel data set on the earnings histories of US males from SSA records. The large sample size allowed them to relax parametric assumptions as well as to examine variations in skewness. They found that the variance of income shocks is stable over the business cycle and is robustly acyclical, whereas the skewness of shocks varies significantly over time in a procyclical fashion. The current paper goes substantially beyond their analysis by studying two new countries and four datasets, shedding light on the contribution on hours versus wages, moving beyond before-tax-and-transfer individual earnings to analyze household earnings with various levels of government provided social insurance, among others.

[Busch and Ludwig \(2020\)](#) adapt the parametric approach of [Storesletten *et al.* \(2004\)](#) to allow for skewness fluctuations and analyze the cyclicity of labor income risk in the United States. They come to the same substantial conclusion as we do, namely, that variation of income risk over the business cycle is asymmetric. In ongoing work, [Angelopoulos, Lazarakis and Malley \(2019\)](#) follow the approach in [Busch and Ludwig \(2020\)](#) to study the cyclicity of higher-order risk in the United Kingdom using panel data from the British Household Panel Survey. They confirm the same finding of strongly procyclical skewness for the UK since the early 1990s. Similarly, [Harmenberg and Sievertsen \(2018\)](#) document procyclical skewness of individual earnings changes in administrative Danish data. In a recent paper, [Pruitt and Turner \(2018\)](#) analyze individual and household-level income dynamics using United States tax records from the IRS. They also document procyclical skewness of income changes for both male and household incomes. Unlike in our four datasets, they find countercyclical dispersion of male (not household) earnings growth.

A couple of recent papers aim at exploring the role played by hours versus wages for the observed cyclical dynamics of earnings changes. In an analysis of administrative unemployment insurance data from Washington State, [Kurman and McEntarfer \(2017\)](#) document procyclical skewness of hourly wage changes. They also explicitly show that the share of workers realizing a wage cut increases substantially in recessions. [Pora](#)

³Using a similar approach, [Bayer and Juessen \(2012\)](#) studied the cyclicity of the variance in Germany, the UK, and the US, and different patterns in Germany and the UK relative to the US and attributed it to differences in institutions.

and Wilner (2017) document in French administrative data that the distribution of earnings changes was more negatively skewed in the 2008 recession than in the directly preceding period. Conditioning on income, they find that for high-income workers, hourly wages account for this change of the distribution, while for low-income workers hours worked are more important. Blass-Hoffman and Malacrino (2016) study data from Italian workers and find the employment margin to play an important role in driving skewness fluctuations in earnings.

Finally, our paper also contributes to a growing literature on the skewness in worker and firm outcomes (beyond labor earnings), such as in firm employment growth (e.g., Decker, Haltiwanger, Jarmin and Miranda (2015), Ilut, Kehrig and Schneider (2018), and Salgado, Guvenen and Bloom (2019)), firm productivity (Kehrig (2011) and Salgado *et al.* (2019)) and stock returns (e.g., Oh and Wachter (2018), and Ferreira (2018), and many others).

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

2 The Data

This section provides an overview of the datasets we use in our empirical analysis, the sample selection criteria, and the variables used in the subsequent empirical analyses. Further details can be found in Appendix A. Briefly, we employ four panel datasets corresponding to three different countries: the Panel Study of Income Dynamics (PSID) for the United States, covering 1976 to 2010;⁴ the Sample of Integrated

⁴The PSID contains information since 1967. We choose our benchmark sample to start in 1976 because of the poor coverage of income transfers before the 1977 wave. We complement our results using a longer period whenever possible and pertinent.

Labour Market Biographies (SIAB⁵) 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 datasets. The PSID has a yearly sample of approximately 2,000 households in the core sample, which is representative of the US population; the SOEP started with about 10,000 individuals (or 5,000 households) in 1984 and, after several refreshments, covers about 18,000 individuals (10,500 households) in 2011.⁶

The SIAB is based on administrative social security records and our initial sample covers on average 370,000 individuals per year. It excludes civil servants, students, and self-employed workers, which make up about 20% of the workforce. From the perspective of our analysis, the SIAB has two caveats: (i) income is top-coded at the limit of income subject to social security contributions, and (ii) individuals cannot be linked to each other, which prohibits identification of households. We deal with (i) by fitting a Pareto distribution to the upper tail of the wage distribution⁷ 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. In addition, for some of the analysis of individual level income dynamics we use the Longitudinal Integrated Database for Health Insurance and Labour Market Studies (LISA), which covers the Swedish population of 16 years and older individuals. In it we are able to identify annual workplace information. Furthermore, we back up the individual level analysis of earnings dynamics for full-time workers using additional data for 1995–2015 from French social security records, the Declaration Annuelle des Donnees Sociales (DADS), which is described in the appendix.

The measure of labor earnings we use is meant to be comprehensive to the extent allowed by each dataset. In all cases, it includes wage and salary income (inclusive of bonuses, overtime, paid time off, and so on) plus the labor portion of self employment

⁵We 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).

⁶These numbers refer to observations after cleaning but before sample selection. Only the representative SRC sample is considered in the PSID. The immigrant sample and high-income sample of the SOEP are not used, because they cover only subperiods.

⁷The 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.4](#) for details.

income. The earnings measure from SIAB does not include self employment income because the dataset lacks information on it. More details of each variable can be found in the data Appendix A.

For each country, we consider three samples: two at the individual level—one for males and one for females—and one at the household level. The samples are constructed as revolving panels: for a given statistic computed based on the time difference between years $t - s$ and t , the panel contains individuals who are ages 25 to 59 in periods $t - s$ and t ($s = 1$ in the case of Sweden and Germany, and $s = 2$ in the case of the United States) and have yearly labor earnings above a minimum threshold in both years. Imposing this threshold allows us to exclude individuals with very weak labor market attachment during the year and also avoids problems with zeroes when dealing with logarithms as we will see below. The threshold is set to the earnings level that corresponds to 520 hours of employment at half the legal minimum wage, which is about \$1,885 US dollars for the United States in 2010.⁸ To avoid possible outliers, we exclude the top 1% of earnings observations in the PSID and SOEP, but not in LINDA (which is from administrative sources). For each individual, we record age, gender, education, and gross labor earnings. By gross earnings we mean a worker’s compensation from his/her employer before any kind of government intervention in the form of taxes or transfers.

Furthermore, 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. In parts of the analysis of individual income dynamics, we use information on 30 occupational categories, which are listed in Appendix I for reference.

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,⁹ that satisfy the age criterion and household income that satisfies the income criteria. At the household level, we analyze pre- and post government earnings. Pre-government earnings is defined as the sum of gross labor earnings

⁸For the United States, we use the federal minimum wage. There is no official minimum wage in Sweden or Germany during this period. For Germany, we follow [Fuchs-Schündeln et al. \(2010\)](#) and take a minimum threshold of 3 euros (in year 2000 euros) for the hourly wage. For Sweden, the effective hourly minimum wage via labor market agreements was around SEK 75 in 2004 ([Skedinger, 2007](#)). For other years, we adjust the minimum wage using the growth rate in the mean real wages.

⁹In PSID and SOEP, the head of a household is defined within the dataset. In LINDA, the head of a household is defined as the sampled male.

earned by the adults in the household. Post-government earnings is constructed by adding taxes and transfers.

3 Empirical Approach

Following the recent literature on higher-order risk discussed above, our empirical approach is nonparametric and flexible. To analyze dynamics, we focus on income changes between two periods and analyze the behavior of this distribution over the business cycle. Specifically, we compute moments $m[\Delta_s y_t]$, where $y_t \equiv \log Y_t$ is the natural log of individual income, Y_t , and $\Delta_s y_t \equiv y_t - y_{t-s}$ is the change or growth rate between years $t - s$ and t .¹⁰ For Germany and Sweden, we consider $s = 1$ and 5 corresponding to short- and long-run changes, respectively. Starting with the 1997 wave, the PSID switches to a biennial structure, so we use $s = 2$ instead of $s = 1$ for the United States through the entire sample period.¹¹

Our primary measures for volatility and skewness are quantile-based, which have important advantages over standardized moments (variance and the skewness coefficient). Some of these advantages are substantive—more below—whereas others are technical or practical: they are more robust to outliers; they allow scrutinizing different parts of the distribution by varying the quantiles used; and they are often easier to interpret than the values of standardized moments. The specific moments we focus on are the log differential between the 90th and 10th percentiles (L9010) as a measure of dispersion, dispersion in the upper (L9050) and lower (L5010) tails, and the Kelley measure of skewness defined as follows:

$$\mathcal{S}_k = \frac{(P90 - P50) - (P50 - P10)}{(P90 - P10)}. \quad (1)$$

The Kelley measure of skewness has a simple interpretation. It measures the difference between the fraction of the overall dispersion, L9010, that is in the right tail, L9050, and the fraction that is in the left tail, L5010. Rearranging (1) gives a simple mapping from a given Kelley value into the fraction of overall dispersion that is in

¹⁰We repeat the main analysis in this paper using the arc-percent measure of growth, $2(Y_t - Y_{t-s})/(Y_t + Y_{t-s})$, which allows us to drop the minimum threshold requirement described above and include observations with zero income in either t or $t - s$. This makes no substantive effect on the conclusions we report in this paper.

¹¹We 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 $s = 2, 4$.

the right tail:

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

which is not possible to do with the skewness coefficient. A second advantage of Kelley as noted above is that it does not suffer from the extreme sensitivity to outliers that the skewness coefficient does. [Kim and White \(2004\)](#) provide a cautionary analysis showing that the skewness coefficient can reveal spurious relationships due to outliers found in some commonly used datasets. This is especially relevant for the survey data from PSID and SOEP that we use in our analysis.

There is also a more substantive benefit of studying quantiles directly: it can reveal a simpler underlying empirical structure or can uncover patterns that are obscured when we focus too closely on standardized moments. Two examples—which will turn out to be empirically relevant—can help illustrate these points. In the first case, suppose that a common negative shock hits all the workers who would have been in the bottom 20% of the income growth distribution, thereby reducing P20 and all percentiles below, leaving the rest of the distribution unchanged. This would register as a rise in the variance but it is not coming from a *symmetric* expansion of the distribution which is customarily associated with a higher variance. Instead, it is directly related to the distribution becoming more left-skewed, without any increase in higher percentiles.

For the second example, suppose that the same negative shock in the first example hits not only the bottom 20% but also the top 20% of the income growth distribution, reducing all percentiles below P20 and above P80, without affecting the percentiles between P20 and P80. In this case, the variance may very well remain unchanged (notice that L9010, L8020, etc., are already constant) but skewness actually becomes more negative. Even though this situation entails a large change in the distribution and a large increase in risk, the variance would not give a hint about this. Furthermore, if we were to describe the changes in the distribution in terms of standardized moments, we would characterize it as a negative shock to the first moment (the mean will fall even though the median is constant), no shock to the second moments, and a negative shock to the third moment. This description obscures the fact that there was actually only one shock that hit both tails in the same direction, and what looked like three separate shocks to three moments—the fall in the mean and skewness and the constant variance—are actually all the consequence of this one tail shock.

Defining Business Cycles

We use two main indicators for business cycles. The first one is based on the official classification of peaks and troughs by the National Bureau of Economic Research (NBER) for the United States and by the Economic Cycle Research Institute (ECRI) for Sweden and Germany.¹² It is well known, however, that some key macroeconomic variables do not perfectly synchronize with expansions and contractions, but their fluctuations might still have an impact on earnings. For example, the US stock market experienced a significant drop in 1987, officially classified as an expansion year, and indeed the skewness of household income growth dips in that year (Figure 2a). Other examples (e.g., 1996) are easy to find for Germany and Sweden. To better capture these more continuous changes in aggregate conditions, we use the (natural) log growth rate of GDP over s years— $\Delta_s(\log GDP_t) \equiv \log(GDP_t) - \log(GDP_{t-s})$ and regress each moment m on a constant, a linear time trend, and this indicator of business cycles:

$$m(\Delta_s y_t) = \alpha + \gamma t + \beta^m \times \Delta_s(\log GDP_t) + u_t. \quad (3)$$

The key parameter of interest is β^m , which measures the cyclicity of moment m . For a quantitative interpretation of the results reported in the next sections, Figure 1 reports annual log GDP growth for each country.

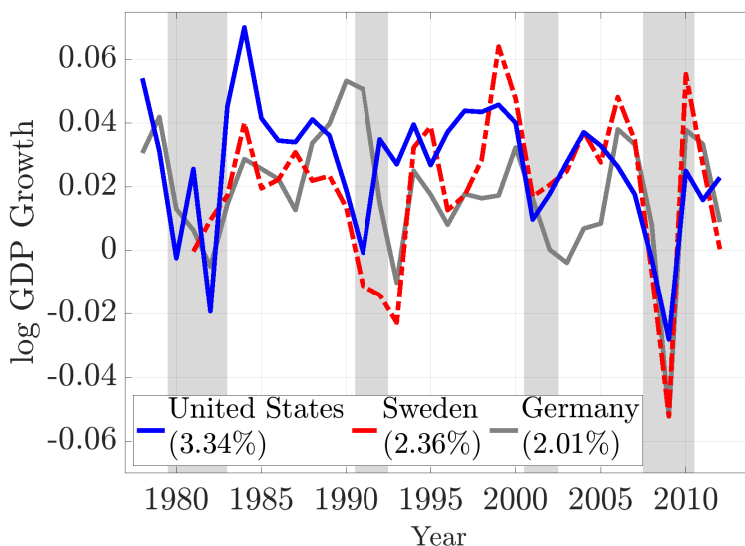
4 Empirical Results: Before-Tax Individual Income

We start our empirical analysis with labor income at the individual level, measured before taxes and government transfers. This is the same income measure used in recent work that found the skewness of the US income growth distribution to be volatile and procyclical and its dispersion to be flat and acyclical (e.g., Guvenen *et al.* (2014)). In this section, we ask three questions that are left unanswered by this earlier work.

First, we ask if these cyclical features are specific to the United States or whether they are robust features of business cycles that are also observed in other countries whose labor markets differ greatly from that in the US. To give an example of these

¹²We make two adjustments to NBER and ECRI classifications. For the US, we classify the 1980–1983 period as a single “double-dip” recession instead of two separate ones. For Sweden, unlike ECRI, we classify the 2001–2003 period as a recession because Swedish GDP fell by a similar magnitude to that in the US and Germany during these years, as seen in Figure 1.

Figure 1: Annual log GPD Growth: US, Germany, and Sweden



Note: The figure plots log GDP growth for each country. The shaded areas indicate US recessions. The numbers in parentheses next to each country indicate the standard deviation of the short-run GDP growth series over the period 1976-2010, where short-run is one-year difference for Germany and Sweden, and two-year difference for the United States, to be consistent with the timing of the PSID data used in our analysis. The corresponding number for the one-year change in the United States is 2.8%. The series for Germany corresponds to West Germany up to and including the 1990–91 change and to (unified) Germany from 1991–92 on.

differences, consider the fact that only 10.7% of US workers are unionized and only 11.9% are covered by trade union agreements, whereas the corresponding fractions are 18.1% and 57.6%, respectively, in Germany, and 67.3% and 89%, respectively, in Sweden.¹³

A second question we ask is whether the patterns of cyclicity found by [Guvenen et al. \(2014\)](#) using US administrative earnings data from the Social Security Administration (SSA) are also borne out in the PSID (survey data), which is widely used in the analysis of income dynamics. The answer is less than obvious because earlier papers that used the PSID and adopted a Gaussian parametric econometric model (which restricts the skewness to zero) found a strongly countercyclical variance of shocks (e.g., [Storesletten et al. \(2004\)](#)). This raises the question: is it the differences in the datasets or in the methodologies (or both) that accounts for these different conclusions? By removing parametric restrictions, our analysis can shed light on this question.

Finally, because of the limited number of covariates available in the SSA data, [Guvenen et al. \(2014\)](#) were not able to explore potential variation in the cyclicity patterns by gender, education, occupation, private versus public sector employees,

¹³OECD (2016). The reported numbers are for 2013.

among others, which we address starting in this section.

4.1 Cyclicalities of Variance and Skewness

We begin in Figure 2 with a simple time-series plot of the standard deviation and the skewness coefficient of the short-run income change distribution for male workers in the United States (biennial), Germany (annual) and Sweden (annual).¹⁴ We start with standardized moments because of their familiarity before we delve into the analysis of quantiles. Recessions are indicated as shaded areas. Two key patterns are clearly evident here. First, in all three countries, the standard deviation varies little over time and the small fluctuations it displays do not typically comove with the business cycle. In contrast, the skewness coefficient shows significant procyclical fluctuations, with skewness dipping consistently in recessions and recovering in expansions. Hence, Figure 2 provides a visual confirmation of the procyclical skewness/flat dispersion pattern in both US Survey data as well as in Germany and Sweden.

Next, to quantify the degree of cyclicalities and compare it across countries, we use the regression framework described above in (3). Table I reports the cyclicalities coefficient, β^m , for four quantile-based moments—L9010, Kelley Skewness, L9050, and L5010—separately for each gender and the three countries.¹⁵ Starting with the United States, the coefficient on L9010 is quantitatively small, slightly negative for men and slightly positive for women, and statistically insignificant with t -stats less than 1.4 for both genders. These estimates confirm the acyclicalities of dispersion in the United States that we saw in Figure 2. The same acyclicalities is also seen in the bottom two panels for Germany and Sweden with small point estimates that are statistically insignificant.¹⁶

Cyclicalities of Skewness

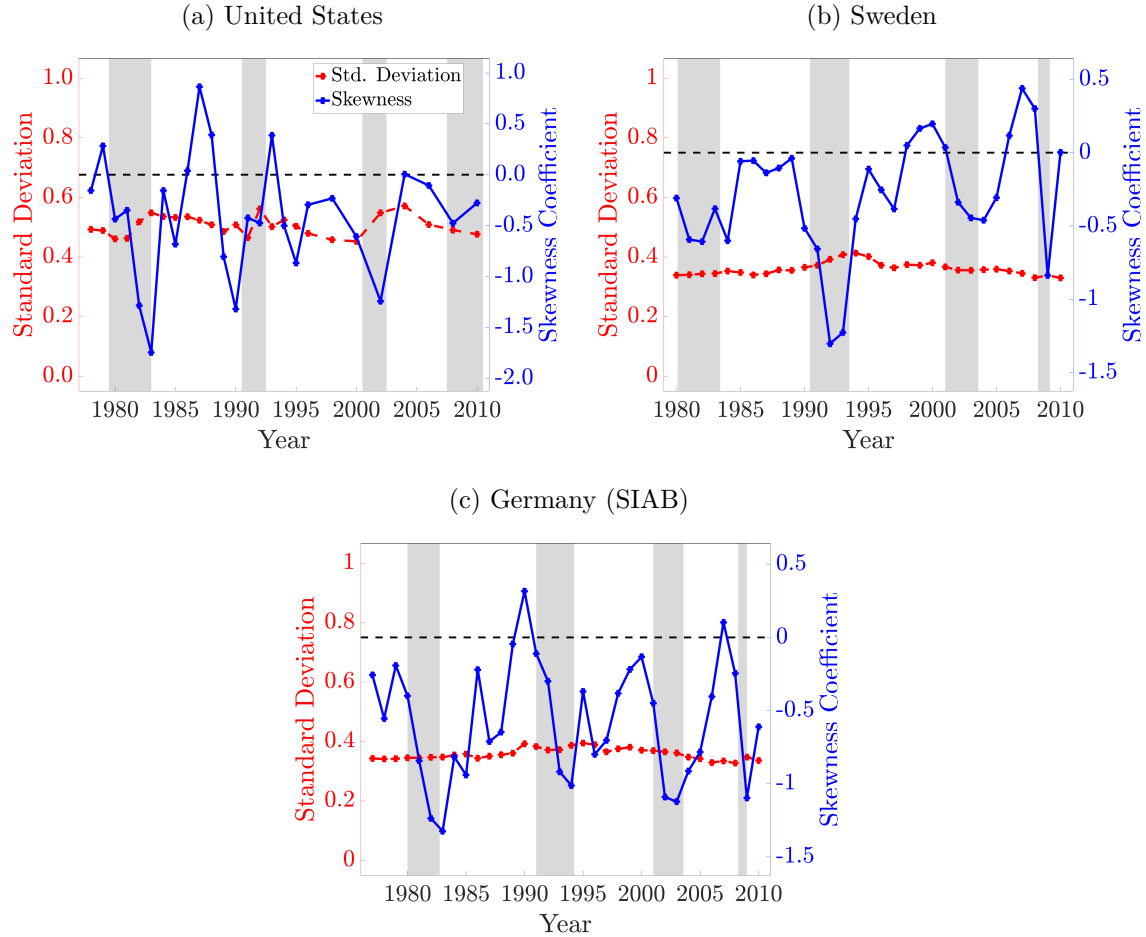
We next turn to the cyclical behavior of skewness. Starting with Figure 3 (left panels), the Kelley skewness for males shows the same procyclical pattern as the skew-

¹⁴The skewness coefficient of random variable x is the third standardized moment: $E(x - E[x])^3 / \sigma^3$, where σ is the standard deviation.

¹⁵We ran two alternative versions of these regressions and obtained the same substantive results. First, we used the arc-percent change rather than log change of income to capture the extensive margin—or zeroes in income (Table C.1 in Appendix C.1). Second, we use a dummy for recessions as a business cycle indicator rather than log GDP change, in the regression (Table C.2 in Appendix C.2). In both cases, we find the same substantive patterns described here.

¹⁶All regression results in the paper based on SIAB data are robust to various sensitivity checks we conducted to address issues of topcoding and a structural break in the wage variable. See Appendix F for details.

Figure 2: Standard Deviation and Skewness of Short-Run Income Growth: Males



Note: Linear trend removed, centered at sample average. Shaded areas indicate recessionary periods (see footnote 12). Year denotes ending year in the growth rate calculations.

ness coefficient in Figure 2, which is probably not surprising but still reassuring.¹⁷ In the PSID, Kelley skewness drops significantly during the 1980s double-dip recession, falling from 0.15 for the 1979–80 change to below -0.2 for the 1982–83 change, as well as the two recessions in the 21st century. There is no drop in skewness during the early 1990s recession, which may be due to potential data issues during the transition PSID went through from 1992 to 1993 or it may be due to the somewhat unusual timing of this recession which appears as two dips in economic activity and skewness in the SSA data analyzed by [Güvenen *et al.* \(2014\)](#).

The synchronization between Kelley skewness and the business cycles is even clearer

¹⁷To reduce the number of figures for readability, we have moved the analogous figure for females to the appendix

Table I: Cyclicalities of Log Annual Income Change Moments: Before-Tax/Transfer Individual Income

	L9010	Kelley	L9050	L5010
United States				
Males	-0.54 (-1.38)	2.25 (4.79)	0.68 (2.49)	-1.23 (-4.27)
Females	0.40 (1.39)	1.17 (3.01)	0.86 (2.57)	-0.47 (-2.38)
Sweden				
Males	-0.26 (-0.64)	3.64 (3.94)	0.78 (4.51)	-1.04 (-2.50)
Females	0.33 (1.84)	1.77 (2.64)	0.65 (2.91)	-0.32 (-1.99)
Germany (SIAB)				
Males	0.15 (0.36)	5.48 (5.80)	0.95 (3.14)	-0.80 (-4.11)
Females	0.34 (0.48)	2.55 (2.05)	0.80 (1.25)	-0.46 (-1.80)

Note: Each cell reports the cyclicalities coefficient β^m (on log GDP change) in a regression of the moment specified in the column header on log GDP change plus a constant and a time trend (eq. 3). Newey-West t -statistics are in parentheses.

in Sweden and Germany (middle and bottom left panels of Figure 3). In particular, Kelley skewness falls significantly during the early 1990s recession, which was much deeper in these countries compared with the United States. In Germany, the Kelley measure swung from 0.31 in 1989–90 down to -0.28 in 1993–94, implying a dramatic shift in the length of the tails: Whereas the upper tail (L9050) accounted for two-thirds ($0.5 + 0.31/2 \approx 0.66$) of the overall L9010 gap in the 1989–90 period, with the remaining one-third accounted for by the lower tail (L5010), these ratios completely flipped by the end of the recession (1993–94), with L5010 growing to account for almost two-thirds (64%) and L9050 shrinking to one-third of L9010. Notice that Kelley skewness falls from 1995 to 1996, which is technically an expansion year for Germany but the GDP growth did in fact fall between those years (see Figure 1). Finally, Sweden experienced a similar but slightly smaller drop with the Kelley going from 0.08 to -0.31 between

the same two years (with the share of the lower half, L5010 rising from 46% to 66%).

In Table I, the cyclicity coefficients for males in column 2 are all positive and statistically significant at the 0.1% level, confirming the strong procyclicality of Kelley skewness. The estimated β^{Kelley} for males is 2.25 for the US, 3.64 for Sweden, and 5.48 for Germany, implying a 2.5-fold larger fall in Kelley in Germany than in the US for the same 1% slowdown in GDP growth. This result is somewhat surprising given the higher prevalence of unions and other worker protection measures in Germany and Sweden relative to the United States, so we will analyze it in greater detail in the next section.¹⁸

To give a quantitative interpretation to these coefficients, consider a two standard deviation decline in log GDP growth in Sweden, swinging from one standard deviation above average to one standard deviation below, which represents a moderate recession. With the estimated $\beta^{\text{Kelley}} = 3.64$, Kelley skewness will fall by $3.64 \times (2 \times 0.0236) \approx 0.17$. For the sake of discussion, if the upper tail to lower tail ratio was 50/50 in an expansion, it would fall to 42/58 in a recession. A severe recession with a four standard deviation swing in GDP growth (such as the 1990–1993 period) would bring the upper-to-lower tail ratio from 50/50 to 33/67. These are very large changes in the relative size of each tail over just a few years, especially in a country like Sweden whose institutions are geared toward social insurance.¹⁹ Finally, skewness is also procyclical for female workers in all countries, with positive and statistically significant coefficients for Sweden and the US at 1% level, and for Germany at 5% level.

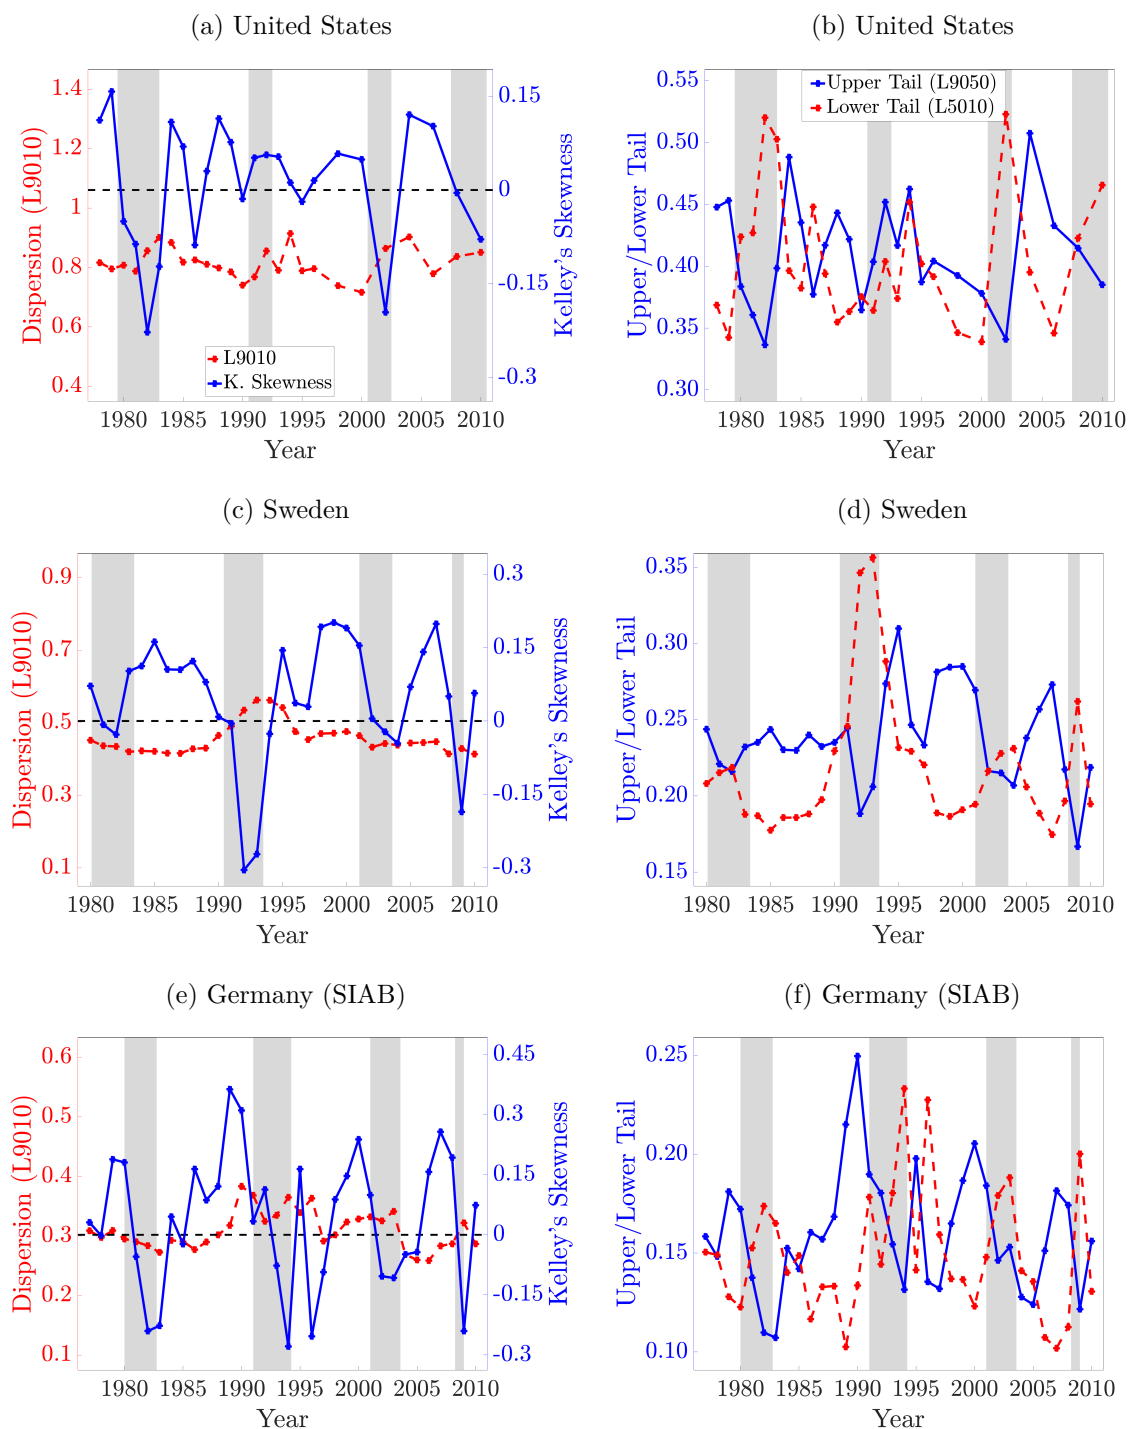
4.2 Inspecting the Tails

Skewness can become more negative from either the compression of the right tail or the expansion of the left tail or both. Each tail is informative about different aspects of labor market outcomes: for example, the compression in the right tail could result from a decline in upward moving opportunities (smaller wage gains with promotions or job changes), whereas the expansion in the left tail is likely to result from larger downside risk (higher likelihood of job losses, increased duration of unemployment, and so on). Furthermore, the government policies that we study below have different effects on each tail. All of these lead to the question: What is the contribution of each tail to the procyclical fluctuations in skewness? And how does this contribution vary across these three countries?

¹⁸Running the regression with the skewness coefficient instead of Kelley measure yields very similar results.

¹⁹The corresponding changes in \mathcal{S}_k for the U.S and Germany are 0.15 and 0.22, respectively.

Figure 3: L9010, Skewness, and Tails of Short-Run Income Growth: All Males



Note: Linear trend removed, centered at sample average. Shaded areas indicate recessionary periods (see footnote 12). Horizontal gray line in the right axis of the left panel indicates zero (symmetry) reference line. Year denotes ending year in the growth rate calculations.

The right panels in Figure 3 plot L9050 and L5010 over time. While the magnitudes somewhat differ, in all three countries both tails contribute significantly to the procyclical skewness. In particular, L9050 starts falling while L5010 starts rising right around the beginning of the recession and they reverse the roles with the start of the expansion. The last two columns of Table I report the cyclicity coefficients for the two tails, which are positive for L9050 and negative for L5010, confirming the pattern we see in Figure 3. The statistical significance of the estimated coefficients is fairly high for men (t -stats between 2.49 and 4.51) and somewhat lower but still significant for women (ranging from 1.80 to 2.91, with the exception of L9050 in Germany with a t -stat of 1.25).

Another point to notice in Table I is that, for all countries, the estimated β 's for each tail are of similar magnitudes to each other. For example, for Sweden, the coefficient for L9050 is 0.78 and for L5010 it is -1.04 . The corresponding coefficients are 0.68 and -1.23 for the US, and 0.95 and -0.80 for Germany. Thus, the shrinking of one tail is largely offset by the expansion of the other tail, making total dispersion, the L9010, move very little over the cycle. As a result, skewness becomes more negative in recessions without any significant change in the variance. One partial exception is also illuminating: L9010 rises slightly during the 1990s recession in Sweden and Germany (less so in the US) because the left tail expands more than the right tail contracts. So, the rise in dispersion is in fact due to a change that is mostly asymmetric in nature, which would not have been apparent by focusing on the variance alone.

These new insights and more nuanced interpretations of income risk over the business cycle underscore the importance of the finer-grain analysis through quantiles undertaken here compared with the simpler analysis of a few standardized moments. In particular, interpreting changes in the variance without considering the changes in skewness delivers an incomplete picture that can be highly misleading.

Turning to the estimates for females in Table I, we observe the same patterns of cyclicity as those of men, whenever the coefficient is significant. In particular, L9050 is procyclical for the US and Sweden, whereas L5010 is countercyclical for all three economies (though only significant at the 10% level for Germany). That said, the magnitudes of coefficients are smaller for women, especially for Kelley skewness, which is largely driven by the much smaller coefficients on L5010 compared with men (about 1/3 that of men's in the US and Sweden and about 1/2 in Germany). In other words, compared with men, the right tail compresses during recession in a comparable fashion, whereas the expansion of the lower tail—or the rise in downside risk—is much smaller.

We will return to this finding when we analyze households earnings.

4.3 Persistence of Skewness Fluctuations

It is well understood that the economic implications of transitory income changes are very different from those of persistent changes. Hence, a natural question is the extent to which the procyclical fluctuations in skewness pertain to the persistent component of earnings. To fix ideas, consider the standard permanent-transitory model of earnings dynamics:

$$y_t = z_t + \varepsilon_t$$

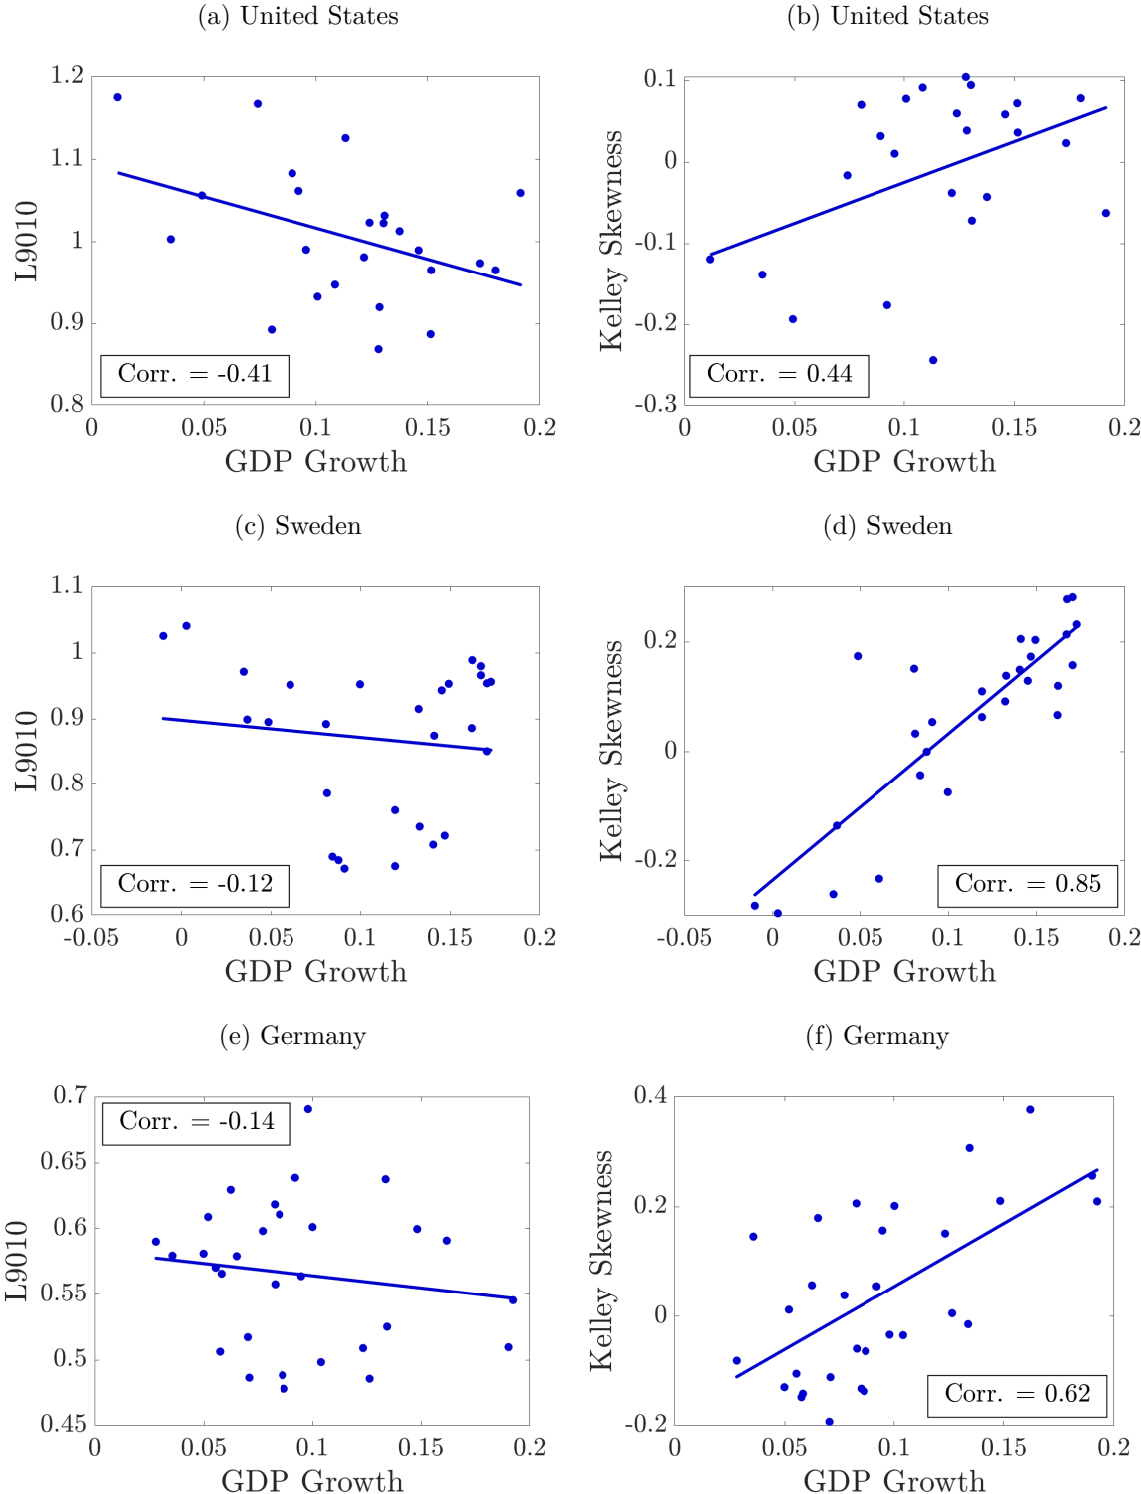
$$z_t = z_{t-1} + \eta_t$$

where η_t and ε_t are zero-mean disturbances, and z_t and ε_t represent the permanent and transitory components, respectively. The s -year difference of log income is $y_t - y_{t-s} = \sum_{j=1}^s \eta_{t-j} + \varepsilon_t - \varepsilon_{t-s}$, which contains s permanent innovations and always 2 transitory ones, so longer-term changes increasingly reflect the properties of permanent shocks. Thus, to investigate the persistence of skewness fluctuations, in this section we study five-year changes for Germany and Sweden and, given the biennial nature of the PSID after 1997, four-year changes for the United States. That said, regressions that use overlapping long-term changes face serious econometric problems in sample sizes found in time series data.²⁰

With these issues in mind, we use more transparent graphical constructs to analyze the properties of persistent changes. Starting in Figure 4, each panel shows a scatterplot of either L9010 or Kelley skewness of longer-run earnings changes for males against five-year log GDP growth. The patterns are fairly easy to discern. For Sweden and Germany, the scatterplots of L9010 are clouds showing no evident relationship with GDP growth as confirmed by the flat fitted line. For the United States, there is some evidence of a downward slope, which is partly attributable to the outlier on the left top corner. The scatterplots for Kelley skewness reveal a stronger positive

²⁰For example, if five-year changes are computed for every year of the sample, the overlap between observations induce strong serial correlation, which makes the autocorrelation consistent standard errors of coefficients to be downward biased, inflating the significance of estimates coefficients (e.g., [Richardson and Stock \(1989\)](#)). This can be an empirically serious problem, for example as has been recognized in the literature on stock return predictability regressions (e.g., [Kirby \(1997\)](#) and references therein). Using only non-overlapping observations reduces the already modest sample size dramatically. We did estimate the cyclical regressions using five-year changes and found the same patterns but do not include them because of the concerns outlined here.

Figure 4: Cyclicity of Dispersion and Skewness of Long-Run Income Changes, Males: United States, Sweden, and Germany (SIAB)



Note: Each figure is a scatterplot of either the L9010 or Kelley skewness of five-year earnings change against five-year log GDP change (four-year change used for the United States).

relationship with GDP growth, which is especially strong in Sweden and Germany.²¹ Appendix G shows the same figures for women, which are again qualitatively telling the same story. It also shows the time-series of moments of five-year income changes.

Additional Evidence from Sub-Populations

We bring additional evidence on the persistence of skewness fluctuations by recognizing that time series data on the entire population is also panel data on subpopulations, and in this particular case, on occupational groups. We conduct this analysis using the SIAB dataset from Germany, the other datasets are either too small to allow this finer grain analysis (the PSID and SOEP) or lack information on occupations (in our version of LINDA)

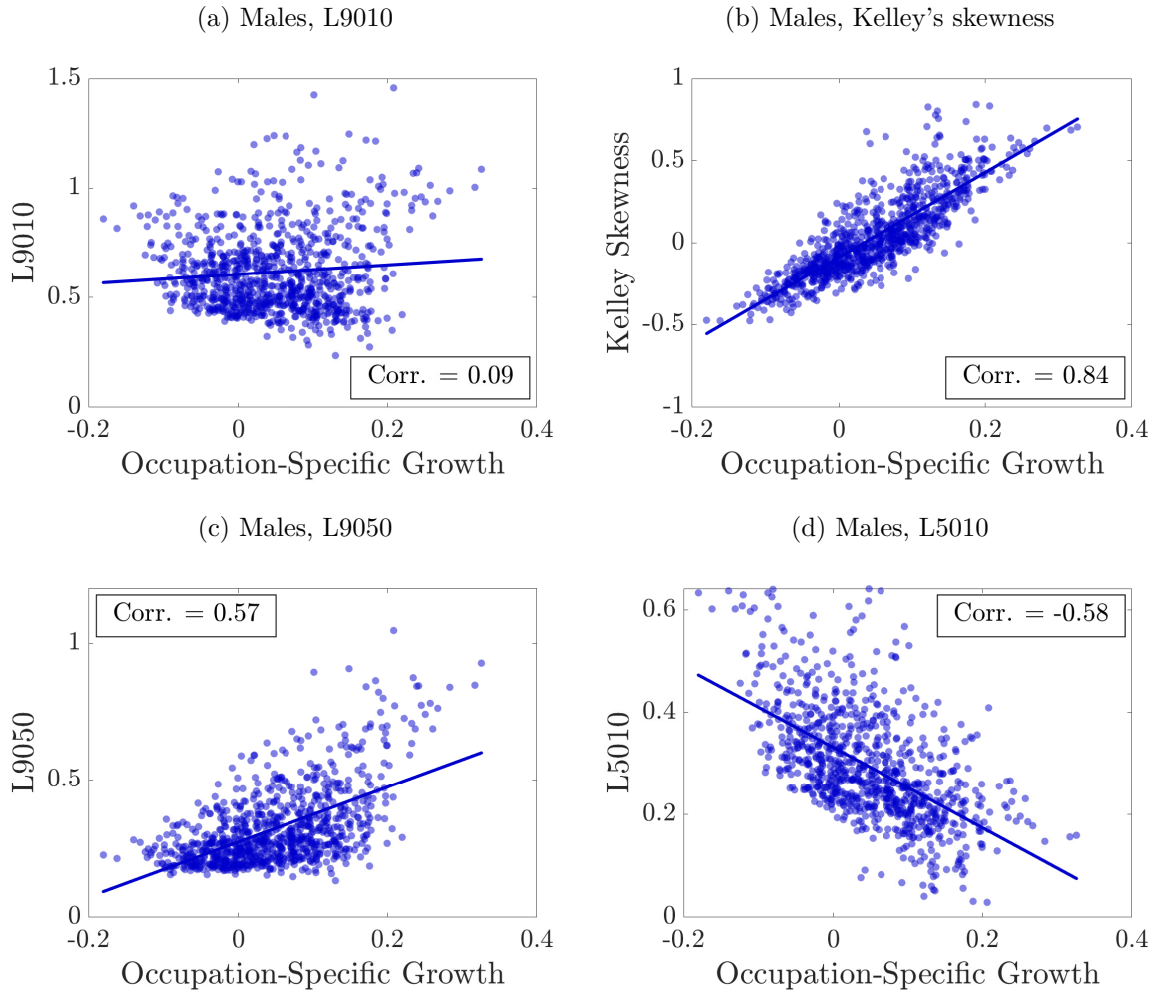
In SIAB, we assign each worker to one of 30 occupational categories in year t based on their occupation in $t - 5$. We compute the same moments of five-year changes as before but now individually for each occupation group. We also construct a business cycle indicator for each occupation by taking the five-year change in average earnings in that occupation. The top left panel of Figure 5 shows the scatterplot of L9010 for each occupation-year cell against average earnings growth for the same cell, which basically shows no relationship, confirming the acyclical nature of dispersion found above. In contrast, the scatter plots for Kelley skewness in the top right panel shows a very clear upward pattern, with substantial range of variation in the magnitude of Kelley skewness (in the y-axis). The bottom two panels make clear that both tails are individually strongly cyclical with L9050 showing a somewhat larger range of variation over the occupation-specific cycle than L5010.

These conclusions are not sensitive to using occupation-specific cycles. Table G.1 in Appendix G reports the raw correlations of each moment with five-year *aggregate* GDP growth. For Kelley skewness the correlations are all positive, with a correlation of 0.49 even at the 10th percentile of correlations. In contrast, the correlations for L9010 range from -0.35 at the 10th percentile to 0.29 at the 90th percentiles with a median of -0.12 .

Overall, these findings corroborate our main results by showing that the same

²¹In an earlier draft of this paper, we have also estimated a more formal econometric process for earnings dynamics featuring permanent and transitory shocks, targeting a large number of moments of short- and long-run earnings changes. The estimated process revealed a strong procyclical variation in the skewness of the permanent component. Similarly, Busch and Ludwig (2020) estimate earnings processes using moments of the cross-sectional income distribution, allowing for state-dependent distributions of income shocks. They find systematic variation of cross-sectional skewness, which can be attributed to procyclical skewness of the persistent component.

Figure 5: Distribution of Five-Year Income Growth; Occupation-Specific Cycles (SIAB): Males



Note: Scatterplot of moments of five-year earnings change against occupation-specific average income growth over the same horizon. There are 900 data points: 30 five-year changes for 30 occupations each.

patterns we observed in the aggregate economy hold, more strongly, at the more disaggregated level. The patterns for females look qualitatively the same; see Appendix G for details.

5 Digging Deeper into the Main Findings

In this section, we extend our analysis of individual earnings in two directions. First, we examine the robustness of our findings in different subgroups of the population, defined by educational attainment, by private/public sector employment, and

occupation. Second, we ask the degree to which the procyclical fluctuations in skewness are explained by changes in hours worked or to changes in wages, or both.

5.1 Heterogeneity across Groups of Workers

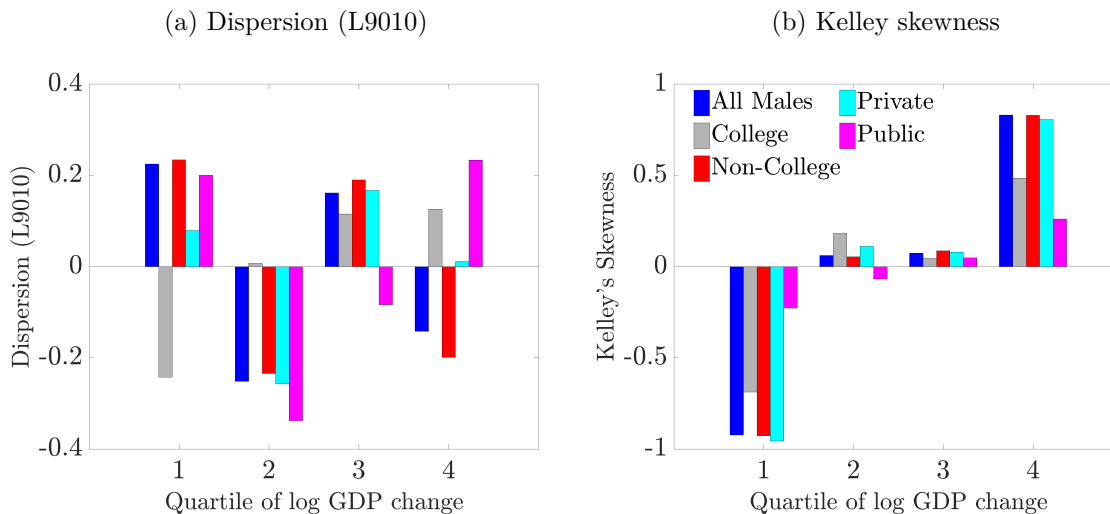
Education and Public vs. Private Sector

We begin by classifying workers (separately for each gender) by educational attainment (college versus non-college graduates) and, separately, by whether they hold a private- or public-sector job. The share of male workers who are college educated is 12%, 16%, and 25%, respectively, in Germany, Sweden, and the United States (the analogous numbers for women are 8%, 17% and 25%). Differences in the size of public sector employment are even larger and also vary significantly between men and women.²² Moreover, public sector jobs are often thought of as less risky, offering generous employment protection and less volatile compensation, so it is interesting to ask if this perception is actually borne out in the data.

To avoid producing too many different figures, we pool the statistics from the three countries as follows: For a given group, we first construct a statistic, say L9010, and log GDP growth for each country-year pair, and assign the statistic to its corresponding quartile of the log GDP growth distribution (pooled over all years), and average within each quartile. Figure 6 shows the L9010 and Kelley skewness for males. The standardization of moments and log GDP changes is performed independently for each country before pooling across countries, which implies that a deviation from zero indicates a standardized deviation from the country-specific mean of the moment. For each quartile, the bars correspond to the average moment for (ordered from the left) the full sample, college graduates, non-college graduates, private employment, and public employment, respectively. Figure 6 shows that the nature of income risk is qualitatively similar across all male subgroups: overall dispersion is acyclical (panel a), whereas Kelley skewness is strongly procyclical (panel b). Furthermore, as Figure B.1 in Appendix B shows, the upper tail is procyclical, and the lower tail is countercyclical. The results for females look qualitatively the same (Figure B.2).

²²For men, the share of public jobs is 23% in Sweden, and 10% and 13% in Germany and the United States. For women, the corresponding figures are 63%, 36% and 18%. For these statistics, we define public sector employment as jobs in public administration, health care, and education (sectors which in Germany and Sweden are dominated by public sector jobs or by jobs funded by the public). Historically, most workers in these sectors were employed by the public; this is less true today.

Figure 6: Higher-Order Moments by Quartiles of Log GDP Change: Males



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

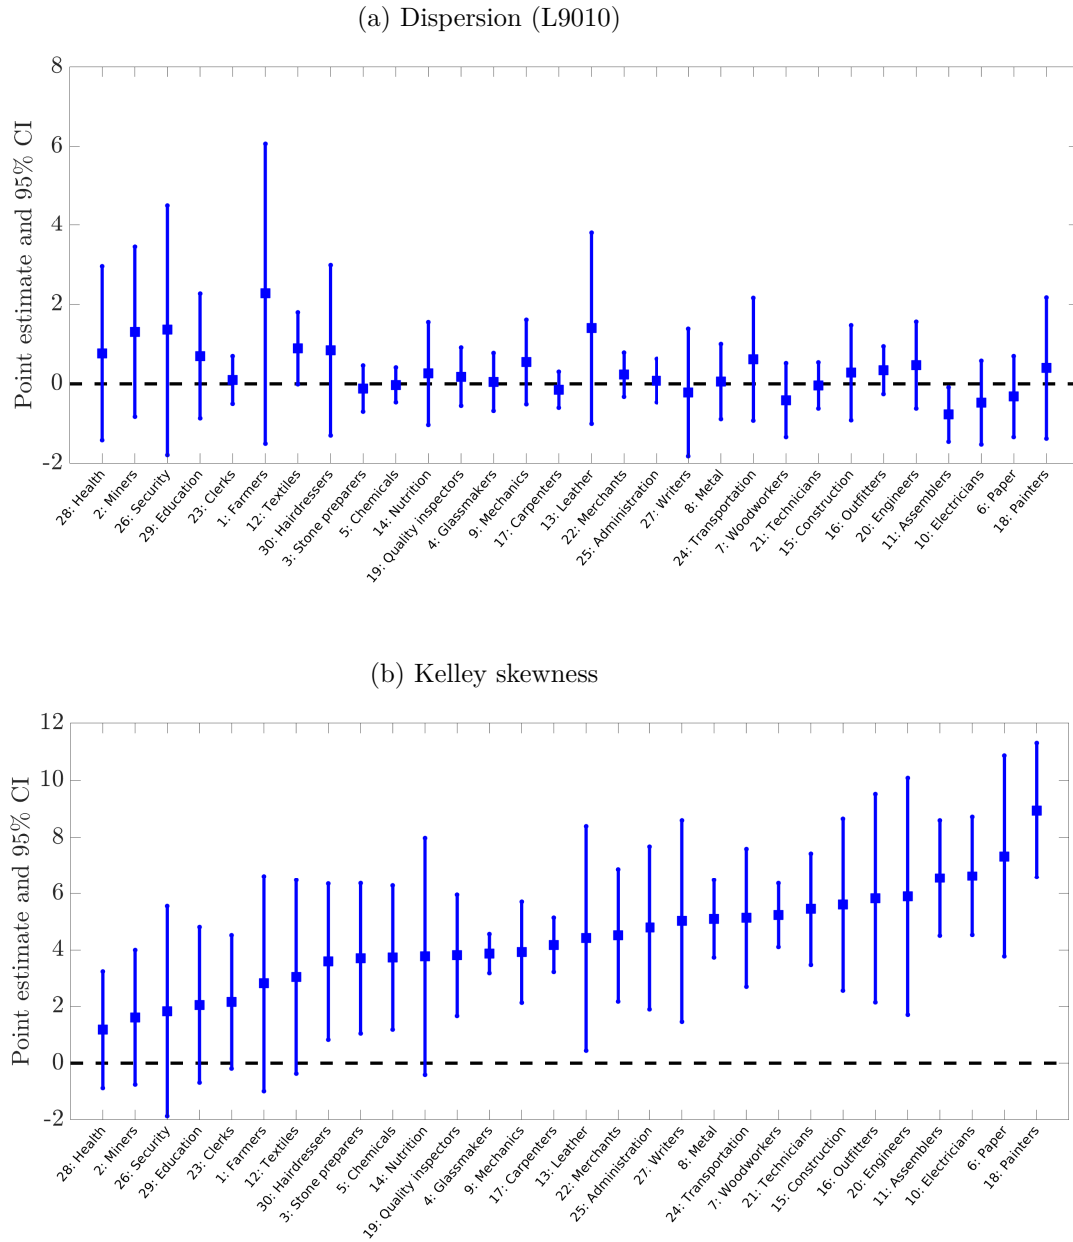
Occupational Groups

We return to the occupational groups in SIAB data for Germany, analyzed in the last section, and focus on annual—rather than five year—changes, which allows us to run the cyclicity regression in (3) separately for each occupation without running into the overlapping observations problem (see footnote 20). Figure 7 shows the estimated β 's for L9010 and Kelley skewness for each occupation. As seen in the bottom panel, the estimated β 's for Kelley skewness are positive for every occupation and statistically significant for the vast majority of them. As before, β 's for dispersion are close to zero for the vast majority of occupations and is not statistically significant for any of them. Further results for the upper and lower tails are in Appendix B.2.

5.2 Earnings versus Wages

A workers' earnings can change because of a change either in hourly wages or in hours worked or a combination of both. So, an important question is to understand whether the fluctuations in the skewness of earnings growth is driven by wages or hours or both. Reliable data on hours worked is scarce because it is often unavailable in administrative data sets (such as the US SSA data, which prevented Guvenen *et al.* (2014) from addressing this question) and measurement error in survey data is a more severe problem for reported hours than for annual earnings (see Bound *et al.* (2001) for a review of evidence from validation studies). This problem is exacerbated by the

Figure 7: Dispersion and Skewness of Short-Run Income Growth by Occupation: Males (Germany (SIAB))



Note: Separate regressions for each of 30 occupation segments. Each marker reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. The confidence bands are based on Newey-West standard errors (maximum lag length considered: 3).

fact that teasing out the *cyclical*ity of the *skewness* of *changes* in a variable essentially involves triple-differencing the data, which amplifies the measurement error in the underlying data.

In this section, we shed light on this question using SIAB for Germany and also bringing additional evidence from two datasets that we did not use so far in the analysis. We start with SIAB, which contains information on the duration of each employment spell and on whether it is a part-time or full-time job. Next, we perform a comparable analysis for Sweden. In particular, we go beyond our baseline dataset LINDA and look at the LISA dataset, which covers the whole population and which has a focus on individuals' labor market experiences rather than on family and transfers. Of main relevance for our analysis is that it has workplace (establishment) information. Neither SIAB nor LISA contain direct information on hours worked. We therefore complement this analysis using French social security data (the DADS) from 1995 to 2014, which contains hours worked for each employment spell as reported by employers (see Appendix D for a description of the data).

Full-Time Workers in Germany and Sweden

We first look at workers with stable employment relationships. To accommodate the different structures of SIAB and LISA, we perform slightly different but comparable analyses. In both datasets, we focus on subsamples of workers whose earnings dynamics are not driven by changes in the extensive margin.

In the SIAB, we define a full-time worker if her full-time spells add up to at least 50 weeks of employment in a given year. (A less strict definition of full-time workers as 45 weeks of employment does not change the results.) The wage variable is the average daily wage rate, where the average is taken over all full-time spells during the year. This is the same measure used in [Dustmann *et al.* \(2009\)](#) and [Card *et al.* \(2013\)](#). We consider the annual change in the average daily wage rate of male workers who are in the full-time sample in both years.

For completeness, the first row of Table II reproduces the estimated β 's for the baseline sample from Table I. Row 2 reports the corresponding β 's using average daily wages for full-time workers instead of annual earnings of all workers. Notice how similar the coefficient on skewness is compared with the baseline sample in the first row (4.73 versus 5.48). Note that 88% of males (73% of women) are in the full-time sample.²³ Naturally, the dispersion of earnings changes is wider than that of wage

²³The sample of full-time female workers contains about 73% of women (who make up only 54%

Table II: Cyclicity of Log Annual Income Change Moments for Males: Individual Income vs. Daily Wages, Germany (SIAB) and Sweden (LISA)

	L9010	Kelley	L9050	L5010
Germany				
Earnings	0.15	5.48	0.95	-0.80
<i>(Baseline Sample)</i>	(0.36)	(5.80)	(3.14)	(-4.11)
Daily Wages	-0.09	4.73	0.30	-0.39
<i>(Full-Time Workers)</i>	(-0.54)	(6.31)	(3.77)	(-3.20)
Daily Wages	-0.12	4.98	0.28	-0.40
<i>(Establishment Stayers)</i>	(-0.81)	(5.78)	(3.29)	(-3.20)
Sweden				
Earnings	-0.06	3.64	0.87	-0.94
<i>(Baseline Sample)</i>	(-0.73)	(4.34)	(4.30)	(-3.81)
Earnings	-0.12	1.78	0.16	-0.28
<i>(Establishment Stayers)</i>	(-4.53)	(7.43)	(5.34)	(-7.90)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3). Full-Time are those that work full time for at least 50 weeks in both years for which the change is calculated.

changes, which is reflected by the point estimates on the tails (last two columns), which are about half as big for wage changes. In the third row, we further restrict the sample by selecting workers who not only work full time but also work at least 50 weeks at the same establishment in two consecutive years. For these workers not only changes in hours but also changes in daily wages should be smaller than for the previous sample.²⁴ Perhaps surprisingly, the estimated β coefficients, including the one on skewness (4.98), barely change.

For the question at hand, there are two shortcomings of LISA relative to SIAB. First, we cannot identify the duration of job spells in LISA, and second, we can only look at total annual earnings, not at daily wages. Still, we can select a sample of workers

of the observations) that contribute to the measures of earnings changes for women. The corresponding figures are 88% of individuals and 82% of observations for males. This implies that part-time employment plays a more important role for the female sample.

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

with minimal room for the extensive margin of labor supply to affect their earnings changes. We do this by selecting workers who earn income from the same establishment in four consecutive years, from $t - 2$ to $t + 1$, and have that establishment as their main employer in $t - 1$ and t . Clearly, this is a more selected group of stayers than the one in SIAB.

The second panel of Table II shows the corresponding estimation results for Sweden. The first row shows the estimates for the full population covered by LISA, which are virtually identical to the estimates based on LINDA. The second row shows the results for the workers staying at their establishment. The coefficient on skewness is about half the size of row 1 but continues to be very significant. An intermediate conclusion is thus that the overall dynamics are not exclusively driven by the extensive margin in either Germany or Sweden. Taken together, this points in a similar direction as recent evidence by Kurmann and McEntarfer (2019), who document in data from Washington that during the Great Recession the incidence of nominal wage cuts for job stayers increased substantially—accompanied by systematic reductions in hours worked, which further decreases earnings.

Hours versus Wages: Additional Evidence from France

While the spell data for Germany allows us to explore the roles of days worked vs. changes in daily wages, part of the variation in daily wages can potentially be attributed to changing hours worked during the day. We thus consider those variables separately, using the same regression framework for France.²⁵ Table III shows results for earnings, hours, and hourly wage changes for males. First, earnings changes display the same patterns we have seen so far, with a cyclicity coefficient on Kelley skewness that is even higher (7.38) than any of the three other countries (Table I), while L9010 is acyclical with small and statistically insignificant coefficients in column 1.

Second, as seen in the second and third rows, the skewness of both changes in hourly wages and hours worked display significant procyclicality with coefficients of 3.37 and 4.27, respectively. However, one important difference between the two is seen in the tails: whereas the cyclicity coefficient for L9050 are similar for wages and hours (0.80 and 0.72), the left tail of the wage growth distribution is much less countercyclical (-0.22) than that of hours (-1.02). This is consistent with the downward rigidity of

²⁵Given the available data from the DADS, we use the years 1995–2015 in the analysis, which gives 20 years for which we can estimate one-year changes. The standard deviation of log GDP growth over that time period is 0.98%.

Table III: Cyclicity of Hours Worked vs. Hourly Wages; France (DADS): Males

	L9010	Kelley	L9050	L5010
Baseline Sample				
Earnings	-0.28 (-0.58)	7.38 (7.60)	1.37 (4.27)	-1.65 (-5.4)
Hourly Wages	0.58 (1.55)	3.37 (2.78)	0.80 (2.49)	-0.22 (-1.2)
Hours Worked	-0.3 (-0.57)	4.27 (4.19)	0.72 (5.04)	-1.02 (-2.23)
Subsample A: Full-Time, Establishment Stayers				
Earnings	-0.04 (-0.14)	5.43 (2.58)	0.46 (1.79)	-0.50 (-2.28)
Hourly Wages	0.47 (1.35)	4.96 (2.63)	0.74 (2.15)	-0.27 (-1.83)
Hours Worked	0.47 (0.98)	-1.69 (-0.22)	0.03 (0.11)	0.44 (0.64)
Baseline Excluding Subsample A				
Earnings	-0.23 (-0.18)	4.63 (5.55)	2.94 (4.33)	-3.16 (-3.24)
Hourly Wages	0.4 (1.04)	2.33 (2.40)	0.77 (2.60)	-0.38 (-1.16)
Hours Worked	-0.02 (-0.01)	4.71 (5.93)	3.13 (3.81)	-3.15 (-3.58)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in the indicated 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). *Full-time Establishment Stayers* are those workers working in Full-time employment for the same establishment for at least 50 weeks in both $t - 1$ and t . *Baseline without Full-time* are those workers who are not in 50 weeks full-time employment in either $t - 1$ or t .

wages, making hours a more elastic margin to adjust for employers. Overall, this evidence confirms that the procyclicality of skewness is driven both by wages and hours.

To gain further insights, we split the baseline sample into full-time workers that stay in the same establishment and the rest of the baseline sample. The results in the middle

and bottom panels of Table III. Earnings changes display strong procyclicality but slightly smaller than the baseline ($\beta^{\text{Kelley}} = 5.43$); however, almost all of it is now due to wages (4.96) and almost none from hours (-1.69 and statistically insignificant). Results are partially flipped for the rest of the sample in the bottom panel: the skewness of earnings changes is still procyclical but now a larger component is coming from hours, which is both volatile and very cyclical (bottom row). Overall, the additional evidence from France confirms and complements our results from Germany and Sweden. Both wages and hours play significant roles in generating skewness fluctuations in earnings. For more strongly attached workers, wages play a more important role and display substantially procyclical skewness driven more by the upper tail, whereas the opposite pattern emerges for less strongly attached workers.

6 Introducing Insurance

So far, our analysis focused on individual labor earnings before taxes and transfers and documented how idiosyncratic risk as measured by this variable varies over the business cycle. While this is an important first step, many questions economists ultimately care about are more directly linked to consumption, which is separated from individual gross earnings by several layers of implicit or explicit insurance. In this section, we study two of these broad sources—insurance within the household and from government social insurance policies—to gauge the extent to which they mitigate downside idiosyncratic risks in recessions.

6.1 Within-Family Insurance

In Table IV, the first row of each panel reports the estimated β 's for the same moments but now using household earnings, which can be compared to their counterparts for individual earnings in Table I. For the United States and Sweden, the cyclical patterns for households are essentially the same as for individuals: procyclical skewness, with each tail's movements almost perfectly canceling out each other, leaving P9010 acyclical. As for magnitudes, the estimated coefficient for Kelley skewness of households falls in between the coefficients reported for males and females in Table I. For example, in the US, $\beta^{\text{Kelley}} = 2.25$ for males and 1.17 for females versus 1.91 for households here, which is not too surprising since the latter combines each spouse's earnings.²⁶ But comparing these coefficients does not tell us much about the

²⁶The comparison between individual and household earnings is less informative for Germany because the former in Table IV uses SIAB data whereas the latter is based on SOEP. It turns out

Table IV: Cyclicalities of Earnings Growth Moments: Actual vs. Synthetic Households

	L9010	Kelley	L9050	L5010
United States				
Actual households	0.04 (0.15)	1.91 (6.57)	0.81 (5.93)	-0.78 (-3.78)
Synthetic households [†]	-0.01 (-0.03)	1.59 (3.88)	0.72 (3.00)	-0.73 (2.52)
Sweden				
Actual Households	-0.02 (-0.08)	2.24 (3.33)	0.50 (4.94)	-0.52 (-2.00)
Synthetic households	-0.24 (-0.83)	1.93 (3.33)	0.35 (3.23)	-0.59 (-2.19)
Germany (SOEP)				
Actual households	-1.17 (-3.33)	1.79 (2.76)	-0.03 (-0.12)	-1.15 (-4.22)
Synthetic households	-0.97 (-3.34)	0.98 (2.09)	-0.17 (-1.06)	-0.80 (3.48)

Note: Each cell reports the coefficient on log GDP change in the cyclical regression (3). Newey-West t -statistics are included in parentheses. [†]Synthetic households are formed by randomly assigning two workers of opposite genders from the sample conditional on certain observables. For the US and Germany, the observables are age and education; for Sweden, the observables are age, region, and average income (binned). The reported parameters are the means of 250 bootstrap estimates, which are also used to compute standard errors.

the extent of smoothing that happens within households. In particular, we want to understand whether each spouse *actively* responds to the earnings shock of their partner (i.e., the added worker effect), and more importantly, whether this response helps dampen the business cycle fluctuations in tail shocks and skewness. In other words, our main focus is not so much on the *level* effect of spousal response but on whether this response changes over the business cycle in a way that mitigates the larger tail shocks in recessions.

To shed light on this question we begin by creating a control group of “synthetic households,” whose composition mimic the baseline sample but in which synthetic spouses have no actual connection to each other and therefore, unlike actual households, cannot respond to each others’ earnings shocks. Thus, to the extent that within-

that even for individuals, the coefficient for Kelley skewness is quite a bit smaller in SOEP than in SIAB (e.g., 1.55 versus 5.48 for men) and L9050 is acyclical for individual earnings as well. So, we need to be cautious in comparing the estimates of β from SOEP to SIAB directly, although SOEP will still be useful for other analyses below.

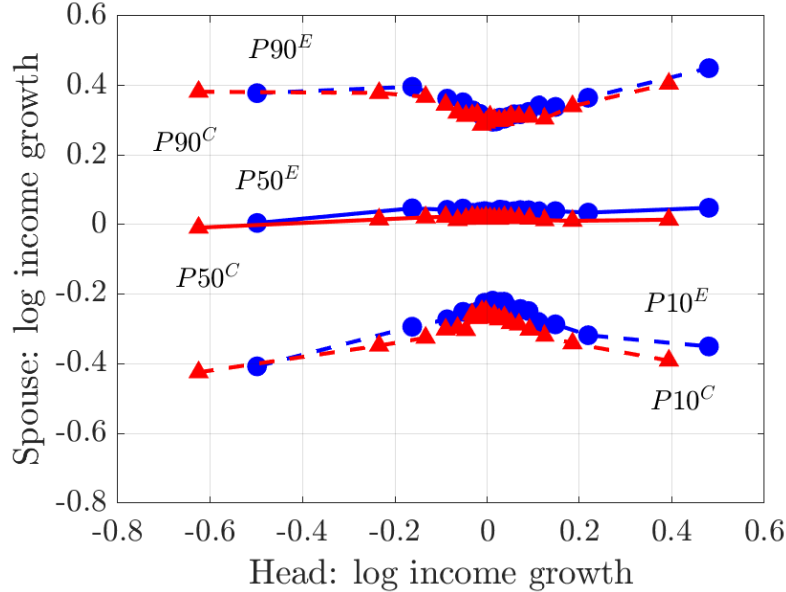
household insurance is present, the cyclicalities for actual households should be smaller than for this control group. We construct this control group by taking each head of household and drawing a synthetic spouse from a subsample of the baseline sample with observable characteristics similar to that of his actual spouse. Specifically, for the US and Germany, we condition on age (seven groups) and education level. For Sweden, we control for age (three groups), region (capital, high-density and low-density regions), and 5-year average income. The pairing is done separately for each $(t - 1, t)$ time period.

The second row of each panel in Table IV reports the cyclicalities coefficients for these synthetic households. Perhaps surprisingly, we see no evidence of within-household insurance. For example, in Sweden, the coefficients for Kelley skewness are 1.93 and 2.24 for synthetic and actual households, respectively. If we were to go strictly by the point estimates, in all three countries the skewness for actual households seems to be more procyclical than two randomly-paired individuals.²⁷ One possible explanation would be the presence of highly correlated shocks: for example, a regional economic shock will hit both spouses of an actual household who live together but not spouses in randomly paired households unless they are not formed by conditioning on region. Similarly, to the extent that couples sort on other job or labor market characteristics—such as industry, firm, education, etc.—their income shocks will have common components. As just discussed, we have conditioned on some of these characteristics when forming synthetic couples to partially control for some of these common shocks, which makes the lack of apparent insurance even more surprising, while still leaving open the possibility that the common component could be based on some other characteristics.

This apparent lack of within-household insurance against idiosyncratic business cycle risk at the population level does not preclude the possibility of such insurance being present within subsets of the population. To investigate this possibility, we take a finer grain approach that requires a larger sample size than what is available in the PSID or SOEP, so we focus on the Swedish LINDA dataset for this analysis. To allow the magnitude of spousal response to vary by household earnings, we first sort households based on their average earnings over the previous five years and split them into three groups: the bottom quintile, the top quintile, and the combined middle three quintiles (P20 to P80). For each group, we sort households by the head’s log annual earnings change, and group them into twenty equally sized bins. Then for

²⁷The point estimate for L5010 for Sweden is slightly smaller (in absolute value) for actual couples than for random couples, however, this difference is not statistically significant.

Figure 8: Spousal Earnings Response to Head’s Earnings Change over the Business Cycle in Sweden: Households with Earnings between P20 and P80



Note: Figure shows spouse’s log earnings growth against household head log earnings growth for households with five-year average earnings between the 20th and 80th percentiles of the distribution. For each marker, the x -axis shows the median earnings growth of heads in that five-percentile wide bin and the y -axis shows the 90th, 50th, or 10th percentile of spouse log earnings growth, the of the corresponding quantile of head earnings growth. Red and blue markers correspond to recession and expansion years, respectively.

each of these twenty groups, we calculate the 10th, 50th, and 90th percentiles of the distribution of spouses’ log earnings growth during the same period. Figure 8 shows the plots for the middle-income household group (P20 to P80). The slope of the spousal response percentile lines tells us about the correlation between spouses’ earnings growth rates. A (large) negative correlation—which would indicate the presence of spousal insurance—would manifest itself as a (large) negative slope, especially in the range where head’s earnings growth is negative. The interpretation is reversed for a positive slope. Because our main interest is in the insurance channel, we focus our discussion on the left half of the figure where head’s earnings growth is negative.

Before getting into the business cycle patterns, let us first discuss the broad patterns we see here. First, the median spousal response is quite flat, which is consistent with the extant literature that focused on the average size of the added worker effect and found it to be small. Having said that, there is also a wide range of spousal responses and they display some systematic variation, with the P90 and P10 lines drawing a

Table V: Summary of Spousal Response

<i>Spousal Response Percentiles</i>						
	P10	P50	P90	P10	P50	P90
<i>Head's earnings change:</i>	Expansion			Recession		
P0–P10	0.34	0.13	0.05	0.20	0.06	–0.01
P10–P50	0.40	–0.06	–0.59	0.43	0.01	–0.40
	Recession – Expansion					
P0–P10	–0.14	–0.07	–0.06			
P10–P50	0.04	0.08	0.19			

Note: Each cell in the top-left panel reports the slope of the fitted line to each spousal response percentile indicated in the column header (P10, P50, P90) over the range of the head's earnings changes in each row (P0–P10 and P10–P50) shown in Figure 8 during expansions. Other panels are interpreted analogously.

bowtie-like shape.

For small to medium size negative changes (on the x-axis), the P90 line is downward sloping, indicating a positive spousal response that offsets part of the decline in head's earnings. To see the magnitudes more clearly, Table V reports the slope of different segments of each line in this figure. For example, during an expansion, for earnings changes of the head that fall between the 10th percentile and the median, the slope of P90 is -0.59 , which corresponds to a substantial spousal response of 0.59 (log) percent for a one (log) percent drop in head's earnings. Notice that these numbers do not imply that household earnings will only fall by 0.41 log percent in this scenario because the two spouses' initial earnings does not have to be the same. We will return to the effect on household earnings below.

At the other end, P10 is sloping upward in the same earnings change range for the head, with a slope of 0.40 (Table V, first column, second row). So for these households, a 1 log percent drop in head's earnings coincides with a 0.40 log percent additional drop in spouse's income. Putting the two pieces together, spouse's earnings systematically vary with changes in head's earnings but not always in a way to offset the change in head's earnings. This is not surprising for reasons discussed above. For example, spousal insurance might be in response to truly idiosyncratic shocks whereas amplifying responses might represent a common component to the shocks of both spouses. Finally, spousal responses to tail shocks have the same sign but are muted compared to what we saw for smaller shocks.

So, how do these spousal response patterns change over the business cycle? First, comparing the lines of recessions to those for expansions in Figure 8, the slowdown in earnings growth for both spouses is easy to see (red lines are shifted both down and to the left). In terms of magnitudes, the bottom panel of Table V reports the changes in slopes (for the same segments of each line discussed above) from recessions to expansions. There are a couple of main takeaways. First, for small- to medium-size negative changes in head’s earnings (P10–P50), the median spousal response goes from a small insurance (slope: -0.06) in expansions to effectively zero (0.01) in recessions, for a net change of 0.08 . However, there is a larger change in the tails of the spousal response distribution with the P90 response falling from -0.59 in expansions to -0.40 in recessions. This is a 19 log percent drop in the spousal offset over the business cycle at the top end. In contrast, the change in P10 is very modest.

Second, for negative tail shocks to the head (P0–P10), the slope of the spousal response lines get flatter, although still positive but small, meaning that spousal earnings growth is less positively correlated in recessions than in expansions. Finally, the counterparts of Figures 8 for households in the top and bottom quintiles show the same qualitative patterns discussed here (see Figure H.2 in Appendix H.1). Overall, the important takeaway from these numbers is that for all but the largest negative shocks to head’s earnings, spousal insurance *declines* in recessions, which is consistent with the highly cyclical skewness in households earnings growth we found above.

While this analysis tells us about spousal response, it does not directly tell us about how household earnings change because that also depends on the relative share of household earnings coming from each spouse. In Appendix H.1, we present a slightly altered analysis, which instead displays the conditional distribution of household income changes.

Finally, our results do not imply that there is no insurance against individual income risk within households. Rather, they show that the channels available to couples to mitigate individual (downside) risk become weaker when the aggregate economy is in a contraction. This is in line with evidence in Pruitt and Turner (2018), who document in tax data from the United States that recessions are times in which female earnings growth is only weakly correlated with male earnings growth, and female labor supply adjustments along the extensive margin in reaction to male earnings losses are less pronounced (i.e., the “added worker effect” is weaker).

6.2 Government: Taxes and Social Insurance Policy

Broadly speaking, government policies affect household earnings through two main channels: one is through income taxation and the other is through transfers or social insurance policies. Here, we first analyze the total effect of these policies by comparing the cyclicity of post-government (i.e., post taxes *and* transfers) household earnings growth and compare with the cyclicity of pre-government (or gross) household earnings growth we documented in the previous subsection. We then disaggregate taxes and each transfer component and assess the role of each policy separately.

The Overall Effect of the Tax and Transfer System

Table VI reports the cyclicity coefficients for pre-government (first row) and post-government (second row) household earnings. Figure H.1 in Appendix H.1 complements this table by plotting the time series of these moments for these two earnings measures. The general pattern we see in the table shows that taxes and transfers have a nontrivial effect on reducing business cycle fluctuations in skewed idiosyncratic risk. Starting with skewness, β^{Kelley} is about half the size for post-government household earnings growth compared with their pre-government counterparts in the United States and Sweden, and shows essentially no cyclicity in Germany.

Looking at each tail separately, the effect of government policies on the lower tail is greatest in Germany (with β^{L5010} shrinking from -1.15 to -0.18) followed by the United States (shrinking from -0.78 to -0.21), with the smallest effect seen in Sweden (-0.52 to -0.38). The ordering is reversed for the upper tail, with β^{L9050} falling the most in Sweden (0.50 to -0.03), followed by the United States (0.81 to 0.55), with the smallest effect seen in Germany where L9050 appears acyclical even in pre-government earnings. (As noted earlier, this is one aspect of SOEP data that differs substantially from SIAB, so this last piece of evidence should be interpreted with care.) Putting these pieces together, we conclude that tax and transfer policies reduce skewness fluctuations by dampening the cyclicity of the lower tail in Germany, the upper tail in Sweden, with the United States falling in between the two. So even though government policies reduce skewness fluctuations, they achieve this by affecting very different parts of the earnings growth distribution.

Unbundling Government Taxes and Transfers

We consider three subcomponents of government transfers that are comparable across countries and are consistently measured in each country over time. These components are: (1) labor-market-related policies, (2) aid to low-income families, and (3)

Table VI: Cyclicalities of Household Annual Earnings Growth Moments: Total Effect of Taxes and Transfers

	L9010	Kelley	L9050	L5010
United States				
Pre-Gov	0.04 (0.15)	1.91 (6.57)	0.81 (5.93)	-0.78 (-3.78)
Post-Gov	0.34 (1.57)	1.09 (3.40)	0.55 (3.20)	-0.21 (-1.43)
Sweden				
Pre-Gov	-0.02 (-0.08)	2.24 (3.33)	0.50 (4.94)	-0.52 (-2.00)
Post-Gov	-0.41 (-2.00)	0.94 (2.38)	-0.03 (-0.44)	-0.38 (-2.33)
Germany (SOEP)				
Pre-Gov	-1.17 (-3.33)	1.79 (2.76)	-0.03 (-0.12)	-1.15 (-4.22)
Post-Gov	-0.36 (-2.04)	-0.00 (-0.00)	-0.18 (-1.09)	-0.18 (-0.98)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in the indicated 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 SOEP and LINDA, 2 for PSID).

pension and disability payments.²⁸ The largest component of labor-market-related policies mainly is unemployment benefit payments, which acts as an automatic stabilizer against rising risk of job/income losses during recessions. The second component, aid to low-income families, consists of several measures of social insurance policies specifically aimed at at-risk households. These would be expected to matter most for low-income households who have a higher likelihood of satisfying at-risk criteria

²⁸The components are measured as follows. “Labor-market-related policies” in all three datasets are unemployment benefits; in LINDA additionally labor market programs; in PSID additionally workers’ compensation. “Aid to low-income families”: LINDA: family support, housing support, cash transfers from the public (no private transfers); SOEP: subsistence allowance, unemployment assistance (before 2005), unemployment benefits II (since 2005); PSID: Supplemental Security Income; Aid to Families with Dependent Children (AFDC); Food Stamps; Other Welfare. “pension payments: LINDA: (old-age) pensions; SOEP: combined old-age, disability, civil service, and company pensions; PSID: combined (old-age) social security and disability (OASI).

during recessions. Although the third component, pension payments and disability insurance, may not seem directly related to business cycles, they provide additional margins for adjustments (through early retirement for eligible workers or disability claims) in response to job losses and the difficulty of finding jobs during recessions.

To assess the contribution of each component, we construct three hypothetical household earnings measures, each obtained by adding one of these three components to pre-government household earnings. The first three rows of each panel in Table VII reports the cyclical coefficients for each of these three measures. The fourth row reports the cyclical coefficient when all three types of transfers are added at once. This measure does not apply income taxes, so the differences between the coefficients in row (4) and those from post-government earnings in Table VI are informative about the effects of income taxes.²⁹

There are a few main takeaways. First, in all countries, labor market related policies seem to be the most effective at reducing the cyclical coefficient of skewness. In the United States, it brings down β^{Kelley} halfway between 1.91 and 1.09 for pre- and post-government household earnings reported in Table VI. The effect is similar in Germany. But the largest effect is seen in Sweden, where adding labor transfers brings β^{Kelley} from 2.24 for pre-government earnings down to 1.14, very close to 0.94 obtained for post-government earnings Table VI. In other words, in Sweden, labor transfers is by far the most important component of government policies, including taxation, for reducing the cyclical coefficient of skewness in household earnings growth. At the other extreme, pension/disability payments are the least effective (again judging on their effect on skewness) with aid-to-low income families falling in between the two.

Finally, comparing the total effect of all transfers (β^{Kelley} in row 4) to its counterparts for post-government earnings in Table VI shows that income taxes play an important role in further reducing the cyclical coefficient of skewness in Germany and the United States but has a much smaller effect in Sweden.³⁰

²⁹Notice that a flat rate income tax will not affect the growth rate of wages, so to the extent that it affects the distribution of earnings growth, it will be indirectly through labor supply. In contrast, a progressive tax will affect both wage growth and labor supply, which suggests that the progressivity of the income tax system is likely to be critical for its impact on smoothing fluctuations in skewed income risk.

³⁰Given the differences noted earlier between SOEP and SIAB, we conducted further investigation using the latter on the effects of taxes versus labor market transfers. In particular, we ran the cyclical regressions at the individual level with and without unemployment benefits added in the income measure and found its effect on cyclical coefficient to be modest, consistent with our results here from SOEP, leaving an important room for the tax system (see Table H.1 in Appendix H.2).

Table VII: Cyclicalities of Household Earnings: Unbundling Transfers

Gross Household Earnings		L9010	Kelley	L9050	L5010
United States					
(1)	+ Labor transfers	0.23	1.56	0.73	-0.50
		(1.12)	(5.73)	(4.84)	(-3.33)
(2)	+ Aid to low-income	0.04	1.86	0.80	-0.75
		(0.19)	(6.09)	(6.25)	(-3.66)
(3)	+ Pensions/Disability	0.04	1.69	0.73	-0.68
		(0.20)	(5.52)	(5.60)	(-3.28)
(4)	+ All transfers (1+2+3)	0.27	1.50	0.68	-0.42
		(1.52)	(3.86)	(3.85)	(-2.75)
Sweden					
(1)	+ Labor transfers	-0.22	1.14	0.13	-0.35
		(-1.23)	(4.23)	(2.04)	(-2.58)
(2)	+ Aid to low-income	-0.07	2.11	0.42	-0.49
		(-0.38)	(3.72)	(4.51)	(-2.47)
(3)	+ Pensions/Disability	-0.07	2.34	0.48	-0.55
		(-0.43)	(3.55)	(4.50)	(-2.68)
(4)	+ All transfers (1+2+3)	-0.29	1.17	0.08	-0.37
		(-1.78)	(3.92)	(0.96)	(-3.30)
Germany (SOEP)					
(1)	+ Labor transfers	-0.92	1.19	-0.11	-0.82
		(-2.60)	(2.64)	(-0.57)	(-3.35)
(2)	+ Aid to low-income	-1.27	1.54	-0.14	-1.13
		(-3.67)	(2.23)	(-0.58)	(-3.94)
(3)	+ Pensions/Disability	-1.15	1.75	-0.04	-1.10
		(-3.36)	(3.06)	(-0.18)	(-4.62)
(4)	+ All transfers (1+2+3)	-0.85	1.30	-0.07	-0.77
		(-2.68)	(4.37)	(-0.43)	(-4.14)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3 for SOEP and LINDA, 2 for PSID). The income measures are calculated by adding the indicated transfer to gross household earnings.

To sum up, the government plays an important role in smoothing the business cycle fluctuations in skewness, which is a crucial aspect of idiosyncratic income risk. The key contributors to this smoothing are unemployment benefit type policies as well as income taxation. One point that we have not discussed so far but is in fact crucial for an overall assessment is the cost at which this smoothing comes. To see an extreme example, a 100% income tax coupled with a lump sum transfer will completely eliminate all idiosyncratic income risk but is clearly not a sensible or practical policy. But it illustrates the point that the potential benefits from dampening of skewness fluctuations delivered by income taxation must be weighed against the myriad other effects of, including the distortions created by, an income tax system. Such an analysis requires a full fledged dynamic model, which is beyond the scope of this paper.

7 Conclusion

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

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

Second, the skewness cyclicity of individual earnings is a robust feature also for full-time workers who are continuously employed at the same establishment over the business cycle in both Germany and Sweden. An increased left-skewness in aggregate

contractions is thus not just driven by a higher occurrence of unemployment periods, but also by income adjustments on-the-job. We complement this analysis by bringing additional evidence from France, where we observe employer-reported hours worked. The cyclical nature of skewness is driven by both hours and wage adjustments, and again the sample of full-time workers displays strong skewness cyclical nature, which is driven by adjustments of hourly wages. For those workers who are not continuously full-time employed, the extensive margin of employment adjustment is important to explain the increased downside risk in contractions.

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

Fourth, government-provided insurance—unemployment insurance, aid to low-income households, social security benefits, among other transfers and taxes—plays an important role in reducing the cyclical nature of downside risk in all three countries. An interesting assessment that is beyond the scope of this paper would be to quantify the relative roles of automatic stabilizers, active expansions of the social safety net, and tax progressivity (which could be an important driver of changes in the upper tail of income changes) through the lens of a structural model.

References

- ANGELOPOULOS, K., LAZARAKIS, S. and MALLEY, J. (2019). *Cyclical income risk in Great Britain*. Working paper no. 7594, CESifo.
- BAYER, C. and JUESSEN, F. (2012). The life-cycle and the business-cycle of wage risk - cross-country comparisons. *Economics Letters*, **117** (3), 831 – 833.
- BLASS-HOFFMAN, E. and MALACRINO, D. (2016). *Employment Time and the Cyclicity of Earnings Growth*. Tech. rep., Department of Economics, Stanford University.
- BOUND, J., BROWN, C. and MATHIOWETZ, N. (2001). Measurement error in survey data. In J. Heckman and E. Leamer (eds.), *Handbook of Econometrics*, 59, Elsevier, pp. 3705–3843.
- BUSCH, C. and LUDWIG, A. (2020). *Higher-Order Income Risk over the Business Cycle*. Cepr dp 14538.
- 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 GHOSH, A. (2014). *Asset Pricing with Countercyclical Household Consumption Risk*. Working paper, University of Chicago.
- DALY, M., HRYSHKO, D. and MANOVSKII, I. (2014). Reconciling estimates of earnings processes in growth rates and levels, manuscript.
- DECKER, R., HALTIWANGER, J., JARMIN, R. and MIRANDA, J. (2015). *Where has all the skewness gone? The decline in highgrowth (young) firms in the US*. Working Paper 21776, National Bureau of Economic Research.
- DUSTMANN, C., LUDSTECK, J. and SCHÖNBERG, U. (2009). Revisiting the german wage structure. *Quarterly Journal of Economics*, **124** (2), 843–881.
- FERREIRA, T. (2018). Stock market cross-sectional skewness and business cycle fluctuations.

- FITZENBERGER, B., OSIKOMINU, A. and VÖLTER, R. (2006). Imputation rules to improve the education variable in the IAB employment subsample. *Schmollers Jahrbuch. Zeitschrift für Wirtschafts- und Sozialwissenschaften*, **Jg. 126**.
- FUCHS-SCHÜNDELN, N., KRUEGER, D. and SOMMER, M. (2010). Inequality trends for germany in the last two decades: A tale of two countries. *Review of Economic Dynamics*, **13** (1), 103–132.
- GOLOSOV, M., TROSHKIN, M. and TSYVINSKI, A. (2016). Redistribution and social insurance. *American Economic Review*, **106** (2), 359–386.
- GUVENEN, F., OZKAN, S. and SONG, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, **122** (3), 621–660.
- HARMENBERG, K. and SIEVERTSEN, H. H. (2018). *The Labor-Market Origins of Cyclical Income Risk*. Working paper, Institute for International Economic Studies.
- HUGGETT, M. and KAPLAN, G. (2016). How large is the stock component of human capital? *Review of Economic Dynamics*, **22**, 21–51.
- ILUT, C., KEHRIG, M. and SCHNEIDER, M. (2018). Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. *Journal of Political Economy*, **126** (5), 2011–2071.
- KAPLAN, G., MOLL, B. and VIOLANTE, G. L. (2016). *Monetary Policy According to HANK*. Working paper, Princeton University.
- KEHRIG, M. (2011). *The Cyclicalities of Productivity Dispersion*. Working Paper 11-15, Center for Economic Studies, U.S. Census Bureau.
- KIM, T.-H. and WHITE, H. (2004). On more robust estimation of skewness and kurtosis. *Finance Research Letters*, **1** (1), 56–73.
- KIRBY, C. (1997). Measuring the predictable variation in stock and bond returns. *Review of Financial Studies*, **10** (3), 579–630.
- KURMAN, A. and MCENTARFER, E. (2017). *Downward Wage Rigidity in the United States: New Evidence from Administrative Data*. Working paper.
- KURMANN, A. and MCENTARFER, E. (2019). *Downward Nominal Wage Rigidity in the United States: New Evidence from Worker-Firm Linked Data*. Working Papers 19-07, Center for Economic Studies, U.S. Census Bureau.

- MANKIW, N. G. (1986). The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics*, **17** (1), 211–219.
- OECD (2016). Economic policy reforms 2016: Going for growth interim report. OECD Publishing.
- OH, S. and WACHTER, J. A. (2018). *Cross-sectional Skewness*. Working paper, Wharton School at University of Pennsylvania.
- PORA, P. and WILNER, L. (2017). The individual dynamics of wage income in france during the crisis. *Economie et Statistique / Economics and Statistics*, (494-496), 179–199.
- PRUITT, S. and TURNER, N. (2018). *The Nature of Household Labor Income Risk*. Working paper.
- RICHARDSON, M. and STOCK, J. H. (1989). Drawing inferences from statistics based on multiyear asset returns. *Journal of Financial Economics*, **25** (2), 323–348.
- SALGADO, S., GUVENEN, F. and BLOOM, N. (2019). *Skewed Business Cycles*. Working Paper 26565, National Bureau of Economic Research.
- SCHMIDT, L. D. W. (2016). *Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk*. Working paper, University of Chicago.
- SHIN, D. and SOLON, G. (2011). Trends in men’s earnings volatility: What does the panel study of income dynamics show?., *Journal of Public Economics*, **95** (7-8), 973–982.
- SKEDINGER, P. (2007). *The Design and Effects of Collectively Agreed Minimum Wages: Evidence from Sweden*. Working Paper Series 700, Research Institute of Industrial Economics.
- STORESLETTEN, K., TELMER, C. I. and YARON, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, **112** (3), 695–717.
- VOM BERGE, P., BURGHARDT, A. and TRENKLE, S. (2013). *Sample-of-Integrated-Labour-Market-Biographies Regional-File 1975-2010 (SIAB-R 7510)*. Tech. rep., Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB).

Appendices

A Data Appendix

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

A.1 PSID

Variables

Demographic and Socioeconomic

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

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

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

Income

Individual Male Wages and Salaries. This is the variable used for individual income in the benchmark case. It is the answer to the question: *How much did (Head) earn altogether from wages or salaries in year t-1, that is, before anything was deducted for taxes or other things?* This is the most consistent earnings variable

over time reported in the PSID, as it has not suffered any redefinitions or change in subcomponents.³¹

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,³² we subtract the labor part of business and farm income before 1993.

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

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

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

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

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

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

Transfers

See Table [A.1](#) below 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

³¹See [Shin and Solon \(2011\)](#) for a comparison of PSID male earnings variables in inequality analyses.

³²In particular, total labor earnings included the labor parts of farm and business income up to the 1993 survey but not in subsequent waves.

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 1970-1993³³ and 2005-2011, the total annual amount of each program is reported for the head, spouse, and others in the family unit. Occasionally, the amount appears combined for several or all members.³⁴ Because in those cases it is impossible to identify separate reciprocity of each member, we consider the benefit amount of the whole family. That is, we add up all available information for all family members, whether combined or separately reported.

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

Detailed Sample Selection

We start with an initial sample of 584,392 SRC individuals interviewed between 1976 and 2011. We then impose the next criteria every year. The number of individuals kept at each stage in the sample selection is listed in Table A.2. 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

³³Our main sample refers to survey years 1977-2011, but complementary results are provided for the annual subsample of the PSID, that is, for 1970-1997. We drop the first two waves in all cases, since benefits such as OASDI, UI, and WC are only reported for the family head and benefits such as SSI are not reported at all.

³⁴This is always the case for Food Stamps.

Table A.1: 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: The table lists the measures used in the three datasets to construct subcomponents of transfers.

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 and then adds information for all family members, where family is defined for tax

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

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

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

purposes. This implies that there is no information about “head of households.” We therefore define the head of a household as the sampled male.

Age. As defined by Statistics Sweden.

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

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

Income

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

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

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

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

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

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

Transfers

Transfers in subcategories 1 and 3 are individual-level transfers. Transfers in subcategory 2 are family level transfers but have been individualized by Statistics Sweden. For each subcategory, we take all transfers received by all members of the households.

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

Detailed Sample Selection

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

A.3 LISA

The LISA database covers all individuals between 16 and 64 years of age for the period 1990-95 and all individuals above age 16 thereafter. Like in LINDA, all income data is based on tax records. Using the same definitions for all variables and the same sample selection procedure as in LINDA results in a sample with around 1.6 million

one-year income changes per year. It contains annual information on employers of individuals, as well as on the establishment. We use this additional information in order to identify workers that stay at the same main establishment for a given t to $t+1$ change. We restrict this sample of stayers to individuals whose main establishment in both t and $t+1$ is the same, and who in addition received income from this establishment in both $t-1$ and $t+2$.

A.4 SIAB

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

Variables

Demographic and Socioeconomic

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

Age. Birth year is reported consistently in SIAB data.

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

Private/Public Employment. An individual is defined as working in the public sector, if he/she works in public administration, health care or education. SIAB contains consistent comparable information for all years of the sample. We use the

classification WZ93 as provided in the data, which aggregates 3-digit codes of the original WZ93 classification into 14 categories. The industry of an employer is registered once a year and assigned to the worker spells of that year. This implies that for some individual spells, there is no information on the industry. For each year, a worker is assigned the industry from the longest spell in that year. We classify as public employment those in sectors 13 (3-digit WZ93 801-804, 851-853: Education, social, and health-care facilities) and 14 (751-753, 990: public administration, social security).

Income

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

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

Imputation of Top-Coded Wages. Before aggregating earnings from all spells, we correct full-time wage spells for the top-coding. We therefore follow [Daly *et al.* \(2014\)](#) and fit a Pareto tail to the cross-sectional wage distribution. The Pareto distribution is estimated separately for each year by age group and sex. We define seven age groups: 25-29, 30-34,...,55-60. As a starting point for the Pareto distribution, we choose the 60th percentile of the subgroup-specific distribution. As in [Daly *et al.* \(2014\)](#), we draw one random number by individual, which we then apply to the annual specific

distributions when assigning a wage to the top-coded workers. We apply the imputation method to the annual distribution of average full-time wages, and hence an individual can be below the cutoff limit if, for example, from two full-time spells in a year only one is top-coded. We therefore define as the top-coding limit the annual specific limit minus 3 DM (1995 DM) as in [Dustmann *et al.* \(2009\)](#).

Transfers

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

Detailed Sample Selection

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

A.5 SOEP

Variables

Demographic and Socioeconomic

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

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

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

Income and Hours

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

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

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

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

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

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

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

Transfers

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

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

- *HH-level transfers subcategory 1 (labor market transfers)*: sum of unemployment benefits received by all members of household.
- *HH-level transfers subcategory 2 (family aid)*: sum of subsistence allowance of all members, + sum of unemployment assistance received by all members (up to 2004), + hh-level measure of unemployment benefits II (since 2005).
- *HH-level transfers subcategory 3 (pensions)*: sum of old-age pensions received by all members of household.

Sample Selection

In order to be in the initial sample for a year, the individual or household head must be between ages 25 and 60 and live in West Germany. In order to have a consistent sample, we drop the immigrant subsample and the high-income subsample. This gives initial sample sizes of 87,582 individual-year observations for the male sample, 76,249 individual-year observations for the female sample, and 76,051 household-year observations for the household sample (see Table A.3). 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 $520 * 0.5 * \text{minimum wage}$, 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 $t-5$

Table A.3: Number of Observations Kept in Each Step: SOEP

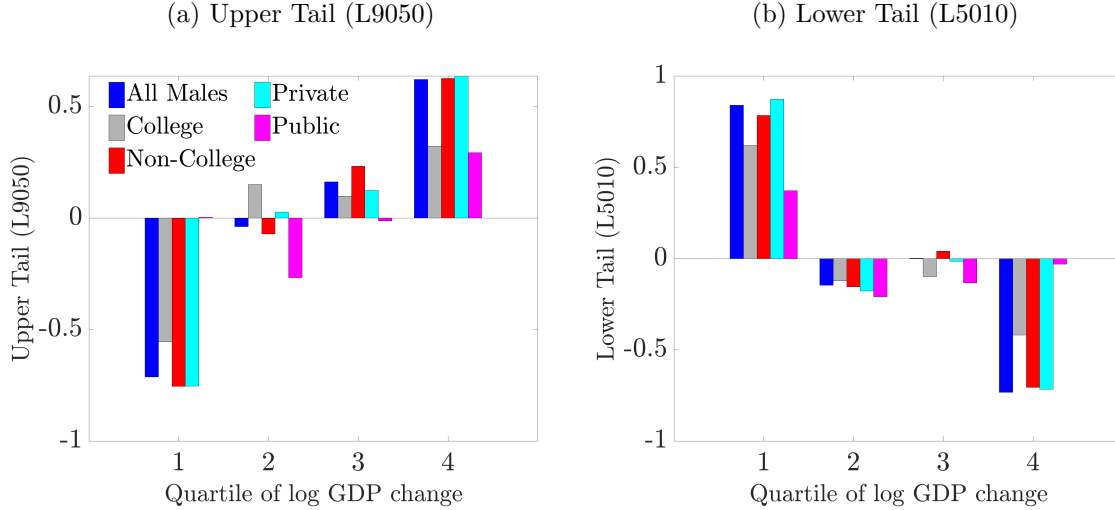
selection step	Male Heads	Households	All Females
initial	87,582	95,982	95,716
drop if no coll. info	86,737	95,008	94,456
drop if in military	86,712	94,990	94,454
drop if no obs on ymin	86,009	94,990	93,960
drop if no obs on hours	86,009	94,990	93,960
keep if ≥ 520 hrs and $ymin > 0$	77,501	87,332	57,979
drop top 1% of ymin per year	76,641	86,379	57,475
drop if $ymin < .5 * 520 * \text{min wage}$	76,460	85,429	56,803
Final #Obs for transitory changes	64,824	71,287	44,805
Final #Obs for persistent changes	38,555	41,048	24,968

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

A.6 DADS

The DADS (“*Declaration Annuelle des Donneés Sociales*”) panel is extracted from exhaustive administrative records of annual employer-employee information with compulsory completion by all firms and establishments within firms. These administrative records are used for taxation and social security purposes. All private sector firms, regardless of the number of employees, are included. Regarding the public sector, it includes dependent workers from semi-public firms. To constitute the panel, a 4% sample of workers was extracted from the records since 1976 and until 2001. These consist of all workers born in October of an even calendar year. Since 2002, workers born in October of all years have been included. This results in samples of approximately 8%. We use data on earnings, number of days paid, start and end date of the spell, and, importantly, number of hours worked, which is recorded since 1993. Due to data processing changes between 1993 and 1994, those years are dropped in the analysis, and thus we end up with usable data from 1995–2015. Given the similar nature of the dataset, we follow the same basic steps of sample selection as for the SIAB.

Figure B.1: L9050 and L5010 against the Quartiles of Log GDP Change: Males



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

B Cyclicity of Individual Earnings by Groups

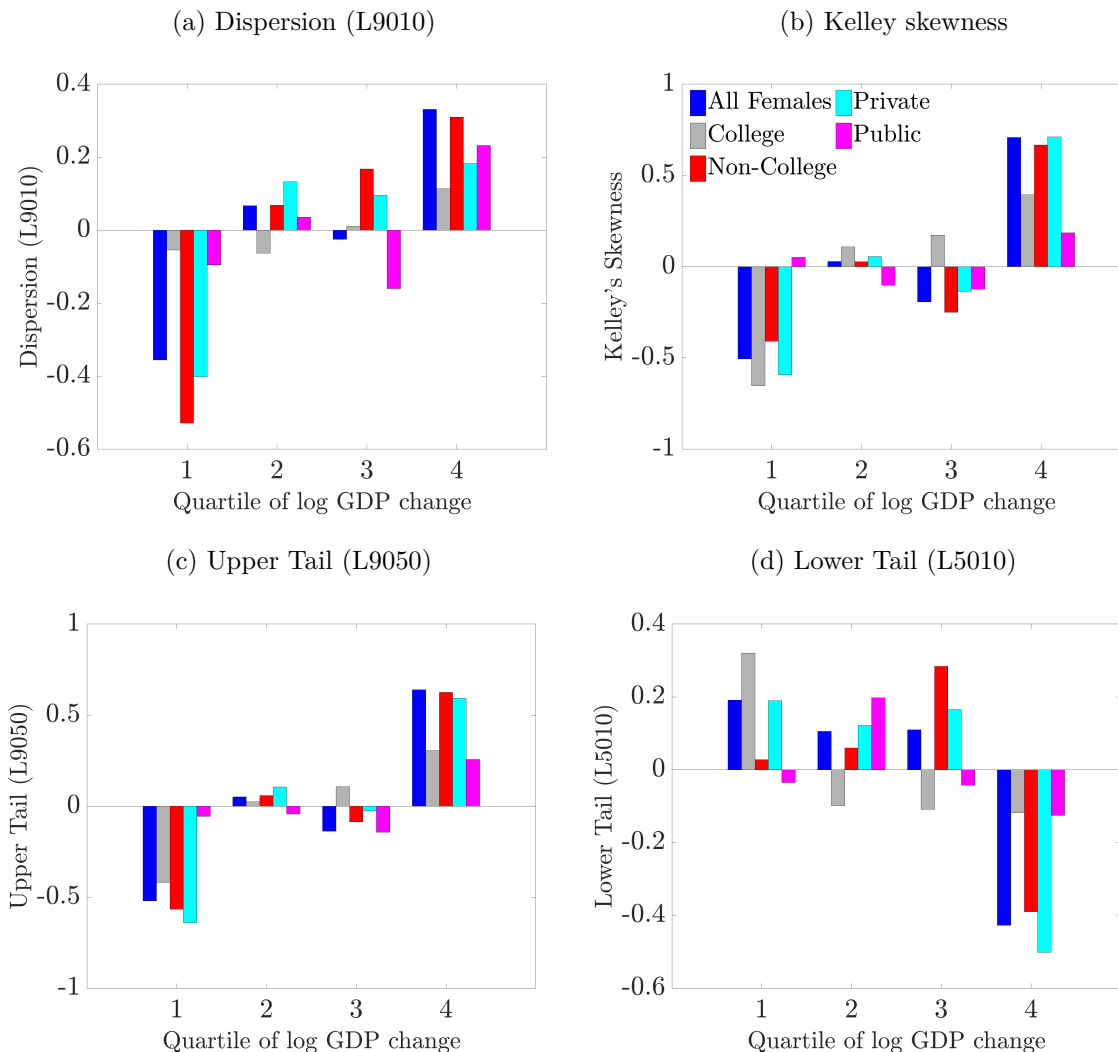
B.1 Education and Sector of Employment

Figure B.1 complements Figure 6 in the text and reports L9050 and L5010 for males in the same format. The next figure (B.2) reports the counterparts of these two figures for females.

Furthermore, we show results of the individual level earnings regressions discussed in Section 4 by subgroups. For each group and country, we estimated our baseline regression (equation 3). The estimated sensitivity coefficients are displayed in Figure B.3, followed by Tables B.1 to B.5 showing the specific estimates, as well as the corresponding t-statistics. Each panel in the figure shows, starting from the left, the regression coefficients along with 95% confidence intervals for males (solid) in Sweden (red, triangles), Germany (green, squares), and the US (blue, bullets), followed by the equivalent regression coefficients for females (dotted). Within each country-gender grouping, the coefficients are (ordered from the left) those from the full sample, college graduates, non-college graduates, private employment, and public employment, respectively.

Figure B.3 confirms the picture that emerged in Figures 6 and B.2: higher-order earnings risk is similar across groups. However, we see some noteworthy differences.

Figure B.2: Higher-Order Moments by Quartiles of Log GDP Change: Females



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

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

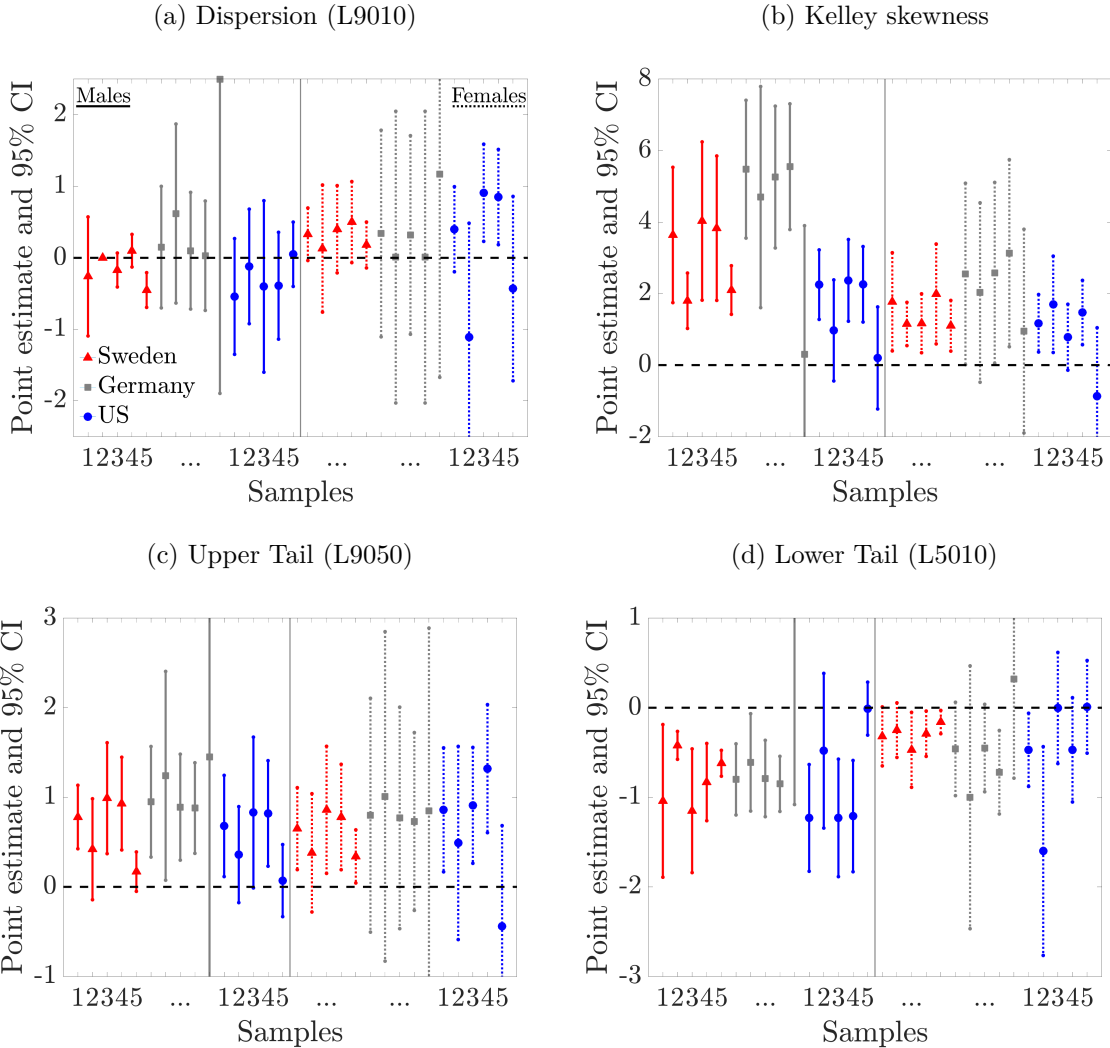
In Sweden, the procyclicality of Kelley’s measure of earnings is lower for the public sector (2.10 for males and 1.10 for females) compared with the private sector (3.83 for

males and 1.99 for females). For males, this is due to differences in the top tail; it compresses strongly for private sector employees, whereas it is acyclical in the public sector. The L5010 gap, on the other hand, fluctuates by comparable magnitudes for both groups. For women, the reduced cyclicality is due to both tails fluctuating slightly less.

Overall, it is somewhat surprising that for workers in the public sector in a country like Sweden with a reputation for high levels of public insurance, there is robust evidence of higher downside risk in recessions—compression of the top and expansion of the bottom of the distribution of income changes—even if the magnitudes are somewhat smaller than in the private sector.³⁵ This finding further strengthens the conclusion in Section 4 that increasing downside (individual) earnings risk appears to be a robust feature of business cycles in developed countries.

³⁵One explanation (suggested by a referee) could be in employment protection that creates concentration of earnings changes. If most workers' earnings cannot fall by statute or contract, then the measured declines of earnings are likely to be concentrated among the set of people who work at firms that are sufficiently hard hit, which would then generate an expansion of the left tail of earnings changes.

Figure B.3: Cyclicity of Earnings for Subgroups (Sweden, Germany (SIAB), and the United States)



Note: Separate regressions for different samples. Each marker reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. The confidence bands are based on Newey-West standard errors (maximum lag length considered: 3). The samples are (1) earnings: full sample, (2) earnings: college graduates, (3) earnings: non-college graduates, (4) earnings: private sector, (5) earnings: public sector. Sweden is marked by red triangles, Germany by green squares, and US by blue circles. In each figure, the left (right) half shows the results for males (females). For details of samples, see text. See tables B.1 to B.5 for specific values and t-statistics.

B.2 Occupations

This section reports additional results to complement the analysis of skewness fluctuations by occupation in Section 5.1. Figure B.4 shows the cyclicity coefficients for

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

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

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

L5010 and L9050 to complement the L9010 and Kelley skewness shown in Figure 7. Table B.5 reports the cyclicalities regressions for five broader occupational categories as instead of the 30 detailed categories in Figure 7). Figures B.5 and B.6 are the analogues of Figures 7 and B.6 for females.

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

	L9010	Kelley	L9050	L5010
United States				
College graduates	-1.11 (-1.44)	1.70 (2.61)	0.49 (0.94)	-1.60 (-2.84)
Non-college	0.91 (2.77)	0.78 (1.75)	0.91 (2.91)	-0.00 (-0.01)
Sweden				
College graduates	0.13 (0.31)	1.15 (4.03)	0.28 (1.22)	-0.25 (-1.74)
Non-college	0.50 (1.96)	1.81 (3.40)	0.75 (2.78)	-0.25 (-2.71)
Germany (SIAB)				
College graduates	0.01 (0.01)	2.03 (1.65)	1.01 (1.12)	-1.00 (-1.39)
Non-college	0.32 (0.47)	2.58 (2.08)	0.77 (1.27)	-0.45 (-1.88)

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

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

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

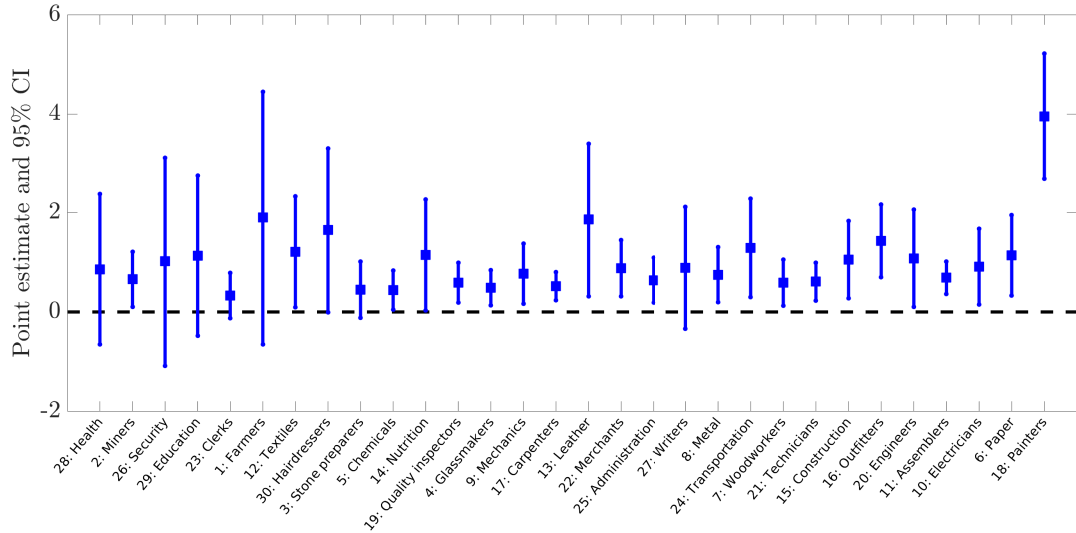
Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3 for SIAB and LINDA, 2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

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

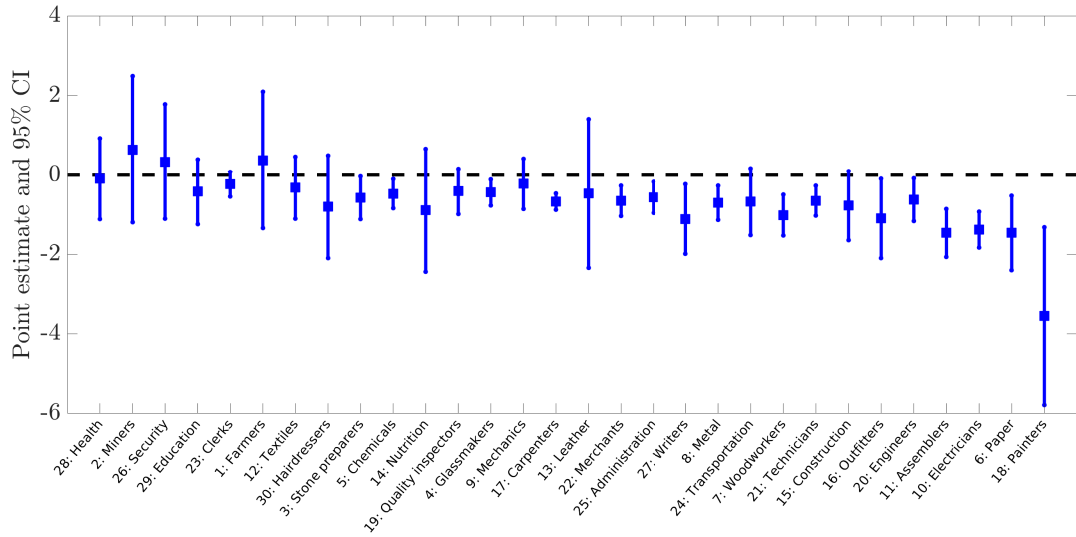
	L9010	Kelley	L9050	L5010
United States				
Private	0.85 (2.64)	1.47 (3.38)	1.32 (3.82)	-0.47 (-1.67)
Public	-0.43 (-0.69)	-0.87 (-0.94)	-0.44 (-0.81)	0.01 (0.04)
Sweden				
Private	0.50 (1.87)	1.99 (3.02)	0.78 (2.81)	-0.29 (-2.43)
Public	0.18 (1.19)	1.10 (3.29)	0.34 (2.43)	-0.16 (-2.61)
Germany				
Private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

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

Figure B.4: Tails of Short-Run Income Growth by Occupation: Males (Germany (SIAB))



(a) Upper Tail (L9050)



(b) Lower Tail (L5010)

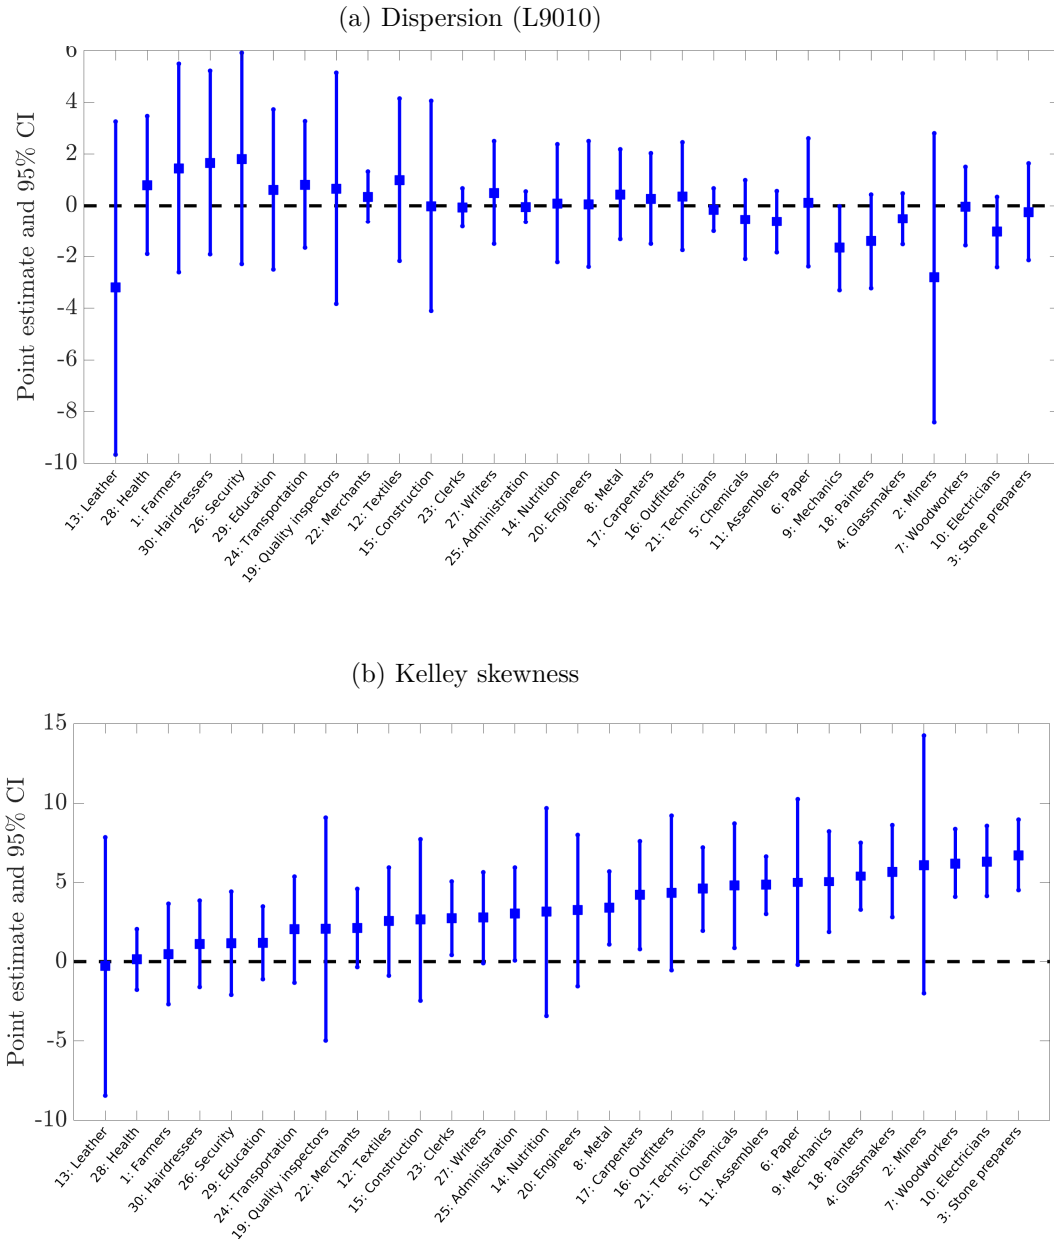
Note: Separate regressions for each of 30 occupation segments. Each marker reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. The confidence bands are based on Newey-West standard errors (maximum lag length considered: 3)

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

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

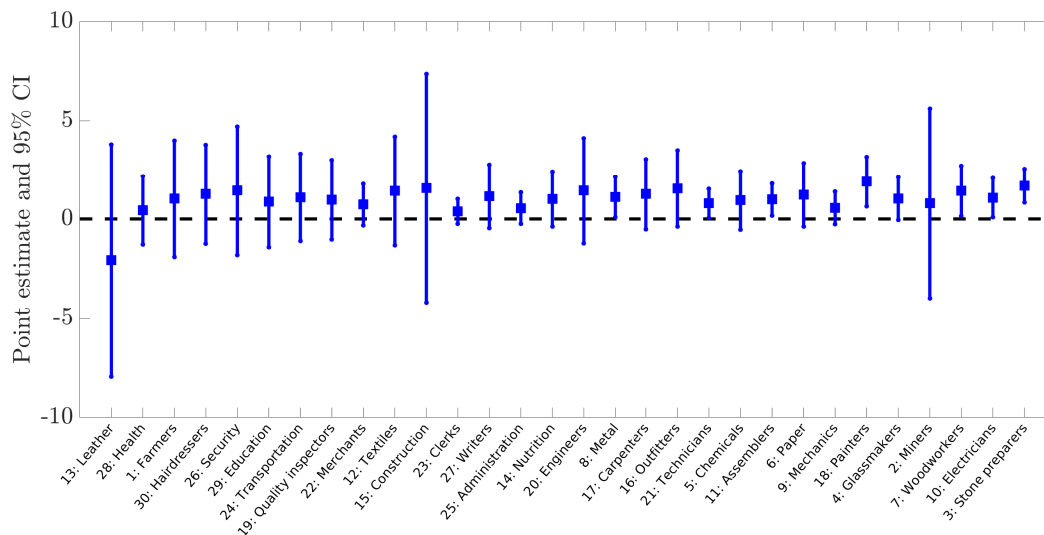
Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3).

Figure B.5: Dispersion and Skewness of Short-Run Income Growth by Occupation: Females (Germany (SIAB))

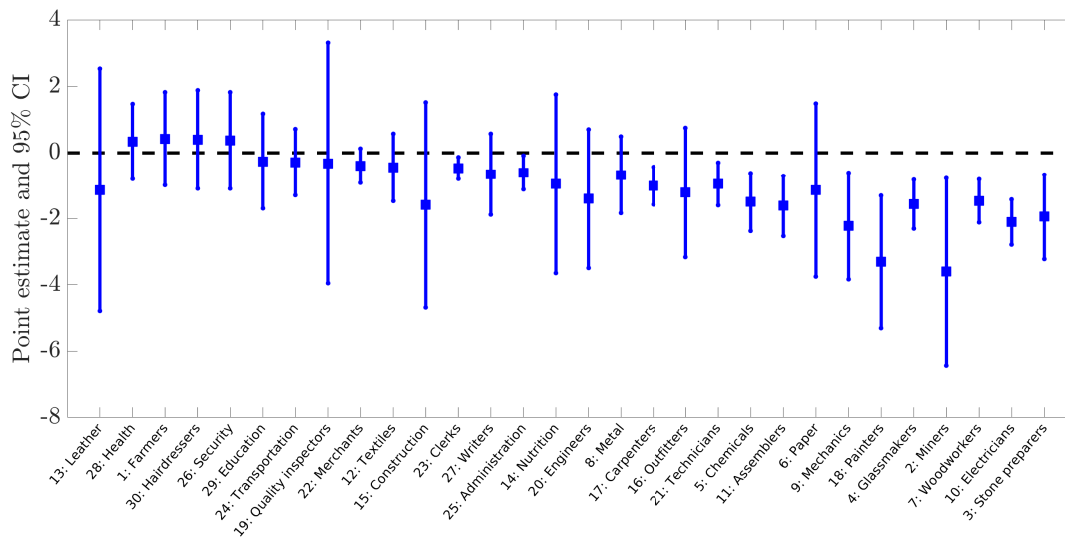


Note: Separate regressions for each of 30 occupation segments. Each marker reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. The confidence bands are based on Newey-West standard errors (maximum lag length considered: 3)

Figure B.6: Tails of Short-Run Income Growth by Occupation: Females (Germany (SIAB))



(a) Upper Tail (L9050)



(b) Lower Tail (L5010)

Note: Separate regressions for each of 30 occupation segments. Each marker reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. The confidence bands are based on Newey-West standard errors (maximum lag length considered: 3)

Table C.1: Cyclicity of Income Growth Moments: Arc Percent Changes

	L9010	Kelley	L9050	L5010
United States				
Males	-2.21 (-1.61)	3.27 (5.43)	1.45 (3.77)	-3.66 (-2.78)
Females	-2.22 (-1.42)	1.99 (3.55)	1.44 (2.14)	-3.66 (-2.80)
Sweden				
Males	0.02 (0.02)	4.86 (3.85)	1.78 (4.10)	-1.76 (-1.95)
Females	1.16 (2.50)	2.67 (2.85)	1.57 (3.22)	-0.41 (-1.42)
Germany (SIAB)				
Males	-0.22 (-0.37)	6.93 (4.50)	1.36 (3.13)	-1.58 (-3.09)
Females	0.11 (0.11)	3.49 (2.10)	1.36 (1.42)	-1.25 (-1.98)

Note: Each cell reports the coefficient on log GDP growth of a regression of each moment of the distribution of income changes as measured by the arc-percent change $(\frac{x_t - x_{t-1}}{(x_t + x_{t-1})/2})$ on log GDP growth, 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).

C Alternative Specifications

C.1 Including Zeros

In this section, we first run an alternative version of the benchmark cyclicity regression: we eliminate the lower threshold in sample selection and calculate earnings changes using the *arc-percent* formula instead of log changes. This allows the inclusion of zero incomes. We then use moments of the distribution of arc-percent changes as the dependent variable in our regression. The main results hold.

C.2 Using a Binary Classification of Years

We now consider an alternative specification to test for the cyclicity of different moments of the distribution of earnings changes, for different types of earnings. Instead of using log GDP change as a measure of the aggregate state of the economy, one can classify years as expansions and contractions, which is what, e.g., [Storesletten *et al.* \(2004\)](#) do. We initiate our classification of years with the NBER dates (as described

in the main text). We extend the measure by considering the average growth of male earnings (as done in [Huggett and Kaplan \(2016\)](#)).³⁶ Our preferred specification remains the continuous measure in the main text, because it is more flexible and simplifies the cross-country comparison. Furthermore, it does not come with the difficulty of exactly timing recessions at the yearly level

Still, there are advantages of resorting to a binary characterization. To begin with, it makes our results more easily comparable to previous work [Storesletten *et al.* \(2004\)](#); [Pruitt and Turner \(2018\)](#); [Guvenen *et al.* \(2014\)](#). More importantly, it tests the robustness of our results in the case we thought of cyclical as a regime switching rather than a linear relationship between the aggregate state and the distribution of earnings.³⁷

The following tables show regression results for an alternative specification where we regress on a dummy that takes the value 1 in expansions instead of regressing on log GDP growth. The results are virtually unchanged, i.e., they tell exactly the same story as our benchmark specification: entering an expansion is associated with a significant increase in skewness and a non-significant change in dispersion. We document the four main types of earnings, but there are no significant changes from any of the results in the main text. We focus on the case of the US since the main issue is comparability with previous work—which is focused on the US—and specification robustness. All in all, we can conclude that, if the cyclical relationship is more a regime-switching model with reasonable noise and realistic dynamics, our regression picks up the relationship.

³⁶This results in contraction years 1970, 1974, 1980, 1981, 1982, 1983, 1990, 1991, 2001, 2002, 2008, 2009, 2010, and 2012.

³⁷We thank one of the referees for suggesting this extension and the thoughtful interpretation.

Table C.2: Cyclicalities of Individual Earnings using binary Measure of Business Cycles: United States

	L9010	Kelley	L9050	L5010
Males (GOS)	0.00	0.15	0.07	-0.07
	(-0.06)	(6.66)	(6.14)	(-4.06)
Males	-0.02	0.15	0.05	-0.08
	(-0.91)	(4.17)	(3.54)	(-3.13)
Females	0.04	0.06	0.06	-0.01
	(2.10)	(2.07)	(2.43)	(-0.85)
HH Earnings	0.01	0.12	0.06	-0.04
	(0.63)	(5.10)	(7.68)	(-2.42)
HH Post-Gov	0.03	0.08	0.04	-0.01
	(2.05)	(4.36)	(4.70)	(-1.23)

Note: Each cell reports the coefficient on a dummy—1 in expansions and 0 in recessions—of a regression of each moment of the distribution of income changes on such dummy, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 2). Males (GOS) uses the moments of [Guvenen et al. \(2014\)](#).

D Hours versus Wages: Females

Table [D.1](#) shows the regressions of full time workers and stayers for German and Swedish females. Table [D.2](#) shows the results for females in SIAB who are not in the full-time sample. The results for females mirror the findings for males in the main text. The systematic variation of skewness of earnings changes over the business cycle cannot be entirely explained by the extensive margin of employment changes. Instead, also the incomes of workers who are continuously full-time employed over the cycle display an increase of left-skewness in aggregate downswings.

E Survey versus Administrative Data

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

Specifically, we have rerun the regression in equation [\(3\)](#) using moments from the SSA data (reported in [Guvenen et al. \(2014\)](#)), and from SOEP data. The resulting coefficients for US males using SSA data for each of the four moments are -0.07 , 2.31 ,

Table D.1: Cyclicalities of Individual Earnings vs. Daily Wages; Germany (SIAB) and Sweden (LISA): Females

	L9010	Kelley	L9050	L5010
Germany				
Earnings (Full population)	0.34 (0.48)	2.55 (2.05)	0.80 (1.25)	-0.46 (-1.80)
Full-Time Daily Wages	0.03 (0.18)	2.12 (5.11)	0.17 (2.61)	-0.14 (-1.58)
Full-Time Daily Wages (Establishment Stayers)	0.02 (0.13)	2.28 (4.84)	0.16 (3.17)	-0.14 (-1.61)
Sweden				
Earnings (Full population)	0.38 (1.85)	1.64 (3.69)	0.64 (2.97)	-0.26 (-3.24)
Establishment Stayers	-0.08 (-1.55)	0.69 (2.55)	0.05 (1.23)	-0.13 (-2.62)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3). Full-Time are those that work full time for at least 50 weeks in both years for which the change is calculated.

Table D.2: Cyclicalities of Days Worked vs. Daily Wages; Germany (SIAB): Females

	L9010	Kelley	L9050	L5010
Earnings	0.37 (0.47)	1.89 (2.34)	0.92 (1.31)	-0.55 (-2.21)
Daily Wages	0.21 (0.86)	1.04 (2.38)	0.26 (1.63)	-0.05 (-0.39)
Days Worked	0.22 (0.22)	2.34 (2.12)	0.81 (0.96)	-0.59 (-2.08)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in the indicated 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). Workers from the baseline sample that are not working full-time \times full-year in either t-1 or t.

1.02, and -1.09 , respectively. These estimates are strikingly similar to those in the first row of the top panel in Table I. The equivalent estimates using SOEP data are -1.27 , 1.55 , -0.23 , and -1.04 . While these numbers differ somewhat from those in the first row of the bottom panel in Table I, they tell the same story. In particular, male earnings changes in both SOEP and SIAB are characterized by asymmetric movements of the tails rather than uniform expansions and contractions of both tails.³⁸ The main difference is that, in the SOEP, there is evidence of countercyclical dispersion, which was not observed in the SIAB.

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

F Robustness of the Empirical Results in SIAB Data

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

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

³⁸We have also run regression 3 using the standard deviation of earnings changes as our measure for overall dispersion instead, and the coefficients are small (0.07 (SIAB), -0.09 (SOEP)) and insignificant (t-stat of 0.42 (SIAB), -0.38 (SOEP)) in both datasets.

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 F.2 shows the results of the respective regressions for earnings, wages, and wages of firm stayers for both males and females. Second, we rerun the regressions completely ignoring top-coding, that is, all individuals from the base sample are in the sample, but with their reported incomes again for earnings, wages, and wages of stayers. Results are in Table F.3.

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

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

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

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3)

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

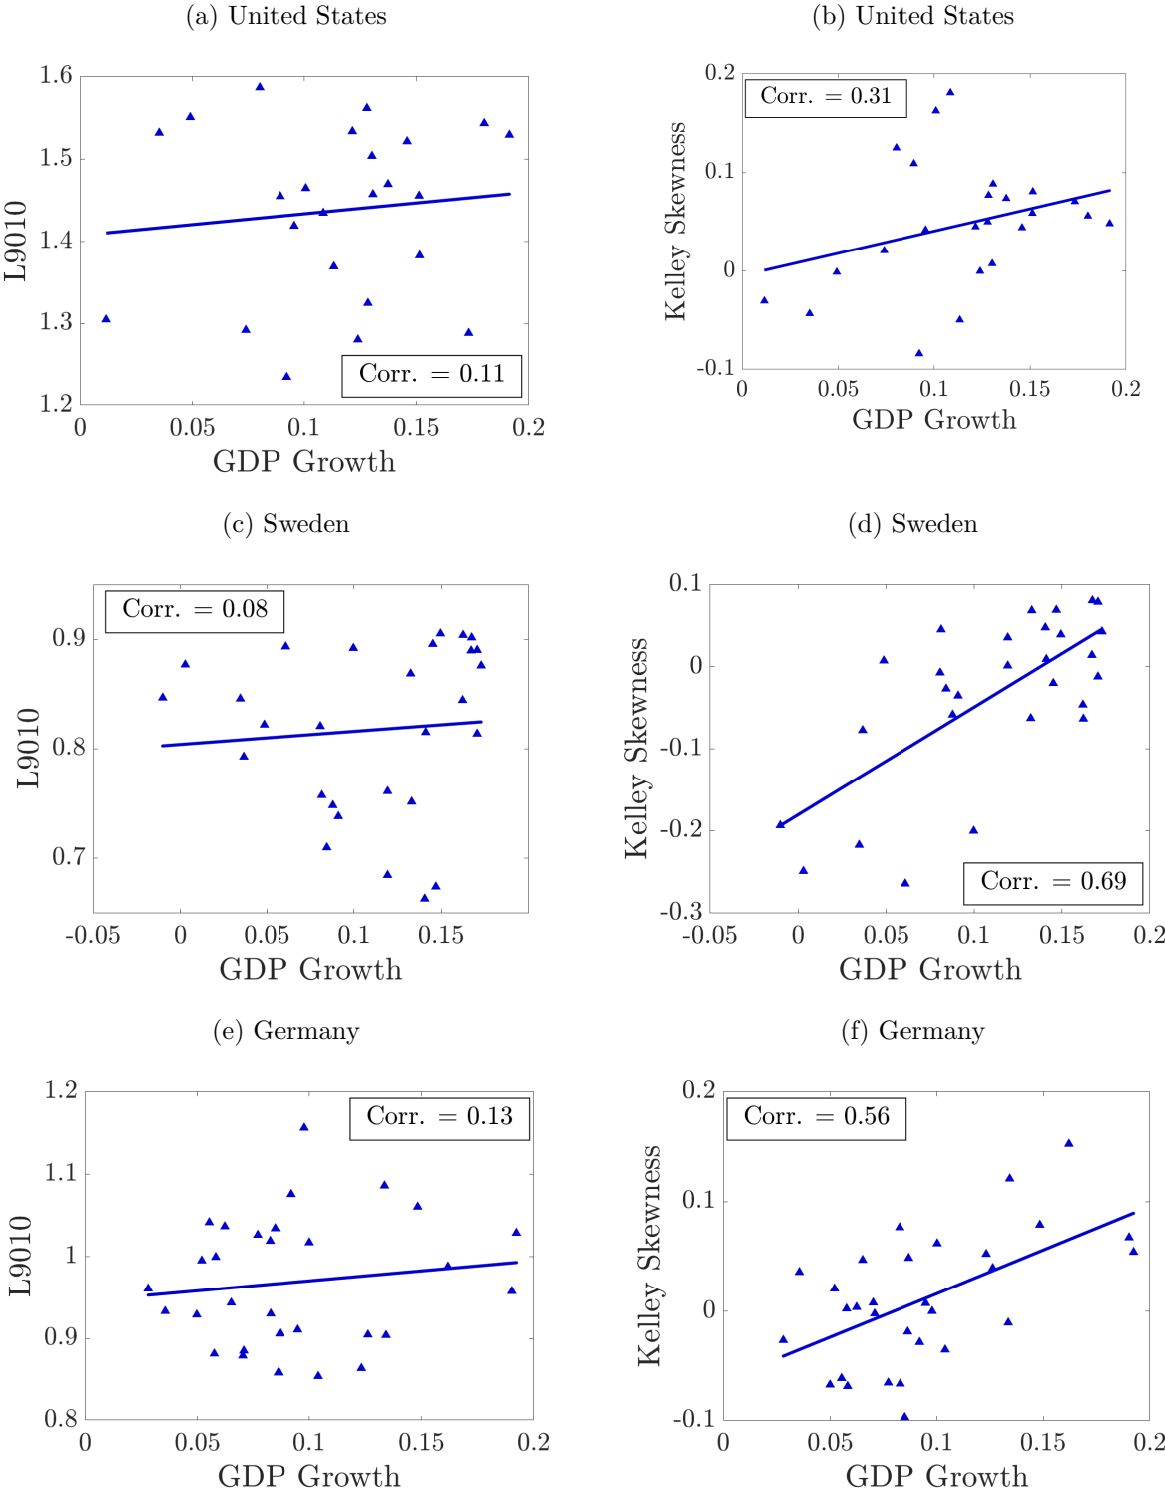
	Std Dev	L9010	Skew	Kelley	L9050	L5010	Hinkley 1	Hinkley 2	L8550	L8050	L5015	L5020
					Not top-coded workers only:							
Male Earnings	0.08 (0.41)	0.26 (0.53)	14.49 (4.26)	4.98 (4.28)	0.96 (2.53)	-0.70 (-3.07)	4.83 (6.66)	4.65 (8.86)	0.48 (3.13)	0.31 (3.40)	-0.44 (-4.06)	-0.28 (-3.08)
Male Wages	-0.01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

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3)

G Long-Run Earnings Growth

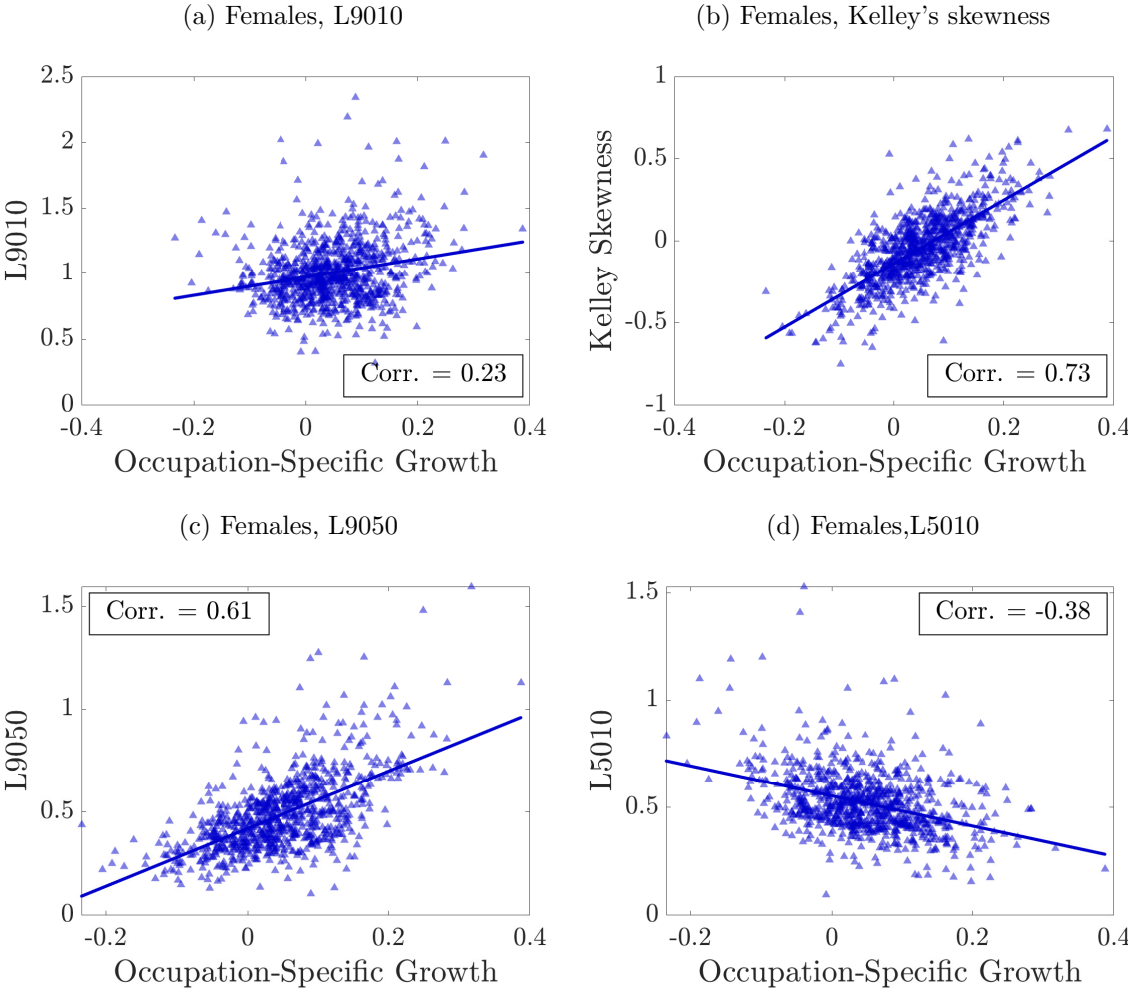
Figure [G.1](#) shows L9010 and Kelley's skewness of long-run earnings growth, that is, five-year changes for Germany and Sweden, and four-year changes for the United States against log GDP growth for females.

Figure G.1: Cyclicity of Five-Year Income Growth; United States, Sweden, and Germany (SIAB): Females



Note: Scatterplot of moment of five-year earnings change against log GDP change over the same horizon. Five-year income changes for Germany and Sweden, four-year income changes for the United States.

Figure G.2: Distribution of Five-Year Income Growth; Occupation-Specific Cycles (SIAB): Females



Note: Scatterplot of moment of five-year earnings change against occupation-specific average growth over the same horizon.

Table G.1 shows the correlation of L9010 and Kelley's skewness of five-year earnings changes with log GDP growth for workers by occupational segment. A worker contributes to the occupation-specific moment if in year $t-5$ he or she is in that occupation. Figure G.2 shows the relationship between occupation-specific average earnings growth and moments of the distribution for females.

Table G.1: Correlation of Five-Year Income Growth with GDP growth: Germany (SIAB)

Occupation	Males		Females	
	L9010	Kelley	L9010	Kelley
1	0.41	0.53	0.15	0.53
2	0.32	0.46	0.10	0.46
3	-0.12	0.53	-0.11	0.53
4	0.24	0.56	0.07	0.56
5	-0.51	0.83	-0.08	0.83
6	-0.29	0.72	0.01	0.72
7	0.32	0.37	-0.02	0.37
8	0.09	0.67	0.12	0.67
9	0.20	0.54	-0.17	0.54
10	0.22	0.61	0.43	0.61
11	-0.24	0.75	0.24	0.75
12	-0.11	0.71	0.36	0.71
13	-0.01	0.66	-0.22	0.66
14	-0.19	0.64	0.12	0.64
15	-0.27	0.63	-0.01	0.63
16	-0.31	0.71	-0.02	0.71
17	-0.16	0.51	0.20	0.51
18	-0.34	0.35	-0.44	0.35
19	-0.08	0.68	-0.02	0.68
20	-0.12	0.60	0.20	0.60
21	-0.26	0.63	-0.25	0.63
22	-0.35	0.68	-0.09	0.68
23	-0.32	0.65	-0.11	0.65
24	-0.10	0.69	0.00	0.69
25	-0.33	0.60	-0.26	0.60
26	0.12	0.55	-0.19	0.55
27	-0.45	0.65	-0.09	0.65
28	0.20	0.52	0.45	0.52
29	0.43	0.58	0.62	0.58
30	-0.08	0.63	0.20	0.63

Note: Each entry is the correlation of occupation-specific moments of the distribution of income growth with log GDP changes over a five-year horizon.

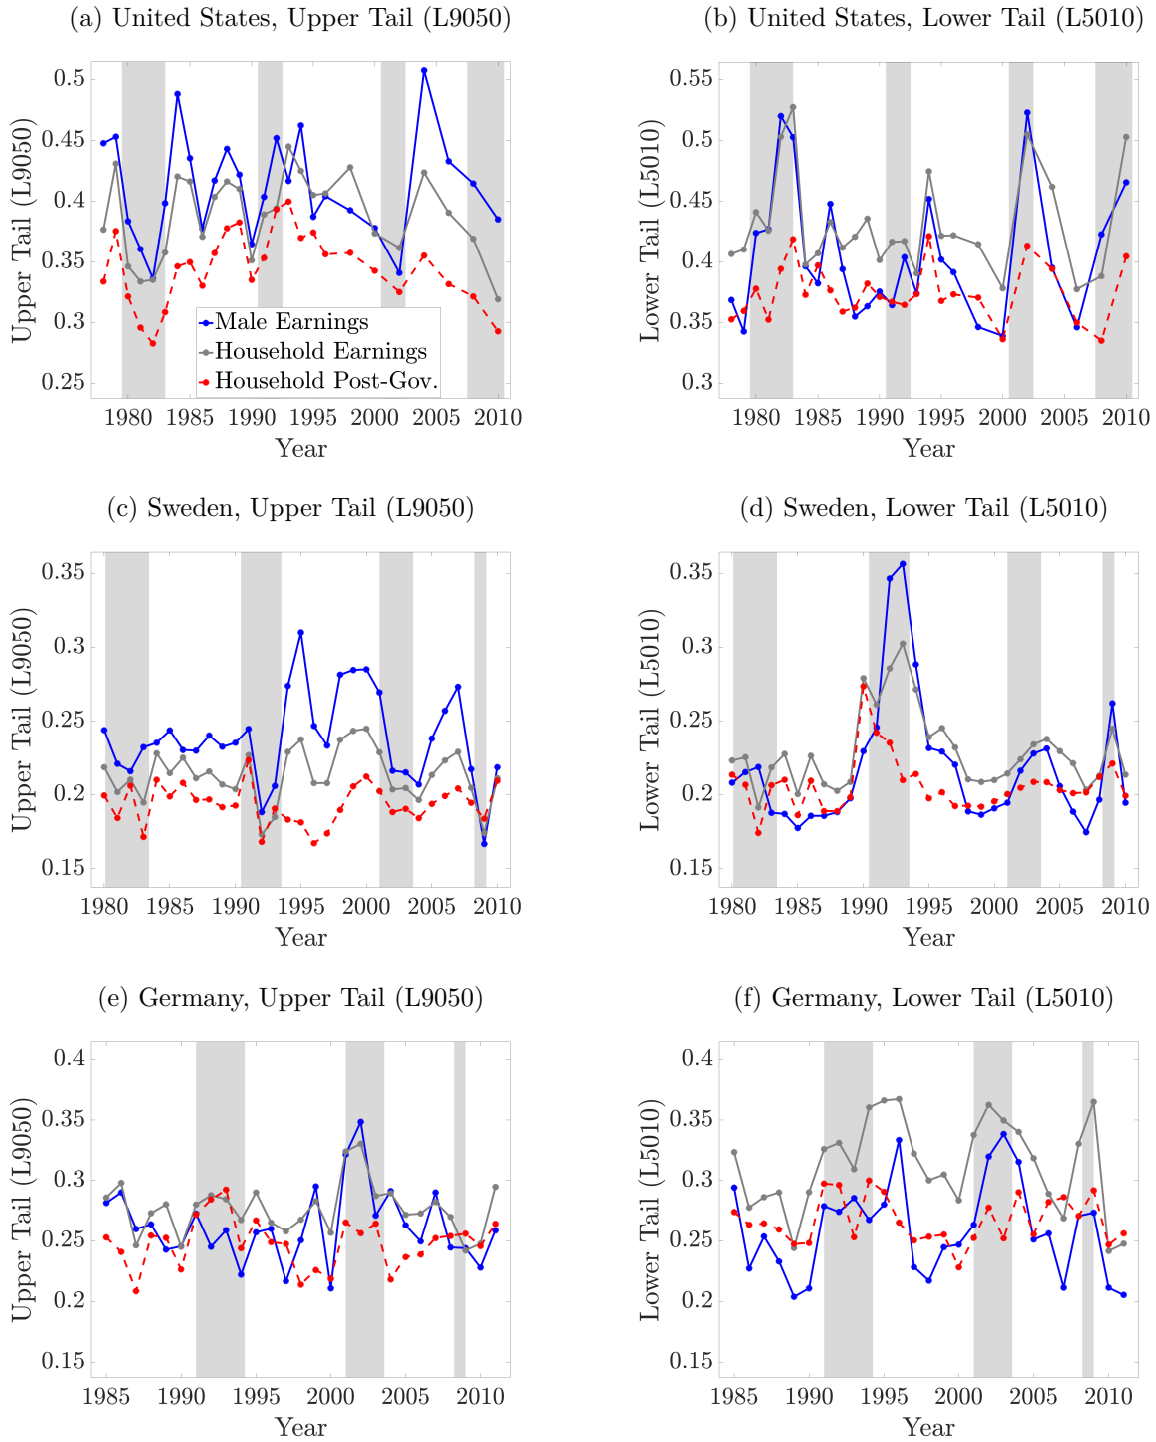
H Insurance

H.1 Household Insurance

For the ranking we first remove year and household head age fixed effects from household earnings. Then, for each year t , we calculate pre-episode earnings as average earnings over the last five years ($t - 5$ to $t - 1$). We only include households where both the head and the spouse separately satisfy the minimum income criteria. We then group households into 3 income groups based on their pre-episode earnings belonging to the 1st - 20th, 21st - 80th, 81st - 100th percentile of the earnings distribution. This ranking is done separately by age (3 groups; <35, 36-45, 46-59) and region (7 regions; capital, south, west, east, central, mid-north and north).

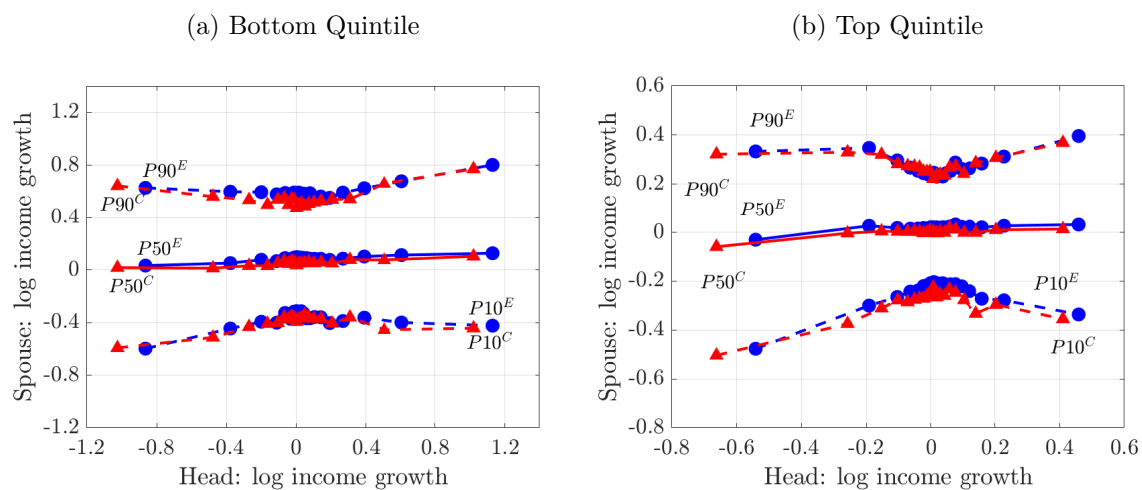
Figure H.3 shows percentiles of the conditional distribution of household log earnings growth for several quantiles of head log earnings growth separately for expansion (blue) and contraction (red) years. The 45°-degree line is the reference point for perfectly correlated incomes of head and spouse. Notice that the larger the income losses (gains) of the head, the more the distribution of household income changes tends to lie above (below) the 45°-degree line, implying that there is insurance at the household level. To the extent that spousal income changes are uncorrelated this insurance is rather passive in the sense that it is explained by having two as opposed to one income, which implies that changes of one of the two incomes less than proportionately translate into changes of the overall income.

Figure H.1: Tails of Short-Run Earnings Growth: United States, Germany (SOEP), and Sweden



Note: Linear trend removed, centered at sample average. Shaded areas indicate recessionary periods (see footnote 12). Horizontal gray line in the right axis of the left panel indicates zero (symmetry) reference line. Year denotes ending year in the growth rate calculations.

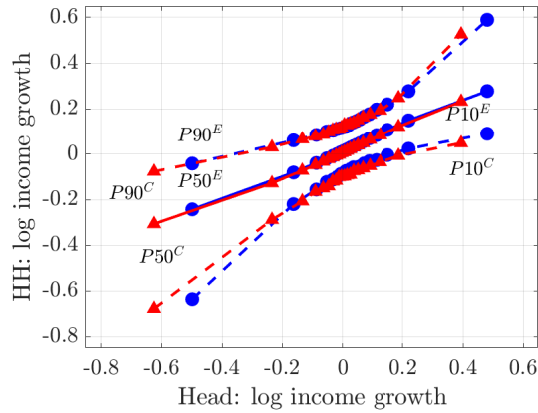
Figure H.2: Spousal Response to Head's Income Change, Bottom and Top Quintiles: Sweden



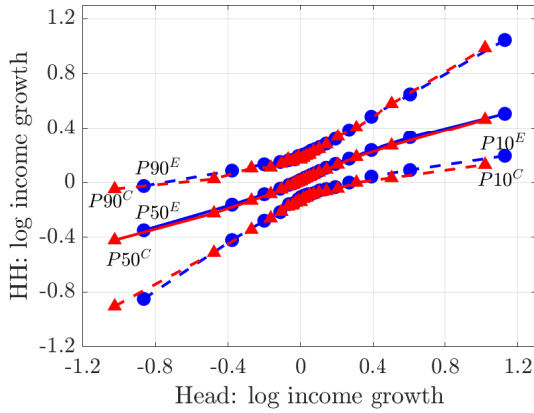
Note: Figure shows percentiles of spouses log earnings growth against household head log earnings growth. The three panels group households based on average household earnings over the last 5 years. For each marker, the y-axis shows the 90th, 50th, or 10th percentile of spouse earnings growth, the x-axis shows the median of the corresponding quantile of head earnings growth. Red and blue markers correspond to recession and boom years, respectively. The dashed line is the 45°-line. For details see text.

Figure H.3: Household Income Change against Head's Income Change: Contractions vs. Expansions; Sweden

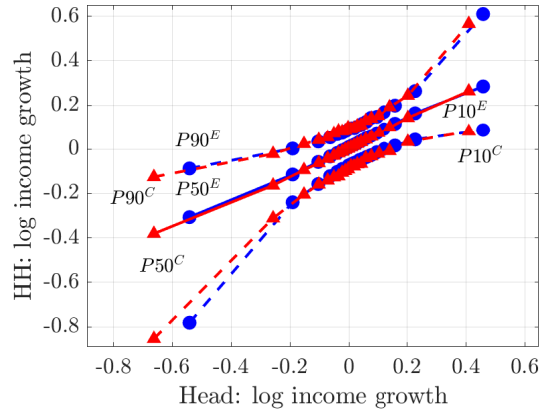
(a) Middle 3 Quintiles



(b) Bottom Quintile



(c) Top Quintile



Note: Figure shows percentiles of household log earnings growth against household head log earnings growth. The three panels group households based on average household earnings over the last 5 years. For each marker, the y-axis shows the 90th, 50th, or 10th percentile of household earnings growth, the x-axis shows the median of the corresponding quantile of head earnings growth. Red and blue markers correspond to recession and boom years, respectively. The dashed line is the 45°-line. For details see text.

Table H.1: Cyclicalitly of Individual Earnings Including Unemployment Benefits in Germany (SIAB)

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

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in an income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3 for SIAB and LINDA, 2 for PSID). The income measures are individual earnings, and individual earnings + unemployment benefits. Differences between estimates in Table I are due to regressions starting in 1981 instead of 1976.

H.2 Tax and Transfer Policies: Additional Results from SIAB

Table H.1 shows additional results for income measures including unemployment benefits from the SIAB.

I Classification of Occupations

The SIAB records 120 occupation groups, which we aggregate to 30 occupational segments according to the KldB88 classification, listed in Table I.1.

Table I.1: Classification of Occupations

Segment	Description
1	Farmers, fishers, gardeners
2	Miners, mineral winners
3	Stone preparers, manufacturers of building materials
4	Potters, glassmakers
5	Chemical workers, plastics processors
6	Paper manufacturers, processors, printers
7	Woodworkers and related
8	Metal producers
9	Mechanics and associated professions
10	Electricians
11	Assemblers and related
12	Textile workers
13	Leather manufacturers, leather and fur processors
14	Nutrition professionals
15	Construction workers
16	Outfitters, decorators, upholsterers
17	Carpenters, modelers
18	Painters and related
19	Quality inspectors and related
20	Engineers, chemists, physicists, mathematicians
21	Technicians, special technical professionals
22	Merchants
23	Clerks, insurance agents, related
24	Traffic and transportation
25	Administration
26	Security
27	Writers, artists
28	Health
29	Social sector, education and related
30	Hairdressers, cleaners, hoteliers

Note: The thirty occupation segments used in the analysis. Segments are based on KldB88 classification of occupations.