American Economic Journal: Macroeconomics 2022, 14(2): 207–242
https://doi.org/10.1257/mac.20190019

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

By Christopher Busch, David Domeij, Fatih Guvenen, and Rocio Madera*

We provide new evidence on business cycle fluctuations in skewed labor income risk in the United States, Germany, Sweden, and France. We document four results. First, in all countries, the skewness of individual income growth is strongly procyclical, whereas its variance is flat and acyclical. Second, this result also holds for continuously employed, full-time workers, 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. (*JEL* E23, E24, E32, H24, I38, J22, J31)

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 the 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 percent—the largest 2-year decline since the Great Depression—whereas the median change in labor earnings was +0.1 percent—slightly positive. The large mean decline was entirely driven by the upper and lower

* Busch: LMU Munich, CESifo (email: christopher.busch@econ.lmu.de); Domeij: Department of Economics, Stockholm School of Economics (email: david.domeij@hhs.se); Guvenen: University of Minnesota, Federal Reserve Bank of Minneapolis, and NBER (email: guvenen@umn.edu); Madera: Southern Methodist University (email: romadera@smu.edu). Benjamin Moll was coeditor for this article. For helpful comments, we thank three anonymous referees, 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), from the Spanish Ministry of Science and Innovation, through the Severo Ochoa Program for Centres of Excellence in R&D (CEX2019-000915-S). Domeij acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation (P16-0252). Guvenen acknowledges financial support from the National Science Foundation (SES-1357874). 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" program (reference: ANR-10-EQPX-17 - Centre d'accès sécurisé aux données, CASD). We also thank Priscilla Fialho for help with the DADS data. Previous versions of this paper were circulated under the title "Asymmetric Business Cycle Risk and Social Insurance."

† Go to https://doi.org/10.1257/mac.20190019 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

tails collapsing during those two years as opposed to a negative aggregate shock pulling down the entire earnings distribution (Guvenen, 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 *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 (*Déclaration Annuelle des Données 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 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 ninetieth percentile and the median (P90–P50) of the earnings growth distribution was twice as large as the gap between the median and the tenth percentile (P50–P10), whereas this ratio had completely flipped by 1994, with the lower tail (P50–P10) growing to twice the size of the upper tail (P90–P50) by 1994 (using equation (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 Guvenen, Ozkan, and Song (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 subpopulations defined by gender, education, occupation, and private/public sector employment. Moreover, the cyclicality of skewness also holds for five-year income changes, which shows that the procyclical swings in skewness are 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 are 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 cyclicality 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 3). 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 cyclicality. 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 downside risk during recessions in each spouse's individual earnings. For example, comparing the cyclicality 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

[1] The PSID and the SOEP are not suitable for the purposes of this analysis because teasing out the *cyclicality* 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, Brown, and Mathiowetz 2001 for evidence from validation studies).

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 postgovernment) household income, 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 postgovernment household earnings measure compared with its pregovernment counterpart. However, this similar effect on skewness in the two countries is driven by different sources: in the United States, the tax-and-transfer system mainly reduces the cyclicality 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 II describes the empirical approach. Section III presents the results for gross individual income for the three main countries. Section IV zooms in on various groups in the population and presents the results on the wages versus hours margin. Section V expands the analysis to households and post-tax-transfer income. Section VI 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 assumption is that shocks to earnings are Gaussian, which implies zero skewness. Restricting attention to the changes in the mean and variance

[2] As we discuss further in Section VB, the results for Germany were mixed. On the one hand, in the SOEP data, the skewness of postgovernment household earnings changes is essentially acyclical. On the other hand, the SOEP data also show some important differences from SIAB data in cyclicality patterns for individuals, which raises some uncertainty about the reliability of this result for Germany.

of income shocks, Storesletten, Telmer, and Yaron (2004) concluded that the variance of income shocks in the US data is countercyclical.[3]

Guvenen, Ozkan, and Song (2014) revisit this question using a large panel dataset on the earnings histories of US males from Social Security Administration (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 three new countries and five 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, Telmer, and Yaron (2004) to allow for skewness fluctuations and analyze the cyclicality 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 cyclicality of higher-order risk in the United Kingdom using panel data from the British Household Panel Survey. They confirm the same finding of strongly procyclical skewness for the United Kingdom since the early 1990s. Similarly, Harmenberg and Sievertsen (2017) document procyclical skewness of individual earnings changes in administrative Danish data. In ongoing work, Friedrich, Laun, and Meghir (2021) corroborate our results and also find procyclical skewness of individual earnings for the case of Sweden, particularly for men and workers in the private sector. In a recent paper, Pruitt and Turner (2020) analyze individual and household-level income dynamics using US tax records from the IRS. They also document procyclical skewness of income changes for both male and household incomes. Unlike in our five 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, Kurmann and McEntarfer (2019) document procyclical skewness of hourly wage changes. They also explicitly show that the share of workers realizing a wage cut increases substantially in recessions. Pora and Wilner (2017) document in French administrative data that the distribution of earnings changes was more negatively skewed in the 2008 recession than in the directly 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. Hoffmann and Malacrino (2019) study data from Italian workers and find the employment margin to play an important role in driving skewness fluctuations in earnings.

[3] Using a similar approach, Bayer and Juessen (2012) studied the cyclicality of the variance in Germany, the United Kingdom, and the United States and found different patterns in Germany and the United Kingdom relative to the United States and attributed it to differences in institutions.

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 et al. 2016; Ilut, Kehrig, and Schneider 2018; Salgado, Guvenen, and Bloom 2019), firm productivity (Kehrig 2015; Salgado, Guvenen, and Bloom 2019), and stock returns (e.g., Oh and Wachter 2018; 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 2016; 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 (2018), 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.

I. 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 online Appendix A. Briefly, in the main analysis we employ four panel datasets corresponding to three different countries: the Panel Study of Income Dynamics for the United States, covering 1976 to 2010;[4] the Sample of Integrated Labour Market Biographies[5] and the German Socio-Economic Panel[6] for Germany, covering 1976 to 2010 and 1984 to 2011, respectively; and the Longitudinal Individual Data Base for Sweden,[7] 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.[8]

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, who make up about 20 percent of the workforce. From

[4] The PSID (PSID 2015) 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.

[5] We use the factually anonymous scientific use file SIAB-R7510, which is a 2 percent draw from the Integrated Employment Biographies data of the Institute for Employment Research (IAB). See vom Berge, Burghardt, and Trenkle (2013) for further description of the data.

[6] See Goebel et al. (2019) for further description of the data (doi: 10.5684/SOEP.v30).

[7] Statistics Sweden (1968–2016)

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

the perspective of our analysis, the SIAB has two caveats: (i) income is top-coded at the limit of income subject to social security contributions, and (ii) individuals cannot be linked to each other, which prohibits identification of households. We deal with (i) by fitting a Pareto distribution to the upper tail of the wage distribution[9] 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),[10] which covers the Swedish population of individuals 16 years old and older. 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 Déclaration Annuelle des Données Sociales,[11] which is described in the online 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 income. The earnings measure from SIAB does not include self-employment income because the dataset lacks information on it.[12] More details of each variable can be found in the online 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 zeros 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 US$1,885 for the United States in 2010.[13] To avoid possible outliers, we exclude the top 1 percent 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

[9] The imputation is done separately for each year by subgroups defined by age and gender. For workers with imputed wages, across years, we preserve the relative ranking within the age-specific cross-sectional wage distribution. The procedure follows Daly, Hryshko, and Manovskii (2015); see online Appendix A.4 for details.

[10] Statistics Sweden (1968–2016)

[11] Insee and Ministère des Finances (2015)

[12] We consider wages in real terms, which we obtain using the CPI. Where not provided with the micro data set, we use the official aggregate series, in the case of the United States, series PCECTPI from FRED, in the case of Germany, from the statistical office; see Destatis (2014b).

[13] For the United States, we use the federal minimum wage. There is no official minimum wage in Sweden or Germany during this period. For Germany, we follow Fuchs-Schündeln, Krueger, and Sommer (2010) and take a minimum threshold of €3 (in year 2000 euros) for the hourly wage. For Sweden, the effective hourly minimum wage via labor market agreements was around SKr75 in 2004 (Skedinger 2007). For other years, we adjust the minimum wage using the growth rate in the mean real wages.

earnings, we mean a worker's compensation from 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 online 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,[14] who satisfy the age criterion and household income that satisfies the income criteria. At the household level, we analyze pre- and postgovernment earnings. Pregovernment earnings is defined as the sum of gross labor earnings earned by the adults in the household. Postgovernment earnings is constructed by adding taxes and transfers.

II. 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 $m\left[\Delta_s y_t\right]$, 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.[15] 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 throughout the entire sample period.[16]

Our primary measures for volatility and skewness are quantile-based, which have important advantages over standardized moments (the 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 ninetieth and tenth 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:

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

[14] In 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.

[15] We 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.

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

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 the right tail:

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

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 percent 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 percent but also the top 20 percent 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 moment, 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.[17] 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 (panel A of Figure 2). 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:

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

The key parameter of interest is β^m, which measures the cyclicality 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.[18]

III. 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, Ozkan, and Song 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 United States. To give an example of these differences, consider the fact that only 10.7 percent of US workers are unionized and only 11.9 percent are covered by trade union agreements, whereas the corresponding fractions are 18.1 percent and 57.6 percent, respectively, in Germany, and 67.3 percent and 89 percent, respectively, in Sweden.[19]

A second question we ask is whether the patterns of cyclicality found by Guvenen, Ozkan, and Song (2014) using US administrative earnings data from the Social Security Administration 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

[17] We make two adjustments to NBER and ECRI classifications. For the United States, 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 United States and Germany during these years, as seen in Figure 1.

[18] The time series of US real GDP is retrieved from FRED (series GDPC1), for Sweden is from Statistics Sweden series GDP: Expenditure Approach (ESA95), for Germany is from Destatis (series Fachserie 18, Reihe 1.2 and Reihe S.27). Later, we also use GDP data for France, which are retrieved from Eurostat (series NAMQ10).

[19] OECD (2016). The reported numbers are for 2013.

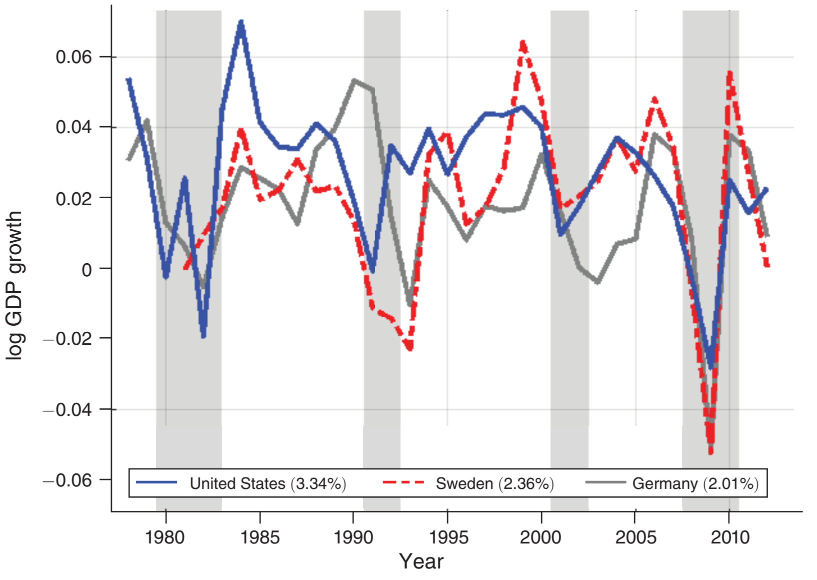

FIGURE 1. ANNUAL LOG GDP GROWTH: UNITED STATES, GERMANY, AND SWEDEN

Notes: 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 percent. The series for Germany corresponds to West Germany up to and including the 1990–1991 change and to (unified) Germany from 1991–1992 on.

countercyclical variance of shocks (e.g., Storesletten, Telmer, and Yaron 2004). This raises the question: is it the differences in the datasets or in the methodologies (or both) that account 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, Ozkan, and Song (2014) were not able to explore potential variation in the cyclicality patterns by gender, education, occupation, private versus public sector employees, among others, which we address starting in this section.

A. *Cyclicality 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).[20] 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

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

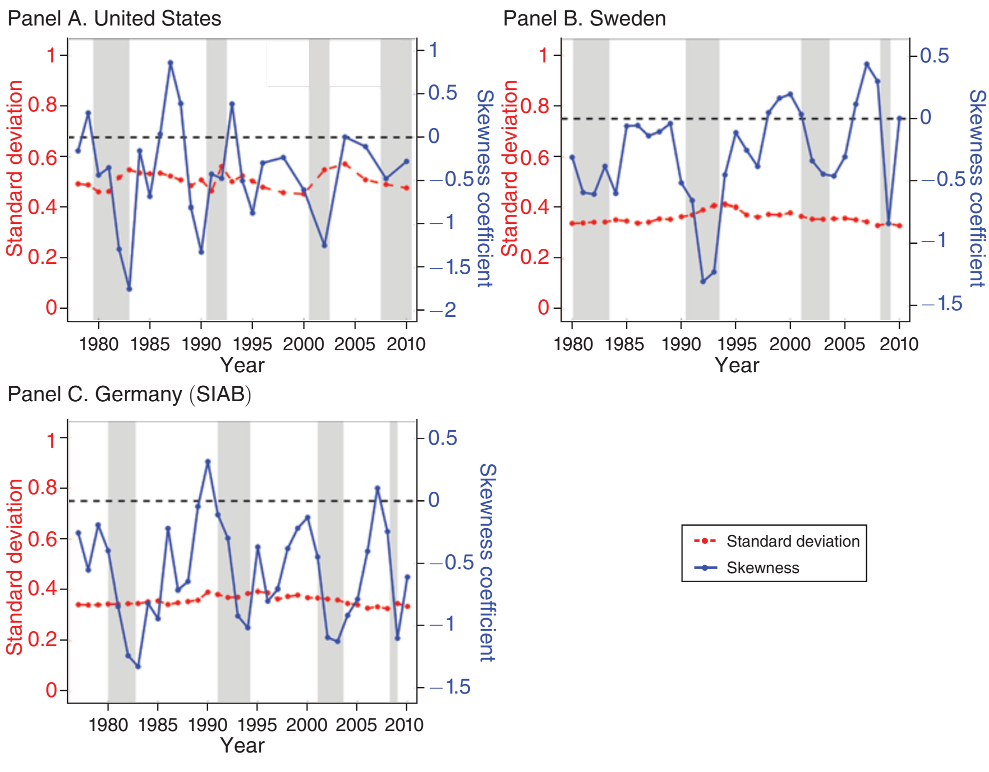

Figure 2. Standard Deviation and Skewness of Short-Run Income Growth: Males

Notes: Linear trend removed, centered at sample average. Shaded areas indicate recessionary periods (see footnote 13). Year denotes ending year in the growth rate calculations.

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 co-move 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 cyclicality and compare it across countries, we use the regression framework described above in (3). Table 1 reports the cyclicality coefficient, β^m, for four quantile-based moments—L9010, Kelley Skewness, L9050, and L5010—separately for each gender and the three countries.[21] 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-statistics

[21] We 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 zeros in income (Table C.1 in online 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 online Appendix C.2). In both cases, we find the same substantive patterns described here.

Table 1—Cyclicality of log Annual Income Change Moments: Before-Tax/Transfer Individual Income

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

Notes: Each cell reports the cyclicality 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 (equation (3)). Newey-West t-statistics are in parentheses.

less than 1.4 for both genders. These estimates confirm the acyclicality of dispersion in the United States that we saw in Figure 2. The same acyclicality is also seen in the bottom two panels for Germany and Sweden with small point estimates that are statistically insignificant.[22]

Cyclicality 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 skewness coefficient in Figure 2, which is probably not surprising but still reassuring.[23] In the PSID, Kelley skewness drops significantly during the 1980s double-dip recession, falling from 0.15 for the 1979–1980 change to below –0.2 for the 1982–1983 change, as well as the two recessions in the twenty-first 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 Guvenen, Ozkan, and Song (2014).

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

[22] All regression results in the paper based on SIAB data are robust to various sensitivity checks we conducted to address issues of top-coding and a structural break in the wage variable. See online Appendix F for details.

[23] To reduce the number of figures for readability, we have moved the analogous figure for females to the online Appendix.

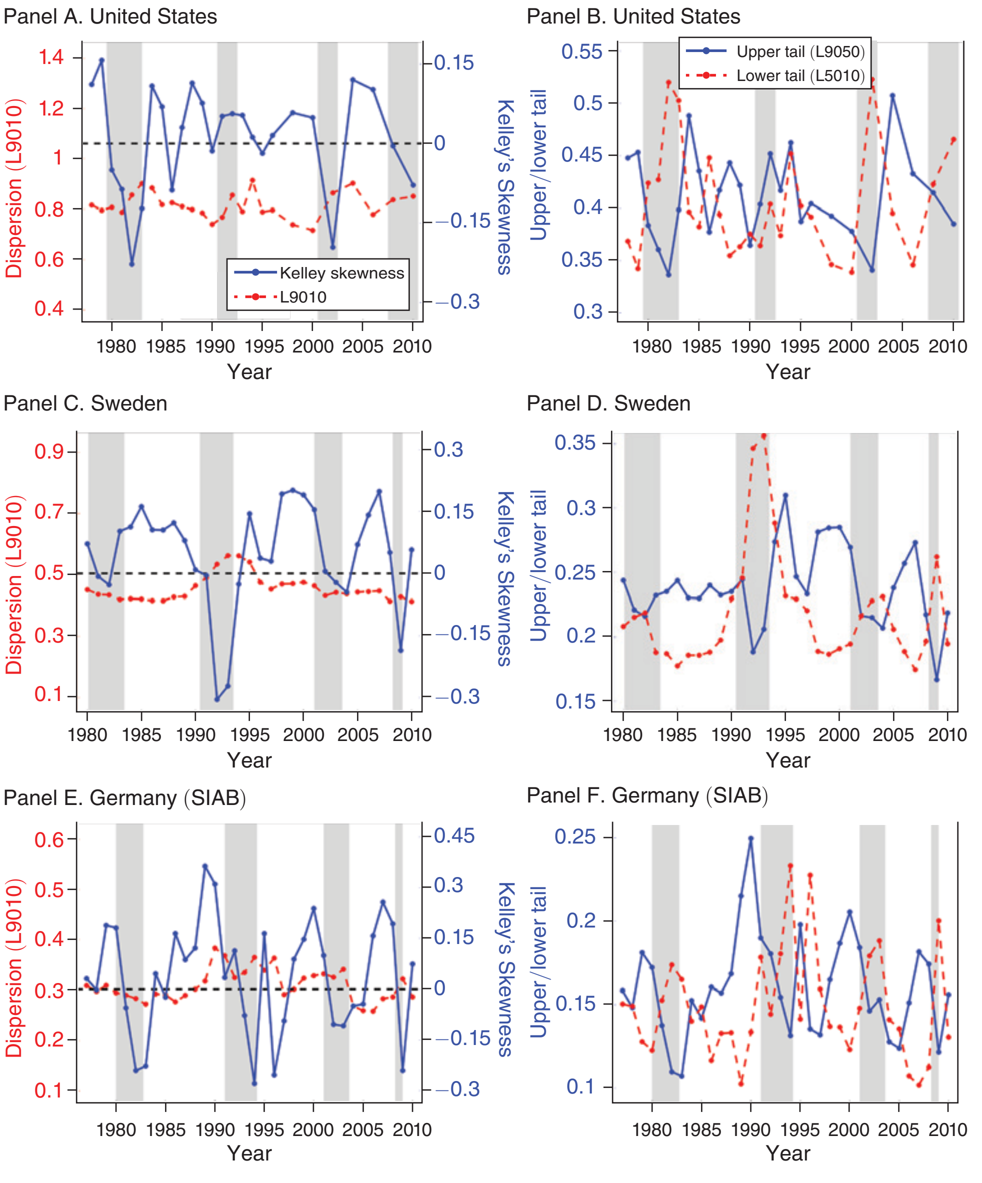

Figure 3. L9010, Skewness, and Tails of Short-Run Income Growth: All Males

Notes: Linear trend removed, centered at sample average. Shaded areas indicate recessionary periods (see footnote 13). 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.

Germany, the Kelley measure swung from 0.31 in 1989–1990 down to –0.28 in 1993–1994, 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 1989–1990 period, with the remaining one-third accounted for by the lower tail (L5010), these ratios completely flipped by the end of the recession (1993–1994), with L5010 growing to account for almost two-thirds (64 percent) and L9050 shrinking to one-third of L9010. Notice that the 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 skewness going from 0.08 to –0.31 between the same two years (with the share of the lower half, L5010, rising from 46 percent to 66 percent).

In Table 1, the cyclicality coefficients for males in column 2 are all positive and statistically significant at the 0.1 percent level, confirming the strong procyclicality of Kelley skewness. The estimated β^{Kelley} for males is 2.25 for the United States, 3.64 for Sweden, and 5.48 for Germany, implying a 2.5-fold larger fall in Kelley in Germany than in the United States for the same 1 percent 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.[24]

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 $\beta^{Kelley} = 3.64$, 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.[25] Finally, skewness is also procyclical for female workers in all countries, with positive and statistically significant coefficients for Sweden and the United States at 1 percent level, and for Germany at 5 percent level.

B. *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?

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

[24] Running the regression with the skewness coefficient instead of Kelley measure yields very similar results.
[25] The corresponding changes in $\mathcal{S}_k$ for the United States and Germany are 0.15 and 0.22, respectively.

right around the beginning of the recession, and they reverse the roles with the start of the expansion. The last two columns of Table 1 report the cyclicality 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-statistics 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-statistic of 1.25).

Another point to notice in Table 1 is that, for all countries, the estimated βs 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 United States 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 United States) 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 1, we observe the same patterns of cyclicality as those of men, whenever the coefficient is significant. In particular, L9050 is procyclical for the United States and Sweden, whereas L5010 is countercyclical for all three economies (though only significant at the 10 percent 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 one-third that of men's in the United States and Sweden and about one-half 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.

C. *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 two 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.[26]

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 relationship with GDP growth, which is especially strong in Sweden and Germany.[27] Online 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 Subpopulations.—We bring additional evidence on the persistence of skewness fluctuations by recognizing that time series data on the entire population are 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 1 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 scatterplot for Kelley skewness in

[26] For example, if five-year changes are computed for every year of the sample, the overlap between observations induces 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 nonoverlapping observations reduces the already modest sample size dramatically. We did estimate the cyclicality regressions using five-year changes and found the same patterns but do not include them because of the concerns outlined here.

[27] In 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.

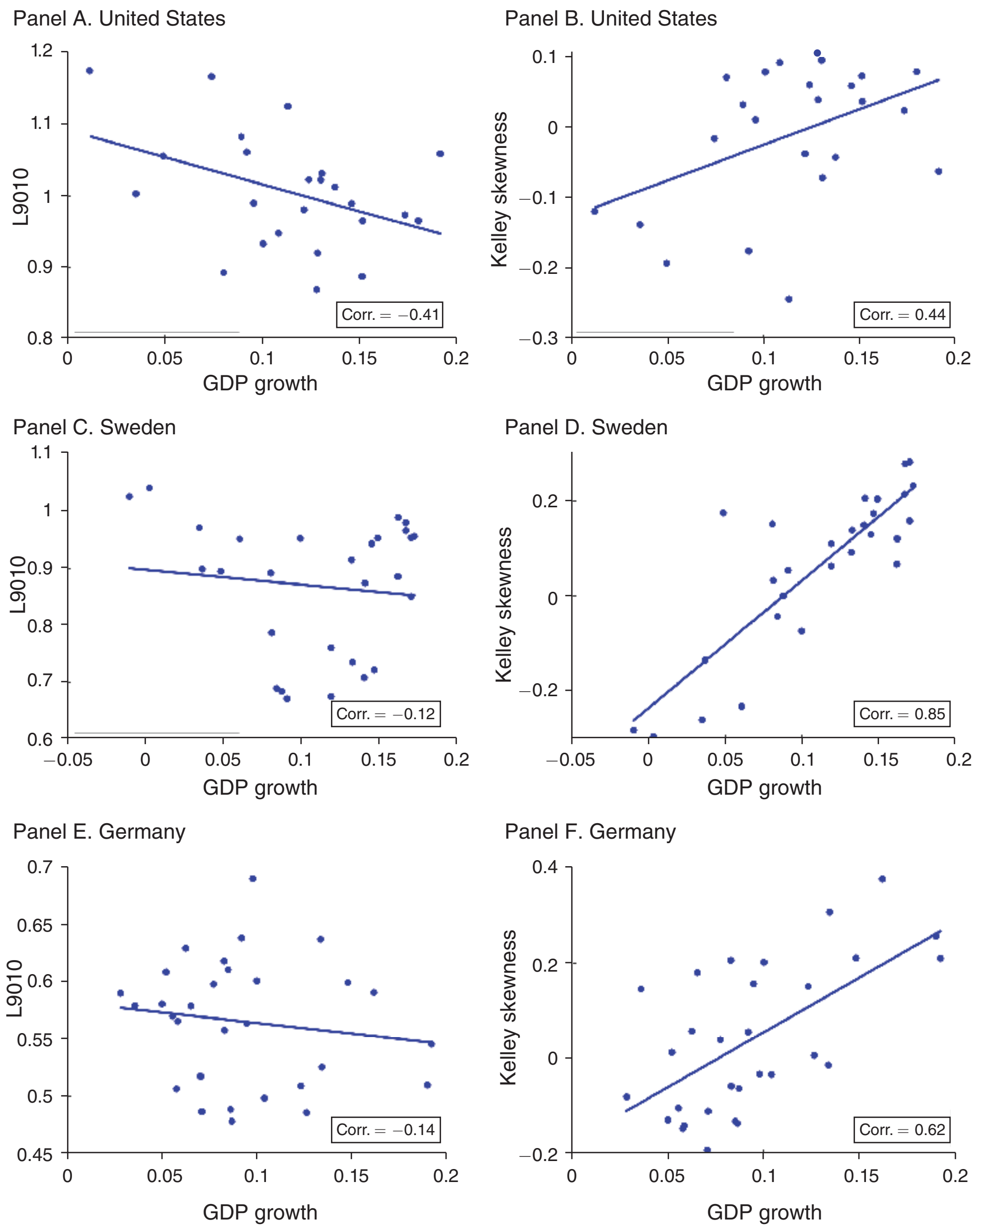

Figure 4. Cyclicality of Dispersion and Skewness of Long-Run Income Changes, Males: United States, Sweden, and Germany (SIAB)

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).

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.

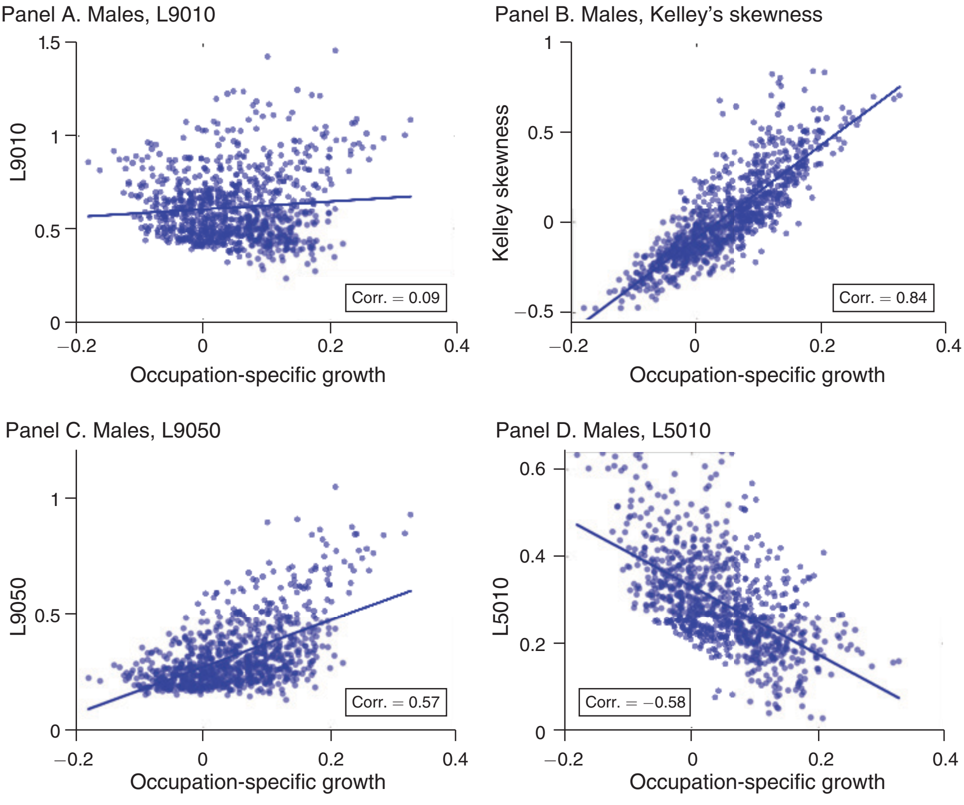

Figure 5. Distribution of Five-Year Income Growth; Occupation-Specific Cycles (SIAB): Males

Notes: 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.

These conclusions are not sensitive to using occupation-specific cycles. Table G.1 in online 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 tenth percentile of correlations. In contrast, the correlations for L9010 range from –0.35 at the tenth percentile to 0.29 at the ninetieth percentiles with a median of –0.12.

Overall, these findings corroborate our main results by showing that the same patterns we observed in the aggregate economy hold, more strongly, at the more disaggregated level. The patterns for females look qualitatively the same; see online Appendix G for details.

IV. 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 to which degree the procyclical fluctuations in skewness are explained by changes in hours worked or by changes in wages, or both.

A. *Heterogeneity across Groups of Workers*

Education and Public versus Private Sector.—We begin by classifying workers (separately for each gender) by educational attainment (college versus noncollege graduates) and, separately, by whether they hold a private or public sector job. The share of male workers who are college educated is 12 percent, 16 percent, and 25 percent, respectively, in Germany, Sweden, and the United States (the analogous numbers for women are 8 percent, 17 percent, and 25 percent). Differences in the size of public sector employment are even larger and also vary significantly between men and women.[28] 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, noncollege 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 online 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).

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 cyclicality regression in (3) separately for each occupation without running into the overlapping observations problem (see footnote 21). Figure 7 shows the estimated βs for L9010 and Kelley skewness for each occupation. As seen in the bottom panel, the estimated βs for Kelley skewness are positive for every occupation and statistically significant for the vast majority of them. As before, βs for dispersion are close to zero for the vast

[28] For men, the share of public jobs is 23 percent in Sweden and 10 percent and 13 percent in Germany and the United States. For women, the corresponding figures are 63 percent, 36 percent, and 18 percent. For these statistics, we define public sector employment as jobs in public administration, health care, and education (sectors that 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.

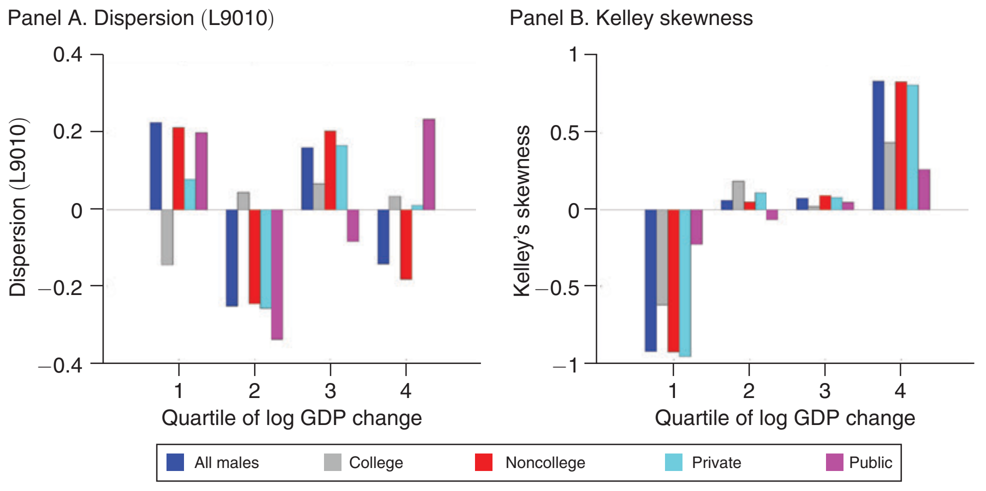

Figure 6. Higher-Order Moments by Quartiles of log GDP Change: Males

Notes: 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.

majority of occupations and are not statistically significant for any of them. Further results for the upper and lower tails are in online Appendix B.2.

B. *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 are scarce because they are often unavailable in administrative datasets (such as the US SSA data, which prevented Guvenen, Ozkan, and Song 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, Brown, and Mathiowetz 2001 for a review of evidence from validation studies). This problem is exacerbated by the fact that teasing out the *cyclicality* 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

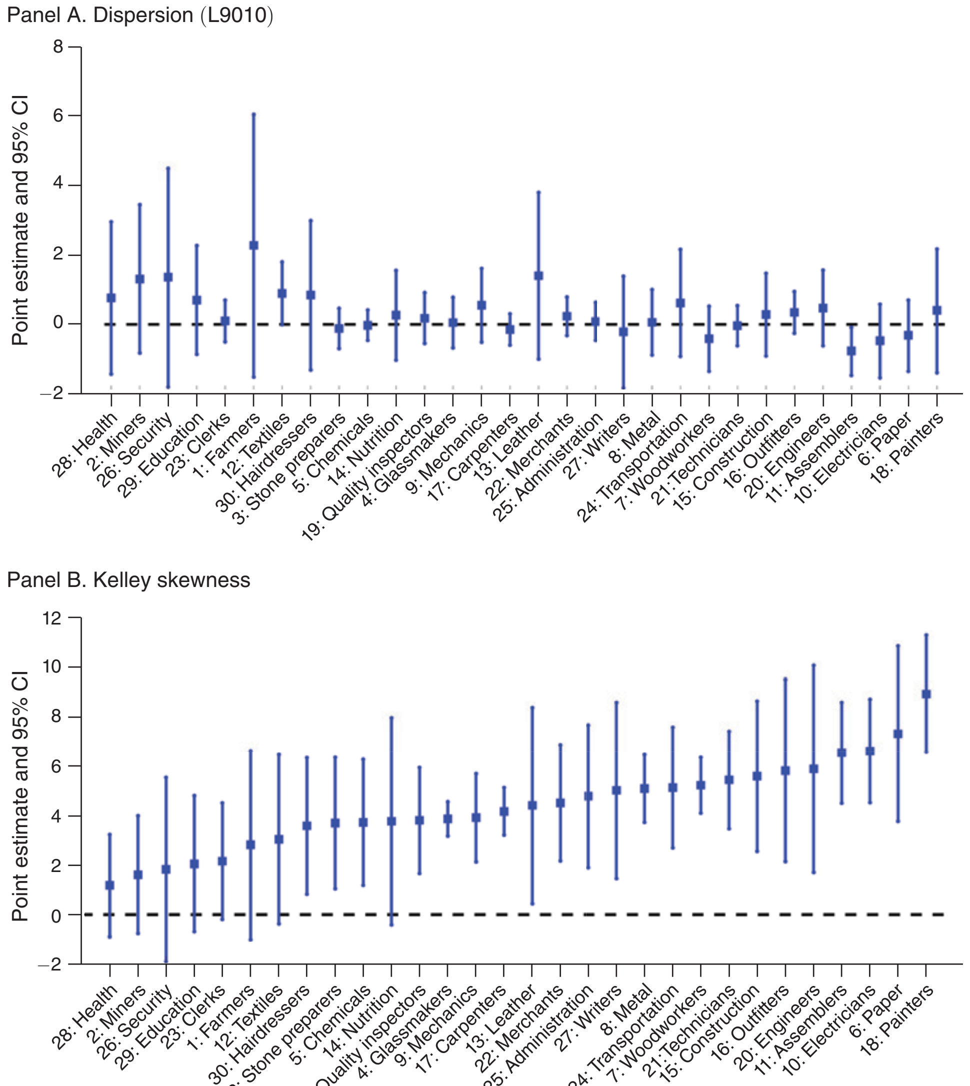

Figure 7. Dispersion and Skewness of Short-Run Income Growth by Occupation: Males (Germany (SIAB))

Notes: 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).

data (the DADS) from 1995 to 2015, which contains hours worked for each employment spell as reported by employers (see online 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, Ludsteck, and Schönberg (2009) and Card, Heining, and Kline (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 2 reproduces the estimated βs for the baseline sample from Table 1. Row 2 reports the corresponding βs 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 percent of males (73 percent of women) are in the full-time sample.[29] Naturally, the dispersion of earnings changes is wider than that of wage changes, which is reflected by the point estimates on the tails (last two columns), which are about half as big for wage changes. 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 2 consecutive years. For these workers, not only changes in hours but also changes in daily wages should be smaller than for the previous sample.[30] 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 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 2 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

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

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

Table 2—Cyclicality of log Annual Income Change Moments for Males: Individual Income versus Daily Wages, Germany (SIAB) and Sweden (LISA)

	L9010	Kelley	L9050	L5010
Germany				
Earnings	0.15	5.48	0.95	−0.80
(Baseline sample)	(0.36)	(5.80)	(3.14)	(−4.11)
Daily wages	−0.09	4.73	0.30	−0.39
(Full-time workers)	(−0.54)	(6.31)	(3.77)	(−3.20)
Daily wages	−0.12	4.98	0.28	−0.40
(Establishment stayers)	(−0.81)	(5.78)	(3.29)	(−3.20)
Sweden				
Earnings	−0.06	3.64	0.87	−0.94
(Baseline sample)	(−0.73)	(4.34)	(4.30)	(−3.81)
Earnings	−0.12	1.78	0.16	−0.28
(Establishment stayers)	(−4.53)	(7.43)	(5.34)	(−7.90)

Notes: 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 who work full-time for at least 50 weeks in both years for which the change is calculated.

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 allow us to explore the roles of days worked versus 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.[31] Table 3 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 cyclicality coefficient on Kelley skewness of similar magnitude (4.81) as for Germany (Table 1), while L9010 is mildly cyclical, driven by L5010 being more cyclical than L9050.

Second, as seen in the second and third rows, the skewness of both changes in hourly wages and hours worked displays significant procyclicality, with coefficients of 2.71 and 2.60, respectively. However, one important difference between the two is seen in the tails: whereas the cyclicality coefficients for L9050 are similar for wages and hours (0.28 and 0.31), the left tail of the wage growth distribution is less countercyclical (–0.45) than that of hours (–0.74). This is consistent with downward rigidity of wages, making hours a more elastic margin to adjust for employers.

[31] Given the available data from the DADS, we use the years 1995–2015 in the analysis, which gives 20 years for which we can estimate 1-year changes. The standard deviation of log GDP growth over that time period is 1.49 percent.

Table 3—Cyclicality of Hours Worked versus Hourly Wages; France (DADS): Males

	L9010	Kelley	L9050	L5010
Baseline sample				
Earnings	−0.50	4.81	0.72	−1.22
	(−1.94)	(6.41)	(2.98)	(−9.60)
Hourly wages	−0.17	2.71	0.28	−0.45
	(−0.93)	(7.57)	(2.31)	(−5.09)
Hours worked	−0.43	2.60	0.31	−0.74
	(−1.65)	(4.34)	(3.54)	(−2.98)
Subsample A. Full-time, establishment stayers				
Earnings	−0.39	3.55	0.09	−0.48
	(−7.96)	(7.94)	(1.56)	(−11.36)
Hourly wages	−0.20	3.23	0.19	−0.38
	(1.11)	(6.24)	(1.54)	(−4.72)
Hours worked	0.00	0.67	0.02	−0.02
	(0.01)	(0.18)	(0.10)	(−0.06)
Baseline excluding full-time workers				
Earnings	−1.02	3.54	1.84	−2.86
	(−2.82)	(6.16)	(3.80)	(−10.17)
Hourly wages	−0.20	1.83	0.34	−0.55
	(−1.26)	(4.09)	(2.70)	(−3.29)
Hours worked	−0.90	3.57	1.96	−2.86
	(−1.92)	(5.73)	(3.25)	(−11.81)

Notes: 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 excluding full-time workers* are those workers who are not in 50 weeks full-time employment in either $t-1$ or t.

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 who stay in the same establishment and the rest of the baseline sample. The results are in the middle and bottom panels of Table 3. Earnings changes display strong procyclicality but slightly smaller than the baseline ($\beta^{Kelley} = 3.55$); however, almost all of it is now due to wages (3.23) and almost none from hours (0.67 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, whereas the opposite pattern emerges for less strongly attached workers.

V. 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.

A. *Within-Family Insurance*

In Table 4, the first row of each panel reports the estimated βs for the same moments but now using household earnings, which can be compared to their counterparts for individual earnings in Table 1. 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 L9010 acyclical. As for magnitudes, the estimated coefficient for Kelley skewness of households falls in between the coefficients reported for males and females in Table 1. For example, in the United States, $\beta^{Kelley} = 2.25$ 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.[32] But comparing these coefficients does not tell us much about 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 mimics the baseline sample but in which synthetic spouses have no actual connection to each other and therefore, unlike actual households, cannot respond to each other's earnings shocks. Thus, to the extent that active within-household insurance is present, the cyclicality 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 United States 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 five-year average income. The pairing is done separately for each $(t-1, t)$ time period.

The second row of each panel in Table 4 reports the cyclicality coefficients for these synthetic households. Perhaps surprisingly, we see no evidence of active within-household insurance. For example, in Sweden, the coefficients for Kelley

[32] The comparison between individual and household earnings is less informative for Germany because the former in Table 1 uses SIAB data, whereas the latter is based on SOEP. It turns out 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.

TABLE 4—CYCLICALITY OF EARNINGS GROWTH MOMENTS: ACTUAL VERSUS SYNTHETIC HOUSEHOLDS

	L9010	Kelley	L9050	L5010
United States				
Actual households	0.04	1.91	0.81	−0.78
	(0.15)	(6.57)	(5.93)	(−3.78)
Synthetic households[a]	−0.01	1.59	0.72	−0.73
	(−0.03)	(3.88)	(3.00)	(2.52)
Sweden				
Actual households	−0.02	2.24	0.50	−0.52
	(−0.08)	(3.33)	(4.94)	(−2.00)
Synthetic households	−0.24	1.93	0.35	−0.59
	(−0.83)	(3.33)	(3.23)	(−2.19)
Germany (SOEP)				
Actual households	−1.17	1.79	−0.03	−1.15
	(−3.33)	(2.76)	(−0.12)	(−4.22)
Synthetic households	−1.06	0.83	−0.22	−0.84
	(−3.33)	(1.64)	(−1.18)	(−3.19)

Notes: Each cell reports the coefficient on log GDP change in the cyclicality regression (3). Newey-West *t*-statistics are included in parentheses.

[a] Synthetic households are formed by randomly assigning two workers of opposite genders from the sample conditional on certain observables. For the United States 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.

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 for two randomly paired individuals.[33] 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

[33] The 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.

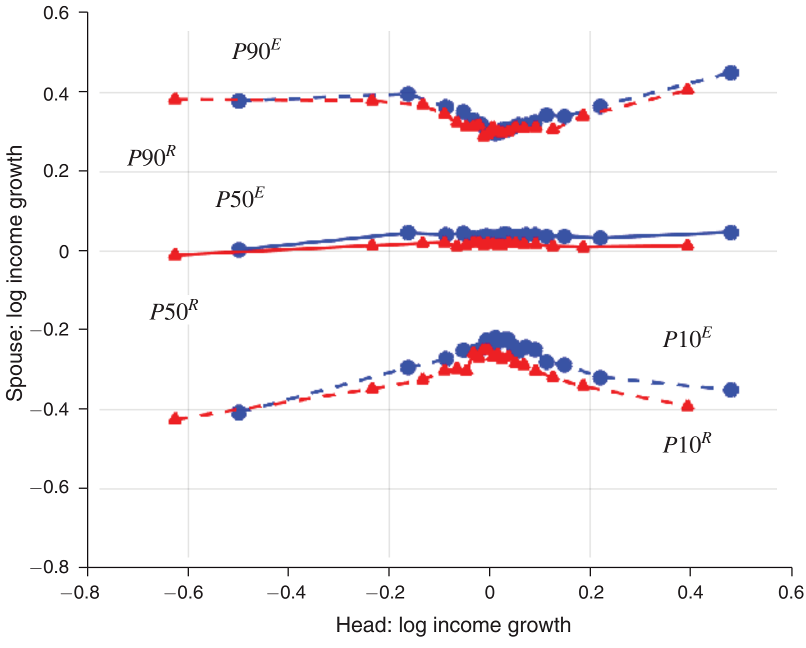

Figure 8. Spousal Earnings Response to Head's Earnings Change over the Business Cycle in Sweden: Households with Earnings between P20 and P80

Notes: Figure shows spouse's log earnings growth against household head log earnings growth for households with five-year average earnings between the twentieth and eightieth 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 ninetieth, fiftieth, or tenth percentile of spouse log earnings growth. Red and blue markers correspond to recession and expansion years, respectively.

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 20 equally sized bins. Then for each of these 20 groups, we calculate the tenth, fiftieth, and ninetieth 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 bow tie–like shape.

Table 5—Summary of Spousal Response

	Spousal response percentiles					
	P10	P50	P90	P10	P50	P90
Head's earnings change	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			

Notes: 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.

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 5 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 tenth 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 1 (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 do 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 5, 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 5 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 gets flatter, although still positive but small, meaning that spousal earnings growth is less positively correlated in recessions than in expansions. Finally, the counterparts of Figure 8 for households in the top and bottom quintiles show the same qualitative patterns discussed here (see Figure H.2 in online 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 online 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 (2020), 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).

B. *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 cyclicality of postgovernment (i.e., post taxes *and* transfers) household earnings growth and compare with the cyclicality of pregovernment (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 6 reports the cyclicality coefficients for pregovernment (first row) and postgovernment (second row) household earnings. 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 postgovernment household earnings growth compared with their pregovernment counterparts in the United States and Sweden and shows essentially no cyclicality in Germany.

Looking at each tail separately, complemented by time series of these moments in Figure H.1 in online Appendix H.1, 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

TABLE 6—CYCLICALITY OF HOUSEHOLD ANNUAL EARNINGS GROWTH MOMENTS: TOTAL EFFECT OF TAXES AND TRANSFERS

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

Notes: 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).

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 pregovernment 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 cyclicality 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 (i) labor market–related policies, (ii) aid to low-income families, and (iii) pension and disability payments.[34] The largest component of labor market–related policies

[34] The 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).

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 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 pregovernment household earnings. The first three rows of each panel in Table 7 report the cyclicality coefficients for each of these three measures. The fourth row reports the cyclicality 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 postgovernment earnings in Table 6 are informative about the effects of income taxes.[35]

There are a few main takeaways. First, in all countries, labor market–related policies seem to be the most effective at reducing the cyclicality of skewness. In the United States, it brings down β^{Kelley} halfway between 1.91 and 1.09 for pre- and postgovernment household earnings reported in Table 6. 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 pregovernment earnings down to 1.14, very close to 0.94 obtained for postgovernment earnings Table 6. In other words, in Sweden, labor transfers is by far the most important component of government policies, including taxation, for reducing the cyclicality 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 postgovernment earnings in Table 6 shows that income taxes play an important role in further reducing the cyclicality of skewness in Germany and the United States but has a much smaller effect in Sweden.[36]

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 that is in fact crucial for an overall assessment is the cost at which this smoothing comes.

[35] Notice 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.

[36] Given 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 cyclicality regressions at the individual level with and without unemployment benefits added in the income measure and found its effect on cyclicality to be modest, consistent with our results here from SOEP, leaving an important room for the tax system (see Table H.1 in online Appendix H.2).

TABLE 7—CYCLICALITY 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)	0.22	1.35	0.64	−0.41
	(1.17)	(4.53)	(4.31)	(−2.70)
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)

Notes: 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 see an extreme example, a 100 percent 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.

VI. 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, thus affecting persistent income changes.

Second, the skewness cyclicality 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 cyclicality of skewness is driven by both hours and wage adjustments, and again the sample of full-time workers displays strong skewness cyclicality, 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 toward 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 semiaggregate 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 cyclicality 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, Konstantinos, Spyridon Lazarakis, and James Malley. 2019. "Cyclical Income Risk in Great Britain." CESifo Working Paper 7594.

Antoni, Manfred, Andreas Ganzer, and Philipp Vom Berge. 1975–2017. "Sample of Integrated Labour Market Biographies (SIAB-R). 1975–2017." Nuremberg: Institute for Employment Research (IAB). https://www.iab.de/389/section.aspx/Publikation/k190903304.

Bayer, Christian, and Falko Juessen. 2012. "The Life-Cycle and the Business-Cycle of Wage Risk—Cross-Country Comparisons." *Economics Letters* 117 (3): 831–33.

Bound, John, Charles Brown, and Nancy Mathiowetz. 2001. "Measurement Error in Survey Data." In *Handbook of Econometrics*, Vol. 5, edited by James J. Heckman and Edward E. Leamer, 3705–843. Amsterdam: Elsevier.

Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera. 2022. "Replication data for: Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E120905V1.

Busch, Christopher, and Alexander Ludwig. 2020. "Higher-Order Income Risk over the Business Cycle." CEPR Discussion Paper 14538.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128 (3): 967–1015.

Constantinides, George M., and Anisha Ghosh. 2016. "Asset Pricing with Countercyclical Household Consumption Risk." *Journal of Finance* 72 (1): 415–60.

Daly, Moira, Dmytro Hryshko, and Iourii Manovskii. 2015. "Reconciling Estimates of Earnings Processes in Growth Rates and Levels." https://www.ecb.europa.eu/pub/conferences/shared/pdf/20151217_4th_conference_hfcs/Session5_Hryshko.pdf.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Where has all the Skewness Gone? The Decline in High-Growth (Young) Firms in the U.S." *European Economic Review* 86: 4–23.

Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg. 2009. "Revisiting the German Wage Structure." *Quarterly Journal of Economics* 124 (2): 843–81.

Eurostat. 2020. "GDP and Main Components (Output, Expenditure and Income)." Eurostat. https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=namq_10_gdp&lang=en (accessed November 4, 2020).

Ferreira, Thiago R.T. 2018. "Stock Market Cross-Sectional Skewness and Business Cycle Fluctuations." Board of Governors of the Federal Reserve System International Finance Discussion Paper 1223.

Friedrich, Benjamin, Lisa Laun, and Costas Meghir. 2021. "Income Dynamics in Sweden 1985–2016." NBER Working Paper 28527.

Fuchs-Schündeln, Nicola, Dirk Krueger, and Mathias Sommer. 2010. "Inequality Trends for Germany in the Last Two Decades: A Tale of Two Countries." *Review of Economic Dynamics* 13 (1): 103–32.

Goebel, Jan, Markus M. Grabka, Stefan Liebig, Martin Kroh, David Richter, Carsten Schröder, and Jürgen Schupp. 2019. "The German Socio-Economic Panel (SOEP)." *Jahrbücher für Nationalökonomie und Statistik* 239 (2): 345–60.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski. 2016. "Redistribution and Social Insurance." *American Economic Review* 106 (2): 359–86.

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122 (3): 621–60.

Harmenberg, Kark, and Hans Henrik Sievertsen. 2017. "The Labor-Market Origins of Cyclical Income Risk." https://www.sv.uio.no/econ/english/research/news-and-events/events/guest-lectures-seminars/job-market/dokumenter/harmenberg-the-labor-market-origins-of-cyclical-income-risk.pdf.

Hoffmann, Eran B., and Davide Malacrino. 2019. "Employment Time and the Cyclicality of Earnings Growth." *Journal of Public Economics* 169 (2): 160–71.

Ilut, Cosmin, Matthias Kehrig, and Martin Schneider. 2018. "Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News." *Journal of Political Economy* 126 (5): 2011–71.

Insee, and Ministère des Finances (DGFiP). 2015. "Panel tous salaries-EDP: 2015". Centre d'Accès Sécurisé aux Données (CASD). datapresentation.casd.eu/10.34724/CASD.86.2507.V1.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3): 697–743.

Kehrig, Matthias. 2015. "The Cyclical Nature of the Productivity Distribution." U.S. Census Bureau Center for Economic Studies Paper CES-WP-11-15.

Kim, Tae-Hwan, and Halbert White. 2004. "On More Robust Estimation of Skewness and Kurtosis." *Finance Research Letters* 1 (1): 56–73.

Kirby, Chris. 1997. "Measuring the Predictable Variation in Stock and Bond Returns." *Review of Financial Studies* 10 (3): 579–630.

Kurmann, André, and Erika McEntarfer. 2019. "Downward Nominal Wage Rigidity in the United States: New Evidence from Worker-Firm Linked Data." U.S. Census Bureau Working Paper CES-19-07.

Mankiw, N. Gregory. 1986. "The Equity Premium and the Concentration of Aggregate Shocks." *Journal of Financial Economics* 17: 211–19.

OECD. 2016. *Economic Policy Reforms 2016: Going for Growth Interim Report.* Paris: OECD Publishing.

Oh, Simon, and Jessica A. Wachter. 2018. "Cross-Sectional Skewness." NBER Working Paper 25113.

Panel Study of Income Dynamics (PSID). 2015. "Public Use Dataset." Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI.

Pora, Pierre, and Lionel Wilner. 2017. "The Individual Dynamics of Wage Income in France During the Crisis." *Economie et Statistique* 494–496: 179–199.

Pruitt, Seth, and Nicholas Turner. 2020. "Earnings Risk in the Household: Evidence from Millions of U.S. Tax Returns." *American Economic Review: Insights* 2 (2): 237–54.

Richardson, Matthew, and James H. Stock. 1989. "Drawing Inferences from Statistics Based on Multiyear Asset Returns." *Journal of Financial Economics* 25 (2): 323–48.

Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom. 2019. "Skewed Business Cycles." NBER Working Paper 26565.

Schmidt, Lawrence D.W. 2016. "Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk." MIT Sloan Working Paper 5500-16.

Skedinger, Per. 2007. "The Design and Effects of Collectively Agreed Minimum Wages: Evidence from Sweden." Research Institute of Industrial Economics Working Paper 700.

Statistics Sweden. 1949–2012. "GDP: Expenditure Approach (ESA95) by Type of Use. Year 1950–2013." Statistics Sweden. http://www.statistikdatabasen.scb.se/pxweb/en/ssd/START__NR__NR0101__NR0101E/NR0102T01A/ (accessed February 6, 2015).

Statistics Sweden. 1968–2016. "Longitudinal Individual Data Base (LINDA)." Statistics Sweden. https://www.scb.se/en/finding-statistics/statistics-by-subject-area/household-finances/income-and-income-distribution/longitudinal-individual-data-base-linda/.

Statistics Sweden. 1990–2019. "Longitudinal Integrated Database for Health Insurance and Labour Market Studies (LISA)." Statistics Sweden. https://www.scb.se/lisa-en.

Statistisches Bundesamt (Destatis). 2006. "Fachserie 18, Reihe S.27, Volkswirtschaftliche Gesamtrechnungen, Inlandsproduktberechnung, Revidierte Vierteljahresergebnisse." Statistisches Bundesamt. https://www.destatis.de/DE/Service/Bibliothek/_publikationen-fachserienliste-18.html (accessed November 5, 2020).

Statistisches Bundesamt (Destatis). 2014a. "Fachserie 18, Reihe 1.2, Volkswirtschaftliche Gesamtrechnungen, Inlandsproduktberechnung Vierteljahresergebnisse." Statistisches Bundesamt. https://www.statistischebibliothek.de/mir/receive/DESerie_mods_00000186 (accessed November 4, 2020).

Statistisches Bundesamt (Destatis). 2014b. "Preise, Verbraucherpreisindizes für Deutschland, Lange Reihen ab 1948." Statistisches Bundesamt. https://www.statistischebibliothek.de/mir/receive/DEHeft_mods_00045773 (accessed November 5, 2020).

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. "Cyclical Dynamics in Idiosyncratic Labor Market Risk." *Journal of Political Economy* 112 (3): 695–717.

TNS Infratest Sozialforschung GmbH. 1984–2013. "German Socio-Economic Panel (SOEP), Data of the Years 1984–2013." Socio-Economic Panel (SOEP). doi:10.5684/soep.v30.

U.S. Bureau of Economic Analysis. 1947–2021. "Personal Consumption Expenditures: Chain-Type Price Index (PCECTPI)." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/PCECTPI (accessed February 6, 2015).

U.S. Bureau of Economic Analysis. 1947–2021. "Real Gross Domestic Product (GDPC1)." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/GDPC1 (accessed February 6, 2015).

U.S. Department of Labor. 1938–2009. "History of Federal Minimum Wage Rates under the Fair Labor Standards Act, 1938–2009." U.S. Department of Labor. https://www.dol.gov/agencies/whd/minimum-wage/history/chart (accessed February 6, 2015).

vom Berge, Philipp, Anja Burghardt, and Simon Trenkle. 2013. "Sample of Integrated Labour Market Biographies Regional File 1975–2010 (SIAB-R 7510)." Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB). http://doku.iab.de/fdz/reporte/2013/DR_09-13_EN.pdf.