Earnings Dynamics and Income Insurance in Germany: A Cohort View*

Niklas Isaak†

Robin Jessen‡

1st February 2024

Abstract

Using data from the Tax Payer Panel (TTP), we document various measures of earnings risk over time and the lifecycle for men and women in Germany during the years 2001-2016. For men, volatility over the lifecycle is U-shaped, while skewness decreases and kurtosis increases with age. In contrast, for women volatility and kurtosis are clearly highest during the childbearing age, while skewness reaches its highest point later in life. During the Great Recession both men and women experienced a large drop in skewness with a larger drop for men. A cohort analysis reveals that, at a given age, women and men born later face higher volatility, but also higher skewness. Strikingly, while men in all cohorts faced a large drop in skewness during the Great Recession, younger women appear not to have been affected by it. We then leverage the German Socio-Economic Panel (SOEP) to study insurance against risk. The household greatly reduces earnings volatility for women and reduces kurtosis for both men and women. In contrast, the tax-transfer system reduces volatility, but does not offer any insurance against higher-order income risk.

JEL: D31, J31, H31

Keywords: Wage risk, Social insurance, Self-insurance, Progressive Taxation

*We are grateful to Almuth Balleer, Charlotte Bartels, Peter Haan, Johannes König, Christoph Schmidt, and seminar participants at RWI, the 2022 Meeting of the European Society for Population Economics in Calabria, the 2022 Annual Congress of the International Institute of Public Finance in Linz, the 2023 IIAE Annual Conference in Olso, and the ECINEQ Meeting 2023 in Aix-en-Provence for helpful comments.

†RWI, Zinnowitzer Str. 1, 10115 Berlin, Germany (e-mail: nisaak@rwi-essen.de); Department of Economics, Freie Universität Berlin.

‡RWI, Zinnowitzer Str. 1, 10115 Berlin, Germany (e-mail: rjessen@rwi-essen.de), IZA

1 Introduction

In this paper we use administrative data, the German Taxpayer Panel (TPP), for the years 2001-2016 to describe the distribution of earnings shocks for men and women from different birth cohorts in Germany. Earnings risk is a major factor in economic decision-making and welfare. Agents consider earnings risk when making decisions about their labor supply, consumption and savings. We then leverage the German Socio-Economic Panel (SOEP) to study to what extent households and the welfare state help to insure against various dimensions of earnings risk.

Our measures of earnings risk are based on the changes in residualized log earnings changes. We document measures of volatility, skewness, and kurtosis of earnings changes over the life cycle and over the business cycle. In line with recent findings for various countries,[1] we find that earnings changes are non-Gaussian. This is important, because deviations from normality have significant welfare implications (De Nardi et al., 2020; Ghosh and Theloudis, 2023). In particular, Our first key finding is that at a given age, for both women and men, earnings changes for younger cohorts, i.e., those born in the late 1970s and later, exhibit higher volatility, which we measure by the P90-P10 differential, than earnings changes of cohorts born earlier. This is in line with recent findings for Italy (Hoffmann et al., 2022), albeit less pronounced. This change in the distribution of shocks alone would imply a welfare cost. However, our second key finding is that differences in higher moments of earnings changes imply a welfare improvement for younger cohorts relative to older cohorts. Earnings changes of younger cohorts exhibit higher skewness, i.e., the upward risk of a large wage increase is higher relative to the downside risk of a large wage decline compared to older cohorts. Moreover, they exhibit lower kurtosis than those of older cohorts, i.e., younger cohorts face fewer extreme tail events.

The life-cycle profiles of earnings risk differ strongly between men and women. For men, we find that earnings volatility is U-shaped, while for women volatility is much higher when entering the labor market and decreases until the age of 40. Skewness has an inverse U-shape over the life-cycle for women and is steadily decreasing for men. Finally, kurtosis somewhat increases over the life-cycle for both genders.

We then show how our measures of earnings risk evolve over the business cycle for the various cohorts. Previous work has shown that skewness of income shocks is strongly

[1] See, e.g., Guvenen et al. (2021) for the US and the articles in the special issue of *Quantitative Economics* on global income dynamics (Volume 13, Issue 4) for various countries.

pro-cyclical (Guvenen et al., 2014; Busch et al., 2022). We document a substantial drop in skewness during the Great Recession, which is larger for men than for women (see also Drechsel-Grau et al., 2022). Importantly, there is no drop in skewness for women born in the 1970s or later. Kurtosis also appears pro-cyclical and, again, the drop in the Great Recession is larger for men.

Next, we show to what extent households and the welfare state help individuals insure against income shocks. To this end we use the Socio-Economic Panel, a long-running representative survey of German households. Its key advantage relative to administrative data is that it contains detailed information on gross and net household income, taking taxes and social security contributions as well as transfers into account. To measure the degree of insurance, we first show to what extent volatility of changes in residualized log household pre-government and post-government incomes differ from the volatility of residualized log earnings changes. For women, we find that the household reduces income earnings volatility substantially and the welfare state further reduces it. For men, the household does not offer substantial income insurance. In line with much of the literature (e.g., Busch et al., 2022; De Nardi et al., 2021) we find no evidence for an added-worker effect. Instead, income insurance via the household happens via income pooling. Due to the absence of an added-worker effect and because women on average face higher earnings risk than men, the household offers little insurance for men. In contrast the welfare state reduces volatility substantially.

Turning to higher-order risk, our key result is that the welfare state offers no protection beyond that of the household. As we show, the tax-transfer system scales earnings shocks and therefore reduces their variance. In contrast, skewness and kurtosis are unit-free, and therefore unaffected by scaling. On the other hand, the household tends to reduce skewness of women's income changes, which by itself has a negative impact on welfare. Further, the household reduces kurtosis of income changes for both women and men, offering protection against tail risks.

As is well understood, persistent income changes have very different implications than transitory income changes, mainly because households can insure against transitory shocks by saving. For this reason it is common to decompose earnings residuals in transitory and persistent components based on a specific income process. A key advantage of this approach is that it allows to quantify to what extent changes in cross-sectional inequality translate into long-run inequality. However, committing to a specific, restrictive income process can heavily bias results (Shin and Solon, 2011). Instead, to gauge to what ex-

tent our results for one-year earnings changes capture trends in persistent rather than transitory shocks, we also show results for five-year changes, which reflect more persistent innovations (Guvenen et al., 2021). We show that our main results hold for both one- and five-year changes.

This paper contributes to a growing strand of recent literature on earnings dynamics. Guvenen et al. (2021) deliver important insights on the non-linearities of male earnings dynamics in the United States. Skewness and kurtosis of earnings changes vary a lot depending on age and income level. Moreover, negative shocks tend to be much more persistent than positive changes for high-income individuals. Guvenen et al. (2022) summarize key facts about income dynamics in various countries.

Pessoa (2021) and Drechsel-Grau et al. (2022) document important non- normalities and non-linearities of individual income dynamics for Germany. Our paper is strongly related, but differs in several key aspects: First, we document in detail how income dynamics differ between cohorts. Second, we extensively investigate insurance against income shocks.

Hoffmann et al. (2022) take a closer look at the development of earnings dynamics over different birth cohorts. Mainly driven by labor market institutions like fixed-term contracts, income risk significantly increased for younger cohorts. For Germany, we do not find a continuing increase in income risk over cohorts. However, for the youngest cohorts in our sample, the millennials, we do find a larger variance in earnings growth at the beginning of their careers compared to previous birth cohorts. We show that this increase in individual earnings risk translates to larger pre- and post-government household income risks.

The remainder of the paper is structured as follows. Section 2 presents the two data sets, the sample selection and the empirical approach used in the analysis and section 3 presents the results. Section 4 concludes.

2 Data and Empirical Approach

We employ two data sets, the German Socio-Economic Panel (SOEP) and the German Taxpayer Panel (TPP). The key advantage of the SOEP is that it contains information on hours and months worked as well as detailed information on the household context and transfers received. The TPP is an administrative data set of tax return data. While it lacks some of the information contained in the SOEP, its main advantage is the larger

number of observations as well as—presumably—less measurement error.

2.1 Socio-Economic Panel (SOEP)

The SOEP is a representative annual household survey of about 15 000 German households. For more information see Goebel et al. (2018) and Schröder et al. (2020). The survey has been conducted since 1984. We use wave 36, which contains retrospective information until 2018. We use information for the years 1991-2018, i.e., all available years since the German reunification. The data contain information on annual earnings, annual hours of work, and months worked. We construct the hourly wage by dividing annual earnings through annual hours and average hours per month by dividing annual hours through months. The construction of the hourly wage leads to denominator bias, a spurious negative correlation between hours of work and hourly wages (Borjas, 1980). Throughout the analysis, we weight observations using household weights.

2.2 German Taxpayer Panel (TPP)

For the cohort analyses, we employ the TPP, an administrative dataset with substantially more observations that allows us to make precise estimates of the moments of earnings growth even after dividing the sample by birth cohort and along the ages or years. The TPP is provided by the Federal Statistical office and uses tax return data in a 5%-sample of the overall collected personal income tax data by the tax authorities. This leaves us with well over 50 million observations for the analysis. However, fewer socioeconomic variables are available and we can only use a shorter time span compared to the SOEP: the TPP covers the years 2001 to 2016. For distributional analyses, the big upside of the TPP is that it contains even the largest incomes without top-coding. On the other hand, though, very small incomes below the income tax threshold do not appear in the TPP unless they have other taxable income from other sources or form a tax unit with a partner over the threshold.

2.3 Variables and income concepts

For our analysis, we use constructed variables from the SOEP's cross-national equivalence file PEQUIV (see Grabka, 2020).

Cohorts — We divide the sample into ten-year (SOEP and TPP) and five-year (TPP) birth cohorts.

Earnings — Earnings refer to an individual's labor earnings in a calendar year. The variable includes bonus payments and payments for overtime work.

Household pre-government income — This variable is the sum of all pre-tax income of all individuals in the household 16 years and older. Besides earnings, it includes capital income and rents as well as payments received from individuals outside of the household such as alimony.

Household post-government income — This variable is the sum of household pre-government income and government transfers minus taxes and social security contributions. Government transfers include both transfers from social security (most importantly public pensions and unemployment benefit) and tax-financed transfers such as child benefits, housing allowance and social assistance.[2] In contrast to administrative tax data such as the TPP, which contain information only on transfers that are relevant for the income tax (child benefits, furlough, pensions, unemployment benefits), the variable additionally includes social assistance. This is a major advantage as it plays an important role for low-income households.

Residualized log income changes — Our measures of risk are based on residualized growth rates. We construct growth rates of variable $x_{i,t}$ for individual i in year t as

$$\Delta_k \ln x_{i,t} = \ln x_{i,t+k} - \ln x_{i,t}, \tag{1}$$

We then construct residualized—or idiosyncratic—growth rates $\Delta_k \ln \tilde{x}_{i,t}$ by regressing growth rates on the full set of interactions of year, age, and sex dummies to account for year and age effects as well as for the different dynamics of women and men. In practice, some of the idiosyncratic changes in earnings might be expected by the individual or might be the result of choices. Then the idiosyncratic variability is an upper bound for actual income risk.

[2] An important component of social assistance is Unemployment Benefit II, a means-tested transfer, which does not depend on previous earnings and covers the so-called social existence minimum.

2.4 Measures of Earnings Risk

Our main measure of income risk is the volatility of idiosyncratic income growth. More robust to outliers than the variance, the volatility, measured as the difference between the 90th and 10th percentile—formally $P90 - P10$—of residualized log income changes, is frequently used in the earnings dynamics literature. Beyond the dispersion, we analyze the third and fourth moment of the earnings growth distribution: skewness and kurtosis. Again, with robustness to outliers in mind, we use measures based on percentiles of the distribution. The Kelley skewness, or formally $S_K = \frac{(P90-P50)-(P50-P10)}{P90-P10}$ describes the relative weight of the upper and lower tail within the $P90 - P10$ volatility (Kelley, 1947). For the kurtosis we use the measure introduced by Crow and Siddiqui (1967). The Crow-Siddiqui measure of kurtosis is calculated using $K_{C-S} = \frac{P97.5-P2.5}{P75-P25}$.

2.5 Persistence

The economic implications of earnings changes and their dispersion may vary quite a lot depending on the persistence of earnings shocks. Individuals can self-insure much better against transitory shocks than against persistent shocks using savings. Hence, we extend the analysis by comparing the earnings risk measures described above of 1-year changes to those of 5-year changes. These short and long term changes can be seen as reflecting transitory and more persistent shocks, respectively.[3] Depending on the horizon of the earnings shocks, the parameter k in Equation 1 is set to either 1 or 5.

In our analysis, we find that the time and life-cycle patterns of one-year and five-year changes are very similar for all measures of risk applied in this study. Thus, the trends that we document appear to be driven to a large extent by the persistent component. Moreover, the insurance via family and government appears to be just as persistent. Our primary results presented in the following sections hence apply to both the 1-year and 5-year changes and are illustrated as 1-year changes. Graphical results for 5-year changes can be found in Appendix A.

[3]Guvenen et al. (2021) provide further intuition on this commonly used approach of investigating persistence and derive the higher-order moments of log change from the random-walk permanent/transitory model.

2.6 Sample Selection and Summary Statistics

Table 1: Descriptive Statistics - SOEP

	Women	Men
Age	43.1	43.5
	(9.14)	(9.10)
HH pre-government income	58 904.3	60 427.6
	(9.14)	(9.10)
HH post-government income	41 147.4	42 216.7
	(21 523.10)	(20 097.46)
Individual labor income	25 743.1	42 774.2
	(16 808.07)	(25 454.10)
Annual hours of work	1 743.3	2 271.2
	(626.48)	(453.27)
Months of work	11.9	11.9
	(.85)	(.65)
Hourly wage	14.8	18.9
	(12.31)	(12.07)
Hours of work per month	147.1	190.7
	(52.22)	(37.26)
N	63 880	77 712

Note: Weigthed with household weights. Means with standard deviations in parentheses. Monetary values in Euro.

We select between 25 and 60 years old individuals. Further, we exclude civil servants, and observations with earnings from self-employment. We drop individuals with unrealistic values for hourly wages and hours of work, i.e., hourly wages below 5 Euro or more than 450 hours of work per month. Finally, we require at least two consecutive observations of hours of work, months worked, earnings and pre- and post-government household income. Table 1 shows summary statistics for the SOEP. In our sample, men have substantially higher average labor incomes than women due to both higher hourly wages and more hours of work per month. The difference in average household incomes of men and women is much smaller because many live in couple households.

3 Results

Using the measure of earnings risk defined in section 3.1, we investigate the dynamics in individual labor income, pre-government household income and post-government household income for women and men.

3.1 Earnings risk over the life cycle and over the years

Before splitting the sample into cohorts, we look at the earnings dynamics of the pooled sample (distinguishing only between women and men) over the life cycle and over the years. Figure 1 shows measures of the second, third and fourth moment of the distribution of residualized earnings growth by age, seperately for women and men. For these moments we use the measures described in Section 2.4 The first panel depicts the volatility. Although men and women start out and end up at very similar levels of volatility of idiosyncratic earnings changes, their life cycle trajectory diverges quite substantially. While men's dynamics follows a U-shaped path with increasing age, women experience a large increase in volatility between age 25 and into the early 30s. This is likely due to childbearing (Kleven et al., 2019), hence the steep decline in volatility between the ages 35 and 40 back towards the path of men.

The upper right panel shows the Kelley skewness. Throughout the working life, men's residualized earnings growth follows a downward trend going from a positive skew early into the career to a skewness of zero around age 50 and then further declining sharply. Women, in turn, reach a negative skew – negative deviations are more pronounced than positive ones – already early in life, which can be explained by mothers who often take a step back from the labor market (dropping out, reducing hours, changing to lower-paid family-friendlier jobs). The flipside of this effect is the strong increase between roughly the ages of 30 and 40, surpassing the skewness in men's earnings dynamics in large parts of the life cycle.

Ultimately, we show the Crow-Siddiqui measure of kurtosis in the lower panel of Figure 1. This measure of excess kurtosis increases slightly throughout men's working life, meaning that with higher age, extreme deviations from the expected earnings path become more likely relative to smaller deviations. Job losses, job changes and switching between full-time and part-time can be reasons for such substantial changes. Similar to the volatility, women's kurtosis of earnings growth increases sharply early in life and

later decreases until age 45, but peaking a little later than the volatility. However, an explanation for this pattern can be found in maternity as well.

Beyond the age of 55, volatility, Kelley skewness and Crow-Siddiqui kurtosis are all very similar between women and men, when children are of less importance and retirement-approaching effects hit both groups likewise.

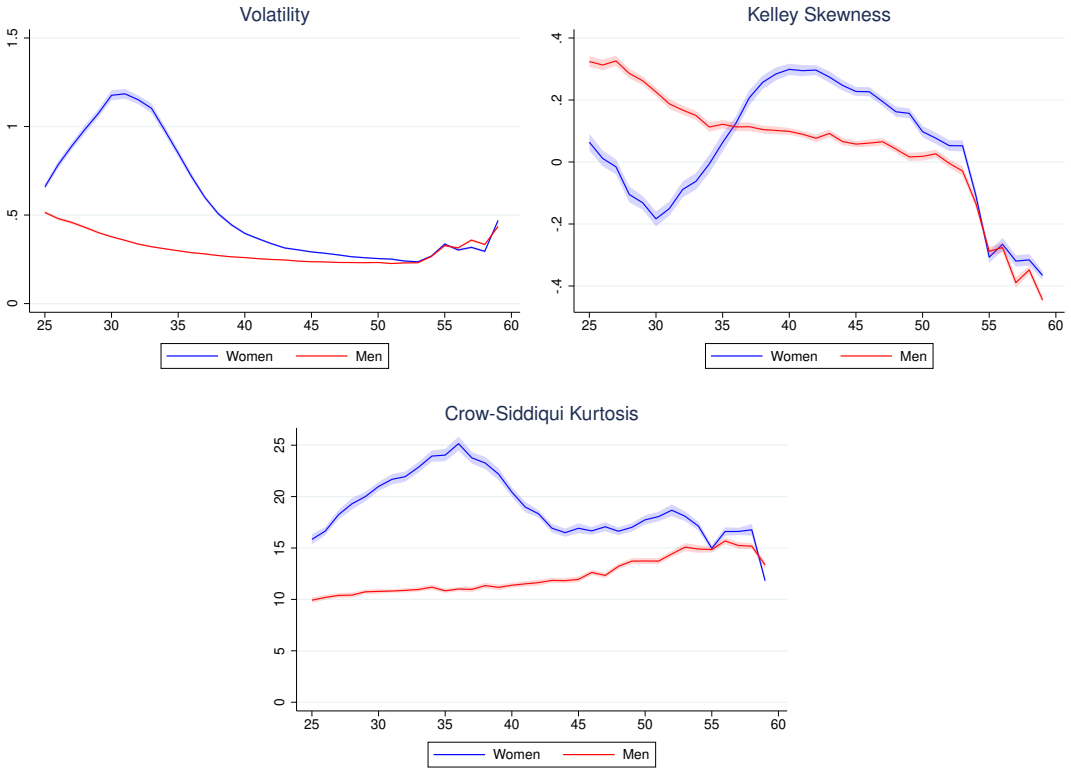

Figure 1: Dynamics of residualized earnings by age

Note: Prime age individuals in the TPP, years 2001-2016

In addition to the life cycle, we analyze the distribution of residualized earnings growth over the years to identify how the business cycle affects its risk. Figure 2 shows the trajectory of our measures for the second, third and fourth moment of residualized labor income growth within the time span that is available in the TPP data. The years 2003 and 2009 are marked as years of recession in Germany. The GDP-loss during the Great Recession in 2009 was far larger than the deficit in 2003 and hence the observable business cycle effects on earnings risk materialize more prominently in 2009.

Due to the typical labor market effects of crises, the volatility increased. Not only did the income stability suffer in the Great recession but, unsurprisingly, the majority of large shocks in the crisis year were negative and led to a negative skew. This drop in skewness that is caused by job-losses and reducing hours or wages (or not increasing them as much as expected) is much sharper for men. In the year following the crisis, the rebound is

similarly strong. In terms of kurtosis, the cyclicality is not as clear as the pro-cyclicality of volatility and skewness. While men experienced a decrease in kurtosis in 2009, women's earnings dynamics became more leptokurtic.

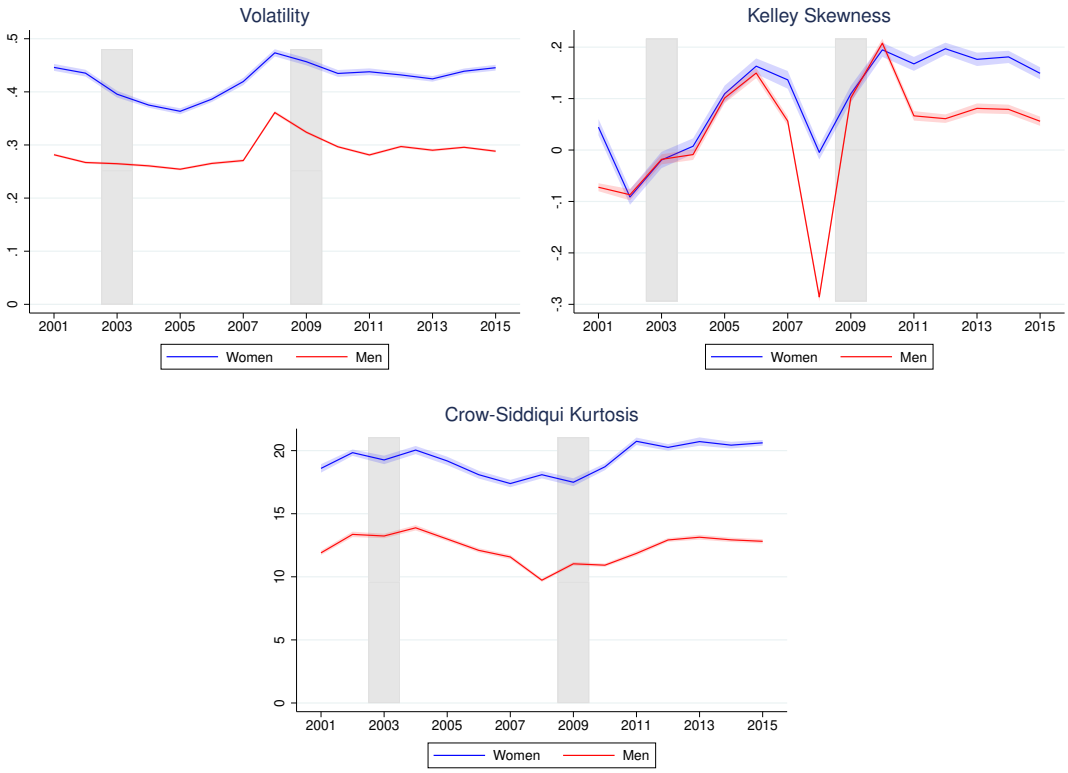

Figure 2: Dynamics of residualized earnings by year

Note: Prime age individuals in the TPP, years 2001-2016

3.2 Earnings Risk by Cohort and Age

The vast data of the TPP allows us to repeat this analysis with the sample split into 10-year-cohorts. Figure 3 shows the second to fourth moments of idiosyncratic changes in individual labor earnings. In the left panels, from top to bottom, the volatility, Kelley Skewness and Crow-Siddiqui Kurtosis are plotted for women, in the right panel for men. Within each panel, every line represents one cohort encompassing 10 birthyears.

Once again, the men's U-shaped earnings volatility trajectory over the life-cycle and the hump early in life of women's earnings volatility are observable. Comparing the cohorts with each other, a visible trend of higher volatility in younger cohorts can be found, as at most ages, the volatility plot of one cohort lies above the next older cohort's. Looking at volatility of earnings growth alone, (risk averse) younger generations in Germany are therefore slightly worse off than older ones, a finding that coincides with a similar result

for Italy in Hoffmann et al. (2022).

Rows two and three of Figure 3 show the higher order risks of earnings growth by cohort and age, each for both women and men, in a similar manner. Conversely to the volatility, it turns out that the earnings growth distribution of younger cohorts tends to exhibit higher skewness than that of older cohorts. Higher skewness implies a lower welfare cost of income risk under CRRA utility. Thus, while younger cohorts face higher income growth variance, it is not clear, if they face higher welfare cost of earnings risk than older ones. Meanwhile, the excess kurtosis of earnings changes follows no clear generational trend and its comparative welfare implications are ambiguous.

3.3 Earnings Risk by Cohort and Year

The different moments of earnings growth risk are affected across all birth cohorts by business cycle effects, see Figure 4. In particular, the Great Recession had a significant impact on the earnings dynamics: The crisis of 2009, which, due to our representation of earnings changes measured at time t as the first difference between $t+1$ and t, appears in 2008, led to a strong increase in the volatility of men's earnings growth across all cohorts. Some cohorts of women did not experience this shock. However, with the surprising exception of the cohort born between 1972 and 1981, women did experience a large decline in skewness during the Great Recession. The procyclicality of skewness is even more prominent for men. Here, every cohort saw a large drop, reducing the welfare. In the following year, the recovery was just as fast, though. The kurtosis of earnings changes was less affected, only the two oldest cohorts of men faced a signifincant decrease.

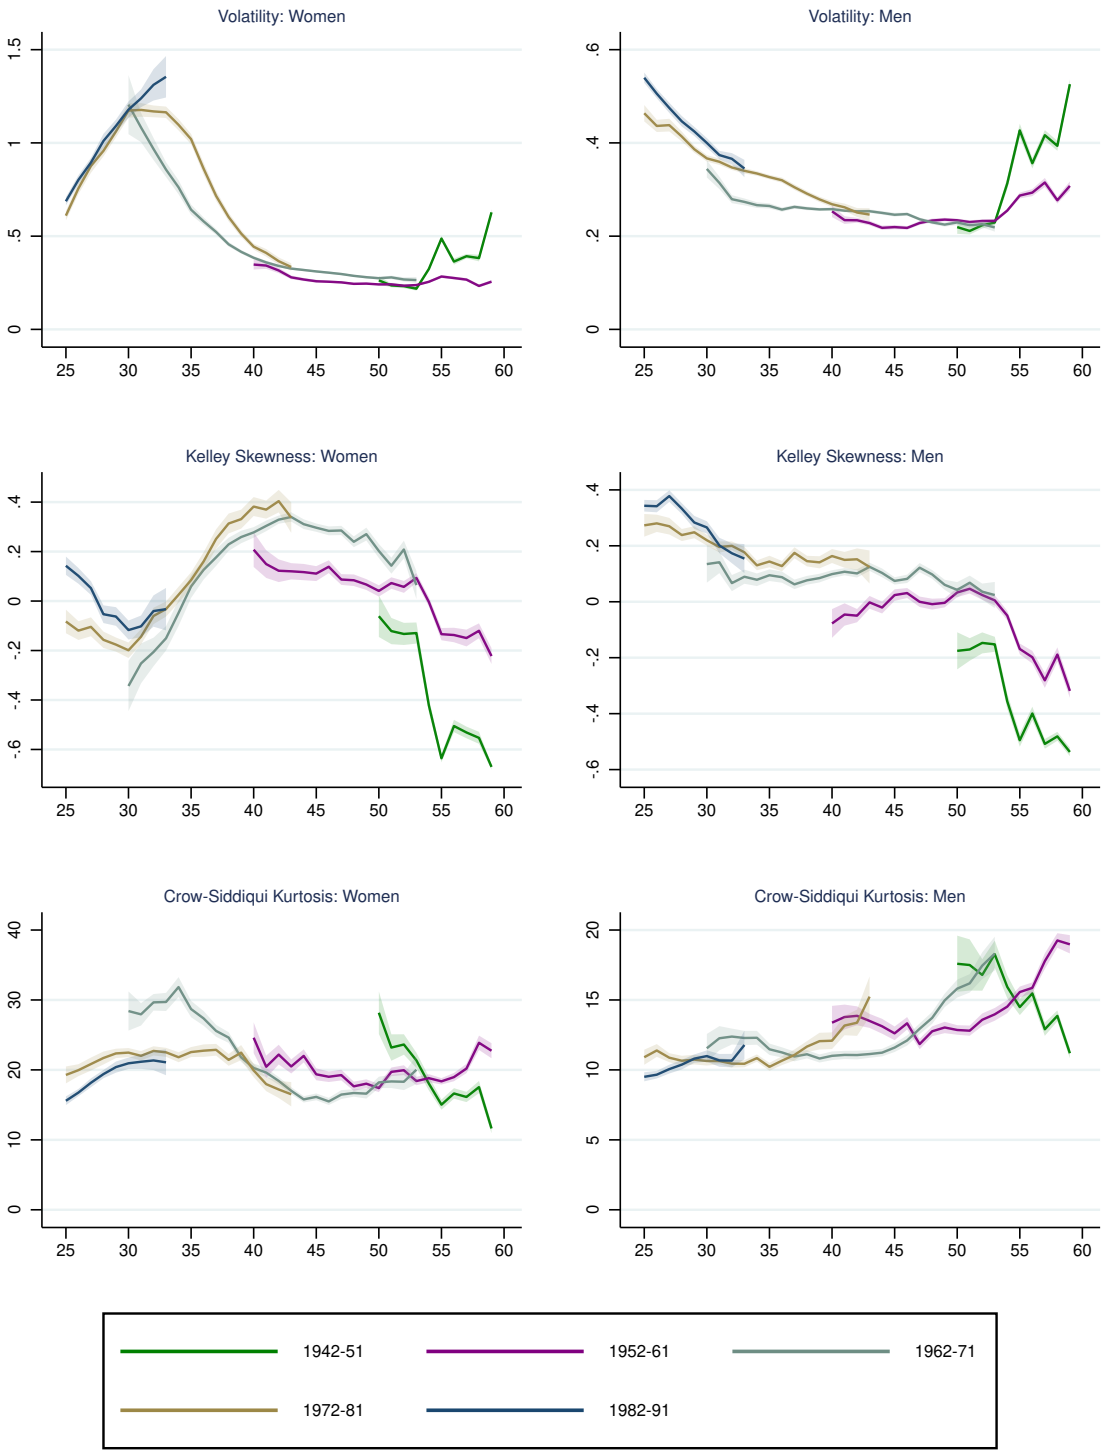

Figure 3: Residualized labor earnings growth risk over the life cycle by cohorts

Note: Prime age individuals in the TPP, years 2001-2016; 95-percent confidence intervals.

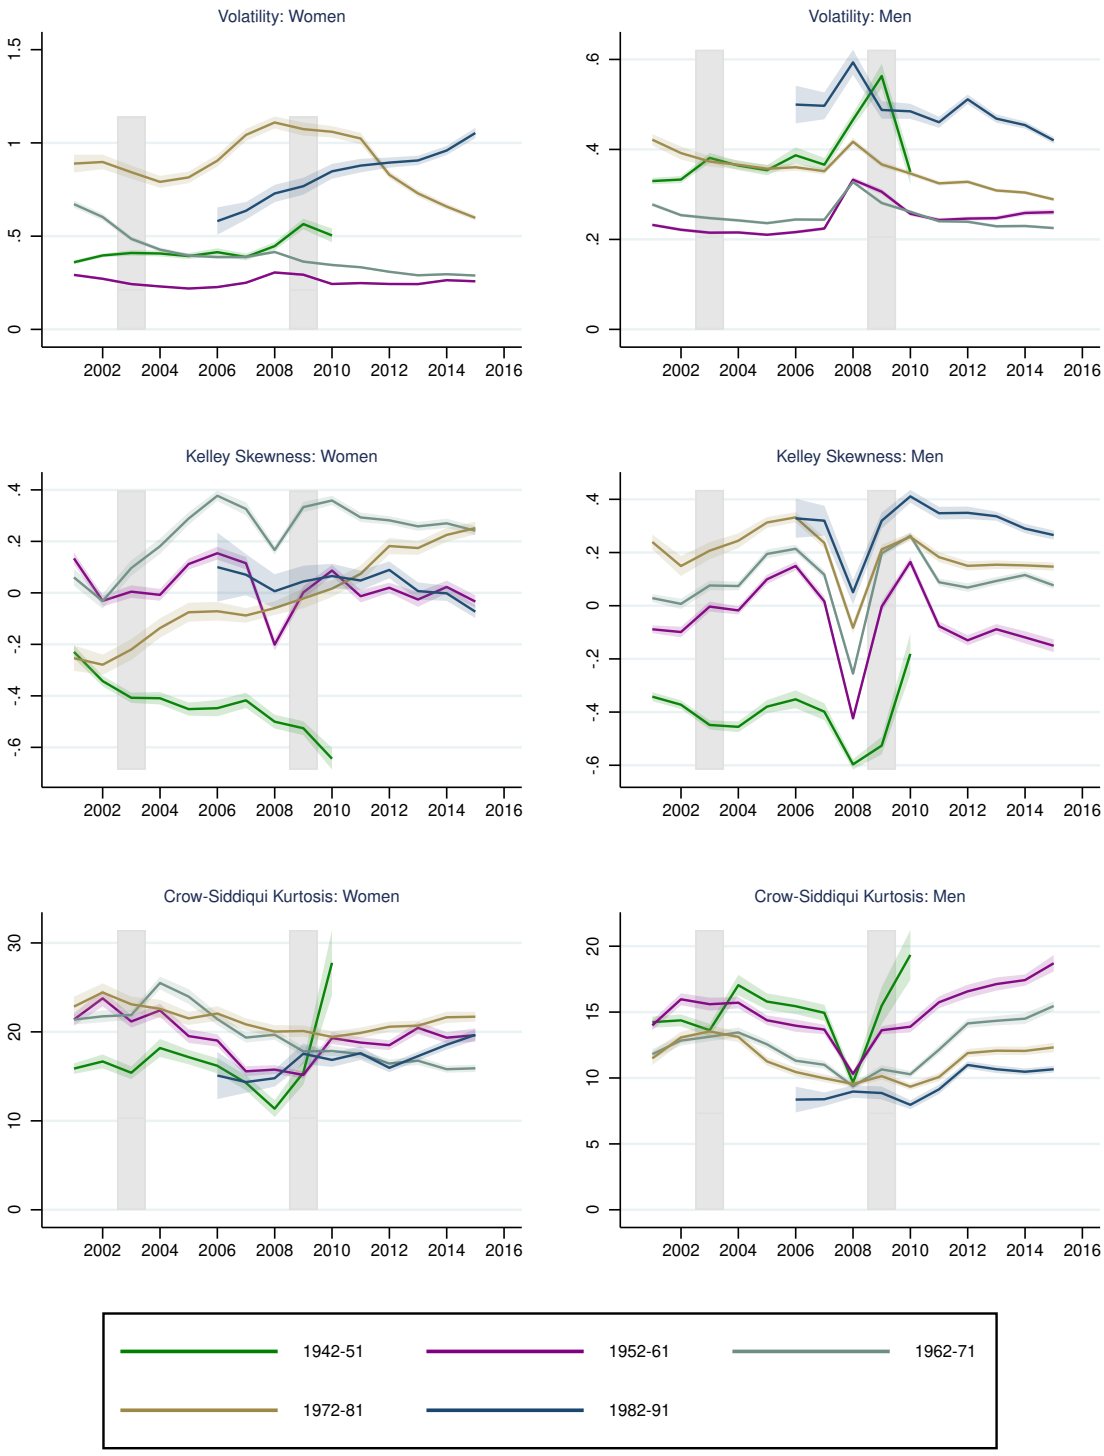

Figure 4: Residualized labor earnings growth risk over time by cohorts

Note: Prime age individuals in the TPP, years 2001-2016; 95-percent confidence intervals.

3.4 Within-family and Government Insurance Against Income Risk

We explored the life-cycle patterns and time trends of changes in individual labor income and found many instances of increased risk in terms of volatility, skewness and kurtosis. However, individuals are often more interested in their household net income rather than the labor income. The former tends to be more stable throughout one's working life and over time thanks to two main insurance channels: The family and the tax- and transfer-system. In this section, we analyse the extent to which the household context and the welfare state can reduce risk and explore trends and channels of these insurance effects.

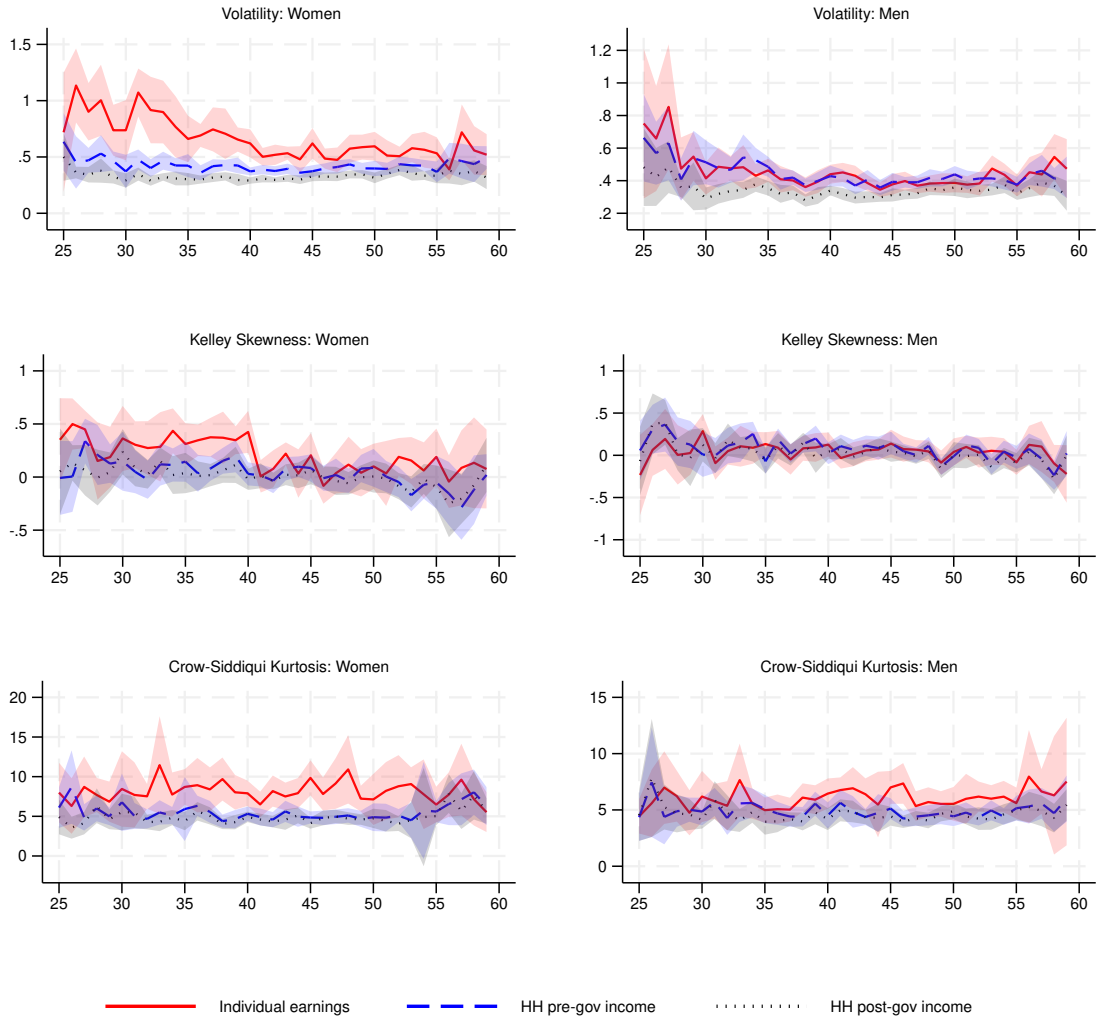

Figure 5: Insurance of one-year earnings changes over the life-cycle

Note: SOEP; Years 2001-2016; bootstrapped 95-percent confidence intervals

This study adds to the theoretical and empirical literature on the insurance of earnings risk. Busch et al. (2022) show empirically that households and the welfare state offer some amount of insurance against skewness of income shocks.

A priori, it appears unlikely that welfare states offer insurance against shocks given that a good approximation of real-world tax-transfer systems is the popular power function (Feldstein, 1969; Heathcote et al., 2017). This type of tax-transfer system scales shocks to log earnings. The power function is

$$x_t = \lambda y_t^{1-\tau}, \tag{2}$$

where x is disposable income, y gross income, λ a parameter that determines the average tax rate and τ a parameter that determines the degree of tax progressivity. Taking logs and first-differencing, we write the change in log disposable income as

$$\Delta lnx_t = (1-\tau)\Delta lny_t. \tag{3}$$

Thus, earnings changes are scaled by $1-\tau$. It follows immediately that $Var(\Delta lnx_t) = (1-\tau)^2 Var(\Delta lny_t)$. In contrast, the third and fourth standardized moments, which are measures of skewness and kurtosis are unit-free, i.e. unaffected by scaling. The same holds for the alternative measures of skewness and kurtosis that we use in this paper.

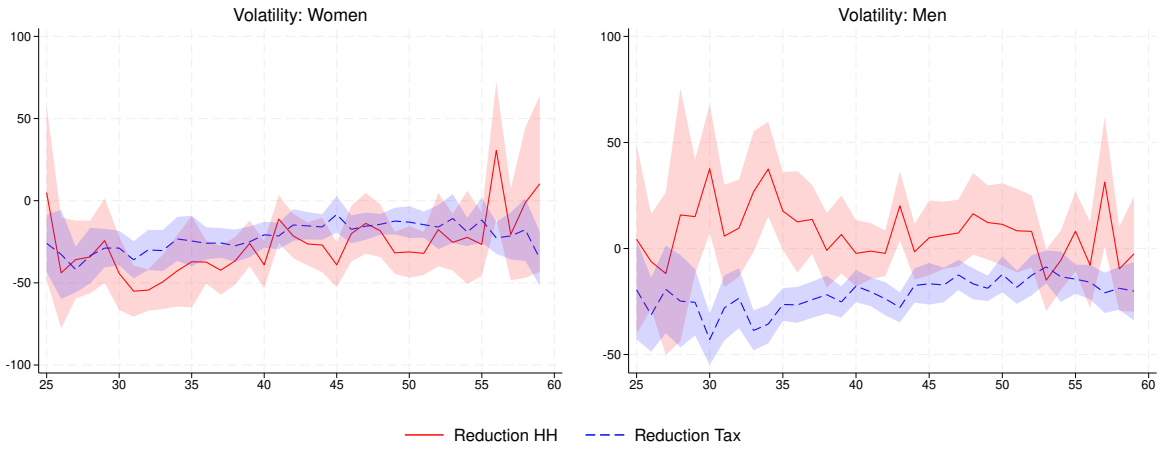

Figure 6: Relative insurance of one-year earnings changes over the life cycle

Note: Percentage reduction in residualized earnings growth volatility. Prime age individuals in the SOEP, years 2001-2016, bootstrapped 95-percent confidence intervals

In this subsection we use the data from the SOEP, which contains detailed information

on the household constellation, all income types and the entirety of taxes and transfers paid and received by the household. For comparison with the TPP, the time horizon is limited to the year 2001-2016. Furthermore, we reduce our sample to individuals who live in couples to show the household insurance effect conditioned on the existence of a second adult with potential income. Figure 5 plots the moments of residualized growth rates of labor income, household income before taxes and transfers and household income after taxes and transfers over the life cycle. In separate panels, volatility, skewness and kurtosis are displayed from the perspective of women and men, and in each panel the individual labor earnings, the household income before and after taxes and transfers are depicted.

For women in couples, the individual earnings growth risk drastically declines with increasing age after a peak at age 31. When including the household (before taxes) however, the earnings risk is reduced to a volatility of below 0.5 at all ages with a rather flat trajectory, meaning that the household insurance channel seems to be the strongest at young ages. The reason is that young women face much higher earnings risk than men. By pooling their earnings with a partner with substantially lower earnings risk, young women reduce their income risk drastically. In all age groups, household income growth risk after taxes and transfers is even lower than household pre-government income risk, indicating the existence of income insurance through the government. However, the 95-percent confidence intervals overlap slightly in some ages.

While we also see decreasing earnings risk for men with increasing age at the beginning of working life, the levels of volatility are notably lower than women's. Moreover, the insurance effect through the household is much smaller and even ambiguous, which hints to the fact that in typical German households, women are the secondary earners with much more volatile labor market biographies. Additional insurance through taxes and transfers is unsurprisingly similar to the results for women with minor differences stemming from the different age constellations within the couples in our sample.

Figure 6 displays the relative change in volatility going from individual labor income to household pre-government income and going from pre-government to post-government income at any given age in percent. Between the ages 26 and 40 the risk reduction through the household ranges mainly between 40% and 60%. At older ages, this effect shrinks until it reaches and even surpasses 0 for the oldest women in the sample. Contrary to this, the insurance effect through taxes and transfers is much more stable over the life-cycle and amounts to an additional 30%. Men, on the other hand, experience a changes in volatility between +35% and -10% depending on the age group when looking at the household

instead of the individual labor income. Income pooling thus increases men's exposure to income risk. The volatility increase is particularly strong in their late 20s to mid 30s when often they have children and their female partner adjusts their labor supply. The insurance via taxes and transfers works in the same way for men as for women in our sample of only couples.

The Kelley skewness as well as the Crow-Siddiqui kurtosis (row two and three of Figure 5) of women's earnings changes is reduced by the household, albeit often to a statistically insignificant amount. For the men in our sample, the skewness remains virtually unchanged and the kurtosis is reduced only slightly and with few exceptions insignificantly by the household. In line with the results from the approximation via power function, the tax- and transfer-system does not affect the third and fourth moment of the earnings changes. While Figure 5 captures the result for short-term effects, Figure A.1 shows the distribution of 5-year changes in income and depicts qualitatively similar results for volatility, skewness, kurtosis and the insurance mechanisms.

Over the years, the volatility of household gross income growth was far more stable than the volatility of women's individual earnings while the two lines for men overlap almost exactly (see Figure 7). The plots of post-government income growth volatility are roughly parallel to those of the pre-government income. In relative terms, the family insurance mechanism accounts for a reduction in volatility of 30% to 40% for female income over the largest stretch (2005 to 2014) while it quite contributes to a constant volatility increase of around 10% for male income, as Figure 8 shows.

The middle row of Figure 7 visualizes that the pro-cyclicality of the skewness is not insured by the household. During the Great Recession, men's household income changes were just as negatively skewed as the labor income and women experienced an much more negatively skewed income in the household due to income pooling. The household does, however, significantly reduce women's excess kurtosis throughout all years. Again, the time trends and business cycle effects do not differ qualitatively between 1-year and 5-year changes in income (see Figure A.2).

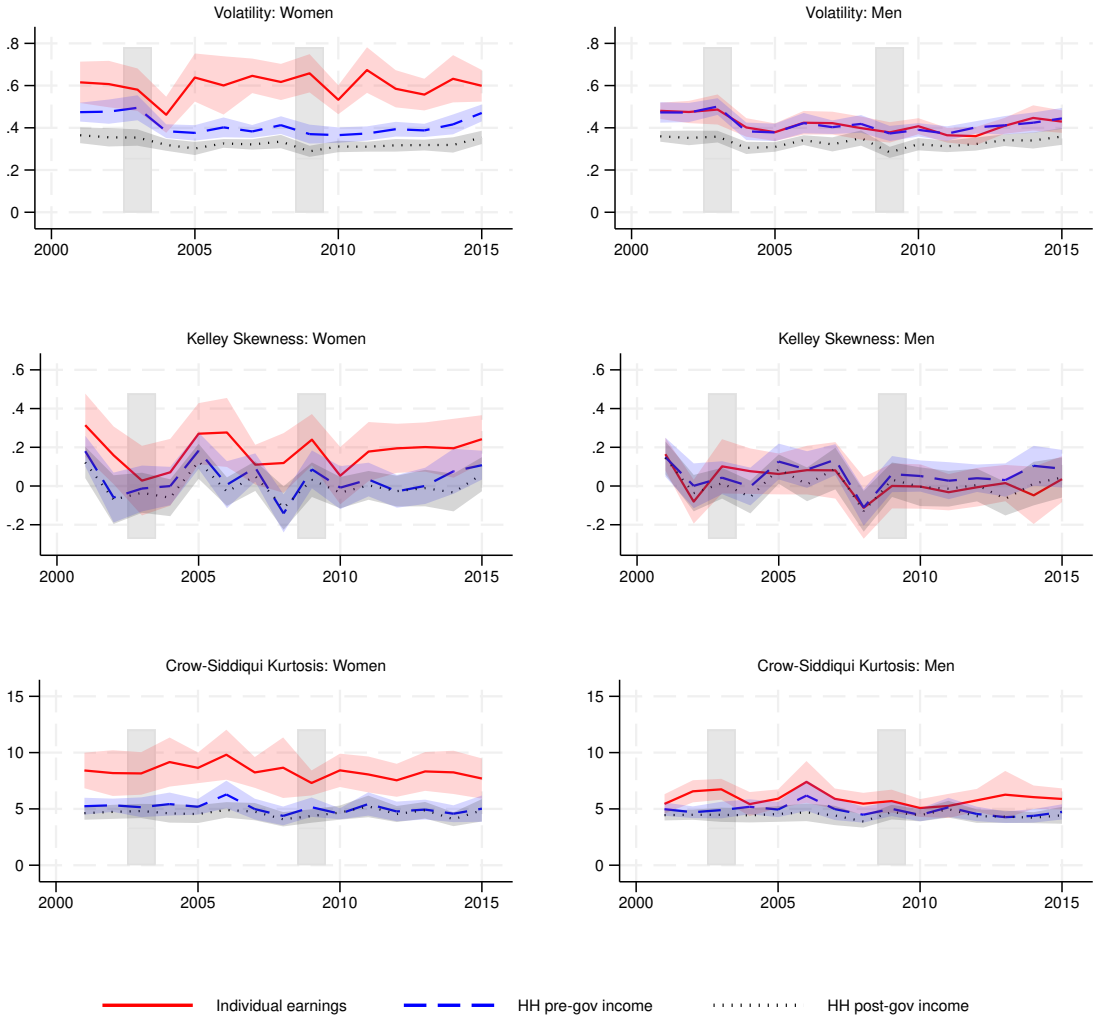

Figure 7: Insurance of one-year earnings changes over time

Note: SOEP; Ages 25-60; bootstrapped 95-percent confidence intervals

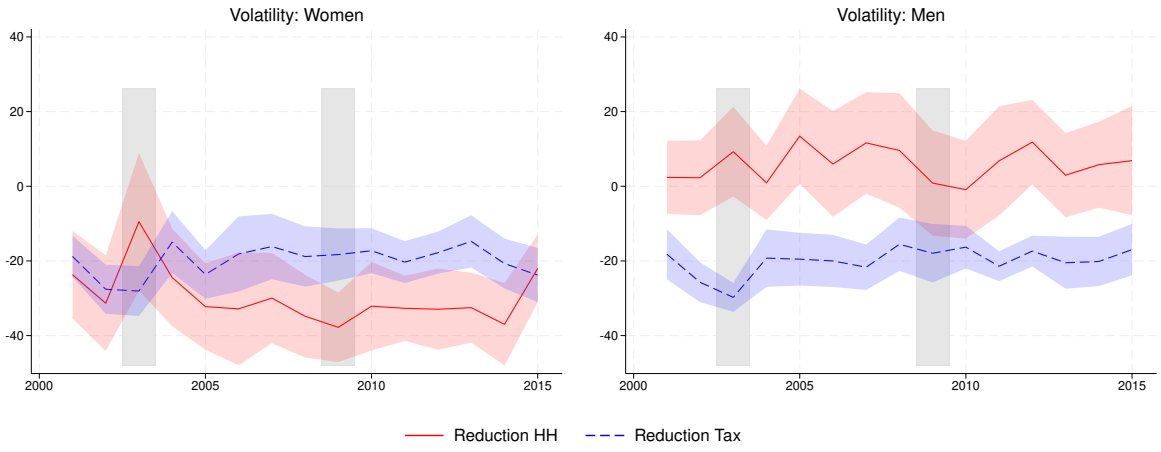

Figure 8: Relative insurance of one-year earnings changes over time

Note: Percentage reduction in residualized earnings growth volatility. Prime age individuals in the SOEP, years 2001-2016, bootstrapped 95-percent confidence intervals

3.5 The Role of the Household

In this subsection, we take a deeper dive into the workings of the family insurance against earnings risk. There are two possible mechanisms of risk reduction through a secondary earner in the household. One is simply income pooling, meaning that adding another (uncertain) income to the total household income can help offset large income shocks to the primary earner. The other channel is the added-worker effect. The secondary earner could adjust their income by changing their hours worked per month in response to an income shock of the primary earner with the goal of maintaining a smooth earnings trajectory over time. Using our full sample of couples in the SOEP data, Figure 9 rejects the existence of such an added-worker effect in our data. Following De Nardi et al. (2021), it plots the deciles in the (positive and negative) shocks to residualized male earnings against the respective change in hours of the woman in the household. The flat horizontal line at zero can be interpreted as the absence of any measurable reactions of women to the man in their partner's change in earnings no matter the size and the direction of the shock. In the presence of a textbook added-worker effects, the graph would be downward sloping. Hence, this channel does not contribute to the reduction in earnings growth risk and the insurance through the household can merely be attributed to income pooling.

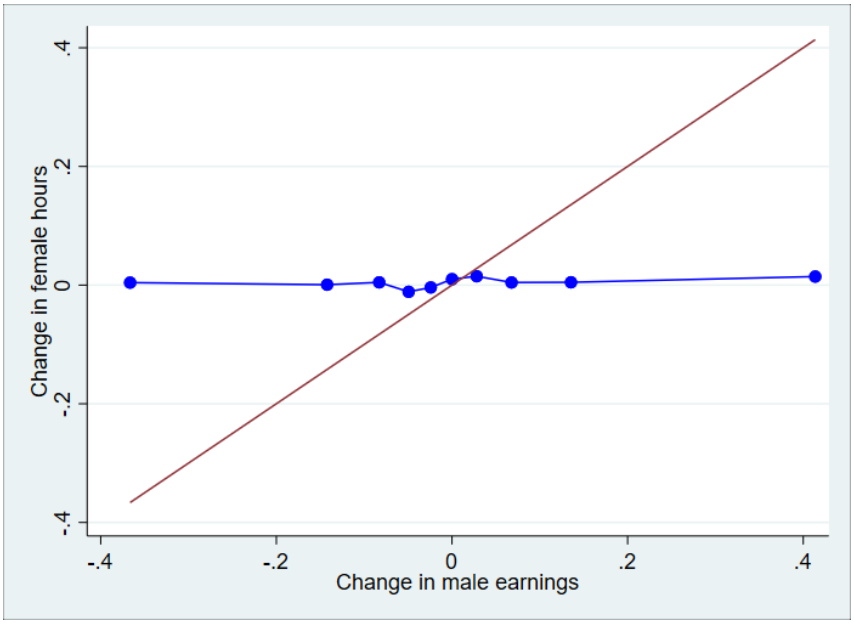

Figure 9: Change in male earnings and female labor supply

Note: Earnings changes t to $t+1$; hours changes t to $t+2$. Each dot represents a decile of changes in male earnings. Prime age males in the SOEP, years 1991-2018

From the perspective of men, for income pooling with their female partner to play a significant role in insuring against income risk, the female's share in household income

must be sizable. Figure 10 shows that, for couples, participation rates increase slightly in age and—at least in the sample of couples—there are no obvious cohort effects. While on aggregate female labor force participation rates have increased over the years, this increase is not observable when restricting the focus on women in couples. Turning to women's share of household income, reported in Figure 11, we observe a clear U-shaped pattern over the life-cycle, while, again, there are no cohort effects.

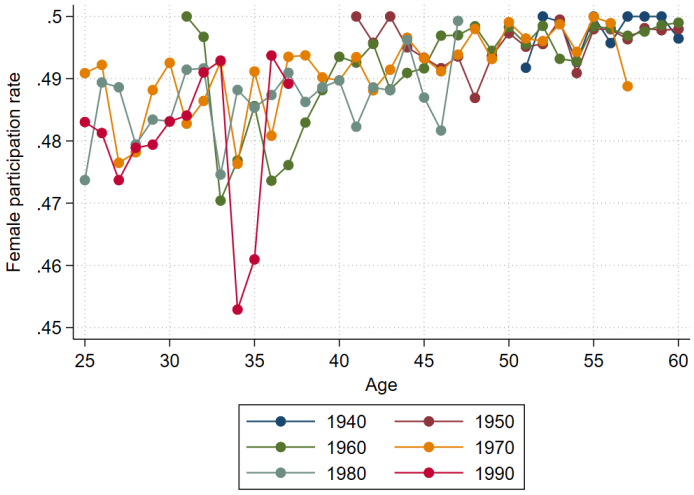

Figure 10: Female labor force participation rates

Note: Prime age females in couples in the SOEP, years 1991-2018

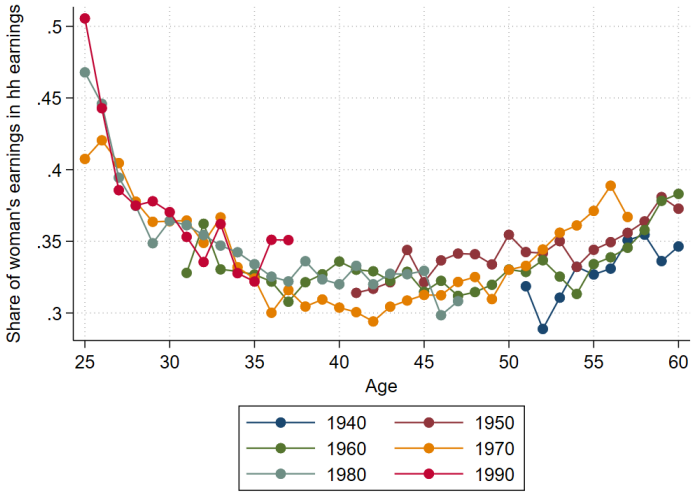

Figure 11: Share of women's earnings in total household earnings

Note: Prime age females in couples in the SOEP, years 1991-2018

The increase in women's share of the household's total earnings at the end of working life is accompanied by a decrease in insurance through the household from the woman's perspective (Figure 6). The reason is that—for the values of individual earnings risk in

our sample, where women face higher earnings risk—women benefit more from income pooling as their share in household income decreases.[4]

3.6 Male Earnings Changes and Insurance

Finally, we analyze the insurance effects of household and government by size and direction of men's unexpected earnings changes as well as along the distribution of net household incomes equivalized using the modified OECD scale. Figure 12 depicts the individual income shocks divided into deciles by size of the earnings change, as well as the resulting change in household income before and after taxes, encompassing the full sample of men aged 25 to 60. The closer the dots—each representing a decile of changes in male earnings—are to the dotted line of no insurance, the smaller the measured insurance effect. In the case of perfect insurance from any earnings risk, the dots would lie on a flat line at *change in household income* $= 0$. Both family and government play an important role in earnings risk mitigation. A 50%-increase in male individual income for example is associated with a 38%-increase in pre-government household income and a mere 25%-increase in post-government household income. Together, these insurance channels cut the earnings risk for the largest decile in earnings changes in half.

We repeat this analysis for three different deciles of the income distribution. The other panels of Figure 12 show the results for the first, fifth and ninth decile, respectively. Naturally, starting out from the lowest income group, we see predominantly large, positive changes in male earnings in the upper right panel. However, the change in post-government household income even after very large income shocks is quite small due to the progressivity of the tax system and high transfer withdrawal rates for this low income earners. For instance, an increase in hours of work by 50 percent leads only to a small increase in disposable income. This example highlights an important drawback of insurance through the state: the efficiency cost of taxation due to adverse labor supply incentives.

The fifth decile, shown in the lower left panel, displays a similar pattern as the full sample: Mostly moderate shocks, which are moderately insured by both the household and the government. In the ninth income decile (lower right panel), insurance against positive earnings shocks is relatively small. The reasons for this are twofold. First, a positive shock to male earnings in this decile usually imply that the already large share

[4]In the benchmark case, where women and men face the same earnings risk and their earnings are uncorrelated, the gain from pooling is maximized when the share of each partner's earnings in household earnings is 0.5. In contrast, when the partners face risks of different magnitude, the gain from pooling is maximized when the share of the more risky earnings is less than half.

of the man's earnings in total household earnings increases even further, reducing the importance of income pooling. Second, the German income tax schedule is esentially flat at the top. In contrast, larger negative shocks are still well-insured and effectively cut in half.

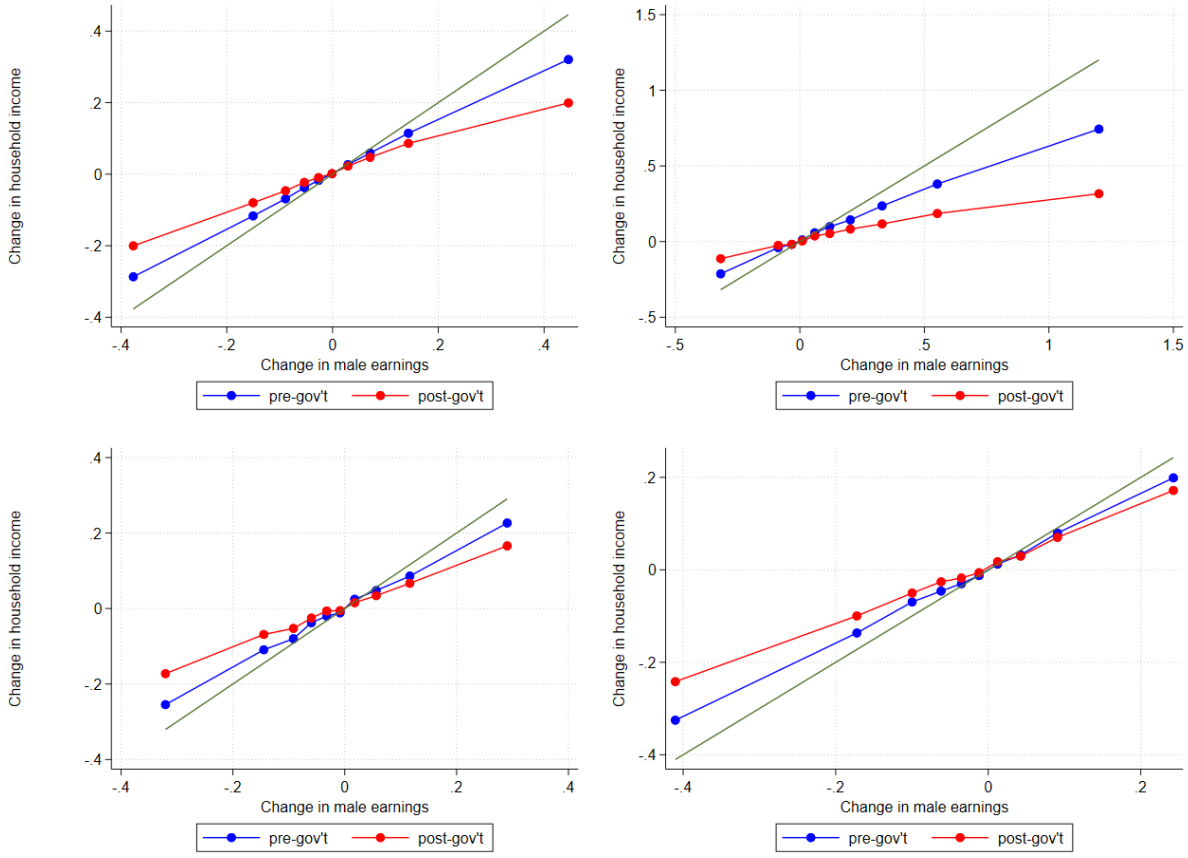

Figure 12: Household before- vs after-tax income

Note: Each dot represents a decile of changes in male earnings. Upper left: Full sample; Upper right: 1st decile; Lower left: 5th decile; Lower right: 9th decile. Prime age males in the SOEP, years 1991-2018

4 Conclusions

We have analyzed the earnings growth risk and mechanisms that mitigate this risk in Germany, adding to a growing literature on the sources of earnings risk and insurance against it. We find that the earnings growth risk is the highest at the beginning of the working life, followed by a decrease until age 40 and from there onwards the risk remains constant. Generally, women face higher earnings risk than men and younger cohorts face higher risk than previous birth cohorts.

Regarding the patterns of income risk over the life cycle, for men, volatility over

the lifecycle is U-shaped, while skewness decreases and kurtosis increases with age. In contrast, for women volatility and kurtosis are clearly highest during the childbearing age, while skewness reaches its highest point later in life. During most years, skewness is positive for men and women, meaning that large positive income changes occur more often than negative ones. During the Great Recession both men and women experienced a large drop in skewness with a larger drop for men. A cohort analysis reveals that, at a given age, women and men born later face higher volatility, but also higher skewness. Strikingly, while men in all cohorts faced a large drop in skewness during the Great Recession, younger women appear not to have been affected by it.

Both the family and the tax- and transfer-system play a large role in the insurance against earnings risk. The household context is especially important for reducing young women's income risk. Taxes and transfers further reduce volatility in household earnings growth by about 30%. The insurance through the household can be attributed almost exclusively to the effect of income pooling since we find no evidence for an added-worker effect. The household greatly reduces earnings volatility for women and reduces kurtosis for both men and women.

The relative size of the insurance effect does not directly depend on the size of the income shock. Nonetheless, splitting the sample into deciles of the income distribution, we observe heterogeneous degrees of insurance, which might reflect undesirable properties of the tax-transfer system. In particular, large, positive shocks at the lowest income group are absorbed by the tax- and transfer-system to a larger extent than the positive shocks in the ninth income decile.

Finally, the tax-transfer system reduces volatility, but does not offer any insurance against higher-order income risk.

References

Borjas, G. J. (1980). The relationship between wages and weekly hours of work: The role of division bias. *The Journal of Human Resources*, 15(3):409–423.

Busch, C., Domeij, D., Guvenen, F., and Madera, R. (2022). Skewed idiosyncratic income risk over the business cycle: Sources and insurance. *American Economic Journal: Macroeconomics*, 14(2):207–42.

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

De Nardi, M., Fella, G., Knoef, M., Paz-Pardo, G., and Van Ooijen, R. (2021). Family and government insurance: Wage, earnings, and income risks in the netherlands and the u.s. *Journal of Public Economics*, 193:104327.

De Nardi, M., Fella, G., and Paz-Pardo, G. (2020). Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. *Journal of the European Economic Association*, 18(2):890–926.

Drechsel-Grau, M., Peichl, A., Schmid, K. D., Schmieder, J. F., Walz, H., and Wolter, S. (2022). Inequality and income dynamics in germany. *Quantitative Economics*, 13(4):1593–1635.

Feldstein, M. S. (1969). The effects of taxation on risk taking. *Journal of Political Economy*, 77(5):755–764.

Ghosh, A. and Theloudis, A. (2023). Consumption partial insurance in the presence of tail income risk. *arXiv preprint arXiv:2306.13208*.

Goebel, J., Grabka, M., Liebig, S., Kroh, M., Richter, D., Schröder, C., and Schupp, J. (2018). The German Socio-Economic Panel (SOEP). *Jahrbücher für Nationalökonomie und Statistik*.

Grabka, M. (2020). SOEP-Core v36 – Codebook for the PEQUIV File 1984-2019: CNEF Variables with Extended Income Information for the SOEP. SOEP Survey Papers, Series D 922, DIW.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk? *Econometrica*, 89:2303–39.

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

Guvenen, F., Pistaferri, L., and Violante, G. L. (2022). Global trends in income inequality and income dynamics: New insights from grid. *Quantitative Economics*, 13(4):1321–1360.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics*, 132:1693–1754.

Hoffmann, E. B., Malacrino, D., and Pistaferri, L. (2022). Earnings dynamics and labor market reforms: The italian case. *Quantitative Economics*, 13(4):1637–1667.

Kelley, T. L. (1947). *Fundamentals of statistics*. Number 2. Harvard University Press.

Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimüller, J. (2019). Child penalties across countries: Evidence and explanations. In *AEA Papers and Proceedings*, volume 109, pages 122–126. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Pessoa, A. S. (2021). Earnings Dynamics in Germany. CESifo Working Paper Series 9117, CESifo.

Schröder, C., König, J., Fedorets, A., Goebel, J., Grabka, M. M., Lüthen, H., Metzing, M., Schikora, F., and Liebig, S. (2020). The economic research potentials of the German Socio-Economic Panel study. *German Economic Review*, 21(3):335–371.

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

Appendix

A Additional Figures

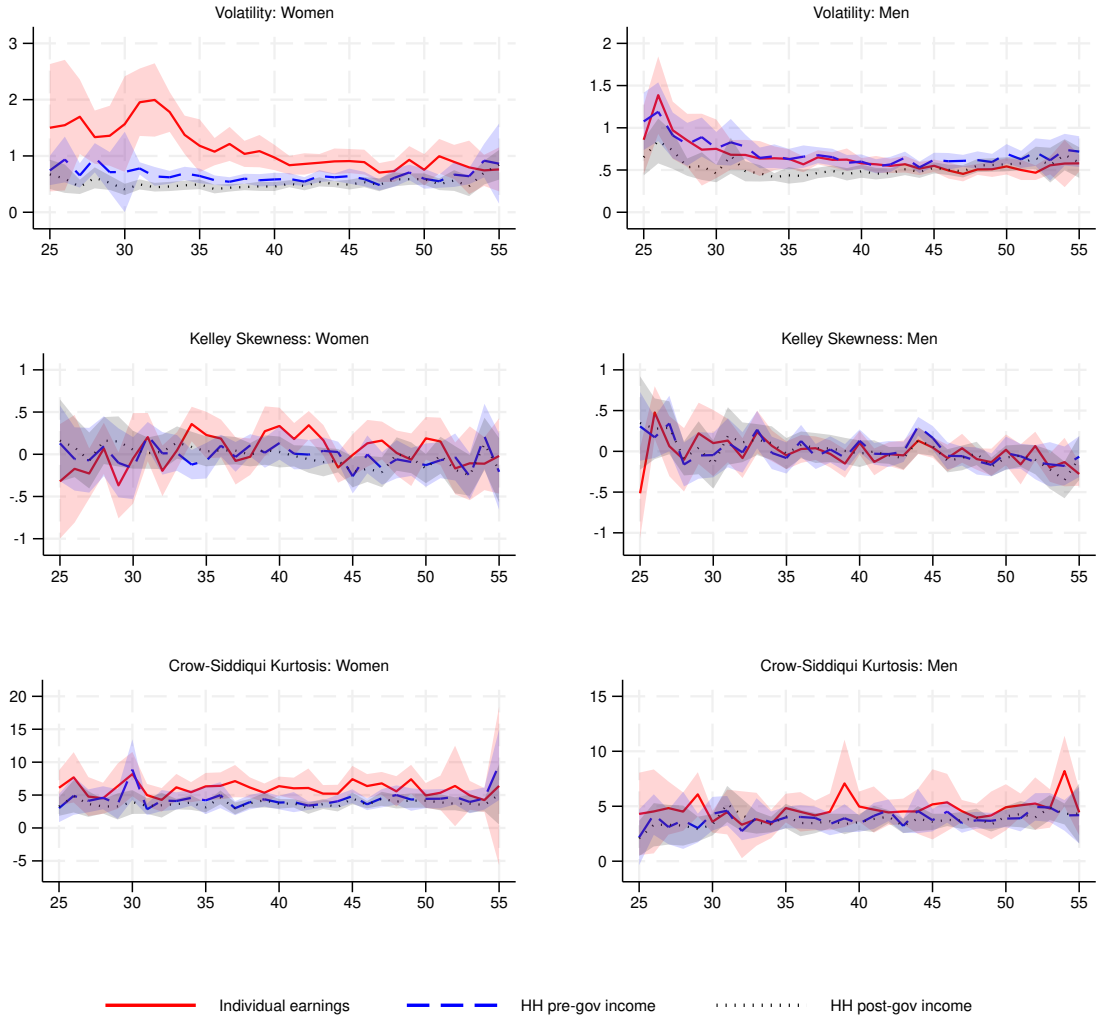

Figure A.1: Insurance of five-year earnings changes over the life-cycle

Note: SOEP; Years 2001-2016; bootstrapped 95-percent confidence intervals

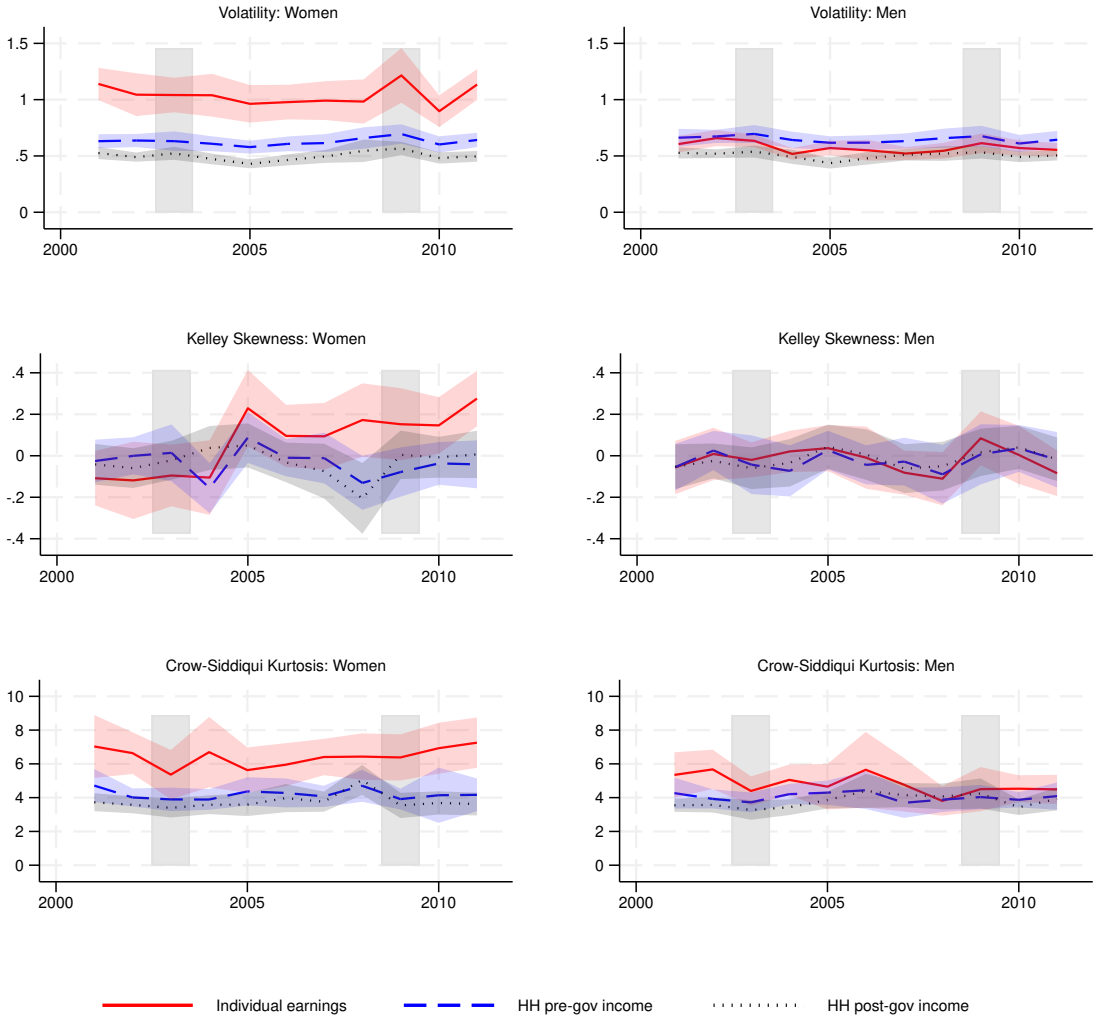

Figure A.2: Insurance of five-year earnings changes over time

Note: SOEP; Ages 25-60; bootstrapped 95-percent confidence intervals

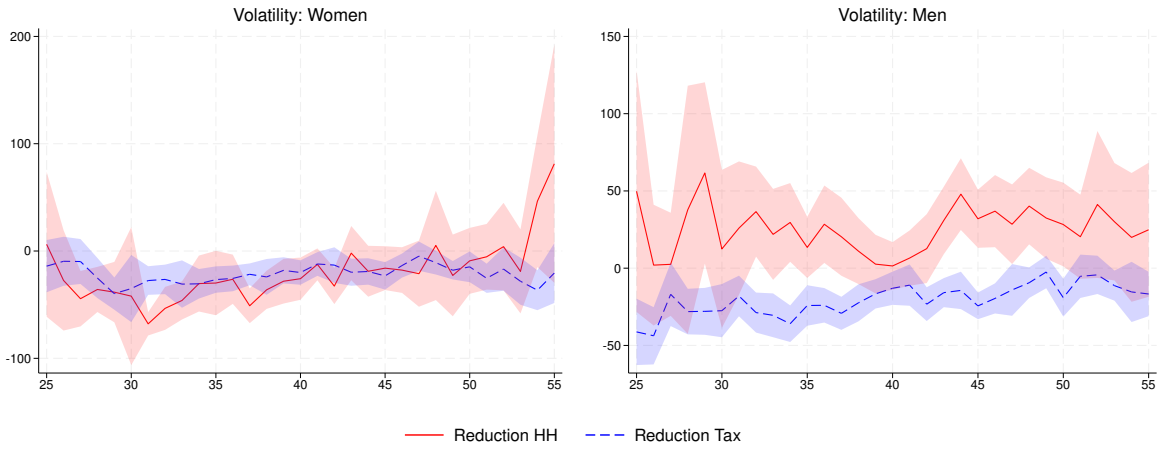

Figure A.3: Relative insurance of five-year earnings changes over the life cycle

Note: Percentage reduction in residualized earnings growth volatility. Prime age individuals in the SOEP, years 2001-2016, bootstrapped 95-percent confidence intervals

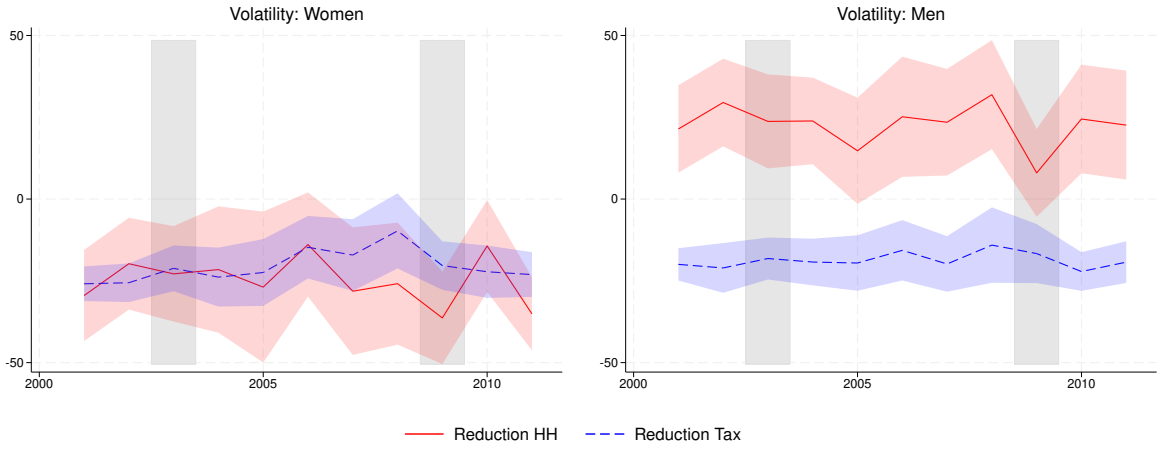

Figure A.4: Relative insurance of five-year earnings changes over time

Note: Percentage reduction in residualized earnings growth volatility. Prime age individuals in the SOEP, years 2001-2016, bootstrapped 95-percent confidence intervals