The Life-Cycle and the Business-Cycle of Wage Risk - Cross-Country Comparisons*

Christian Bayer†
Falko Juessen‡
Universität Bonn and IGIER
Technische Universität Dortmund and IZA

This version: January 4, 2011

Abstract

This paper provides a cross-country comparison of life-cycle and business-cycle fluctuations in the dispersion of household-level wage innovations. We draw our inference from household panel data sets for the US, the UK, and Germany. First, we find that household characteristics explain about 25% of the dispersion in wages within an age group in all three countries. Second, the cross-sectional variance of wages is almost linearly increasing in household age in all three countries, but with increments being smaller in the European data. Third, we find that wage risk is procyclical in Germany while it is countercyclical in the US and acyclical in the UK, pointing towards labor market institutions being pivotal in determining the cyclical properties of labor market risk.

KEYWORDS: Life-cycle risk, uncertainty fluctuations, business cycle, heterogeneity, wages
JEL-codes: E24, D31, D91, J31

---

*We would like to thank Per Krussell, Dirk Krüger, Andreas Schabert and seminar and conference participants at the University of Dortmund, NASM 2008, LAMES 2008, EEA 2009, ZEI Summer-School 2009, ESWC 2010, for helpful comments and suggestions.
Financial support from the German Science Foundation (DFG), SFB 823, is gratefully acknowledged.

†Universität Bonn, Kaiserplatz 7-9, 53113 Bonn, Germany; Tel. +49-228-734073; email: christian.bayer@uni-bonn.de

‡Technische Universität Dortmund, Department of Economics, 44221 Dortmund, Germany; phone: +49-231-755-3291; fax: +49-231-755-3069; email: falko.juessen@tu-dortmund.de

1 Introduction

A growing macroeconomic literature has documented that economic uncertainty negatively comoves with the business cycle, i.e. uncertainty increases in recessions and decreases during booms. For example, Bloom (2009) and Bloom et al. (2009) show that, in the US, stock market volatility spikes at times of economic downturn. Gilchrist et al. (2009) find a similar result looking at the link of credit spreads, uncertainty, and investment. Alexopoulos and Cohen (2009) find a close negative correlation between the frequency of using the word 'uncertainty' in the 'New York Times' and economic activity in the US. Bachmann and Bayer (2009a) find for a large panel of German firms that the dispersion of firm-level productivity growth rates is significantly countercyclical. Finally, Storesletten et al. (2004) document from US PSID data a negative correlation of the cycle with the dispersion of innovations to household income. In summary, there is growing evidence that microeconomic uncertainty, i.e. the cross-sectional variance of idiosyncratic shocks, is strongly linked to business cycle movements.

This paper adds to this literature in two aspects: first, by analyzing the cyclicality of the cross-sectional dispersion of wage-innovations instead of innovations in earnings. This is important as wages are more closely linked to productivity and are less subject to endogenous household decisions than earnings are. Second, we contribute by providing a cross-country perspective on the cyclicality of the dispersion of wage innovations, comparing the US, the UK, and Germany. This way, we provide results for a range of labor market setups, in which the US labor market is the most liberal and the German labor market the most regulated one.

We find that the estimated wage processes have a similar structure overall. However, we find the ranking of labor market regulation reflected empirically in the size and cyclicality of wage risk. In the US, wage risk is highest and strongly countercyclical. In the UK, wage risk is mild and acyclical. In Germany, wage risk is smallest and procyclical.

Why are these cross-country differences important? One can interpret our findings in two ways. First, one can think of wages as reflecting marginal productivity as in a neoclassical model. Given this interpretation, our findings suggest that there are significant structural differences in the way productivity risk behaves between the three large economies studied. This might be important for thinking about the driving forces of aggregate fluctuations in various countries, for example as in Bloom et al. (2009), Bachmann and Bayer (2009a), or Gilchrist et al. (2009). Second, one can think of wages as rather an outcome of complex economic processes, such as bargaining with unions, search-and-matching, or sticky wage-setting. In this second case, the cyclicality of wage risk may help discriminating between (quantitatively) different labor market models, similar to the point made by Bachmann and Bayer (2009b) for the cyclicality of the investment rate dispersion, which is found to be a strong overidentifying restriction in heterogenous firm RBC models. This point of view is linked to the general point made by Gomme et al. (2004) that taking into account subgroup heterogeneity may well enhance our understanding of business-cycle fluctuations at the aggregate level. Examples of this approach are Heathcote et al. (2009) or Ljungqvist and Sargent (2008), who emphasize the role of income risk in a general equilibrium model with heterogeneous households. Finally, fluctuations in wage risk are important factors in determining the costs of business cycles, see for example Storesletten et al. (2001) or Krebs (2003, 2007).

Our paper is also related to the literature that studies the life-cycle of earnings inequality in general and without an explicit business cycle perspective.[1] In particular, our paper can be understood as an extension of a recent cross-country project on the study of economic inequality, see Krueger et al. (2010) for an overview article. In contrast to our paper, the primary focus of this cross-country project is to document time-trends in inequality in wages, labor earnings, income, consumption, and wealth, while we focus on the cyclicality of wage risk. The papers of this project also provide some information on the link between inequality and the business cycle.[2] However, these studies only measure the effects of recessionary episodes that are observed during the sample period of the micro data. Thus, inference is based on a very limited number of business cycles. Our identification strategy for the cyclicality of wage risk is substantially different, exploiting the interaction between age, time, and cross-sectional variance and using a much longer time span for identification by conditioning on the macroeconomic history of a cohort like in Storesletten et al. (2004).

The remainder of this paper is organized as follows. Section 2 presents our statistical model setup. Section 3 describes the data we use. Section 4 presents our GMM estimation results followed by a set of robustness checks. The last section concludes the paper. A separate appendix providing more details on the estimation technique and

[1] For the three countries we focus on - UK, US, and Germany - examples of this literature are: (for the UK) Dickens (2000), Gossling et al. (2000), Blundell et al. (2008), Blundell and Etheridge (2010), (for the US) Guvenen (2007, 2009), Krueger and Perri (2006), Primiceri and van Rens (2007), Heathcote et al. (2010), (for Germany) Biewen (2005), Fuchs-Schündeln et al. (2010).

[2] The overall conclusion by Krueger et al. (2010) is that the dynamics of wage inequality during downturns are not uniform across countries nor across recession episodes. They conclude that the impact of recessions on wage inequality is likely to depend on the specific causes of the recession and on the structure of the labor market.

sample selection is available online.[3]

2 Wage process and estimator

To describe the life-cycle development of wage dispersion and the business-cycle dependence of wage risk, we specify a statistical time-series model of a household's average wage rate. We assume that the average wage rate of a household i that has h years of labor market experience at time t can be described by

$$\begin{aligned} w_{it}^h &= \beta x_{it}^h + \omega_{it}^h; \\ \omega_{it}^h &= \alpha_i + z_{it}^h + \varepsilon_{it}; \\ z_{it}^h &= \rho z_{it-1}^{h-1} + \sqrt{\phi_t}\nu_{it}, \end{aligned} \tag{1}$$

where x_{it}^h is a vector of observable household characteristics and ω_{it}^h is the wage residual. The wage residual is composed of a household fixed effect α_i, a transitory i.i.d. (or measurement error) term ε_{it}, and a persistent component z_{it}^h. This component z_{it}^h follows an AR-1 process with autocorrelation ρ and innovations $\phi_t \nu_{it}$, with variance $\sigma_\nu^2 = 1$. In particular, we assume that the scaling parameter ϕ_t depends on the business cycle, such that $\phi_t = \phi(Y_t)$, where Y_t is some measure of aggregate activity, e.g. deviations from trend GDP growth. This specification is close to the one used in Storesletten et al. (2004), where ϕ_t can take two values, ϕ_L and ϕ_H, depending on whether GDP growth is below or above average, but generalizes their framework by allowing for more than two regimes. We parameterize ϕ as

$$\phi(Y_t) = \frac{\underline{\phi} + \bar{\phi}}{2} + \left(\bar{\phi} - \underline{\phi}\right)\left(\frac{1}{1+\exp(-Y_t)} - \frac{1}{2}\right),$$

which allows the effect of business cycles to vary according to the strength and not just to the sign of the cycle. We impose $\bar{\phi}, \underline{\phi} > 0$ in order to ensure that $\phi_t > 0$.

While the variance of persistent shocks ν_{it} is potentially time varying, we restrict the variance of fixed effects as well as the variance of measurement error, σ_α^2 and σ_ε^2, to be constant over time. We checked the importance of this restriction and estimated an alternative specification that removes cohort-varying variances of fixed individual effects α_i, $\sigma_\alpha^2(t-h)$, and time-varying variances of transitory shocks ε_{it}, $\sigma_\varepsilon^2(t)$. The estimation results qualitatively coincide with our baseline results obtained assuming constant variances.[4]

[3] This appendix is provided as referee appendix.
[4] (SEE REFEREE APPENDIX)

We choose Y_t to be standardized GDP deviations from trend, such that $\frac{\underline{\phi}+\bar{\phi}}{2}$ is the variance at trend growth and the sign of $\left(\bar{\phi}-\underline{\phi}\right)$ reflects whether the dispersion of wage innovations is pro- or countercyclical. We estimate the parameters of the model using a generalized method of moments estimator, where – following Storesletten et al. (2004) – we derive moment conditions by conditioning on the business cycle history of a cohort $c = t-h$ at time t. The assumed wage process allows us to exploit business cycle history beyond the actual time-span of our panel data on wages. The variance structure implies moment conditions that we can exploit for estimation.

For a household with labor market experience h the variance of ω_{it}^h is given by

$$\mu_{1,t,h} = var\left(\omega_{it}^h | t, h\right) = \sigma_\alpha^2 + \sigma_\varepsilon^2 + \sum_{s=0}^{h-1} \rho^{2s} \phi\left(Y_{t-s}\right). \tag{2}$$

Due to the persistency parameter ρ the wage dispersion in an age-year cell $\mu_{1,t,h}$ memorizes past business cycle episodes. This can be thought of as the annual rings of a tree capturing information on historic climatic conditions. In addition, the persistency implies that $\mu_{1,t,h}$ displays a clear life-cycle pattern, as has been discussed in detail for example in Deaton and Paxson (1994). The profile of $\mu_{1,t,h}$ in h is the closer to linear the closer ρ is to 1, i.e. the higher the persistency of shocks to wages.

A second set of moment conditions is given by the autocovariances of wage residuals:

$$\mu_{2,t,h} = cov\left(\omega_{it}^h, \omega_{it+1}^{h+1} | t, h\right) = \sigma_\alpha^2 + \rho \sum_{s=0}^{h-1} \rho^{2s} \phi\left(Y_{t-s}\right). \tag{3}$$

We use this second set of moment conditions to discriminate between the variance of fixed effects σ_α^2 and the variance of transitory effects or measurement error σ_ε^2. Moreover, this second set of moments helps to identify the autocorrelation coefficient ρ more directly. Yet, exploiting these moment conditions requires observing a household in two consecutive years, such that it requires the use of panel data.

We estimate the parameter vector $\theta = \left(\rho, \sigma_\alpha^2, \sigma_\varepsilon^2, \bar{\phi}, \underline{\phi}\right)$ by a generalized method of moments, minimizing the distance between $\left(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2\right), \boldsymbol{\mu}_i := \left(\mu_{i,1,1}, \ldots, \mu_{i,T,H}\right)$, and their empirical counterparts $(\mathbf{m}_1, \mathbf{m}_2)$:

$$\hat{\theta} = \arg\min \left[(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) - (\mathbf{m}_1, \mathbf{m}_2)\right] \mathbf{W} \left[(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) - (\mathbf{m}_1, \mathbf{m}_2)\right]', \tag{4}$$

where $\mathbf{W}$ is a positive definite weighting matrix that captures the correlation structure across moments.[5]

An advantage of our identification strategy is that the direct selection effect of some households becoming unemployed during recessions is negligible in shaping the business cycle dependence of wage risk. This is because we use evidence on wage dispersions after the recession is over and most households have become employed again. We can do so because the wage process memorizes innovations to wages, even if these innovations may be latent during periods of temporary unemployment. Note that a downside of our empirical strategy exploiting the time series behavior of the wage dispersion of a given cohort is that we cannot say much about what creates wage risk and its cyclicality. However, we conduct a large set of robustness tests to exclude that our results are driven by particular characteristics of workers, sample selection, or ways to measure the cycle.[6]

3 Data

3.1 Sample selection

Our data consists of three large household panels as well as aggregate real GDP data for all three countries. We draw our inference from the PSID (for the US), the BHPS (for the UK), and the GSOEP (for Germany). All three data sets are constructed in a similar manner and include information on pre-tax labor income, hours worked, and a set of household characteristics, such as education, household size etc. The three data sets cover different periods of time: the PSID data we use provides annual information for the years 1968-1997, the BHPS for 1991-2007, and the GSOEP for 1984-2006. Despite covering different time periods, we can use these data sets for a cross-country comparison because our identification strategy relies on the fact that, in each year, all three data sets cover a rich cross-section of cohorts that have worked through different working histories.

To compute a household's average wage rate, we use information on pre-tax labor income and hours worked. The labor income of a household is defined as the annual income from employment or self employment of household head and spouse. Analogously, we define the total market hours of work a household supplies. We abstract from hours worked in home production. By pooling information on household head and spouse we take the view that household composition is a fundamental risk of the household. We select households as the unit of interest in our analysis, because labor market decisions of each individual household member are subject to the risk-sharing agreements within

[5] Details on the weighting matrix $\mathbf{W}$ can be found in the separate appendix. (SEE REFEREE APPENDIX)

[6] Results and further details are available in a separate appendix. (SEE REFEREE APPENDIX)

the household and will thus differ significantly according to household composition. In our baseline estimations, we do not further restrict the samples to male household heads. As stated above household composition in our view is one source of risk over the life-cycle which we do not want to censor. In any case, we run a robustness check in which we consider males only and find that results do not change qualitatively.[7] Our baseline sample selection removes all observations with missing income, education or hours data, and removes observations where the household head is below 25 or above 56 years of age. In a separate appendix we provide details on how we treat the three individual data sets.

3.2 Descriptive statistics

3.2.1 Importance of observable characteristics by age

Since we are concerned with the evolution of residual wage dispersion and the cyclicality of wage risk, we need to filter from the observed wages the influence of fixed household characteristics. We do so by estimating (1) in two steps. First we estimate β by OLS and then use the residuals of this estimation, $\hat{\omega}_{it}^{h}$, to calculate their variance and autocovariance for a given cohort in a given year, see (2) and (3). In the first-stage regression, we use dummies for age, education, gender, and household size to control for fixed household characteristics plus year dummies to control for time effects on the wage level. Since deterministic wage profiles in age may be different across education groups, we add third-order polynomials in age for each education class (classes being roughly equivalent to: CSE or below, O-lvl / GCSE, A-lvls or occupational training, some tertiary, first university degree, higher degree).[8]

To obtain a first impression of (a) how much of wage variability is predictable and (b) whether there are substantial differences between the three data sets already at this stage, we calculate the ratio between the variance due to deterministic differences in wages and the variance of the "unfiltered" wage data (controlling only for time effects) within each age group (similar to R^2 statistics conditional on age), $\frac{\sigma^2_{raw}(h)-\sigma^2_{filtered}(h)}{\sigma^2_{raw}(h)}$. The results for these statistics are displayed in Figure **??**(a). As one can see, deterministic differences explain for no age group more than 30% of the variance of wages. In all three data sets, household characteristics become more important over the life-cycle in explaining wage dispersion, with a peak of the importance around 40 years of age, which then levels out. While the profiles are similar overall, a difference appears for younger ages in the GSOEP, where the first stage regression has particularly low explanatory power.

[7] (SEE REFEREE APPENDIX)

[8] Results of the first stage regression are available on request.

Figure 1

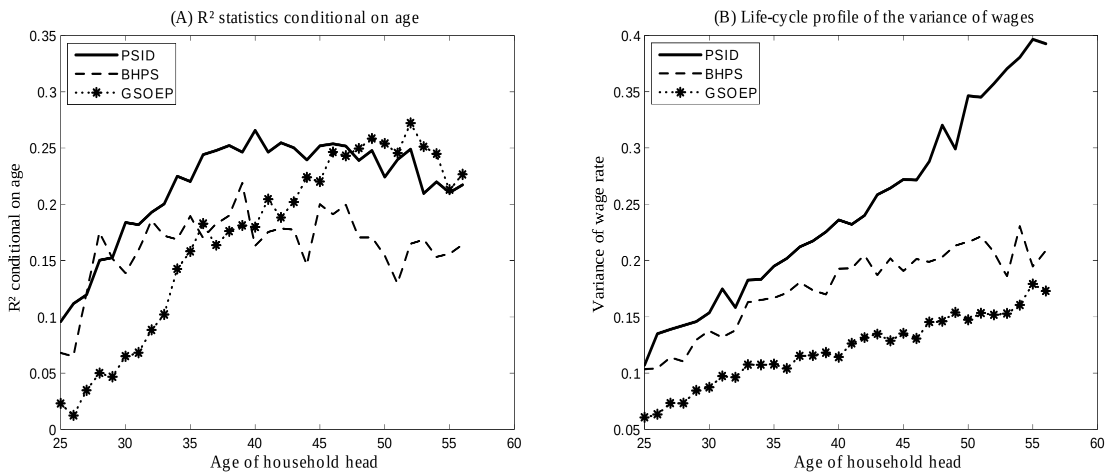

One possible reason could be the larger variability of the age of labor market entry in Germany.

3.2.2 Age profiles of the wage dispersion

Next, we consider a reduced-form life-cycle pattern of the residual wage dispersion as has become standard in this literature since Deaton and Paxson (1994). For this purpose, we regress the empirical variance of the wage residual on age and time dummies. For identification, we assume that time effects are mean zero. Figure **??**(b) displays the corresponding age profiles.

The figure shows that residual wage dispersion is the highest in the US at almost all ages and the lowest in Germany. Moreover, all three countries display age-profiles that are (almost linearly) increasing in age, which suggests high levels of autocorrelation (for an in-depth discussion see e.g. Deaton and Paxson, 1994). Finally, the slope of the age-profile is steepest for the US, which implies that the difference in the variance of wages between countries increases in age.[9]

Measuring age-profiles of dispersion is challenging because of the interplay of time and cohort effects, which cannot be identified separately, see Heathcote et al. (2005) or Krueger et al. (2010). Note, however, that our parametric specification of the wage process with business cycle effects on wage risk and long memory, as we estimate it in the next section, can be interpreted as estimating time-varying cohort effects with correlated

[9] These findings are consistent with Krueger et al. (2010), who argue that one would expect the level of wage disparity to be larger in countries where institutional constraints in the labor market are less severe.

innovations across cohorts. As such it can be regarded as a compromise view.

4 Estimation of wage processes

Having illustrated that the wage processes are likely to have a long memory, we can exploit the business cycle history of each cohort as has been proposed by Storesletten et al. (2004). Table 1 summarizes our estimation results. As a measure of the aggregate cycle, we use linearly detrended standardized first differences of log GDP. To calculate the moments used in the GMM estimation, i.e. the variance and autocovariance for a given cohort at a given point in time, we construct sequences of two-year balanced panels from the original data. We calculate the reported standard errors by running 2,000 bootstrap replications. In each bootstrap we redraw new two-year panels by drawing households of the corresponding original balanced panels.[10] Specification (I) in Table 1 shows our baseline estimation results, while specifications (II) - (VI) provide robustness checks.

Across all specifications, we find that for the US, wage risk is strongly and significantly counter-cyclical, $\bar{\phi}-\underline{\phi} < 0$. This confirms empirical evidence for income risk in the US provided by Storesletten et al. (2004). Wage risk at trend growth, $\frac{\bar{\phi}+\underline{\phi}}{2}$, is larger in the US than for the other two countries, while it is smallest for Germany. Also in stark difference to the US, wage risk is procyclical in Germany. A one standard deviation decrease in the business cycle component of GDP growth increases wage risk in the US by roughly 46% while it decreases wage risk by the same fraction in Germany. For the UK we find procyclicality in wage risk in our baseline specification, but the estimate is not significant and quantitatively much smaller than for Germany. Moreover, it changes sign in some of the robustness checks, so that wage risk can overall be described as acyclical for the UK.

In our baseline specification, the importance of transitory shocks (or measurement error) is similar across data sets as we obtain very similar estimates. Also the persistence of wage shocks is surprisingly similar across data sets with $\hat{\rho} \approx .92$. What differs though is the size of the variance of fixed effects. Yet these variances are not precisely estimated.

In order to check how robust our results are, we run alternative specifications. First, we exclude low skilled workers (II) or high skilled workers (III) to check whether results

[10] We disregard the fact that the wage residuals are results of a first stage regression, such that regression uncertainty in theory adds to the estimation uncertainty in the second step. Yet, with the number of observations being as large as in our samples the variation introduced by regression uncertainty in the first stage is negligibly small. Since some of the age-year cells are very sparsely filled with households, we run into the problem when doing sample splits that the standard bootstrap algorithm may end up with no observation in a given age-year cell. Therefore, we increase the sample size drawn in the bootstrap by factor 2 resp. 4 for education and income splits (always factor 4 for BHPS) and correct the standard errors by premultiplying $\sqrt{2}$ (repectively 2).

are driven by a subgroup in the labor market. Second, we define the cycle based on HP-filtering of log GDP instead of first differences (IV). Third, we assume labor market entry at age 22 instead of 25 (V). Finally, we use individual male wages instead of household average wages (VI). Results with respect to cyclicality do not change across these alternative specifications.

While the robustness checks are reaffirming that there are differences in the cyclicality of wage risk across countries, they are disappointing in hinting towards an explanation of these cross country differences as we find them pervasive across education groups, gender and how we measure wages. Rather our results suggest that differences in labor market institutions might be driving our results. Yet, these differences cannot be measured within the micro data but only in their differences across countries.

5 Conclusion

This paper has provided a cross-country comparison of life-cycle wage dispersions taking into account the business cycle fluctuations in wage risk. We find that wage risk increases in recessions in the US while it decreases in recessions in Germany. For the UK we find wage risk to be by-and-large acyclical. There are two interpretations to this. First, we can think of wages as reflecting marginal productivity as in a neoclassical world. In this case, there would be significant structural differences in the way productivity risk behaves between the three large economies studied. Second, we can think of wages as rather an outcome of more complex economic processes, such as bargaining with unions, search-and-matching, or sticky wage-setting. In this second case, given the structural labor market differences of the three economies studied, our findings suggest new restrictions on the identification of structural labor market models, in the sense that these models should also speak to the cyclicality of wage risk.

Table 1: Summary of estimation results

	PSID	BHPS	GSOEP	PSID	BHPS	GSOEP	PSID	BHPS	GSOEP
	(I) All households			(II) Low skilled excluded			(III) High skilled excluded		
ρ	0.9146	0.9195	0.9045	0.9407	0.9588	0.9106	0.9158	0.9558	0.9050
	(0.032)	(0.0232)	(0.0125)	(0.0238)	(0.0144)	(0.0087)	(0.0228)	(0.0337)	(0.0071)
σ^2_α	0.0230	0.0012	0.0116	0.0620	0.0000	0.0185	0.0419	0.0638	0.0061
	(0.0363)	(0.0244)	(0.0079)	(0.0324)	(0.0036)	(0.0061)	(0.0296)	(0.0331)	(0.0047)
σ^2_ε	0.0405	0.0390	0.0293	0.0487	0.0439	0.0300	0.0408	0.0436	0.0267
	(0.0065)	(0.0044)	(0.0021)	(0.0054)	(0.0041)	(0.0014)	(0.0052)	(0.0049)	(0.0013)
$(\bar{\phi}+\underline{\phi})/2$	0.0322	0.0218	0.0145	0.0186	0.0164	0.0126	0.0272	0.0062	0.0153
	(0.0113)	(0.0069)	(0.0025)	(0.0095)	(0.0017)	(0.0018)	(0.0094)	(0.0092)	(0.0015)
$\bar{\phi}-\underline{\phi}$	-0.0644	0.0110	0.0290	-0.0372	-0.0226	0.0252	-0.0545	-0.0003	0.0306
	(0.0226)	(0.0099)	(0.0050)	(0.0189)	(0.0104)	(0.0035)	(0.0188)	(0.0065)	(0.0030)
	(IV) GDP HP(100)-filtered			(V) Labor market entry at age 22			(VI) Individual male wages		
ρ	0.8992	0.9194	0.8971	0.9347	0.9445	0.9220	0.9206	0.9595	0.9243
	(0.0458	(0.0256)	(0.0756)	(0.0249)	(0.0146)	(0.0104)	(0.0383)	(0.1827)	(0.0440)
σ^2_α	0.0000	0.0000	0.0535	0.0000	0.0000	0.0000	0.0344	0.0459	0.0140
	(0.0364)	(0.0036)	(0.0251)	(0.0303)	(0.0206)	(0.0033)	(0.0378)	(0.0303)	(0.0140)
σ^2_ε	0.0351	0.0391	0.0329	0.0423	0.0423	0.0297	0.0448	0.0452	0.0294
	(0.0086)	(0.0043)	(0.0051)	(0.0050)	(0.0033)	(0.0018)	(0.0071)	(0.0079)	(0.0028)
$(\bar{\phi}+\underline{\phi})/2$	0.0423	0.0219	0.0067	0.0282	0.0155	0.0138	0.0279	0.0080	0.0108
	(0.0142)	(0.0070)	(0.0092)	(0.0070)	(0.0043)	(0.0016)	(0.0118)	(0.0095)	(0.0039)
$\bar{\phi}-\underline{\phi}$	-0.0845	-0.0134	0.0135	-0.0564	0.0150	0.0276	-0.0557	-0.0046	0.0215
	(0.0286)	(0.0092)	(0.0083)	(0.0140)	(0.0091)	(0.0033)	(0.0244)	(0.0174)	(0.0078)

Estimation results from a generalized method of moments estimation as given by eq. (4). Standard errors (in parentheses) obtained from 2,000 bootstrap replications. The business cycle indicator is linearly detrended first differences in log GDP unless stated otherwise. Education splits are twice over-sampled (factor 4 for BHPS) and the standard errors corrected by premultiplying $\sqrt{2}$ (BHPS: 2). Low skilled: GCSE and below, High skilled: first university degree and above.

References

[1] Alexopoulos, Michelle and Jon Cohen (2009): "Uncertain Times, uncertain measures," Working Papers tecipa-352, University of Toronto.

[2] Bachmann, Rüdiger and Christian Bayer (2009a): "Firm-Specific Productivity Risk over the Business Cycle: Facts and Aggregate Implications", Bundesbank Discussion Paper Series 1: 17/2009.

[3] Bachmann, Rüdiger and Christian Bayer (2009b): "The Cross-section of Firms over the Business Cycle: New Facts and a DSGE Exploration", Bundesbank Discussion Paper Series 1: 15/2009.

[4] Biewen, Martin (2005): "The covariance structure of East and West German incomes and its implications for the persistence of poverty and inequality", German Economic Review, 6, 445–469.

[5] Blundell, Richard, Luigi Pistaferri and Ian Preston (2008): "Consumption Inequality and Partial Insurance," American Economic Review, 98(5), 1887-192.

[6] Blundell, Richard and Ben Etheridge (2010): "Consumption, Income, and Earnings Inequality in the UK", Review of Economic Dynamics, 13(1), 76-102.

[7] Bloom, Nick, Max Floetotto and Nir Jaimovich (2009): "Really Uncertain Business Cycles", mimeo, Stanford University .

[8] Bloom, Nick (2009): "The Impact of Uncertainty Shocks," Econometrica, 3(5), 623-685.

[9] Deaton, Angus and Christina Paxson (1994): "Intertemporal Choice and Inequality," Journal of Political Economy, 102(3), 437-467.

[10] Dickens, Richard (2000): "The Evolution of Individual Male Earnings in Great Britain: 1975-95," Economic Journal, 110(460), 27-49.

[11] 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), 103-132.

[12] Gilchrist, Simon, Jae Sim, and Egon Zakrajšek (2009): "Uncertainty, Credit Spreads, and Investment Dynamics", mimeo.

[13] Gomme, Paul, Richard Rogerson, Peter Rupert and Randall Wright (2004): "The Business Cycle and the Life Cycle," NBER Macroeconomics Annual 2004.

[14] Gosling, Amanda, Stephen Machin and Costas Meghir (2000): "The Changing Distribution of Male Wages in the U.K," Review of Economic Studies, 67(4), 635-66.

[15] Guvenen, Fatih (2007): "Learning Your Earning: Are Labor Income Shocks Really

Very Persistent?," American Economic Review, 97(3), 687-712.

[16] Guvenen, Fatih (2009): "An Empirical Investigation of Labor Income Processes," Review of Economic Dynamics, 12(1), 58-79.

[17] Heathcote, Jonathan, Kjetil Storesletten and Gianluca Violante (2005): "Two Views of Inequality Over the Life-Cycle," Journal of the European Economic Association, 3(2-3), 543-52.

[18] Heathcote, Jonathan, Kjetil Storesletten and Gianluca Violante (2009): "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," working paper.

[19] Heathcote, Jonathan, Fabrizio Perri and Gianluca Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006", Review of Economic Dynamics, 13(1), 15-51.

[20] Krebs, Tom (2003): "Human Capital Risk And Economic Growth," The Quarterly Journal of Economics, 118(2), 709-744.

[21] Krebs, Tom (2007): "Job Displacement Risk and the Cost of Business Cycles," American Economic Review, 97(3), 664-686.

[22] Krueger, Dirk and Fabrizio Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," Review of Economic Studies, 73(1), 163-193.

[23] Krueger, Dirk, Fabrizio Perri, Luigi Pistaferri, Giovanni L. Violante (2010): "Cross-sectional facts for macroeconomists", Review of Economic Dynamics, 13(1), 1-14.

[24] Ljungqvist, Lars and Thomas J. Sargent (2008): "Two Questions about European Unemployment," Econometrica, 76(1), 1-31.

[25] Primiceri Giorgio E. and Thijs van Rens (2009): "Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality," Journal of Monetary Economics, 56 (1), 20-39.

[26] Storesletten, Kjetil, Chris I. Telmer and Amir Yaron (2001): "The Welfare Costs of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk," European Economic Review, 45 (7), 1311-1339.

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