

IZA DP No. 4402

**The Life-Cycle and the Business-Cycle
of Wage Risk: A Cross-Country Comparison**

Christian Bayer
Falko Juessen

September 2009

The Life-Cycle and the Business-Cycle of Wage Risk: A Cross-Country Comparison

Christian Bayer

IGIER, Università Bocconi

Falko Juessen

*Technische Universität Dortmund
and IZA*

Discussion Paper No. 4402
September 2009

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

The Life-Cycle and the Business-Cycle of Wage Risk: A Cross-Country Comparison^{*}

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.

JEL Classification: E20, D31, D91, J31

Keywords: life-cycle risk, uncertainty fluctuations, business cycle, heterogeneity, wages

Corresponding author:

Falko Juessen
Department of Economics
Technische Universität Dortmund
44221 Dortmund
Germany
E-mail: falko.juessen@tu-dortmund.de

^{*} 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, for helpful comments and suggestions.

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.¹ 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 cyclicity of the cross-sectional dispersion of wage-innovations instead of innovations to earnings. This is important as wages are more closely linked to productivity and are less subject to endogenous household decisions than earnings. Second, we contribute by providing a cross-country perspective on the cyclicity of the dispersion of wage-innovations, comparing the US, the UK, and Germany. This means we provide results across a set of different labor market setups, in which the US labor market is the most liberal and the German labor market the most regulated one, while the UK labor market can be considered the intermediate case. Interestingly, we find this ranking to be also reflected empirically in the cyclicity of wage risk: the US displays countercyclical, the UK acyclical, and Germany procyclical dispersions in wage innovations.

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) or

¹The life-cycle of earnings inequality without explicitly linking wage risk to the business cycle has been studied in various papers. For the three countries we focus on - UK, US, Germany - examples are: (for the UK) Dickens (2000), Gossling et al. (2000), Blundell et al. (2006), and Blundell and Etheridge (2009) (for the US) Guvenen (2007a,b), Krueger and Perri (2006), Primiceri and van Rens (2007), and Heathcote et al. (2009), (for Germany) Krueger et al. (2009).

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 cyclical nature of wage risk may help discriminating between (quantitatively) different labor market models, similar to the point made by Bachmann and Bayer (2009b) for the cyclical nature of the investment rate dispersion, which is found to be a strong overidentifying restriction in heterogeneous 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. (2007) or Sargent and Ljungqvist (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 Krebs (2003, 2007).

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. Appendices follow.

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

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

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 $\sigma_\nu^2 = 1$. We assume that the variance of fixed effects as well as the variance of measurement error, σ_α^2 and σ_ε^2 , are constant over time, while the variance of persistent shocks is potentially time varying. 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

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

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 $\phi_t > 0$. We choose Y_t to be standardized GDP deviations from trend, such that $\frac{\phi + \bar{\phi}}{2}$ is the variance at trend growth and the sign of $(\bar{\phi} - \underline{\phi})$ 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 . This allows us to exploit business cycle history beyond the actual time-span of our panel data on wages. For a household with labor market experience h the variance of ω_{it}^h is given by

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

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:

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

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.

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

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

where \mathbf{W} is a positive definite weighting matrix that captures the correlation structure across moments, see Appendix A for details. We evaluate the properties of this estimator in a number of Monte-Carlo experiments and find that the estimator is capable of identifying the cyclicity of wage risk in data sets of a size comparable to the data sets at hand, see Appendix B.

3 Data

3.1 Sample Selection

Our data base consists of three large household panel data sets 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. Yet, 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.

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 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 we understand household composition as one fundamental source of income and productivity risk over the life-cycle, e.g. due to death of the main earner. Yet, we run a robustness check in which we consider males only.

We restrict the data sets as follows. We remove all observations with missing income, education or hours data, and remove observations where the household head is below 25 or above 56 years of age. For the PSID, we drop all households that belong to the Survey of Economic Opportunity (SEO) or the Latino sample. For the BHPS, we keep only households living in England.² For the GSOEP, we restrict the sample to West German households and drop observations from the migrant and high income sample. Moreover, we remove all households where the household head is a migrant who has immigrated to Germany after the age of ten. We expect the productivity life cycle of migrants to be governed by other institutional factors (e.g. education) than those of non-migrants.³

We use the age h of the household head to attach a working and business cycle history to each household in a given year t . We define an age-year (h, t) cell as the sub-sample of households in year t with the same year of birth $c = t - h$, irrespective of the exact date of the survey interview relative to the exact date-of-birth of the interviewee. From the information of the household's labor income y_{it} and hours worked η_{it} we calculate the household's wage rate $w_{it} := y_{it}/\eta_{it}$. We restrict the sample to those households that supply at least 520 hours of market work per year, being equivalent to at least a quarter of full-time employment of one household member (assuming 40 hours per week). In addition, we apply a three standard deviation criterion to the wage rate for each age-year cell of households to identify outliers, which we then remove from the sample. We experimented with alternative outlier selection criteria which in general lead to similar results, see Section 4.2. Finally, a household enters the sample only if we observe the household at least in two consecutive years since we want to calculate autocovariances. This means we generate a set of two-year overlapping panels from the original data, which we then use for the analysis. Tables 8 and 7 in Appendix C provide information

²The BHPS started with mainly households living in England. In later sample waves, households from Wales, Scotland, and Northern Ireland were added to the BHPS, which implies that these economically diverse parts of the UK do not have a constant sampling weight.

³We treat the GSOEP data differently to the other data sources in this respect, because in the GSOEP the year of entry into the country of a household member is easily available. Controlling for the status of a migrant could be important for three reasons: First, it is not clear how comparable years of schooling are for migrants doing their schooling abroad. Second, migrants who enter the country in working age have gone through a different business cycle history. Third, migration often comes in waves that are likely to be related to the business cycle.

on the number of observations (i.e. households) by year in each two-year panel as well as on the number of observations we loose due to the sample selection.

3.2 Descriptive Statistics

3.2.1 Importance of Observable Characteristics by Age

Since we are concerned with the evolution and cyclicity 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 the moments as in (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. 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).⁴

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). The results for $\frac{\sigma_{raw}^2(h) - \sigma_{filtered}^2(h)}{\sigma_{raw}^2(h)}$ are displayed in Figure 1. 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. 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 $m_{1,t,h}$ on age and time dummies. For identification, we assume that time effects are mean zero. Figure 2 displays the corresponding age profiles.

⁴Results of the first stage regression are available on request.

Figure 1: R2 Statistics for the Wage Regression on Household Characteristics





Figure 2: Life-Cycle Profile of the Variance of Wages



Results from a year-time decomposition of the age-year specific variances $m_{1,t,h}$ as in Deaton and Paxson (1994). Time effects are assumed to be mean zero for identification. The graph displays the estimated age coefficients.

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.

4 Estimation of Wage Processes

4.1 Baseline Specification

Having illustrated that the wage process is 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). Yet, we need to decide on the subset of moments we want to match from the vector of moments \mathbf{m} . In their baseline specification, Storesletten et al. (2004) restrict the set of moments to those from a subset of age groups, a theoretical reason for which could be that further moment conditions provide little extra information but only increase estimation uncertainty (i.e. moment conditions being "weak"). However, experimenting with such selected age groups in a set of Monte-Carlo experiments did not show any evidence of a problem of weak moment conditions, see Appendix B. Correspondingly, increasing the number of moment conditions in general increases the precision of the estimator. Therefore, we use all ages between 30 and 55 and all available years for each sample. We restrict ourselves to this prime-age group in order to eliminate the effect of country-specific differences in household formation (age below 30) and retirement (at ages 55+) that may otherwise influence our results. We assume that households enter the labor market - more specifically, accumulate labor market shocks - at the age of 25.

Table 1 presents the estimation results from our baseline specification. As a measure of the aggregate cycle we use linearly detrended standardized first differences of log-GDP. We calculate standard errors by running 1,000 bootstrap replications drawing households in each two-year panel.⁵

We find that for the US, wage risk is strongly and significantly counter-cyclical, $\bar{\phi} - \underline{\phi} < 0$. This confirms 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 countries, while it is smallest for Germany. Also in stark difference to the US, wage

⁵We 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.

Table 1: Baseline Estimation Results

	PSID	BHPS	GSOEP
ρ	.9245 (.0300)	.9246 (.0255)	.9190 (.0136)
σ_α^2	.0443 (.0383)	.0096 (.0257)	.0280 (.0089)
σ_ε^2	.0440 (.0063)	.0404 (.0046)	.0316 (.0019)
$(\bar{\phi} + \underline{\phi}) / 2$.0252 (.0109)	.0192 (.0074)	.0102 (.0025)
$\bar{\phi} - \underline{\phi}$	-.0504 (.0218)	.0107 (.0091)	.0204 (.0049)

Estimation results from a generalized method of moments estimation as given by eq. (4), see appendix for details. Standard errors (in parentheses) obtained from 1,000 bootstrap replications.

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, but the estimate is not significant and quantitatively much smaller than for Germany, so that wage risk can be described as acyclical for the UK. In fact, the point estimate for the UK changes sign when we use different measures of cyclicity as in the robustness checks below. By contrast, the findings for Germany and the US turn out to be robust.

Interestingly, the importance of transitory shocks (or measurement error) seems to be 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.

4.2 Robustness Checks

Next, we check the robustness of these findings to changes in the measure of cyclicity and sample construction. First, we check the robustness to alternative measures of the cycle. As alternatives, we capture the cyclical component of GDP by HP-filtering with a smoothing parameter of 6.25 following Uhlig and Ravn (2002) and HP-filtering with

Table 2: Robustness to Cyclical Indicator

	GDP-HP(6.25)			GDP-HP(100)		
	PSID	BHPS	GSOEP	PSID	BHPS	GSOEP
ρ	.9138 (.0469)	.9275 (.0240)	.8991 (.0362)	.9000 (.0389)	.9189 (.0264)	.9097 (.0148)
σ_α^2	.0319 (.0484)	.0147 (.0242)	.0441 (.0178)	.0000 (.0369)	.0000 (.0257)	.0383 (.0085)
σ_ε^2	.0413 (.0085)	.0413 (.0041)	.0319 (.0037)	.0354 (.0078)	.0391 (.0044)	.0319 (.0019)
$(\bar{\phi} + \underline{\phi})/2$.0306 (.0151)	.0177 (.0066)	.0087 (.0064)	.0419 (.0135)	.0220 (.0070)	.0092 (.0021)
$\bar{\phi} - \underline{\phi}$	-.0613 (.0218)	-0.0234 (.0091)	.0175 (.0049)	-.0839 (.0270)	-.0151 (.0097)	.0183 (.0043)

See notes to Table 1.

Table 3: Robustness to Sample Selection

	Percentile-Outlier-Criterion			Labor Market Entry at Age 22		
	PSID	BHPS	GSOEP	PSID	BHPS	GSOEP
ρ	.9319 (.0209)	.9352 (.0249)	.9218 (.0114)	.9349 (.0304)	.9450 (.0153)	.9328 (.0090)
σ_α^2	.0640 (.0389)	.0248 (.0269)	.0253 (.0079)	.0000 (.0393)	.0000 (.0214)	.0098 (.0081)
σ_ε^2	.0496 (.0059)	.0420 (.0044)	.0316 (.0018)	.0424 (.0055)	.0424 (.0034)	.0313 (.0018)
$(\bar{\phi} + \underline{\phi})/2$.0213 (.0102)	.0147 (.0070)	.0102 (.0022)	.0279 (.0086)	.0154 (.0044)	.0112 (.0018)
$\bar{\phi} - \underline{\phi}$	-.0426 (.0202)	.0014 (.0088)	.0204 (.0044)	-.0558 (.0173)	.0133 (.0077)	.0224 (.0037)

See notes to Table 1.

Table 4: Sample Split According to Skill

	Low and Medium Skilled			High and Medium Skilled		
	PSID	BHPS	GSOEP	PSID	BHPS	GSOEP
ρ	.9256 (0.0386)	.9555 (0.0662)	.9165 (0.0119)	.9512 (0.0330)	.9572 (0.0137)	.9263 (0.0134)
σ_α^2	.0612 (0.0407)	.0637 (0.0641)	.0196 (0.009)	.0771 (0.0422)	.0000 (0.0095)	.0343 (0.0081)
σ_ε^2	.0439 (0.0067)	.0436 (0.0097)	.0287 (0.0019)	.0507 (0.0064)	.0438 (0.0084)	.0322 (0.0018)
$(\bar{\phi} + \underline{\phi}) / 2$.0211 (0.0119)	.0062 (0.0179)	.0118 (0.0025)	.0143 (0.0112)	.0166 (0.0035)	.0086 (0.0022)
$\bar{\phi} - \underline{\phi}$	-.0422 (0.0234)	-.0003 (0.0108)	.0237 (0.0049)	-.0287 (0.0226)	-.0146 (0.0220)	.0172 (0.0044)

Estimation results from a generalized method of moments estimation as given by eq. (4), see appendix for details. Standard errors (in parentheses) obtained from 1,000 bootstrap replications. Since some of the age-year cells are very sparsely filled with households, we run into the problem 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 (factor 4 for BHPS) and correct the standard errors by premultiplying $\sqrt{2}$ (respectively 2 for the BHPS).

Low skilled: GCSE and below, *Medium skilled:* A-lvls and some tertiary, *High skilled:* first university degree and above.

the more standard smoothing parameter of 100, see Table 2. As one can see, our results are qualitatively and quantitatively robust to the choice of the cyclical indicator. The only apparent difference is that using the standard HP-filter for aggregate GDP, we run into a corner solution for the estimated variance of fixed effects.

Finally, we check the robustness to our sample selection criteria. First, we try with dropping only those households that fall in the top-bottom 0.5% percentiles of each age-group, respectively (instead of using a three standard deviation criterion in an age-year cell). Second, we check to which extent our results are sensitive to the assumption that households start to accumulate wage shocks at the age of 25. Instead, we set this age of labor-market entry to 22, but keep the set of age groups used for estimation constant (30-55). Results are displayed in Table 3. Third, we look at a subset of households either excluding the very high skilled or the very low skilled, see Table 4. We do so in order

Table 5: Individual male wages

	Labor Market Entry at Age 25			Labor Market Entry at Age 22		
	PSID	BHPS	GSOEP	PSID	BHPS	GSOEP
ρ	0.9668 (.0809)	0.9846 (.0506)	0.9228 (.0404)	0.9553 (.0739)	0.9749 (.0385)	0.9372 (.0204)
σ_α^2	0.1262 (.0545)	0.0951 (.0358)	0.296 (.0133)	0.0741 (.0647)	0.0725 (.0356)	0.0100 (.0105)
σ_ε^2	0.0572 (.0095)	0.0388 (.0061)	0.0271 (.0032)	0.0534 (.0083)	0.0377 (.0051)	0.0268 (.0026)
$(\bar{\phi} + \underline{\phi}) / 2$	0.0083 (.0162)	0.0025 (.0107)	0.0090 (.0040)	0.0149 (.0137)	0.0045 (.0085)	0.0102 (.0025)
$\bar{\phi} - \underline{\phi}$	-0.0165 (.0321)	-0.0050 (.0142)	0.0180 (.0080)	-0.0299 (.0083)	0.0071 (.0142)	0.0204 (.0050)

See notes to Table 1.

to eliminate the effects of the very low skilled workers dropping out of the sample into unemployment along the cycle. Finally, we restrict the sample to individual wage data for males in order to exclude cyclical effects from within-household risk sharing, in particular the possibility of females dropping out of the labor force at economic downturns. The results of this specification are presented in Table 5.

Again, our results turn out to be robust. Moving to individual data increases the estimated autocorrelation somewhat and correspondingly decreases the variance at trend growth, $(\bar{\phi} + \underline{\phi}) / 2$, and the cyclicity parameter, $\bar{\phi} - \underline{\phi}$. Notwithstanding, a one standard deviation in detrended GDP changes risk by 46% for both the US and Germany. Thus, the documented patterns on cyclicity can also be found for male prime-age workers, a group of workers which is less subject to cyclical nonemployment. In summary, our previous findings do not only reflect the cyclical behavior of women's labor market participation. Yet, the results become less statistical significant.

In a similar vein, we can read the results in Table 4, where we focus on households with high or medium levels of education - thus leaving out unskilled workers. Results change little for this group, although fluctuations in unemployment are typically concentrated among the unskilled.

5 Final Thoughts

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 cyclicity of wage risk.

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", mimeo, U Michigan
- [3] Bachmann, Rüdiger and Christian Bayer (2009b): "The Cross-section of Firms over the Business Cycle: New Facts and a DSGE Exploration", mimeo, U Michigan
- [4] Blundell, Richard, Luigi Pistaferri and Ian Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887-192.
- [5] Blundell, Richard and Ben Etheridge (2009): "Consumption, Income, and Earnings Inequality in the UK", mimeo.
- [6] Bloom, Nick, Max Floetotto and Nir Jaimovich (2009): "Really Uncertain Business Cycles", mimeo, Stanford University .
- [7] Bloom, Nick (2009): "The Impact of Uncertainty Shocks," forthcoming *Econometrica*.
- [8] Deaton, Angus and Christina Paxson (1994): "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102(3), 437-467.
- [9] Dickens, Richard (2000): "The Evolution of Individual Male Earnings in Great Britain: 1975-95," *Economic Journal*, 110(460), 27-49.
- [10] Gilchrist, Simon, Jae Sim, and Egon Zakrajšek: "Uncertainty, Credit Spreads, and Investment Dynamics", mimeo.
- [11] Gomme, Paul, Richard Rogerson, Peter Rupert and Randall Wright (2004): "The Business Cycle and the Life Cycle," NBER Macroeconomics Annual 2004.
- [12] 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.
- [13] Guvenen, Fatih (2007): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?," *American Economic Review*, 97(3), 687-712.
- [14] Guvenen, Fatih (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12(1), 58-79.
- [15] Heathcote, Jonathan, Kjetil Storesletten and Gianluca Violante (2007): "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," working paper.
- [16] Jonathan Heathcote, Fabrizio Perri and Gianluca Violante (2009): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-

- 2006", mimeo.
- [17] Krebs, Tom (2003): "Human Capital Risk And Economic Growth," *The Quarterly Journal of Economics*, 118(2), 709-744.
 - [18] Krebs, Tom (2007): "Job Displacement Risk and the Cost of Business Cycles," *American Economic Review*, 97(3), 664-686.
 - [19] Krueger, Dirk and Fabrizio Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73(1), 163-193.
 - [20] Krueger, Dirk, Nicola Fuchs-Schündeln and Mathias Sommer (2009): "Inequality Facts for Germany: Levels, Trends, and their Interpretation", mimeo.
 - [21] Ljungqvist, Lars and Thomas J. Sargent (2008): "Two Questions about European Unemployment," *Econometrica*, 76(1), 1-31.
 - [22] Primiceri Giorgio E. and Thijs van Rens (2007): "Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality," *IZA Discussion Papers* 3239.
 - [23] Ravn, Morten O. and Harald Uhlig (2002): "On adjusting the Hodrick-Prescott filter for the frequency of observations," *The Review of Economics and Statistics*, 84(2), 371-375.
 - [24] Storesletten, Kjetil, Chris I. Telmer and Amir Yaron (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112(3), 695-717.

Appendices

A Details on the Estimation Procedure

The GMM estimator used in this paper employs a weighting matrix to gain efficiency over the simple unweighted minimum distance estimator. In principle, one should weight the moment-distances $(\mu - m)$ with the inverse covariance matrix of each moment condition. In practice, such approach involves estimation of the covariance matrix and can lead to more imprecision in small samples relative to the use of a unitary weighting matrix.

We follow Storesletten et al. (2004) and take an intermediate approach, where we estimate the most apparent part of the covariance matrix that results from the definition of our moment conditions themselves. Yet, we ignore the covariance that results from the fact that the sequence of two year panels that we construct from the data partly overlap in the set of households they contain. Instead, we treat the two year panels as if each of them was sampled independently.

This implies that we presume variance ($k = 1$) and autocovariance ($k = 2$) moments, $m_{k,t,h}$ and $m_{k,t+s,h+j}$, are uncorrelated as long as $s \neq 0$ or $j \neq 0$, since then the wage information that they summarize stems from different, independent households. For $s \neq j$, this assumption holds as each household ages by one year each period. For $s = j$, we have to invoke the assumption that the overlap of the two year panels is negligible. As argued, we make this assumption to avoid small sample biases from imprecise covariance estimation.

By contrast, moments $m_{1,t,h}$ and $m_{2,t,h}$ exhibit some clear cut correlation, as they both exploit information of the wages of the same set of individuals at time t (as well as time $t+1$). Moreover, they are also correlated from the fact that the future wage in period $t + 1$ of these households that is used to construct $m_{2,t,h}$ is correlated with the wage in period t . The covariance $\sigma_{12}(t, h) = cov(m_{1,t,h}, m_{2,t,h}) \in R^{2 \times 2}$ can be easily estimated within sample. The moment conditions refer to the product $\hat{\omega}_{i,t,h}^2$ and $\hat{\omega}_{i,t,h}\hat{\omega}_{i,t+1,h+1}$ at the household level, such that we estimate $\sigma_{12}(t, h)$ as the sample covariance

$$\hat{\sigma}_{12}(t, h) = \frac{1}{N_{th}} \sum_{i=1}^{N_{t,h}} \begin{pmatrix} \hat{\omega}_{i,t,h}^4 & \hat{\omega}_{i,t,h}^3 \hat{\omega}_{i,t+1,h+1} \\ \hat{\omega}_{i,t,h}^3 \hat{\omega}_{i,t+1,h+1} & \hat{\omega}_{i,t,h}^2 \hat{\omega}_{i,t+1,h+1}^2 \end{pmatrix} \\ - \frac{1}{N_{th}^2} \sum_{i=1}^{N_{t,h}} \begin{pmatrix} \hat{\omega}_{i,t,h}^2 & \\ \hat{\omega}_{i,t,h} \hat{\omega}_{i,t+1,h+1} & \end{pmatrix} \sum_{i=1}^{N_{t,h}} \begin{pmatrix} \hat{\omega}_{i,t,h}^2 & \hat{\omega}_{i,t,h} \hat{\omega}_{i,t+1,h+1} \end{pmatrix},$$

where N_{th} is the number of observations in each age-year cell. With the correct ordering

of moments $\left(\mathbf{m} := \text{vec} \begin{pmatrix} \mathbf{m}'_1 \\ \mathbf{m}'_2 \end{pmatrix}\right)$ the weighting matrix \mathbf{W} then takes block diagonal form with $N_{th}^2 \hat{\sigma}_{12}^{-1}(t, h)$ as diagonal elements. In particular, this means we weight moment conditions more that stem from age-year cells that hold more observations and that are less kurtotic.

B Monte Carlo Results for the GMM Estimator

To check for potential small sample bias that might be inherent in the GMM method and to understand how the estimator would be affected by reducing the number of moment restrictions by selecting only a subset of age-groups, we run a set of Monte-Carlo experiments. We simulate the data generating process (DGP) for residual wages $\omega_{i,h,t}$ as given in (1) and use an AR(1) process with autocorrelation ρ_Y to simulate GDP deviations from trend. We do so for $N = 60$ households for birthcohorts $c = 1, \dots, 58$ (in line with the birthcohorts in the PSID - 1913 until 1971). We simulate each of these cohorts for $H = 32$ years (again in line with the ages 25-56 we consider). Finally, we restrict the sample to the years $t = c + h = 55, \dots, 81$. This way we obtain a sample similar to the PSID (though perfectly balanced).

For the DGP we set $\rho = 0.93$, $\sigma_a^2 = 0.04$, $\sigma_\varepsilon^2 = 0.04$, $(\bar{\phi} + \underline{\phi})/2 = 0.025$, $(\bar{\phi} - \underline{\phi}) = -0.045$, and $\rho_Y = 0.6$ roughly in line with the PSID estimates. We assume normally distributed innovations. We compare the estimator that uses all age groups $h = 6, \dots, 31$ to an estimator that uses only information from $h = 6, 16, 26$ (corresponding to ages 30, 40, 50 in the actual data). Table 6 displays the results of these experiments.

The results show that there is no strong small sample bias on any of the estimated parameters. Yet, average estimates and true parameters differ slightly. The standard errors we obtain and the small sample bias are lower when all moment conditions are imposed, showing that there is no apparent problem from weak moment conditions. The efficiency gain from using *more* moment conditions is most pronounced for the cyclicity parameter $\bar{\phi} - \underline{\phi}$ — which is intuitive as using information on more age groups implies more business cycle information being exploited.

Overall, the standard errors in this experiment are comparable to the ones obtained by running the bootstrap on the actual data as in Tables 1 to 3. Standard errors are slightly smaller in the simulation experiment but in general of the same order of magnitude. The difference is likely due to the fact that the actual data is not balanced in the size of age-year cells as is the experimental data. In particular, we observe that also in the simulated data, the variance of fixed effects cannot be estimated very precisely — i.e. it is only weakly identified.

Table 6: Results from a Monte Carlo experiment for the GMM estimator employed

	ρ	σ_α^2	σ_ε^2	$(\bar{\phi} + \underline{\phi}) / 2$	$\bar{\phi} - \underline{\phi}$
Full Set of Moment Conditions	0.9272 (0.0110)	0.0311 (0.0221)	0.0386 (0.0036)	0.0277 (0.0067)	-0.0452 (0.0068)
Only ages 6,16,26	0.9265 (0.0168)	0.0299 (0.0259)	0.0382 (0.0054)	0.0285 (0.0089)	-0.0450 (0.0123)
True Parameters	0.9300	0.0400	0.0400	0.0250	-0.0450

Results from 5,000 replications of a Monte Carlo experiment of the estimator described in the paper and appendix A using an age-year cell size of $N = 60$. The numbers refer to the average estimate, numbers in brackets give standard deviations.

C Sample Selection

The following tables display the sample selection.

Table 7: Sample Selection

	PSID	BHPS	GSOEP
Initial Number of Observations	68219	38733	60488
Eliminated Number of Observations:			
– Due to elimination of outliers, hours, age, ≥ 2 obs per HH, only HH w/ less than 14 HH members	5269	6734	6995
– Due to non-consecutive observations	7999	7470	8359
Final sample	54951	24529	45134

Table 8: Number of Observations by Year

Year	1968	1969	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980
PSID	1300	1312	1335	1357	1367	1448	1508	1564	1612	1653	1725	1797	1836
BHPS													
GSOEP													

Year	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993
PSID	1878	1913	1966	1991	2043	2079	2126	2184	2202	2226	2257	2164	2212
BHPS											1509	1485	1498
GSOEP				1595	1597	1678	1699	1682	1676	1695	1716	1706	1679

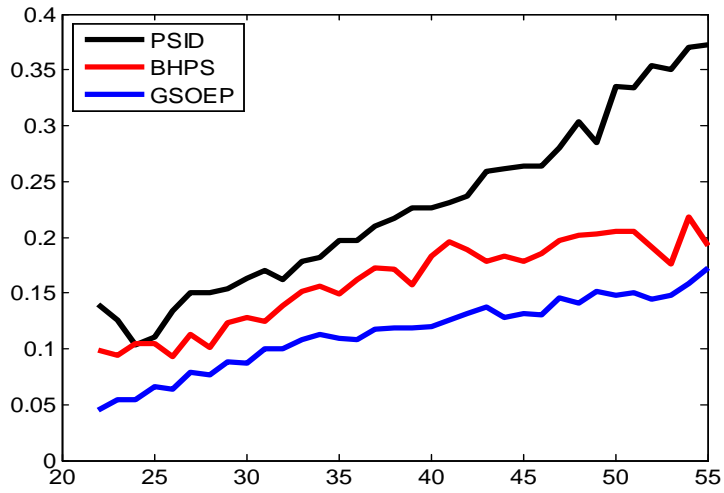
Year	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006
PSID	2616	2692	2588										
BHPS	1487	1538	1570	1735	1621	1702	1685	1492	1485	1486	1432	1415	1389
GSOEP	1670	1662	1641	1608	1877	1809	3428	3191	3085	2963	2814	2663	

The number of observations refers to the number of households in each two-year subpanel with consecutive observations in years $t, t + 1$.

D Further Graphs


Extending the age of entry to 22 does not change the general slope of the life-cycle profile, see figure 3. In line with Krueger et al. (2009) we find a slight J-shape which is more pronounced in incomes, see figure 4

Figure 3: Life-Cycle Profile of the Variance of Wages starting at age 22



See notes to Figure 2.

Figure 4: Life-Cycle Profile of the Variance of Income starting at age 22



See notes to Figure 2.