Hours risk, wage risk, and life-cycle labor supply *

Robin Jessen[†]
Johannes König[‡]

January 22, 2019

Abstract

We decompose permanent earnings risk into contributions from hours and wage shocks. In order to distinguish between hours shocks and labor supply reactions to wage shocks we use a life-cycle model of consumption and labor supply. Estimating our model with the Panel Study of Income Dynamics (PSID) shows that both permanent wage and hours shocks play an important role in explaining earnings risk, but wage shocks have greater relevance. Allowing for hours shocks improves the model fit considerably. The empirical strategy allows for the estimation of the Marshallian labor supply elasticity without the use of consumption or asset data. We find this elasticity on average to be negative, but small. The degree of consumption insurance implied by our results is in line with recent estimates in the literature.

Keywords Earnings Risk · Wage Risk · Frisch Elasticity · Marshallian Elasticity · Consumption Insurance

JEL Classification D31 · J22 · J31

*We thank Richard Blundell, Flora Budianto, Roland Döhrn, Steffen Elstner, Giulio Fella, Michael Graber, Thorben Korfhage, Jack Light, Maria Metzing, Ian Preston, Itay Saporta-Eksten, Carsten Schröder, Viktor Steiner, Alexandros Theloudis, Guglielmo Weber, seminar participants at the Freie Universität Berlin, the EALE conference in 2018 in Lyon, and ESWM 2018 in Naples for valuable comments.

[†]RWI, Invalidenstraße 112, 10115 Berlin , Germany (e-mail: robin.jessen@rwi-essen.de).

[‡]DIW Berlin; Department of Economics, Freie Universität Berlin, Boltzmannstr. 20, 14195 Berlin, Germany (e-mail: j.koenig@fu-berlin.de).

1 Introduction

What drives the riskiness of earnings? A glance at the recent literature on life-cycle consumption, saving and labor supply suggests an implicit consensus: shocks to wages are the central source of risk. In this paper we re-open this discussion by starting from the natural decomposition of earnings into hours worked and wages. Thus, the main contribution of the paper is a decomposition of earnings risk along these lines: We tailor a structural model of life-cycle labor supply to feature earnings risk from both wage and hours shocks and assess the strength of their contributions to total earnings risk.

Knowing the extent to which hours and wage shocks contribute to total income risk is of general interest as it should inform future modeling decisions when estimating permanent earnings risk. Further, it is an indicator of the effectiveness of specific policy measures aimed at reducing income risk. For instance, if income risk was driven almost entirely by hours risk, devising policies to reduce the impact of shocks to wages would not be a fruitful endeavor. Moreover, we find that the estimate of the Marshallian labor supply elasticity is sensitive to this modeling choice.

In our model individuals face idiosyncratic shocks to their productivity of market work, which result in wage shocks, as is standard. For instance, a promotion is a positive permanent wage shock and loss of human capital a negative one. Our concise extension of the standard life-cycle model of consumption and labor supply is to model hours shocks as innovations in the disutility of work. These shocks are conceptualized in an analogous fashion to wage shocks, for instance as shocks to home production. Consider the case of an elderly parent falling ill or being in need of around-the-clock care; this increases the opportunity costs of market work sharply. Depending on the nature of the illness, the shock is permanent or fades out. In terms of observed choices, one would then notice a shock to hours of work. The findings in Geyer and Korfhage (2015) indicate that individuals who provide informal care would engage in more market work in the absence of care recipients' need.

In most models wage shocks are the sole drivers of earnings risk. Wage shocks can be recovered using moments of residuals of a wage equation alone. This does not hold for hours shocks. In our setting hours residuals contain hours shocks *in addition* to labor supply reactions to wage shocks. Net of Frisch labor supply adjustments, these reactions are governed by a single transmission parameter, which measures the impact of permanent income shocks on the marginal utility of wealth. Thus, without identifying this parameter separating the two shock types is impossible. The identification of the parameter serves an additional purpose. It enables us to calculate the Marshallian labor supply elasticity eschewing consumption or asset data, the reliability of which has been hotly debated (Attanasio and Pistaferri 2016). Thus, the second main contribution of our paper is the estimation of the transmission parameter without using consumption data, offering a new method to estimate the Marshallian elasticity. Finally, the parameter can be linked to consumption insurance. The larger this parameter, the larger is the impact of permanent shocks on

consumption, and the lower is the degree of insurance against risk. We show that the comovement of consumption and earnings implied by our estimate is in line with estimates by Blundell et al. (2008).

We apply our framework to observations on married men in the US from the Panel Study of Income Dynamics (PSID). Our samples starts in 1970 and ends in 1997, when the survey frequency turned biennial. We focus on this group because the extensive labor supply margin plays a small role in their labor supply behavior and in order to allow comparisons to the previous literature.

Our labor supply estimation yields a Frisch labor supply elasticity of 0.36 and our estimate of the average Marshallian elasticity is -0.08, which is close to recent estimates in Blundell et al. (2016) and Heathcote et al. (2014). We find that the standard deviation of permanent wage shocks is larger than the standard deviation of transitory shocks. The same holds for hours shocks, where the standard deviation of permanent shocks is about twice as large as that of transitory shocks. For most samples, the standard deviation of permanent hours shocks is larger than that of permanent wage shocks.

However, the respective impact on earnings risk cannot directly be inferred from this evidence, as the reaction to shocks depends on the degree of insurance. The main exercise with the key components of earnings risk in hand is a variance decomposition. We shut down each of the stochastic components except for one in order to quantify their respective contributions to overall earnings risk. At the sample mean of the transmission parameter, permanent wage shocks explain about 18 percent of earnings growth risk, while permanent hours shocks explain 13 percent. Transitory wage shocks dominate transitory hours shocks. While transitory shocks are responsible for the lion's share of earnings growth risk, only permanent shocks have a substantial impact on life-time earnings. At the mean of the transmission parameter, a positive permanent hours shock of one standard deviation at age 30 increases life-time earnings by 124 000 Dollar compared to 150 000 Dollar for a permanent wage shock of one standard deviation. Thus, both types of shocks play an important role for life-time earnings.

We also consider a set of alternative models that resemble those applied in the literature. Crucially, a model abandoning hours shocks fits the data worse and leads to a substantial overestimation—in absolute terms—of the Marshallian elasticity. Finally, we show how our estimate of the transmission parameter can be used to calculate the pass-through of permanent wage shocks to consumption. Calibrating the parameter of relative risk aversion to two, we find that these pass-through parameters for different samples are roughly in line with those estimated in Blundell et al. (2008), who use consumption and earnings data to estimate this parameter. For the full sample this calculation implies that—on average—a permanent increase in wages by one percent leads to an increase in consumption by .76 percent.

Our paper is related to studies that decompose total income risk into persistent and transitory components, which derive from ideas by Friedman (1957) and Hall (1978) (see MaCurdy 1982; Abowd and Card 1989; Meghir and Pistaferri 2004; Guvenen 2007; Blundell et al. 2008; Guvenen

2009; Hryshko 2012; Heathcote et al. 2014; Blundell et al. 2016). Abowd and Card (1989) were pioneers in analyzing the covariance structure of earnings and hours of work. They find that most of the idiosyncratic covariation of earnings and hours of work occurs at fixed wage rates.

In contrast, more recent papers have focused on insurance mechanisms rather than shock sources and restrict the source of risk to wage shocks. In a rich model of family labor supply and consumption, Blundell et al. (2016) estimate the Marshallian and Frisch consumption and labor supply elasticities using hours, income, asset, and consumption data. Similar to them, we allow for partial insurance of permanent wage shocks, but we depart from their approach by specifying an autoregressive process for partially insured shocks to the disutility of work and using hours and earnings data alone.

With a similar focus, Heathcote et al. (2014) analyze the transmission of wage shocks to hours in a setting where shocks are either fully insurable or not insurable at all (island framework). They derive second hours-wage moments from which they identify variances of shocks, the Frisch elasticity of labor supply, and the coefficient of relative risk aversion. Our study differs in two important aspects: First, we assume that shocks are partially insurable as indicated by a consumption insurance parameter similar to Blundell et al. (2008, 2013, 2016). This parameter may differ between individuals. Second, we introduce hours shocks and estimate their variance. While Heathcote et al. (2014) allow for initial heterogeneity between agents in the disutility of work, they hold this parameter constant over the life-cycle.

There are papers that do focus on shock sources more explicitly and for this purpose employ dynamic programming techniques. Low et al. (2010) quantify the contributions of productivity shocks, job losses, and job offers to overall earnings risk. They find that wage risk is much more important than job destruction risk due to the transitory nature of the latter. They model labor supply as a discrete decision with fixed hours of work and the possibility of job loss, while we focus on the intensive margin of work hours and allow for hours adjustment and permanent and transitory shocks to hours.[1] Similarly, Kaplan (2012) models consumption and hours of work and allows for involuntary unemployment shocks. These shocks along with nonseparable hours preferences on the extensive and intensive margin aid in the modeling of the declining inequality in hours worked over the first half of the life-cycle. Altonji et al. (2013) quantify the earnings variance contributions of i.i.d. wage and hours shocks in addition to employment and job changes, but they do not allow for individual-specific permanent hours shocks. Moreover, they do not work with a structural model, but rather approximate economic decisions of agents in their account of the dynamics of earnings and wage profiles. We are the first to model individual-specific hours shocks as a combination of permanent and transitory shocks.

The next section outlines the life-cycle model of labor supply and consumption, section 3 describes how the magnitudes of shock variances and labor supply elasticities are estimated. In

[1]For our PSID sample the variable for annual weeks worked shows little variation and thus offers scant evidence for the importance of intrayear unemployment spells.

section 4 we present results for the parameters of wage and hours processes and the Frisch and Marshallian labor supply elasticities. Then we offer decompositions of residual earnings variance and risk, which quantify the importance of wage and hours shocks. Further we calculate the influence of the two shock types on life-time earnings. In section 5 we give a characterization of permanent hours shocks, show results when varying the modeling assumptions, discuss the model fit and benchmark our results by relating them to consumption insurance. Section 6 concludes.

2 The Model

Individuals maximize the discounted sum of utilities over their lifetime, which runs from t_0 to T:[2]

$$\max_{c_t,h_t} E_{t_0}\left[\sum_{t=t_0}^{T}\rho^{t-t_0}\mathrm{v}(c_t,h_t,b_t)\right], \tag{1}$$

where c_t is annual consumption and h_t denotes annual hours of work, while b_t contains taste shifters. ρ denotes a discount factor and $\mathrm{v}(\cdot)$ an in-period utility function.

The budget constraint is

$$\frac{a_{t+1}}{(1+r_t)} = (a_t + w_t h_t + N_t - c_t), \tag{2}$$

where a_t represents assets, r_t the real interest rate, and N_t non-labor income.

Instantaneous utility takes the additively-separable constant relative risk aversion (CRRA) form

$$\mathrm{v}_t = \frac{c_t^{1-\vartheta}}{1-\vartheta} - b_t\frac{h_t^{1+\gamma}}{1+\gamma}, \qquad \vartheta \geq 0, \gamma \geq 0. \tag{3}$$

We specify $b_t = \exp(\varsigma\Xi_t - \upsilon_t)$. Ξ_t is a set of personal characteristics. υ_t is an idiosyncratic disturbance with mean zero, where the innovations to this term are the hours shocks. They capture unexpected changes in disutility of labor supply, e.g., childcare or spousal needs, sickness, and other unexpected changes in home production. Thus b_t captures taste-shifters in a very broad sense. The set-up of the model is standard, but our approach differs from the previous literature by introducing dynamics to the taste shifter.

Wage and hours shock processes — Denote by Δ the first difference operator. Wage growth is determined by human capital related variables X_t, which contains $\Delta\Xi$, and an idiosyncratic error ω:

$$\Delta \ln w_t = \alpha X_t + \Delta\omega_t \tag{4}$$

[2]We omit individual-specific subscripts.

Idiosyncratic hours (υ_t) and wage (ω_t) components consist of permanent and transitory components, p_t and τ_t, that follow a random walk and an MA(1)-process respectively. For $x \in \{\upsilon, \omega\}$ these are given by:

$$
\begin{aligned}
x_t &= p_t^x + \tau_t^x \\
p_t^x &= p_{t-1}^x + \zeta_t^x \\
\tau_t^x &= \theta_x \epsilon_{t-1}^x + \epsilon_t^x \\
E\left[\zeta_t^x \zeta_{t-l}^x\right] &= 0, \quad E\left[\epsilon_t^x \epsilon_{t-l}^x\right] = 0 \quad \forall l \in \mathbb{Z}_{\neq 0}
\end{aligned}
$$

Permanent and transitory shocks have mean zero and variances $\sigma^2_{\zeta,x}$ and $\sigma^2_{\epsilon,x}$, respectively. Permanent and transitory hours and wage shocks are uncorrelated. Note that hours and wage shocks contain only individual-specific shocks as X_t contains year dummies.

Labor supply — An approximation of the first order condition with respect to consumption yields the intertemporal labor supply equation (see MaCurdy 1981, Altonji 1986, and Appendix A):

$$
\Delta \ln h_t = \frac{1}{\gamma}\left[-\ln(1 + r_{t-1}) - \ln \rho + \Delta \ln w_t - \varsigma \Delta \Xi_t + \eta_t + \Delta \upsilon_t\right], \tag{5}
$$

where $\frac{1}{\gamma}$ is the Frisch elasticity of labor supply and η_t is a function of the expectation error in the marginal utility of wealth.[3] γ is identified by estimating equation (5) using instrumental variables for $\Delta \ln w_t$. The next objective is separating υ_t and η_t as the latter contains adjustments from wage shocks. To this end we first make explicit how wage and hours shocks transmit into changes in permanent income, which, in turn, result in changes in the marginal utility of wealth.

Denote by $\widehat{\Delta x}$ idiosyncratic changes in x. Then $\widehat{\Delta lny_t}$ are changes in log earnings that result from wage and hours shocks. It is useful to group these into transitory and permanent changes, distinguished by the superscripts *per* and *tra* respectively:

$$
\widehat{\Delta lny_t} = \widehat{\Delta lnw_t^{per}} + \widehat{\Delta lnw_t^{tra}} + \widehat{\Delta lnh_t^{per}} + \widehat{\Delta lnh_t^{tra}} \tag{6}
$$

The expressions for transitory and permanent wage changes in terms of shocks are obtained directly from the wage process:

$$
\widehat{\Delta lnw_t^{tra}} = \epsilon_t^\omega + (\theta_\omega - 1)\epsilon_{t-1}^\omega - \theta_\omega \epsilon_{t-2}^\omega \tag{7}
$$

$$
\widehat{\Delta lnw_t^{per}} = \zeta_t^\omega. \tag{8}
$$

[3] $\eta_t = \frac{\varepsilon_{\lambda_t}}{\lambda_t} - O\left(-1/2(\varepsilon_{\lambda_t}/\lambda_t)^2\right)$, i.e., it contains the expectation error of marginal utility of wealth and the approximation error.

Note that in the case of transitory wage changes, everything apart from ϵ_t^{ω} is known to the agent at $t-1$. In contrast, the idiosyncratic wage change due to permanent shocks is entirely surprising. Write idiosyncratic hours changes as

$$\widehat{\Delta \ln h_t} = \frac{1}{\gamma}\left[\widehat{\Delta \ln w_t} + \eta_t + \Delta \upsilon_t\right]. \tag{9}$$

We make the simplifying assumption that transitory shocks do not impact η_t.[4] Thus, the expressions for transitory hours changes in terms of shocks follow immediately from the stochastic processes of transitory shock components and the Frisch labor supply equation (9):

$$\widehat{\Delta lnh_t^{tra}} = \frac{1}{\gamma}\left(\epsilon_t^{\upsilon} + (\theta_{\upsilon}-1)\epsilon_{t-1}^{\upsilon} - \theta_{\upsilon}\epsilon_{t-2}^{\upsilon} + \epsilon_t^{\omega} + (\theta_{\omega}-1)\epsilon_{t-1}^{\omega} - \theta_{\omega}\epsilon_{t-2}^{\omega}\right). \tag{10}$$

In our model the expectation error η_t is a linear function of unexpected permanent changes to income. This is in line with models that approximate the life-time budget constraint like Blundell et al. (2016). The expression is

$$\eta_t = -\phi_t^{\lambda}\left(\widehat{\Delta lnw_t^{per}} + \widehat{\Delta lnh_t^{per}}\right), \quad \ln \phi_t^{\lambda} \sim N\left(\mu_{\phi}, \sigma_{\phi}^2\right). \tag{11}$$

Note that the first term in parentheses, the wage change, is exogenous, while the change in hours worked is an endogenous choice. The parameter ϕ_t^{λ} measures how shocks to income transmit to η_t. It is a measure of consumption insurance; perfectly insured individuals do not adjust their consumption as a response to a permanent shock and thus their marginal utility of consumption is unchanged. For instance, for individuals who have accumulated substantial assets, remaining life-time earnings only play a relatively small role in total life-time income. These individuals do not adjust their consumption by much in response to a wage shock. Blundell et al. (2016) study in detail what governs the transmission of shocks to consumption and hours worked. The parameter is individual-specific since it depends—among other things—on the amount of assets currently held in relation to the total stock of human wealth (see Blundell et al. 2016, p. 396, for a specification of the related consumption-insurance parameter). In the case of no insurance, a one percentage change in income leads to a one percentage change in consumption and $\phi_t^{\lambda} = \vartheta$.[5] In the case of full consumption insurance, $\phi_t^{\lambda} = 0$ and income changes do not translate into changes in consumption at all. It is reasonable to expect that there is at least some degree of insurance, which implies that the estimate of $E[\phi_t^{\lambda}]$ is a lower bound for the average degree of relative risk aversion.

[4] A long time horizon implies that transitory shocks have a negligible impact on the marginal utility of wealth. Blundell et al. (2008) show that this result holds empirically for PSID data.

[5] This can be seen by taking logs of the first derivative of equation (3) with respect to c_t. See also equation (13) below. ϕ_t^{λ} might exceed ϑ if shocks to life-time income not captured by the model are positively correlated with permanent hours and wage shocks.

Positive income shocks lead to a decrease in the marginal utility of wealth, therefore ϕ_t^λ is positive and should follow a distribution with no support on negative values. Hence, we estimate the model under the assumption that ϕ_t^λ is lognormally distributed. An equivalent transmission parameter for income shocks to the marginal utility of wealth is estimated in Alan et al. (2018). Alan et al. (2018) also assume a lognormal distribution for this parameter.

Plugging equation (8) into (11) and subsequently (11) into (9) and solving for $\widehat{\Delta lnh_t}$ yields the expression for idiosyncratic permanent changes in hours of work:

$$\widehat{\Delta lnh_t^{per}} = \frac{1-\phi_t^\lambda}{\gamma+\phi_t^\lambda}\zeta_t^\omega + \frac{1}{\gamma+\phi_t^\lambda}\zeta_t^\upsilon \tag{12}$$

The term $\kappa = \frac{1-\phi_t^\lambda}{\gamma+\phi_t^\lambda}$ gives the uncompensated reaction to a permanent wage change, i.e. the Marshallian labor supply elasticity. If $\phi_t^\lambda = 0$, the case of perfect insurance, the Marshallian collapses to the Frisch elasticity, the reaction to a transitory shock. The transmission coefficient for a permanent hours shock, $\frac{1}{\gamma+\phi_t^\lambda}$, has the same property.[6] The higher ϕ_t^λ, the more are hours shocks cushioned. Regarding wage shocks, an increase in ϕ_t^λ leads to a decrease in $\frac{1-\phi_t^\lambda}{\gamma+\phi_t^\lambda}$. Thus, for positive Marshallian elasticities, ϕ_t^λ cushions the labor supply reaction to wage shocks, while it amplifies the negative labor supply reaction for negative Marshallians.

With equations (6)-(12) we have now fully specified how wage and hours shocks affect the labor supply equation and in turn income. Finally, we can derive how they affect consumption.

Consumption — The equation for consumption growth can be obtained analogously to equation (5) (see, e.g., Altonji 1986):

$$\Delta \ln c_t = \frac{1}{\vartheta}\left[\ln(1+r_{t-1}) + \ln\rho - \eta_t\right] \tag{13}$$

Thus income shocks are directly related to consumption growth by $-\eta_t/\vartheta$. The direct estimation of equation (13) using consumption data is beyond the scope of this study. Nonetheless, we benchmark our results by calculating the reaction of consumption to wage shocks by calibrating ϑ.

Figures 1 and 2 show how each type of permanent shock propagates through the various quantities of interest. The major distinction for the two types is that wage shocks do not only have a direct effect on income, but also affect the choice of hours through the Marshallian elasticity.

[6] A further property of the hours shock transmission coefficient is that it equals the Hicksian elasticity for a permanent *wage* shock. This peculiarity arises from the way b affects marginal trade-offs: in a static version of the model with no unearned income the marginal optimality condition (MRS = price ratio) is given by $b\frac{h^\gamma}{c^{-\vartheta}} = w$. In the static model, the consumer maximizes utility, $V = \frac{c^{1-\vartheta}}{1-\vartheta} - b\frac{h^{1+\gamma}}{1+\gamma}$, subject to the budget constraint, $c = wh$. When we hold the level of consumption constant, a change in w and a change in b cause the same type of adjustment in h, although differently signed. However, when we let consumption adjust and derive the Marshallian demand for h, we find that $\ln h = \frac{1}{\gamma+\vartheta}\left[(1-\vartheta)\ln w - \ln b\right]$. Through the effect on the budget constraint, a change in w causes an income effect of size $\frac{-\vartheta}{\gamma+\vartheta}$ in elasticity form. b does not affect the budget constraint and therefore does not have the same effect.

Figure 1: Transmission of Permanent Wage Shock

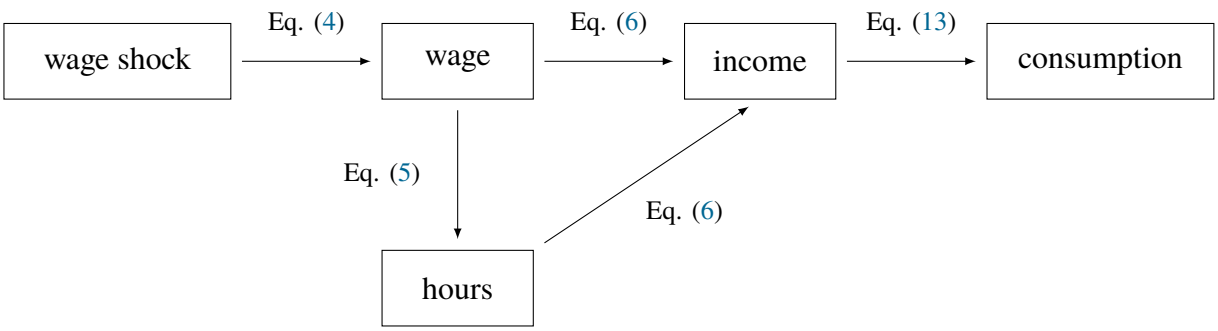

Note: Labels of arrows indicate corresponding equations.

Figure 2: Transmission of Permanent Hours Shock

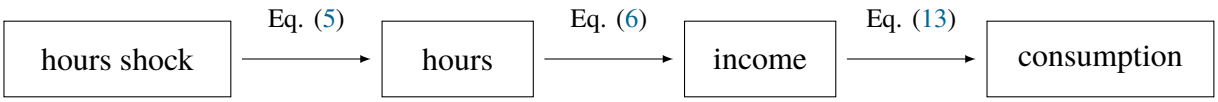

Note: Labels of arrows indicate corresponding equations.

3 Recovering Labor Supply Elasticities, Wage Shocks, and Hours Shocks

In this section we detail how the labor supply elasticities as well as the standard deviations of permanent and transitory parts of idiosyncratic wage components, ω_t, and hours components, υ_t, are recovered in the estimation. We proceed through three stages. First, we use OLS to obtain residuals from the wage equation (4) and IV to obtain residuals from the hours equation (5) as well as an estimate for the Frisch labor supply elasticity (Step 1). Second, we estimate the variances of transitory and permanent shocks to wages by fitting three theoretical autocovariance moments of the wage residuals to the data (Step 2). Third, we estimate hours shock variances by fitting the corresponding autocovariance moments for the hours residuals as well as the covariance of hours and wage residuals to the data (Step 3). Step 3 builds on steps 1 and 2 as it uses estimates for the wage process parameters and the Frisch elasticity.

The data — We use annual data from the PSID for the survey years 1970 to 1997, which gives 27 years usable for first-differenced estimations. After this point in time the PSID is biennial. In total we have 46 340 observations across individuals and years. Annual hours of work and earnings refer to the previous calendar year. Hours in the PSID are calculated as weeks worked times usual hours of work per week. Figure 4 in the appendix shows the kernel densities of these two variables. Weeks have a very peaked distribution, while weekly hours spread more. Thus, we can conclude that variations in annual hours of work are not substantially driven by short spells of unemployment.

Earnings consist of wages and salaries from all jobs and include tips, bonuses, and overtime. We calculate the hourly wage by dividing earnings through hours of work. As hours and earnings are measured with error, a negative correlation between measured hours and wages is induced, which we correct for as described in the next paragraph. Our sample consists of working, married males aged 28 to 60, who are the primary earners of their respective households. Table 1 shows summary statistics of the main sample. Monetary variables are adjusted to 2005 real values using the CPI-U.

Table 1: Descriptives

	mean	s.d.
Age	41.35	8.66
Annual hours of work	2220.28	530.11
Hourly wage	26.86	22.83
Number of kids in household	1.64	1.39
N	46340	

Note: Own calculation based on the PSID. Monetary values inflated to 2005 real dollars.

Measurement errors — We state the variance-covariance moments with measurement error. Measurement error is modeled as having no intertemporal and cross-sectional correlation, but we do allow for correlation between the types of measurement error. Denote by

$$\ln \tilde{x}_t = \ln x_t + me_{x,t} \tag{14}$$

the observed value for the log of variable x, where $me_{x,t}$ is the mean zero measurement error with variance $\sigma^2_{me,x}$. The variances encountered in the moment conditions are $\sigma^2_{me,h}, \sigma^2_{me,w}$, and $\sigma^2_{me,h,w}$, which are the variances of measurement errors in log hours, log wages and their covariance respectively.

Following Meghir and Pistaferri (2004) and Blundell et al. (2016) we use estimates from the validation study by Bound et al. (1994) for the signal-to-noise ratios of wages, hours, and earnings. As in Blundell et al. (2016), we assume that the variance of the measurement error of hours is $\sigma^2_{me,h} = 0.23var(\ln h)$, the variance of the measurement error of wages is $\sigma^2_{me,w} = 0.13var(\ln w)$, and the variance of the measurement error of earnings is $\sigma^2_{me,y} = 0.04var(\ln y)$, where $var(\ln h)$, $var(\ln w)$, and $var(\ln y)$ denote the variances of the levels of log wages, log hours, and log earnings. The covariance of the measurement errors of log wages and hours is given by $\sigma^2_{me,h,w} = (\sigma^2_{me,y} - \sigma^2_{me,w} - \sigma^2_{me,h})/2$. We correct the theoretical moments using these estimates for the parts that are attributable to error.

Step 1: Frisch elasticity, hours residuals, and wage residuals — The augmented empirical labor supply equation containing measurement errors is

$$\Delta \ln \tilde{h}_t = \frac{1}{\gamma}\left[-\ln(1+r_{t-1}) - \ln\rho + \Delta \ln \tilde{w}_t - \varsigma \Delta \Xi_t + \eta_t + \Delta \upsilon_t\right] - \frac{1}{\gamma}\Delta me_{w,t} + \Delta me_{h,t}. \tag{15}$$

The error term of equation (15) is correlated with differenced log wages because wage shocks impact the marginal utility of wealth and because of measurement error. To obtain the Frisch elasticity $1/\gamma$ from equation (15) we use human capital related instrumental variables for $\Delta \ln \tilde{w}_t$ following MaCurdy (1981). These instruments predict the expected part of wage growth. Thus, the instruments are uncorrelated with innovations in the marginal utility of wealth and measurement error. Hours residuals $(\eta + \Delta\tilde{\upsilon}_t)/\gamma = (\eta + \Delta\upsilon_t - \Delta me_{w,t})/\gamma + \Delta me_{h,t}$ are obtained by running IV on differenced log hours using differenced year, child, disability and state dummies as covariates. The instruments for the differenced log wage are interactions of age and years of education, i.e., age, education, education^2, age $\times$ education, age $\times$ education^2, age^2 $\times$ education and age^2 $\times$ education^2. Wage residuals $\Delta\tilde{\omega}_t = \Delta\omega_t + \Delta me_{w,t}$ are obtained by estimating equation (4) augmented by an error term, i.e. regressing differenced log wages on the same exogenous regressors as in the hours equation as well as the excluded instruments.

The results of Step 1 are an estimate for $1/\gamma$, a matrix of hours residuals, and a matrix of wage residuals. These are used in Steps 2 and 3.

Step 2: Wage shocks — After recovering $\Delta\tilde{\omega}_t$, all parameters of the autoregressive process, $(\theta, \sigma^2_{\epsilon,\omega}, \sigma^2_{\zeta,\omega})$, are identified through combinations of the autocovariance moments. Label the k-th autocovariance moment by $\Lambda_{\tilde{\omega},k}$:

$$\Lambda_{\tilde{\omega},0} = E\left[(\Delta\tilde{\omega}_t)^2\right] = 2\left(1 - \theta_\omega + \theta_\omega^2\right)\sigma^2_{\epsilon,\omega} + \sigma^2_{\zeta,\omega} + 2\sigma^2_{me,w} \tag{16}$$

$$\Lambda_{\tilde{\omega},1} = E\left[\Delta\tilde{\omega}_t \Delta\tilde{\omega}_{t-1}\right] = -(\theta_\omega - 1)^2\sigma^2_{\epsilon,\omega} - \sigma^2_{me,w} \tag{17}$$

$$\Lambda_{\tilde{\omega},2} = E\left[\Delta\tilde{\omega}_t \Delta\tilde{\omega}_{t-2}\right] = -\theta_\omega \sigma^2_{\epsilon,\omega} \tag{18}$$

Net of $\sigma^2_{me,w}$, dividing $\Lambda_{\tilde{\omega},2}$ by $\Lambda_{\tilde{\omega},1}$ identifies the parameter θ_ω. Successively, the variance of the transitory shock is identified from $\Lambda_{\tilde{\omega},1}$ and the variance of the permanent shock from $\Lambda_{\tilde{\omega},0}$ (see Hryshko 2012).

Step 3: Hours shocks — The residual obtained from estimating the labor supply equation contains both hours shocks υ_t and the expectation error, η_t. The variance of the residual of the labor supply equation contains both the mean and the variance of $\frac{\phi_t^\lambda}{\gamma+\phi_t^\lambda}$ and the variance of the permanent hours shocks, which causes an identification problem. The procedure for wage moments does not carry over. We use the contemporaneous covariance of hours and wage residuals to identify the mean of $1-\frac{\gamma}{\gamma+\phi_t^\lambda}$, which is equivalent to $\frac{\phi_t^\lambda}{\gamma+\phi_t^\lambda}$. To arrive at the theoretical variance moment, substitute equations (8) and (12) into (11) and subsequently (11) into (15) to find the following expression for the hours residual

$$\begin{aligned}\frac{\eta+\Delta\tilde{\upsilon}_t}{\gamma} =& \frac{1}{\gamma}\Bigg[-\left(1-\gamma\frac{1}{\gamma+\phi_t^\lambda}\right)\zeta_t^\upsilon-(1+\gamma)\left(1-\gamma\frac{1}{\gamma+\phi_t^\lambda}\right)\zeta_t^\omega \\ &+\zeta_t^\upsilon+\epsilon_t^\upsilon+(\theta\upsilon-1)\epsilon_{t-1}^\upsilon-\theta\upsilon\epsilon_{t-2}^\upsilon\Bigg] \\ &-\frac{1}{\gamma}\Delta me_{w,t}+\Delta me_{h,t},\end{aligned} \tag{19}$$

where the first line on the right hand side equals η_t/γ, i.e. the labor supply reaction due to the impact of shocks on the marginal utility of wealth. The two terms inside the brackets give a wealth effect, i.e. the Marshallian net of the Frisch reaction, of a permanent shock to the disutility of work or to the hourly wage respectively. Note that in the case of full insurance ($\phi_t^\lambda=0$) these terms equal zero. The second line contains immediate, i.e. Frisch, reactions to shocks in the disutility of work. The third line contains the terms due to measurement error. The variance can be written as

$$\begin{aligned}\Lambda_{\tilde{\upsilon},0}=E\left[\left(\frac{\eta_t+\Delta\tilde{\upsilon}_t}{\gamma}\right)^2\right] =& \frac{1}{\gamma^2}\left((1+\gamma)^2(1-2\gamma\mathcal{M}_1+\gamma^2\mathcal{M}_2)\sigma_{\zeta,\omega}^2\right)+\mathcal{M}_2\sigma_{\zeta,\upsilon}^2 \\ &+\frac{1}{\gamma^2}\left(\sigma_{\zeta,\upsilon}^2+2\left(\theta_\upsilon^2-\theta_\upsilon+1\right)\sigma_{\epsilon,\upsilon}^2\right) \\ &+2\sigma_{me,h}^2+\frac{2\sigma_{me,w}^2}{\gamma^2}-\frac{4\sigma_{me,h,w}^2}{\gamma},\end{aligned} \tag{20}$$

where $\mathcal{M}_1$ and $\mathcal{M}_2$ denote the first and second non-central moments of $1/(\gamma+\phi_t^\lambda)$, the random component in $1-\frac{\gamma}{\gamma+\phi_t^\lambda}$. As no analytical expression exists for these moments, we find them numerically as described in Appendix B. The interpretation is analogous to (19): the first line captures the part of the variance that is due to marginal utility of wealth effects, while the second captures the part of the variance due to direct labor supply reactions to shocks, and the third line is due to measurement error.

The autocovariance moments of the hours residual $\Lambda_{\tilde{\upsilon},1}$ and $\Lambda_{\tilde{\upsilon},2}$ are analogous to their wage process counterparts:

$$\Lambda_{\tilde{\upsilon},1} = E\left[\frac{(\eta_t + \Delta\tilde{\upsilon}_t)(\eta_{t-1} + \Delta\tilde{\upsilon}_{t-1})}{\gamma^2}\right] = -\frac{(\theta_\upsilon - 1)^2\sigma^2_{\epsilon,\upsilon}}{\gamma^2} - \sigma^2_{me,h} - \frac{\sigma^2_{me,w}}{\gamma^2} + \frac{2\sigma^2_{me,h,w}}{\gamma} \tag{21}$$

$$\Lambda_{\tilde{\upsilon},2} = E\left[\frac{(\eta_t + \Delta\tilde{\upsilon}_t)(\eta_{t-2} + \Delta\tilde{\upsilon}_{t-2})}{\gamma^2}\right] = -\frac{\theta_\upsilon\sigma^2_{\epsilon,\upsilon}}{\gamma^2} \tag{22}$$

To estimate the variance of permanent hours shocks, we need to identify $\mathcal{M}_1$ using the contemporaneous covariance of hours and wage residuals:

$$\Lambda_{\tilde{\omega},\tilde{\upsilon},0} = E\left[\frac{(\eta_t + \Delta\tilde{\upsilon}_t)\Delta\tilde{\omega}_t}{\gamma}\right] = -\frac{(\gamma+1)(1-\gamma\mathcal{M}_1)\sigma^2_{\zeta,\omega}}{\gamma} - \frac{2\sigma^2_{me,w}}{\gamma} + 2\sigma^2_{me,h,w}. \tag{23}$$

This covariance is larger in absolute value the smaller γ and the smaller $\mathcal{M}_1$, which is due to a larger $E[\phi^\lambda_t]$. When γ goes to infinity, the effect of permanent wage shocks on income is only mechanical and not through labor supply reactions. It is important to stress that, given estimates for γ and $\sigma^2_{\zeta,\omega}$ only equation (23) is needed for the estimation of $E[\phi^\lambda_t]$. The estimation of the parameter does not hinge on the identification of hours shocks.

$\mathcal{M}_1$ and $\mathcal{M}_2$ contain both μ_ϕ and σ_ϕ. Theoretically, σ_ϕ is identified through the cokurtosis moments of the wage and hours residuals. However, cokurtosis moments are very noisy, hence σ_ϕ can only be estimated to a reasonable degree of reliability when using several million observations.[7] Therefore, we calibrate σ_ϕ to 1.023 based on results in Alan et al. (2018). In that paper, the parameter follows a log-normal distribution and is identified through a separate moment condition based on consumption data. Using this calibration, once $\mathcal{M}_1$ is estimated, the mean of ϕ^λ_t, $E[\phi^\lambda_t] = e^{\mu_\phi + \frac{\sigma^2_\phi}{2}}$, can be recovered. In section 5 we show the robustness of our results to alternative values of this parameter. Halving the standard deviation σ_ϕ increases the variance of permanent hours shocks only slightly and the estimate for $E[\phi^\lambda_t]$ is qualitatively the same.

To sum up, the parameters $\theta_\upsilon, \sigma_{\epsilon,\upsilon}, \sigma_{\zeta,\upsilon}$, and $E[\phi^\lambda_t]$ are identified through the moments $\Lambda_{\tilde{\upsilon},0}$, $\Lambda_{\tilde{\upsilon},1}$, $\Lambda_{\tilde{\upsilon},2}$, and $\Lambda_{\tilde{\omega},\tilde{\upsilon},0}$.

Marshallian elasticity — The term multiplied with $\sigma^2_{\zeta,\omega}$ in equation (23) can be rewritten as $E\left[\frac{1-\phi^\lambda_t}{\gamma+\phi^\lambda_t}\right] - \frac{1}{\gamma}$, the average Marshallian minus the Frisch elasticity of labor supply. Thus, the

[7] Simulations evidencing this are available upon request from the authors.

Marshallian can directly be calculated from the parameter estimates. The Marshallian elasticity is the uncompensated reaction to a permanent wage shock.[8]

The Marshallian elasticity is the relevant concept for the evaluation of tax reforms, which are best described as unanticipated, permanent shifts in net-of-tax wages (Blundell and Macurdy 1999). Using similar considerations as in our study, the Marshallian elasticity has been estimated using the covariance of earnings and wages, the shares of individuals' human wealth over family human wealth, and the ratio of assets to total (human and non-human) wealth in Blundell et al. (2016, eq. A2.23). Heathcote et al. (2014) use the covariance of hours and consumption as well as of wages and consumption to estimate the Marshallian elasticity. In contrast, we rely only on hours and wage data.[9]

Estimation — We estimate the parameters of the autoregressive processes and the transition of wage shocks by fitting the theoretical moments $\{\Lambda_{\omega,k}, \Lambda_{\upsilon,k}, \Lambda_{\omega,\upsilon,k}\}$ to those of the data. The vector of parameters, denoted Θ, is estimated using the method of minimum distance and an identity matrix serves as the weighting matrix.[10] The distance function is given by

$$DF(\Theta) = [m(\Theta) - m^d]' I [m(\Theta) - m^d], \tag{24}$$

where $m(\Theta)$ indicates theoretical moments and m^d empirical moments. An outline of the entire estimation procedure is detailed in Hryshko (2012). Standard errors are obtained by block bootstrap with 200 replicates.

4 Main Results

Throughout this section we show results for the full sample as well as three sub-samples, which might differ with respect to their labor supply behavior and exposure to shocks: a sample excluding young workers under 40, one consisting of individuals with more than high school education and a sample of individuals without children younger than seven years in the household.

Standard deviations of wage shocks — Table 2 reports the standard deviations of permanent and transitory wage shocks as well as the parameter of transitory shock persistence. First, while the magnitude of the standard deviation of permanent shocks ($\sigma_{\zeta,\omega}$) is similar in the four samples, excluding young workers leads to a decline of this figure. This is in line with the finding of

[8] See Keane (2011, p.1008) for a discussion of why reactions to permanent shocks equal the Marshallian elasticity.

[9] Heathcote et al. (2014) also estimate a variant that does not rely on consumption data. Their approach differs because their specific island framework implies that the marginal utility of wealth is constant across individuals in the same age-year cell.

[10] Altonji and Segal (1996) show that the identity weighting matrix is preferable for the estimation of autocovariance structures using micro data.

slightly higher variances of permanent wage shocks at younger ages as in Blundell et al. (2016) and Meghir and Pistaferri (2004). Second, for all samples the standard deviation of transitory shocks ($\sigma_{\epsilon,\omega}$) is smaller than that of permanent shocks. Third, the highly educated face a substantially lower standard deviation of transitory shocks than the full sample. Lastly, for those without young children permanent and transitory shocks are slightly lower than for the full sample.

Table 2: Wage Variances

	I	II	III	IV
	Full sample	Age≥40	High educ	No children <7
θ_ω	0.2701	0.3450	0.2737	0.1832
	(0.0090)	(0.0241)	(0.0325)	(0.0212)
$\sigma_{\epsilon,\omega}$	0.1337	0.1382	0.0772	0.1166
	(0.0017)	(0.0030)	(0.0015)	(0.0025)
$\sigma_{\zeta,\omega}$	0.1770	0.1554	0.1765	0.1639
	(0.0009)	(0.0014)	(0.0007)	(0.0011)
N	46340	20607	19831	24547

Note: Own calculation based on the PSID. Bootstrapped standard errors in parentheses.

Standard deviations of hours shocks — The first three rows in Table 3 show the parameters of the process of shocks to the disutility of work. For ease of interpretation, the parameters are reported as they enter the hours equation (5), i.e. multiplied with the Frisch elasticity $1/\gamma$. The estimates for the hours process are of comparable size to those of the wage process. The standard deviation of permanent hours shocks drops when we consider only the highly educated. Otherwise, permanent shocks to the disutility of work are of a fairly consistent size across the samples. For all three subsamples the standard deviation of transitory shocks is lower than in the full sample, but persistence of the shocks also differs across samples. The magnitude of the standard deviation of permanent hours shocks is a first indicator that these shocks play a significant role for overall earnings risk. However, as described in section 2, the effect of innovations in the marginal disutility of work on earnings depends on the degree of consumption insurance.

Table 3: Hours Variances and Labor Supply Elasticities

	I	II	III	IV
	Full sample	Age>=40	High educ	No children <7
θ_υ/γ	0.1515	0.4013	0.1140	0.2463
	(0.0039)	(0.0059)	(0.0065)	(0.0053)
$\sigma_{\epsilon,\upsilon}/\gamma$	0.1114	0.0730	0.0709	0.0790
	(0.0011)	(0.0012)	(0.0014)	(0.0012)
$\sigma_{\zeta,\upsilon}/\gamma$	0.1990	0.2102	0.1648	0.1914
	(0.0010)	(0.0327)	(0.0010)	(0.0058)
$1/\gamma$	0.3614	0.4020	0.2851	0.3148
	(0.0856)	(0.3778)	(0.0975)	(0.1080)
$E[\phi_t^\lambda]$	1.8918	1.4084	0.5668	0.9565
	(0.1117)	(4.0920)	(0.0436)	(1.2774)
$E[\kappa]$	-0.0767	-0.0023	0.1302	0.0631
	(0.0105)	(0.0254)	(0.0078)	(0.0164)

Note: Own calculation based on the PSID. Clustered standard errors for $1/\gamma$, bootstrapped standard errors for other coefficients in parentheses.

Frisch elasticity — Row 4 in Table 3 reports the estimates of the Frisch elasticity. In contrast to the most closely related papers (Blundell et al. 2016; Heathcote et al. 2014), we obtain the Frisch elasticity directly through IV estimation and not through covariance moments.[11] The estimated Frisch elasticity for the main sample is 0.36, which is in line with the literature (Keane 2011). The point estimate of the Frisch elasticity increases when excluding younger individuals. This result is expected as younger individuals could be less willing to reduce their hours of work in the case of a decrease in the hourly wage because the accumulation of human capital impacts their opportunity costs of time (Imai and Keane 2004). Similarly, human capital considerations are more important for the highly educated, where the Frisch elasticity is lower than that of the main sample. The estimate for those without young children is fairly close to that of the main sample, but slightly smaller.

Transmission parameter — Row 5 in Table 3 shows the estimated mean of the parameter that measures the transmission of shocks to the marginal utility of wealth, $E[\phi_t^\lambda]$. The smaller this parameter, the more individuals are insured against shocks. A value of zero indicates that permanent shocks do not impact the marginal utility of wealth at all. We expect households with larger accumulation of assets relative to human wealth to exhibit smaller values of $E[\phi_t^\lambda]$. The point estimate drops only slightly relative to the full sample, when excluding young workers, but is

[11] Table 10 in the appendix additionally displays the Kleibergen and Paap (2006) F statistic, indicating that only sample II might suffer from weak instrument bias and should therefore be interpreted with caution.

substantially smaller when focusing on those without young children and especially on the highly educated. These drops with respect to the main sample are expected.

Marshallian elasticity —Table 3 reports the average of the Marshallian elasticity defined in equation (12) as the reaction to a permanent wage shock.[12] The wealth effect outweighs the substitution effect, leading to a negative (but small) estimate for the main sample, in line with the recent literature.[13] The negative Marshallian implies that hours move in the opposite direction of wages and thus function as a consumption smoothing device. When excluding younger workers, the estimate edges even closer to zero, signifying no long-term adjustment in hours for older workers. The smaller the average transmission parameter, the closer is the average Marshallian to the Frisch elasticity because the shock has a smaller effect on the marginal utility of wealth. The smaller wealth effect for older workers is expected because individuals close to the end of their life-cycle experience the same change to their marginal utility of wealth from either a transitory or a permanent shock. In the sample without young children in the household the estimate is positive, making the substitution effect the dominant force as the average transmission parameter is relatively small for this sample. The highly educated show the highest positive Marshallian elasticity driven by their very small transmission parameter.

Importance of hours and wage shocks — Using our estimates for the variances of hours and wage shocks allows us to quantify their contribution to the cross-sectional variance of overall earnings growth. The stochastic component of earnings net of measurement error is given by the sum of hours and wage residuals plus the Frisch reactions to idiosyncratic wage changes, which we had removed from the hours residual by detrending with wages, see equation (15). The variance of stochastic earnings growth is thus given by

$$E\left[\left(\widehat{\Delta \ln y}\right)^2\right] = \frac{1}{\gamma^2}\Big[\gamma^2 \mathcal{M}_2 \left(\sigma^2_{\zeta,\upsilon} + (\gamma+1)^2\sigma^2_{\zeta,\omega}\right) + 2(\gamma+1)^2((\theta_\omega - 1)\theta_\omega + 1)\sigma^2_{\epsilon,\omega} + 2((\theta_\upsilon - 1)\theta_\upsilon + 1)\sigma^2_{\epsilon,\upsilon}\Big] \tag{25}$$

$\mathcal{M}_2$ is obtained numerically using the estimates of the underlying parameters. Note that $\mathcal{M}_2$ depends on the mean and the variance of the transmission parameter ϕ^λ_t, which are known to individuals. Additionally the realization of transitory components of wage and hours growth are partially known in advance, see equation (7). Therefore this overall variance is not a measure of

[12] We calculate κ as the numerical expectation $E\left[\frac{1-\phi^\lambda_t}{\gamma+\phi^\lambda_t}\right]$.

[13] Blundell et al. (2016) and Heathcote et al. (2014) find Marshallian elasticities for men of -0.08 and -0.16 respectively. The latter number is obtained by inserting the obtained parameter estimates in the formula for the labor supply reaction to an uninsurable shock. Altonji et al. (2013) report a coefficient that determines "the response to a relatively permanent wage change" of -0.08.

risk. The first row of Table 4 shows the cross-sectional variance of the stochastic component of earnings growth for the four samples. Rows 2-5 show the contributions of shock components, i.e., the variance of earnings growth when the variances of all other shock components are set to zero. Earnings growth variances for the highly educated and for the sample excluding households with young children are substantially lower than that for the main sample. For all samples, except that of the highly educated, transitory wage shocks play the largest role in explaining earnings growth variance. The main source of the variance in income growth for the highly educated are permanent wage shocks. Their contribution is roughly double that of permanent hours shocks. For all samples the share of the cross-sectional hours variance due to permanent hours shocks is at least half the share due to permanent wage shocks. For the sample excluding younger workers, the shares of earnings growth variance explained by permanent hours and wage shocks are of similar magnitude. For older workers the seniority dynamics of their wage become less relevant, which should drive down the contribution of permanent wage shocks to the variance of earnings. A similar dynamic does not necessarily follow for shocks to home production. The highly educated experience the bulk of their variation through permanent wage shocks, while all other sources of the variance lose relevance compared to the main sample. The earnings and wage profile of the highly educated is much steeper; human capital can still increase substantially through labor market experience (see Imai and Keane 2004). These potential increases in productivity during working life leave room for a greater variation across individuals and over time.

Table 4: Decomposition of Variance of Earnings Growth

	I	II	III	IV
	Full sample	Age>=40	High educ	No children <7
$V\left(\widehat{\Delta \ln y}\right)$	0.1464	0.1293	0.0857	0.1071
$\sigma_{\epsilon,\omega}$	0.0532	0.0581	0.0158	0.0400
$\sigma_{\zeta,\omega}$	0.0426	0.0326	0.0398	0.0319
$\sigma_{\epsilon,\upsilon}$	0.0216	0.0082	0.0090	0.0100
$\sigma_{\zeta,\upsilon}$	0.0290	0.0304	0.0210	0.0252

Note: Variance of $\widehat{\Delta \ln y}$ when all other shock variances are set to zero. First line: actual variance of $\widehat{\Delta \ln y}$ given by equation (25). Own calculation based on the PSID.

In the following we consider earnings risk. Note that permanent changes in income lead directly to changes in consumption, only mediated through the consumption insurance parameter. Thus earnings risk translates to consumption risk and has immediate welfare implications for risk averse individuals. When evaluating *risk* of idiosyncratic earnings growth instead of its cross-sectional *variance*, everything that is known to an agent at $t-1$ must be excluded from (25) and ϕ_t^λ must be treated as non-stochastic. Denote by I_{t-1} the agent's information set at $t-1$. At that point in time the agent knows ϕ_t^λ and the realization of shocks in $t-1$. Thus, $E\left[\widehat{\Delta \ln y_t}|I_{t-1}\right]$ includes

the transitory components from the previous two periods. The resulting equation for earnings risk conditional on the information set in $t-1$ is

$$E\left[\left(\widehat{\Delta \ln y_t} - E\left[\widehat{\Delta \ln y_t}|I_{t-1}\right]\right)^2 \middle| I_{t-1}\right] = \frac{\sigma^2_{\zeta,\upsilon} + (\gamma+1)^2\sigma^2_{\zeta,\omega}}{(\gamma+\phi^\lambda_t)^2} + \frac{1}{\gamma^2}\left(\sigma^2_{\epsilon,\upsilon} + (\gamma+1)^2\sigma^2_{\epsilon,\omega}\right) \quad (26)$$

In Table 5 ϕ^λ_t is set to the sample mean. Thus, the cross-sectional variance of unexpected earnings growth can be interpreted as earnings risk for a typical individual in each sample. A comparison of the first lines in Tables 4 and 5 shows that for the full sample earnings growth risk at the mean is about 55% of the cross-sectional idiosyncratic earnings growth variance; a substantial fraction. The degree to which the size of contributions of transitory shocks decreases depends on the parameter θ of the respective MA(1) process. The smaller θ, the larger is the share of risk in the total variance of idiosyncratic earnings growth. The importance of permanent shocks decreases relative to Table 4 because ϕ^λ_t is non-stochastic. The importance of permanent wage shocks decreases for all samples and to a large degree for the first two samples, where the Marshallian labor supply elasticity is negative at the mean of ϕ^λ_t. Similarly, the importance of permanent hours shocks for total earnings risk is much smaller for these two samples as ϕ^λ_t cushions the reaction to innovations in the marginal disutility of work. For all samples, permanent hours shocks explain at least 17 percent of earnings risk. Nonetheless, wage risk is more important in all samples, although for older individuals the magnitudes are very close.

Table 5: Decomposition of Earnings Risk at Mean

	I	II	III	IV
	Full sample	Age>=40	High educ	No children <7
$V(\Delta \ln y)$	0.08	0.0803	0.0731	0.0788
$\sigma_{\epsilon,\omega}$	0.0331	0.0375	0.0098	0.0235
$\sigma_{\zeta,\omega}$	0.0205	0.0194	0.0381	0.0275
$\sigma_{\epsilon,\upsilon}$	0.0124	0.0054	0.005	0.0062
$\sigma_{\zeta,\upsilon}$	0.014	0.018	0.0201	0.0217

Note: Earnings growth risk with ϕ^λ_t set to sample mean. First line: Total earnings risk at mean given by equation (26).
Own calculation based on the PSID.

While transitory shocks are an important driver of cross-sectional earnings growth variance, only permanent shocks have a large impact on the present value of life-time earnings. A back-of-the-envelope calculation[14] using the average coefficients of the main sample shows that for an

[14]The impact of a typical permanent wage shock at the mean of ϕ^λ_t is given by the geometric series $y\left(1+\left(1-\left(E[\phi^\lambda_t]\right)/\left(\gamma+E[\phi^\lambda_t]\right)\right)\sigma_{\zeta,\omega}\right)(1-1/r^{65-age})/(1-1/r)$ and the impact of a typical permanent hours shock is $y\left(1/(\gamma+E[\phi\lambda_t])\right)\sigma_{\zeta,\upsilon}(1-1/r^{65-age})/(1-1/r)$, where annual earnings y are set to the sample mean of 57 267 Dollar and the real interest rate r is 1.0448 based on World Bank figures for our period. This abstracts from deterministic earnings growth, i.e. it makes the simplifying assumption that earnings would remain constant without shocks.

individual aged 30 and retiring at 65, a typical positive permanent wage shock of one standard deviation increases the present value of life-time earnings by about 150 000 Dollar, while a typical positive permanent hours shock increases it by 124 000 Dollar.[15] Typical permanent wage and hours shocks at age 50 for the same individual increase life-time income by 92 000 and 76 000 Dollar respectively. Thus, both types of permanent shocks have a substantial impact on life-time earnings.

The impact of hours shocks depends largely on the degree of insurance. In the benchmark case of full insurance with $\phi_t^\lambda = 0$ individuals adjust their hours of work much more in response to a shock to the disutility of work. In this case, the impact of a typical permanent wage shock at age 30 is 252 000 Dollar because in this case the Frisch labor supply reaction amplifies the wage shock. The analogous impact of a typical hours shock is 208 000 Dollar. Clearly, the impact of a permanent shock on life-time income varies widely between individuals.

5 Discussion

Characterizing hours shocks — In order to investigate and understand the sources of permanent hours shocks, we estimate their standard deviation in alternative samples. Column I in Table 6 reports the value for the full sample. Column II contains results for a sample of individuals in "blue collar" industries. Individuals in advanced technical sectors, like electrical and mechanical engineering or skilled service jobs like legal or medical services are excluded.[16] One could expect that the demand for these more regularized and routine jobs only allows for very limited variation in hours. However, this does not seem to be the case, as the estimate of the permanent shocks hardly changes. In column III we exclude the years 1981 and 1982, when a global recession hit the US. The estimate of the standard deviation of permanent hours shocks hardly changes, which shows that the results are not driven by the crisis. Finally, in column IV the sample is restricted to individuals who have stayed in their current job for at least 3 years consecutively.[17] Given that the estimate for the permanent hours shock variance is very close to that of the main sample, we can infer that permanent hours shocks do not reflect changes in occupation or job instability. In contrast, the variance of transitory hours shocks in the stayers sample drops sharply ($\sigma_{\epsilon,\upsilon}/\gamma = 0.0386$, s.e.: 0.0011), attesting to the transitory nature of the effect of job changes on hours. The upshot of all of these results is that permanent hours shocks play an important role throughout all samples and are not restricted to very specific adjustments or at-risk groups. The fact that hours

[15]Note that the ratio of the impacts of typical permanent hours and wage shocks on lifetime earnings equals the square root of the corresponding ratio of contributions to earnings risk reported in Table 5.

[16]In particular, individuals in education, sport, legal, health, and other services including service industries, mechanical and electrical engineering, financial institutions, and insurance as well as public administration, social security, private households, volunteering and churches are excluded.

[17]Note that this means that all moments rely only on information from the current job.

shocks do not seem to be driven by job changes or possibly unwanted changes in hours of work during crises suggests an interpretation of permanent hours shocks as shocks to home production.

Table 6: Permanent Hours Shock Variances in Alternative Samples and Models

	I	II	III	IV
	Main	Blue collar	Exclude years 81-82	Only stayers
$\sigma_{\zeta,\upsilon}/\gamma$	0.1990	0.2067	0.2066	0.1818
	(0.0010)	(0.0022)	(0.0025)	(0.0014)
N	46340	32313	40999	29211

Note: Own calculation based on the PSID. Bootstrapped standard errors in parentheses.

Hours shocks and transmission in alternative models — In Table 7 we report the parameters of the hours shock process and the transmission parameter as well as the implied Marshallian elasticity for the main sample under various restrictions or alternative assumptions. Further we display a measure of overall fit of these alternative models, namely the value of the distance function $DF(\Theta)$, to develop an idea about the value of the main model in describing the data. The estimates of the main model are repeated for comparison in column I. In column II, the variance of $\ln(\phi_t^\lambda)$ is calibrated to half the value of our main specification. All estimated coefficients except for the standard deviation of permanent hours shocks and the mean of the transmission parameter are unchanged. The standard deviation of the permanent hours shocks is slightly larger. The reason is that the variance of the transmission parameter interacts with the variance of hours shocks in explaining the variance of hours growth (equation (20)). The fit of the model changes just slightly, since the variance of permanent hours shocks can freely adjust. The exercise demonstrates that the results only depend to a small degree on this calibration. Columns III and IV illustrate the importance of allowing for hours shocks. In column III the variance of permanent hours shocks is set to zero. While the estimated variance of transitory hours shocks increases only slightly, the estimated mean of the transmission parameter increases to roughly 2.47. The fit of this model is substantially worse with an increase of the distance function by six orders of magnitude. The implied Marshallian elasticity doubles. In column IV both transitory and permanent hours shocks are restricted to zero. In this case the estimated average transmission parameter increases to roughly 20.8 and the implied Marshallian elasticity is about -0.7. These extreme estimates are caused by the fact that the transmission of wage shocks is now the only channel to explain hours variance. Naturally, the fit takes another hit from this restriction, although it is not as severe as the first jump. The order of magnitude of the distance function increases by one. That the final change in fit is relatively smaller than the one in model III further underlines the fact that permanent hours shocks are an important part of the picture in the attempt to explain the variance of the hours residual.

Table 7: AR Hours Estimation in Alternative Models

	I	II	III	IV
	Main Model	σ_ϕ halved	$\sigma_{\zeta,\upsilon} = 0$	$\sigma_{\zeta,\upsilon} = 0$ & $\sigma_{\epsilon,\upsilon} = 0$
θ_υ/γ	0.1515	0.1515	0.1454	
	(0.0039)	(0.0039)	(0.0013)	
$\sigma_{\epsilon,\upsilon}/\gamma$	0.1114	0.1114	0.1501	
	(0.0011)	(0.0011)	(0.0005)	
$\sigma_{\zeta,\upsilon}/\gamma$	0.1990	0.2116		
	(0.0010)	(0.0026)		
$E[\phi_t^\lambda]$	1.8918	1.4317	2.4705	20.7997
	(0.1117)	(0.0691)	(0.1784)	(3.2642)
$E[\kappa]$	-0.0767	-0.0767	-0.1450	-0.6952
	(0.0105)	(0.0105)	(0.0111)	(0.0093)
$DF(\Theta)$	1.9398×10^{-11}	7.0055×10^{-11}	3.8247×10^{-05}	0.0005

Note: Own calculation based on the PSID. Bootstrapped standard errors in parentheses.

Model Fit — The model attempts to fit the first three autocovariance moments of the hours and wage residuals and the covariance of the two. In Table 8 we show how the data estimates of these moments stack up against the values fit by the model. We use the main sample to evaluate the fit. As expected, the model fits the empirical moments very well.

Table 8: Model Fit

	Var. wages	1. AutoCov wages	2. AutoCov wages	Var. hours	1. AutoCov hours	2. AutoCov hours	Cov. hours & wages
emp.	0.1530	-0.0560	-0.0048	0.1443	-0.0529	-0.0052	-0.1010
mod.	0.1530	-0.0560	-0.0048	0.1443	-0.0529	-0.0052	-0.1010

Note: Variance moments of residuals obtained from the regressions of equations (5) and (4) for the main sample. Own calculation based on the PSID.

The current model does not allow for variation in these targeted variances over age groups and thus imposes that their pattern is essentially flat over the life-cycle. Figures 3a, 3b and 3c show the two residual variance series and the covariance over age. The figures show that these variances do not vary substantially over the life-cycle.

Partial consumption insurance — The parameter ϕ_t^λ is directly related to consumption growth, see equation (13) and Alan et al. (2018). In our model with endogenous labor supply permanent wage shocks translate into changes in consumption by $\phi_t^\lambda/\vartheta \times (1 + (1 - \phi_t^\lambda)/(\gamma + \phi_t^\lambda))$. We set $\vartheta = 2$, which is close to the estimates of related papers[18] and calculate the resulting pass-through at mean values of ϕ_t^λ, reported in Table 9. For the full sample we find that on average a permanent

[18] Blundell et al. (2016) estimate a parameter of relative risk aversion of 2.4 and Heathcote et al. (2014) estimate 1.7.

Figure 3: Fit of Variance and Covariance Moments over the Life-Cycle

(a) Variance of Hours Residuals

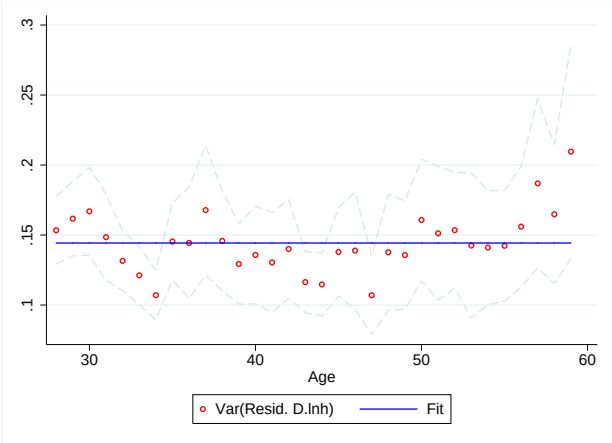

(b) Variance of Wage Residuals

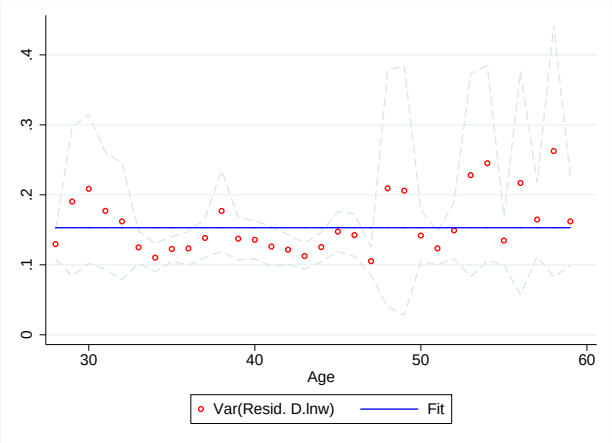

(c) Covariance of Hours and Wage Residuals

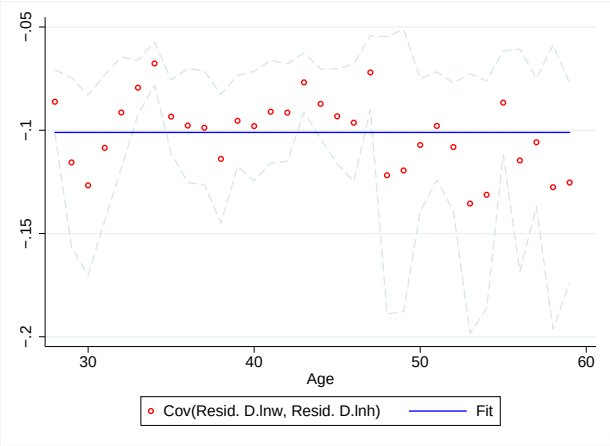

Note: Own calculation based on the PSID. Empirical and theoretical variance and covariance moments of residuals obtained from the estimation of equations (4) and (5) for the main sample with bootstrapped 95 percent confidence interval.

wage shock of one percent leads to an increase in consumption by 0.76 percent. This figure can be compared to studies that use consumption data to obtain similar parameters. Blundell et al. (2016) use 1999-2009 PSID data and find that the Marshallian response of consumption to the male's wage shock is 0.58, when female labor supply is held constant. We obtain a slightly smaller pass-through parameter for the older sample than for the main sample, but find a substantially smaller pass-through of wage shocks to consumption for the highly educated, for whom a permanent wage increase by one percent leads to an increase in consumption of just 0.31 percent. Using a similar data set to ours, 1978-1992 PSID data, Blundell et al. (2008) estimate the pass-through of permanent *income* shocks to consumption, which is given by ϕ_t^λ/ϑ in our model. With a Marshallian labor supply elasticity close to zero—as the one we have estimated—this parameter comes close to the pass-through of permanent wage shocks. Their estimate for the full sample is 0.64 and the estimate for their college sample is 0.42. We confirm the finding that the highly educated are much better insured against income shocks than is the case for the whole population.

Table 9: Pass-through of Permanent Wage Shocks to Consumption

I	II	III	IV
Full sample	Age>=40	High educ	No young children
0.7648	0.6304	0.3135	0.4820

Note: $E[\phi_t^\lambda]/\vartheta \times (1 - E[\phi_t^\lambda])/(\gamma + E[\phi_t^\lambda])$ with ϑ=2.
Own calculation based on the PSID.

Much of the literature on consumption insurance makes use of moment conditions involving consumption data. We obtain comparable estimates from labor supply and earnings data alone. Similarly, Heathcote et al. (2014) estimate their model with and without moment conditions of consumption. Their estimates of the share of insurable wage dispersion are essentially the same. A simple back-of-the-envelope calculation based on our results for the pass-through parameter to the marginal utility of wealth yields consumption insurance parameters that are broadly comparable to those obtained in previous papers using consumption data. This adds to the notion that much can be learned about consumption insurance from earnings and labor supply data alone.

6 Conclusion

At the outset we asked a simple question: What are the drivers of the riskiness of earnings? To get at the answer, we have decomposed idiosyncratic earnings uncertainty into contributions of transitory and permanent wage and hours shocks. This is a departure from extant work, where unexplained earnings volatility is entirely due to wage shocks. In order to separate hours shocks from labor supply reactions to wage shocks, we build on a life-cycle model of labor supply and consumption

and estimate a transmission parameter that captures the impact of shocks on the marginal utility of wealth and varies between individuals. This parameter is directly related to consumption insurance. We find that both wages and hours are subject to permanent shocks. At the mean, permanent wage shocks have a stronger impact on life-time earnings. Using the mean of the transmission parameter and mean annual earnings, a positive permanent wage shock of one standard-deviation at age 30 increases life-time earnings by 150 000 Dollar, while the effect of a permanent hours shock of one standard-deviation is 124 000 Dollar. Both permanent hours and wage shocks are an important source of cross-sectional earnings growth variance and earnings risk. Ergo, the data tell a story beyond wage risk.

Along the way to this result, we have shown an alternative way to calculate the Marshallian elasticity of labor supply, which we find to be negative, but small, at -0.08. There is more insurance against permanent wage shocks among the highly educated, for whom we estimate a small positive Marshallian elasticity. Setting the variance of both transitory and permanent hours shocks to zero, we estimate a Marshallian of -0.70 for the main sample, which demonstrates the importance of modeling hours shocks.

Our investigation of the sources of permanent hours shocks leads us to believe that they can be characterized as shocks to home production, but we cannot rule out that other restrictions affect the variance of hours. However, we do rule out two potential major sources of hours shocks. Permanent hours shocks persist as a phenomenon when restricting the sample to individuals who stay in their respective jobs and when excluding the years 1981-82, during which a major economic crisis hit the US. These tests, along with the results from our four main samples, strongly suggest that hours shocks are a phenomenon that is not restricted to specific, one-off adjustments or only relevant for narrowly defined groups.

Calibrating the coefficient of relative risk aversion, we calculate the pass-through of permanent wage shocks to consumption and find reasonable figures in the same range as those reported in Blundell et al. (2008, 2016). These results are encouraging as they show that comparable estimates of consumption insurance can be obtained using consumption or labor supply data.

Natural extensions of our framework include modeling family labor supply and the extensive labor supply margin. Moreover, the sources of hours shocks merit further research. One promising avenue would be to explicitly model and then separate out shocks to home production from other restrictions to labor supply.

Appendix

A Derivation of the Labor Supply Equation

The residual in the labor supply equation consists of in-period shocks and expectations corrections in the marginal utility of wealth due both to wage and hours shocks.

The first order condition of the consumer's problem w.r.t. h_t is:

$$\frac{\partial \mathcal{L}}{\partial h_t} = E_t\left[\left(-b_t h_t^{\gamma}\right) + \lambda_t w_t\right] = 0, \tag{27}$$

where $\lambda_t = \frac{\partial u(c_t, h_t, b_t)}{\partial C_t}$ denotes the marginal utility of wealth. The Euler equation of consumption is given by

$$\frac{1}{\rho(1+r_t)}\lambda_t = E_t[\lambda_{t+1}]. \tag{28}$$

Expectations are rational, i.e., $\lambda_{t+1} = E_t[\lambda_{t+1}] + \varepsilon_{\lambda_{t+1}}$, where $\varepsilon_{\lambda_{t+1}}$ denotes the mean-zero expectation correction of $E_t[\lambda_{t+1}]$ performed in period $t+1$. Expectation errors are caused by innovations in the hourly wage residual ω_{t+1} and innovations in hours shocks υ_{t+1}, which, as implied by rational expectations, are uncorrelated with $E_t[\lambda_{t+1}]$. Rational expectations imply that $\varepsilon_{\lambda_{t+1}}$ is uncorrelated over time, so that regardless of the autocorrelative structure of the shock terms, $\varepsilon_{\lambda_{t+1}}$ will only be correlated with the innovations of the shock processes.

Resolving the expectation operator in equation (27) yields

$$b_t h_t^{\gamma} = \lambda_t w_t. \tag{29}$$

Taking logs of both sides we arrive at the structural labor supply equation

$$\ln h_t = \frac{1}{\gamma}\left(\ln \lambda_t + \ln w_t - \ln b_t\right). \tag{30}$$

To find an estimable form for $\ln h_t$, we take logs of (28) and resolve the expectation:

$$\ln \lambda_t = \ln(1+r_t) + \ln \rho + \ln\left(\lambda_{t+1} - \varepsilon_{\lambda_{t+1}}\right)$$

A first order Taylor-expansion of $\ln\left(\lambda_{t+1} - \varepsilon_{\lambda_{t+1}}\right)$ gives $\ln(\lambda_{t+1}) - \frac{\varepsilon_{\lambda_{t+1}}}{\lambda_{t+1}}$, leading to the expression

$$\ln \lambda_t = \ln(1+r_t) + \ln \rho + \ln(\lambda_{t+1}) - \frac{\varepsilon_{\lambda_{t+1}}}{\lambda_{t+1}} + O\left(-\frac{1}{2}\left(\frac{\varepsilon_{t+1}}{\lambda_{t+1}}\right)^2\right). \tag{31}$$

Accordingly, when we backdate (31), we can insert it in (30) and remove $\ln \lambda_t$ by first differencing.

B Distribution of the Shock Pass-Through on Hours

The moments of the term $\frac{\phi_t^\lambda}{\gamma+\phi_t^\lambda}$ are not as tractable as the rest of the random variables in the variance moment estimation, since we assume $\ln \phi_t^\lambda \sim N(\mu_\phi, \sigma_\phi)$. We can refine the expression to find a more basic expression:

$$\frac{\phi_t^\lambda}{\gamma + \phi_t^\lambda} = 1 - \gamma \frac{1}{\gamma + \phi_t^\lambda}$$

The only random term in this expression is $\frac{1}{\gamma+\phi_t^\lambda}$. We can find its distribution by re-expressing its CDF in terms of the underlying normal distribution of $\ln \phi$. Let $\frac{1}{\gamma+\phi_t^\lambda} = Z$. Then

$$P(Z \leq z) = P\left(\frac{1}{\gamma + \phi_t^\lambda} \leq z\right) \tag{32}$$

$$P\left(\ln \phi_t^\lambda \leq \ln\left(\frac{1}{z} - \gamma\right)\right) = \int_{-\infty}^{\ln\left(\frac{1}{z}-\gamma\right)} \frac{\exp\left(-\frac{(x-\mu_\phi)^2}{2\sigma_\phi^2}\right)}{\sqrt{2\pi\sigma_\phi^2}} dx \tag{33}$$

Integrating this CDF, we find the CDF for the random variable Z.

$$F(z) = 1/2\left(1 - \mathrm{Erf}\left(\frac{\ln\left(\frac{1}{z} - \gamma\right) - \mu_\phi}{(2\sigma_\phi^2)^{1/2}}\right)\right)$$

Here $\mathrm{Erf}(\cdot)$ is the Gaussian error function. To generate the first and second noncentral moments, we take the derivative to find the PDF of Z.

$$f(z) = -\frac{\exp\left(-\frac{\left(\ln\left(\frac{1}{z}-\gamma\right)-\mu_\phi\right)^2}{2\sigma_\phi^2}\right)}{\sqrt{2\pi\sigma_\phi^2} z(1 + z\gamma)}$$

The first and second noncentral moments are $\mathcal{M}_1 = \int_0^{1/\gamma} z f(z) dz$ and $\mathcal{M}_2 = \int_0^{1/\gamma} z^2 f(z) dz$. These are calculated via numerical integration, as there is no closed form solution. We implement these formulas in our moment conditions. In estimation we restrict the values of μ_ϕ not to exceed 5, as the moments of $\frac{\phi_t^\lambda}{\gamma+\phi_t^\lambda}$ asymptote beyond that point.

C Tables

Table 10: Frisch Labor Supply Equation Estimation

	I Full sample	II Age>=40	III High educ	IV With children
$\Delta \ln(wage)$	0.3614	0.4020	0.2851	0.3148
	(0.0856)	(0.3778)	(0.0975)	(0.1080)
N	46340	20607	19831	24547
Kleibergen and Paap (2006) F stat	18.4680	1.2408	11.7317	11.6739

Note: Own calculation based on the PSID. Clustered standard errors in parentheses.

D Distributions of Weekly Hours of Work and Weeks Worked

Figure 4: Distributions of Weekly Hours of Work and Weeks Worked

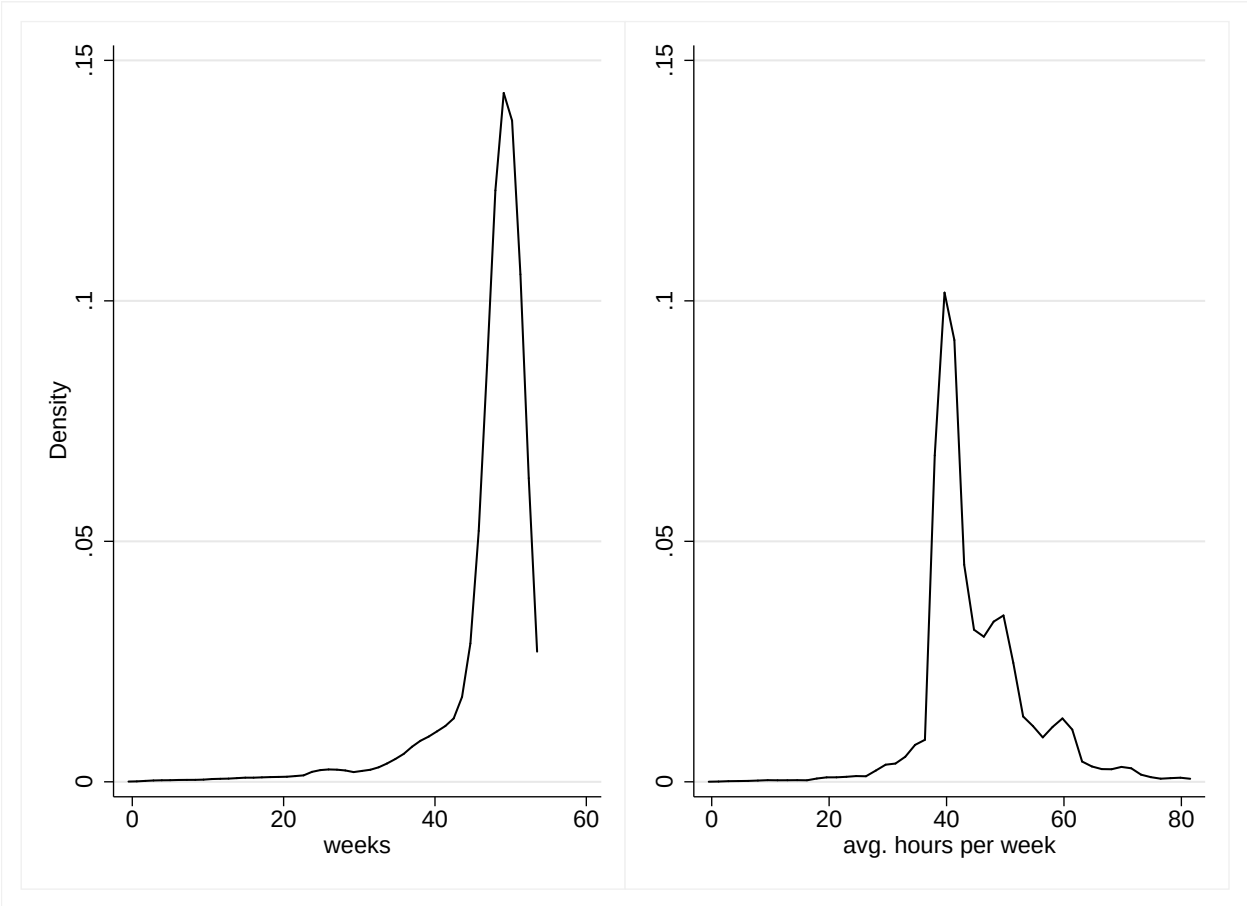

Note: Own calculation based on the PSID. Epanechnikov kernel densities of weekly hours of work and hours worked per year for the pooled cross-section, bandwith 1.5.

References

Abowd, J. M., and D. Card (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57(2), 411–445. Cited on pages 3 and 4.

Alan, S., M. Browning, and M. Ejrnæs (2018): "Income and Consumption: a Micro Semistructural Analysis with Pervasive Heterogeneity," *Journal of Political Economy*, 126(5). Cited on pages 8, 13, and 22.

Altonji, J. G. (1986): "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy*, 94(3), S176–S215. Cited on pages 6 and 8.

Altonji, J. G., and L. M. Segal (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 14(3), 353–366. Cited on page 14.

Altonji, J. G., A. A. Smith Jr., and I. Vidangos (2013): "Modeling Earnings Dynamics," *Econometrica*, 81(4), 1395–1454. Cited on pages 4 and 17.

Attanasio, O. P., and L. Pistaferri (2016): "Consumption Inequality," *Journal of Economic Perspectives*, 30(2), 3–28. Cited on page 2.

Blundell, R., H. Low, and I. Preston (2013): "Decomposing changes in income risk using consumption data," *Quantitative Economics*, 4(1), 1–37. Cited on page 4.

Blundell, R., and T. Macurdy (1999): "Labor supply: A review of alternative approaches," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3 of *Handbook of Labor Economics*, chap. 27, 1559–1695. Elsevier. Cited on page 14.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921. Cited on pages 3, 4, 7, 24, and 25.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016): "Consumption inequality and family labor supply," *The American Economic Review*, 106(2), 387–435. Cited on pages 3, 4, 7, 10, 14, 15, 16, 17, 22, 24, and 25.

Bound, J., C. Brown, G. J. Duncan, and W. L. Rodgers (1994): "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data," *Journal of Labor Economics*, 12(3), 345–368. Cited on page 10.

Friedman, M. (1957): *A Theory of the Consumption Function*, NBER Book. National Bureau of Economic Research. Cited on page 3.

Geyer, J., and T. Korfhage (2015): "Long-term Care Insurance and Carers' Labor Supply–A Structural Model," *Health economics*, 24(9), 1178–1191. Cited on page 2.

Guvenen, F. (2007): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?," *American Economic Review*, 97(3), 687–712. Cited on page 3.

——— (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12(1), 58–79. Cited on page 3.

Hall, R. E. (1978): "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86(6), 971–987. Cited on page 3.

Heathcote, J., K. Storesletten, and G. L. Violante (2014): "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, 104(7), 2075–2126. Cited on pages 3, 4, 14, 16, 17, 22, and 24.

Hryshko, D. (2012): "Labor income profiles are not heterogeneous: Evidence from income growth rates," *Quantitative Economics*, 3(2), 177–209. Cited on pages 4, 11, and 14.

Imai, S., and M. P. Keane (2004): "Intertemporal Labor Supply and Human Capital Accumulation," *International Economic Review*, 45(2), 601–641. Cited on pages 16 and 18.

Kaplan, G. (2012): "Inequality and the life cycle," *Quantitative Economics*, 3(3), 471–525. Cited on page 4.

Keane, M. P. (2011): "Labor supply and taxes: A survey," *Journal of Economic Literature*, 49(4), 961–1075. Cited on pages 14 and 16.

Kleibergen, F., and R. Paap (2006): "Generalized reduced rank tests using the singular value decomposition," *Journal of Econometrics*, 133(1), 97–126. Cited on pages 16 and 28.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100(4), 1432–67. Cited on page 4.

MaCurdy, T. E. (1981): "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 89(6), 1059–85. Cited on pages 6 and 11.

MaCurdy, T. E. (1982): "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics*, 18(1), 83–114. Cited on page 3.

Meghir, C., and L. Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32. Cited on pages 3, 10, and 15.