Happiness and the Persistence of Income Shocks

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

This version: July 27, 2012

Abstract

We reassess the empirical effect of income and employment on self-reported well-being. Our analysis makes use of a novel two-step estimation procedure that allows applying instrumental variable regressions with ordinal observable data. As suggested by the theory of incomplete markets, we differentiate between the effects of persistent and transitory income shocks. In line with this theory, we find that persistent shocks have a significant impact on happiness while transitory shocks do not. This has consequences also for inference about the happiness effect of employment. We find that employment per se is rather associated with a decline in happiness.

KEYWORDS: incomplete markets, happiness, income persistence
JEL-codes: E21, D12, D60

*We would like to thank Andrew Oswald and Paul Frijters for helpful comments. Financial support by the ERC under ERC-StG -282740 "HUMP" and the DFG under Sonderforschungsbereich 823, project "Household Economics, Labor Markets, and Macroeconomic Dynamics" is gratefully acknowledged.

†Universität Bonn, Department of Economics, Adenauerallee 24-42, 53113 Bonn, Germany; phone: +49-228-73 4073; email: christian.bayer@uni-bonn.de.

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

1 Introduction

Do individual economic conditions contribute to a person's well-being or happiness? Frey and Stutzer (2002) and Clark et al. (2008) survey the large body of literature that has devoted itself to this question. For our purpose we may summarize the findings of this literature as follows:

1. in the cross-section there is a small but significantly positive correlation between household income and happiness, and
2. employment contributes substantially to happiness beyond the gain in income.[1]

Our paper contributes to this literature by qualifying both results. We show that the contribution of income shocks on a person's happiness depends crucially on the persistence of these shocks and that taking this into account changes what one infers about the contribution of employment to happiness. This is important as the findings of the happiness literature are sometimes viewed as a challenge to standard, neoclassical economics.

In a simple model of a world with complete markets, the findings as surveyed by Frey and Stutzer (2002) and Clark et al. (2008) are indeed a puzzle for neoclassical economics. In such world, households should face complete consumption insurance with respect to their idiosyncratic income shocks. Therefore, the coefficient of a cross-sectional regression of happiness on income should actually be zero (and not positive). Moreover, with additive separability in consumption and leisure, standard preferences suggest the coefficient of employment on happiness should be negative (and not strongly positive).[2]

Parts of the happiness literature (e.g. when calculating income compensation for, say, airport noise, cf. van Praag and Baarsma (2005)), however, seem to start out from an alternative – not less simplifying – model of autarky when interpreting the empirical results. Also in this setup, where insurance markets are entirely absent, the empirical results of the happiness literature rest surprising. Under autarky, an increase in income translates into an equal increase in consumption and hence should lead to a significant increase in happiness if market consumption is an important determinant of

[1] See in particular Clark and Oswald (1994), Clark et al. (2001), and Clark and Delta (2003).

[2] Yet, already moving away from additively separable preferences makes it arbitrary what to expect as an effect of employment on happiness. Take for example Hansen's (1985) model of indivisible labor. There, it depends on the cross derivative of utility in consumption and leisure whether a worker or an unemployed achieves a higher utility level. If less leisure increases the marginal utility from consumption then the employed would obtain higher consumption levels under full insurance, maybe even being overcompensated for the direct utility loss from less leisure.

life-satisfaction. Also, the positive coefficient on employment remains surprising if one thinks of leisure as providing utility.

Now, the modern (macro)economic literature has to a good part moved away from either two extreme assumptions – neither are insurance markets complete nor completely absent. The workhorse model of heterogeneous-agent macroeconomics, the standard incomplete markets model, see Bewley (1980), Huggett (1993), and Aiyagari (1994), assumes that households face idiosyncratic labor market risk against which they can only self-insure by accumulating or de-cummulating uncontingent assets. Importantly, in this setting, households result differently capable to insure against income shocks with different persistence. Kaplan and Violante (2010), applying a method developed by Blundell et al. (2008), show that both in the standard incomplete markets model as well as in consumption data, households are much better able to insure against transitory than against persistent shocks to their incomes.

We show that this point has important consequences for the interpretation of usual happiness regressions.[3] First and foremost, the coefficient on income in a happiness regression is a mixture of the effects of transitory and persistent income shocks. Second and consequently, this introduces a bias on the estimated coefficients of other variables if these variables are not just equally correlated with persistent and transitory shocks to income. Employment is an example for such variable. Non-employment spells are correlated with long term, i.e. persistent, decreases in income, see e.g. Arulampalam (2001).

In fact, we show, applying Blundell et al. (2008)'s framework to happiness data from the German Socio-Economic Panel (SOEP), that: first, persistent income shocks translate significantly stronger to happiness than average income shocks (more than twice as much), second, that transitory shocks do not significantly contribute to happiness, i.e. are perfectly insured, and third, that this leads to a strong upward bias in the coefficient estimate on employment in a happiness regression.

The remainder of this paper is organized as follows: Section 2 develops the econometric model and methodology, Section 3 introduces the data set, Section 4 presents results, Section 5 relates our partial insurance results to some existing findings about differently strong effects of income and employment on happiness in different economic environments. Section 6 concludes.

[3]The only paper we are aware of that argues for an insurance interpretation of the low income coefficient in a happiness regression is Dehejia et al. (2007).

2 Econometric Model and Methodology

The economic analysis of self-reported happiness data usually starts off with applying an ordered probit model to the data and then discussing the effects of various (control) variables on happiness, most importantly income and the employment status.[4] We deviate from this tradition only in one, but as it turns out, important point which we borrow from the consumption / incomplete markets literature. We assume that income shocks can be partially insured by the household and shocks of different persistence can be differently insured.

2.1 Income and latent utility

To distinguish between shocks of different persistence, we need to put some structure on the income process of a household. We assume a household i's log income y_{it} at time t is composed of a component $g(z_{it})$ that reflects deterministic effects of household characteristics z_{it}, a transitory stochastic component ψ_{it}, a persistent stochastic component x_{it}, and a fixed component μ_i^y, such that

$$y_{it} = g(z_{it}) + y_{it}^* \quad (1)$$

$$y_{it}^* = x_{it} + \psi_{it} + \mu_i^y \quad (2)$$

$$x_{it} = \rho x_{it-1} + \epsilon_{it} \quad (3)$$

where ψ and ϵ are i.i.d. shocks. For most of this paper, we will assume for simplicity the persistent income component to follow a random walk ($\rho = 1$) but we will check the sensitivity of our results with respect to this assumption.

Next, we assume that a household has a felicity function that translates leisure and market consumption into utility. This utility u_{it}^{**} is latent and we assume in line with Blundell et al. (2008) that plugging in the household's consumption choice into this felicity function yields

$$u_{it}^{**} = u_{it}^* + f(z_{it}) \quad (4)$$

$$u_{it}^* = \alpha_x x_{it} + \alpha_\psi \psi_{it} + \mu_i^u + \xi_{it}. \quad (5)$$

The coefficients α_x and α_ψ measure how much the household values consumption of

[4] The literature has discussed a number of potential econometric problems in this setup, in particular those stemming from fixed effects in happiness and income, and has proposed solutions thereof, see e.g. Frijters et al. (2004b).

market goods and importantly how well insured a household is against income fluctuations. For example, we would expect $\alpha_x = \alpha_\psi > 0$ under autarky, $\alpha_x = \alpha_\psi = 0$ under complete markets, and under incomplete markets, with self-insurance, $\alpha_x > \alpha_\psi \approx 0$, see Kaplan and Violante (2010). The fixed effect μ_i^u measures permanent utility differences between households and $f(z_{it})$ captures the effects of other household characteristics, e.g. marital status, health, age, schooling, etc. The importance of market consumption for felicity can be grasped from comparing the size of α_x and α_ψ to the coefficients of other observables in $f(z_{it})$.

Equation (5) cannot be directly estimated for two reasons. First, the latent utility level is not observable and second, we do not observe separately the persistent and the transitory income components. Instead we only observe income y. For now assume, however, that we could in fact observe latent utility (we come back to this point in Section 2.2). Moreover, assume we know f and g. Under these assumptions, we can identify (for $\rho = 1$) α_x and α_ψ from the regression equations

$$\Delta u_{it}^* = \alpha_x \Delta y_{it}^* - (\alpha_x - \alpha_\psi)\Delta\psi_{it} + \Delta\xi_{it}, \tag{6}$$

$$\Delta u_{it}^* = \alpha_\psi \Delta y_{it}^* + (\alpha_x - \alpha_\psi)\epsilon_{it} + \Delta\xi_{it}. \tag{7}$$

by the instrumental variable regressions suggested in Blundell et al. (2008). To identify α_x, we instrument Δy_{it}^* by $y_{it+1}^* - y_{it-2}^*$ and to identify α_ψ, we instrument Δy_{it}^* by Δy_{it+1}^*. These instruments are uncorrelated with $\Delta\psi_{it}$ respectively ϵ_{it}, but correlate with Δy_{it}^*, see Appendix A for further details. Intuitively, the first instrument, the three year growth rate in income, captures the persistent movements in income, whereas the latter, the next year's growth rate in income is negatively correlated to the current growth rate through this year's transitory income shock.

While the instrumental variable regressions identify the two parameters α_x and α_ψ separately, a simple OLS regression

$$\Delta u_{it}^* = \alpha \Delta y_{it}^* + \upsilon_{it}$$

of latent utility on income y yields an estimated coefficient $\hat{\alpha}$ that can be understood as a weighted mean of α_ψ and α_x that ensures

$$E\left\{\Delta y_{it}^* \left[(\alpha_\psi - \alpha)\Delta\psi_{it} + (\alpha_x - \alpha)\epsilon_{it} + \Delta\xi_{it}\right]\right\} = 0, \tag{8}$$

as with (1) and (5)

$$\upsilon_{it} = (\alpha_\psi - \alpha)\Delta\psi_{it} + (\alpha_x - \alpha)\epsilon_{it} + \Delta\xi_{it}.$$

Plugging in $\Delta\psi_{it} + \epsilon_{it}$ for Δy^*_{it} we obtain for the OLS estimate

$$plim\ \hat{\alpha} = \alpha_\psi \frac{2\sigma^2_\psi}{2\sigma^2_\psi + \sigma^2_\epsilon} + \alpha_x \frac{\sigma^2_\epsilon}{2\sigma^2_\psi + \sigma^2_\epsilon}. \tag{9}$$

Hence, the estimate for α is a weighted sum of the effects of permanent and transitory shocks to income, where the weights are equal to the contribution of permanent and transitory shocks to the variance of income growth.

Importantly, the different effects of permanent and transitory income shocks imply that any coefficient estimate on additional variables we may include in (5) is likely biased – we omit a variable, the difference between observed income growth and the growth of the permanent income component. Let us assume there is another observable variable e_{it}, not included in z_{it}. Concretely, assume this is the employment status, whose effect on utility is γ. Now consider again the regression equation and assume we are satisfied with estimating the average income effect as in (8)

$$\Delta u^*_{it} = \alpha\Delta y^*_{it} + \gamma\Delta e_{it} + \upsilon_{it}, \tag{10}$$

$$\begin{aligned} \upsilon_{it} &= (\alpha_\psi - \alpha)\Delta\psi_{it} + (\alpha_x - \alpha)\epsilon_{it} + \Delta\xi_{it} \\ &= (\alpha_\psi - \alpha_x)\left[\frac{\sigma^2_\epsilon}{2\sigma^2_\psi + \sigma^2_\epsilon}\Delta\psi_{it} - \frac{2\sigma^2_\psi}{2\sigma^2_\psi + \sigma^2_\epsilon}\epsilon_{it}\right] + \Delta\xi_{it}, \end{aligned} \tag{11}$$

where the last equation assumes α to be the weighted mean from before. As we omit the separation of average and permanent income growth, it enters in the composed error term υ_{it}, which contains both permanent and transitory shocks. As a result the estimates on α and γ will be generically biased. They will only unbiased in the special case where

$$\frac{cov(\Delta e_{it}, \Delta\psi_{it})}{cov(\Delta e_{it}, \epsilon_{it})} = \frac{2\sigma^2_\psi}{\sigma^2_\epsilon},$$

i.e. when employment change is associated with permanent and transitory income fluctuations that have the same relative size as average permanent and transitory income shocks.

To illustrate the bias, suppose we observe a person who moves out of employment. Then, she faces a decrease in her income. Yet, to obtain an unbiased estimate of the

effect of employment on happiness, γ, we need to assume that this decrease in income induced by unemployment is just as persistent as any average random income fluctuation. However, this is unlikely to be the case if the person faces generous unemployment benefits that insure well the short-run income loss. In this case the unemployment-induced decrease in income reflects mostly persistent income components. Consequently, there is a positive correlation between employment and the composed error term υ, which leads to an upwards bias in γ.

However, we can estimate the parameter γ if augmenting (6) by Δe_{it} and instrumenting it by e_{it-2} and Δy^*_{it} by $y^*_{it+1} - y^*_{it-2}$ as before. e_{it-2} should be a valid instrument, as the employment status two periods ago should be uncorrelated to a *transitory* income shock in period t and correlated with Δe_{it} if e is a stationary but persistent process. By contrast, it is not straightforward to find an instrument for e in (7) from the time-series of employment histories. Future and past employment histories may well correlate with persistent income shocks.[5]

2.2 Constructing latent utility from observed happiness

So far, we have established an instrumental variable regression to estimate the effects of persistent and transitory income shocks and employment on latent utility assuming this latent utility is observable. While latent utility is not observable, we do observe self-reported life-satisfaction in the data we use. This variable is reported on a scale from 0 to 10. We assume that this happiness variable is generated from an ordered probit model, where happiness h_{it} is determined by

$$h_{it} = j \text{ if } u^{**}_{it} \in (\overline{c}_j, \overline{c}_{j+1}] \tag{12}$$

The latent u^{**}_{it} is determined as in (4) and all error terms $\mu^y_i, \mu^u_i, \epsilon_{it}, \psi_{it}, \xi_{it}$ are normally distributed. Moreover, we assume that u^{**}_{it} is scaled such that residual utility u^*_{it} has unit variance.

Under these assumptions we can estimate the cutoff values $\overline{c}_j$ and the statistical (not necessary causal) effect of controls $f(z_{it})$ by a standard ordered probit estimator. Note that we should not give causal interpretation to these estimates as they will also include correlations of controls with fixed effects and income shocks.[6] The cutoff values

[5] Examples are skill depreciation in past unemployment (leading to a correlation between current persistent income and lagged employment) and persistent productivity causing lower firing probabilities (leading to a correlation between current persistent income and future employment).

[6] Frijters et al. (2004b) suggest an estimator to obtain consistent estimates in the presence of fixed effects for the ordered probit setup. We do not employ their estimator in our first-stage regression as

Table 1: Sample Selection

Initial number of observations	224,127
After constraining to ages 25-55	134,494
After accounting for missings in	
education	132,554
income	132,524
happiness	132,173
health satisfaction	131,989
Final number of observations for first-stage regressions	131,989
Final number of observations for which all instruments can be constructed	82,803

are scaled appropriately to be compatible with u^*_{it} having a unit variance. Since we obtain for each household-year an estimate $\widetilde{f(z_{it})}$, we can infer an interval $U_{it} = (\tilde{\bar{c}}_{h_{it}} - \widetilde{f(z_{it})}, \tilde{\bar{c}}_{h_{it}+1} - \widetilde{f(z_{it})}]$ in which u^*_{it} must have been fallen. Together with the normality assumption for u^*_{it} this means we can calculate the conditional expected value $\overline{u^*_{it}}$ for residual latent utility

$$\overline{u^*_{it}} = \frac{\int_{u \in U_{it}} u\phi(u)}{\Phi(U_{it})}, \tag{13}$$

where ϕ is the density of a standard normal and $\Phi(U)$ the probability of U for a standard normal distribution.

Replacing u^*_{it} by $\overline{u^*_{it}}$ in the estimation equations derived in Section 2.1 renders the previously derived estimators feasible. It introduces measurement error, but only to the dependent variable, which does not bias estimations. The huge advantage of our procedure is that we can apply standard linear regression techniques once $\overline{u^*_{it}}$ is estimated and hence can e.g. use first differences to control for fixed effects.

we are not interested in obtaining structural estimates in this first stage. While their estimator is more efficient in the presence of fixed effects than our estimation procedure, the advantage of the latter is that it is easily extended to the IV regressions we need to do.

Table 2: Descriptive Statistics

Variable	Mean	Std. Dev.	Min	Max
Happiness	7.145	1.740	0	10
Income	10.431	0.534	2.594	13.770
Age	40.564	8.358	25	55
Schooling (semesters)	24.150	5.192	14	36
Household size	3.032	1.249	1	6
Satisfaction with health	6.953	2.113	0	10
No. of children	0.892	1.015	0	4

Fraction of respondents who are ...

Disabled	6%	Employed	78%
Males	52%	Living with a spouse	83%

3 Data

We use data on subjective well-being from the German annual socio-economic panel (SOEP). The SOEP is a representative longitudinal study of private households and individuals and covers information on household composition, employment, earnings, health and satisfaction indicators. The time period used in our analysis is 1984-2010. In the baseline specification, we restrict the sample to household heads and spouses between 25-55 years of age, consider West German households only and drop observations from the migrant and high income samples. In an alternative specification, we split the sample by gender.[7]

Individual happiness is measured on an integer scale from 0 to 10.[8] To measure

[7] Throughout, we take the variables from the 100%-sample version of the Cross-National Equivalent File of the SOEP ($PEQUIV-files). Variable keys are provided in Appendix C.

[8] The survey question is: How satisfied are you with your life, all things considered? 0 means completely dissatisfied, 10 means completely satisfied.

income, we use post-government income in real terms, which represents the combined income after taxes and government transfers in the previous year of all individuals in the household. We define a person as being employed when supplying more than 520 hours of market work per year, being equivalent to more than a quarter of full-time employment. Alternatively, we use a broader definition of employment, where persons having positive wages in the previous year and working at least 52 hours are classified as employed. Table 1 provides information on the number of observations as well as on the number of observations we loose due to sample selection.

To control for outliers, we drop those households that fall in the top-bottom .25% percentiles of residual incomes from a first stage regression (see below) in each year. We then re-estimate the first-stage income regression for the cleaned sample. Table 2 shows summary statistics of the variables used in the final estimation sample.

4 Results

The first step of our analysis is to regress household incomes on a large set of control variables z_{it} that include dummies for each year of schooling, dummies for age, for marital status, for living with a spouse, for the number of children,[9] for the various levels of self-reported health status, for disability and interaction terms of schooling coded in 5 levels with a second order age polynomial. We use the same set of variables for our first step ordered probit regression of happiness. This gives us estimates of $\overline{u_{it}^*}$ and y_{it}^* as defined in Section 2.

4.1 Happiness and income

We can use these data to estimate the effect of income on happiness. Table 3 reports the main results of this exercise. While the simple OLS regression (Column i) suggests some significant positive effect of income on happiness (a 50% increase in income has roughly the same effect as being married), this coefficient drops significantly when using first differences to control for fixed household effects (Column ii).[10] This finding is in line with what other researchers have found, cf. Ferrer-i Carbonell and Frijters (2004). In other words, there are households that are both permanently more happy and permanently earn more. As these differences are fixed, we cannot identify what causes what and

[9] We top-code 379 household observations with more than 4 children.

[10] Table 7 in Appendix B shows that the coefficient estimates from the OLS estimator are not significantly different from a one-step ordered probit estimation including not only the controls but also income (and employment).

Table 3: Happiness, Income and Employment

	i OLS	ii FD	iii FD	iv IV transitory	v IV permanent	vi IV permanent
Income shocks						
- all	.316	.211	.201			
	(.010)	(.015)	(.015)			
- transitory				.036		
				(.043)		
- permanent					.421	.460
					(.064)	(.077)
Employment			.063			-.109
			(.012)			(.113)

Standard errors in parenthesis. OLS refers to an OLS estimation of $\overline{u_{it}^*}$ on y_{it}^*, FD to the same regression using first-differences to control for fixed household effects, in Column iii we include employment e_{it} as an additional regressor. IV permanent and IV transitory refer to the IV regressions derived in (6) and (7) respectively to identify the effects of permanent and transitory income shocks. In iv we instrument Δy_{it}^* by $y_{it+1}^* - y_{it}^*$, in v and vi we instrument Δy_{it}^* by $y_{it+1}^* - y_{it-2}^*$ to identify the effect of persistent income shocks and in vi we additionally instrument Δe_{it}^* by e_{it-2} to identify employment effects. The IV regressions all control for fixed effects by first differencing. Both the income and employment variables have been regressed on the same set of controls we included in the first-stage ordered probit regression for happiness. Residuals from these regression are used as regressors here.

causation may be reversed. A household may earn more because its members are happy. Yet, in light of the incomplete markets and consumption literature, there is another explanation. Permanent income shocks can be insured much less and hence translate stronger into utility. Differentiating them out therefore decreases the coefficient estimate on income.

The instrumental variable regressions in Table 3 show that, in fact, permanent shocks influence happiness more strongly than transitory ones (Columns iv and v). When instrumenting in order to identify the effect from permanent income shocks (Column v), the income coefficient is with .421 twice as large as for the average income shock (column ii) while transitory income shocks have no (significant) impact on happiness (Column iv). Note that our IV regressions still control for fixed effects by first differencing.

4.2 Happiness and employment

Since we find a strong difference in the effects of transitory and persistent income shocks on happiness, our theoretical considerations from Section 2 suggest that augmenting the happiness regression by employment as an additional variable leads to a biased estimate of the effect of employment on happiness. The sign of this bias depends on the correlation of employment with transitory and persistent income shocks. As we work with German data, where the replacement rate in unemployment insurance is relatively high, the main income effect of an unemployment spell is likely the persistent one, which suggests an upward bias in the estimate.

In fact, this is the case as Columns iii and vi in Table 3 show, which present the estimates from a regression augmented by employment. A naive interpretation of the estimation in first differences suggests that a household suffers from loosing employment (above the caused drop in income) just as much as from a $\hat{\gamma}/\hat{\alpha} \approx 30\%$ drop in income.[11]

In other words, a non-employed household would be indifferent between working and earning 70% of the unemployment benefits or not working and earning full unemployment benefits. Yet, as our IV procedure in Column (vi) shows, this finding is just an artifact of not controlling for the correlation of non-employment spells with permanent income shocks. Once we do so, the effect of employment on happiness becomes negative (but insignificant). The point estimate suggests that a household needs to be compensated by an income that exceeds unemployment benefits by 20% in order to be indifferent between working and not working.

[11] Compared to Winkelmann and Winkelmann (1998) we find much smaller effects of unemployment on happiness, yet this might well be due to our richer set of control variables that also include health satisfaction and non-parametric age and schooling effects.

Table 4: Robustness Check I: Pseudo-differences

	i IV-p	ii IV-p	iii IV-t
Income [permanent (p) or transitory (t) shock]	.428 (.056)	.498 (.127)	.015 (.047)
Employment		-.211 (.345)	

See notes to Table 3. In i and ii, we instrument $y^*_{it} - \rho y_{it-1}$ by $y_{it+1} - \rho^3 y^*_{it-2}$, in iii we instrument $y^*_{it} - \rho y_{it-1}$ by $y^*_{it+1} - \rho y_{it}$.

Table 5: Robustness Check II: Alternative Measures of Employment and Happiness

	w/o health satisfaction in first stage regression			lower empl. cutoff	log hours
	FD	IV-t	IV-p	IV-p	IV-p
Income (all / t. / p.)	.111 (.013)	-.131 (.038)	.505 (.066)	.439 (.073)	.534 (.095)
Employment	.033 (.011)		-.129 (.099)	-.050 (.090)	-.132 (.101)

See notes to Table 3. IV-t refers to the transitory shock instrumentation, IV-p to the persistent shock instrumentation In the penultimate column, we define a person to be employed if working more than 52 hours in the reporting year, in the last column we replace the employment indicator by log-hours worked. Note that this restricts the sample to persons who have worked at least one hour.

Table 6: Robustness Check III: Sample Splits

	Men			Women			Public Employees		
	FD	IV-t	IV-p	FD	IV-t	IV-p	FD	IV-t	IV-p
Income	.237	.102	.334	.222	.085	.635	.210	.249	.294
(all / t. / p.)	(.023)	(.066)	(.121)	(.023)	(0.66)	(.127)	(.050)	(.149)	(.192)
Employment	.142		-.097	.038		-.156			
	(.027)		(.269)	(.014)		(.142)			

See notes to Table 3. IV-t refers to the transitory shock instrumentation, IV-p to the persistent shock instrumentation. In the last three columns we restrict the sample to households where all income is from employment as a public employee.

4.3 Robustness

As a first robustness check we relax the random-walk assumption for income and assume $\rho = .9$ in line with what Bayer and Juessen (2009) report as an estimate from SOEP data and construct pseudo-differences of $\overline{u^*_{it}}$, y^*_{it} and e_{it}. We then instrument with the instruments for income suggested in Kaplan and Violante (2010) and again with e_{it-2} for employment. Results are shown in Table 4. Using pseudo-differences yields basically the same picture as under the unit-root assumption for income.

As a second set of robustness checks, we exclude the health satisfaction variable from the first stage regression, use an average annual employment of 1 hour per week as indicator of being employed, or replace the employment measure by the number of hours worked. We find that persistent income shocks translate more into happiness than do transitory ones and the effect of employment on happiness is negative throughout when using the IV estimator. This holds also true when using a lower cutoff for considering a person employed or when regressing on log-hours as continuous variable.

In a third set of robustness checks, we split the sample by men and women, and look at public employees only. The results are reported in Table 6. Qualitatively, the results for the sample split according to gender do not differ from our estimates when pooling men and women.

What is interesting from a theory point of view is that women not only exhibit a

higher dis-utility from work (in line with data on female labor market participation), but also seem to be able to insure permanent income shocks comparatively less, i.e. they have a much higher pass-through of persistent income shocks on happiness. Being able to differentiate between gender in the pass-through from income to happiness shows a further strength of the happiness data - we have information on the individual level, which typically is not the case for consumption data where some consumption goods are public goods within the household. In substance the finding of women being less insured to permanent income shocks might call for a careful future analysis as this may relate to differences in insurance abilities within or outside the household (e.g. due to differences in household bargaining power) which should be reflected in consumption data or deep psychological differences where women suffer more from fluctuations in available economic resources.

The results for public employees are reconfirming our identification idea. For civil servants, transitory income shocks can be considered fairly unimportant given the nature of the compensation schemes in the German civil service. Hence, we expect the IV and FD estimates to be similar and this is what we find.

5 Discussion

How do our findings relate to the previous economic happiness literature? Most closely related is Dehejia et al. (2007) who show that households who have access to informal insurance markets through religious organizations show both a weaker consumption-income and a weaker happiness-income relationship. More broadly speaking, there is a number of papers showing that the coefficient of income on happiness is larger in less developed countries, see e.g. Graham and Pettinato (2002). With respect to these papers, our results suggest that some part of the difference between countries is related to differences in capital-market development that lead to different insurance against income shocks. Similarly, we can reinterpret the findings of Frijters et al. (2004a, 2006) and Caporale et al. (2009) that show for Russia and East Germany, respectively the Czech Republic, Poland, and Hungary, relatively strong income effects on happiness. In particular, Caporale et al. (2009) explicitly compares to Western Europe where they find lower income effects. Our results suggest that these cross-country differences might also be driven by: first, income changes as a result of economic transition (e.g. human capital that was valuable in communist times might have become obsolete, overall productivity catches up, etc.) being particularly persistent and second, transition economies still having underdeveloped capital markets, and third, the within-country asset distribution

being not yet in steady state, such that self-insurance abilities are below their long-run level. In particular, the results by Lelkes (2006) are reassuring for this interpretation. Lelkes (2006) provides evidence for the happiness-income relation in Hungary in 1992 and 1998 and shows that during transition the income-happiness relation declined.

With respect to our results concerning the happiness effects of being employed, some comments may be in order. First, one should not read the results as "unemployment does not harm". First, the strong difference between IV and FD estimate points towards important long-run income effects of unemployment. Second, our regressions, at least in the baseline specification, control for happiness effects of health and given that there is a literature discussing unemployment effects on health, there may be indirect effects of unemployment on happiness through health, which we keep constant. Third, our employment definition does not allow us to discriminate between non-participation and unemployment, while both may have different effects on happiness.

6 Conclusion

This paper has reassessed the link between household income, employment and happiness in light of an incomplete markets setup, where households can only self insure against income shocks. This limited ability to insure predicts a positive income-happiness relation. More importantly, it predicts that shocks with different persistence have different impact on happiness. This is exactly what we find in the happiness data we analyze. While persistent income shocks have an impact on happiness, transitory income shocks do not and are hence perfectly insured. Besides this point, we show that disregarding the differential impact of income shocks with different persistence also biases inference on the impact of other factors on happiness, in particular employment. We show that once one controls for the differential effects of persistent and transitory income shocks, employment per se no longer contributes to a household's well being. These findings give support to two important assumptions in modern macroeconomics: insurance markets are incomplete and households suffer from a dis-utility of work.

References

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Arulampalam, W. (2001). Is unemployment really scarring? Effects of unemployment experiences on wages. *The Economic Journal*, 111(475):585–606.

Bayer, C. and Juessen, F. (2009). The life-cycle and the business-cycle of wage risk: A cross-country comparison. IZA Discussion Papers 4402, Institute for the Study of Labor (IZA).

Bewley, T. (1980). Interest bearing money and the equilibrium stock of capital. *mimeo.*

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *The American Economic Review*, 98(5):1887–1921.

Caporale, G., Georgellis, Y., Tsitsianis, N., and Yin, Y. (2009). Income and happiness across Europe: Do reference values matter? *Journal of Economic Psychology*, 30(1):42–51.

Clark, A. and Delta (2003). Unemployment as a social norm: Psychological evidence from panel data. *Journal of Labor Economics*, 21(2):323–351.

Clark, A., Frijters, P., and Shields, M. (2008). Relative income, happiness, and utility: An explanation for the Easterlin paradox and other puzzles. *Journal of Economic Literature*, 46(1):95–144.

Clark, A., Georgellis, Y., and Sanfey, P. (2001). Scarring: the psychological impact of past unemployment. *Economica*, 68(270):221–241.

Clark, A. and Oswald, A. (1994). Unhappiness and unemployment. *The Economic Journal*, 104(424):648–659.

Dehejia, R., DeLeire, T., and Luttmer, E. F. P. (2007). Insuring consumption and happiness through religious organizations. *Journal of Public Economics*, 91(1-2):259–279.

Ferrer-i Carbonell, A. and Frijters, P. (2004). How Important is Methodology for the estimates of the determinants of Happiness?*. *The Economic Journal*, 114(497):641–659.

Frey, B. and Stutzer, A. (2002). What can economists learn from happiness research? *Journal of Economic literature*, 40(2):402–435.

Frijters, P., Geishecker, I., Haisken-DeNew, J., and Shields, M. (2006). Can the Large Swings in Russian Life Satisfaction be Explained by Ups and Downs in Real Incomes?*. *The Scandinavian Journal of Economics*, 108(3):433–458.

Frijters, P., Haisken-DeNew, J., and Shields, M. (2004a). Investigating the patterns and determinants of life satisfaction in Germany following reunification. *Journal of Human Resources*, 39(3):649–674.

Frijters, P., Haisken-DeNew, J., and Shields, M. (2004b). Money does matter! Evidence from increasing real income and life satisfaction in East Germany following reunification. *The American Economic Review*, 94(3):730–740.

Graham, C. and Pettinato, S. (2002). *Happiness and hardship: Opportunity and insecurity in new market economies*. Brookings Inst Pr.

Hansen, G. (1985). Indivisible labor and the business cycle. *Journal of monetary Economics*, 16(3):309–327.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969.

Kaplan, G. and Violante, G. L. (2010). How Much Consumption Insurance beyond Self-Insurance? *American Economic Journal: Macroeconomics*, 2(4):53–87.

Lelkes, O. (2006). Tasting freedom: Happiness, religion and economic transition. *Journal of Economic Behavior & Organization*, 59(2):173–194.

van Praag, B. M. S. and Baarsma, B. E. (2005). Using happiness surveys to value intangibles: The case of airport noise. *Economic Journal*, 115(500):224–246.

Winkelmann, L. and Winkelmann, R. (1998). Why are the Unemployed so Unhappy. *Economica*, 65:1–15.

A Moment restrictions

To derive (6), take first-differences of (5),

$$\Delta u_{it}^* = \alpha_x \Delta x_{it} + \alpha_\psi \Delta \psi_{it} + \Delta \xi_{it}, \tag{14}$$

and replace $\Delta x_{it} = \Delta y^*_{it} - \Delta\psi_{it}$ (see 2)

$$\begin{aligned}\Delta u^*_{it} &= \alpha_x \left[\Delta y^*_{it} - \Delta\psi_{it}\right] + \alpha_\psi \Delta\psi_{it} + \Delta\xi_{it} \\ &= \alpha_x \Delta y^*_{it} - \left(\alpha_x - \alpha_\psi\right) \Delta\psi_{it} + \Delta\xi_{it}.\end{aligned}$$

Analogously, to derive (7), replace $\Delta\psi_{it} = \Delta y^*_{it} - \Delta x_{it} = \Delta y^*_{it} - \epsilon_{it}$ (for $\rho = 1$) in (14)

$$\begin{aligned}\Delta u^*_{it} &= \alpha_x \epsilon_{it} + \alpha_\psi \left[\Delta y^*_{it} - \epsilon_{it}\right] + \Delta\xi_{it}, \\ &= \alpha_\psi \Delta y^*_{it} + \left(\alpha_x - \alpha_\psi\right) \epsilon_{it} + \Delta\xi_{it}.\end{aligned}$$

We thus derived for latent utility the following estimation equations:

$$\Delta u^*_{it} = \alpha_x \Delta y^*_{it} - (\alpha_x - \alpha_\psi)\Delta\psi_{it} + \Delta\xi_i t, \tag{15}$$

$$\Delta u^*_{it} = \alpha_\psi \Delta y^*_{it} + (\alpha_x - \alpha_\psi)\epsilon_{it} + \Delta\xi_i t. \tag{16}$$

As

$$y^*_{it+1} - y^*_{it-2} = \psi_{it+1} - \psi_{it-2} + \epsilon_{t-1} + \epsilon_t + \epsilon_{t+1} \tag{17}$$

$$y^*_{it+1} - y^*_{it} = \psi_{it+1} - \psi_{it} + \epsilon_{t+1}, \tag{18}$$

the former is not correlated with the residual $(\alpha_x - \alpha_\psi)\Delta\psi_{it} + \Delta\xi_{it}$ from (15) while the latter is not correlated with the residual $(\alpha_x - \alpha_\psi)\epsilon_{it} + \Delta\xi_{it}$ from (16). At the same time, since

$$y^*_{it} - y^*_{it-1} = \psi_{it} - \psi_{it-1} + \epsilon_t \tag{19}$$

both instruments correlate with $y^*_{it} - y^*_{it-1}$.

B Similarity of second step OLS and single step ordered probit estimators

We check the robustness of our two-step estimation procedure by comparing two-step OLS estimates on income and a standard single-step ordered probit regression. As Table 7 shows running a two-step estimation procedure instead of a single-step one does not significantly change results.

Table 7: Similarity of Two-step and One-step Estimation

	i OLS	ii O-probit
Income	.316 (.009)	.325 (.008)

Standard errors in parentheses. OLS refers to the two-step estimation described in the main text, where we first estimate a model for happiness using ordered probit and a model for income using an OLS estimator using the same set of control variables in both regressions. We then generate residuals that we regress on each other linearly. O-probit refers to a single step ordered-probit estimation that includes income along with the control variables. The table reports only the coefficient estimate on income.

C Variables and their keys in the SOEP data

Variable	Key
Overall life satisfaction	p11101
HH post-government income	i11102
Employment status of individual	e11102
Annual work hours of individual	e11101
Relation to HH head	d11105
Number of persons in HH	d11106
Number of children in HH	d11107
Number of years of education	d11109
Indicator - wife/spouse in HH	h11112
Subjective satisfaction with health	m11125
Disability status	m11124
Age of individual	d11101
Marital status	d11104

The variables are from the 100%-sample version of the Cross-National Equivalent File of the SOEP ($PEQUIV-files).