

NCER Working Paper Series

The mystery of the U-shaped relationship between happiness and age.

Paul Frijters and Tony Beatton

Working Paper #26

July 2008

The mystery of the U-shaped relationship between happiness and age.

Paul Frijters and Tony Beatton*

January 7, 2008

Abstract

In this paper we address the puzzle of the relation between age and happiness. Whilst the majority of psychologists have concluded there is not much of a relationship at all, the economic literature has unearthed a possible U-shape relationship. In this paper we replicate the U-shape for the German SocioEconomic Panel (GSOEP), and we investigate several possible explanations for it.

JEL-Codes: C23, C25, I31.

Key-Words: Happiness methodology, unobservables, latent variable models, age effects, cohort effects.

*Frijters: School of Economics and Finance, Queensland University of Technology, and RSSS, Australian National University, Canberra, Australia; Tony Beatton: School of Economics and Finance, Queensland University of Technology. We would like to thank conference and staff seminar participants for useful comments and suggestions. The usual disclaimer applies.

1. Introduction

What is the relationship between happiness and age? Do we get more miserable as we get older, or are we perhaps more or less equally happy throughout our lives with only the occasional special event (marriage, birth, promotion, health shock) that temporarily raises or reduces our happiness, or do we actually get happier as life gets on and we learn to be content with what we have?

The answer to this question in the recent economic literature on the subject is that the age-happiness relationship is U-shaped¹. This finding holds for the US, Germany, Britain, Australia, Europe, and apparently even South Africa. The stylised finding is that individuals gradually get unhappier after their 18th birthday, with a dip around 50 followed by a gradual upturn in old age. The predicted effect of age can be quite large, i.e. the difference in average happiness between an 18 year old and a 50 year old can be as much as 1.5 points on a 10 point scale.

In the psychological literature, the consensus opinion is that age matters little. The unimportance of age for happiness was already found in the early cross-country study by Cantril (1965) and has been replicated several times since. Palmore and Luikart (1972) for instance comment in their review that ‘Several variables thought to be related to life satisfaction had little or no relationship: age, sex, total social contacts, ...’. Diener and Suh (1998) give a reason for this by stating that ‘positive and negative affect clearly replicates across age cohorts’ where positive and negative affect are seen as major personality factors involved in having high or low life satisfaction. More recently, Dear et al. (2002) conclude that the prevalence of high life satisfaction simply becomes less common at higher ages. From this reading, it is clear that either the psychologists have overlooked something important for a long time or that the economists have somehow gotten it wrong recently. This paper intends to find out.

¹Recent papers on this in the economic literature include Bell and Blanchflower (2007), Blanchflower and Oswald (2001, 2004, 2007), Clark (2006), Dear et al. (2002), Di Tella et al. (2001), Ferrer and Frijters (2004), Gerdtham and Johannesson (2001), Hayo and Seifert (2003), Helliwell (2003), Oswald (1997), Powdthavee (2005), Seifert (2003), Senik (2004), Theodossiou (1998) and Winkelmann and Winkelmann (1998).

An introduction to the found effects of correlates of happiness can be found in Frey and Stutzer (2002). For a recent general introduction to the economic literature on happiness, see Clark et al. (2008). For a full list of nearly all papers in the field of happiness, see Veenhoven’s Database of Happiness (introduced in Veenhoven et al. 1994).

We re-examine the age-happiness relationship and delve into the methodological aspects of the problem. We essentially want to know if the U-shape that economic scholars find is an artefact or real, and what the actual relationship between age and life satisfaction is.

We re-examine the age-happiness relationship in an often-used dataset, the German Socio Economic Panel which has an extensive set of variables on the individual level. This data-richness allows us to not only replicate the findings of other studies based on cross-sectional data, but furthermore allows us to explore the dynamic interplay between age, covariates, unobserved heterogeneity, and happiness.

The format of the paper is to let the puzzle of the age-happiness relation unfold. We first briefly review the recent literature where we summarise the main findings of others, as well as their methodology. Then we present the data we have and show that we can indeed also generate a U-shape in happiness when we run similar regressions to those in the literature. We then go through successive explanations of the U-shape, including the possibility that it is an artefact due to missing fixed unobservables, that it might be a panel selection effect, that it might be due to the inclusion of particular untrustworthy age ranges, that it might be due to supposed cohort effects, or that the U-shape finding is indeed ‘the truth’. In the conclusions we summarise our findings and what they mean for future research on the economics of happiness.

One may wonder what the age-happiness relation has to do with economics. One main area in which it matters to economic decision making is in the utilitarian calculus of the benefits of living longer and of keeping individuals at various ages alive for longer. If it were truly the case that the very old are happier than the middle-aged, as the current status-quo would have one believe, then an additional year of life of a very old person is worth more to a utilitarian than an additional year of life of a middle-aged person, and health costs should reflect this. If it is alternatively the case that the very old are unhappier than the middle-aged, then the utilitarian maximiser would think an additional year of a middle-aged person to be worth more than that of a very old person and would let health-care decisions reflect this.

1.1. Literature review

Whilst a lot of the economic literature on the age-happiness relation is recent, there have been earlier discussions of it (see Theodossiou 1998 for a discussion

of the history of this issue). Up till the early 2000s, the opinion of economists about the effect of age was still divided. Clark and Oswald (1994) found a U-shaped pattern for the UK, whilst Winkelmann and Winkelmann (1998) on the other hand found no U-shape in happiness but simply a very strong negative effect of age. Easterlin and Schaeffer (1999), using 20 years of the US General Social Survey even concluded that life satisfaction is almost flat in age, with neither a U-shape nor a negative slope. Alesina et al. (2001) and Van Praag et al. (2000) even found an inverted U-shape.

Despite this early controversy, nearly all recent papers come down on the side of a U-shape relation between happiness and age. Blanchflower and Oswald (2001, 2004) simply state that ‘Wellbeing is U-shaped in age’. Gerdtham and Johannesson (2001) also report a U-shape in age with a minimum around the age of 55. Hayo and Seifert (2003) and Seifert (2003) also report a U-shape and call the U-shaped age effect a ‘typical finding in happiness regressions’.

The most comprehensive study to date is Blanchflower and Oswald (2007) who combine cross-sectional data for the US, Europe, and the World Value Survey. In total, they have about 800,000 respondents in over 60 countries for which they all report a U-shape in happiness and age. Clark (2006) claims some robustness with respect to methodology for this finding when he concludes that ‘Panel analysis controlling for fixed effects continues to produce a U-shaped relationship between well-being and age’.

In order to get a feeling for the role of methodology in these findings, we reproduce in the next Table the main findings of the recent economic studies on the U-shape between age and happiness. We show the found coefficients on age and age-squared and detail the source of the data and the estimation method. We may mention already that all the studies included in this table also use other personal variables in the same regression. The controls mainly include measures for employment, income, partnerships, the number of children, education, and sometimes indicators of where someone lives.

Table 1: Results from Life Satisfaction regressions
(t-values) from recent economic studies

Author, date	Sample (size & name)	Coefficients - Pooled (t-value)		Coefficients - Fixed Effects (t-value)		Dependent variable and controls
		Age	Age Squared	Age	Age Squared	
(Blanchflower & Oswald, 2007)	USA: General Social Survey 1972-2006	Ordered Logit (men+women averaged) USA: -0.0211 (4.39) <i>N</i> = 45,474	Ordered Logit (men+women averaged) USA: 0.0003 (5.92) <i>N</i> = 45,474			Dep: Happiness Controls: yes (specification without cohort)
(Blanchflower & Oswald, 2007)	Europe: Eurobarometer 1976-2002	Ordered Logit (men+women averaged) Eur: -0.045 (31.31) <i>N</i> = 589,446	Ordered Logit (men+women averaged) Eur: 0.00052 (10.1) <i>N</i> = 589,446			Dep: Life Satisfaction Controls: yes (specification without cohort)
(Blanchflower & Oswald, 2007)	World Value Survey 1981- 2004	Ordered Logit (men+women averaged) WVS: -0.0505 (10.1) <i>N</i> = 163,852	Ordered Logit (men+women averaged) WVS: 0.0003 (5.92) <i>N</i> = 163,852			Dep: Life Satisfaction Controls: yes (specification without cohort)
(Blanchflower & Oswald, 2004)	UK: Eurobarometer Survey 1975-1998	Ordered Logits - All UK: -0.0424 (2.84) <i>N</i> = 54,549	Ordered Logits - All UK: 0.0005 (15.38) <i>N</i> = 54,549			Dep: Life Satisfaction Controls: yes
(Clark, 2006)	British Household Panel Survey (BHPS) waves 1 to 14	-0.075 (-25) <i>N</i> = 82,096	0.00091 (30.33) <i>N</i> = 82,096	Applied age cohorts to derive fixed effect coefficients		Dep: Life Satisfaction Controls: yes
(Di Tella, MacCulloch, & Oswald, 2001)	Eurobarometer Survey Series 1975-1991	OLS -0.02 (20.0) <i>N</i> = 264,710	OLS 0.0002 (33.33) <i>N</i> = 264,710			Dep: Life Satisfaction Controls: yes
(Powdthavee, 2005)	Statistics South Africa OHS study of 1997	-0.011 (z-stat: -2.38) <i>N</i> = 20,634	0.0001 (z stat: 2.03) <i>N</i> = 20, 634			Dep: Life Satisfaction Controls: yes
(Senik, 2004)	Russian longitudinal monitoring survey (RLMS).	Ordered Probit (2) -0.050 (8.33) <i>N</i> = 17,897	Ordered Probit (2) .001 (p < .01) <i>N</i> = 17,897			Dep: Life Satisfaction Controls: yes
(Winkelmann & Winkelmann, 1998)	German Socio- Economic Panel 1984-89 waves of the GSOEP	-0.098 (-9.8) <i>N</i> = 20,944	0.0012 (12) <i>N</i> = 20,944	fixed effects logit model 2 -0.118 (-3.19) <i>N</i> = 20,944	fixed effects logit model 2 -0.0001 (0.25) <i>N</i> = 20,944	Dep: Life Satisfaction Controls: yes

The table confirms the very strong effect that age is found to have upon life satisfaction in recent studies, and that the effect of linear age is always negative, whilst that of age-squared is positive, indicating a U-shape. Bearing in mind that the age at which the minimum occurs is given by the coefficient of linear age divided by twice the coefficient of age-squared, it indeed appears that the majority

of the studies find an age of around 55 as the age at which the minimum occurs. The Table also underscores that the effects are mainly found in cross-sections when controls are added for individual socio-economic variables.

2. The data

We use the full 1984-2002 waves of the German Socio-Economic Panel (GSOEP), described in Wagner et al. (1993)². We use only the information on West Germany in order to be able to abstract from the importance of the 1990 German reunification, which had a tremendous impact on the lives and satisfaction levels of East Germans (see eg. Frijters et al. 2004). The appendix gives a table with summary statistics.

3. Analysis of the puzzle

3.1. Is there also a U-shape for Germany?

For all analyses that follow it holds that the full regression tables are shown in the appendix but that we tell the story in graphs in the main text. We experimented with using both simple least-squares and latent-variable analyses (for cross-sectional as well as fixed-effects analyses) but we found, as in Ferrer and Frijters (2004), that there is no qualitative difference. The appendix shows the least-square results for they are most closely aligned with the methods used by others, though results for the latent-variable analyses are available on request.

We begin with showing a picture of the raw pooled cross-sectional relationship between age and aggregate happiness for the GSOEP, with the predicted lines overlaid for least-squared regressions that include either just age or age and age-squared. The shown intercepts are normalised such that satisfaction at age 20 is always the same.

²The GSOEP is a longitudinal household survey sponsored by the Deutsche Forschungsgemeinschaft. It is organized by the German Institute for Economic Research (Berlin), and the Center for Demography and Economics of Aging (Syracuse University). We thank these institutes and the director Dr. G. Wagner for making the data available.

Age and Life Satisfaction: pooled regression




Figure 1: life satisfaction in the GSOEP for the pooled sample.

These findings are quite typical of those found in the psychological literature and in the datasets used by other scholars in the economics literature: there is a bit of a happiness decline in the years after 18 and close to death (where there are not many individuals left), but for the age range 20 to 60 there is no strong relation between age and happiness to be seen. If we overlay the regression results when we include linear age as the only variable we get a significantly negative coefficient.

If we overlay a regression line with both age and age-squared, we find a significant U-shaped pattern. The regression results are

$$\begin{aligned}
 LifeSat_i &= 7.747 - 0.0217 * age_i + 0.00016 * age_i^2 \\
 &\quad (249.9) \quad (15.7) \quad (11.6) \\
 R^2 &= 0.0037, N = 176770
 \end{aligned}$$

Age-squared is highly significant, but the age at which the minimum occurs is about 90 with this simple specification, implying that for the vast majority of the sample, there is not so much a U-shape but rather a horizontal j-shape.

What if we now add additional regressors to this simple specification? In the next graphs we show the predicted age-happiness profiles when we successively add additional variables. The exact specifications are in Table 3 of the Appendix.




Figure 2: life satisfaction in the GSOEP for the pooled sample with added controls.

In the first added line, called ‘the usual suspect’ we include the socio-economic variables commonly found in the happiness regressions of Table 1. They are log-income, gender, education in years, the number of kids, a marriage dummy, and 3 indicators of work-status (employed, non-participant and unemployed). We see a dramatic deepening of the U-shape, with the predicted happiness decline from 18 to 70 year old being about 1.6 for Germany. We also see the upswing occurring earlier, i.e. at age 75.

When we also include indicators of health and measures of wealth, we find the strongest U-shape yet, with the predicted happiness decline from 18 to 70 year old being about 1.8 for Germany. When we finally throw in a large set of indicators of life events (including the loss of a spouse, being fired, and birth of a child), the age at which the minimum occurs becomes earlier and the U-shape becomes less deep but it is still strongly significant. The main regression with covariates we will talk about in the remainder is the specification with the socio-economic variables and health because it has the strongest U-shape.

The graphs above are perfectly in line with the findings of other scholars in the field: when standard regressors are added, a very strong U-shape effect emerges with predicted age effects far bigger than anything observable in the raw data. The puzzle is thus indeed in the data we have available, which makes it likely that an explanation of this puzzle for our data will carry over to the datasets used by others. We next address a variety of possible explanations.

3.2. Potential explanation I: it is all about the very young and the very old

A naive first-thought when looking at the initial graphs is that there is a particular issue with the early ages, i.e. age 18 to 22, and with high ages, i.e. those above 80. This is because the happiness decline is particularly steep for the early years and erratic at the later years, which makes one wonder if the young are being overly optimistic about their actual levels of happiness and that the the happiness of the very old is hard to tell from the few data points in that range.

To examine this possibility, the graph below show the results of the regressions when we simply drop the under-22 year olds, as well as those over 80 from the data (about 9.3% of the panel). The exact specifications are in Table 4 of the appendix.




Figure 3: life satisfaction in the GSOEP for the pooled sample for the mid-age range.

Figure 3 shows that there is actually a strengthening of the U-shape when we only focus on the middle range in age: predicted life satisfaction at the minimum is now around 5.5 whilst it was around 6 for the whole sample. There is no clear qualitative difference between the results however. Hence the U-shape cannot be explained by the extremities of the age range and must be due to relations in large parts of the age range.

3.3. Potential explanation 2: it is all about reverse causality

An important finding in the literature so far is that happiness is strongly affected by stable personality traits (see Frey and Stutzer 2002, Ferrer and Frijters 2004, and Argyle 1999). These are fixed individual traits that are usually part of the error term. A stylised finding from both the economic and the psychological literature is that accounting for fixed traits has a very strong impact on the coefficients found for socio-economic variables (see Ferrer and Frijters 2004, but also the review of Clark et al. 2008). A leading explanation for this is the possibility of reverse causality, i.e. that personality traits that make you happier also make it more likely that you will have a high income, a job, a partner, good health, high wealth, and a high education.

Could the problem of reverse causality caused by unobserved fixed traits explain something about the U-shape? At first glance one would think not because fixed personality traits are by design uncorrelated with age. However, personality traits can be correlated with variables that are correlated with age, such as income, a job, a partner, good health and wealth. How would this work?

Consider the problem in its simplest form. Suppose for the purposes of this subsection the truth is that the following relationship holds

$$y_{it} = \alpha_1 * age_{it}^2 + x_{it}\beta + f_i + u_{it}$$

$$f_i \perp age_{it}^2, cov(f_i, x_{it}) > 0, cov(age_{it}, x_{it}) > 0. E[u_{it}|age_{it}, x_{it}, f_i] = 0$$

where we have for simplicity dispensed with a linear age term and all variables are normalised to have expectation 0 implying there is no constant term either; there are individual fixed traits f_i unrelated to age-squared but related to a composite time-varying socio-economic variable called x_{it} . There is an error term u_{it} orthogonal to everything else. What are now the estimated coefficients if we mistakenly run a regression without accounting for fixed-effects? The asymptotic values are

$$p \lim \hat{\beta} = \beta + \frac{cov(f_i, x_{it})}{var(x_{it})} + (\alpha_1 - p \lim \hat{\alpha}_1) \frac{cov(age_{it}^2, x_{it})}{var(x_{it})}$$

$$p \lim \hat{\alpha}_1 = \alpha_1 + (\beta - p \lim \hat{\beta}) \frac{cov(x_{it}, age_{it}^2)}{var(age_{it})}$$

which shows that even though age_{it}^2 is not correlated with the omitted fixed-effect, the coefficient on age_{it}^2 can nevertheless be biased when it is related to

included time-varying variables that are correlated with the omitted fixed-effect. The equations become rather elaborate if we add a linear age term and a constant but the basic principle remains that a bias in the age-term can occur if the added variables are correlated with age and with the omitted fixed-effect.

To explore this possibility we run simple fixed-effect analyses to see how this changes the found U-shape. The following graph shows what a fixed-effect regression begets. The exact specifications are in Table 5 of the appendix.




Figure 4: life satisfaction in the GSOEP for the balanced panel.

The results for this graph are both confirming and surprising. The graph shows the raw relation between age and happiness and has overlaid three lines. The U-shaped line is the same one as we showed previously and is the pooled regression with the preferred specification including health and wealth. Overlaid are two lines from fixed-effect regressions. The dashed line is the result of running the same regression as for the pooled regression but including fixed effects. As one can see, the U-shape completely disappears, i.e. the age-squared coefficient becomes tiny and insignificant. It however replaces the U-shape by a similarly puzzling effect, which is a very strongly significant negative linear relation. The third line, which shows the result of just running a fixed effect regression with only age and age-squared as regressors confirms this. The U-shape actually slightly reverses into an inverted U shape, but a very strong negative relation emerges.

Before we turn to explain the new puzzle of the strong negative linear relation, we first want to confirm that the disappearance of the U-shape is indeed because of reverse causality. For one, the table in the appendix re-iterates the finding of many other studies (reviewed in Ferrer 2005) that the coefficients of most socio-

economic variables become much smaller when one adds fixed effects. The income coefficient drops by more than 40% and the importance of marriage, a job, health and education all reduce.

Do these variables correlate with age though? The following graph gives the simple averages by age of the 4 most significant socio-economic variables in the simple regressions with health. The variables have been re-scaled to fit on the same graph.

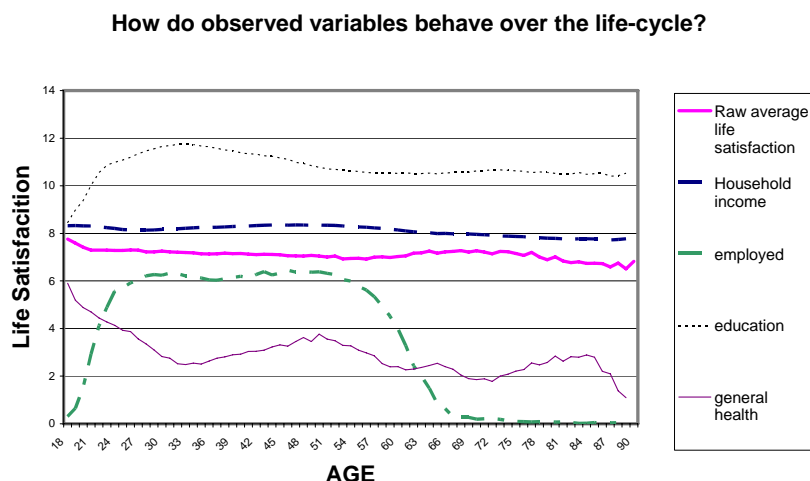


Figure 5: age and observed correlates in the GSOEP.

This graph indeed shows a strong relation between age and employment (reverse U-shaped), education (reverse J-shaped), and health (which has a reverse U-shape in the main age range between 30 and 50). Household income doesn't vary much with age though there is a slight inverted U-shape for the middle age range. These relations all go in the direction we anticipated above: an artificially high coefficient for employment, education, income and health would all give rise to a false U-shape in age.

Whilst we have found a solution to the original puzzle, i.e. the U-shape is an artefact of reverse causality with the included covariates³, it has been succeeded by a possibly even greater puzzle. This new puzzle is the strong negative relation between age and happiness over time in fixed-effects regressions even though the relation in the pooled cross-sectional data is much less pronounced. Where does

³If there are adaptation effects (habit formation), such as hypothesised by Easterlin (2001), then even the found coefficients from fixed-effects regressions (which are based on short-run effects) are overestimates of the true long-run effect of covariates.

this strong slope come from? We have seen that it is not due to any other included variable because it remains when we don't include other variables. We know for the same reason that it is not due to reverse causality. It also cannot be due to some simple missing variable, like average income, because that would go in the opposite direction (average income rose). What candidate explanations then remain?

3.4. Explanation for the negative slope I: time and cohort effects

A popular explanation in the recent literature is that there are important time and cohort effects on happiness (e.g. Cribier 2005). Right at the outset, one should say that such explanations are a little unsatisfactory because both time and cohort effects are in a sense 'aggregate unobservables'. For instance, a cohort effect is just a missing aggregate variable specific to an age-group but where we don't know what the missing variable is. The missing variable could be the mental experience of a particular war, or the effect of a particular diet in a certain era, or some cultural trait particular to an era (like expectations), etc. It would be preferable to measure the supposed elements making up a cohort effect before becoming convinced cohort effects actually exist. We also need to know what the cohort effects consist of if we are to make any policy-relevant inferences about how 'happy' cohort effects might be created.

Another problematic aspect of the notion of cohort-effects and time-effects is that they are statistically hard to identify. It has been known for a long time that it is not possible to simultaneously identify both age, year, and cohort effects simply because one can write one as a function of the other two. Only by introducing somewhat arbitrary restrictions on cohort effects can they be separated from age and time-effects.

What one can do is to dispense with age effects and simply presume that the effects of age pick up cohort effects. One can then divide up the age distribution into different cohorts and label the effect of being born in a particular interval as due to a cohort effect. Blanchflower and Oswald (2007) for instance define cohorts by age-intervals, i.e. 10-year intervals, and label the effects of these age-intervals as cohorts. This practically means that there is a sharp dividing line in the influence of whatever causes a cohort effect on particular days in the century. Someone born on January 5 1910 for instance could be in one cohort, and someone born the day after in another. If one does not include age as a regressor, this procedure is akin to using a semi-parametric function of age (e.g. Clark 2006). However, if one

then also adds linear and quadratic age variables to a regression containing these age-intervals and proceeds with assuming that the age-intervals pick up something very different from age effects, then these arbitrary age-intervals become binding: the assumption that cohort effects jump at particular ages is then what separately identifies age effects from time and cohort effects.

So far in this paper we have ignored the possibility of cohort or time effects. Insofar as there are linear time effects, then these would be indistinguishable in a fixed-effect framework from age effects. This is because we can write $age_{it} = age_{i0} + t$ where the first term (age_{i0}) is fixed and absorbed by the fixed-effect term, whilst the second term (t) is a straightforward time effect.

In order to ascertain whether there are likely to be time or cohort effects, we next look at the evolution over time of aggregate life satisfaction of the whole pooled panel. If there are strong time or cohort effects capable of explaining the large decline that we saw under the fixed-effects regressions, then we should see such a decline in the aggregate data.




Figure 6: year and life satisfaction in the GSOEP for the pooled sample.

The graph above shows the evolution over time for the GSOEP. As we can see, there is indeed a strong decline. Bearing in mind that the standard deviation of the mean life satisfaction in a year with so many observations (10,000 per year) is less than 0.02, the year on year changes are highly significant. The overall decline also fits somewhat, though it is not quite enough: in 19 years, the aggregate drop is only 0.4 whilst the drop predicted by the fixed-effect regression is about 0.6 for 19 years.

There might hence be cohort or time effects⁴ responsible for the found drop by age, but the predicted future drop in satisfaction would be enormous, i.e. at current trends a predicted drop by more than 1 point in the next 30 years. Also, this predicted drop doesn't tally with what we know from other surveys (like the Eurobarometer Survey) where it was found that aggregate life satisfaction is quite constant in Western countries over time, including Germany⁵. There is hence still something not quite right about this 'explanation' because whilst it fits the GSOEP as a whole, it seems to violate what we know to hold at the aggregate for happiness cross-sections. What other explanation is left?

3.5. Explanation for the negative slope II: there is something wrong with the panel data

Forced to reject or seriously doubt all other reasonable explanations, we now turn to the most uncomfortable potential explanation, which is that there is something wrong with repeat happiness responses in this panel data. What if individuals who remain in the panel and who keep answering the survey are 'different' from those who drop out of the panel or who don't keep answering the satisfaction question?

The GSOEP has large numbers of drop-outs each year. For instance, of the roughly 10,000 individuals in the original 1984 GSOEP sample, only about 4,000 remain in our data for the full 19 years. The GSOEP replaces those who no longer answer with new respondents, based on a desire to keep a representative sample in terms of variables included in a census, such as gender, age, income, and education.

What kind of selection could cause the large decline seen in the fixed-effect regressions and in the aggregate data shown above? A naive thought would be that only those who are unhappy keep answering the GSOEP, i.e. a selection based on the fixed effects f_i . This is not a valid possibility however, because fixed-effects drop out in the fixed-effect specification, making it irrelevant whether there is a selection on fixed traits.

A second naive thought is that there could be a selection on particularly unpleasant observed life events, i.e. only those with bad events happening to them

⁴Time and cohort effects cant be meaningfully separated if we also include non-linear age effects.

⁵Clark et al. (2008) show that life-satisfaction profiles have been virtually flat in cross-sections in the last 30 years for Germany, France the US, and many other Western countries.

keep answering the panel. Whilst there is limited support for this in the graphs above (when adding some life events, the predicted life satisfaction decline reduces), there is no full support for this: as far as observed life events are concerned, the negative age effect remains when they are taken into account.

A third naive thought is that there could be a selection on transitory unobserved negative shocks. If this were the case though, then there should only be a one-off drop in life satisfaction and not the sustained decline we see from the fixed-effect regression.

The most problematic thought is that there could be a selection on unobserved strongly persistent negative shocks. This implies selection on a strongly persistent (but not fixed) part of u_{it} . If we are for instance thinking of writing $u_{it} = \sum \rho^t e_{it}$ with e_{it} being i.i.d. shocks, then a high ρ (close to 1) would indicate a strong persistence in shocks and the selection we would worry about is on e_{it} . If, wave after wave, it is the case that individuals are more likely to stay another year in the GSOEP when e_{it} is lower, then we would indeed be able to get the strong negative age slope observed in the fixed-effects regression.

How can we verify this possibility? We verify it by comparing the answers of those who stay in the panel with those who enter for the first time. Every year, there are several thousand new entrants in the GSOEP who have never answered the questionnaire before. Some of these are new samples, some are partners of regular respondents, and some are children becoming old enough to be in the sample. What is the relation between age and life satisfaction for them and what does aggregate satisfaction over the years look like for them?

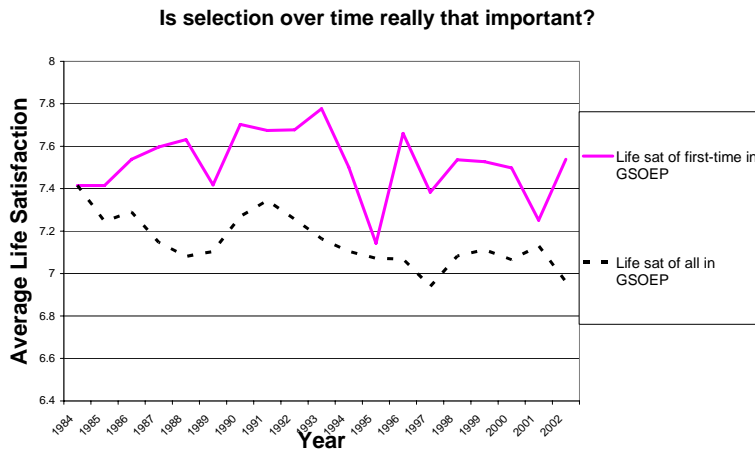



Figure 7a and 7b: life satisfaction in the GSOEP for first-time respondents.

Figure 7a (top) shows the aggregate level of satisfaction by age of those who answer for the first time, as well as for the entire pool. As we can see, there is still a decline in happiness at a very young age and at very old ages (which has very few observations and therefore looks erratic), but there is no decline at all from age 25 to 75. If we then overlay the predicted regression lines, then the predicted regression line for the whole sample shows the horizontal j-shape, but the predicted regression line for the first-time panel entrants is almost flat. Not quite flat, because there is still a significantly negative age trend, but the coefficient is about 85% smaller than that for the full sample; when adding a quadratic term all age effects become insignificant for the new entrants (the exact specifications are in Table 6 of the appendix). Hence there is a small age effect on happiness,

concentrated at the very young and the very old, but it is not U-shaped. Rather, it is simply a decline.

Figure 7b (bottom) confirms the impression of the top graph: when looking at the average satisfaction over the years for first-time respondents, there suddenly is no time profile of happiness anymore, in line with what is found for cross-sectional studies which by design only question individuals once. This is consistent with the notion that there are no age, time, or cohort effects, or that they at least miraculously cancel each other out. The decline in the aggregate panel is thus indeed most likely due to selection on time-varying unobservables.

How bad is the decline in satisfaction for repeat respondents? As a final graph we show aggregate satisfaction depending on the number of years someone has answered the questionnaire




Figure 8: the degree of selection in the GSOEP for stayers in the panel.

The graph indeed confirms the running hypothesis above: there is a truly large decline in reported satisfaction as an individual is in the panel for longer. We now know this is not age related or related to observables, but this is a selection on unobservables, and given that there is no discernible ‘bounce-back’, it furthermore has to be a relatively persistent time-varying unobservable responsible for this decline. The decline is indeed large enough to explain the fixed-effect profile: in 18 years, there is a 0.64 reduction in life satisfaction, which translates to a decline of 2.2 over 60 years, almost exactly the predicted amount from the fixed-effect regression.

3.5.1. Is there simple corroborating information of a dynamic selection on negative shocks?

A natural question to ask is whether there is any information outside the panels themselves that can be used to verify if the panel is indeed retaining the ‘unfortunate’, i.e. those who have experienced unobserved negative shocks. What we can look at to verify such a possibility is to look at the observed negative happiness-relevant shocks that are observed in the panel for which we can find some official outside statistics to check them against.

For the vast majority of the time-varying happiness-related variables in the panel, there is no reliable outside information to check the panel’s selection with. The one negative shock for which there is some information is divorce rates. According to UN statistics, the yearly divorce rate (official annulments of registered marriages) per 1000 in the population is 5.2 in Germany⁶. If we take into account that the panel does not contain individuals below 18 (the 0-17 year olds make up some 18% in Germany), then this would imply that divorce rates should be 6.2 per 1000 in the GSOEP.

As it is, the yearly self-reported divorce rate in the GSOEP is about right, i.e. about 6 per 1000 on average and rising over time in the panel. Hence self-reported divorce rates are not higher for the sample than for the population which means we cannot claim firm outside evidence of negative dynamic selection.

3.6. Alternative interpretations?

Can we think of alternative interpretations for the findings above not based on selection on time-varying unobservables, which is the worst kind of selection possible in panel analyses because there is not much that can be done about it and one relies on a leap of faith that it doesn’t affect the coefficients of time-varying observables?

One alternative is that we are not looking at unobservable negative errors at all, but rather that we are looking at the disappearance of positive errors. It might be the case that Germans over time become more comfortable talking about their levels of happiness and other personal matters because they have not noticed any breach of privacy or other adverse effects of answering the questionnaire. Singer and von Thurn (1995) suggest these are important factors in the quality

⁶The UN numbers refer to the number of granted annulments. Since one annulment affects the marital status of 2 people, we have doubled this to arrive at the number of individuals getting divorced.

of responses. Hence perhaps the GSOEP respondents are becoming more truthful by not glossing over their actual level of happiness as much as they do in the earlier waves. This possibility ties in with the notion of panel learning of which Juster (1986, page 401) observes that ‘later interview waves appear to have higher quality than earlier ones’.

The weak point in this explanation is that one would think the process of getting comfortable with responding levels off after a few years. This would imply that there would eventually have to be some ‘bottoming out’ of the happiness decline for stayers in the panel. This is not really noticeable in the graph above.

Whilst the interpretation that we are looking at more truthful responses over time is probably more soothing to the collectors of the GSOEP because it takes away the suspicion that the representativeness of the panel is in doubt, the implications for the happiness literature of this second interpretation is worse than the first interpretation. If it were just the case that panels suffer from non-random attrition based on time-varying unobservable happiness determinants then there is still a lot of useful information in cross-sectional surveys and one could have the hope that the selection does not involve cross-terms between observable time-varying variables and unobservable ones, allowing useful interpretations of panel coefficients.

If it is alternatively the case that we cannot trust the responses on happiness questions of the first 19 years of responses, then we effectively cannot trust over 99% of the data in this field. And the changes over time are big, certainly big enough to get seriously worried: an overstatement of happiness by 0.7 could drop a country for instance from being one of the happiest country in the world to being one of the unhappier countries in the world, for instance nullifying the validity of all the rank-tables. Now, of course, if the overstatement is a ‘universal constant’ and doesn’t differ by country, then the implications are less strong but the hypothesis of a universal constant overstatement can’t be claimed a priori. If a possible tendency to overstate is furthermore correlated with observables, then cross-sectional and panel analyses all become highly suspect.

4. Conclusions and discussion.

This paper started out with the puzzling findings of other researchers of a U-shaped relationship between age and happiness. We replicated this relationship for Germany using a well-known panel, the GSOEP. We showed that the age decline is in the raw data and the U-shape becomes most pronounced when adding

commonly used socio-economic variables.

We firstly showed that the inclusion of the usual socio-economic variables in a cross-section leads to a U-shape in age that results from indirectly-age-related reverse causality. Putting it simply: good things, like getting a job and getting married, appear to happen to middle aged individuals who were already happy. This reverse causality shows up in cross-sections as inflated coefficients for income, marriage, and getting a job. In order to fit the actual age profile of happiness, the bias in coefficients for socio-economic variables forces the predicted age profile to become U-shaped. When one controls for fixed-effects, the non-linearities all but disappear. This explained the U-shape puzzle but replaced it with an even more glaring anomaly, i.e. the exceedingly large negative linear effect of age on happiness in fixed-effect regressions.

In trying to explain the new puzzle of the negative age slope, we were able to discount the likelihood of cohort and time effects: there simply is no sufficiently strong time profile in the aggregate responses to explain the fixed-effect results and indeed there is no time trend at all for first-time responses.

This left sample selection as the reason for the anomalously high age slope. We confirmed that those who answer for the first-time showed a 85% smaller age-profile in happiness. There indeed is a small decrease in happiness after age 18 and after about age 80, but no significant change between the ages of 25 and 75 for first-time responses. On the other hand, those who stayed in the panel continuously reported lower levels of life satisfaction with the decline being 0.64 for those in the GSOEP for 19 years. This perfectly fit the fixed-effects regressions.

The two possible interpretations of this are that there is either a selection of stayers on somewhat persistent time-varying happiness unobservables or that those who keep answering the questionnaire become progressively more honest about their actual, lower, level of satisfaction. If there is a selection on unobservables then this is problematic for the reliability of long panels and analyses based on them. If individuals are untruthful for the first 19 years of responding to questionnaires then this more or less affects all the analyses in the field, especially cross-sectional studies which make up the bulk of the literature. The exact time profile of the happiness of the stayers (a virtually continuous decline) slightly favours the notion of selection on time-varying unobservables because one would have expected the effect of becoming open and honest to the interviewer to gradually level off before 19 years.

Can we think of a reasonable third alternative explanation which is neither damning for the collection of long representative panels or for the happiness field

as a whole? We, frankly speaking, cannot think of one. The found effect of age in fixed-effect regressions is simply too large and too out of line with everything else we know to be believable. The difference between first-time respondents and stayers and between the number of years someone stays in the panel doesn't allow for explanations based on fixed traits or observables. There has to be either a problem on the left-hand side (i.e. the measurement of happiness over the life of a panel) or on the right-hand side (selection on time-varying unobservables). Future research will have to investigate whether similar problems arise with other variables (such as, perhaps, health) and whether this issue is of similar importance in other panel data sets.

References

1. Alesina, A., Di Tella, R., MacCulloch, R. (2001), 'Inequality and happiness: are Europeans and Americans different?', *NBER* 8198.
2. Argyle, M. (1999). "Causes and correlates of happiness:", in: Kahneman, D., Diener, E. and Schwarz N. (eds.). *Well-Being: The Foundations of Hedonic Psychology*. Russell Sage Foundation, New York. Chapter 18. Beckett, S., Gould, W., Lillard, L., & Welch, F. (1983). Attrition From the PSID. Santa Monica, CA: Union Research Corporation.
3. Bell, D., & Blanchflower, D. G. (2007). The Scots may be brave but they are neither healthy nor happy. *Scottish Journal of Political Economy*, 54(2), 166-194.
4. Blanchflower, D. G., & Oswald, A. J. (2004). Well-being over time in Britain and the USA. *Journal of Public Economics*, 88(7/8), 1359.
5. Blanchflower, D. G., & Oswald, A. J. (2007). Is Well-Being U-shaped over the life cycle?, IZA Discussion Paper 3075, forthcoming in *Social Science and Medicine*.
6. Cantril, H. (1965), *The pattern of human concerns*, New Brunswick, Rutgers University Press.
7. Clark, A. E. (2006). Born to be mild? Cohort effects don't explain why well-being is U-shaped in age (Vol. *Working Paper* N° 2006 - 35): Centre National de la Recherche Scientifique – École des Hautes Études en Sciences Sociales, École Nationale des Ponts et Chaussées – École Normale Supérieure.
8. Clark, A., Frijters, P., Shields, M.A. (2008), 'A survey of the income happiness gradient', forthcoming in the *Journal of Economic Literature*, also IZA, NCER, and DELTA discussion paper.
9. Cribier, F. (2005), 'Changes in the experience of life between two cohorts of Parisian pensioners born in circa 1907 and 1921', *Ageing and Society*, pp. 1-18.

10. Dear, K., Henderson, S., & Korten, A. (2002). Well-being in Australia. *Social Psychiatry and Psychiatric Epidemiology*, 37(11), 503-509.
11. Di Tella, R., MacCulloch, R. J., & Oswald, A. J. (2001). Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness. *American Economic Review*, 91(1), 335-341.
12. Diener, E., & Suh, M. E. (1998). Subjective well-being and age: An international analysis. In K. W. Schaie & M. P. Lawton (Eds.), *Annual review of gerontology and geriatrics*, Vol. 17: Focus on emotion and adult development. (pp. 304-324). New York, NY, US: Springer Publishing Co.
13. Easterlin, R.A. (2001). Income and happiness: Towards a unified theory. *Economic Journal*, vol. 111, pp. 465-484
14. Easterlin, R.A., and Schaeffer (1999). Income and subjective well-being over the life-cycle. In C.D. Ryff & Marshall, V.W. (eds), *The Self and Society in Ageing processes*. New York: Springer.
15. Ferrer-i-Carbonell, A. (2005). Income and well-being: An empirical analysis of the comparison income effect. *Journal of Public Economics*, vol. 89, pp. 997-1019.
16. Ferrer-i-Carbonel, A., Frijters, P. (2004), 'The effect of methodology on the determinants of happiness', *Economic Journal* 114, 641-659
17. Frey, B.S. and Stutzer, A. (2002). What can economists learn from happiness research? *Journal of Economic Literature*, vol. 40, pp. 402-435.
18. Frijters, P., Shields, M.A. and Haisken-DeNew, J.P. (2004a). Money does matter! Evidence from increasing real incomes in East Germany following reunification. *American Economic Review*, vol. 94, pp. 730-741.
19. Gerdtham, U.-G., & Johannesson, M. (2001). The relationship between happiness, health, and social economic factors: Results based on Swedish microdata. *Journal of Socio-Economics*, 30(6), 553-557.
20. Hayo, B., & Seifert, W. (2003). Subjective economic well-being in Eastern Europe. *Journal of Economic Psychology*, 24(3), 329.

21. Helliwell, J. F. (2003). How's life? Combining individual and national variables to explain subjective well-being. *Economic Modelling*, 20(2), 331.
22. Juster, F. T. (1986). Response Errors in the Measurement of Time Use. *Journal of the American Statistical Association*, 81.
23. Oswald, A. J. (1997). Happiness and Economic Performance. *Economic Journal*, 107(445), 1815-1831.
24. Palmore, E., & Luikart, C. (1972). Health and social factors related to life satisfaction. *Journal of Health and Social Behavior*, 13(1), 68-80.
25. Plug, E.J.S. (1997), *Leyden welfare and beyond*, Ph.D. thesis, Thesis Publishers, Amsterdam.
26. Plug, Erik J. S. & van Praag, Bernard M. S. & Hartog, Joop, 1999. "If we knew ability, how would we tax individuals?," *Journal of Public Economics*, Elsevier, vol. 72(2), pages 183-211
27. Powdthavee, N. (2005). Unhappiness and Crime: Evidence from South Africa. *Economica*, 72(3), 531-547.
28. Seifert, W. (2003). Subjective economic well-being in Eastern Europe. *Journal of Economic Psychology*, 24(3), 329-348.
29. Senik, C. (2004). When information dominates comparison: Learning from Russian subjective panel data. *Journal of Public Economics*, 88(9/10), 2099.
30. Singer, E., & von Thurn, D. R. (1995). Confidentiality assurances and response: a quantitative review of experimental literature. *Public Opinion Quarterly*, 59(1), 66-77.
31. Theodossiou, I. (1998). The effects of low-pay and unemployment on psychological well-being: a logistic regression approach. *Journal of Health Economics*, 17(1), 85-104.
32. Van Praag, B.M.S., Frijters, P. and Ferrer-i-Carbonell, A. (2000) 'A structural model of well-being'. *Tinbergen Institute Discussion paper* TI 2000-053/3, the Netherlands.
33. Veenhoven, R., et al. (1994), *World Database of Happiness: correlates of happiness*. Rotterdam: Erasmus University.

34. Wagner, G.G., Burkhauser, R.V., and Behringer, F., (1993). 'The English language public use file of the German Socio-Economic Panel', *Journal of Human Resources*, 28(2): 429-433.
35. Winkelmann, L., & Winkelmann, R. (1998). Why Are the Unemployed So Unhappy? Evidence from Panel Data. *Economica*, 65(257), 1-15.

Appendix A: descriptive tables and regression tables.

The German Socio-Economic Panel (GSOEP) is a representative panel of the German population. The first wave was only in the Federal Republic of Germany in 1980, and includes former East Germany since 1990. It currently tracks about 20.000 individuals and 12.000 households. See Wagner et al. (1993) or Plug (1997) for more detailed descriptions of the data.

Table 2 shows the sample means and standard deviations of all the variables used in the paper. Tables 3 to 5 show the regression results for the 4 groups of respondents looked at, differentiated by specification.

Table 2: Sample averages in the GSOEP;
N = 176,770 (all) and N=18,821 (first-time)

	Entire Sample		First-time Respondents	
	Mean	s.d.	Mean	s.d.
overall life satisfaction (self-assessed)	7.16	1.85	7.45	2.01
age	44.26	16.91	38.35	17.28
age*age	2244.67	1659.88	1769.01	1577.34
ln(household income)	8.20	0.53	8.16	0.51
male (1=yes)	0.49	0.50	0.49	0.50
level of education (years)	10.93	2.46	10.51	2.41
number of children in family	0.65	0.99	0.71	1.01
married (1=yes)	0.65	0.48	0.55	0.50
employed (1=yes)	0.47	0.50	0.44	0.50
non-participant in the labour-force (1=yes)	0.33	0.47	0.37	0.48
unemployed (1=yes)	0.04	0.20	0.04	0.21
average regional income (Euro)	4149.99	477.88	3980.37	332.64
own or purchasing dwelling (1=yes)	0.42	0.49	0.34	0.48
asset income (Euro)	2359.80	10700.08	1258.23	6706.26
imputed rent (Euro)	1484.61	2910.48	921.85	2134.08
current state of health (stated)	3.12	5.83	6.39	4.76
invalid (1=yes)	0.04	0.20	0.04	0.19
household member died this year (1=yes)	0.006	0.08	0.003	0.06
divorced (1=yes)	0.05	0.22	0.04	0.19
separated from partner (1=yes)	0.01	0.12	0.01	0.11
partner dead (1=yes)	0.06	0.24	0.05	0.22
just married (1=yes)	0.02	0.15	0.05	0.21
just divorced (1=yes)	0.005	0.07	0.004	0.06
just separated (1=yes)	0.012	0.11	0.006	0.08
partner just died (1=yes)	0.003	0.06	0.001	0.03
just had a baby (1=yes)	0.04	0.19	0.02	0.14
pregnant (1=yes)	0.01	0.11	0.01	0.01
just fired from job (1=yes)	0.02	0.12	0.00	0.07