

CEP Discussion Paper No 784

March 2007

(Revised January 2008)

The Marginal Utility of Income

R. Layard, G. Mayraz and S. Nickell

Abstract

In normative public economics it is crucial to know how fast the marginal utility of income declines as income increases. One needs this parameter for cost-benefit analysis, for optimal taxation and for the (Atkinson) measurement of inequality. We estimate this parameter using four large cross-sectional surveys of subjective happiness and two panel surveys. Altogether, the data cover over 50 countries and time periods between 1972 and 2005. In each of the six very different surveys, using a number of assumptions, we are able to estimate the elasticity of marginal utility with respect to income. We obtain very similar results from each survey. The highest (absolute) value is 1.34 and the lowest is 1.19, with a combined estimate of 1.26. The results are also very similar for subgroups in the population. We also examine whether these estimates (which are based directly on the scale of reported happiness) could be biased upwards if true utility is convex with respect to reported happiness. We find some evidence of such bias, but it is small—yielding a new estimated elasticity of 1.24 for the combined sample.

JEL classification: I31, H00, D1, D61, H21

Keywords: Marginal utility, income, life satisfaction, happiness, public economic, welfare, inequality, optimal taxation, reference-dependent preferences.

This paper was produced as part of the Centre's Wellbeing Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

Acknowledgements

We thank the Esmee Fairbairn Foundation for financial support to the Centre for Economic Performance's Wellbeing Programme.

Richard Layard is Director of the Wellbeing Programme at the Centre for Economic Performance, London School of Economics. Guy Mayraz is an Occasional Research Student for the Wellbeing Programme at CEP, and Research Student in the Department of Economics, LSE. Stephen Nickell is Professor of Economics and Warden of Nuffield College, Oxford. He is also an Associate of the Labour Markets Programme, CEP.

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© R. Layard, G. Mayraz and S. Nickell, submitted 2007

ISBN 978 0 85328 159 7

1 Introduction

In normative public economics it is crucial to know how fast the marginal utility of income declines as income increases. One needs this parameter for cost-benefit analysis, for optimal taxation and for the (Atkinson) measurement of inequality. For example, in cost benefit analysis a central problem is how to aggregate the costs and benefits that accrue to different people. A natural way to do this is to weight each person's change in income by his or her marginal utility of income¹.

The key issue for public economics is not how strongly income affects utility but how this effect changes with income. To focus on this question we assume that the elasticity, ρ , of marginal utility with respect to income is constant. Thus, utility, u , is given by

$$u = \begin{cases} \frac{y^{1-\rho}-1}{1-\rho} & \rho \neq 1 \\ \log y & \rho = 1 \end{cases} \quad (1)$$

where y is income. It follows that if we take two people, A and B , the ratio of their marginal utilities is given by

$$\frac{\partial u^B / \partial y}{\partial u^A / \partial y} = \left(\frac{y^A}{y^B}\right)^\rho \quad (2)$$

Thus, for example, if $\rho = 1$ (and so $u = \log y$) the marginal utilities are inversely proportional to income: someone with an income of \$10,000 has ten times the marginal utility of someone getting \$100,000. Both [Bernoulli \(1738, 1954\)](#), who invented the mathematical idea of utility, and [Dalton \(1920\)](#), who pioneered the measurement of inequality, made just this logarithmic assumption².

What is needed, in this context, is a cardinal measure of utility where unit intervals have the same meaning at all points along the scale. As is well-known, such a measure is assumed whenever decisions are explained on the basis of a decision function involving the weighted addition of utility in different states (choice under uncertainty) or in different time periods (intertemporal choice). Thus, studies of such behaviour have been used to obtain estimates of ρ . However, these estimates have two key problems for our present purpose: (1) We are interested in utility as actually experienced *ex-post* rather than in utility as evaluated *ex-ante* for purposes of decision-making. As [Kahneman \(Kahneman, 1999\)](#) points out, the forecasts of utility that are used for decision making often turn out to be inaccurate. (2) Moreover, the estimates of ρ obtained from choice behaviour are very indirect, involving many extraneous assumptions, and it is not therefore surprising that they yield a wide range of estimates. Those

¹This corresponds to a utilitarian social welfare function, which is a simple average of the utility of all individuals. If the social welfare function is concave in the individual utilities, a further weighting is needed for the diminishing marginal social value of an increment to individual utility ([Atkinson and Stiglitz, 1980](#), Part 2).

²Dalton measured inequality by W_e/W , where W_e is the welfare that would be obtained if everybody received the average income. This measure is only invariant with respect to equi-proportional changes in all incomes if $\rho = 1$. By contrast, Atkinson measured equality by y_e/\bar{y} where y_e is the equal income which would yield W if everybody had it. This measure is invariant to equi-proportional changes in all incomes, even if $\rho \neq 1$.

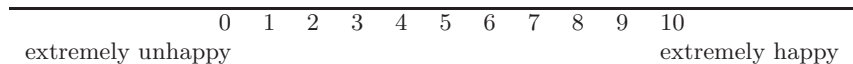
based on choice under uncertainty range from 0 to 10 or even higher³. Estimates based on intertemporal choice vary less⁴, but in any case the assumption of intertemporal additivity is problematic⁵. For both these reasons a new estimate could be helpful, especially if it yields consistent estimates from a wide variety of data sources.

Our approach is based not on inferences from behaviour, but on the direct measurement of experienced happiness in six major surveys which also asked questions about income and other variables. Most of the surveys have been analysed before to examine the impact of income on happiness⁶, but as far as we are aware, none of these studies address the crucial question for public economics of the curvature of the relationship—in other words the value of ρ .

That is the focus of this paper. In Section 2 we discuss some key issues which arise in the measurement of utility. In Section 3 we describe the survey data in more detail. In Sections 4 and 5 we give our results, and conclude in Section 6.

2 Measuring utility

In the surveys that we analyse, a typical happiness question is "Taking all things into account, how happy are you these days?". The respondent then chooses one of a number of values, for example:



In most surveys, individuals are surveyed only once, but in the German Socio-Economic Panel (GSOEP) and the British Household Panel Study (BHPS) they are surveyed in a number of successive years. The datasets we use are described in detail in Table 1 and the happiness questions in Table 2. Histograms of the replies are in Figure 1.

In order to make use of these data, we need to investigate how they are generated. First, we assume that individual i has a level of experienced utility u_i which is cardinal and is a function of a set of observable variables

³Lanot et al. (2006) includes a survey of measures showing a wide variety of estimates. For example, Metrick (1995) studied game shows and estimated that risk behaviour is close to risk neutrality, while Cohen and Einav (2005) investigated insurance purchases, and estimated a coefficient of risk aversion of more than 10. As is well known, explaining the equity premium within expected utility theory requires a coefficient of the order of 30 or even more (Mehra and Prescott, 1985, 2003)

⁴See Blundell et al. (1994) and Chetty (2006) for examples of estimates of ρ based on intertemporal choices.

⁵Many economists believe that habit formation plays an important role in intertemporal substitution decisions (Attanasio, 1999).

⁶See, for example, Clark and Oswald (1996), Easterlin (1995), Winkelmann and Winkelmann (Winkelmann and Winkelmann, 1998), Di Tella et al. (2003), Helliwell (2003), Van Praag (2004), and Blanchflower and Oswald (2004). For a recent survey of the happiness and income literature see Clark et al. (forthcoming).

(e.g. income, work, etc.). We suppose u_i is comparable across individuals. In order to generate an answer to a happiness question, individual i applies an idiosyncratic, strictly increasing, function f_i to u_i . Thus, reported happiness for individual i , h_i , is given by

$$h_i = f_i(u_i) \tag{3}$$

Under this assumption, reported happiness is an ordinal non-comparable measure of true utility, because each individual applies their own monotonic transformation to u_i . It is our intention to make a further restrictive assumption, namely that f_i is common to all individuals up to a random additive term, that is

$$h_i = f_i(u_i) = f(u_i) + v_i \tag{4}$$

with v_i representing an error term that is independent of the circumstances affecting true utility. There is a considerable body of evidence to justify this assumption. First, evidence suggests that respondents report their level of happiness in a way that is consistent with other meaningful measures of their utility. For example we can ask a person’s friend to rate the subject’s happiness, or even ask an independent observer who has never met the person before. The reports of these ”third parties” turn out to be correlated across people with the subject’s own report (Diener and Suh, 1999). Second, neuropsychologists can now measure the level of activity in the areas of the brain where positive and negative affect are experienced. These levels of activity are well correlated with the subject’s self-report (Davidson, 1992, 2000; Davidson et al., 2000) both across people and over time. Third, another important exercise in validation is to ask whether the resulting measure relates to external factors in an expected manner. All the evidence that we have is that it does. For example, income increases reported happiness, separation and divorce decrease it, etc. (Kahneman et al., 1999). Likewise, measured happiness predicts quitting, marital break-up, and other behaviours⁷. Overall, in light of this evidence, we feel quite justified in making the assumption embedded in Equation 4, namely that individuals’ reported happiness reflects this experienced utility in a manner which is at least partially comparable across individuals and, in particular, is independent of income.

Nevertheless, there remains a problem with Equation 4 arising from the possible non-linearity of the function, $f(\cdot)$. If f is non-linear, then equal intervals on the reported happiness scale do not reflect equal intervals of true utility. For example, given our interest in estimating the decline of marginal utility, a significant concern is that happiness reports may be concave in true utility. We investigate this potential problem in Section 5. We cannot reject the possibility of such a non-linear relationship, but show that it would only have a small effect on our conclusions.

⁷A good survey of papers studying the predictive value of satisfaction measures can be found in Clark et al. (2006).

Table 1: Description of surveys as used. Details relate only to the subset of observations used in the regressions. The happiness variable was categorical when levels are indicated, and numerical when a range is indicated.

Survey Name	Type	Countries	Years	Obs.	Happiness variable	Household income variable
General Social Survey	Cross-section	United States	1972-2004 (25 waves)	17,603	Happiness (3 levels)	Yearly gross
World Values Survey	Cross-section	Worldwide (50)	1981-2003 (4 waves)	37,288	Life satisfaction (1-10)	Varies
European Social Survey	Cross-section	Europe (22 in 1st wave, 26 in 2nd wave)	2002 and 2004 (2 waves)	26,687	Both happiness and life satisfaction (0-10)	Monthly net
European Quality of Life Survey	Cross-section	Europe (28)	2003	8,175	Both happiness and life satisfaction (1-10)	Monthly net
German Socio-Economic Survey	Panel	Germany	1984-2005 (22 waves)	78,877	Life satisfaction (0-10)	Monthly net
British Household Panel Survey	Panel	Britain	1996-2004 (7 waves)	43,484	Life satisfaction (1-7)	Monthly net

3 Data and strategy

Our aim is to estimate a direct utility function. The dependent variable is happiness, h , and the independent variables are income, y (acting as a proxy for consumption), hours of work (measured in various ways), sex, age, education, marital status, and employment status. A few comments are needed on these variables.

3.1 The reported happiness variable

The happiness questions are listed in Table 2. As Figure 1 shows, allowing for the different ranges the histograms look similar. For our subsequent statistical analysis we want to compare surveys, so we normalise all the happiness data to a 0 to 10 scale using the appropriate affine transformation⁸. When both happiness and satisfaction with life as a whole

⁸For example, 1-7 was shifted and stretched to run from 0 to 10. The qualitative scale in the General Social Survey (not very happy, fairly happy, very happy) was first turned into a numerical scale: 0, 1, 2 and then stretched.

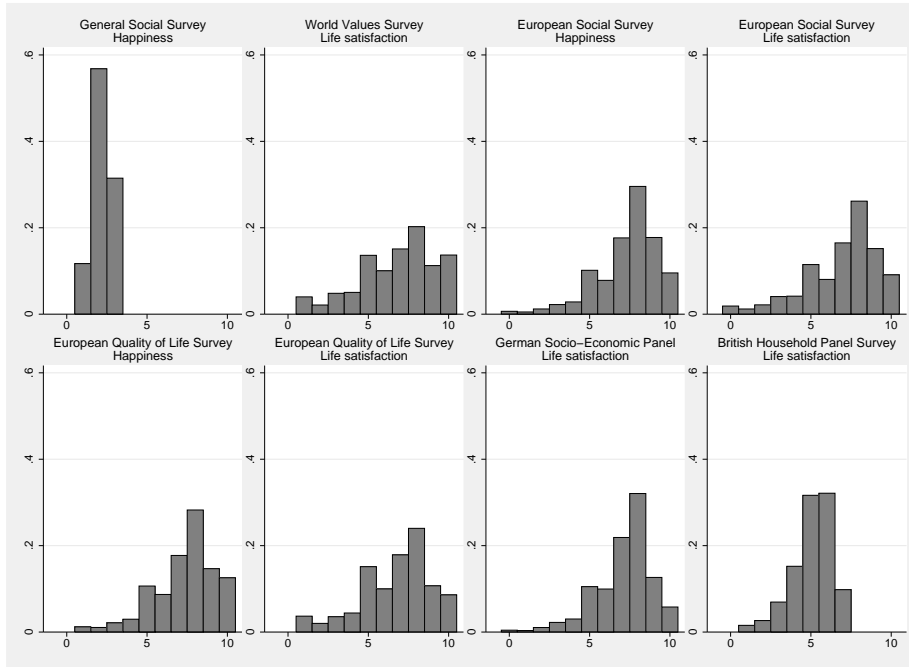


Fig. 1: Histograms of answers to happiness and life satisfaction questions.

were recorded, the responses were averaged to produce a single dependent variable. The averaged variable was typically more closely related to external circumstances than either happiness or life satisfaction on their own, consistent with the possibility that each of these variables includes an independent measurement error.

3.2 The income variable

We elected to use total household income, rather than the respondent's own income. Household income was adjusted to units of constant purchasing power with the exception of the World Values Survey⁹. It was not, however, normalised to income per equivalent adult as we see children as a choice variable, so that income normalisation is misleading for our purposes. Exact income data are available for the panel surveys (German Socio-Economic Panel and British Household Panel Survey). In the cross-section surveys only income bands are available, and these we converted into numerical values using the mid point of each band. For respondents

⁹The period of time and the currency varied by country and wave, and we did not have sufficient documentation available to ensure correct normalisation. This does not matter for the estimation of ρ in Section 4.3 because we used country/wave and income interaction terms.

Table 2: The survey questions for the reported happiness variables we used.

Survey	Variable	Question
General Social Survey	Happiness	Taken all together, how would you say things are these days would you say that you are very happy, pretty happy, or not too happy?
World Values Survey	Life sat.	All things considered, how satisfied are you with your life as a whole these days? Please use this card to help with your answer.[range of 1-10 with 1 labelled "Very Dissatisfied" and 10 labelled "Very Satisfied"]
European Survey	Social Happiness	Taking all things together, how happy would you say you are? Please use this card [range of 0-10 with 0 labelled "Extremely unhappy", and 10 labelled "Extremely happy"]
European Survey	Social Life sat.	All things considered, how satisfied are you with your life as a whole nowadays? Please answer using this card, where 0 means extremely dissatisfied and 10 means extremely satisfied. [range of 0-10 with 0 labelled "Extremely dissatisfied" and 10 labelled "Extremely satisfied"]
European Quality of Life Survey	Quality of Life Happiness	Taking all things together on a scale of 1 to 10, how happy would you say you are? Here 1 means you are very unhappy and 10 means you are very happy.
European Quality of Life Survey	Quality of Life Life sat.	All things considered, how satisfied would you say you are with your life these days? Please tell me on a scale of 1 to 10, where 1 means very dissatisfied and 10 means very satisfied.
German Socio-Economic Panel	Socio-Economic Panel Life sat.	In conclusion, we would like to ask you about your satisfaction with your life in general. Please answer according to the following scale: 0 means 'completely dissatisfied', 10 means 'completely satisfied'. How satisfied are you with your life, all things considered?
British Household Panel Survey	Household Panel Life sat.	How dissatisfied or satisfied are you with your life overall? [range of 1-7 with 1 labelled "Not satisfied at all" and 7 labelled "Completely satisfied".]

in the lowest income band we assumed an income of two thirds of the upper limit of the band, and for respondents in the highest income band we assumed an income of 1.5 of the lower income limit of the band.

3.3 Other causal variables

We also have to control for other variables affecting happiness—especially hours of work. Unfortunately, hours data can in many cases be unreliable, and we proceed in two ways.

In our main analysis we measure employment by the following set of dummies: inactive (the omitted variable), unemployed job seeker, full-time work, and part-time work (the distinction between full-time and part-time work is not available in all the surveys). In a separate analysis we combine the last two categories into a single dummy, and also interact this dummy with a cubic polynomial in hours of work.

We also include the following control variables in all our analysis:

- Sex
- Age (including a quadratic)
- Education (dummies for degrees of achievement, or years of education and years of education squared)
- Marital status (dummies)

In addition to these standard controls we include country or wave dummies, and country/wave interaction terms when there is more than one wave per country. The inclusion of these dummy variables enables us to control for common fixed effects to do with country/wave. Most importantly, it allows us to control for systematic additive reporting biases. For example, in the European Social Survey a given level of income in Denmark is associated with a higher reported happiness than a similar level of income in Poland. This may have to do with the higher quality of life in Denmark, but it may also be the case that people in Denmark will report any given level of utility more positively than would people in Poland. By using country/time dummies we avoid this problem altogether and obtain *ceteris paribus* results that are valid across different countries and time periods¹⁰.

3.4 Sample

In order to have a relatively homogeneous population, we confined ourselves to people aged 30-55, for whom annual income tends to be well correlated with permanent income. For the same reason we excluded the 5% at either extreme of the distribution of fitted residuals from a linear regression of $\log y$ on a set of standard regressors. We believe that many of these observations are the result of measurement error, or else reflect transitory deviations from usual income. Taking such observations at face

¹⁰The flip side is that our evidence can shed no light on the source of inter-country variability in reported happiness, or the stability of reported happiness across time (The Easterlin Paradox).

value could lead to misleading conclusions about the long-run relationship between utility and income (consistent with this the plots of Figure 3 show that no clear functional relationship with happiness exists for these extreme income observations, while an excellent fit exists for all the other observations).

3.5 Strategy of analysis

For our analysis we make the following assumptions:

1. Reported happiness, h , is linked to true utility, u , via a fixed transformation f , as in Equation 4.
2. The experienced utility of person i at time t in country c can be written as follows:¹¹

$$u_{it} = \alpha_{ct} \frac{y_{it}^{1-\rho} - 1}{1-\rho} + \sum_j \beta_j x_{jit} + \gamma_i + \gamma_{ct} + \epsilon'_{it} \quad (5)$$

where y_{it} is person's i 's income at time t , which enters the equation via a constant elasticity of marginal utility function with parameter ρ . In this formulation ρ corresponds to minus the elasticity of marginal utility with respect to income. $\rho = 0$ corresponds to a linear relationship between utility and income, with higher values of ρ denoting increasing concavity. Marginal utility is constant for $\rho = 0$, and is diminishing for $\rho > 0$. The coefficient on income, α_{ct} , is assumed to be the same for all people within a given country at a given point in time, but it can vary between countries or at different time points. x_{jit} are other controls, γ_i is a person fixed effect, and γ_{ct} is a country/wave dummy. Finally, ϵ'_{it} is the error.

Writing the model in this way allows for the possibility that income is not measured in comparable units in different countries or at different points in time. This is a particular problem in the World Values Survey in which the documentation of the measurement of income is incomplete, but is also a potential concern in other surveys due to the limitations of purchasing power parity calculations. In the model of Equation 5 y_{it} can be multiplied by an arbitrary constant μ_{ct} with no change to the fit of the model, and hence no bias in the maximum likelihood estimate of ρ ¹².

Note, the function explaining u_{it} is additive in form. This is a crucial assumption. Without imposing some such functional form on the relationship between u and y , we cannot identify the curvature of this relationship. We have more to say on this issue in Section 5.

¹¹Changes in income (and not only the level) are believed to affect utility, but in our two panel surveys income changes come out as insignificant in regressions. The same is true if additional lags are included.

¹²If the coefficient on income had not been allowed to vary between countries or different points in time, the model would have a special status for $\rho = 1$. This is because when $\rho = 1$ ($\log y$) any normalisation errors are absorbed by the country/wave dummies and do not contribute to the error. Allowing α to vary by country/wave eliminates the effect of normalisation errors for any ρ .

3. ρ is the same for different people. This assumption is not essential for our estimates to be useful, but it allows a particularly clear interpretation of the results. We relax this assumption by estimating ρ separately in different surveys, and also in selected population subgroups.

Our goal is to obtain an estimate of ρ . In the main part of our analysis we assume that the transformation f in Equation 4 is linear. That is,

$$h_{it} = u_{it} + v_{it}, \quad \text{cov}(v_i, y_i) = 0 \quad (6)$$

Combining this assumption with Equation 5 we obtain

$$h_{it} = \alpha_{ct} \frac{y_{it}^{1-\rho} - 1}{1-\rho} + \sum_j \beta_j x_{jit} + \gamma_i + \gamma_{ct} + \epsilon_{it} \quad (7)$$

where $\epsilon_{it} = \epsilon'_{it} + v_{it}$. The main analysis includes the following three steps:

1. We investigate the logarithmic hypothesis ($\rho = 1$) using cross-sectional analysis.
2. Panel fixed effects analysis: we test whether including person fixed effects in the regressions substantially changes the estimated relationship between income and utility.
3. Maximum likelihood estimates of the curvature parameter, ρ , based on Equation 7.

Finally, in Section 5 we relax the linearity assumption in Equation 6. We investigate whether the function relating reported happiness to true utility is concave, and compute the implied correction to the estimate of ρ obtained in our main analysis.

4 Main results

Figure 2 shows the simple cross-sectional relationship between happiness and income in one of the surveys. Each of the 20 points corresponds to 5% of the people arrayed according to income. The body of our results is an attempt to analyse the curvature of this relationship in greater detail.

4.1 The logarithmic model: cross-section

We begin by exploring the hypothesis that $\rho = 1$. In other words, that happiness depends linearly on log income. We therefore estimate

$$h_{it} = \alpha \log y_{it} + \sum_j \beta_j x_{jit} + \gamma_{ct} + \epsilon_{it} \quad (8)$$

The fit is quite good and is plotted in Figure 3 where both income and happiness have been adjusted for the effect of all the control variables. As can be seen, if we ignore the 5% at each end of the income distribution, the plots lie quite well along a straight line. For the moment we are focusing on the cross-sectional fit, i.e. on the six upper panels in the

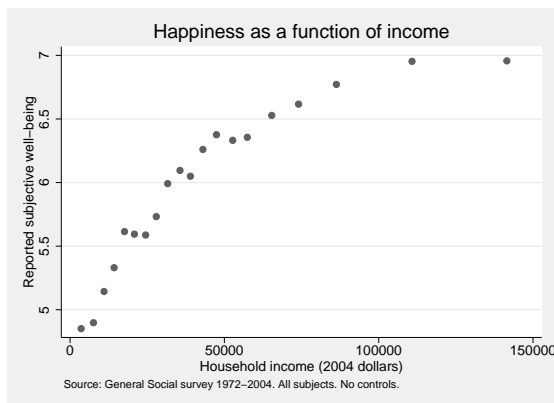


Fig. 2: The simple cross-sectional relationship between reported happiness and income in the U.S. General Social Survey.

figure. Considering the variety of countries, the similarity of the slopes is interesting.

However, the obvious question is whether the relationship with log income is truly linear. We therefore add a term in log income squared. The estimated equation is as follows:

$$h_{it} = \alpha_1 \log(y_{it}/\bar{y}_{ct}) + \alpha_2 \log^2(y_{it}/\bar{y}_{ct}) + \sum_j \beta_j x_{jit} + \gamma_{ct} + \epsilon_{it} \quad (9)$$

As can be seen, we have now normalised income by the average income in that country and time period. This is because log income and log income squared are highly correlated across the whole sample, which makes the interpretation of the estimated coefficients quite difficult. The country dummy does not eliminate this problem, which is why we normalise income by average income in the country and year concerned.

The results are shown in Table 3. As can be seen, in all the surveys the coefficient on the squared term is negative—meaning that the relationship between happiness and income is more concave than implied by the log function alone, i.e. ρ is greater than 1. To estimate ρ we need to proceed to maximum likelihood estimation.

4.2 The logarithmic model: panel fixed-effects

However, first we can exploit further the longitudinal data from the German Socio-Economic Panel and British Household Panel Survey by including person fixed effects. In this way, we can eliminate any bias caused by correlations between income and unmeasured personal characteristics which also directly affect happiness. Such correlations are the main source of endogeneity in the income variable. We therefore estimate Equations 8

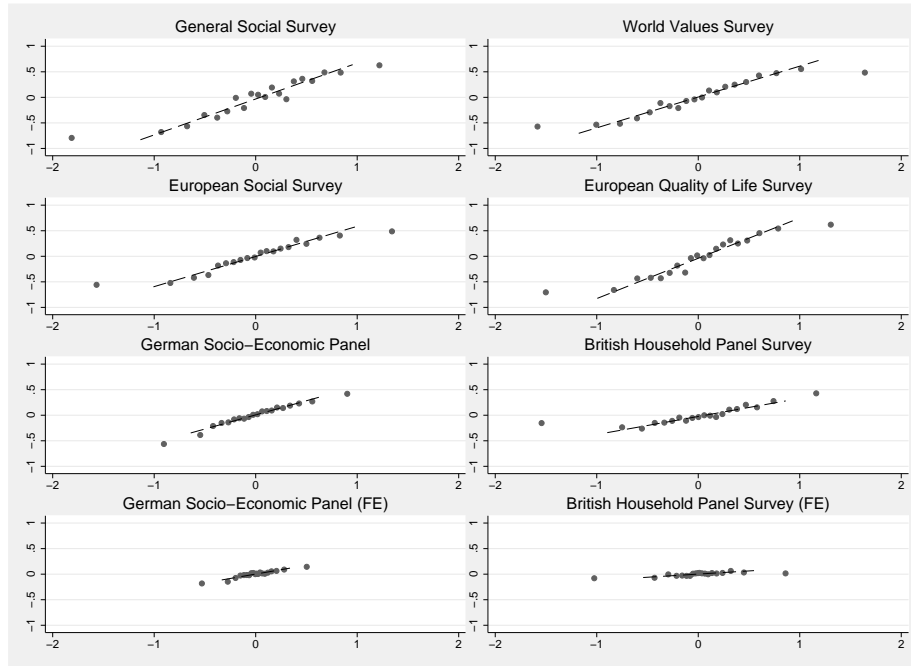


Fig. 3: The partial relationship between reported happiness (y -axis) and log income (x -axis). FE indicates person fixed-effects were included in the regression. The graphs show a consistent near-linear relationship, with some variation in the slopes.

Table 3: Partial regression coefficients of happiness on log relative income and log relative income squared (t-scores in parentheses). Relative income is relative to mean income in country/year. Country/year dummies and other personal controls are included but not shown.

Survey	$\log(y/\bar{y})$	$\log^2(y/\bar{y})$
General Social Survey	0.70 (14.0)	-0.07 (1.7)
World Values Survey	0.62 (26.1)	-0.09 (3.7)
European Social Survey	0.59 (24.9)	-0.15 (5.8)
European Quality of Life Survey	0.81 (17.4)	-0.08 (1.5)
German Socio-Economic Panel (no fixed effects)	0.55 (27.2)	-0.10 (2.7)
British Household Panel Survey (no fixed effects)	0.35 (14.0)	-0.06 (2.6)
German Socio-Economic Panel (panel fixed effects)	0.36 (11.0)	-0.08 (0.1)
British Household Panel Survey (panel fixed effects)	0.25 (6.7)	-0.10 (3.5)

and 9 with a person fixed effect added in, in order to get a better handle on causality.

The results, reported in the bottom two panels of Figure 3 and Table 3, are similar to those of the cross-sectional analysis of the previous section, a result is consistent with the causal interpretation¹³. As compared with the cross-section, the slopes in the panel fixed effects regressions are somewhat lower. The downward bias resulting from measurement error in income is exacerbated in fixed effects models (Griliches and Hausman, 1986), so some reduction in the slope coefficient is to be expected¹⁴

4.3 Maximum likelihood estimation of ρ

We now come to the central analysis of this paper, where we directly estimate the curvature of the happiness-income relationship. We use maximum likelihood to estimate the parameters of Equation 7, that is

$$h_{it} = \alpha_{ct} \frac{y_{it}^{1-\rho} - 1}{1-\rho} + \sum_j \beta_j x_{jit} + \gamma_{ct} + \epsilon_{it} \quad (10)$$

To find the maximum likelihood estimate we compute the log likelihood for values of ρ between 0 and 3 in steps of 0.1, and then use a quadratic approximation in the vicinity of the maximum. As before observations with extreme income values in the top or bottom 5% are excluded.

The results are shown in the first column of Table 4. In each of these six surveys, using data on different countries and time periods, the maximum likelihood estimate of ρ falls in the narrow range of 1.19-1.34. This similarity is in spite of the great differences between the countries, cultures, and languages in the surveys we used, as well as the differences between the questionnaires in the different surveys.

We can now ask what *single* value of ρ best explains the data in the six surveys we have. To obtain the best overall estimate we add up for each value of ρ the log likelihood terms from each of the six surveys, enabling us to find how good each value of ρ is at explaining the data from all the surveys put together. The overall maximum likelihood estimate we obtain is $\rho = 1.26$, with a 95% confidence interval of 1.16-1.37.

The critical question, however, is whether this overall estimate is in fact consistent with all the separate sets of data. The test for that is how $\rho = 1.26$ compares with the confidence intervals in the different surveys, and the answer is that it falls well within the 95% confidence interval of all the individual surveys. Thus, unlike the case with choice under

¹³In an annual time series, variations in happiness are unlikely to be a major cause of variation in income. Another potential concern is time-dependent common causes, such as a promotion, which may have a direct effect on reported happiness, as well as on income. To demonstrate conclusively that these effects are secondary it is necessary to utilise exogenous events which affect income, but have no direct effect on utility. A start in this direction has been made by Gardner and Oswald (2007), who studied winners of medium sized lottery prizes.

¹⁴Income is measured with some error not only because of simple misreporting, but also because income at the time of the survey is not always representative of the person's typical income.

Table 4: Maximum likelihood estimates of ρ in different surveys using first the linear dependent variable model, and then ordered logit. 95% confidence intervals in parentheses.

Survey	Standard estimate	Ordered logit estimate
General Social Survey	1.20 (0.91-1.48)	1.26 (0.96-1.55)
World Values Survey	1.25 (1.05-1.45)	1.26 (1.06-1.46)
European Social Survey	1.34 (1.12-1.55)	1.25 (1.02-1.49)
European Quality of Life Survey	1.19 (0.87-1.52)	1.05 (0.71-1.38)
German Socio-Economic Panel	1.26 (0.90-1.63)	1.15 (0.81-1.49)
British Household Panel Survey	1.30 (0.97-1.62)	1.32 (0.99-1.65)
Combined estimate	1.26 (1.16-1.37)	1.23 (1.12-1.34)

uncertainty, we find that estimates of ρ based on happiness surveys yield a single estimate that is consistent with very different data sets.

4.4 The robustness of the estimate of ρ

Even if a single value of ρ is consistent with data from different cultures, it can still be the case that significantly different values of ρ may be found in different subgroups of the same population. To investigate this possibility we looked at subgroups defined by sex, age, education level, and by marital or cohabiting status¹⁵. After defining the different subgroups we proceeded to find the estimate of ρ that is combined across the different surveys, but computed for each subgroup separately.

The results are reported in Table 5. As in the case of the different surveys, we find high degree of consistency between the estimates for different subgroups. The formal test for this is that the overall maximum likelihood value of 1.26 lies not only within the 95% confidence interval of the different surveys, but also well within the 95% confidence interval for all the different subgroups. Taken together, these results suggest that ρ does not vary systematically between different subgroups of the population defined by sex, age, education, or marital status. Our results, of course, say nothing about the possibility of non-systematic variation in ρ between different individuals. Furthermore, these results are consistent with systematic differences in choice behaviour¹⁶, but imply that such differences are likely to be due to factors other than underlying differences in experienced utility.

¹⁵In the case of education we defined a low, medium and high education subgroups in each survey, so that roughly one third of the sample is in each of the three categories. In the case of relationship status we split the sample between couples, singles, and other (separated, divorced, or widowed). In four of the surveys the marital status variable contained a living with partner category which we combined with the married category to create the couples subgroup. In the two surveys in which living with partner was not part of the marital status variable the couples category equals the married category.

¹⁶There is some evidence for systematic group differences in risk attitudes (Dohmen et al., 2006) and in intertemporal choice (Blundell et al., 1994).

Table 5: Maximum likelihood estimates of ρ in different subgroups (using a standard linear dependent variable model and with 95% confidence intervals in parentheses).

Sub-group	ML estimate
Men	1.22 (1.06-1.39)
Women	1.26 (1.11-1.40)
Age 30-42	1.27 (1.12-1.42)
Age 43-55	1.26 (1.10-1.41)
Low education	1.13 (0.85-1.40)
Mid education	1.21 (1.01-1.42)
High education	1.26 (1.16-1.37)
Couples	1.27 (1.11-1.43)
Singles (never married)	1.44 (1.13-1.77)
Widowed/Separated/Divorced	1.34 (0.85-1.83)
All subjects	1.26 (1.16-1.37)

As a further test of robustness we tested the alternative specification for hours that we described in Section 3.3, namely using a single dummy for being at work, and adding its interaction with $hours$, $hours^2$, and $hours^3$. In so doing we also kept a dummy for being unemployed, as distinct from being outside the labour force. The result of this alternative specification is a maximum likelihood estimate of $\rho = 1.24$ with a 95% confidence interval of 1.14-1.35, as compared with $\rho = 1.26$ with a 95% confidence interval of 1.15-1.37 when using our standard specification. The estimate of ρ thus appears robust to different treatments of working hours. Finally, some readers may be interested in the estimated coefficient when, in addition to controlling for hours worked, only gender and age controls are used. The resulting estimate in this case is $\hat{\rho} = 1.32$, which lies within the 95% confidence interval for the standard specification.

5 Compression of the reported happiness scale

Finally we need to raise the following question: Is the scale which respondents use in reporting their utility the *true* scale (up to an arbitrary positive affine transformation), or some non-linear transformation of it? In the terms of Equation 4, is f a linear function?¹⁷

In particular, there comes from psycho-physics the suspicion that reports using bounded scales (e.g. 0-10) may be a concave function of the

¹⁷Loosely speaking, of course. Since utility reports are discrete, whereas true utility is continuous, the function f cannot strictly be linear. But, if we think of the reports as the rounding off to the nearest point on the scale of a continuous function of utility, we can give precise meaning to the notion of f being linear or else a non-linear concave function.

true scale of utility. In reports of physical sensations this often seems to happen, since replies on bounded scales are generally concave with respect to the use of unbounded scales. For example respondents have been asked to report the duration of noises of differing length (Stevens, 1975, p. 229). In one experiment, a bounded scale of 1 to 7 was offered and respondents were exposed to different durations (the shortest and longest durations being exhibited first). Respondents assigned a category to each duration. These values turned out to be a concave function of the true durations. By contrast when respondents were offered an unbounded scale, their replies were roughly proportional to the true durations.

One could argue that findings relating to the sensation of external physical stimuli are irrelevant to the reporting of utility. When people report their utility they are reporting not on their experience of an external stimulus, but directly on their internal state. Thus, there may well be a non-linear relationship between external circumstances and utility (as we suggest for the case of income), but there is no obvious reason why the relationship between utility and reported happiness should be non-linear. In fact, there is evidence from the neuroscience of physical sensation (Johnson et al., 2002) that the relationship between neural activity and verbal reports is linear. If we are willing to assume that subjective experience is linear in neural activity this evidence implies that verbal reports are linear in subjective experience. Finally, while many people would agree on the meaning of intervals in reported happiness, there is less sense in which we can talk of the concept of zero happiness. So presenting people with an interval scale seems to be a natural way to proceed.

Our strategy to explore this issue is as follows. Consider first a general model of the relationships between reported happiness, h_i , true experienced utility, u_i , and income, y_i . Thus we might have

$$h_i = f(u_i), \quad f' > 0 \quad (11)$$

$$u_i = g\left(\frac{y_i^{1-\rho} - 1}{1-\rho}, z_i, \epsilon_i\right) \quad (12)$$

where z_i reflects other variables and ϵ_i is an unobserved error term. The only investigation we can undertake is of the function

$$h_i = f\left(g\left(\frac{y_i^{1-\rho} - 1}{1-\rho}, z_i, \epsilon_i\right)\right) \quad (13)$$

We are interested in the curvature of u_i , and hence that of the function g with respect to y_i . Any analysis of Equation 13 can tell us nothing about this without making a prior assumption about the function g . So we continue to maintain the assumption of linearity, that is

$$u_i = \alpha \frac{y_i^{1-\rho} - 1}{1-\rho} + z_i + \epsilon_i \quad (14)$$

Conditioning on this assumption, our strategy is then to investigate the function f in the equation

$$h_i = f\left(\alpha \frac{y_i^{1-\rho} - 1}{1-\rho} + z_i + \epsilon_i\right) \quad (15)$$

In particular, we see if we can find any evidence that $f'' < 0$. We use three approaches based on this framework:

1. Ordered logit: if f is concave and if the error term ϵ_i is symmetric, then the estimated cut points of the ordered logit will tend to be convex with respect to reported happiness (implying a concave transformation from true utility to reported happiness).
2. Variance of the residuals: if we estimate a linear version of the happiness regression (Equation 15), then if f is concave, the fitted residuals will tend to have a lower variance at higher values of predicted happiness, assuming ϵ_i is homoscedastic. This is essentially because at high values of predicted happiness, variations in actual happiness generated by variation in ϵ_i tend to be lower because of the concavity of f in the true model.
3. Spline: assuming f is concave, then if we combine the regressors into a single prediction variable by running a linear regression, and then run separate regressions of reported happiness on each half of the sample, the slope should be higher for low values of \hat{h}_i and lower for high values of \hat{h}_i .

In the next three sections we pursue these three approaches and compute a corrected estimate of ρ . This turns out to be quite close to the estimate generated in Section 4.3. We then give the intuition for this outcome.

5.1 Ordered logit

Ordered logit makes it possible to allow for a non-linear f in the estimate of ρ . The ordered logit cut points are displayed in the top left panel of Figure 4, and are consistent with the proposed concave relationship between reported happiness and true utility. The maximum likelihood estimate of ρ obtained using ordered logit is reported in the right column of Table 4. The resulting estimate of $\hat{\rho} = 1.23$ is only marginally lower than the linear regression estimate of 1.26 (Section 4.3).

This calculation relies on the assumption of a symmetric error distribution. A negatively skewed error distribution could also account for convex cut points even if f is linear.

5.2 Variance of the residuals

Another approach involves a fundamental assumption about what it is we want to measure in public economics. We are concerned with the impact of the outside world upon how people feel. If a person notices no difference in how he feels, that should count as no difference in happiness. And the natural units in which to measure changes in happiness is the "just noticeable difference" (JND). This has been much used in the study of physical sensation, and is variously defined — for example, as the difference that in repeated observations would be noticed 50% of the time (Stevens, 1975).

However, we cannot assume that, when people use a response scale from 0 to 10, they make the differences between each point on the scale

represent a given number of JNDs. A way to investigate this is by repeatedly asking an individual about his or her happiness (with no change of circumstances or mood available to explain any change in the answer). In such a test-retest situation we would not expect the same answer every time unless the JND was vanishingly small. But we would want the variance of the replies to be independent of the person’s position on the scale of happiness — implying that units of the scale reflect the same number of JNDs at all points on the scale.

We would therefore like to retest individuals in unchanging circumstances. But we have no such data for closely repeated observations. We do, however, have repeated annual observations on the happiness of the same individual from two panel surveys. We can extract from these data a person fixed effect, as well as the effect of changes in all observable independent variables, leaving an error that approximates a test-retest error. In the top right panel of Figure 4 we group individuals into 20 groups with ascending average predicted reported happiness, and plot for each group the average value of their root mean squared prediction error σ_i . This shows clearly that the errors are smaller for higher values of reported happiness.

If we assume this error reflects test-retest error and possibly other homoscedastic error, then the error variance for true utility should be the same at all points on the scale¹⁸. The negative slope in the top right panel of Figure 4 would then imply that at the upper end the reported happiness scale is a compressed version of the true utility scale.

In producing a corrected scale of true utility, our psychological assumptions are:

1. True utility, u , should be measured using a scale with constant standard errors.
2. Reported happiness, h , has variable units which are proportional to the measured standard errors.

Thus, to convert intervals on the h scale to units on the u scale we deflate them by $\sigma(h)$, or for purposes of normalisation by $\sigma(h)/\sigma_0$, where σ_0 is the error variance at a reference point h_0 , where true utility will be defined as equal to reported happiness:

$$du = \frac{dh}{\sigma(h)/\sigma_0} \tag{16}$$

We can now use the individual data underlying the top right panel of Figure 4 to estimate the following relationship:

$$\frac{\sigma(h)}{\sigma_0} = 1 - c(h - h_0) \tag{17}$$

Where $c > 0$. Substituting Equation 17 into Equation 16 and integrating we obtain

¹⁸We use panel fixed effects regression to eliminate stable personality differences from the error. The error is thus made of test-retest error, and other time-varying factors affecting utility other than the ones we can control for. We assume here that this error is homoscedastic.

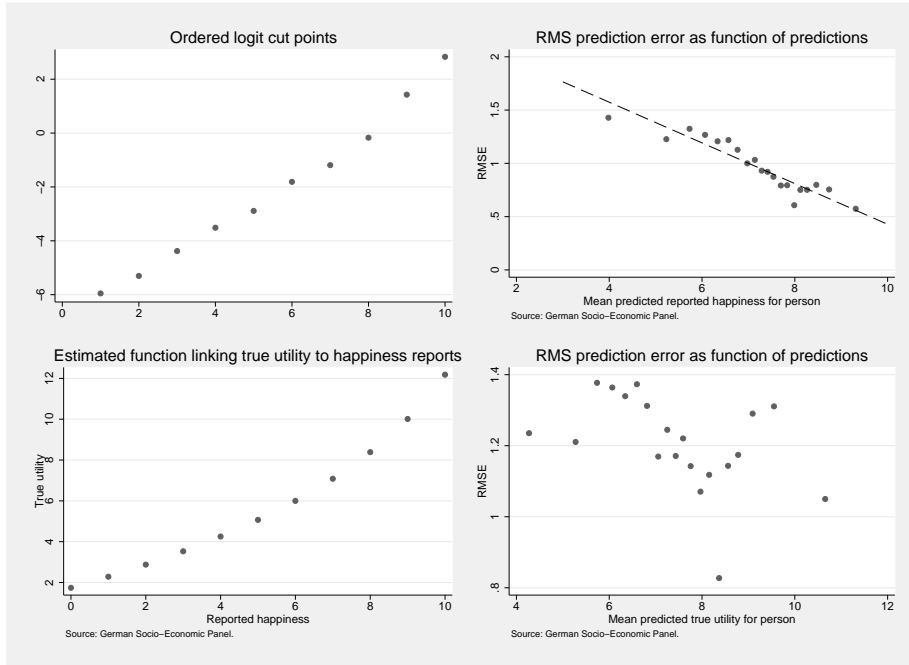


Fig. 4: Evidence for the concavity of $h = f(u)$ using the German Socio-Economic panel. (i) The top left panel shows the cut points in an ordered logit model. The x -axis denotes the number of steps in the reported happiness variable. The y -axis is arbitrary, so only the curvature is significant. The cut points of ordered probit are virtually indistinguishable from those of ordered logit. (ii) The top right panel shows the root mean squared prediction error as a function of mean prediction in a panel fixed effects regression. (iii) The bottom-left panel shows the function relating true utility to reported happiness that can be estimated from the information in the top-right panel by assuming homoscedastic test-retest errors in the true utility domain. (iv) The bottom-right panel shows the test-retest errors transformed into the true-utility domain.

$$u - u_0 = -\frac{\log(1 - c(h - h_0))}{c} \quad (18)$$

This function is plotted in the bottom left panel of Figure 4. The bottom, right panel repeats the plot of root mean square prediction errors against predictions, but this time in the u (true utility) domain. The lack of correlation is consistent with the assumption of homoscedastic test-retest error in the true utility domain.

To estimate the implied correction for ρ we first differentiate Equation 18 with respect to income:

$$\frac{\partial u}{\partial y} = \frac{1}{1 - c(h - h_0)} \frac{\partial h}{\partial y} \quad (19)$$

$$\frac{\partial^2 u}{\partial y^2} = \frac{1}{1 - c(h - h_0)} \frac{\partial^2 h}{\partial y^2} + \frac{c}{(1 - c(h - h_0))^2} \left(\frac{\partial h}{\partial y}\right)^2 \quad (20)$$

Denoting the elasticity of true marginal utility with respect to income by ρ_u , and that of marginal reported happiness by ρ_h we can write

$$\begin{aligned} \rho_u &= -\frac{(\partial^2 u / \partial y^2)}{\partial u / \partial y} y \\ &= -\frac{(\partial^2 h / \partial y^2)}{\partial h / \partial y} y - \frac{c}{1 - c(h - h_0)} \frac{\partial h}{\partial y} y \Big|_{h=h_0} \\ &= \rho_h - cy \frac{\partial h}{\partial y} \Big|_{h=h_0} \end{aligned} \quad (21)$$

Here c has units of $1/h$, so the correction term is dimensionless. As ρ_h was estimated using Equation 10 we have $\partial h / \partial y = y^{-\rho_h}$. The correction is therefore given by the following equation:

$$\rho_u - \rho_h = cy^{1-\rho_h} \Big|_{h=h_0} \quad (22)$$

Using the overall maximum likelihood value of $\rho_h = 1.26$ we estimated Equation 17 using a panel fixed effects regression on data from the German Socio-Economic Panel, finding $c = 0.153$. The value of $y^{1-\rho_h}$ for $h = h_0$ was found by taking the sample mean in the range $h_0 \pm 0.2$. Plugging in the numbers into Equation 22 results in an estimated correction term of -0.021 , or $\rho_u \approx 1.26 - 0.02 = 1.24$. An analogous estimate using the British Household Panel resulted in a correction term of -0.016 . Note that in principal, if c had been large enough, we could have obtained a corrected estimate of $\rho_u < 0$, that is a convex relationship between u and income (Oswald, 2005). The evidence, however, implies that in practice the relationship between true utility and income is concave and is not much different from that between reported happiness and income.

5.3 The spline test

The spline test is based on the idea that if the function f linking reported utility, h , with true utility, u , is linear, then the relationship between h and the linear regression prediction, \hat{h} , should have the same slope of 1

Table 6: Spline test results showing the estimated $\hat{\delta}$ in Equation 23 (t -score in parentheses), the F -score of the joint restriction $\beta = 0$ and $\delta = 0$, and the implied correction to $\hat{\rho}$.

	$\hat{\delta}$	F -test	Correction to $\hat{\rho}$
General Social Survey	-0.13 (-1.18)	2.08	-0.01
World Values Survey	-0.04 (-1.16)	0.67	-0.00
European Social Survey	-0.23 (-5.27)	13.97	-0.02
European Quality of Life Survey	-0.07 (-1.17)	0.84	-0.01
German Socio-Economic Panel	-0.30 (-6.05)	18.42	-0.09
British Household Panel Survey	-0.17 (-1.79)	10.46	-0.03
Combined	-0.08(-4.71)	13.17	-0.01

independently of the value of \hat{h} , and a constant term of zero. On the other hand, if f is concave, the slope estimates based on low values of \hat{h} should be higher than the slope estimates based on high values of \hat{h} .

In each survey we first estimate on the entire sample Equation 10 with ρ imposed at the overall maximum likelihood estimate of 1.26. We then identify the median value of \hat{h} , and create a dummy, d_i , which equals 1 if \hat{h}_i exceeds that median value, and 0 otherwise. We then estimate the following equation:

$$h_i = \alpha + \beta d_i + (\gamma + \delta d_i) \hat{h}_i \quad (23)$$

The values reported in Table 6 show that $\delta < 0$ and $\beta \neq 0$, consistent with f being concave. This result is not statistically significant in all surveys, but is strongly statistically significant in the combined sample (t -score = 4.71 for $\delta < 0$). The implied correction to $\hat{\rho}$ is -0.01, bringing the estimate down from 1.26 to 1.25. The calculations involved are explained below.

Using Equations 11 and 14 we have the following relationship between the elasticity of true marginal utility, u , and that of marginal reported happiness, h :

$$\rho_h = -y \frac{\partial^2 h / \partial y^2}{\partial h / \partial y} = \rho_u - \alpha y^{1-\rho_h} \frac{f''(\cdot)}{f'(\cdot)} \quad (24)$$

where α is the coefficient in Equation 14 and y is income. We estimate the correction to $\hat{\rho}$ computing the derivatives at the median prediction point, \hat{h}_0 , and using the average value of $y^{1-\rho_h}$ within a 0.2 distance of \hat{h}_0 .

To estimate the required derivatives we perform the following regression separately below and above median reported happiness:

$$h_i^p = \alpha^p + \beta^p \hat{h}_i^p \quad (25)$$

where $p \in \{l, h\}$ denotes whether \hat{h} is in the lower or upper part of the sample with respect to the median. We use the following linear approximation for the two halves of the sample:

$$h_i^p = f(\hat{h}_i^p) \approx \alpha^p + \beta^p \hat{h}_i^p \quad (26)$$

Using this approximation we obtain

$$\begin{aligned}
 f'(\hat{h}_0) &\approx \frac{f(\tilde{h}^h) - f(\tilde{h}^l)}{\tilde{h}^h - \tilde{h}^l} \\
 &= \frac{(\hat{\alpha}^h + \hat{\beta}^h \tilde{h}^h) - (\hat{\alpha}^l + \hat{\beta}^l \tilde{h}^l)}{\tilde{h}^h - \tilde{h}^l}
 \end{aligned} \tag{27}$$

where \tilde{h}^p denotes the median predicted happiness in sample half p .
 Similarly,

$$\begin{aligned}
 f''(\hat{h}_0) &\approx \frac{f'(\tilde{h}^h) - f'(\tilde{h}^l)}{\tilde{h}^h - \tilde{h}^l} \\
 &= \frac{\hat{\beta}^h - \hat{\beta}^l}{\tilde{h}^h - \tilde{h}^l}
 \end{aligned} \tag{28}$$

5.4 Why the scale compression matters so little

For the non-linear scale compression to have a big effect two conditions must hold:

1. Rich and poor people should typically be found at significantly different points of the utility scale.
2. The effect of utility on reported happiness should be substantially different at these different points.

In reality, however, neither of these conditions hold. First, income is only moderately correlated with utility. For example, in the GSOEP¹⁹ the regression coefficient of reported happiness on log income is 0.7, so that a quadrupling of income is necessary to lift a person one unit on the 0-10 utility scale. For it to have a significant effect over common income differences, the slope of the relationship between real utility and its reports would thus have to be substantially different at intervals of less than 1 on the utility scale, but, second, the slopes at these different points are in fact similar. Thus, although we have found evidence for a non-linear relationship between real utility and its reports, this non-linear relationship can only have a modest impact on the results of Section 4.

6 Conclusion

A key parameter for policy analysis is the elasticity (ρ) of the marginal utility of income with respect to the level of income. To estimate this, we have used six different surveys, including country studies of USA, Britain and Germany, and multi-country studies involving most first-world countries and in one case third-world countries also.

We have found a striking uniformity in the estimates obtained from these totally different surveys. The lowest estimate was 1.19 and the

¹⁹The simple correlation between reported happiness and log income is about 0.2.

highest 1.30. This uniformity is in contrast to the wide range of estimates obtained from behavioural studies.

Our overall estimate is 1.26 (with a 95% confidence interval of 1.16–1.37). This estimate falls well-within the 95% confidence interval of each of the individual surveys. It is also well-within the 95% confidence interval of each sub-group of the population (sex, age, education, marital status), so that any behavioural differences between these groups cannot be attributed to differences in the curvature of the experienced utility-income relationship.

To interpret this value of ρ , it is useful to start from the logarithmic ($\rho = 1$) where the marginal utility of income is inversely proportional to income. Our higher estimate of ρ implies that marginal utility falls somewhat faster than that.

However, there is the obvious question, Do the answers correspond to the true level of utility, or do respondents in their replies use a happiness scale which is convex or concave with respect to their true utility? We investigate this in three ways, with similar results. In each case we find under certain assumptions that true utility is convex with respect to reported happiness, so that true marginal utility declines less fast than reported marginal utility. However, we can also compute the size of the difference. It turns out that the maximal implied correction would only reduce our estimate of ρ to 1.24.

We have thus confirmed the (cardinalist) assumption of nineteenth century economists that marginal utility of income declines with income. Given a number of assumptions, we have been able to estimate a numerical value for the rate at which this occurs. This value can be used to inform the distributional weights in cost-benefit analysis and the analysis of optimal tax.

References

- Atkinson, A. and Stiglitz, J. (1980). *Lectures in Public Economics*, McGraw Hill.
- Attanasio, O. (1999). Consumption, *Handbook of Macroeconomics* **1**: 741–812.
- Bernoulli, D. (1738). Specimen theoriae novae de mensura sortis, *Commentarii academiae scientiarum imperialis Petropolitanae*(for 1730 and 1731) **5**: 175–192.
- Bernoulli, D. (1954). Exposition of a New Theory on the Measurement of Risk, *Econometrica* **22**: 23–36.
- Blanchflower, D. and Oswald, A. (2004). Well-being over time in Britain and the USA, *Journal of Public Economics* **88**(7-8): 1359–1386.
- Blundell, R., Browning, M. and Meghir, C. (1994). Consumer Demand and the Life-Cycle Allocation of Household Expenditures, *The Review of Economic Studies* **61**(1): 57–80.

- Chetty, R. (2006). *A Bound on Risk Aversion Using Labor Supply Elasticities*, National Bureau of Economic Research Cambridge, Mass., USA.
- Clark, A., Frijters, P. and Shields, M. (2006). Income and Happiness: Evidence, Explanations and Economic Implications, *unpublished paper*, 8th August .
- Clark, A., Frijters, P. and Shields, M. (forthcoming). Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles, *Journal of Economic Literature* .
- Clark, A. and Oswald, A. (1996). Satisfaction and comparison income, *Journal of Public Economics* **61**(3): 359–381.
- Cohen, A. and Einav, L. (2005). *Estimating Risk Preferences from Deductible Choice*, National Bureau of Economic Research Cambridge, Mass., USA.
- Dalton, H. (1920). The measurement of the inequality of incomes, *Economic Journal* **30**(119): 348–61.
- Davidson, R. (1992). Emotion and affective style: Hemispheric substrates, *Psychological Science* **3**: 39–43.
- Davidson, R. (2000). Affective style, psychopathology and resilience: Brain mechanisms and plasticity, *American Psychologist* **55**: 1196–1214.
- Davidson, R., Jackson, D. and Kalin, N. (2000). Emotion, plasticity, context and regulation: Perspectives from affective neuroscience, *Psychological Bulletin* **126**: 890–906.
- Di Tella, R., MacCulloch, R. and Oswald, A. (2003). The macroeconomics of happiness, *The Review of Economics and Statistics* **85**(4): 809–827.
- Diener, E. and Suh, E. (1999). National differences in subjective well-being, in D. Kahneman, E. Diener and N. Schwartz (eds), *Well-Being: the foundations of hedonic psychology*, Russel Sage Foundation, New-York.
- Dohmen, T. et al. (2006). *Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-validated Survey*, Centre for Economic Policy Research.
- Easterlin, R. (1995). Will Raising the Incomes of All Increase the Happiness of All?, *Journal of Economic Behavior and Organization* **27**(1): 35–47.
- Gardner, J. and Oswald, A. (2007). Money and Mental Wellbeing: A Longitudinal Study of Medium-Sized Lottery Wins, *Journal of Health Economics* **26**: 49–60.
- Griliches, Z. and Hausman, J. (1986). Errors in variables in panel data, *Journal of econometrics* **31**(1): 93–118.

- Helliwell, J. (2003). How's life? Combining individual and national variables to explain subjective well-being, *Economic Modelling* **20**(2): 331–60.
- Johnson, K. O., Hsiao, S. S. and Yoshioka, T. (2002). Neural coding and the basic law of psychophysics, *The Neuroscientist* **8**(2): 111–121.
- Kahneman, D. (1999). Objective happiness, in D. Kahneman, E. Diener and N. Schwartz (eds), *Well-Being: the foundations of hedonic psychology*, Russel Sage Foundation, New-York.
- Kahneman, D., Diener, E. and Schwarz, N. (1999). *Well-being: The Foundations of Hedonic Psychology*, Russell Sage Foundation.
- Lanot, G., Hartley, R. and Walker, I. (2006). Who really wants to be a millionaire: estimates of risk aversion from game show data, *Warwick Economic Research Paper No. 719*.
- Mehra, R. and Prescott, E. (1985). The Equity Premium, *Journal of Monetary Economics* **15**(2): 145–61.
- Mehra, R. and Prescott, E. (2003). The Equity Premium in Retrospect, in G. Constantinides, M. Harris and S. R. (eds), *Handbook of the Economic of Finance*, 1 edn, Vol. 1, Elsevier, chapter 14, pp. 889–938.
- Metrick, A. (1995). A Natural Experiment in "Jeopardy!", *The American Economic Review* **85**(1): 240–253.
- Oswald, A. (2005). On the common claim that happiness equations display diminishing marginal utility of income, *Iza Discussion Paper* (1781).
- Stevens, S. (1975). *Psychophysics: Introduction to its perceptual, neural, and social prospects*, John Wiley and Sons.
- Van Praag, B. (2004). *Happiness Quantified: A Satisfaction Calculus Approach*, Oxford University Press.
- Winkelmann, L. and Winkelmann, R. (1998). Why Are the Unemployed So Unhappy? Evidence from Panel Data, *Economica* **65**(257): 1–15.

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

- | | | |
|-----|--|---|
| 783 | Gustavo Crespi
Chiara Criscuolo
Jonathan Haskel | Information Technology, Organisational Change
and Productivity Growth: Evidence from UK Firms |
| 782 | Paul Castillo
Carlos Montoro
Vicente Tuesta | Inflation Premium and Oil Price Volatility |
| 781 | David Metcalf | Why Has the British National Minimum Wage Had
Little or No Impact on Employment? |
| 780 | Carlos Montoro | Monetary Policy Committees and Interest Rate
Smoothing |
| 779 | Sharon Belenzon
Mark Schankerman | Harnessing Success: Determinants of University
Technology Licensing Performance |
| 778 | Henry G. Overman
Diego Puga
Matthew A. Turner | Decomposing the Growth in Residential Land in the
United States |
| 777 | Florence Kondylis | Conflict-Induced Displacement and Labour Market
Outcomes: Evidence from Post-War Bosnia and
Herzegovina |
| 776 | Willem H. Buiter | Is Numéraireology the Future of Monetary
Economics? <i>Unbundling numéraire and medium of
exchange through a virtual currency and a shadow
exchange rate</i> |
| 775 | Francesco Caselli
Nicola Gennaioli | Economics and Politics of Alternative Institutional
Reforms |
| 774 | Paul Willman
Alex Bryson | Union Organization in Great Britain
<i>Prepared for symposium for the <u>Journal of Labor
Research</u> on "The State of Unions: A Global
Perspective"</i> |
| 773 | Alan Manning | The Plant Size-Effect: Agglomeration and
Monopsony in Labour Markets |
| 772 | Guy Michaels | The Effect of Trade on the Demand for Skill –
Evidence from the Interstate Highway System |
| 771 | Gianluca Benigno
Christoph Thoenissen | Consumption and Real Exchange Rates with
Incomplete Markets and Non-Traded Goods |
| 770 | Michael Smart
Daniel M. Sturm | Term Limits and Electoral Accountability |
| 769 | Andrew B. Bernard
Stephen J. Redding
Peter K. Schott | Multi-Product Firms and Trade Liberalization |

- | | | |
|-----|--|---|
| 768 | Paul Willman
Alex Bryson | Accounting for Collective Action: Resource Acquisition and Mobilization in British Unions |
| 767 | Anthony J. Venables | Shifts in Economic Geography and their Causes |
| 766 | Guy Michaels | The Long-Term Consequences of Regional Specialization |
| 765 | Fabrice Murtin | American Economic Development Since the Civil War or the Virtue of Education |
| 764 | Carlo Rosa
Giovanni Verga | The Impact of Central Bank Announcements on Asset Prices in Real Time: Testing the Efficiency of the Euribor Futures Market |
| 763 | Benjamin Aleman-Castilla | The Effect of Trade Liberalization on Informality and Wages: Evidence from Mexico |
| 762 | L. Rachel Ngai
Roberto M. Samaniego | An R&D-Based Model of Multi-Sector Growth |
| 761 | Mariano Bosch | Job Creation and Job Destruction in the Presence of Informal Labour Markets |
| 760 | Christian Hilber
Frédéric Robert-Nicoud | Owners of Developed Land Versus Owners of Undeveloped Land: Why Land Use is More Constrained in the Bay Area than in Pittsburgh |
| 759 | William Nickell | The CEP-OECD Institutions Data Set (1060-2004) |
| 758 | Jean Eid
Henry G. Overman
Diego Puga
Matthew Turner | Fat City: the Relationship Between Urban Sprawl and Obesity |
| 757 | Christopher Pissarides | Unemployment and Hours of Work: the North Atlantic Divide Revisited |