

The marginal utility of income[☆]

R. Layard^a, S. Nickell^b, G. Mayraz^{c,*}

^a *Centre for Economic Performance, London School of Economics*

^b *Department of Economics and Centre for Economic Performance, London School of Economics*

^c *Nuffield College, Oxford, and Centre for Economic Performance, London School of Economics*

Received 30 March 2007; received in revised form 17 January 2008; accepted 21 January 2008

Available online 31 January 2008

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. Thus, on the basis of our estimates, the marginal utility of income declines somewhat faster than in proportion to the rise in income.

© 2008 Elsevier B.V. All rights reserved.

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

Keywords: Marginal utility; Income; Life satisfaction; Happiness; Public economic; Welfare; Inequality; Optimal taxation; Reference-dependent preferences

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 income.¹

[☆] We thank the Esmee Fairbairn Foundation for financial support to the CEP well-being programme.

* Corresponding author.

E-mail address: g.mayraz@lse.ac.uk (G. Mayraz).

¹ This 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).

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 assumption.²

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 based on choice under uncertainty range from 0 to 10 or even higher.³ Estimates based on intertemporal choice vary less,⁴ but in any case the assumption of intertemporal additivity is problematic.⁵ 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 happiness,⁶ 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 Section 4 we give our results, and conclude in Section 5.

² Dalton 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 $\rho \neq 1$.

³ [Lanot 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](#)).

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

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

⁶ See, 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. in press](#).

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:

	0	1	2	3	4	5	6	7	8	9	10
extremely unhappy											extremely happy

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 Fig. 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 (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

$$h_i = f_i(u_i) \quad (3)$$

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

$$h_i = f_i(u_i) = f(u_i) + v_i \quad (4)$$

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

Table 1
Description of surveys as used

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 Panel 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

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.

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 Social Survey	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 Social Survey	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	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	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	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	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".]

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 behaviours.⁷ Overall, in light of this evidence, we feel quite justified in making the assumption embedded in Eq. (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 Eq. (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.⁸ Here, we simply assume this possibility away, and suppose that f is linear. Thus we assume

$$h_i = u_i + v_i \quad (5)$$

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 Fig. 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 transformation.⁹ When both happiness and satisfaction with life as a whole

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

⁸ In Layard et al. (2007) we pursue this issue at some length under a variety of different assumptions. To summarise, we reject the linearity of f but the extent of non-linearity is very low, and the impact on the measured rate of decline of marginal utility with respect to income is negligible.

⁹ For 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 Survey.¹⁰ 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 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.

¹⁰ The 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.

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 periods.¹¹

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 value could lead to misleading conclusions about the long-run relationship between utility and income (consistent with this, the plots of Fig. 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 Eq. (4).
2. This transformation is linear, as in Eq. (5). That is, $h_{it} = u_{it} + v_{it}$.
3. The experienced utility of person i at time t in country c can be written as follows:¹²

$$u_{it} = \alpha_{ct} \frac{y_{it}^{1-\rho} - 1}{1-\rho} + \sum_j \beta_j x_{jit} + \gamma_i + \gamma_{ct} + \varepsilon'_{it} \quad (6)$$

where y 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

¹¹ The 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).

¹² Changes 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.




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

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 Eq. (6) 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 .¹³

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.

4. 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 . First, we combine Eq. (6) with the assumed linear relationship between reported happiness and experienced utility (Eq. (5)) to 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 $\epsilon_{it} = \epsilon'_{it} + v_{it}$. Our strategy now includes 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 Eq. (7).

¹³ If 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 $\rho=1$, 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 .

4. Main results

Fig. 2 shows the simple cross-sectional relationships between happiness and income, and between happiness and log 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

$$h_{it} = \alpha \log y_{it} + \sum_j \beta_j x_{jit} + \gamma_{ct} + \varepsilon_{it} \tag{8}$$

The fit is quite good and is plotted in Fig. 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 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:

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

As can be seen, we have now normalized 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




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)

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)

Relative income is relative to mean income in country/year. Country/year dummies and other personal controls are included but not shown.

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 Eqs. (8) 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 Fig. 3 and Table 3, are similar to those of the cross-sectional analysis of the previous section, a result consistent with the causal interpretation¹⁴. 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 expected.¹⁵

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 Eq. (7), that is

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

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

¹⁴ In 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.

¹⁵ Income 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

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)

95% confidence intervals in parentheses.

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 uncertainty, we find that estimates of ρ based on happiness surveys yield a single estimate that is consistent with very different data sets. In order to estimate the influence of the actual happiness scores on our results, we repeated the estimation using an ordered logit model which simply makes use of the happiness rankings implicit in the scores. The overall ML estimate of ρ is now 1.23, very close to our previous estimate of 1.26.

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 status¹⁶. 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 a 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 behaviour¹⁷, but imply that such differences are likely to be due to factors other than underlying differences in experienced utility.

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, hours², and hours³. In so doing we also kept a dummy for being unemployed, as distinct from being outside the labour force. The result of this

¹⁶ In 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.

¹⁷ There 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
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)

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 $\rho=1.32$, which lies within the 95% confidence interval for the standard specification.

5. 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 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 subgroup 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.

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., Stiglitz, J., 1980. Lectures in Public Economics. Mc-Graw 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., 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., Meghir, C., 1994. Consumer Demand and the Life-Cycle Allocation of Household Expenditures. *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., Shields, M., in press. Relative income, happiness and utility: an explanation for the Easterlin Paradox and other puzzles. *Journal of Economic Literature*.

- Clark, A., Oswald, A., 1996. Satisfaction and comparison income. *Journal of Public Economics* 61 (3), 359–381.
- Cohen, A., 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), 348–361.
- 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., Kalin, N., 2000. Emotion, plasticity, context and regulation: Perspectives from affective neuroscience. *Psychological Bulletin* 126, 890–906.
- Di Tella, R., MacCulloch, R., Oswald, A., 2003. The macroeconomics of happiness. *Review of Economics and Statistics* 85 (4), 809–827.
- Diener, E., Suh, E., 1999. National differences in subjective well-being. In: Kahneman, D., Diener, E., Schwartz, N. (Eds.), *Well-being: The Foundations of Hedonic Psychology*. Russell 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., Oswald, A., 2007. Money and Mental Wellbeing: A Longitudinal Study of Medium-Sized Lottery Wins. *Journal of Health Economics* 26, 49–60.
- Griliches, Z., 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), 331–360.
- Kahneman, D., 1999. Objective happiness. In: Kahneman, D., Diener, E., Schwartz, N. (Eds.), *Well-being: The Foundations of Hedonic Psychology*. Russell Sage Foundation, New York., pp. 213–242.
- Kahneman, D., Diener, E., Schwarz, N., 1999. *Well-being: The Foundations of Hedonic Psychology*. Russell Sage Foundation.
- Lanot, G., Hartley, R., Walker, I., 2006. Who really wants to be a millionaire: estimates of risk aversion from game show data. Warwick Economic Research Paper vol. 719.
- Layard, R., Mayraz, G., Nickell, S., 2007. The marginal utility of income. Centre for Economic Performance Discussion Paper no. 784.
- Mehra, R., Prescott, E., 1985. The Equity Premium. *Journal of Monetary Economics* 15 (2), 145–161.
- Mehra, R., Prescott, E., 2003. The equity premium in retrospect. In: Constantinides, G., Harris, M., Stulz, R. (Eds.), 1st ed. *Handbook of the Economics of Finance*, vol. 1, pp. 889–938. chapter 14. Elsevier.
- Metrick, A., 1995. A Natural Experiment in “Jeopardy!”. *American Economic Review* 85 (1), 240–253.
- Van Praag, B., 2004. *Happiness Quantified: A Satisfaction Calculus Approach*. Oxford University Press.
- Winkelmann, L., Winkelmann, R., 1998. Why Are the Unemployed So Unhappy? Evidence from Panel Data. *Economica* 65 (257), 1–15.