

Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox*

Betsey Stevenson

Wharton, University of Pennsylvania
CESifo and NBER

betsey.stevenson@wharton.upenn.edu
<http://bpp.wharton.upenn.edu/betseys>

Justin Wolfers

Wharton, University of Pennsylvania
CEPR, CESifo, IZA and NBER

jwolfers@nber.org
www.nber.org/~jwolfers

Abstract

The “Easterlin paradox” suggests that there is no link between a society’s economic development and its average level of happiness. We re-assess this paradox analyzing multiple rich datasets spanning many decades. Using recent data on a broader array of countries, we establish a clear positive link between average levels of subjective well-being and GDP per capita across countries, and find no evidence of a satiation point beyond which wealthier countries have no further increases in subjective well-being. We show that the estimated relationship is consistent across many datasets and is similar to the relationship between subject well-being and income observed within countries. Finally, examining the relationship between changes in subjective well-being and income over time within countries we find economic growth associated with rising happiness. Together these findings indicate a clear role for absolute income and a more limited role for relative income comparisons in determining happiness.

This draft: August 19, 2008

First draft: 2/15/2008

Keywords: Happiness, subjective well-being, Easterlin Paradox, life satisfaction, economic growth, well-being-income gradient, hedonic treadmill

JEL codes: D6, I3, J1

* The authors would like to thank Gary Becker, David Blanchflower, Angus Deaton, Richard Easterlin, Carol Graham, Daniel Kahneman, Alan Krueger, David Laibson, Andrew Oswald, and Luis Rayo for useful discussions, and Gale Muller and his colleagues at Gallup for access to and help with the Gallup World Poll. Rohak Doshi and Michael L. Woodford provided excellent research assistance. We would like to thank the Zicklin Center for Business Ethics Research and the Zell/Lurie Real Estate Center for generous research support. The data and Stata programs used in this paper are available for download from the authors’ homepages.

I. Introduction

Economic growth has long been considered an important goal of economic policy, yet in recent years some have begun to argue against further trying to raise the material standard of living, claiming that such increases will do little to raise well-being. These arguments are based on a key finding in the emerging literature on subjective well-being, called the “Easterlin paradox,” which suggests that there is no link between the level of economic development of a society and the overall happiness of its members. In several papers Richard Easterlin has examined the relationship between happiness and GDP both across countries and within individual countries through time (1974, 1995, 2005a, 2005b). In both types of analysis he finds little significant evidence of a link between aggregate income and average happiness.

In contrast, there is robust evidence that *within* countries those with more income are happier. These two seemingly discordant findings—that income is an important predictor of individual happiness, yet apparently irrelevant for average happiness—have spurred researchers to seek to reconcile them through models emphasizing reference-dependent preferences and relative income comparisons (Easterlin, 1973, p. 4).² Layard offers an explanation: “people are concerned about their relative income and not simply about its absolute level. They want to keep up with the Joneses or if possible to outdo them” (Layard, 2005a, p. 45). While leaving room for absolute income to matter for some people, Layard and others have argued that absolute income is only important for happiness when income is very low. Layard argues, for example, that “once a country has over \$15,000 per head, its level of happiness appears to be independent of its income per head” (2003, p. 17).³

The conclusion that absolute income has little impact on happiness has far-reaching policy implications. If economic growth does little to improve social welfare, then it should not be a primary goal of government policy. Indeed, Easterlin argues that his analysis of time trends in subjective well-being “undermine the view that a focus on economic growth is in the best interests of society” (2005c, p. 441). Layard argues for an explicit government policy of maximizing subjective well-being (2005a).⁴ Moreover, he notes that relative income comparisons imply that each individual’s labor

² Easterlin summarizes his findings:

“In all societies, more money for the individual typically means more individual happiness. However, raising the incomes of all does not increase the happiness of all. The happiness-income relation provides a classic example of the logical fallacy of composition—what is true for the individual is not true for society as a whole.

The resolution of this paradox lies in the relative nature of welfare judgments. Individuals assess their material well-being, not in terms of the absolute amount of goods they have, but relative to a social norm of what goods they ought to have (italics in original).”

Layard (1980, p. 737) is more succinct: “a basic finding of happiness surveys is that, though richer societies are not happier than poorer ones, within any society happiness and riches go together.” For a recent review of the use of reference-dependent preferences to explain these observations, see Clark, Frijters, and Shields (2008).

³ For other arguments proposing a satiation point in happiness, see Veenhoven (1991) Clark, Frijters, and Shields (2008); Frey and Stutzer (2002).

⁴ For a concurring view from the positive psychology movement, see Diener and Seligman (2004).

effort imposes negative externalities on others (by shifting their reference points) and that these distortions would be best corrected by higher taxes on income or consumption.

Evaluating these strong policy prescriptions demands a robust understanding of the true relationship between income and well-being. Unfortunately, the present literature is based on fragile and incomplete evidence about this relationship. At the time the Easterlin paradox was first identified, few data were available to allow an assessment of subjective well-being across countries and through time. The difficulty of identifying a robust GDP-happiness link from scarce data led some to confound the absence of evidence of such a link with evidence of its absence.

The ensuing years have seen an accumulation of cross-country data recording individual life satisfaction and happiness. These recent data (and a reanalysis of earlier data) suggest that the case for a link between economic development and happiness is quite robust. The key to our findings is a resolute focus on the magnitude of the subjective well-being-income gradient estimated within and across countries at a point in time as well as over time, rather than its statistical significance or insignificance.

Our key result is that the estimated subjective well-being-income gradient is not only significant but also remarkably robust across countries, within countries, and over time. These comparisons between rich and poor members of the same society, between rich and poor countries, and within countries through time as they become richer or poorer all yield similar estimates of the well-being-income gradient. Our findings both put to rest the earlier claim that economic development does not raise subjective well-being and undermine the possible role played by relative income comparisons.

These findings invite a sharp reassessment of the stylized facts that have informed economic analysis of subjective well-being data. Across the world's population, variation in income explains a sizable proportion of the variation in subjective well-being. There appears to be a very strong relationship between subjective well-being and income, which holds for both rich and poor countries, falsifying earlier claims of a satiation point at which higher GDP per capita is not associated with greater well-being.

The rest of this paper is organized as follows. The first section provides some background on the measurement of subjective well-being and economic analysis of these data. Subsequent sections are organized around alternative measurement approaches to assessing the link between income and well-being. Thus, the second section compares average well-being and income across countries. Whereas earlier studies focused on comparisons of small numbers of industrialized countries, newly available data allow comparisons across countries at all levels of development. These comparisons show a powerful effect of national income in explaining variation in subjective well-being across countries. In the third section we confirm the earlier finding that richer people within a society are typically happier than their poorer brethren. Because these national cross sections typically involve quite large samples, this finding is extremely statistically significant and has not been widely disputed. However, Easterlin (1974) and others have argued strongly that the positive relationship between income and subjective well-being *within* countries is much larger than that seen across countries. This argument is not borne out by the data: the well-being-income gradient measured within countries is similar to that measured between countries. The paper's fourth section extends our analysis to assessing national time-series movements

in average well-being and income. Consistent time series measuring subjective well-being data are scarce, and the existing data is noisy. These factors explain why past researchers have not found a link between economic growth and growth in happiness. We reexamine three of the key case studies from previous research and find that a more careful assessment of the experiences of Japan, Europe, and the United States does not undermine the claim of a clear link between economic growth and happiness, a finding supported by repeated international cross-sections. Our point estimates suggest that the link may be similar to that found in cross-country comparisons, although substantial uncertainty remains around these estimates. The fifth section briefly explores alternative measures of well-being.

II. Some Background on Subjective Well-Being and Income

Our strategy in this paper is to use all of the important large-scale surveys now available to assess the relationship between subjective well-being and happiness. These surveys typically involve questions probing happiness or life satisfaction. The World Values Survey, for example, asks, “Taking all things together, would you say you are: very happy; quite happy; not very happy; not at all happy?” and, “All things considered, how satisfied are you with your life as a whole these days?” Other variants of the question, such as that in the Gallup World Poll, employ a ladder analogy: interviewees are asked to imagine a ladder with each rung representing a successively better life. Respondents then report the “step” on the ladder that best represents their life.

These questions (and many other variants) are typically clustered under the rubric of “subjective well-being” (Diener, 2006, pp. 399-400).⁵ Although the validity of these measures remains a somewhat open question, a variety of evidence points to a robust correlation between answers to subjective well-being questions and more objective measures of personal well-being. For example, answers to subjective well-being questions have been shown to be correlated with physical evidence of affect such as smiling, laughing, heart rate measures, sociability, and electrical activity in the brain (Diener, 1984). Measures of individual happiness or life satisfaction are also correlated with other subjective assessments of well-being such as independent evaluations by friends, self-reported health, sleep quality, and personality (Diener, Lucas, and Scollon, 2006; Kahnman and Krueger, 2006). Subjective well-being is a function of both the individual’s personality and his or her reaction to life events. One would therefore expect an individual’s happiness to be somewhat stable over time, and accurate measurements of subjective well-being to have high test-retest correlations, which indeed they do (Eid and Diener, 2006). Self-reports of happiness have also been shown to be correlated in the expected direction with changes in life circumstances. For example, an individual’s subjective well-being typically rises with marriage and income growth and falls while going through a divorce.

⁵ Diener suggests that “subjective well-being refers to all of the various types of evaluations, both positive and negative, that people make of their lives. It includes reflective cognitive evaluations, such as life satisfaction and work satisfaction, interest and engagement, and affective reactions to life events, such as joy and sadness. Thus, subjective well-being is an umbrella term for the different valuations people make regarding their lives, the events happening to them, their bodies and minds, and the circumstances in which they live.”

Although the results from each of these approaches suggest that cross-sectional comparisons of people within a population have some validity, there is less evidence about the validity of comparisons across populations, which can be confounded by translation problems and cultural differences. Many researchers have argued for the possibility of a biologically based set of emotions that are universal to humans and appear in all cultures (Diener and Tov, 2008). Research has found that people across cultures clearly recognize emotions such as anger, sadness, and joy when displayed in others' facial expressions (Ekman and Friesen, 1971; Ekman, et al., 1987). Studies have also found that when people around the globe are asked about what is required for more happiness or life satisfaction, the answers are strikingly uniform: money, health, and family are said to be the necessary components of a good life (Easterlin, 1974). Diener and Tov (2008) argue that it is this possibility of biologically based universal emotions that suggests that well-being can be compared across societies.

A similar argument applies to making comparisons of subjective well-being within countries over time. One difficulty with time-series assessments is the possibility that small changes in how people perceive or answer questions about their happiness may be correlated with changes in the outcomes—such as income—whose relationship with subjective well-being one wishes to assess. The evidence regarding aggregate changes in happiness over time is inconsistent. Aggregate happiness has been shown to fall when unemployment and inflation rise, and to move in the expected direction with the business cycle (Di Tella, MacCulloch, and Oswald, 2003; Wolfers, 2003). However, on average, women in both the United States and Europe report declining happiness relative to men over recent decades, a finding that is difficult to reconcile with changes in objective conditions (Stevenson and Wolfers, 2007). Finally, this current paper is motivated by a desire to better understand the failure of past studies to isolate a link between happiness and economic growth.

A largely under-acknowledged problem in making intertemporal comparisons is simply the difficulty in compiling sufficiently comparable data. For instance, Smith (1986) shows that small changes in the ordering of questions on the U.S. General Social Survey led to large changes in reported happiness. These same data seem to show important day of week and seasonal cycles as well. Another difficulty with intertemporal comparisons is that attempts to cobble together long time series (such as for Japan, the United States, or China) often involve important coding breaks. Many of these issues simply add measurement error, making statistically significant findings more difficult to obtain. However, when scarce data are used to make strong inferences about changes in well-being over decades, even small amounts of measurement error can lead to misleading inferences.

To date, much of the economics literature assessing subjective well-being has tended to use measures of “life satisfaction” and “happiness” interchangeably. The argument for doing so is that these alternative measures of well-being are highly correlated and have similar covariates. However, they capture somewhat different concepts, with happiness more related to affect whereas satisfaction is more evaluative. The psychology literature has tended to treat questions probing affect as distinct from more evaluative assessments. We will consider both the income-happiness and income-satisfaction links in parallel. A subtle measurement issue is also involved, in that many of the surveys asking individuals

about their happiness provide a shorter scale of answers (such as “very happy,” “pretty happy,” and “not so happy”) than do those asking typical life satisfaction questions (which often use the “ladder” technique described above).

A final measurement issue to consider is the likely functional form of the relationship between subjective well-being and income. Most early studies considered the relationship between the level of absolute income and the level of happiness, and thus often found a curvilinear relationship. In some cases the lack of evidence of a clear linear relationship between GDP per capita and happiness led to theories of a satiation point, beyond which more income would not increase happiness. A more natural starting point might be to represent well-being as a function of the logarithm of income rather than absolute income. And indeed, recent research has shown that *within* countries “the supposed attenuation at higher income levels of the happiness-income relation does not occur when happiness is regressed on log income, rather than absolute income” (Easterlin, 2001, p. 468). However, if happiness is linearly related to log income in the within-country cross section, then cross-country studies should also examine the relationship between average levels of subjective well-being and average levels of log income. If economic development raises individual incomes equiproportionately, then average log income will rise or fall in tandem with the log of average income. Thus, most of our analysis assesses the relationship across countries between well-being and the log of GDP per capita, which is (surprisingly enough) a departure from much of the literature.⁶ Throughout our analysis we make heavy use of bivariate scatterplots and nonparametric regression techniques in order to allow the reader to assess the appropriate functional forms visually.

Finally, as in the existing literature, our analysis of the relationship between happiness and income involves an assessment of correlations rather than an attempt to establish tight causal links. Thus, our aim is simply to sort out the stylized facts about the link between income and well-being. Several interesting variants of the question could be asked—such as whether it is GDP, broader measures of economic development, or alternatively, changes in output or in productivity that drive happiness. Unfortunately, we lack the statistical power to resolve these questions.

III. Cross-Country Comparisons of Income and Well-Being

In his seminal 1974 paper Easterlin asked whether “richer countries are happier countries” (1974, p. 104). Examining two international datasets, he found a relationship across countries between aggregate happiness and income that he described as “ambiguous” and, although perhaps positive, small (*ibid.*, p.108). Subsequent research began to show a more robust positive relationship between a country’s income and the happiness of its people, leading Easterlin to later conclude that “a positive happiness-income relationship typically turns up in international comparisons” (1995, p. 42). However, this relationship has been argued as prevailing only over low levels of GDP per capita; once wealthy countries have satisfied basic needs, they have been described as on the “‘flat of the curve,’ with additional income buying little if any extra happiness” (Clark, Frijters, and Shields, 2008, p. 96). Although the literature has largely settled on the view that

⁶ Previous authors examining the relationship between well-being and log GDP include Easterlin (1995), Leigh and Wolfers (2006), and Deaton (2008); however, explicit discussion of the appropriate functional form is quite rare.

aggregate happiness rises with GDP for low-income countries, there is much less consensus on the magnitude of this relationship, or on whether a satiation point exists beyond which further increases in GDP per capita are associated with no change in aggregate happiness (Deaton, 2008).⁷

The early cross-country studies of income and happiness tended to be based on only a handful of countries, often with rather similar income per capita, and hence did not lend themselves to definitive findings. In addition, as the relationship between subjective well-being and the log of income is approximately linear, the analysis in terms of absolute levels of GDP per capita likely contributed to the lack of clarity around the relationship between income and happiness among wealthier countries. As we will show, new large-scale datasets covering many countries point to a clear, robust relationship between GDP per capita and average levels of subjective well-being in a country. Furthermore, we find no evidence that countries become satiated—the positive income-happiness relationship holds for both developed and developing nations.

Our macroeconomic analysis focuses on measures of real GDP per capita measured at purchasing power parity. For most countries we use the most recent data from the World Bank’s World Development Indicators database; where we are missing data, we refer to the Penn World Tables (version 6.2) and, failing that, the CIA Factbook. For earlier years we use data from Maddison (2007).⁸ The average of log income per person may be a more desirable aggregate than the log of average income, and so in some specifications we also account for the difference between these measures (also known as the mean log deviation).

Measuring average levels of subjective well-being is somewhat more difficult, as this typically involves aggregating individual responses to a qualitative question. Moreover, we wish to make comparisons across surveys that contain subjective well-being questions with varying numbers of categories for the responses. To do this, we need to convert the subjective well-being measures to a normalized measure, which we do through the use of ordered probit regressions of happiness on a series of country (or country-year) fixed effects (with no other controls), and then treat these fixed effects as average levels of well-being within a country (or country-year).⁹ Appendix A compares our ordered probit index with four alternative approaches to cardinalizing both life satisfaction and happiness, demonstrating that these alternatives yield highly correlated well-being aggregates. The distinct advantage of the ordered probit is that coefficients can be interpreted relative to the dispersion of the distribution of latent well-being in the population. As such, our ordered probit index should be interpreted as highlighting differences in average levels of happiness or life satisfaction between countries, relative to the pooled within-country standard deviation.

⁷ Deaton finds no evidence of a satiation point. His analysis of the 2006 Gallup World Poll finds a strong relationship between log GDP and happiness that is, if anything, stronger among high-income countries.

⁸ When filling in missing years, we interpolate using the annual percentage changes listed in the Penn World Tables. When filling in missing countries, we apply the ratio of a country’s GDP per capita to U.S GDP per capita, using data from the Penn World Tables or the CIA Factbook, to the World Bank data.

⁹ Throughout, we use suggested surveys weights to ensure our estimates are nationally representative for each country in each wave.

We present our analysis chronologically, so that the reader may see how the literature has progressed. To allow easy visual comparisons, we use a similar scale when graphing happiness and GDP and try to keep this scale consistent throughout the paper.

The top row of graphs in Figure 1 shows the three earliest cross-country comparisons of subjective well-being of which we are aware. Each of these comparisons is based on only four to nine countries, which were similar in terms of economic development. As a consequence, these comparisons yield quite imprecise estimates of the link between happiness and GDP. We have provided two useful visual devices to aid in interpretation: a dashed line showing the ordinary least squares regression line (our focus), and a shaded area that shows a central part of the happiness distribution, with a width equal to the cross-sectional standard deviation.

The graphs in the second row of Figure 1 show the cross-country comparisons presented by Easterlin (1974).¹⁰ Analyzing the 1960 data, Easterlin argues that “the association between wealth and happiness indicated by Cantril’s international data is not so clear-cut.... The inference about a positive association relies heavily on the observations for India and the United States” (ibid., p. 108).¹¹ Turning to the 1965 World Survey III data, Easterlin argues that “The results are ambiguous.... If there is a positive association between income and happiness, it is certainly not a strong one” (ibid.). Rather than highlighting the positive association suggested by the regression line, he argues that “what is perhaps most striking is that the personal happiness ratings for 10 of the 14 countries lie virtually within a half a point of the midpoint rating of 5 [on the raw 0-10 scale].... The closeness of the happiness ratings implies also that a similar lack of association would be found between happiness and other economic magnitudes” (ibid., p. 106). The clustering of countries within the shaded area on the chart gives a sense of this argument. However, the ordered probit index is quite useful here in quantifying the differences in average levels of happiness across countries relative to the within-country variation. Unlike the raw data, the ordered probit suggests quite large differences in well-being relative to the cross-sectional standard deviation. Similarly, the use of log income rather than absolute income highlights the linear-log relationship. Finally, Easterlin mentions briefly the 1946 and 1949 data shown in the top row of Figure 1, noting that “the results are similar... if there is a positive association among countries between income and happiness it is not very clear” (ibid., p. 108).

Although the correlation between income and happiness in these early surveys is not especially convincing, this does not imply that income has only a minor influence on happiness, but rather that other factors (possibly including measurement error) also affect the national happiness aggregates. Even so, three of these five datasets suggest a statistically significant relationship between happiness and the natural logarithm of GDP per capita. More important, the point estimates reveal a positive relationship between well-being and income, and a precision-weighted average of these five regression coefficients is 0.45, which is comparable to the sort of well-being-GDP gradient suggested in cross-sectional comparisons of rich and poor people within a society (a theme we explore further below).

¹⁰ We plot the ordered probit index, whereas Easterlin graphs the mean response..

¹¹ Following Cantril (1965), Easterlin also notes that “the values for Cuba and the Dominican Republic reflect unusual political circumstances—the immediate aftermath of a successful revolution in Cuba and prolonged political turmoil in the Dominican Republic.”

We have also located several other surveys from the mid-1960s through the 1970s that show a similar pattern. In particular, the ten-nation “Images of the World in the Year 2000” study, conducted in 1967, and the twelve-nation Gallup-Kettering Survey, from 1975, both yield further evidence consistent with an important and positive well-being-GDP gradient. Subsequent cross-country data collections have become increasingly ambitious, and analysis of these data has made the case for a linear-log relationship between subjective well-being and GDP per capita even stronger, while also largely confirming that the magnitudes suggested by these early studies were quite accurate.

Figure 2 presents data on life satisfaction from each wave of the World Values Survey separately; illustrating the accumulation of new data through time (We turn to the data on happiness from this survey below, in Figure 5).¹² In the early waves of the survey the sample consisted mostly of wealthy countries; given the limited variation in income, these samples yielded suggestive, but not definitive, evidence of a link between GDP and life satisfaction. As the sample expanded, the relationship became clearer. In each wave the regression line is upward sloping, and the estimated coefficient is statistically significant and similar across the four waves, with its precision increasing in the later waves. We also plot estimates from locally weighted (or lowess) regressions, to get a sense of whether there are important deviations from the linear-log functional form.¹³ In the earliest waves the small number of countries and limited heterogeneity in income across countries made it difficult to make robust inferences about the relationship between life satisfaction and economic development. Nonetheless, pooling data from all four waves and allowing wave fixed effects yields an estimate of the satisfaction-income gradient of 0.40 (se=0.04, clustering by country), and an F-test reveals that wave-specific slopes are jointly statistically insignificant relative to a model with a common slope term ($F_{3,78}=1.98$).

In some cases the expansion of the World Values Survey to include poorer countries resulted in explicitly unrepresentative samples.¹⁴ For example, Argentina was included in the 1981-84 wave, but the sample was limited to urban areas and was not expanded to become representative of the country overall until the 1999-2004 wave. Chile, China, India, Mexico, and Nigeria were added in the 1989-93 wave, but their samples largely consisted of the more educated members of society and those living in urban areas. These limitations are spelled out clearly in the survey documentation, but have been ignored in most subsequent analyses. The non-representative samples typically came from poorer countries, and involved sampling richer (and hence likely happier) respondents. As such, inclusion of these observations imparts a downward bias on estimates of the well-being-income gradient. We therefore exclude from our analysis countries which the survey documentation suggests are clearly not representative of the entire population. Observations for these countries are plotted in Figure 2 using hollow squares. As expected, these observations typically sit above the regression line. Appendix B provides a comparison of our results when these countries are included in the analysis, along with greater detail regarding sampling in the World Values Survey.

Subsequently, the 2002 Pew Global Attitudes Survey interviewed 38,000 respondents in forty-four countries across the development spectrum. The subjective well-being question is a form of Cantril’s (1965) “Self-Anchoring

¹² In order to make these data collections consistent, we analyze only adult respondents.

¹³ The lowess estimator is a local regression estimator that plots a flexible curve.

¹⁴ We thank Angus Deaton for alerting us to these limitations in the World Values Survey.

Striving Scale”. Respondents were shown a picture and told, “Here is a ladder representing the ‘ladder of life.’ Let's suppose the top of the ladder represents the best possible life for you; and the bottom, the worse possible life for you. On which step of the ladder do you feel you personally stand at the present time?” Respondents were asked to choose a step along a range of 0 to 10. As before, we run an ordered probit of the ladder ranking on country fixed effects to estimate average levels of subjective well-being in each country, and we compare these averages with the log of GDP per capita in Figure 3. These data show a linear relationship similar to that seen in Figure 2.

The most ambitious cross-country surveys of subjective well-being come from the 2006 Gallup World Poll. This is a new survey designed to measure subjective well-being consistently across 132 countries. Similar questions were asked in all countries, and the survey contains data for each country that are nationally representative of people aged 15 and older. The survey asks a variety of subjective well-being questions, including a ladder question similar to that used in the 2002 Pew survey. As Figure 4 shows, these data yield a particularly close relationship between subjective well-being and the log of GDP per capita. Across the 131 countries for which we have usable GDP data (we omit Palestine), the correlation exceeds 0.8. Moreover, the estimated coefficient on log GDP per capita, 0.42, is similar to those obtained using the World Values Survey, the Pew survey and from the earlier surveys, including those assessed by Easterlin. These findings are also quite similar to those found by Deaton (2008), who also shows a linear-log relationship between subjective well-being and GDP per capita using the Gallup World Poll.¹⁵ Deaton emphasizes that the clearer relationship in the Gallup data reflects the inclusion of surveys from a greater number of poor countries.

As discussed previously, the economics literature has tended to treat measures of happiness and life satisfaction as largely interchangeable, whereas the psychology literature distinguishes between the two. We now turn to assessing the relationship between measures of happiness and income and comparing these estimates with the estimates of the relationship between measures of life satisfaction and income considered thus far. (We consider additional measures of subjective well-being and their relationship to income in a later section.) Figure 5 investigates both the happiness-GDP link and the life satisfaction-GDP link estimated using the latest wave of the World Values Survey—the life satisfaction data are those discussed above. Happiness is measured using the following question: “Taking all things together, would you say you are: ‘very happy,’ ‘quite happy,’ ‘not very happy,’ ‘not at all happy?’” The results suggest that these measures may not be as synonymous as previously thought: happiness appears to be somewhat less strongly correlated with GDP than is life satisfaction.¹⁶ Although much of the sample shows a clear relationship between log income and happiness, these data yield several particularly puzzling outliers. For example, the two poorest countries in the sample,

¹⁵ We estimate a well-being-income gradient that is about half that estimated by Deaton because we have standardized our estimates through the use of ordered probits, whereas Deaton is estimating the relationship between the raw life satisfaction score and log income. Putting both on a similar scale yields similar estimates. Appendix A compares our ordered probit approach with other possible cardinalizations of subjective well-being.

¹⁶ The contrast in Figure 5 probably overstates this divergence, as it plots the data for the 1999-2004 wave of the World Values Survey, whereas Table 1 shows that earlier waves yielded a clearer happiness-GDP link.

Tanzania and Nigeria, have the two highest levels of average happiness, yet both have much lower average life satisfaction—indeed, Tanzania reported the lowest average satisfaction of any country.¹⁷

This apparent noise in the happiness-GDP link partly explains why earlier analyses of subjective well-being data have yielded mixed results. We reran both the happiness and life satisfaction regressions with Tanzania and Nigeria removed, and it turns out that these outliers explain at least part of the puzzle. In the absence of these two countries, the well-being-GDP gradients, measured using either life satisfaction or happiness, turn out to be very similar. Equally, in these data the correlation between happiness and GDP per capita remains lower than that between satisfaction and GDP per capita.

To better understand whether the happiness-GDP gradient systematically differs from the satisfaction-GDP gradient, we searched for other data collections that asked respondents about both happiness and life satisfaction. Figure 6 brings together two such surveys: the 1975 Gallup-Kettering survey and the First European Quality of Life Survey, conducted in 2003. In addition, the bottom panels of Figure 6 show data from the 2006 Eurobarometer, which asked about happiness in its survey 66.3 and life satisfaction in survey 66.1. In each case the happiness-GDP link appears to be roughly similar to the life satisfaction-GDP link, although perhaps, as with the World Values Survey, slightly weaker.

Table 1 formalizes all of the analysis discussed thus far with a series of regressions of subjective well-being on log GDP per capita, using data from the Gallup World Poll, all four waves of the World Values Survey, and the Pew Global Attitudes Survey. The coefficient on log GDP per capita is reported along with its standard error. The first column reports coefficient estimates from ordered probit regressions of individual well-being on the natural log of real GDP per capita, with robust standard errors clustered by country; the second column adds controls for gender and a quartic in age and its interaction with gender. The third column reports the results of a two-stage process: in the first stage we aggregated the data to the country level by running an ordered probit regression of subjective well-being on country fixed effects, which we interpret as a measure of average national happiness. In the second stage we estimated an ordinary least squares regression of these country fixed effects on log GDP per capita. The coefficient from this second regression is reported in the third column of Table 1. In all the data sets examined, estimates of the relationship obtained from the respondent-level analysis are similar to that obtained through the two-stage process. Moreover, each of these datasets yields remarkably similar estimates of the subjective well-being-GDP gradient, typically centered around 0.4.

The regressions reported in the first three columns of Table 1 are performed on the complete sample of countries for each survey; the samples in the remaining two columns consist only of countries with GDP per capita above or below \$15,000 (in 2000 dollars), using the same two-stage process as in the third column, to allow us to assess whether the well-

¹⁷ One might suspect that survey problems are to blame, and indeed, the survey notes for Tanzania suggest (somewhat opaquely) that “There were some questions that caused problems when the question was translated, especially questions related to... Happiness because there are different perceptions about it.” We are not aware of any other happiness data for Tanzania, but note that in the 2002 Pew survey Tanzania registered the second-lowest level of average satisfaction among forty-four countries (Figure 3). The high levels of happiness recorded in Nigeria seem more persistent: Nigeria also reported the eleventh-highest happiness rating in the 1994-99 wave of the World Values Survey, although it was around the mean in the 1989-93 wave.

being-GDP gradient differs for rich and poor countries. It has been argued that income is particularly important for happiness when the basic needs of food, clothing, and shelter are not being met, but that beyond this threshold happiness is less strongly related to income. In its stronger form, this view posits a satiation point beyond which more income no longer raises the happiness of a society. For instance, Layard claims that “if we compare countries, there is no evidence that richer countries are happier than poorer ones—so long as we confine ourselves to countries with incomes over \$15,000 per head.... At income levels below \$15,000 per head things are different....” (2005b, p. 149) Frey and Stutzer offer a similar assessment of the literature, suggesting that “income provides happiness at low levels of development but once a threshold (around \$10,000) is reached, the average income level in a country has little effect on average subjective well-being” (2002, p. 416).

Employing Layard’s cutoff, we find that the relationship between subjective well-being and log GDP per capita is, if anything, stronger, rather than weaker, in the wealthier countries, although this difference is statistically significant only in a few cases. The point estimates are, on average, about three as large for those countries with incomes above \$15,000 compared to those countries with incomes below \$15,000.¹⁸ We thus find no evidence of a satiation point. Indeed, a consistent theme across the multiple datasets shown in Table 1 and Figures 1-6 appears to be that there is a clear positive relationship between subjective well-being and GDP per capita, even when the comparison is among developed economies only.

The fact that the coefficient on log GDP per capita may be larger for rich countries should be interpreted carefully. Taken at face value, the Gallup results suggest that a 1 percent rise in GDP per capita would have about three times as large an effect on measured well-being in rich as in poor nations.¹⁹ Of course, a 1 percent rise in U.S. GDP per capita is about ten times as large as a 1 percent rise in Jamaican GDP per capita. Consider instead, then, the effect of a \$100 rise in average incomes in Jamaica and the United States. Such a shock would raise log GDP per capita by ten times more in Jamaica than in the United States, and hence would raise measured well-being by about three times as much in Jamaica as in the United States. For the very poorest countries, this difference is starker. For instance, GDP per capita in Burundi is about one-sixtieth that in the United States, and hence a \$100 rise in average income would have a twenty-fold larger impact on measured well-being in Burundi than in the United States.²⁰

One explanation for the difference between our findings and earlier findings of a satiation point may be differences in the assumed functional form of the relationship between well-being and GDP. In particular, whereas we have analyzed well-being as a function of log GDP per capita, several previous analyses have focused on the absolute

¹⁸ This finding is consistent with Deaton (2008).

¹⁹ Higher income yields a larger rise in the happiness index, but not necessarily a larger rise in happiness, since we do not know the “reporting function” that translates true hedonic experience into our measured well-being index (Oswald, forthcoming).

²⁰ Using the Gallup World Poll data, we can check whether the log GDP-well-being gradient differs for the very poorest countries. When restricting the sample to countries with GDP per capita below \$3,000, we obtain estimates very similar to those for countries with GDP per capita between \$3,000 and \$15,000. This is also evident in the nonparametric fit shown in Figure 4.

level.²¹ Figure 7 shows that the log specification yields a better fit, although the difference is small (Deaton, 2008, p. 58). Viewed either way, there remains robust evidence of a strongly positive well-being-income link for rich countries. We have re-estimated the well-being-GDP relationship using levels of GDP per capita as the independent variable and found that the well-being-GDP gradient is about twice as steep for poor countries as for rich countries. That is, consistent with our earlier findings, a rise in income of \$100 is associated with a rise in well-being for poor countries that is about twice as large as for rich countries. (A 1 percent rise in GDP per capita is associated with much larger income gains, and hence much larger well-being gains for rich countries.)

Thus, our conclusion that there is strong evidence against a satiation point is robust to whether one conceives of well-being as rising with log GDP per capita or with its absolute level. As Figure 7 demonstrates, even with observations on 131 countries, we have insufficient data to draw particularly strong inferences about the appropriate functional form, although the evidence is certainly suggestive of a linear-log well-being-income relationship. In the next section we turn to within-country comparisons, and given the much larger samples involved, it will be clear that—at least at the individual level—well-being is best thought of as rising in log income. It is this finding that guides our choice of the appropriate functional form for between-country comparisons.

IV. Income and Happiness: Comparing Within-Country and Between-Country Estimates

A very simple benchmark for assessing the magnitude of the between-country well-being-GDP gradient measured in the previous section (typically centered around 0.4) is the within-country well-being-income gradient. In particular, Easterlin argued that “the happiness differences between rich and poor countries that one might expect on the basis of the within country differences by economic status are not borne out by the international data” (1974, pp. 106-107). Thus, we now turn to comparing the happiness of richer and poorer members of the same society at a single point in time.

On this question there is a clear consensus in the literature, aptly summarized by Easterlin: “As far as I am aware, in every representative national survey ever done a significant bivariate relationship between happiness and income has been found” (2005, p. 67). And indeed, we have made similar comparisons in over 100 countries and have yet to find a (statistically significant) exception. Although there has been some debate about the magnitude of this effect, income is clearly an important correlate with happiness. For example, Frank argues for the importance of income for happiness as follows: “When we plot average happiness versus average income for clusters of people in a given country at a given time...rich people are in fact a lot happier than poor people. It’s actually an astonishingly large difference. There’s no one single change you can imagine that would make your life improve on the happiness scale as much as to move from the bottom 5 percent on the income scale to the top 5 percent” (2005, p. 67).

²¹ In a levels specification, the subjective well-being-income gradient is curvilinear and thus is less steep among wealthier countries. While the slope is never zero, the flattening out of the curve may be more easily misinterpreted as satiation.

In this spirit we examine the relationship between happiness and income in the United States from 1972 through 2006 using the General Social Survey (GSS), produced by the National Opinion Research Center. Figure 8 plots the coefficients from an ordered probit regression of happiness on income category by year fixed effects against family income. The x-axis, family income, is converted from income categories to income by fitting interval regressions to the income data on the assumption that income follows a log-normal distribution.²² Each circle in the Figure represents an income category in a particular year; the diameter of each circle is proportional to the population in the income category in that year. The statistical significance of this relationship is not in doubt, largely because each round of the GSS (as with most happiness surveys) involves over 1,000 respondents. This plot also leaves very little doubt about the functional form: the linear-log relationship between our happiness index and family income is clearly evident throughout the income distribution.²³

To further investigate the functional form relationship, we investigated the relationship in other datasets for other countries and find similar evidence pointing to a linear-log relationship between subjective well-being and income. In Figure 9, we use the Gallup World Poll (as it covers the most countries of any of our data sets) and show estimates from a regression of life satisfaction on separate income category fixed effects for each country, controlling for country fixed effects. We have usable household income data for 113 countries.²⁴ The coefficient estimates on the individual country-household income category fixed effects are plotted against the log of household income, normalized by subtracting off the country average. This figure also points strongly to a linear relationship between subjective well-being and the log of family income, with no evidence of satiation.

It is the juxtaposition of these statistically significant cross-sectional findings with statistically insignificant cross-country and time-series results that gave rise to the Easterlin paradox. Theories emphasizing relative income comparisons would suggest that the between-country well-being-income gradient would be smaller than the within-country well-being-income gradient (if relative income comparisons are made intranationally). Yet, the suggestive comparison of the gradients in Figure 4 with Figure 9 points to the opposite conclusion, with the gradient estimated between countries larger than that seen within the countries.

While Figure 9 plots the gradient seen when examining all of the countries together, it is worth estimating the within-country well-being-income gradient separately for individual countries to see the range of estimated within-country gradients. Thus, for each country we estimate an ordered probit regression of life satisfaction on the natural log of household income, controlling for gender and a quartic in age, entered separately for men and women. The coefficient estimates obtained in each regression (rounded to the nearest 0.05) are displayed in Figure 10 as a histogram summarizing

²² We thank Angus Deaton for this suggestion.

²³ Because the GSS retained the nominal income categories used in 1973, some very low income cells are somewhat off the regression line (the circles to the far left of the graph), reflecting both the fact that small cells yield imprecise happiness estimates and the difficulties in imputing appropriate incomes to the bottom-coded group.

²⁴ We drop Kenya because it lacks labels for income groups, Laos because it contains clearly implausible income groupings, and Uzbekistan because the income categories listed in the data involve overlapping ranges. Respondent-level income data are unavailable for Egypt, Iran, Iraq, Jordan, Kuwait, Latvia, Lebanon, Morocco, Pakistan, Palestine, the Philippines, Saudi Arabia, Sri Lanka, Turkey, the United Arab Emirates, and Yemen. This leaves us with valid household income data for 113 countries.

the entire sample.²⁵ Overall, the average well-being-income gradient is 0.38, with the majority of the estimates between .25 and .45 and 90 percent are between 0.07 and 0.72. In turn, much of the heterogeneity likely reflects simple sampling variation: the average country-specific standard error is 0.07, and 90 percent of the country-specific regressions have standard errors between 0.04 and 0.11.

As an alternative representation of these data, Figure 11 directly compares within-country and between-country estimates of the well-being-income gradient. Each solid circle plots the GDP per capita and average well-being of a single country (hence the circles suggests the between-country well-being-GDP gradient), and the slope of the arrows, fitted to each circle, represents the slope of the well-being-income gradient estimated within that country. Not only are the slopes of the arrows remarkably similar across countries; they are also typically quite close to the between-country well-being-GDP slope (the thick dashed line). Figure 12 repeats this exercise using data from the 1999-2004 wave of the World Values Survey. The household income data in that survey are not as uniform as those in the Gallup World Poll, requiring us to omit several countries.²⁶ However, for the countries with sufficient data, the pattern that emerges is similar to that seen in the Gallup data. Repeating the same exercise for the Pew data also yields similar findings (not shown).

Table 2 pools the various national surveys so as to arrive at a summary estimate of the within-country well-being-income gradient. Thus, for each international dataset, we perform an ordered probit of subjective well-being on log household income, controlling for country (or, for the World Values Survey, country by wave) fixed effects, which serve to control for not only the between-country variation in GDP per capita, but also variation in measured income due to differences in exchange rates, purchasing power or other country-specific factors. The first column shows the results from a simple ordered probit of well-being on log household income, controlling for these fixed effects; the second column adds controls for gender, a quartic in age, and the interaction of these variables. Comparison of these results with the corresponding between-country estimates in Table 1 shows them to be roughly similar in magnitude, although as seen in the figures, in most cases the between-country estimates are larger than the within-country estimates, which are centered around 0.3.

An important issue in considering the within-country cross-sectional relationship between income and subjective well-being is the extent to which measured income differences at a point in time reflect differences in permanent income versus transitory shocks. If people are able to smooth their consumption, then subjective well-being should change little with transitory income changes, and permanent shocks should have a much larger impact. The variation in GDP per capita between countries is likely dominated by variation in permanent income, whereas the variation in annual income within a population likely reflects both permanent and transitory shocks.

²⁵ As with the GSS, our various data sources typically report income in categories, rather than as a continuous variable. We follow the same method for each of our datasets, fitting interval regressions to these income data on the assumption that income follows a log-normal distribution. If a dataset contains a bottom income category of zero, we combine it with the succeeding income category. We perform these regressions separately for each country-wave of each dataset.

²⁶ In many cases, particularly in earlier waves of the World Values Survey, household income is reported only as an ordinal variable with no information regarding the underlying cardinal measure.

A simple back-of-the-envelope calculation can help determine an upper bound on the extent to which these issues are distorting the comparisons in tables 1 and 2. If all cross-country variation in GDP per capita is assumed permanent and people are perfect permanent-income consumers, then the coefficients in Table 1 can be interpreted as the response of well-being to a shock to consumption. Standard estimates for the United States suggest that around half the variation in annual income in a national cross-section is transitory, and a \$1 shock to transitory income typically translates into around a \$0.05 shock to permanent income. Thus, a \$1 change in measured income translates to a roughly \$0.525 change in permanent income. In this case the estimates in Table 2 need to be adjusted upward by around 90 percent ($1/0.525$) to be interpreted as the relationship between well-being and permanent income or consumption. If, instead of assuming perfect smoothing, we accept Campbell and Mankiw's (1990) estimate that 50 percent of income is earned by "rule-of-thumb" consumers whose propensity to consume from current income is equal to their propensity to consume from permanent income, the relevant adjustment is closer to 30 percent. This adjustment would make the within- and between-country estimates roughly similar.

We can also address this issue empirically. In an effort to isolate the response of well-being to permanent income, the regression reported in the last column of Table 2 instruments for income using educational attainment, entered separately for each country (Rivers and Vuong, 1988).²⁷ Although we are confident that these instruments isolate variation in permanent rather than transitory income, we do not hold much faith that the exclusion restriction holds—that education does not have an effect on well-being beyond that mediated by income.²⁸ Given that these omitted effects are likely positive, our instrumental variables estimates may overstate the within-country income-well-being gradient. Indeed, in most cases the instrumental variables estimates are larger than the ordered probit estimates of well-being on income. In the largest dataset—the Gallup World Poll—the estimated gradient is .6.

The discussion above has been premised on the straightforward view that transitory income shocks yield smaller impacts on well-being than do permanent shocks. Yet the most direct evidence we have on this point—the movement of well-being over the business cycle—in fact suggests the opposite. Figure 13 shows that business-cycle variation in the output gap produces quite large affects on subjective well-being. Indeed, the estimated well-being-transitory income gradient suggested by these shocks is about five times larger than the well-being-GDP gradient estimated in Table 1. If this sort of variation is representative of the response of happiness to transitory income, then, paradoxically enough, our findings in Table 2 may substantially overstate the within-country well-being-permanent income link.

Although our analysis provides a useful measurement of the bivariate relationship between income and well-being both within and between countries, there are good reasons to doubt that this corresponds to the *causal* effect of income on well-being. It seems plausible (perhaps even likely) that the within-country well-being-income gradient may be biased

²⁷ We follow their approach to estimating an instrumental variables ordered probit. Thus, the first stage involves a regression of log household income on indicator variables for each level of education in each country, controlling for country fixed effects as well as gender and a quartic in age, entered separately for each gender and indicators for missing age or gender. The second stage involves an ordered probit regression of well-being on the predicted values and residuals from the second stage, as well as the same controls, including country fixed effects.

²⁸ For instance, Lleras-Muney (2005) shows positive impacts of compulsory schooling on health.

upward by reverse causation, as happiness may well be a productive trait in some occupations, raising income. A different perspective, from offered by Kahneman, et al. (2006), suggests that within-country comparisons overstate the true relationship between subjective well-being and income because of a “focusing illusion”: the very nature of asking about life satisfaction leads people to assess their life relative to others, and they thus focus on where they fall relative to others in regard to concrete measures such as income. Although these specific biases may have a more important impact on within-country comparisons, it seems likely that the bivariate well-being-GDP relationship may also reflect the influence of third factors, such as democracy, the quality of national laws or government, health, or even favorable weather conditions, and many of these factors raise both GDP per capita and well-being (Kenny, 1999).²⁹ Other factors, such as increased savings, reduced leisure, or even increasingly materialist values may raise GDP per capita at the expense of subjective well-being. At this stage we cannot address these shortcomings in any detail, although, given our reassessment of the stylized facts, we would suggest an urgent need for research identifying these causal parameters.

V. Economic Growth and Happiness

The last two sections have shown that wealthier societies have greater subjective well-being than poorer societies and that, to a similar degree, wealthier members of a society are happier than their poorer counterparts. This then leads to our final question: do societies get happier through time as they become richer? Easterlin argues that the possibly confounding “cultural influences on international happiness comparisons underscore the importance of national time series evidence... for inferring the relationship between subjective well-being and economic development” (1995, pp. 43-44). Indeed, the core of the Easterlin paradox lies in Easterlin’s failure to isolate statistically significant relationships between average levels of happiness and economic growth through time. Easterlin’s 1974 and 1995 papers contain three important datasets, tracking the time series of happiness within Europe, Japan, and the United States.

Our analysis is based on three observations about the inferences that existing datasets can support. First, absence of evidence should not be confused with evidence of absence. This is particularly important given both the variability of happiness aggregates between surveys and the limited range of variation in time series rather than cross-national comparisons of GDP per capita. Second, when we reanalyze these data, we find that happiness has in fact risen in Japan and Europe. The failure of happiness to rise in the United States remains a puzzling outlier, although the extent to which it constitutes a sharp exception should not be overstated. Third, as more data have become available, in the form of both extended national time series and observations from new countries, evidence that happiness rises with GDP per capita has started to accumulate.

Indeed, the World Values Survey has been running since 1981, and across its four waves we now have repeated observations on a large number of countries, spread across several decades. Figure 14 shows the movement of both life satisfaction and real GDP per capita across the waves for all countries for which this survey offers repeated observations.

²⁹ Kenny argues directly for reverse causation running from happiness to income.

As before, we estimate average well-being in a country-wave as the coefficient from an ordered probit regression of well-being on a saturated set of country by wave fixed effects. Arrows link each individual country's change in well-being–GDP space over time, and so the slope of each arrow corresponds to the well-being-income gradient derived from two consecutive observations in a country's national time series (dotted arrows connect points where the sampling frame changed, and hence valid time series comparisons cannot be made).

Several points are evident from Figure 14. First, there appears to be a general tendency for economic growth to be accompanied by growth in subjective well-being (arrows tend to point northeast), and economic decline, which is most visible in the former Eastern bloc, has been accompanied by a decline in well-being (arrows pointing southwest). Of the 89 changes shown in the left-hand panel of Figure 14, happiness and GDP per capita change in the same direction in 62 cases (53 show growth in both; 9 show declines), whereas they move in opposite directions in 27 (of which 20 reflect economic growth unaccompanied by growth in happiness, and 7 reflect growing happiness despite economic decline). The life satisfaction data in the right-hand panel yield much weaker results, with satisfaction and GDP per capita moving in the same direction in only 46 of 90 cases, reflecting generally weaker measured life satisfaction in the two most recent waves of the survey (more on this below).

Second, when we average across these country-specific estimates, the well-being-income link within countries through time appears to be roughly similar to that estimated from the pooled cross-country, cross-time variation (shown as the dashed line in each panel). Third, substantial heterogeneity remains in these estimated responses, although this may reflect the influence of other factors on measured well-being.

Finally, these time-series changes are strongly influenced by the result of common patterns across countries observed in specific waves. We suspect that the trend in life satisfaction has been distorted by changes in question ordering. In particular, in the 1994-99 and 1999-2004 waves, the life satisfaction question was preceded by a question asking "How satisfied are you with the financial situation of your household?" Respondents typically rate their financial satisfaction substantially lower than their life satisfaction (on the same 1-10 scale, responses average about one point lower), and hence this question may have influenced how respondents subsequently reported their life satisfaction. To check this, we assess the (raw) correlation between life satisfaction and financial satisfaction for the eight countries with representative samples in each round of the World Values Survey; this correlation was 0.53 and 0.57 in the two most recent waves, significantly above previous levels (0.45 in the first wave and 0.43 in the second). The happiness question was never proximate to the financial satisfaction question, and the correlation of happiness with financial satisfaction was quite stable across each of the waves (it was recorded as 0.29, 0.30, 0.32 and 0.29 from the earliest to the latest wave). Similarly, in the 1994-99 and 1999-2004 waves, the happiness question was by a battery of questions probing the importance of friends, family, leisure, politics and religion, and a similar analysis reveals that the correlation of measured happiness with these variables rose. If these questions prime positive thoughts, this question-order change may have inflated measured happiness in the past two waves.

These question-order changes make direct comparisons of country's well-being levels across successive waves problematic. However, to the extent that these influences are common across countries, first differencing will yield useful estimates. Thus, in Figure 15, we analyze changes in life satisfaction and log GDP between each wave of the panel. The first row shows differences between adjacent waves, while the second row shows longer differences, with the last panel showing differences between the first and last wave. (Because of the uneven participation through time of many countries, these longer difference panels contain information not shown in the first row). In each comparison of pairs of waves, we find that larger rises in per capita GDP are associated with larger rises in life satisfaction, and the magnitude of these gradients tends to be centered around 0.4. A parallel analysis of the happiness data (not shown) yielded roughly similar results (the slope was positive in five of six panels and statistically significant in only one case).

Panel regressions provide an alternative and statistically more efficient way to combine this information, and so in Table 3 we turn to analyzing both life satisfaction and happiness measures in the World Values Survey as a country-wave panel dataset. The first column reports the results of respondent-level ordered probit regressions of well-being on log GDP per capita, while the second column aggregates the data to the country-wave level; these are ordinary least squares regressions of our well-being index on log GDP per capita. The first row reports results for the simple bivariate well-being-GDP relationship and hence pools both within-country and between-country variation. The estimated coefficients are 0.4 for life satisfaction and 0.2 for happiness. To isolate the within-country time-series variation, the second row adds controls for country fixed effects. Consistent with Figure 14, the well-being-GDP gradient estimated from this time-series variation is similar to that estimated from the point-in-time between-country comparisons.

The next row adds controls for each wave of the World Values Survey, these controls partial out the changes in well-being that reflect differences in surveys across waves. As might be expected, in light of the previous discussion of question order effects, the inclusion of these controls increases the estimate of the time-series life satisfaction-GDP gradient to nearly 0.6, and lowers the estimate of the time-series happiness-GDP gradient to a bit more than 0.2.

In subsequent rows we take first differences of consecutive country-wave observations, as well as long differences (subtracting the first from the last observation for each country). Consistent with the analysis in Figure 15, there is a clear, and statistically significant, relationship between changes in life satisfaction and log GDP per capita over time in these countries. The estimates for happiness are similar, albeit smaller and less precisely estimated. These repeated international cross-sections yield estimates of the time series well-being-income gradient centered roughly around 0.4, with larger estimates for life satisfaction than happiness.

Equally, it is worth emphasizing that these estimates are both somewhat imprecisely estimated, and fragile. Although the large cross-country datasets allow for useful comparisons between populations in abject poverty with those in industrialized powerhouses, the within-country time series variation is simply less impressive. Indeed, it is worth noting that the standard deviation of log GDP per capita across countries (in the 1999-2004 wave) is 1.0, whereas the standard deviation of between-wave first differences in log GDP per capita (across all waves) is only 0.2, and hence strong inferences are difficult to draw. Moreover, the inferences one draws from these data are particularly sensitive to the quite

unusual economic trajectories of a small number of countries, such as the rapid growth in Korea and Ireland, and the decline of the former Eastern bloc countries (Figure 15). Even so, most of our approaches to these data yield suggestive evidence falsifying the Easterlin hypothesis that the time-series well-being-GDP gradient is zero. Moreover, even in those cases in which the data fail to falsify the null that the gradient is zero, they also fail to falsify the alternative null hypothesis that this gradient is equal to 0.4, similar to that obtained from our between-country or within-country analyses.

How should our findings be reconciled with earlier reports suggesting no link between changes in GDP per capita over time and life satisfaction? We suspect that the key is simply that our analysis of the satisfaction-income gradient based on both within- and between-country comparisons gives us a specific quantitative yardstick for assessing the importance of (even imprecisely estimated) trends in subjective well-being.

Europe

We turn next to the other major set of repeated international cross-sectional data, the Eurobarometer Survey, a data collection intended to track public opinion across the European Union. We draw our data from the Mannheim Eurobarometer Trendfile, which collects available microdata from 1970 to 2002, supplemented with data extracted from print editions of the Eurobarometer Reports series from 2002 through 2007. These surveys initially asked respondents in what were then the nine member states of European Union about their life satisfaction. A life satisfaction question has been asked at least annually (and often semi-annually) from 1973 onward (except in 1974 and 1996). The survey has expanded as the E.U. expanded, covering fifteen countries by 2002 (with separate surveys for East and West Germany), and it presently includes thirty countries (including three candidate countries), yielding a broad but unbalanced panel. A happiness question was also briefly asked (from 1975 through 1986, except in 1980 and 1981, and in a different format in 2006); given these gaps in the data, we focus on life satisfaction. For the purposes of our analysis, we keep West Germany separate from East Germany, which permits us to analyze a continuous sample of well-being among West Germans over thirty-five years.

We begin by analyzing the relationship between life satisfaction and GDP for the nine countries that constituted the original 1973 sample. Easterlin analyzed these same nine countries (through to 1989), concluding that “satisfaction drifts upward in some countries, downward in others. The overall pattern, however, is clearly one of little or no trend in a period when real GDP per capita rises in all of these countries from 25 to 50 percent” (1995, p. 38). In a subsequent update, Easterlin maintains that “I think the evidence continues to support my generalization in the 1995 study” (2005a, p. 434).

Figure 16 updates this analysis, adding a further eighteen years of data (shown with hollow circles). In eight of the nine countries, rising GDP per capita has been associated with rising life satisfaction, six of which are statistically significant ($p < .10$, assessed using Newey-West standard errors, accounting for first-order autocorrelation). This figure also suggests a couple of puzzles: a significant declining trend in satisfaction is observed in Belgium, and declining life satisfaction in Ireland during the 1970s and 1980s, although this was quickly followed by rising satisfaction during the rapid economic growth associated with the “Irish Miracle”. (Satisfaction appeared to be anomalously high in the very first

Irish survey; dropping this observation yields a statistically significant coefficient on log GDP per capita of 0.14, with a standard error of 0.05 for the entire sample period.) Our point is not to count up the number of statistically significant responses one way or the other, but rather to suggest that across the nine large European countries for which we have a long time series, life satisfaction has typically risen with GDP per capita. Moreover, estimates of the satisfaction-GDP gradient based on these national time series, although quite variable, average around 0.25, with some estimates larger and some smaller.

The upward trend in life satisfaction across the European Union is not widely understood, and in Figure 17, we provide some intuition for why this has not been obvious. The simplest approach to building an E.U.-wide time series of life satisfaction involves taking a population-weighted average of the satisfaction levels of whichever countries happen to be member states at any point in time. Through this period the E.U. has systematically expanded to incorporate poorer countries, which have lower average life satisfaction. This expansion has thus pushed downward measured average life satisfaction, despite the fact that satisfaction rose within most countries. This can most easily be seen by simply examining the nine countries (the EU-9) which have been in the E.U., and hence the Eurobarometer, since 1973. This analysis, shown with short-dashes, takes a population-weighted average of the satisfaction indices of the EU-9 countries and shows rising life satisfaction through time.

In order to use the data from all countries without having the resulting time series driven by compositional changes, we also construct a regression-adjusted series by running an OLS regression of national satisfaction indices on time (survey round) fixed effects, weighting by population, and controlling for country fixed effects (thereby adjusting for different average well-being levels among new entrants to the E.U.). These time fixed effects are also plotted in Figure 17 and clearly suggest a rising trend in life satisfaction similar to that seen in the EU-9 average. Finally, we create a spliced series by summing through time first differences in the broadest available fixed-weight average of satisfaction in E.U. member nations.³⁰ This series is quite similar to the regression-adjusted measure. Clearly the simple average disguises much of the rise in satisfaction occurring within E.U. member nations.

Even accounting for these compositional changes it would be difficult to infer that a positive trend either did or did not exist on the basis of only Easterlin's 1973-89 sample. But over the entire 1973-2007 period, the magnitude of the trend rise in satisfaction turns out to be quite close to what might be expected given underlying GDP trends. Fitting a simple time trend to the composition-adjusted aggregates shown in Figure 17 suggests that life satisfaction in Europe rose at a (statistically significant) average rate of about 0.006 per year, compared with a trend rise in log GDP per capita of around 0.020 per year. Considered jointly, these trends point to a long-run satisfaction-GDP gradient of about 0.3

³⁰ The spliced series analyzes the EU-9 (from the beginning of the survey in summer 1973), expanding to the EU-10 (adding Greece, from 1981 onward), then the EU-12 (adding Portugal and Spain from 1986 onward), the EU-12+ (German reunification, adding East German surveys from the Fall 1990 survey onward), the EU-16 (adding Austria, Finland and Sweden from 1995 onward), the EU-26 (adding Cyprus, the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Slovakia and Slovenia from the Fall 2004 survey onward), and the EU-27 (adding Bulgaria and Romania from 2007). Population weights for each index reflect the average population share of each country in that aggregate.

(=0.006/0.020), which both falsifies the null hypothesis of no positive relationship and is roughly consistent with the magnitudes seen in our within- and between-country assessments.

To further examine these patterns, Table 4 formalizes our findings with a series of panel regressions exploiting all of the Eurobarometer Survey observations across all countries; the top panel analyzes life satisfaction and the bottom panel happiness. (Due to the very limited happiness data available, estimates in the bottom panel are extremely imprecise.) As with our panel analysis of the World Values Survey, we begin by including no fixed effects and subsequently add in country fixed effects and then both country and year fixed effects. The latter two estimates focus on the time series relationship between satisfaction and GDP and yield coefficients of about 0.2.

The last two rows of each panel attempt to minimize the potential influence of high-frequency variation, by averaging well-being and GDP over five-year periods, or over entire decades. These first-difference regressions yield somewhat larger estimates for life satisfaction, although the decadal differences are imprecisely estimated. Because happiness was only included in the early years of the survey there are many fewer observations yielding extremely imprecise estimates. Nonetheless, in all cases the estimated coefficients are positive, in some specifications we can falsify the null hypothesis that the well-being-GDP gradient is zero, and in no case can we falsify that it is 0.4.

Japan

Arguably the most persuasive evidence in favor of the Easterlin paradox has come from Japan, which provides a striking case study both because of its dramatic growth in the postwar period (real GDP has risen by a factor of six since World War II), and because it was believed that consistent data on subjective well-being had been continuously collected by the government since 1958 in the “Life in Nation” surveys. Previous researchers have analyzed the simple summary of these questions provided by Veenhoven (1993), observing that average levels of well-being had remained flat even in the face of this spectacular growth (Easterlin, 1995, pp. 39-40).³¹

Upon closer inspection, however, these Japanese data are neither as persuasive as many thought, nor is the trend flat. We returned to the original codebooks and had the questions translated.³² This exercise was quite revealing, suggesting several important series breaks. Accounting for these breaks yields a very different perspective. We provide a full accounting in Table 5, which shows both the literal and idiomatic translations of the survey questions as they have changed.

Three important findings emerge from this table. First, in 1964 the response categories changed dramatically. The top category was changed from the catch-all “Although I am not innumerablely satisfied, I am generally satisfied with life now” to the more demanding “Completely satisfied.” Not surprisingly, the proportion reporting their well-being in this

³¹ For instance, he notes that “Between 1958 and 1987 real per capita income in Japan multiplied a staggering five-fold, propelling Japan to a living level equal to about two-thirds that of the United States.... Despite this unprecedented three decade advance in level of living, there was no improvement in mean subjective well-being.” These observations have been cited approvingly by Layard (2005a), Frank (2005), and Kahneman and coauthors (2006), among dozens of others.

³² We thank Michael L. Woodford for his patient assistance with these translations.

highest category declined from 18.3 percent to 4.4 percent. The second category from the top also became more demanding, changing from “Although I can’t say that I am satisfied, if life continues in this way, it will be okay,” to “Although I can’t say I am completely satisfied, I am satisfied.” In parallel, the bottom category changed from “Life now is very unbearable” to “Completely dissatisfied,” but the proportion choosing this lowest category changed little. Second, the questions asked from 1958 to 1969 focused on feelings about “life at home,” whereas the focus of the relevant question from 1970 onward was on global life satisfaction. Third, the survey question—and the allowable responses—changed again in 1992.

Properly viewed, this leaves us with four periods within which we can make useful assessments of trends in subjective well-being in Japan. A cursory inspection of Table 5 suggests an upward trend in well-being in 1958-63, continuing when a new question was asked for the 1964-69 period, followed by a slower rise from 1970 to 1991. This roughly parallels the path of Japanese GDP through these periods. From 1992 until 2007, life satisfaction fell, but this coincides with the end of the Japanese growth miracle and indeed the onset of an economic slump. All told, these findings suggest that subjective well-being in Japan has largely risen with GDP per capita, and that it rose most sharply during the period of rapid growth.

Having established that these data appear qualitatively consistent with a positive satisfaction-GDP gradient, we now turn to a quantitative assessment of the magnitude of this link. One simple approach involves treating these data as four separate datasets and following our earlier style of analysis. Thus, within each continuous subseries we create a time series of average well-being by performing an ordered probit of subjective well-being on time fixed effects.³³ By construction, the levels of these four series are not comparable, and hence comparisons within, but not between, series are valid. Figure 18 shows the home/life satisfaction-GDP gradient within each of these periods, and it is clear that throughout the period in which Japan moved from poor to affluent (the first three panels), subjective well-being rose with GDP per capita. The far right panel shows that since 1992 the Japanese economy has shown very little growth, and life satisfaction has fallen sharply.

Figure 19 is a time-series plot of economic progress and subjective well-being in Japan. The top panel shows roughly three episodes in Japanese economic history, corresponding roughly to the changes in the survey questions regarding subjective well-being discussed above: spectacular growth during the period 1958-69, spanning one series break; slower growth from 1970 to 1991; and then anemic growth from 1992 onward, which coincided with the emergence of large-scale unemployment. The symbols in the bottom panel of the Figure show the corresponding (and incomparable across periods) movements in subjective well-being within each of the periods for which consistent data exist.

³³ Table 5 shows the proportions coded as “not sure,” “don’t know,” or “none of the above,” but we simply drop these observations from the rest of the analysis.

In an attempt to create a consistent series across the last fifty years, we pool each of these time series and run the following regression to estimate the extent of the relevant series breaks, while controlling for secular and cyclical influences:

$$\begin{aligned} \text{Well-being}_t = & -1.67 - 0.26 * I(1964 \leq \text{year} \leq 1969) - 0.54 * I(1970 \leq \text{year} \leq 1991) - 0.59 * I(1992 \leq \text{year}) \\ & (0.49) (0.07) \qquad (0.11) \qquad (0.14) \\ & -0.06 * \text{Unemployment rate}_t + 0.24 * \log(\text{GDP per capita}_t) \qquad n=51 \\ & (0.02) \qquad (0.06) \end{aligned}$$

The coefficients on each of the three dummy variables reveal that the changes in the survey question did in fact yield statistically significant (and clearly economically important) changes in estimated well-being. Making adjustments for the series breaks suggested by this regression results in the gray line in the bottom panel of Figure 19. This time series suggests that subjective well-being did in fact grow strongly in Japan, at least through the period in which GDP grew most strongly. The regression also finds an important role for unemployment, and this factor explains most of the sharp decline in subjective well-being through the 1990s, as well as the reversal over the past few years as unemployment has started to decline.³⁴ The unemployment coefficient is roughly comparable to, although somewhat larger than, estimates for other OECD countries (Wolfers, 2003). We can also use this coefficient to back out a “cyclically adjusted” well-being series for Japan, also shown in the bottom panel of Figure 19. As should be clear, this series bears a strong relationship with GDP per capita, and indeed, the estimated coefficient, 0.24, is again roughly consistent with our other time-series findings.

Finally, other data also suggest that well-being in Japan has tracked the country’s economic development. For instance, from 1974 through 1991 the same survey also asked, “How do you feel about your life now?” and the proportion answering “perfectly complete” or “somewhat complete” trended strongly upward. A somewhat different version of the question was asked from 1992 through 2007, and the proportions feeling perfectly or somewhat complete show a slow decline over this later period. The World Values Survey also provides useful time-series comparisons: in 1981, 16 percent of Japanese respondents reported being very happy, rising to 18 percent in 1990 and then 34 percent in 1995, before falling slightly to 29 percent in 2000. Life satisfaction data from that survey yield a less clear trend, but given the impact of changes in question ordering, it is worth noting that the decline in life satisfaction in Japan was smaller than that experienced in most other countries. Other early assessments of well-being were shown in Figure 1: in each of the comparisons in which Japan is included (the 1960 “Patterns of Human Concerns” surveys, the 1965 World Survey, and the 1975 Kettering survey, shown in Figure 6), subjective well-being in Japan was consistent with its moderate level of economic development. More recent surveys (such as the World Values Survey or the Gallup World Poll) show that Japan’s well-being is now at a level consistent with its status today as an affluent country.

³⁴ If instead we estimated this equation without controlling for unemployment, the estimated coefficient on log GDP per capita would be 0.16 (with a standard error of 0.06).

The most widely used dataset for analyzing happiness in the United States is the General Social Survey, conducted on a nationally representative sample of about 1,500 respondents each year from 1972 through 1993 (except 1992), rising to around 3,000 respondents every second year from 1994 through 2004, and 4,500 respondents in 2006. These repeated cross-sectional surveys ask, “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?” Many researchers have examined the trend in US happiness over this period and all have come to the same conclusion: the US has not gotten any happier over this time period and has even experienced a mild decline in happiness (Easterlin, 1995).³⁵ Our analysis turns up similar findings. The top panel of Figure 20 plots the coefficients from an ordered probit regression of happiness on year fixed effects.³⁶ These data suggest a very mildly declining happiness trend through this period (slope = -0.0010, with a standard error of 0.0008), which suggests that our happiness index declined by about 0.035 points between 1972 and 2006 (with a 95 percent confidence interval around this decline ranging from -0.09 to +0.02).

The middle panel of Figure 20 shows that log real GDP per capita rose by 0.66, (or 66 log points) over the same period, and the juxtaposition of this income growth with a roughly flat happiness trend appears to provide useful support for the Easterlin paradox. Indeed, a happiness-income gradient of 0.4 would have led one to expect the happiness index to have risen by 0.26 points. Translating this to the individual happiness categories, U.S. GDP growth from 1972 to 2006 was enough to suggest that by the end of the sample, another 10 percent of the population should have been “very happy,” and the proportions “not too happy” and “fairly happy” should have been about 4 and 6 percentage points lower than actually observed, respectively. Moreover, there is clear evidence of the absence of a time series happiness-income relationship here—the data clearly reject the view that happiness grew as predicted by the happiness-income gradient estimated within the United States or across countries. Although the U.S. time series is thus a data point supporting the Easterlin paradox, it should be regarded as an interesting exception warranting further scrutiny.

To better understand trends in happiness within the United States and its relationship to recent income growth, we look more closely at the patterns of income growth. In particular, the fruits of economic growth through this period were quite unequally distributed (Stevenson and Wolfers, 2008).³⁷ From 1972 through 2005, data from the Current Population Survey (CPS) suggest that average real household income grew by only 15 to 20 percent in each of the three bottom quintiles; the fourth quintile experienced growth of 30 percent, and only the top quintile realized income growth of 59

³⁵ He shows a slight decline in happiness between 1972 and 1991 in the United States and in Easterlin (2005c) finds no trend in happiness between 1972 and 2002. Blanchflower and Oswald (2004) report that well-being declined between 1972 and 1998. Stevenson and Wolfers (2007) find that well-being has been flat for men between 1972 and 2006 and has declined among women over this period.

³⁶ We have corrected these data for the biases due to changes in question ordering noted by Smith (1979, 1990). Stevenson and Wolfers (2008) provide a detailed explanation of how these corrections are made and show their impact on individual years.

³⁷ They examine how happiness has changed across demographic and socioeconomic groups in the US. They find that some groups—non-whites—have gotten happier, while others—those with less than a college degree—have gotten less happy. Overall, however the variance of happiness has declined both within and between groups.

percent (DeNavas-Walt, Proctor, and Lee, 2006). In turn, the top two quintiles of the household income distribution experienced mild growth in happiness, while happiness actually declined for the bottom three quintiles.

The family income data recorded in the GSS suggest roughly similar real income growth: an average increase of about 32 percent over the full sample, which was quite unequally distributed, with real declines reported in the bottom quintile. Although the CPS data reported above are surely a more reliable indicator of national trends in the income distribution, the family income data collected in the GSS may speak to the characteristics of the particular sample for whom we have happiness data.

Given these unbalanced gains, it is worth asking how the income-happiness link at the micro level aggregates to yield the macroeconomic income-happiness link. In the simple case in which income gains accrue proportionally across the distribution, individual happiness-log income functions aggregate to a macro-level linear relationship between average log income and happiness aggregates. However, the sharp rise in inequality over recent decades drives a large wedge between the rise in the log of average income (which is what we typically observe in macro data) and the average of log income (which is the relevant aggregate for predicting average happiness).

We computed the rise in income inequality in both the CPS and the GSS samples. From 1972 through 2006, the CPS measure of the log of average real household income rose by 41 log points, while inequality—as measured by the mean log deviation—rose by 19 log points.³⁸ Together these numbers imply that the average of log household income rose by only 22 log points over the full sample. For the GSS, the rise in the log of average family income is slightly smaller, at 32 log points, and the measured rise in inequality (again measured as the mean log deviation) is 15 log points.

Thus, within the GSS sample, the average of the log of family income has risen by only around 17 log points since 1972 (equivalent to an annual rate of growth of only around 0.5 percent a year).³⁹ Based on a happiness-income gradient of around 0.4, it seems reasonable to expect that happiness in the United States would have been basically flat over the past thirty-five years (or more precisely, to have risen by only $0.4 \times 0.17 = 0.07$ points). Thus, by refocusing our attention on the appropriate macroeconomic aggregate (in the bottom panel of Figure 20), it can be seen that the U.S. experience could be roughly consistent with the accumulated evidence of a robust happiness-income link.

Moreover, many other societal trends beyond trends in GDP per capita may influence trends in happiness. As such, no single national case study can be dispositive in our effort to understand how national well-being changes with

³⁸ It is worth being a bit more explicit about how reasonably robust economic growth translates into weaker growth in the average log income. From 1972 through 2006, real GDP per capita grew by 93 percent, or 66 log points, and disposable personal income per capita rose by a similar amount. Beyond these aggregate data from the Bureau of Economic Analysis (BEA), the Census Bureau also calculates income per household from the March CPS. These alternative data suggest that income per capita (in 2005 dollars) rose by 66 percent, or by 51 log points. Much of the gap between the BEA and the CPS measures reflects differences in deflators. From 1972 through 2006 the CPI-U-RS (the version of the consumer price index used by the Census Bureau to deflate the CPS data) rose 11 log points more than the GDP deflator. This difference would be even larger (22 log points) if we deflated instead by the official CPI-U series. On a per household basis, the rise in the log of average income was even less impressive, at only 41 log points.

³⁹ We deflate the GSS income data using the CPI-U-RS rather than the CPI-U-X. If instead we used the official deflator, the average log of family income would have registered barely any growth at all.

economic development. The Easterlin paradox suggests that on average countries will not get happier as they get richer. The evidence from the countries for which we have time series observations, on balance, casts doubt on the Easterlin paradox, with Europe, Japan, countries in the World Values Survey becoming happier, on average, as income grows.

VI. Alternative Measures of Subjective Well-Being

Our discussion so far has analyzed three basic measures of subjective well-being: reports of happiness, of life satisfaction, and of well-being expressed in terms of a “ladder” with the best and worst possible lives at top and bottom. Yet this still leaves a lot unsaid about the subjectively experienced lives of the rich and the poor. Fortunately, recent major advances in cross-country data collections have started to paint a broader picture of subjective well-being.

We begin by analyzing the battery of ten questions typically grouped as the Bradburn Affect Balance Scale, which were included in the first two waves of the World Values Survey (Bradburn, 1969). This scale is intended to separately assess both positive and negative affect, by probing direct reports of whether various pleasant and unpleasant feelings have been experienced recently. This battery of questions asks the respondent whether, during the past few weeks, he or she has had any of five positive experiences (“pleased about having accomplished something,” “proud because someone had complimented you on something you had done,” “particularly excited or interested in something,” “things going your way,” or “on top of the world/feeling that life is wonderful”) and five negative experiences (“bored,” “upset because somebody criticized you,” “so restless you couldn’t sit long in a chair,” “very lonely or remote from other people,” or “depressed or very unhappy”) (Bradburn, 1969, chapter 4).⁴⁰

We analyze each question separately in Table 6, and because our dependent variable is binary (whether or not the respondent reported experiencing each feeling), we use probit regressions. To separately isolate the between-country and the within-country variation, we run one regression in which log GDP per capita is the only independent variable, and then another substituting log household income for GDP per capita, and controlling for country fixed effects. To maintain some consistency in the units, we report actual probit coefficients rather than the elasticity of predicted probabilities.

The first panel of Table 6 shows that in both within-country and between-country comparisons, measures of positive affect are all positively, and measures of negative affect negatively, associated with income. While the within-country comparisons are typically statistically significant, the limited number of country observations gives us less precise and hence less significant estimates (we cluster the standard errors in these regressions by country). The magnitudes of the estimated within- and between-country coefficients are roughly similar, and the within-country estimates typically lie in the 95% confidence interval surrounding the between-country estimates. Putting these together into a measure of net affect (the average number of positive experiences less the average number of negative experiences) yields a measure that

⁴⁰ He found that among his U.S. sample, within the group of positive or negative questions, responses tended to be highly correlated, but that responses between questions probing “positive affect” and “negative affect” were not closely related. Moreover, individual evaluations of happiness appear to reflect positive and negative affect in roughly equal measure.

is strongly related to both log household income and log GDP per capita, and these estimates reflect the impact of income on positive and negative affect in roughly equal measure.⁴¹

Figure 21 presents the cross-country comparisons graphically. The top row reveals that in richer countries a larger proportion of the population is more likely to report each positive experience (except feeling “particularly excited or interested in something”), and the bottom row shows that a smaller proportion of the population in richer countries typically reports negative experiences. The regressions reported in the figure show how the proportion of the population agreeing with each statement rises or falls with log GDP per capita (and hence these estimates are scaled differently than the probit coefficients reported in Table 6). Interestingly, as in the analysis of self-assessed happiness, Nigeria (the poorest country in this figure) is an outlier for all of the measures of positive affect, with Nigerians reporting a much higher likelihood of experiencing positive feelings than residents of other low-income countries. The bottom row shows the relationship between each of the five measures of negative feelings and GDP per capita. In all cases the negative affect-GDP gradient is negative, with a higher proportion of people in poor countries experiencing negative feelings. (These measures of negative affect suggest that Nigerians have a more typical experience for their income.)

We next turn to a particularly rich series of well-being questions contained in the Gallup World Poll. Respondents are asked to report whether they experienced “the following feeling during a lot of the day yesterday.” The feelings include enjoyment, physical pain, worry, sadness, boredom, depression, anger, and love. The top of the bottom panel of Table 6 shows that among the positive emotions, the enjoyment-income gradient is positive and similar for both the between- and the within-country estimates. More income is clearly associated with more people having enjoyment in their day. Love is less clearly related to income, although within countries, more income is associated with being more likely to experience love. Among the negative emotions, physical pain, boredom, depression, and sadness are all lower at higher levels of income, at both the national and the individual level. The within-country estimates also reveal that worry and anger fall with income.

Figure 22 allows a fuller examination of the proportion of people in a country experiencing these emotions and GDP per capita. The percent of people in a country who enjoyed the previous day rises from an average of 65 percent in low income countries to 80 percent in the wealthiest countries. Depressions, pain, boredom, and anger all appear to fall linearly with rises in log GDP per capita. The magnitudes of these relationships are large: compared with the poorest countries those in the wealthiest countries are a third less likely to experience pain or depression and a fifth less likely to report boredom.

The final set of regressions analyze the relationship between income and some more specific experiences in people’s lives, such as feeling respected, smiling, engaging in interesting activities, feeling proud, and learning. Income is positively related to wanting to have more days like yesterday, with feeling well rested, with feeling treated with respect, with being able to choose how to spend your time, with smiling or laughing, with feeling proud, with having done

⁴¹ The summary measure of net affect is computed by adding up the positive and negative measures, with zero indicating an equal number of positive and negative experiences.

something interesting, and with eating good tasting food. For most of these assessments the within-country coefficients are larger than the between-country coefficients, although there are some notable exceptions. Figure 23 plots the proportion of people having these experiences in each country against GDP per capita and feeling respected, or having eaten good food display particularly strong relationships with GDP per capita. The proportion of people smiling or laughing rises with income both within and between countries, although the coefficient estimates point to a somewhat stronger relationship to income within countries. This last measure is particularly interesting, as smiling has been shown to be correlated with reported levels of happiness or life satisfaction. Indeed, in these data, people who report smiling more also tend to report greater life satisfaction. Finally, both Table 6 and Figure 23 point to an increasing ability to choose how you spent your time during the day as income rises.

All told, these alternative measures of well-being paint a somewhat more nuanced picture of the different experiences of rich and poor people within countries, and between rich and poor countries. Moreover, these data point to a robust relationship between greater income and greater reported well-being. We suspect that these intriguing new cross-country data collections will launch a productive research program aimed at better understanding the drivers of the robust well-being–income gradient we have identified.

VII. Discussion

This paper has revisited—and where appropriate, revised—the stylized facts regarding the link between subjective well-being and income. Our analysis encompasses virtually all of the extant data linking happiness or life satisfaction to income. Moreover, we have endeavored to place this analysis in a single coherent framework that allows us to make meaningful comparisons across different surveys and different ways of asking about subjective well-being. We were motivated to better understand the Easterlin paradox, and so we have analyzed separately three relationships between income and happiness: that obtained from contrasting rich and poor members of a society, that obtained from contrasting rich and poor countries, and that obtained from observing the paths of average happiness as the average income of countries change. Our measurement framework allows us to assess the extent to which these relationships may differ.

Our key contribution is the finding that the relationship between subjective well-being and income within countries (that is, contrasting the happiness of rich and poor members within a country) is similar to that seen between countries, which in turn is similar to the time-series relationship (comparing the happiness of countries at different points in time as they get richer or poorer). In multiple datasets from several decades and covering various populations, we estimate well-being–income gradients that tend to be centered around 0.4. We estimate slightly steeper gradients when comparing well-being between countries, although reading across datasets and taking account of sampling error, we can reject neither the hypothesis that the gradients are the same within and between countries, nor the hypothesis that there are small differences between the two.

The time-series part of our analysis is necessarily only suggestive: repeated (and comparable) surveys of subjective well-being data are both noisy and scarce, and hence they speak less clearly. In many cases we find happiness within a country rising during periods of economic growth and rising most rapidly when economic growth is more rapid. The United States stands out as a notable exception: Americans have experienced no discernable increase in happiness over the past thirty-five years (and indeed, happiness among U.S. women has declined). In contrast, Japan stands out as a remarkable success story, recording rising happiness during its period of rapid economic growth. So, too, life satisfaction has trended upward in Europe, and this trend has been most evident in those countries in which economic growth has been most robust. All told, our time-series comparisons, as well as evidence from repeated international cross-sections, appear to point to an important relationship between economic growth and growth in subjective well-being. Quantitatively, the time-series well-being–GDP gradient yields a role for income similar to that seen in our within- and between-country contrasts. Taken as a whole, the time-series evidence is difficult to reconcile with earlier claims that economic growth yields no boost to happiness.

Easterlin and others have argued that comparisons of rich and poor people within a country yield starker happiness differences than comparisons of rich and poor countries, and have cited this as evidence that relative income differences are a key driver of happiness. Graham notes that “a common interpretation of the Easterlin paradox is that humans are on a ‘hedonic treadmill’: Aspirations increase along with income and, after basic needs are met, relative rather than absolute levels matter to well-being” (2008, p. 77). In its strong form this hypothesis suggests that people (and public policy) are powerless to deliver lasting gains in happiness, because individual happiness returns inexorably to one’s set point of happiness. Our findings clearly falsify this strong form of adaptation: we find that those enjoying materially better circumstances also enjoy greater subjective well-being and that ongoing rises in living standards have delivered higher subjective well-being. However, milder forms of adaptation are potentially consistent with our findings.

Our findings point to an important role for absolute levels of income in shaping happiness and a lesser role for relative income comparisons than was previously thought. Equally, our findings are sufficiently imprecise that they may still admit a role for relative income comparisons in shaping subjective well-being. We find that estimates of the within- and between-country well-being-income gradient tend to cluster around 0.4, and we lack sufficient evidence to say that these gradients are clearly different. Thus, our evidence is consistent with the view that only absolute income matters to happiness (which would imply that the within- and between-country estimates are identical). Indeed, whereas previous analyses of the link between income and happiness suggested a *prima facie* case for relative income playing a dominant role, our updated reanalysis finds no such case.

Equally, our findings do admit the possibility of an interesting role for relative income comparisons. For instance, the within-country coefficient is typically about 0.3 and might be biased downward by the influence of transitory income in the cross section. Thus, perhaps the true within-country coefficient is 0.45, and our estimates are consistent with a view that the between-country coefficient is about 0.36 (with the time-series coefficient a bit weaker still). This is consistent with both absolute and relative income impacting well-being, with the former having a weight about four times as large as

the latter.⁴² Thus, our findings should not be interpreted as falsifying the view that relative income plays a role in shaping happiness, although they do bound the extent to which relative income may matter.

In light of this range of possible interpretations, we would suggest that more fine-grained evidence on the role of relative income should come from direct evidence of relative income shocks, such as that investigated by Luttmer (2005). In particular, a comparison of the between and within well-being-income gradients cast doubt on a large role for intranational relative income comparisons in determining happiness. While these comparisons do not speak to the role of international relative income comparisons, the earliest surveys—conducted in the 1940s—yielded similar estimates of the between country well-being-income gradient to those seen today. However, these early surveys involved sufficiently few data points that it is impossible to know with precision whether the gradient between subjective well-being and income across countries has changed over time to reflect the growing awareness of the opportunities available to others around the world. In short, the most compelling evidence for the importance of absolute income over relative income in determining happiness may eventually come from the times series evidence.

Finally, we should note that our analysis has largely focused on establishing the magnitude of the bivariate relationship between subjective well-being and income, rather than tracing the causal effects of income on happiness. We believe that further research aimed at better understanding the causal pathways will be fruitful.

⁴² For instance, these findings are consistent with a simple happiness function: $Happiness = 0.36 \log(\text{individual income}) + 0.09 \log(\text{individual income/average income})$

VIII. Works Cited

- Blanchflower, David, and Andrew Oswald. "Well-Being Over Time in Britain and the USA." *Journal of Public Economics* 88, no. 7-8 (2004): 1359-1386.
- Bradburn, Norman M. *The Structure of Psychological Well-Being*. Chicago: Aldine Publishing Company, 1969.
- Buchanan, W., and H. Cantril. *How Nations See Each Other*. Urbana: University of Illinois Press, 1953.
- Campbell, John Y., and N. Gregory Mankiw. "Permanent income, current income, and consumption." *Journal of Business and Economic Statistics* 8 (1990): 269-279.
- Cantril, H. *The pattern of human concerns*. New Brunswick, New Jersey: Rutgers University Press, 1965.
- Cantril, Hadley. *Public Opinion, 1935-1946*. Princeton, New Jersey: Princeton University Press, 1951.
- Clark, Andrew E., Paul Frijters, and Michael A. Shields. "Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles." *Journal of Economic Literature* 46, no. 1 (2008): 95-144.
- Deaton, Angus. "Income, Health and Well-being around the World: Evidence from the Gallup World Poll." *Journal of Economic Perspectives*, 22(2) (2008): 53-72.
- DeNavas-Walt, Carmen, Bernadette D. Proctor, and Cheryl Hill Lee. *Income, Poverty and Health Insurance Coverage in the United States: 2005*. Current Population Reports, Washington, D.C.: U.S. Census Bureau, 2006.
- Di Tella, Rafael, Robert J MacCulloch, and Andrew J Oswald. "The Macroeconomics of Happiness." *Review of Economics and Statistics* 85, no. 4 (2003): 809-827.
- Diener, Ed. "Subjective Well-Being." *Psychological Bulletin* 95, no. 3 (1984): 542-575.
- Diener, Ed. "Guidelines for National Indicators of Subjective Well-Being and Ill-Being" *Journal of Happiness Studies*, vol 7, no. 4 (2006), 397-404.
- Diener, Ed, Richard E. Lucas, and Christie Napa Scollon. "Beyond the Hedonic Treadmill: Revising the Adaptation Theory of Well-Being." *American Psychologist* 61, no. 4 (2006): 305-314.
- Diener, E., & Seligman, M. E. P. (2004). "Beyond Money: Toward an Economy of Well-Being". *Psychological Science in the Public Interest*, 5, no. 1, 1 – 30.
- Diener, E., & Tov, W. (in press). "Culture and subjective well-being". In S. Kitayama & D. Cohen (Eds.), *Handbook of cultural psychology*. New York: Guilford.
- Easterlin, Richard A. "Diminishing Marginal Utility of Income? Caveat Emptor." *Social Indicators Research* 70, no. 3 (2005b): 243-255.

- Easterlin, Richard A. "Does economic growth improve the human lot? Some empirical evidence." In *Nations and Households in Economic Growth: Essays in Honor of Moses Abramowitz*, by Paul A David and Melvin W. Reder. New York: Academic Press, Inc., 1974.
- Easterlin, Richard A. "Does Money Buy Happiness?" *The Public Interest* 30 (1973): 3-10.
- Easterlin, Richard A. "Feeding the Illusion of Growth and Happiness: A Reply to Hagerty and Vennhoven" *Social Indicators Research*, 2005c, 74:3, 429 – 443.
- Easterlin, Richard A. "Income and Happiness: Towards a Unified Theory." *The Economic Journal* 111, no. 473 (2001): 465-484.
- Easterlin, Richard A. "Will Raising the Incomes of All Increase the Happiness of All?" *Journal of Economic Behavior and Organization* 27, no. 1 (1995): 35-48.
- Eid, Michael, and Ed Diener. "Global Judgments of Subjective Well-Being: Situational Variability and Long-Term Stability." *Social Indicators Research* 65, no. 3 (2004): 245-277.
- Ekman, P., & Friesen, W.V. (1971). Constants across cultures in the face and emotion. *Journal of Social and Personality Psychology*, 17, 124-129.
- Ekman, P., Friesen, W.V., O'Sullivan, M., Chan, A., Diacoyanni-Tarlatzis, I., Heider, K., Krause, R., LeCompte, W.A., Pitcairn, T., Ricci-Bitti, P.E., Scherer, K., Tomita, M., & Tzavaras, A. (1987). Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of Social and Personality Psychology*, 53, 712-717.
- Frank, Robert H. "Does Absolute Income Matter?" In *Economics and Happiness*, edited by P.L. Porta and L. Bruni, 65-91. Oxford University Press, 2005.
- Frey, Bruno S., and Alois Stutzer. "What Can Economists Learn from Happiness Research?" *Journal of Economic Literature* 40 (2002): 402-435. Gallup, George H. "Human Needs and Satisfaction: A Global Survey." *The Public Opinion Quarterly* 40, no. 4 (1976): 459-467.
- Hagerty and Veenhoven." *Social Indicators Research* 74, no. 3 (2005a): 429-443.
- Kahneman, Daniel, Alan B. Krueger, David Schkade, Norbert Schwarz, and Arthur A. Stone. "Would You Be Happier If You Were Richer? A Focusing Illusion." *Science* 312 (June 2006): 1908-1910.
- Kahneman, Daniel, and Alan B. Krueger. "Developments in the Measurement of Subjective Well-Being." *Journal of Economic Perspectives* 20, no. 1 (2006): 3-24.
- Kenny, Charles, "Does Growth Cause Happiness, or Does Happiness Cause Growth?" *Kyklos*, vol. 52, 3-26 (1999).
- Layard, Richard. "Happiness: Has Social Science a Clue" Lionel Robbins Memorial Lectures 2002/3. Delivered on 3, 4, 5 March 2003 at the London School of Economics.
<http://cep.lse.ac.uk/events/lectures/layard/RL030303.pdf>
- Layard, Richard. *Happiness: Lessons from a New Science*. London: Penguin, 2005a.

- Layard, Richard. "Human Satisfaction and Public Policy." *The Economic Journal*, 1980: 737-750.
- Layard, Richard. "Rethinking Public Economics: The Implications of Rivalry and Habit." In *Economics and Happiness*, edited by P.L. Porta and L. Bruni, 147-169. Oxford University Press, 2005b.
- Leigh, Andrew and Justin Wolfers, "Happiness and the Human Development Index: Australia is Not a Paradox", *Australian Economic Review*, 39(2) 176-184, June 2006.
- Luttmer, Erzo F.P. "Neighbors as Negatives: Relative Earnings and Well-Being," *Quarterly Journal of Economics*, 120(3), pp. 963-1002, August 2005.
- Lleras-Muney, Adriana. "The Relationship Between Education and Adult Mortality in the U.S." *Review of Economic Studies* 72, no. 1 (2005).
- Maddison, Angus. *Historical Statistics for the World Economy: 1-2003 AD*. 2007.
- Oswald, Andrew J. "On the Curvature of the Reporting Function from Objective Reality to Subjective Feelings." *Economics Letters*, Forthcoming.
- Rivers, Douglas, and Quang H. Vuong. "Limited Information Estimators and Exogeneity Tests for Simultaneous Probit Models." *Journal of Econometrics* 39 (1988): 347-366.
- Smith, Tom W. "Time Trends, Seasonal Variations, Intersurvey Differences, and Other Mysteries." *Social Psychology Quarterly* 42, no. 1 (1979): 18-30.
- Smith, Tom W. *Timely Artifacts: A Review of Measurement Variation in the 1972-1989 GSS*. GSS Methodological Report, NORC, University of Chicago, 1990.
- Smith, Tom W. *Unhappiness on the 1985 GSS*. GSS Methodological Report, Chicago: NORC, 1986.
- Stevenson, Betsey, and Justin Wolfers. "Happiness Inequality in the United States." *Journal of Legal Studies*, forthcoming 2008
- Stevenson, Betsey, and Justin Wolfers. *The Paradox of Declining Female Happiness*. mimeo, University of Pennsylvania, 2007.
- Strunk, Mildred. "The Quarter's Polls." *Public Opinion Quarterly* 14, no. 1 (1950): 174-192.
- van Praag, B.M.S. and A. Ferrer-i-Carbonell, *Happiness Quantified. A Satisfaction Calculus Approach*, Oxford University Press, New York, USA, 2004.
- Veenhoven, Ruut. *Happiness in Nations, Subjective Appreciation of Life in 56 Nations 1946-1992*. Rotterdam: Erasmus University, 1993.
- Veenhoven, Ruut. *World Database of Happiness, Trend in Nations*. Rotterdam: Erasmus University.
- Veenhoven, Ruut "Is Happiness Relative?" *Social Indicators Research* 24, 1991, pp. 1-34
- Wolfers, Justin. "Is Business Cycle Volatility Costly? Evidence from Surveys of Subjective Well-being." *International Finance* 6, no. 1 (2003): 1-26.

Appendix A Cardinalizing Happiness and Life Satisfaction

Our approach to constructing an index of average well-being in a country-year (or country-wave) is to report the coefficient from an ordered probit regression of subjective well-being on country by year (or country by wave) fixed effects. This appendix tries to make this approach more transparent, thereby demonstrating how to reconcile our results with alternative approaches.

A simple approach to aggregating data on subjective well-being involves arbitrarily assigning to qualitative categories scores equal to their rank order. Thus, in the World Values Survey, a response of “not at all happy” is given a value of 1, “not very happy” a value of 2, “quite happy” a value of 3, and “very happy” a value of 4. Average well-being is then calculated as the simple average of these values. This appears to be the most common approach in the literature.

A key difficulty with this approach is that the scaling of well-being measures from different surveys will vary, depending on whether the question used a three-, four-, five-, seven-, ten-, or eleven-point scale (others also occur). In turn, this approach yields estimates of the well-being–income gradient that are neither comparable across surveys nor have an obvious economic interpretation.

Thus, a somewhat more satisfying index might be constructed by normalizing the dependent variable (subtracting its mean and dividing by its standard deviation), which would yield a common metric. Moreover, this metric would have an economic interpretation, scaling differences in well-being relative to its cross-sectional standard deviation. (This approach yields results very close to our approach.)

Even so, the limitation of this approach is that it imposes a linear structure, implying, for example, that the difference between being “not very happy” and “not at all happy” is equal to the difference between being “quite happy” and “not very happy.” Although psychologists have often been willing to accept that the subjective distances between successive points on category scales are similar, we can use data on the proportions of the population who report themselves as being in each category to relax (or test) this assumption.

To make use of these population proportions, the ordered probit (Figure A1) makes a parametric assumption, imposing normality on the distribution of the underlying latent “well-being” measure. Two normalizations are also imposed: that the latent variable has a mean of zero and that it has a standard deviation of one. The country or country by wave fixed effects we estimate (and interpret as well-being) are simply shifts in the mean of this distribution.

There is a very simple mapping between our results and the simple approach described above: whereas the “value” of each categorical answer is simply imposed in the simple approach, in our approach it is equal to the expected value of a standard normal variable, conditional on being between the estimated upper and lower cutoff points. Van Praag and Ferrer-i-Carbonell (2004) describe this as “probit-adapted OLS.” Table A1 reports the mapping between the underlying categorical responses, the standardized categorical responses, and our scaling derived from these ordered probits.

As the table shows, our method yields a cardinalization that is very similar to that obtained simply by standardizing the variables used in the usual approach. This provides a useful approximation: to map results in other studies to ours, one can simply divide the estimates of the well-being–income gradient estimated in those studies by the standard deviation of well-being. These results are graphed in Figure A2, which shows the cardinalization imposed by our ordered probit procedure in each of three key datasets. As should be clear, our procedure is well approximated by a linear transformation of the simpler approach, which simply analyzes the ordered categories directly.

Next, it is worth assessing this approach relative to four alternative metrics, of which three are typically used in the literature; the fourth is an interesting extension of our approach.

Means: Continuing with the most common approach in the literature, the simplest (and most transparent) approach is to take the ordinal ranking of alternatives as cardinal measures of happiness. This approach may make more sense when analyzing questions that ask respondents to give a cardinal response (such as the World Values Survey life satisfaction question, which asks for a response on a scale of 1 to 10).

Population proportions: An alternative involves reporting the proportion of the population reporting themselves as, say, “quite happy” or “very happy.” This approach has the advantage that it yields a natural scaling (from 0 to 1) and is directly interpretable. One difficulty is that this approach may lead changes in the dispersion of happiness to be interpreted as changes in the average level of happiness. To minimize this possible confound, one typically chooses a cutoff near the median response. However, the median response in poor countries can turn out to be a far more common response in rich countries.

Ordered logits: The ordered logit is similar to our ordered probit approach but assumes a slightly different (fatter-tailed) distribution of the latent “happiness” in the population. The logistic function also imposes a standard deviation on the latent variable of $\pi/\sqrt{3}$, which makes the coefficients somewhat differently scaled than with the ordered probit.

Heteroscedastic ordered probit: The ordered probit imposes an equal variance in residual happiness, whereas the heteroscedastic ordered probit allows both the mean and the variance of happiness to vary by country-year. Alternatively phrased, this approach relaxes the assumption of similar cutoff points for each country and year, allowing proportional shifts in these cutoff points, by country-year.⁴³

Figures A3 through A5 compare these alternative aggregators with our ordered probit approach, analyzing separately the satisfaction ladder from the Gallup World Poll and the life satisfaction and happiness data, by country and wave, in the World Values Survey. These figures suggest that alternative methods of aggregating subjective well-being all tend to yield highly correlated estimates.

⁴³ Stevenson and Wolfers (2008) provide greater detail on this method.

Appendix B Comparing Countries in the World Values Survey

A few low-income countries in certain waves of the World Values survey were explicitly not designed to be representative of their entire population. These selected samples add measurement error that is, in many cases, correlated with income, education, and other factors that are related to subjective well-being. In most cases these non-representative samples will lead average subjective well-being to be over-estimated relative to the population mean. Moreover, non-representative sampling typically occurred in countries with low per capita GDP. For many of these countries, the sampling frame changed in successive waves to become representative of the entire population (and this change occurred in parallel with rising GDP). As such, we should expect that for these countries average subjective well-being in the population will decline over time as more rural, low-income, and low-education citizens are included in the sampling frame.

In the results presented throughout the paper, we have excluded a few countries in particular waves because the survey notes that the sampling frame was specifically not representative of the entire country and no compensatory sampling weights are provided.⁴⁴ In this appendix we detail the reasons why these observations were excluded and show how our results are impacted when these country-wave observations are included. We begin by documenting the sampling issues specific to countries which are impacted:

Argentina was surveyed in all four waves; however in the first 3 waves sampling was limited to the urbanized central portion of the country and resulted in a wealthier, more educated sample of Argentineans compared with the population average. In the 1999-2004 the sample was designed to be representative of the entire country. We include in our analysis only observations from Argentina in the 1999-2004 wave.

Bangladesh was included in two waves. In the 1994-99 wave the survey oversampled men and people in urban areas "to reflect the fact that awareness is more widespread in the urban areas".⁴⁵ Sampling weights are not provided and it is therefore not possible to correct for the oversampling. The 1999-2004 wave was designed to be representative. We include in our analysis only observations from Bangladesh in the 1999-2004 wave.

Chile was included in three waves; however in the 1989-93 and 1994-99 waves the sample was limited to the central portion of the country, which covers slightly less than two-thirds of the population and has an average income about 40% higher than the national average. In the 1999-2004 survey the sample was drawn from 29 selected cities. As a result of these partial-country samples, we exclude Chile from all of our analysis.

China was included in three waves. For the 1989-93 wave, the survey notes state that they "undersampled the illiterate portion of the public and oversampled the urban areas and the more educated strata." Moreover, the survey notes explicitly state that "the oversampled groups tend to have orientations relatively similar to those found in industrial societies" and that "the data probably underestimate the size of cross-national differences."⁴⁶ In the 1994-

⁴⁴ Many samples over-sample specific groups, but sampling weights are provided in order to yield nationally representative estimates. Sampling weights cannot adjust for the fact that some groups were not samples at all.

⁴⁵ <http://www.worldvaluessurvey.com/> Survey notes for Bangladesh (BD_WVS 1995).

⁴⁶ <http://www.worldvaluessurvey.com/> Survey notes for China (CN_WVS 1990).

99 wave, a random sample of central China was done, which contains about two-thirds of the population. Sampling weights are not provided for any of the waves. These two surveys are quite different from each other: in the first wave only 1% of sample was from a town with less than 50,000 people, while in the second wave 63% were. In the first wave, 60% of respondents were men, the proportion falls to 53% in the second wave. In the 1999-2004 wave the sampling frame was drawn from a previous nationally representative survey and was conducted throughout the entire country, with the exception of six remote provinces which comprise 5.1% of the total population. The sample was also limited to those ages 18-65. Despite some limitations, we believe that the last wave is approximately representative and include observations from this last wave in our analysis. Observations from the earlier waves are excluded. If we were to also exclude the final wave, there is no notable impact on our analysis.

The **Dominican Republic** was included only in the 1995-99 wave. The sample included only 18-49 year olds and only 4 communities were chosen to be surveyed. We exclude observations from the Dominican Republic from our analysis.

Egypt was included only in the 1999-2004 wave. The survey notes that a disproportionately high percentage of housewives were included; examining the survey we find that women are also disproportionately from large urban areas, in particular Cairo. Since no sample weights are provided, we exclude Egypt from our analysis.

India was included in three waves. In the 1989-93 wave the sample was designed such that 90 percent of respondents were literate (compared to a population average of less than 50 percent). Interviews were carried out in the eight most widely-spoken languages of India, but the rural 10 percent of the sample was confined to the five (out of 14) Hindi-speaking states in the sample. In the 1994-99 wave, the survey was only conducted in Hindi (which is the language of less than half of the general population), the sample was stratified to allocate 90% of the interviews to urban areas and 10% to rural areas. In 1999-2004, the survey was designed to be representative of 97% of the population and was conducted in 10 languages. Sample weights were not provided for any of the waves. We include only this last wave in our analysis.

Nigeria was included in three waves. The 1989-93 and 1994-99 waves focused on the literate and urban portion of the population. As such, over 40% of the respondents in the first wave had attended university. This proportion falls to 23% in the second wave, which included a larger rural sample, and to 12% in the 1999-2004 wave which was designed to be representative of the population. We include only this last wave in our analysis.

Northern Ireland was included in the 1999-2004 wave, but was excluded due to missing GDP data.

Pakistan was included in two waves; however in the 1994-99 wave, sampling was only done in Punjab (a little over half of Pakistan). In 1999-2004, the sampling frame included the entire country. We include only the 1999-2004 observations in our analysis.

South Africa was included in three waves; however the first wave, 1989-93, over-represents minority races and blacks were only sampled in certain areas. Sampling weights were not provided. The next two waves, 1994-99 and 1999-2004 were designed to be representative of the population. We exclude observations from the first wave only.

Table B1 compares the main coefficient estimates from Tables 1-3 with that which we obtain by including all country-wave observations. The first panel shows the between-country analysis, with columns 1 and 3 reproducing the results shown in column 3 of Table 1 for life satisfaction and happiness respectively. Since the excluded observations typically represent a group with above average income and education (and hence, likely higher happiness), our expectation is that incorporating these countries will yield lower estimates of the well-being-income gradient. These estimates are shown in column 2 for life satisfaction, and column 4 for happiness. The first and the last wave show little impact on the estimated coefficient, as the samples are largely the same (only Argentina was excluded from the first wave and only Chile and Egypt from the last). The 1989-93 and 1994-99 waves yield larger differences as 6 countries were excluded from the former and 8 from the latter. As expected, including these biased samples attenuates the estimated coefficients substantially. Yet in all cases the estimated coefficient remains positive and statistically significant.

The second panel examines the impact of including the unrepresentative national samples on the within-country cross-sectional estimates. The first and third column reproduces the coefficients from column 2 of Table 2. Despite the truncation of the poor in these samples, the fact that subjective well-being is linearly related to log income suggests that excluding a portion of the income distribution will not bias the coefficient estimates. Moreover, as we show in Figure 10, most countries have a subjective well-being income gradient of around 0.4 and there is little systematic variation in that gradient. As such we should expect little difference in the Table 2 estimates that are obtained when we include observations from all of the country waves. Indeed, the estimated coefficients with the excluded samples, again shown in columns 3 and 4, are little different from those obtained without these countries.

Finally, the last panel reproduces the estimates shown in Table 3, column 2 in which we analyze the World Values Survey as a country-wave panel dataset using the country aggregated (macro) data. The Table 3 estimates are shown in columns 1 and 3, while the comparison estimates, in which unrepresentative country-wave observations are included, are shown in columns 2 and 4. The first row shows the simple bivariate well-being-GDP relationship, and hence pools both within-country and between country variation. These results are little impact by the inclusion of the unrepresentative country-wave observations. The second row includes country fixed-effects in these regressions and therefore isolates the within-country time series variation. The inclusion in columns 2 and 4 of countries whose sample becomes more representative as GDP grows, reduces the estimated coefficient. The third row adds controls for each wave of the World Values Survey in addition to the country controls. Again, the inclusion of the non-representative samples reduces the estimated coefficients. Finally, the last two columns considers first differences of consecutive country-wave observations and long-differences, including only the first and last observation for each country. Excluding differences involving countries where the survey frame changed yields positive robust estimates of a positive relationship between life satisfaction and income and happiness and income over time. Not surprisingly, including countries whose samples are becoming increasingly representative of the poor over time decreases these estimates substantially. Including the non-comparable intertemporal variation in well-being also yields less precise estimates. Even when these countries are included, the results are still roughly consistent with the null that the time series well-being-income gradient is close to the 0.4 range obtained from our between-country and within-country analyses.

Table 1. Cross-Country Regressions of Subjective Well-Being on GDP per Capita

Survey	Ordered probit regressions, micro data		OLS regressions, national data			Sample size
	Without controls	With controls	All countries	GDP per capita > \$15,000	GDP per capita < \$15,000	
Gallup World Poll, 2006: Ladder question	0.396*** (0.023)	0.422*** (0.023)	0.418*** (0.022)	1.076***† (0.211)	0.348*** (0.037)	139,051 (113 countries)
World Values Survey: Life Satisfaction						
1981-84 wave	0.525** (0.263)	0.291 (0.331)	0.498* (0.252)	1.677** (0.703)	0.722 (0.582)	23,537 (19 countries)
1989-93 wave	0.551*** (0.096)	0.551*** (0.096)	0.558*** (0.096)	0.504 (0.467)	0.391 (0.256)	50,553 (35 countries)
1994-99 wave	0.408*** (0.054)	0.418*** (0.054)	0.462*** (0.051)	0.327 (0.421)	0.394*** (0.084)	65,779 (45 countries)
1999-2004 wave	0.321*** (0.041)	0.329*** (0.041)	0.346*** (0.046)	0.455** (0.223)	0.208** (0.090)	94,224 (67 countries)
Combined, with wave fixed effects	0.373*** (0.038)	0.377*** (0.037)	0.398*** (0.040)	0.477** (0.198)	0.280*** (0.073)	234,093 (79 countries)
World Values Survey: Happiness						
1981-84 wave	0.650*** (0.250)	0.523*** (0.263)	0.569** (0.230)	1.662 (0.987)	0.550 (0.688)	22,294 (18 countries)
1989-93 wave	0.710*** (0.130)	0.725*** (0.128)	0.708*** (0.123)	0.328 (0.475)	0.144 (0.309)	49,281 (35 countries)
1994-99 wave	0.319*** (0.056)	0.335*** (0.056)	0.354*** (0.058)	0.248 (0.235)	0.212** (0.082)	63,785 (46 countries)
1999-2004 wave	0.118* (0.062)	0.138** (0.061)	0.126* (0.073)	0.766***† (0.218)	-0.146 (0.117)	92,799 (66 countries)
Combined, with wave fixed effects	0.229*** (0.055)	0.245*** (0.055)	0.244*** (0.063)	0.612***† (0.170)	-0.015 (0.100)	228,159 (79 countries)
Pew Global Attitudes, 2002: Ladder question	0.223*** (0.041)	0.242*** (0.040)	0.224*** (0.041)	0.466** (0.191)	0.168** (0.082)	37,974 (44 countries)

Notes: Table reports results of regressions of the indicated measure of well-being on log real GDP per capita. Numbers in parentheses are robust standard errors, clustered by country. Asterisks indicate statistically significant from zero at the *10 percent, **5 percent, and ***1 percent level; † denotes that the coefficient estimate for rich countries is statistically significantly larger than that for poor countries, at the 1 percent level.

Micro data: Ordered probit regressions, using data by respondent, of subjective well-being on log real GDP per capita for the respondent's country, weighting observations to give equal weight to each country × wave.

National data: National well-being index is regressed on log real GDP per capita. The well-being index is calculated in a previous ordered probit regression of well-being on country × wave fixed effects.

Controls: Include a quartic in age, interacted with sex, and indicators for missing age or sex.

Sample size: Only nationally representative samples are analyzed, which eliminated seventeen country-wave observations from ten countries in the World Values Survey (see appendix B for further details).

Gallup World Poll, 2006: Respondents were asked, "Please imagine a ladder with steps numbered from zero at the bottom to ten at the top. Suppose we say that the top of the ladder represents the best possible life for you and the bottom of the ladder represents the worst possible life for you. If the top step is 10 and the bottom step is 0, on which step of the ladder do you feel you personally stand at the present time?"

World Values Survey, Life satisfaction: Respondents were asked, "All things considered, how satisfied are you with your life as a whole these days?" Possible answers range from 1 (dissatisfied) to 10 (satisfied).

World Values Survey, Happiness: Respondents were asked, "Taking all things together, would you say you are: (1) very happy; (2) quite happy, (3) not very happy, (4) not at all happy?"

Pew Global Attitudes Survey, 2002: Respondents were shown a picture of a ladder with ten steps and asked, "Here is a ladder representing the 'ladder of life.' Let's suppose the top of the ladder represents the best possible life for you, and the bottom, the worse possible life for you. On which step of the ladder do you feel you personally stand at the present time?" Answers are scored from 1 (bottom rung) to 10 (top rung).

Table 2. Within-Country Ordered Probit Regressions of Subjective Well-Being on Income

Survey	Without controls	With controls	Instrumental variables	Sample size
Gallup World Poll, 2006	0.321***	0.318***	0.592***	102,583
Ladder question	(0.005)	(0.005)	(0.014)	(113 countries)
World Values Survey: Life satisfaction				
1981-84 wave	0.167***	0.199***	n.a.	12,198
	(0.019)	(0.022)		(10 countries)
1989-93 wave	0.130***	0.153***	0.001	32,371
	(0.011)	(0.011)	(0.041)	(26 countries)
1994-99 wave	0.225***	0.243***	0.233***	11,924
	(0.012)	(0.013)	(0.021)	(9 countries)
1999-2004 wave	0.277***	0.286***	0.305***	60,988
	(0.007)	(0.007)	(0.018)	(52 countries)
Combined, with country	0.232***	0.249***	0.258***	117,481
× wave fixed effects	(0.007)	(0.007)	(0.013)	(62 countries)
World Values Survey: Happiness				
1981-84 wave	0.324***	0.281***	n.a.	12,021
	(0.021)	(0.023)		(10 countries)
1989-93 wave	0.198***	0.188***	0.064	31,475
	(0.012)	(0.013)	(0.047)	(26 countries)
1994-99 wave	0.208***	0.209***	0.269***	13,176
	(0.013)	(0.013)	(0.022)	(10 countries)
1999-2004 wave	0.259***	0.248***	0.292***	60,627
	(0.008)	(0.008)	(0.020)	(52 countries)
Combined, with country	0.244***	0.234***	0.266***	117,299
× wave fixed effects	(0.008)	(0.008)	(0.015)	(62 countries)
Pew Global Attitudes Survey, 2002:	0.320***	0.324***	0.451***	32,463
Ladder question	(0.008)	(0.008)	(0.016)	(43 countries)

Notes: Table reports results of ordered probit regressions of the indicated measure of well-being on log household income, controlling for country fixed effects or country × wave fixed effects where noted. See the notes to table 1 for wording of survey questions.

Observations are weighted to give equal weight to each country × wave. Numbers in parentheses are robust standard errors, clustered by country. Asterisks indicate statistical significance at the *10 percent, **5 percent, and ***1 percent level.

Controls: include sex, a quartic in age, and their interaction, and indicators for missing age or sex.

Instrumental variables regression: The first stage instruments for log household income using indicator variables for levels of education, entered separately for each country, controlling for a quartic in age, interacted with gender, and country fixed effects. The second stage is an ordered probit regression of well-being on the predicted values, the residuals, the same controls, and country fixed effects.

Samples size: Samples are restricted to observations with valid household income data, from nationally representative samples (see appendix B). Instrumental variables regressions are further restricted to those countries with valid education data, which further restricts the World Values Survey samples, as valid education data were available for zero, then three, ten, and fifty-two countries in successive waves.

Table 3. Panel Regressions of Subjective Well-Being on GDP per Capita: World Values Survey

Dependent variable and specification	Micro data estimates	Macro data estimates	Sample
Life satisfaction, 1981-2004			
Levels	0.386*** (0.039)	0.414*** (0.041)	234,093 (166 country-waves)
Levels with country fixed effects	0.307*** (0.065)	0.301*** (0.091)	234,093 (166 country-waves)
Levels w/ country & wave fixed effects	0.579*** (0.088)	0.552*** (0.118)	234,093 (166 country-waves)
First differences	n.a	0.596*** (0.082)	87 differences
Long differences	0.575*** (0.116)	0.314*** (0.072)	133,900 (98 country-years = 49 differences)
Happiness, 1981-2004			
Levels	0.213*** (0.056)	0.230*** (0.064)	228,159 (165 country-waves)
Levels with country fixed effects	0.388*** (0.093)	0.363*** (0.131)	228,159 (165 country-waves)
Levels w/ country & wave fixed effects	0.263** (0.111)	0.216 (0.187)	228,159 (165 country-waves)
First differences	n.a	0.215 (0.136)	86 differences
Long differences	0.305** (0.147)	0.114 (0.103)	132,662 (98 country-waves = 49 differences)

Sources: World Values Surveys, 1981-2004.

Notes: Table reports results of regressions of the indicated measure of well-being on log real GDP per capita. Sample pools observations from all nationally representative samples in the four waves of the World Values Survey. Numbers in parentheses are robust standard errors, clustered by country. Asterisks indicate statistical significance at the *10 percent, **5 percent, and ***1 percent level.

Micro data: Ordered probit regression of subjective well-being, using data by respondent, on log real GDP per capita for the respondent's country, weighting observations to give equal weight to each country × wave. Standard errors are clustered by country-wave.

Macro data: National well-being index, using data by country-wave, is regressed on log real GDP per capita. The index is calculated in a previous ordered probit regression of well-being on country × wave fixed effects.

Life satisfaction: See table 1, for wording of survey question.

Long differences: Difference between first and last observation for each country.

Happiness: See table 1, for wording of survey question.

Table 4. Panel Regressions of Subjective Well-Being on GDP per Capita: Eurobarometer

	Micro data estimates	Macro data estimates
Dependent variable: Life satisfaction, 1973-2007		
Levels	0.737*** (0.181)	0.769*** (0.177)
Levels and country fixed effects	0.192*** (0.066)	0.194*** (0.059)
Levels and country & wave fixed effects	0.208** (0.099)	0.193** (0.094)
First differences, five year averages	n.a.	.579*** (.181)
First differences, decadal averages	n.a.	.333 (.231)
Dependent variable: Happiness, 1975-86		
Levels	0.422 (0.517)	0.448 (0.489)
Levels and country fixed effects	0.554 (0.351)	0.626* (0.346)
Levels and country & wave fixed effects	1.037 (0.993)	1.262 (0.904)
First differences, five year averages	n.a.	0.107 (.840)
First differences, decadal averages	n.a.	2.108 (1.678)

Sources: 1973-2002 data are drawn from Eurobarometer Trendfile, and 2002-07 from biannual Eurobarometer reports.

Notes: Table reports results of regressions of the indicated measure of well-being on log real GDP per capita. Sample pools observations from all Eurobarometer samples, using sample weights to typically yield around 1,000 nationally representative respondents in each country and wave (keeping East and West Germany separate). Numbers in parentheses are robust standard errors, clustered by country. Asterisks indicate statistical significance at the *10 percent, **5 percent, and *1 percent level.

Micro data: Ordered probit regression of subjective well-being, using data by respondent, on log real GDP per capita for the respondent's country, weighting observations to give equal weight to each country $\times$ wave. Standard errors are clustered by country-wave.

Macro data: National well-being index, using data by country-wave, is regressed on log real GDP per capita. The index is calculated in a previous ordered probit regression of well-being on country $\times$ wave fixed effects.

Life satisfaction: Respondents were asked, "On the whole, are you (1) very satisfied, (2) fairly satisfied, (3) not very satisfied, or (4) not at all satisfied with the life you lead?" Sample yields 850,153 respondents from 776 country $\times$ wave observations in 31 countries; 77 five-year first differences, and 37 decadal differences.

Happiness: Respondents were asked, "Taking all things together, how would you say things are these days—would you say that you're (3) very happy, (2) fairly happy, or (1) not too happy these days?" Sample yields 134,590 respondents from 139 country $\times$ wave observations in 12 countries; 19 five-year first differences, and 9 decadal first differences (all 1980s compared with 1970s).

First differences, five year averages: Data were averaged by country in five-year periods (1973-77, 1978-82, 1983-87, 1988-92, 1993-97, 1998-2002, and 2003-07), and first differences of well-being were regressed against first differences in the log of average real GDP per capita.

First differences, decadal averages: Data from the 1970s, 1980s, 1990s, and 2000s were averaged separately, and first differences of well-being were regressed against first differences in the log of average real GDP per capita.

Table 5. Subjective Well-Being in Japan: Life in Nation Survey

	Very satisfied	Fairly satisfied	Not very satisfied	Not at all satisfied	Unsure	DK/NA	Sample
Japanese	ところであなたは、お宅の暮らし向きについてどう思っていますか、あなたの気持は、この中〔回答票イ〕ではどれに近いですか。						
<i>Literal</i>	<i>By the way, how do you feel about the way your life is going at home? Which of these is your feeling close to?</i>						
Idiomatic	How do you feel about your circumstances at home? Please choose one of the following.						
	上をみればきりがな いが、大体において 今の生活に満足して いる	満足とはいえないが 、今程度の生活が続 けられれば、まあま あだと思う	今の生活ではまだまだ 不満だ	今のままの生活 ではとてもやり きれない	不明		
	<i>Although I am not innumerablely satisfied, I am generally satisfied with life now</i>	<i>Although I can't say that I am satisfied, if life continues in this way, it will be okay.</i>	<i>Somewhat dissatisfied with life now</i>	<i>Life now is very unbearable</i>	<i>Unclear</i>		
	Satisfied	Not satisfied, not dissatisfied	Somewhat dissatisfied	Extremely dissatisfied	Not sure		
Feb 1958	16%	44%	29%	9%	2%		15941
Jan 1959	17%	49%	25%	6%	3%		16897
Jan 1960	15%	45%	28%	6%	6%		17291
Jan 1961	14%	47%	29%	5%	5%		17103
Jan 1962	16%	45%	29%	5%	5%		16709
Jan 1963	18.3%	45.3%	26.1%	4.8%	5.4%		16007
Japanese	あなたはお宅の暮らしについてどう思っていますか、この中であなたの気持に一番近いものを選んで下さい。						
<i>Literal</i>	<i>How do feel about your life at home? Please choose the thing that is closest to how you feel.</i>						
Idiomatic	How do you feel about your life at home? Please choose one of the following.						
	充分/十分満足して いる	充分/十分とはいえな いが一応満足してい る	まだまだ不満だ	きわめて不満だ	不明		
	<i>Completely satisfied</i>	<i>Although I can't say I am completely satisfied, I am satisfied</i>	<i>Somewhat dissatisfied</i>	<i>Completely dissatisfied</i>	<i>Unclear</i>		
	Completely satisfied	Satisfied	Somewhat dissatisfied	Completely dissatisfied	Not sure		Sample
Jan 1964	4.4%	56.6%	33.5%	3.4%	1.9%		16698
Jan 1965	4.5%	55.7%	33.8%	4.2%	1.8%		16145
Jan 1966	4.5%	53.9%	34.4%	4.9%	2.3%		16277
Feb 1967	5.2%	55.4%	33.1%	4.2%	2.2%		16358
Jan 1968	6.2%	57.9%	29.8%	4.0%	2.0%		16619
Jan 1969	5.7%	57.8%	31.0%	4.0%	1.5%		16848

Japanese	あなたは、現在の暮らしについてどう思っていますか、この中ではどうでしょうか。					
Literal	<i>How do you feel about your life now? Which of the following?</i>					
Idiomatic	How do you feel about your life now? Please choose one of the following.					
	十分満足している	十分満足していると はいえないが、一応 満足している	まだまだ不満だ	きわめて不満だ	不明	
	<i>Completely satisfied</i>	<i>Although I can't say I am completely satisfied, I am satisfied.</i>	<i>Somewhat dissatisfied</i>	<i>Completely dissatisfied</i>	<i>Unclear</i>	
	Completely satisfied	Satisfied	Somewhat dissatisfied	Completely dissatisfied	Not sure	Sample
Jan 1970 ^(b)	6.0%	58.9%	29.4%	3.8%	2.0%	16739
Jan 1971	4.8%	52.6%	36.0%	4.8%	1.8%	16399
Jan 1972	5.4%	54.1%	34.8%	4.5%	1.2%	16985
Jan 1973	10.0%	50.5%	32.4%	5.5%	1.6%	16338
Jan 1974	3.5%	50.4%	38.0%	6.7%	1.3%	16552
Nov 1974	3.8%	46.6%	39.9%	8.0%	1.6%	8123
May 1975	5.5%	54.8%	33.6%	4.7%	1.4%	8145
Nov 1975	4.4%	53.9%	35.1%	5.2%	1.4%	8188
May 1976	5.8%	55.4%	33.2%	4.6%	1.1%	8343
Nov 1976	4.7%	55.6%	33.9%	4.5%	1.4%	8225
May 1977	9.1%	55.1%	29.7%	4.7%	1.4%	8219
May 1978	5.4%	58.9%	30.6%	3.8%	1.3%	8116
May 1979	7.1%	60.4%	28.5%	3.1%	0.9%	8239
May 1980	5.4%	57.2%	31.7%	4.5%	1.1%	8373
May 1981	5.4%	58.5%	30.5%	4.5%	1.1%	8348
May 1982	5.7%	60.1%	29.0%	4.0%	1.2%	8303
May 1983	5.8%	59.0%	30.2%	4.0%	0.9%	8106
May 1984	5.8%	59.6%	29.8%	3.9%	0.9%	8031
May 1985	7.3%	63.3%	25.0%	3.6%	0.9%	7878
May 1986	6.2%	62.0%	26.9%	4.0%	0.9%	7857
May 1987	6.0%	58.6%	30.5%	4.1%	0.9%	7981
May 1988	6.2%	58.4%	30.4%	4.1%	0.9%	7711
May 1989	5.4%	57.7%	30.8%	5.1%	1.0%	7735
May 1990	7.1%	59.7%	27.8%	4.3%	1.1%	7629
May 1991	6.7%	60.4%	28.4%	3.7%	0.8%	7639

Japanese	あなたは、全体として、現在の生活にどの程度満足していますか。この中ではどうでしょうか。						
Literal	<i>Overall, to what degree are you satisfied with your life now? Which of the following?</i>						
Idiomatic	Overall, to what degree are you satisfied with your life now? Please choose one of the following.						
	満足している	まあ満足している	やや不満だ	不満だ	どちらともいえない	わからない	
	<i>Satisfied</i>	<i>You might say I'm satisfied</i>	<i>Somewhat dissatisfied</i>	<i>Dissatisfied</i>	<i>I can't say any of the above</i>	<i>Don't know</i>	
	Satisfied	Somewhat satisfied	Somewhat dissatisfied	Dissatisfied	None of the above	Don't know	Sample
May 1992	9.3%	59.9%	21.0%	6.3%	2.9%	0.5%	7504
May 1993	10.3%	59.5%	20.9%	6.4%	2.6%	0.4%	7327
May 1994	8.3%	57.0%	23.3%	7.7%	3.3%	0.4%	7608
May 1995	10.4%	62.4%	19.8%	4.8%	2.4%	0.3%	7347
Jul 1996	10.3%	59.6%	21.6%	6.2%	2.0%	0.4%	7303
May 1997	9.8%	56.7%	22.8%	7.8%	2.5%	0.4%	7293
Dec 1999	9.5%	54.2%	23.8%	10.4%	1.8%	0.3%	7022
Sep 2001	8.1%	53.4%	26.1%	10.2%	1.8%	0.4%	7080
Jun 2002	7.9%	52.9%	26.1%	10.7%	2.1%	0.3%	7247
Jun 2003	7.2%	50.9%	28.1%	11.5%	2.0%	0.2%	7030
Jun 2004	7.2%	52.6%	26.8%	10.5%	2.4%	0.4%	7005
Jun 2005	7.7%	51.8%	27.0%	10.5%	2.6%	0.3%	6924
Oct 2006	9.4%	57.1%	25.1%	7.4%	0.9%	0.1%	5941
Jul 2007	8.3%	54.4%	26.6%	9.4%	1.0%	0.3%	6086

Source: Life in Nation surveys, 1958-2007, www8.cao.go.jp/survey/index-ko.html (in Japanese).

Notes: Literal translations emphasize the grammatical and lexical form of the source text.

Table 6. Probit Regressions of Affect on Income: World Values Survey & Gallup World Poll

Affect	Percent reporting affect	Probit coefficients	
		Between-country estimates	Within-country estimates
World Values Survey: Bradburn Affect Balance Scale “During the past few weeks, did you ever feel...?”			
Measures of positive affect			
Pleased	70.5%	0.294** (0.150)	0.140*** (0.012)
Proud	40.8%	0.493*** (0.179)	0.084*** (0.011)
Excited or interested	53.2%	0.054 (0.124)	0.164*** (0.011)
On top of the world	35.1%	0.538*** (0.186)	0.071*** (0.011)
Things going your way	49.6%	0.148 (0.155)	0.206*** (0.011)
Total positive (sum)	2.50 (s.d. 1.54)	0.571** (0.251)	0.234*** (0.011)
Measures of negative affect			
Bored	23.4%	-0.223* (0.120)	-0.129*** (0.012)
Upset/criticized	17.8%	-0.157*** (0.059)	-0.050 (0.013)
Restless	30.3%	-0.112 (0.095)	-0.092*** (0.011)
Lonely	17.0%	-0.160 (0.118)	-0.205*** (0.013)
Depressed	20.6%	-0.102 (0.117)	-0.173*** (0.012)
Total negative (sum)	1.14 (s.d. 1.30)	-0.182 (0.135)	-0.188*** (0.012)
Net affect balance (total positive-total negative)	1.36 (s.d. 1.97)	0.876*** (0.201)	0.421*** (0.017)
Gallup World Poll, 2006 “Did you experience the following feelings during a lot of the day yesterday?”			
Enjoyment	72.0	0.154*** (0.021)	0.187*** (0.007)
Physical pain	26.7	-0.125*** (0.015)	-0.139*** (0.007)
Worry	34.6	0.009 (0.020)	-0.123*** (0.007)
Sadness	21.7	-0.040*** (0.015)	-0.181*** (0.007)
Boredom	23.9	-0.036* (0.019)	-0.120*** (0.007)
Depression	14.7	-0.094*** (0.023)	-0.182*** (0.008)
Anger	19.8	-0.021 (0.017)	-0.072*** (0.007)
Love	66.3	0.050 (0.029)	0.128*** (0.007)

Gallup World Poll, 2006

“Now, please think about yesterday, from the morning until the end of the day. Think about where you were, what you were doing, who you were with, and how you felt.”

Would you like to have more days like yesterday?	66.9	0.032** (0.016)	0.120*** (0.007)
Did you feel well rested?	65.7	0.027* (0.014)	0.067*** (0.006)
Were you treated with respect?	84.6	0.146*** (0.028)	0.135*** (0.008)
Were you able to choose how you spent your time all day?	70.0	0.035* (0.018)	0.030*** (0.006)
Did you smile or laugh a lot yesterday?	70.6	0.103*** (0.017)	0.148*** (0.007)
Were you proud of something you did?	59.3	0.012 (0.023)	0.120*** (0.007)
Did you learn or do something interesting?	52.5	0.029 (0.022)	0.149*** (0.007)
Did you eat good tasting food?	74.1	0.194*** (0.021)	0.222*** (0.007)

Notes: The (binary) dependent variable in each regression is the answer (yes or no) to the survey question. All regressions control for respondent gender, a quartic in age, their interaction, and indicators for missing age or gender. Numbers in parentheses are robust standard errors. Asterisks indicate statistical significance at the *10 percent, **5 percent, and ***1 percent level.

Between-country estimates: Probit regression of affect measure on log real GDP per capita, clustering standard errors by country. Sample sizes vary by question, but the Gallup World Poll typically yielded around 134,000 respondents from 130 countries, while the World Values Survey yielded around 66,000 respondents from 31 or 32 countries with nationally representative samples.

Within-country estimates: Probit regression on log household income, further controlling for country fixed effects (and hence exploiting only within-country income comparisons). Because these regressions also require valid household income data, the sample size was smaller: typically the Gallup World Poll yielded around 100,000 respondents from 113 countries, while the World Values Survey yielded around 42,000 respondents from 24 countries.

Table A1. Alternative Scaling of Survey Responses

World Values Survey: Happiness				World Values Survey: Life satisfaction			Gallup World Poll: Life satisfaction		
<i>Verbal description</i>	<i>Simple coding</i>	<i>Standard -ized</i>	<i>Our method</i>	<i>Simple coding</i>	<i>Standard- ized</i>	<i>Our method</i>	<i>Simple coding</i>	<i>Standard -ized</i>	<i>Our method</i>
Not at all happy	1	-2.70	-2.41	1	-2.24	-2.27	0	-2.37	-2.65
Not very happy	2	-1.35	-1.32	2	-1.84	-1.73	1	-1.92	-2.06
Quite happy	3	-0.01	-0.05	3	-1.44	-1.41	2	-1.48	-1.65
Very happy	4	1.34	1.33	4	-1.04	-1.12	3	-1.03	-1.20
				5	-0.64	-0.72	4	-0.59	-0.75
				6	-0.24	-0.33	5	-0.14	-0.17
				7	0.16	0.02	6	0.30	0.39
				8	0.56	0.48	7	0.75	0.83
				9	0.96	0.98	8	1.19	1.36
				10	1.36	1.70	9	1.64	1.88
							10	2.08	2.46

Simple coding: Coding gives each category a score equal to its ordered rank.

Standardized values: Values take the simple coding, subtract its mean, and divide by the standard deviation

Our method: Our method involves running an ordered probit of well-being on country $\times$ wave fixed effects. Estimates shown are the expected value of a latent happiness index, conditional on being in each category.

Table B1. Influence of Non-Representative Samples in the World Values Survey

<i>Dependent variable:</i>		<i>Life Satisfaction</i>		<i>Happiness</i>	
Sample:		Representative samples	All observations	Representative samples	All observations
Panel 1: Comparisons between countries, coefficient on log real GDP per capita					
(Comparison of results in Table 1 with the full sample)					
1981-84 wave		0.498*	0.510**	0.569**	0.596***
		(0.252)	(0.230)	(0.230)	(0.193)
1989-93 wave		0.558***	0.210***	0.708***	0.260**
		(0.096)	(0.073)	(0.123)	(0.098)
1994-99 wave		0.462***	0.323***	0.354***	0.214***
		(0.051)	(0.069)	(0.058)	(0.069)
1999-2004 wave		0.346***	0.347***	0.126*	0.0125*
		(0.046)	(0.045)	(0.073)	(0.072)
Combined, with wave fixed effects		0.398***	0.318***	0.244***	0.181***
		(0.040)	(0.052)	(0.063)	(0.063)
Panel 2: Comparisons within countries, coefficient on log household income, controlling for country fixed effects					
(Comparison of results in Table 2 with the full sample)					
1981-84 wave		0.199***	0.199***	0.281***	0.281***
		(0.022)	(0.022)	(0.023)	(0.023)
1989-93 wave		0.153***	0.145***	0.188***	0.190***
		(0.011)	(0.010)	(0.013)	(0.011)
1994-99 wave		0.243***	0.250***	0.209***	0.217***
		(0.013)	(0.012)	(0.013)	(0.013)
1999-2004 wave		0.286***	0.274***	0.248***	0.245***
		(0.007)	(0.007)	(0.008)	(0.008)
Combined, with country × wave effects		0.249***	0.237***	0.234***	0.231***
		(0.007)	(0.007)	(0.008)	(0.008)
Panel 3: Comparisons between countries and over time, coefficient on log real GDP per capita					
(Comparison of results in Table 3 with the full sample)					
Levels		0.414***	0.339***	0.230***	0.181***
		(0.041)	(0.053)	(0.064)	(0.061)
Levels with country fixed effects		0.301***	0.151	0.363***	0.283**
		(0.091)	(0.136)	(0.131)	(0.128)
Levels w/ country & wave fixed effects		0.552***	0.253*	0.216	0.009
		(0.118)	(0.192)	(0.187)	(0.170)
Short differences		0.596***	0.407***	0.215	0.081
		(0.082)	(0.122)	(0.136)	(0.122)
Long differences		0.314***	0.172	0.114	-0.025
		(0.072)	(0.118)	(0.103)	(0.117)

Notes: Table reports results of regressions of well-being on the indicated income variable, both in the nationally-representative World Values Survey samples analyzed in the main text and in the entire sample. Numbers in parentheses are robust standard errors. Asterisks indicate statistically significant from zero at the *10 percent, **5 percent, and ***1 percent level.

Representative samples: Columns 1 and 3 include only nationally representative country-wave samples, yielding 234,093 life satisfaction respondents (and 228,159 happiness respondents) in 166 (165) country-waves, from 79 countries. Results in the middle panel also require household income data, reducing sample to 117, 481 (or 117,299) respondents, 93 (94) country-waves, and 59 countries.


All observations: Columns 2 and 4 include all country-waves, including non-nationally representative samples described in this appendix. This broader sample yields an extra 25,582 satisfaction respondents (and 26,365 happiness respondents), from 17 (18) more country-waves, adding 1, 7, 7, and 2 countries to the satisfaction (and 1, 7, 8, and 2 to the happiness) samples, respectively, in waves 1-4), yielding 13 (14) more short first differences and 6 (7) more long first differences and resulting in a total sample of 82 countries. Results in the middle panel 2 also require household income data, and so the broader sample yielded only 9,968 extra satisfaction respondents (and 9,946 extra happiness respondents), from 6 country-waves for a total of 61 countries.

Panel 1: National well-being index is regressed on log real GDP per capita. The well-being index is calculated in a previous ordered probit regression of well-being on country × wave fixed effects. Standard errors are clustered by country. See the notes to table 1 for further details.

Panel 2: Results of ordered probit regressions of respondent-level well-being on log household income, controlling for country fixed effects, or country × wave fixed effects plus a quartic in age, gender, their interaction, and indicators for missing age or sex. See notes to table 2 for details.

Panel 3: National well-being index is regressed on log real GDP per capita, pooling together observations from all four waves of the survey. National well-being index is calculated in a previous ordered probit regression of well-being on country × wave fixed effects. Standard errors are clustered by country. See the notes to table 3 for further details.

Figure 1. Early Cross-Country Surveys of Subjective Well-Being



Notes: Well-being data are aggregated into an index by running an ordered probit regression of happiness or satisfaction on country fixed effects separately for each survey. Income data were extracted from Maddison (2007) and reflect estimates of real GDP per capita at purchasing power parity in 1990 U.S. dollars. Dashed lines are fitted from OLS regressions of this well-being index on log GDP. Country abbreviations in all figures are standard ISO country codes.

Gallup 1946, Happiness: Data were extracted from Cantril (1951), who reports on polls by four Gallup affiliates. Countries included are Canada, France, the United Kingdom, and the United States. Respondents were asked, “In general, how happy would you say you are—very happy, fairly happy, or not very happy?”


Tension Study 1948, Satisfaction: Data were extracted from Buchanan and Cantril (1953), reporting on a UNESCO study of “Tensions Affecting International Understanding.” Countries included are Australia, France, and Germany. Italy, the Netherlands, Norway, Mexico, the United Kingdom, and the United States. Respondents were asked, “How satisfied are you with the way you are getting on now?—very, all right, or dissatisfied?”

Gallup 1949, Happiness: Data were drawn from Strunk (1950). Countries included are Australia, Canada, France, the Netherlands, Norway, the United Kingdom, and the United States. Respondents were asked the same question as the 1946 Gallup surveys.

Patterns of Human Concerns 1960, Satisfaction: Data were extracted from tabulations by Cantril (1965), as reported in Veenhoven (2007). Countries include Brazil, Cuba, the Dominican Republic, Egypt, Germany, India, Japan, Nigeria, Panama, Poland, United States, and Yugoslavia; data from the Philippines are missing. Data for the United States were tabulated from the Interuniversity Consortium for Political and Social Research. Surveys were run from 1957 to 1963 using Cantril’s “Self-Anchoring Striving Scale,” which probes about the best and worst possible futures, then shows a picture of a ten-step ladder and asks, “Here is a picture of a ladder. Suppose that we say the top of the ladder [pointing] represents the best possible life for you and the bottom [pointing] represents the worst possible life for you. Where on the ladder [moving finger rapidly up and down ladder] do you feel you personally stand at the *present* time?”

World Survey III 1965, Happiness: Data were extracted from Easterlin (1974, table 7), who reported cross-tabulations for France, Germany, Italy, Malaysia, the Philippines, Thailand, and the United Kingdom from the World Survey III and added data for the United States from the October 1966 AIPO poll and for Japan from the 1958 survey of Japanese national character. Respondents were asked the same question as in note b. Easterlin reports only the proportion “not very happy” for Japan; hence we infer the well-being index based only on the lower cutpoint of the ordered probit regression run on the eight other countries.


Figure 2. Life Satisfaction and Real GDP per Capita: World Values Survey



Sources: World Values Survey. Sources for GDP per capita are described in the text.

Notes: Sample includes twenty (1981-84), forty-two (1989-93), fifty-two (1994-99), or sixty-nine (1999-2004) countries; see text for details of the sample. Observations represented by hollow squares are drawn from countries in which the World Values Survey sample is not nationally representative; see appendix B for further details. Respondents are asked, "All things considered, how satisfied are you with your life as a whole these days?" and asks respondents to choose a number from 1 (completely dissatisfied) to 10 (completely satisfied). Data are aggregated into a satisfaction index by running an ordered probit regression of satisfaction on country $\times$ wave fixed effects. Dashed lines are fitted from an OLS regression; dotted lines are fitted from lowess regressions. These lines and the reported regressions are fitted only from nationally representative samples. Real GDP per capita is at purchasing power parity in constant 2000 international dollars.

Figure 3. Life Satisfaction and Real GDP per Capita: Pew Global Attitudes Survey



Sources: Pew Global Attitudes Survey, 2002. Sources for GDP per capita are described in the text.

Notes: Sample includes forty-four developed and developing countries. Respondents are shown a picture of a ladder with ten steps and asked, “Here is a ladder representing the ‘ladder of life.’ Let’s suppose the top of the ladder represents the best possible life for you; and the bottom, the worse possible life for you. On which step of the ladder do you feel you personally stand at the present time?” Data are aggregated into a satisfaction index by running an ordered probit regression of satisfaction on country fixed effects. Dashed line is fitted from an OLS regression; dotted lines are fitted from lowess regressions. Real GDP per capita is at purchasing power parity in constant 2000 international dollars.


Life satisfaction (ordered probit index)

Real GDP per capita, (thousands of dollars, log scale)

$y = -3.592 + 0.418 \cdot \ln(x)$ [se=0.026]
Correlation=0.82

Notes: Sample includes 131 developed and developing countries. Respondents are asked, “Please imagine a ladder with steps numbered from zero at the bottom to ten at the top. Suppose we say that the top of the ladder represents the best possible life for you and the bottom of the ladder represents the worst possible life for you. If the top step is 10 and the bottom step is 0, on which step of the ladder do you feel you personally stand at the present time?” Dashed line is fitted from the reported ordinary least squares regression; dotted line is fitted from a lowess estimation. GDP per capita is at purchasing power parity in constant 2000 international dollars.

Figure 5. Subjective Well-Being and Real GDP per Capita: 1999-2004 World Values Survey




Sources: World Values Survey, 1999-2004 wave. Sources for GDP per capita are described in the text.

Notes: Sample includes sixty-nine developed and developing countries. Observations represented by hollow squares are drawn from countries in which the World Values Survey sample is not nationally representative; see appendix B for further details. Dashed lines are fitted from the reported OLS regression; dotted lines are fitted from lowess regressions; both regressions are based only on nationally representative samples. GDP per capita is at purchasing power parity in constant 2000 international dollars.

Life satisfaction: Question asks, “All things considered, how satisfied are you with your life as a whole these days?” and asks respondents to choose a number from 1 (dissatisfied) to 10 (satisfied). Data are aggregated into a satisfaction index by running an ordered probit regression of satisfaction on country × wave fixed effects.

Happiness: Question asks, “Taking all things together, would you say you are: ‘very happy,’ ‘quite happy,’ ‘not very happy,’ [or] ‘not at all happy?’” Data are aggregated into a satisfaction index by running an ordered probit regression of happiness on country × wave fixed effects.

Figure 6. Subjective Well-Being and Real GDP per Capita in Selected Surveys




Notes: Well-being data are aggregated separately for each indicator in each survey by running an ordered probit regression of happiness or satisfaction on country fixed effects. Dashed lines are fitted from OLS regressions of this well-being index on log GDP. Real GDP per capita is at purchasing power parity in constant 2000 international dollars.

Kettering-Gallup Survey 1975: Data were extracted from Veenhoven (2007). Sample includes eleven developed and developing countries. Happiness question asks, “Generally speaking, how happy would you say you are: ‘very happy,’ ‘fairly happy,’ [or] ‘not too happy?’” Life satisfaction question asks, “Now taking everything about your life into account, how satisfied or dissatisfied are you with your life today?” and asks respondents to choose a number from 0 (dissatisfied) to 10 (satisfied).

First European Quality of Life Survey 2003: Sample includes twenty-eight European countries. Happiness question asks, “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.” The life satisfaction question asks, “All things considered, how satisfied or dissatisfied are you 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.”

Eurobarometer 2006: Happiness sample includes thirty European countries drawn from Eurobarometer 66.3. Happiness question asks, “Taking all things together would you say you are: ‘very happy,’ ‘quite happy,’ ‘not very happy,’ [or] ‘not at all happy?’” Life satisfaction sample includes twenty-eight European countries drawn from Eurobarometer 66.1 (missing Croatia and Turkey). The life satisfaction question asks, “On the whole, are you very satisfied, fairly satisfied, not very satisfied or not at all satisfied with the life you lead?”


Figure 8. Assessing the Functional Form of the Happiness-Family Income Gradient: General Social Survey



Source: General Social Survey (USA), 1972-2006.

Notes: Each circle aggregates income and happiness for one GSS income category in one year, and its diameter is proportional to the population of that income category in that year. The vertical axis in each panel plots the coefficients from an ordered probit regression of happiness on family income category $\times$ year fixed effects; the horizontal axis plots real family income, deflated by the CPI-U-RS.. In each panel the short- and long-dashed lines are fitted from regressions of happiness on family income and the log of family income, respectively, weighting by the number of respondents in each income category $\times$ year. Survey question asks, “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?”


Figure 9. Within-Country Comparisons of Life Satisfaction and Household Income: Gallup World Poll



Source: Gallup World Poll, 2006.

Notes: Each circle aggregates satisfaction in one income category in one country, and its diameter is proportional to the population of that income category in that country. The vertical axis plots the coefficients from an ordered probit regression of life satisfaction on indicator variables for each income category in each country, controlling for country fixed effects; the horizontal axis plots the logarithm of average real household income in each country $\times$ income category, less the country average. The dashed line is from an OLS regression, weighting by the number of respondents in each income category $\times$ country.


Figure 10. Distribution of Estimates of the Within-Country Life Satisfaction-Income Gradient



Source: Gallup World Poll, 2006.

Notes: Figure plots the distribution of regression coefficients for 113 developed and developing countries from country-specific ordered probit regressions of satisfaction on log household income, controlling for gender, a quartic in age, and their interaction.


Figure 11. Within-Country and Between-Country Estimates of the Life Satisfaction-Income Gradient: Gallup World Poll



Source: Gallup World Poll, 2006. Sources for GDP per capita are described in the text.

Notes: Each round dot plots life satisfaction against GDP per capita for one of 131 developed and developing countries. The slope of the arrow represents the satisfaction-income gradient estimated for that country from a country-specific ordered probit of satisfaction on the log of annual real household income, controlling for gender, a quartic in age, and their interaction. Usable household income data were unavailable for eighteen countries. The dashed line represents the between-country satisfaction-income gradient estimated from an OLS regression of the satisfaction index on the logarithm of real GDP per capita. GDP per capita is at purchasing power parity in constant 2000 international dollars.


Figure 12. Within-Country and Between-Country Estimates of the Life Satisfaction-Income Gradient: 1999-2004 WVS



Source: World Values Survey, 1999-2004 wave. Sources for GDP per capita are described in the text.

Notes: Each round dot plots life satisfaction against GDP per capita for one of sixty-nine developed and developing countries; hollow squares denote samples that are not nationally representative. The slope of the arrow represents the satisfaction-income gradient estimated for that country from a country-specific ordered probit regression of satisfaction on the log of household income, controlling for gender, a quartic in age, and their interaction, as well as indicator variables for missing age or gender. Usable household income data were unavailable for eighteen countries. The dashed line represents the between-country satisfaction-income gradient estimated from an OLS regression of the satisfaction index on the logarithm of real GDP per capita. GDP per capita is at purchasing power parity in constant 2000 international dollars.

Figure 13. Happiness and the Output Gap in the United States



Sources: General Social Survey, 1972-2006; Bureau of Economic Analysis.

Notes: “Output gap” is the difference between real GDP per capita and its trend, estimated using a Hodrick-Prescott filter on annual data on the logarithm of real GDP per capita, with the smoothing parameter set to 6.25. Happiness data are aggregated into a happiness index by running an ordered probit regression of happiness on year fixed effects. See figure 8 for wording of the question. See text for details of the sample.


Figure 14. Subjective Well-Being-Income Gradients across Time: World Values Survey



Source: World Values Survey, waves 1-4. Sources for GDP per capita are described in the text.

Notes: Arrows show the evolution of measured well-being and real GDP for each country. Dotted arrows join observations based on noncomparable sampling frames (see appendix B). Dashed line is fitted from an OLS regression of the well-being measure on the natural log of real GDP, estimated from pooling all four waves. See notes to figure 3 for question details and construction of aggregate well-being indices. Real GDP is at purchasing power parity in constant 2000 international dollars.


Figure 15. Change in Life Satisfaction and Economic Growth: World Values Survey



Source: World Values Surveys, waves I (1981-84), II (1989-93), III (1994-99), and IV (1999-2004). Sources for GDP per capita are described in the text.

Notes: Round dots show changes in life satisfaction and real GDP per capita between various waves of the World Values Survey; hollow squares reflect changes based on noncomparable sampling frames (see appendix B). Dashed lines show the fit from the reported OLS regression of changes in the life satisfaction index on the percent change in GDP, based only on comparable changes in life satisfaction. Graphs in the first row show nineteen, ten, and seventeen comparable short first differences, and those in the second row twenty-five, thirty-two, and thirty-three long first differences. GDP per capita is at purchasing power parity in constant 2000 international dollars.


Figure 16. Change in Life Satisfaction and Economic Growth in Europe: Eurobarometer Survey



Sources: Eurobarometer Trendfile (for 1973-2002); biannual Eurobarometer reports (for 2002-07). Sources for GDP per capita are described in the text.

Notes: Solid circles represent separate observations from each round of the Eurobarometer survey from 1973 to 1989; these were the data analyzed in Easterlin (1995); open circles extend the sample from 1990 to 2007. Each panel shows data for one of the nine countries analyzed in Easterlin (1995). Data are aggregated into a satisfaction index by running an ordered probit regression of satisfaction on country $\times$ wave fixed effects and subtracting country averages. Dashed lines are fitted from the reported OLS regression; Newey-West standard errors (se) are reported, accounting for first-order autocorrelation. The life satisfaction question asks, "On the whole, are you very satisfied, fairly satisfied, not very satisfied, or not at all satisfied with the life you lead?" GDP per capita is at purchasing power parity in constant 2000 international dollars.


Figure 17. Trends in Life Satisfaction in the European Union



Source: Eurobarometer, 1973-2007.

Notes: Lines depict alternative aggregations of semiannual time series of life satisfaction for each country, derived by running an ordered probit of satisfaction on country $\times$ wave fixed effects. “Average of a changing EU” is calculated by taking a population-weighted average of the satisfaction indices for the set of countries that were members of the European Community or the European Union at the indicated point in time; and hence the average is affected by changes in the group’s composition. “EU-9 average” is calculated by taking a fixed-weight average of the satisfaction indices of the nine members of the European Community at the beginning of the sample: Belgium, Denmark, France, West Germany, Ireland, Italy, Luxembourg, the Netherlands, and the United Kingdom; weights reflect the average population share of each country in the group. “Spliced fixed-weight series” simply sums through time first differences in the broadest available fixed-weight average of satisfaction in the member nations; consistent fixed-weight indices were calculated separately for each constellation through the sample: EU-9 (summer 1973-fall 2007), EU-10 (adding Greece, from 1981 onward), EU-12 (adding Portugal and Spain, from 1986 onward), EU-12+ (adding East German surveys, from the fall 1990 survey onward), EU-16 (adding Austria, Finland, and Sweden from 1995 onward), EU-26 (adding Cyprus, the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Slovakia, and Slovenia from the fall 2004 survey onward), and EU-27 (adding Bulgaria and Romania, from 2007 onward). Population weights for each index reflect the average population share of each country in that aggregate. “Regression-adjusted” index is the series of time fixed effects estimated by running a population-weighted OLS regression of individual country satisfaction indices on survey time fixed effects, controlling for country fixed effects


Figure 18. Life Satisfaction and GDP per Capita over Time in Japan



Source: Life in Nation surveys, 1958-2007.

Notes: The series in each of the four panels reports responses to a different life satisfaction question, and therefore comparisons should be made only within each panel. GDP per capita is at purchasing power parity in constant 2000 international dollars.

Figure 19. Economic Conditions and Life Satisfaction in Japan



Source: Life in Nation surveys, 1958-2007; BLS Foreign Labor Statistics; authors' calculations.


Notes: GDP trend growth rates are estimated by regressing log real GDP per capita on time trends. GDP growth trends are calculated for each period in which comparable satisfaction data were collected (1958-63, 1964-69, 1970-91, 1992-2007).

Raw series, with breaks: These are four separate raw series, each calculated from an ordered probit run on survey time fixed effects, for the sample periods in which the same survey question was asked (see figure 18). Because the questions changed substantially, these raw indices are only comparable within each period (these levels are normed so that each series begins at zero). Vertical lines indicate series breaks.

Adjusted for series breaks: The raw life satisfaction indices were pooled, and an OLS regression of satisfaction on series fixed effects, the unemployment rate, and log GDP per capita was run. The adjusted series subtracts the estimated of the series fixed effects from the raw series to yield a regression-adjusted continuous series.

Adjusted for series breaks and unemployment: The above series is further adjusted for cyclical influences by subtracting the product of the estimated coefficient on unemployment and the contemporaneous unemployment rate.


Figure 20. Happiness and Income over Time in the United States



Sources: General Social Survey; Current Population Survey; Bureau of Economic Analysis.

Notes: Happiness index is described in figure 13.


Figure 21. Cross-Country Measures of Recent Feelings and GDP: World Values Survey



Source: World Values Survey, 1981-84 and 1989-93 waves. Sources for GDP per capita are described in the text.

Notes: Questions are those typically grouped as the Bradburn Affect Balance Scale. Respondents were asked, "We are interested in the way people are feeling these days. During the past few weeks, did you ever feel...?" Top and bottom rows show responses to questions relating to positive and negative affect, respectively. Each observation represents one of fifty-four nationally representative country-wave samples drawn from forty developed and developing countries. Dashed lines are fitted from OLS regressions of the percent agreeing with the statement on log real GDP per capita; dotted lines are fitted from lowess estimations. GDP per capita is at purchasing power parity in constant 2000 international dollars.


Figure 22. Cross-Country Measures of Recalled Feelings and GDP: Gallup World Poll



Source: Gallup World Poll, 2006. Sources for GDP per capita are described in the text.

Notes: Respondents were asked, “Did you experience [insert feeling here] during a lot of the day yesterday?” Each observation represents one of up to 130 developed and developing countries in the sample (not asked in Iraq). Dashed lines are fitted from ordinary least squares regressions of the percent agreeing with the statement on log real GDP per capita; dotted lines are fitted from lowess estimations. GDP per capita is at purchasing power parity in constant 2000 international dollars.


Figure 23. Cross-Country Measures of Daily Experiences and GDP: Gallup World Poll



Source: Gallup World Poll, 2006. Sources for GDP per capita are described in the text.

Notes: Questions were prefaced as follows: "Now, please think about yesterday, from the morning until the end of the day. Think about where you were, what you were doing, who you were with, and how you felt." Each observation represents one of up to 130 developed and developing countries in the sample (not asked in Iraq). Dashed lines are fitted from OLS regressions of the percent agreeing with the statement on log real GDP per capita; dotted lines are fitted from lowess estimations. GDP per capita is at purchasing power parity in constant 2000 international dollars.


Figure A1. Fitting an Ordered Probit



Source: World Values Survey, 1981-2004.

Notes: Happiness question asks, “Taking all things together, would you say you are: ‘very happy,’ ‘quite happy,’ ‘not very happy,’ [or] ‘not at all happy?’” Figure shows the assumed normal distribution of a latent happiness variable when running an ordered probit regression; by a standard normalization, this distribution has a mean of zero and a standard deviation of one. Dashed lines represent the cutpoints estimated from running an ordered probit regression of happiness on country $\times$ wave fixed effects. Thus the country $\times$ wave fixed effects are equivalent to scoring “not at all happy” as -1.460, “not very happy” as -0.313, “quite happy” as 0.839, and “very happy” as 2.256.


Figure A2. Cardinal and Ordinal Measures of Well-Being



Sources: Gallup World Poll, 2006; World Values Surveys, 1981-2004.

Notes: The horizontal axis shows each ordered well-being category; the vertical axis shows the cardinal value implicitly given to these values by the ordered probit index, derived from an ordered probit regression on country $\times$ wave fixed effects (see figure A1). Dashed lines show the cardinalization derived by simply treating the ordinal rankings as cardinal, and standardizing by subtracting the mean and dividing by standard deviation. In the World Values Survey, happiness responses are coded as 1, not at all happy; 2, not very happy; 3, quite happy; 4, very happy; life satisfaction responses are coded on 1 to 10 scale, with 10 representing most satisfied. In the Gallup World Poll, respondents are shown a picture of a ladder with ten rungs, with the top rung (most satisfied) coded as 10.


Figure A3. Alternative Estimates of Life Satisfaction: Gallup World Poll, 2006



Sources: Gallup World Poll, 2006.

Notes: Each point shows the estimated average level of satisfaction in one country, comparing one of four alternative metrics on the vertical axis with the ordered probit index on the horizontal index. Dashed lines are fitted from the reported OLS regression; dotted lines are fitted from lowest estimation.


Figure A4. Alternative Estimates of Average Happiness: World Values Survey



Sources: World Values Surveys, 1981-2004.

Notes: Each point shows the estimated average level of happiness in one country, comparing one of four alternative metrics on the vertical axis with the ordered probit index on the horizontal index. Dashed lines are fitted from the reported OLS regression; dotted lines are fitted from lowess estimation.

Figure A5. Alternative Estimates of Average Life Satisfaction: World Values Survey



Sources: World Values Surveys, 1981-2004.

Notes: Each point shows the estimated average level of satisfaction in one country, comparing one of four alternative metrics on the vertical axis with the ordered probit index on the horizontal index. Dashed lines are fitted from the reported OLS regression; dotted lines are fitted from lowess estimation.