

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@wharton.upenn.edu
www.nber.org/~jwolfers

Abstract

The “Easterlin Paradox” suggests that there is no link between the level of economic development of a society and average levels of happiness. We return to Easterlin’s question: “Will raising the incomes of all increase the happiness of all?” and analyze multiple rich datasets spanning recent decades and a broader array of countries. We establish a clear positive link between GDP and average levels of subjective well-being across countries with no evidence of a satiation point beyond which wealthier countries have no further increases in subjective well-being. Moreover, we show that this relationship is consistent with the relationship between income and happiness within countries, suggesting a minimal role for relative income comparisons as drivers of happiness. Finally, we examine the relationship between changes in subjective well-being and income over time within countries, finding that economic growth has been associated with rising happiness.

This draft: May 9, 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, Danny Blanchflower, Angus Deaton, Richard Easterlin, Doug Elmendorf, Carol Graham, Danny Kahneman, Alan Krueger, Greg Mankiw, Andrew Oswald, Luis Rayo and Larry Summers 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 Zicklin Center for Business Ethics Research and the Zell/Lurie Real Estate Center for generous research support.

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, 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 Easterlin (1974, 1995, 2005a, 2005b) has examined the relationship between happiness and GDP both across countries, and within individual countries through time. In both of these exercises, he finds no significant evidence of a link between income and happiness.

In contrast, there is robust evidence that within countries those with more income are happier. These two findings—that income is an important predictor of individual happiness and yet apparently irrelevant for aggregate happiness—have spurred researchers to seek a reconciliation through evidence of reference-dependent preferences and the importance of relative-income comparisons.² Layard (2003, p. 8) 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.” While leaving room for absolute income to matter for some people, Layard and others have argued for a limit to the extent that absolute income matters with Layard (2003, p. 19) arguing that “once a country has over \$15,000 per head, its level of happiness appears to be independent of its income per head.”³

The conclusion that absolute income has little impact on happiness invites 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 (2005, p.441) argues that the types of subjective well-being data we will analyze “undermine the view that a focus on economic growth is in the best interests of

² Easterlin (1973, p.4) summarizes his findings:

“In all societies, more money for the individual typically means more 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.”

Layard (1980) 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 To see other arguments proposing a satiation point in happiness see Veenhoven (1991) and Frey and Stutzer (2002).

society.” Layard (2005) argues for an explicit government policy of maximizing subjective well-being.⁴ Moreover, he notes that relative income comparisons imply that each individual’s labor effort exerts negative externalities on others (by shifting their reference points), and that these distortions would be best corrected by higher taxes on income or consumption.

These strong policy prescriptions demand a robust understanding of what the true relationship between income and well-being is. Unfortunately the present literature is based on fragile and incomplete evidence about the relationship between income and subjective well-being. At the time the Easterlin Paradox was identified there was little data available to assess 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.

In the ensuing years there has been an accumulation of cross-national data recording individual life satisfaction and happiness. These recent data (and a re-analysis of earlier data) suggests 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 income-happiness gradient estimated within and across countries at a point in time as well as over time, rather than its statistical (in)significance.

Our key result is that the estimated income-happiness 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 undermines the possible role played by relative income comparisons.

These findings invite a sharp re-assessment of the “stylized facts” that have informed economic analysis of subjective well-being data. Across the world’s population, variation in income explains a sizeable 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 is not associated with greater well-being.

The rest of this paper is organized as follows. Section II 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-

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

being. Thus, section III compares average well-being and income across countries. While earlier studies had 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 section IV 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 there has not been much dispute over this claim. 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 income-happiness gradient measured *within* countries is similar to that measured *between* countries. In section V we extend our analysis to assessing national time series movements in average well-being and income. Subjective well-being data are both noisy, and scarce, and these factors that explain why past researchers have not found a link between economic growth and growth in happiness. Our point estimates suggest that the link is probably similar to that found in cross-country comparisons, although there remains substantial uncertainty around these estimates. We also re-examine three of the key case studies that had been raised in prior research, finding that a more careful assessment of the experiences of Japan, Europe and the United States does not undermine the claim that there is a clear link between economic growth and happiness.

II. Background: Subjective Well-Being and Income

Our focus 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 (the World Values survey asks: “Taking all things together, would you say you are: very happy; quite happy; not very happy; not at all happy”), or life satisfaction (“All things considered, how satisfied are you with your life as a whole these days?”). Other variants of the question (such as the Gallup World Poll) employ a ladder technique where interviewees are asked to imagine a ladder with each rung representing a successively better life for the individual. Respondents then report the “step” on the ladder that represents their life.

These questions (and many other variants) are typically clustered under the rubric of “subjective well-being”.⁵ While the validity of these measures remains a somewhat open question, a variety of

⁵ Diener (2005) 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,

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 and Kahneman and Krueger, 2006). A person's subjective well-being is a function both of an individual's personality and their reaction to life events. As such we would expect individual's happiness to be somewhat stable over time and that accurate measurements of subjective well-being would have high rates of test-retest correlation, which indeed they do (Eid and Diener 2004). 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 rises with marriage and income growth and falls while going through a divorce.

While each of these approaches suggests a level of validity in the cross-sectional comparisons of people within a population, 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 across cultures there is clear recognition of emotions such as anger, sadness, and joy displayed in facial expressions (Ekman & Friesen, 1971; Ekman et al., 1987). Moreover, studies have found that when people are asked about what is required for more happiness or life satisfaction there is a striking uniformity around the globe with people responding that money, health, and family are 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 times 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—one wishes to assess using subjective well-being data. The evidence regarding aggregate time series movements of happiness is inconsistent. Aggregate happiness has been shown to fall when unemployment and inflation rise (Di Tella, MacCulloch and Oswald 2003), and to move in the expected direction with business cycle volatility (Wolfers 2003). However, Stevenson and Wolfers

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.”

(2007) show that on average, women in both the U.S. and Europe report declining happiness relative to men over recent decades, a finding they argue is difficult to reconcile with changes in objective conditions. And this 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 inter-temporal comparisons is simply the difficulty in compiling sufficiently comparable data. For instance, Smith (1986) shows that small changes in question ordering on the U.S. General Social Survey led to large changes in reported happiness. These same data also seem to show important day of week and seasonal cycles, as well. As we shall also see, the attempts to cobble together long time series (such as for Japan, the U.S. 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 treating them synonymously is that these alternative measures of well-being are highly correlated, and have similar covariates. However, they capture somewhat different concepts with happiness capturing affect and satisfaction more evaluative. The psychology literature has tended to treat questions probing affect (such as “happiness”) as distinct from more evaluative assessments (“life satisfaction”). As such we will consider both the income-happiness and income-satisfaction links in parallel. There is also a subtle measurement issue involved here, as much of the data asking individuals about their happiness provides them with a shorter scale (such as “very happy”, “pretty happy”, and “not so happy”), while typical life satisfaction questions allow for a larger range of possible answers (for example such questions are more likely to be asked using 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. As a result, a curvilinear relationship was often found. In some cases, the lack of evidence of a clear linear relationship between GDP and happiness led to theories of a “satiation point”—a point beyond which more income would not increase happiness. A more natural starting point might be to represent well-being as a function of log 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 this suggests that aggregate studies should examine the relationship between average levels of subjective well-being and average levels of log income. If economic development raises individual incomes equi-proportionately, then changes in the average level of log income rise 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, which is (surprisingly enough) a departure from much of the literature.⁶ In later sections we also assess the role that changes in the income distribution play in driving a wedge between average log income and the log of average income. Throughout our analysis, we will make heavy use of bivariate scatter plots and non-parametric regression, so as to allow the eye to assess the appropriate functional forms.

Finally, it should be noted that—as with the existing literature—our analysis of the relation 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. We should add that there are several interesting variants of the question at hand—such as whether it is GDP, broader measures of economic development, or alternatively whether it is 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 Happiness

In 1974, Easterlin (p.104) asked whether “richer countries are happier countries?” Examining two international datasets he found a relationship between aggregate happiness and income across countries that he described as “ambiguous” (p. 108) and, while perhaps positive, small. Subsequent research began to show a more robust positive relationship between a country’s income and the happiness of its people, leading Easterlin (1995, p.42) to conclude that “a positive happiness-income relationship typically turns up in international comparisons.” However, this relationship has been argued as only prevailing over low levels of GDP, and once wealthy countries have satisfied basic needs, they have been described as “‘flat of the curve’ with additional income buying little if any extra happiness” (Clark et al, 2006). While the literature has largely settled on the view that aggregate happiness rises with GDP for low income countries, there is much less consensus on the magnitude of this relationship, or on whether

⁶ Previous authors examining the relationship between well-being and log GDP include Easterlin (1995), Leigh and Wolfers (2006) and Deaton (2008), however consideration of the correct functional form is typically over-looked.

there is a satiation point at which further increases in national GDP are associated with no change in aggregate happiness.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 levels, 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 of absolute income levels of GDP likely contributed to the lack of clarity around the relationship between income and happiness among wealthier countries. As we shall see, new large scale datasets involving many countries point to a clear, robust relationship between GDP and average levels of subjective well-being in a country. Furthermore, we find no evidence that countries become satiated—the positive income-happiness relation holds for both developed and developing nations.

Our macroeconomic analysis will focus on measures of real GDP per capita measured at purchasing power parity. For most countries we use the most recent data from the 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 (otherwise 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. As such, 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-by-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). 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

⁷ Deaton (2008) also 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, if anything, is stronger among high income countries.

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

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.

We present our analysis chronologically, so that the reader may see how the literature has progressed. In order to allow easy visual comparisons, we shall use a similar scale when graphing happiness and GDP, and try to keep this scale consistent throughout the paper. We have provided two useful visual devices to aid in interpretation: a dashed line showing the OLS regression line (our focus), and a shaded area that shows a central part of the happiness distribution (with a width equal to one times the cross-sectional standard deviation).

The top row of Figure 1 shows the three earliest cross-national comparisons of subjective well-being of which we are aware. Each of these comparisons is based on only four to nine countries, which are similar in terms of economic development. As such, these comparisons yield quite imprecise estimates of the link between happiness and GDP.

The second row of Figure 1 shows the cross-national comparisons presented in Easterlin (1974). 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.”⁹ Turning to the 1965 World Survey III data, Easterlin argues (p.108) that “The results are ambiguous... If there is a positive association between income and happiness, it is certainly not a strong one.” Rather than highlighting the positive association suggested by the regression line, he argues (p.106) that “[w]hat 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.” The clustering of countries within the shaded areas on the chart give a sense of this argument. However, the ordered probit index is quite useful here in quantifying the extent of differences in average levels of happiness across countries relative to the within-country variation and the ordered probit, unlike the raw data, suggests quite large differences in well-being. Similarly, the use of log income rather than absolute income highlights the linear-log relationship. Finally, Easterlin gives brief mention to the 1946 and 1949 data shown in the top row, noting (p.108) that “the results are similar... if there is a positive association among countries between income and happiness it is not clear.”

⁹ Following Cantril (1965), he 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.”

While 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, simply 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 log of GDP per capita. More importantly, the point estimates reveal a positive relationship between well-being and income, and a precision-weighted average of these five regressions coefficients yields a coefficient of 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 further explore in section IV).

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, run in 1967, and the twelve nation Gallup-Kettering Survey, run in 1975, both yield further evidence consistent with an important GDP-well-being gradient. Subsequent cross-national data collections have become increasingly ambitious, and these have both made the case for a linear-log relationship between subjective well-being and GDP per capita even stronger, and largely confirmed that the magnitudes suggested by these early studies were quite accurate.

In Figure 2 we show each wave of the World Values Survey separately, so as to show the accumulation of new data through time. We begin by highlighting the data on life satisfaction (“All things considered, how satisfied are you with your life as a whole these days?”), and turn to data on happiness below (Figure 5). In the early waves of the World Values Survey the sample was comprised mostly of wealthy countries, and with limited variation, this yielded suggestive, but not definitive, evidence of a link between GDP and life satisfaction. As the sample expanded it incorporated more poor countries; however, these early expansions were explicitly not representative of the entire population in these countries.¹⁰ 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 1989-1993 wave, but the samples largely consisted of the more educated members of society and those living in urban areas (and thus the wealthiest) members of the poorest countries. The sampling frame for these countries became representative in the 1999-2004 wave, with the exception of Mexico which became representative in the 1994-99 wave.¹¹ Despite these limitations, as additional countries were added to the WVS the

¹⁰ We thank Angus Deaton for alerting us to these limitations in the World Values Survey.

¹¹ In Chile the sample covers the central portion of the country, which contains 63% of the total population and has an income level that is about 40% higher than the national average. In China and Nigeria the urban and educated

relationship became clearer. In each wave, we see an upward sloping regression line that is statistically significant and the estimated regression coefficient is similar across the four waves with precision increasing in the later waves. We also plot estimates from locally-weighted (or lowess) regressions, in order to get a sense of whether there are important deviations from the linear-log functional form.¹²

In the earliest waves the small number of countries and limited heterogeneity in income across the countries in the samples made it difficult to make robust inferences about the relationship between life satisfaction and levels of development. Nonetheless, pooling data from all four waves yields an estimate of the satisfaction-income gradient of 0.34 (se=0.06, clustering by country), and an F-test reveals that wave-specific intercepts and slopes are jointly statistically insignificant relative to a model with a common intercept and slope term ($F_{6,174}=1.72$).

Subsequently, the 2002 Pew Global Attitudes Survey was fielded across around 38,000 respondents in 44 countries across the development spectrum. The subjective well-being question differs from that used in the World Values Survey, utilizing a form of Cantril's (1965) "Self-Anchoring Striving Scale," asking respondents: "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? [0-10 steps]." As before, we use an ordered probit to estimate average levels of subjective well-being, by country, and compare these averages with the log of per capita GDP in Figure 3. These data show a tightly estimated linear relationship between life satisfaction and the log of per capita GDP.

The most ambitious cross-national 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 is 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 per capita GDP. 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, 0.42, is similar to that obtained using the World Values Survey, and comparable to that obtained from the Pew survey, as well as the

population was oversampled. In India, 90% of the sample was limited to the literate population (which comprised less than 50% of the population). In Mexico only those in urban areas were surveyed. There was non-random sampling in developed countries as well. For example, in many developed countries the youth were over-sampled.

¹² The locally-weighted scatter plot smoothing (LOWESS) estimator is a local regression estimator which places a flexible curve with no parametric restrictions on the form of the curve.

earlier surveys assessed by Easterlin. These findings are also quite similar to those of Deaton (2008), who also finds that well-being in the Gallup World Poll has a linear-log relationship with GDP per capita, and emphasizes that the clearer evidence in the Gallup data reflects the inclusion of surveys from a greater number of poor countries.¹³

In the data that we have explored thus far, the measure of subjective well-being has largely consisted of “life satisfaction” assessments. As discussed previously, the economics literature has tended to treat measures of happiness and life satisfaction as largely interchangeable, while the psychology literature distinguishes between the two, with happiness considered an affective measure and satisfaction thought of as more evaluative. We now turn to considering to assessing the relationship between measures of happiness and income and comparing these estimates with that seen for measures of life satisfaction.¹⁴ In Figure 5 we consider the happiness-GDP link and the life-satisfaction-GDP link estimated using the latest wave of the World Values Survey—the aforementioned life satisfaction data and a question probing happiness (“Taking all things together, would you say you are—Very happy; Quite happy; Not very happy; Not at all happy”). These data suggest that these measures may not be as synonymous as previously thought, as it appears that happiness is somewhat less correlated with GDP than is life satisfaction.¹⁵ While 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, Tanzania and Nigeria, reported both the two highest levels of happiness, yet both have much lower life satisfaction—indeed Tanzania reported the lowest average satisfaction of any country.¹⁶

This apparent noise in the happiness-GDP link partially explains why earlier analyses of subjective well-being data have yielded mixed results. We have also re-run both the happiness and life satisfaction regressions, removing Tanzania and Nigeria, and it turns out that these outliers explain at least part of the puzzle. That is, in the absence of these two countries, the well-being-GDP gradient, measured using either life satisfaction or happiness turns out to be very similar. Equally, in these data,

¹³ 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.

¹⁴ We consider additional measures of subjective well-being and their relationship to income in Section VI.

¹⁵ The contrast in Figure 5 probably overstates this divergence as it plots the data for the 1999-2004 wave of the WVS, while Table 1 shows that earlier waves yielded a clearer happiness-GDP link.

¹⁶ While one might be tempted to suspect that sampling problems are to blame, it is worth noting that Nigeria also reported the 11th highest happiness rating in the 1994-99 wave of the WVS, although it was around the mean in the 1989-93 wave. Unfortunately we are not aware of any other happiness data for Tanzania, but it is worth noting that in the 2002 Pew survey, Tanzania registered the second lowest level of average satisfaction among 44 countries (See Figure 3).

the correlation between happiness and GDP remains smaller than the correlation between satisfaction and GDP.

In order 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 on 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 run in 2003. In addition, the bottom panel of Figure 6 shows data from the 2006 Eurobarometer which asked about happiness in 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.

In Table 1 we formalize all of the analysis discussed thus far with a series of regressions of subjective well-being on log income, 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 is reported along with its standard error. In the first column, we estimate ordered probit regressions of individual well-being against the natural log of real GDP per capita, clustering our standard errors 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 which we aggregate the data to the country level in the first stage by running an ordered probit regression of subjective well-being on country fixed effects. In the second stage, we estimate an ordinary least squares regression of the country fixed effects on log GDP per capita and it is this coefficient we report. In all the data sets examined estimates of the relationship obtained from the respondent level analysis is 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, of around 0.2 to 0.4.

The final two columns in Table 1 allow us to assess whether the well-being-GDP gradient differs for rich and poor countries. Veenhoven (1991) argues that income is particularly important for happiness when the basic needs of food, clothing, and shelter are not being met and 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 will no longer raise the happiness of society. For instance, Layard (2005b, p.149) 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...” Frey and Stutzer (2002) provide 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.”

We employ Layard’s cutoff, ranking countries with annual GDP per capita above \$15,000 (in year 2000 \$US) as “rich”. Strikingly, we find that the relationship between subjective well-being and log GDP is stronger, rather than weaker in the rich-country sample. Indeed, across the Gallup, World Values and Pew datasets, the coefficient on log income tends to be nearly three times larger for rich countries than for poor countries, a finding consistent with Deaton (2008). As such, we find no evidence of a satiation point. Indeed, a consistent theme across multiple datasets shown in Figure 1 and Figure 6 appears to be that there is a clear positive relationship between subjective well-being and GDP, even when making comparisons only between advanced economies.

The fact that the coefficient on log GDP appears to be larger for rich countries should be interpreted carefully. Taken at face value, these results suggest that a 1% rise in GDP would have about three times as large of an effect on measured well-being in rich than in poor nations.¹⁷ Of course a 1% rise in U.S. GDP is about ten times larger than a 1% rise in Jamaican GDP (Jamaican per capita GDP is about one-tenth that in the U.S.). Consider instead the effect of a \$100 rise in average incomes in Jamaica and the United States. This shock would raise log(GDP) per capita by ten times more in Jamaica than 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, per capita GDP in Burundi is about one-sixtieth that in the United States, and hence a \$100 rise in average incomes would have a twenty-fold larger impact on measured well-being in Burundi than the United States.¹⁸

One explanation for the difference in our findings with 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 level of per capita GDP. However, Figure 7 shows both

¹⁷ Greater 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” which translates true hedonic experience into our measured well-being index (Oswald, forthcoming).

¹⁸ 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 the very poor ($GDP < \$3,000$) we obtain very similar estimates as when restricting the sample to the less-developed ($\$3,000 < GDP < \$15,000$). This is also evident in the non-parametric fit shown in Figure 4.

approaches, and the log specification yields a better fit, although the difference is small.¹⁹ 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 as the independent variable, finding the well-being-GDP gradient is about two times 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 (while a 1% rise in GDP 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, or its absolute level. As Figure 7 demonstrates, we have data on too few countries to draw particularly strong inferences about the appropriate functional form. In the next section we turn to within-country comparisons, and given the much larger samples involved, it shall 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 to Across Countries Estimates

A very simple benchmark for assessing the magnitude of the between-country well-being-GDP gradient measured in the previous section (typically about 0.2-0.4) would be the within-country well-being-income gradient. In particular, Easterlin (1974, p.106-107) argued that “the happiness difference 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.” 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 (2001, p.468): “As far as I am aware, in every representative national survey ever done a significant bivariate relationship between happiness and income has been found.” And indeed, we have made similar comparisons in over one hundred countries, and are also yet to find a (statistically significant) exception. While there has been somewhat more debate about the magnitude of such an effect, income is clearly an important correlate with happiness. For example, Frank (2005) argues the importance of income for

¹⁹ Deaton’s (2008, p.58) assessment of the functional form for the bivariate well-being-GDP relationship led him to conclude that “the relationship between the log of income and life satisfaction offers a reasonable fit for all countries, high-income and low-income, and if there is any evidence for deviation, it is small and in the direction of the slope being higher among the richer countries.”

happiness: “When we plot average happiness versus income for clusters of people in a given country at a given time, we see that rich people are in fact much 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.”

In this spirit, we examine the relationship between happiness and income in the United States from 1972 through 2006 using the General Social Survey. In Figure 8 we plot the coefficients from an ordered probit regression of happiness on income category by year fixed effects. Each circle in this figure represents an income category in a particular year with the diameter proportional to the population of the income category that year.²⁰ The statistical significance of this relationship is not in doubt, largely because each round of the GSS (as with most happiness surveys) involves over a thousand 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.²¹ We have checked this relationship in other datasets for other countries, confirming this finding. Finally, it is clear from this plot that income is indeed, a powerful force shaping the distribution of happiness.

It is the juxtaposition of these *statistically* significant cross-sectional findings with statistically insignificant cross-country or time series results that yielded the Easterlin Paradox. Whereas theories emphasizing relative income comparisons would suggest that the between-country GDP-happiness gradient would be smaller than the within-country income-happiness gradient (if relative income comparisons are made intra-nationally), the suggestive comparison of the gradients in Figure 7 and Figure 8 yields the opposite conclusion.

It is important to see whether these findings about the magnitude of the within-country gradient simply reflect the peculiarities of the United States. As such, we turn to the estimating the within-country well-being-income gradient for each country using the Gallup World Poll data. That is for each country

²⁰ As this section turns to examining individual income data, it is worth noting that 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 run these regressions separately for each country-wave of each dataset. We thank Angus Deaton for this suggestion.

²¹ Because the GSS retained the nominal income categories used in 1973, there are some very low income cells that are somewhat off the regression line, reflecting both the fact that small cells yield imprecise happiness estimates, and the difficulties imputing appropriate incomes to the bottom-coded group.

with usable income data²² we estimate an ordered probit regression of life satisfaction on the natural log of household income. Rather than listing 113 separate coefficients, Figure 10 lists these coefficients (rounded to the nearest 0.05) in a histogram summarizing the entire sample. The average well-being-income gradient was 0.37, and 90% of the estimates lie between 0.09 and 0.70. In turn, much of this heterogeneity likely reflects simple sampling variation: the average country-specific standard error was 0.06, with 90% of the country-specific regressions yielding standard errors between 0.04 and 0.11. We have also re-run these regressions controlling for gender, and a quartic in age, entered separately for men and women, obtaining very similar results.

We provide an alternative representation of these data in Figure 11 which provides a direct comparison of within-country and between-country estimates of the well-being-income gradient. Thus, each dot shows the GDP and average well-being of each country (and hence the cloud of dots suggests the between-country well-being-GDP gradient), while the slope of the arrows corresponds to the slope of the well-being-income gradient estimated within that country. Not only are the slopes of these arrows remarkably similar across countries, they are also typically quite close to the between-country well-being-GDP slope. Figure 12 repeats this exercise using data from the 1999-2004 round of the World Values Survey. The household income data in the WVS are not as uniform as the Gallup World Poll leading to the omission of several countries. However, for the countries for which there is sufficient data, a similar pattern emerges to that seen in the Gallup World Poll. Repeating the same exercise for the Pew data also yields similar findings.

In Table 2, we pool 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 run an ordered probit of subjective well-being on log family income, controlling for country (or country*wave) fixed effects, which serve to control for not only the between-country variation in GDP, but also variation in measured income due to changes in exchange rates or purchasing power, or other survey-specific changes. The first column shows the results from a simple ordered probit of well-being on log household income, while in the second column we add controls for gender, a quartic in age, and the interaction of these variables. Comparing the results in these columns with the corresponding between-country estimates in Table 1, we see roughly comparable magnitudes, although the between-country estimates are slightly larger than the within-country estimates in most cases. Relative income comparisons would lead between-country

²² We dropped Kenya because it lacked labels for income groups, Laos because it contained clearly implausible income groupings, and Uzbekistan because the income categories listed in the data involved overlapping ranges. Respondent-level income data were unavailable for Egypt, Iran, Iraq, Jordan, Kuwait, Latvia, Lebanon, Morocco, Pakistan, Palestine, Philippines, Saudi Arabia, Sri Lanka, Turkey, UEA and Yemen. This leaves us with valid household income data for 113 countries.

estimates to be smaller than within-country comparisons. Similarly, a focusing illusion as suggested by Kahneman et al (2008) suggests that relationships between income and well-being within countries are overstated and should therefore be larger than the relationship seen between countries.

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 there should be little change in subjective well-being with transitory income changes, while permanent shocks should have a much larger impact. The variation in GDP 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.

A very simple back-of-the-envelope calculation can help provide an upper bound on the extent to which these issues are distorting the comparisons in Table 1 and Table 2. If all cross-country variation in GDP is 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 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, is comprised of a roughly \$.55 change in permanent income. In this case, the estimates in Table 2 need to be adjusted upward by around 80% ($=1/0.55$) to be interpreted as the relationship between well-being and permanent income or consumption. If instead of assuming perfect smoothing, we account for Campbell and Mankiw's (1990) estimate that 50% 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%. This latter 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 last column of Table 2 instrument for income using education, entered separately for each country.²³ While we are confident that these instruments isolate variation in permanent rather than transitory income, we do not believe that the exclusion restriction holds—that education does not have an effect on well-being beyond that mediated by income. (For instance, Lleras-Muney (2005) shows

²³ We follow Rivers and Vuong (1988) in their approach to estimating an IV ordered probit. Thus, the first stage involves a regression of log household income on indicator variables for each level of education, country fixed effects (and their interaction). The second stage involves an ordered probit regression of well-being on the predicted values and residual from the second stage, as well as country fixed effects.

positive impacts of compulsory schooling on health). Given that these omitted effects are likely positive, our IV estimates likely overstate the within-country income-happiness gradient.

The discussion above has been premised on the straightforward view that transitory income shocks yield smaller impacts on well-being than 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 illustrates, showing that business-cycle variation in the output gap yields quite large effects on subjective well-being. Indeed, the estimated happiness-transitory income gradient suggested by these shocks is about five times larger than GDP-happiness 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.

While 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 income-well-being gradient may be biased upward by reverse causation, as happiness may well be a productive trait in some occupations, raising income. A different perspective from Kahneman et al (2006) suggests that within-country comparisons overstate the true relationship between subjective well-being and income due to 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 regards to concrete measures such as income. While these specific biases may not affect the between-country comparisons (although Kenny, 1999 argues for reverse causation), it seems likely that the bivariate well-being-GDP relationship may 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 and well-being. Other factors, such as increased savings, reduced leisure, or even increasingly materialist values may raise GDP at the expense of subjective well-being. At this stage, we are not going to be able to 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 brethren. This then leads to our final question: do societies get happier through time as they become richer? Easterlin (1995) 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.” Indeed, the core of the Easterlin Paradox lies in his failure to isolate statistically significant relationships between average levels of happiness through time, and economic growth. Easterlin (1974, 1995) contains 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 can be supported by existing datasets. 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 GDP within, rather than between countries. Second, we re-analyze these data, finding that happiness has in fact risen in Europe, and a re-analysis of the Japanese data also invites a re-examination of earlier conclusions. The failure of happiness to rise in the United States remains a puzzling outlier, although the extent to which it provides a sharp exception should not be overstated. Third, as more data have become available, both extended national time series, and adding observations from new countries, evidence that happiness rises with GDP has started to accumulate.

Indeed, the World Values Survey has now been running since 1981, and across its four waves, we now have repeated observations on a large number of countries, spread across several decades. However, as noted earlier, the sampling frame changed considerably for many of these countries across the waves. Figure 14 shows the evolution of both life satisfaction and real GDP per capita, for all of those countries in which this survey offers repeated observations (including those countries whose sampling frame changed). 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*wave fixed effects. Arrows link each individual country’s evolution in well-being–GDP space through time, and so the slope of these arrows corresponds to the well-being-income gradient derived from each of two consecutive observations in a country’s national time series.

Several points should be evident from this chart. First, there appears to be a general tendency for economic growth to be accompanied by growth in subjective well-being (arrows tend to point to the north-west), and economic decline, which is most visible in the former eastern bloc, has been accompanied by a decline in well-being (arrows pointing south-east). Of the 101 changes shown in Figure 14, in 67 cases happiness and GDP change in the same direction (58 show growth in both; 9 show declines), compared with 34 cases where they move in opposite directions (of which 25 reflect economic growth unaccompanied by growth in happiness, and 9 reflect growing happiness despite economic

decline). The life satisfaction data yield much weaker results, with satisfaction and GDP moving in the same direction in only 51 of 101 cases.

Second, averaging 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). Third, there still remains substantial heterogeneity in these estimated responses, and there are many important exceptions. For instance, while the right panel shows sharp declines in life satisfaction in China and India, which are not evident in the happiness numbers. Both of these countries surveyed only the urban, educated population in the early years, with the sampling frame shifting in the last wave to become representative of the broader population. Naturally we cannot be sure whether this is due to true heterogeneity, or measurement error in our well-being indices (a problem that is likely exacerbated by first-differencing the data), although at least in the case of China, Kahneman and Krueger (2006) provide supporting evidence of declining well-being.

Finally, these time series changes are strongly influenced by the result of common patterns across countries: Rising GDP in most nations, accompanied by a trend rise in happiness, and lower life satisfaction in the two most recent waves of the WVS. We suspect that the trend in life satisfaction has been distorted by changes in question ordering. In particular, in the 1994-1999 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 change how respondents subsequently report their life satisfaction. To check this, we assess the (raw) correlation between life satisfaction and financial satisfaction for the nine countries which have been represented in each round of the WVS; this correlation was 0.51 and 0.52 in the most recent two waves, well above previous levels (0.46 in the first wave and 0.43 in the second wave). The happiness question was never proximate to financial satisfaction, and the correlation of happiness with financial satisfaction was quite stable across each of the waves (it was recorded as 0.30, 0.30, 0.31 and 0.29 from earliest to latest wave).

We formalize our analysis in Table 3, in which we analyze the World Values Survey as a country-wave panel dataset, analyzing both life satisfaction and happiness. In the first column, we report the results of ordered probit regressions of well-being on log GDP, and as we move down the rows, we add controls for country and then wave fixed effects. The second column aggregates the data to the country*wave level, running OLS regressions of our well-being index on log GDP. The first row shows the simple bivariate well-being-GDP relationship, and hence pools both within-country and between

country variation. In order to isolate the within-country time series variation, the second row includes country fixed-effects in these regressions. Reflecting the limited variation in levels of economic development within countries, these estimates are much less precisely estimated than those derived from the cross-country variation. Nonetheless, consistent with Figure 14, the well-being-GDP gradient estimated from this time series variation is similar to that estimate from the point-in-time between-country comparisons, albeit smaller in the case of life satisfaction, and larger for happiness.

Adding further controls for each wave of the World Values Survey partials out the changes in well-being that reflect differences in surveys across waves, and as might be expected in light of the previous discussion of question order effects, this increases the estimate time-series life satisfaction-GDP gradient to around 0.25. With the happiness question, these controls account for the aggregate rise in happiness that occurred in partnership with rising GDP in most countries, and the remaining variation appears largely uncorrelated with per capita GDP. The imprecision of these estimates is noteworthy though, and just as we cannot falsify a null hypothesis that the happiness-GDP gradient is zero, we cannot falsify that it is 0.25. In subsequent rows, we take first differences of consecutive country-wave observations, and also long-differences, including only the first and last observation for each country. Running a line through these numbers it appears that our best estimate of the time series well-being-GDP gradient is probably around 0.2.

It is more important to emphasize how fragile these inferences are. While the large cross-country datasets allow for useful comparisons between populations in abject poverty with industrialized powerhouses, the within-country 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, while the standard deviation of between-wave first differences in log GDP (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 small number of countries with quite unusual economic trajectories, such as the rapid growth in the Korea, Ireland, and China and the decline of former eastern bloc countries (see Figure 15). Even so, most of our approaches to these data yield suggestive evidence falsifying the view that the time series well-being-GDP gradient is zero. By contrast, these data largely fail to falsify the alternative null that this gradient falls in the 0.2-0.4 range obtained from our between-country or within-country analyses.

Europe

We turn next to the other major set of repeated cross-section data, the Eurobarometer Survey, drawing our data from the Mannheim Eurobarometer Trendfile, which collects available microdata from

1970-2002. We have supplemented these files with data from 2002-2007, extracted from print editions of the Eurobarometer Reports series. These surveys initially asked the population of nine European countries about their life satisfaction at least annually from 1973 onward (except 1974 and 1996). Subsequently, the survey expanded to cover 18 countries by 2002, and presently includes 31 countries, yielding a broad, but unbalanced panel. A happiness question was also briefly asked (from 1975-1986, except 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 34 year sample of well-being among West Germans.

We begin by analyzing the evolution of life satisfaction and GDP for the nine countries which constituted the original 1973 sample. Easterlin (1995, p.38) also analyzed these 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.” In a subsequent update, Easterlin (2005) maintains that “I think the evidence continues to support my generalization in the 1995 study.”

In Figure 16 we update this analysis, adding a further 18 years of data (shown with hollow circles). In eight of these nine countries rising GDP has been associated with rising life satisfaction, and in six of these cases, the correlation is 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 satisfaction trend in Belgium, and declining life satisfaction in Ireland during the first phase of the “Irish miracle”, quickly followed by rising satisfaction during its second phase. (Satisfaction appeared to be anomalously high in the very first Irish survey; dropping this observation yields a statistically significant coefficient on log GDP 0.14 with a standard error of 0.05.) Our point is not to count up the number of statistically significant responses one way or the other, but rather to suggest that across these nine large European countries, life satisfaction has typically risen with GDP. Moreover, estimates of the satisfaction-GDP gradient based on these national time series, while quite variable, are centered around 0.2, with some being larger, and some smaller.

Figure 17 provides an alternative way of assessing European satisfaction trends, using data from all countries in the Eurobarometer survey. The simplest way to construct a time series for “average” satisfaction would be to report average levels of satisfaction across countries at each date. However, the Eurobarometer panel is extremely unbalanced, as this data collection (and indeed, the European Union) has gradually expanded to include a rising number of poorer (and less happy) nations. It is important that our satisfaction aggregate not be affected by these compositional changes, and as such, we construct our

series by running an ordered probit on time (survey-round) fixed effects, controlling for country fixed effects. These time fixed effects are shown in the figure, and they clearly suggest a mildly rising trend in life satisfaction. Note that it would be difficult to infer that such a trend either did or did not exist on the basis only of Easterlin's 1973-89 sample. The magnitude of the European trend in life satisfaction, 0.005 units per year, can be assessed relative to the trend in GDP per capita, which rose by around 2½ percent per year. Considered jointly, these trends point to a satisfaction-GDP gradient of about 0.2, which both falsifies the null effect of no positive relationship, and is roughly consistent with the magnitudes seen in our within- and between-country assessments. Figure 17 also shows a similarly-estimated happiness series, which proves to be both much more volatile, and to have a somewhat sharper upward trend. Again, the average response of happiness to economic growth appears roughly consistent (or at least not inconsistent) with our findings in previous sections.

In order to further examine these patterns, we formalize our findings with a series of panel regressions in Table 4, exploiting all of the observations across all countries, and analyzing life satisfaction and happiness separately in the upper and lower panels. Our initial regression includes no fixed effects, and hence the coefficient reflects both comparisons between countries at different levels of development, and changes in GDP through time. The large coefficient reflects the important influence of between-country variation, and the estimate is consistent with our earlier finding of a particularly steep well-being–income slope when comparing rich countries. In order to focus on the time series movements, we control for country fixed effects, and as suggested above, find a coefficient of about 0.2 for life satisfaction, and a larger, although less precise estimate for happiness. Finally, in order to focus on comparisons between those countries whose per capita GDP grew faster with those which grew slower, we also account for (usually biannual) wave fixed effects in our final estimates and once again, this comparison yields estimates of the well-being–GDP gradient consistent with previous sections.

We suspect that the key to reconciling our findings with earlier reports suggesting no link between movements through time in GDP and life satisfaction 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.

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 post-war period (real GDP rose 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.²⁴

Upon closer inspection, these Japanese data are neither as persuasive as many thought, nor is the trend flat. We have returned to the original codebooks, and had the questions re-translated.²⁵ This was rather revealing, suggesting several very important series breaks. Accounting for these series breaks yields a very different perspective. We provide a full accounting in Table 5, and show both literal and idiomatic translations of the survey questions as they have changed.

There are three important findings in this table. First, in 1964, the response categories changed dramatically. Specifically, 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”, leading the proportion reporting their well-being in this highest category to decline from 18.3% to 4.4%. So too, the second top category becomes somewhat 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”, although the proportions in this lowest category did not much change. Second, questions asked from 1958-69 focused on feelings about “life at home”, rather than global life satisfaction (which was the focus of the relevant question from 1970 onward). Third, the survey question—and the allowable responses—change 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, and a slower rise from 1970-91. In turn, this roughly parallels the path of Japanese GDP through these periods. From 1992-2007, life satisfaction has fallen, although this coincides with the end of the Japanese growth

²⁴ For instance, Easterlin (1995, p.39-40) 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 (2005), Frank (2005), Kahneman, Krueger, Schkade, Schwarz and Stone, 2006, among dozens of others.

²⁵ We would like to thank Michael L. Woodford for his patient assistance with these translations.

miracle, and indeed, the onset of an economic slump. All told, this suggests that subjective well-being in Japan has largely risen with GDP, and that it rose most sharply during the period of rapid growth.

Having established that qualitatively these data appear consistent with a positive GDP-satisfaction 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 sub-series, we create a time series of average well-being by an ordered probit of subjective well-being on survey fixed effects.²⁶ By construction, the levels of these series are not comparable, and hence comparisons within, but not between series are valid. Figure 19 shows the satisfaction-GDP gradient within each of these periods, and it is clear that throughout the period in which Japan moved from poor to affluent (shown in the first three panels), subjective well-being rose with per capita GDP. The right-most panel shows that since 1992 the Japanese economy has shown very little growth, and subjective well-being has fallen sharply.

Figure 20 shows 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 changes in the well-being question asked: spectacular growth during the period 1958-69 spanning one series break in the well-being question; slower growth from 1970-91; and then anemic growth from 1992 onward, which coincided with the emergence of large-scale unemployment. The dotted markers in the bottom panel of the figure show the corresponding (and non-comparable) movements in subjective well-being within each of the periods for which consistent data exist.

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

$$\begin{aligned}
 \text{Well-being}_t = & -1.65 - 0.39 * I(1964 \leq \text{year} \leq 1969) - 0.57 * I(1970 \leq \text{year} \leq 1991) - 0.53 * I(1992 \leq \text{year}) \\
 & (0.49) \quad (0.07) \quad (0.11) \quad (0.14) \\
 & -0.063 * \text{Unemployment rate}_t + 0.24 * \log(\text{GDP per capita}_t) \quad n=51 \\
 & (0.02) \quad (0.06)
 \end{aligned}$$

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 these adjustments yields the gray line, shown in the bottom panel of Figure 20. This time series suggests that subjective well-being did in fact grow strongly in Japan—at least

²⁶ While Table 5 shows the proportions coded as “not sure”, “don’t know” or “none of the above”, we simply drop these observations from the rest of the analysis.

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 happiness through the 1990s, and the reversal over the past few years as unemployment has started to decline. The unemployment coefficient is roughly comparable—although somewhat larger than estimates for other OECD countries in Wolfers (2003). We can also use this coefficient to back out a “cyclically-adjusted” well-being series for Japan, shown in the lower-panel of Figure 20. As should be clear, this bears a strong relationship with per capita GDP, and indeed, the estimated coefficient, 0.24, is again roughly consistent with our other time series findings.

Finally, it is worth noting that other data also suggest that well-being in Japan has tracked its rising level of economic development. For instance, from 1974-91 the same survey also asks “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 1992-2007, 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, and in 1981 16% of Japanese respondents reported being very happy, rising to 18% in 1990, then 34% in 1995, before falling to 29% 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 are shown in Figure 1: In each of the comparisons shown (the 1960 “Patterns of Human Concerns” surveys; the 1965 World Survey; or the 1975 Kettering survey, shown in Figure 6) we see that 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) now show that Japan’s well-being is at a level consistent with its modern status as an affluent country.

United States

The most widely used dataset for analyzing happiness in the United States is the General Social Survey, a nationally representative sample of about 1,500 respondents each year from 1972-1993 (except 1992), continuing with around 3,000 respondents every second year from 1994 through to 2004, rising to 4,500 respondents in 2006. These repeated cross-sections 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?” We report the coefficients from an ordered probit regression of these responses on year fixed effects in Figure

18.²⁷ These data suggest a very mildly declining happiness trend through this period (slope = -0.0011, with a standard error of 0.0009), and this estimated trend suggests that our happiness index declined by about 0.037 points between 1972 and 2006 (with a 95% confidence interval around this decline ranging from -0.10 to +0.02 points).²⁸

The second panel of Figure 18 shows that real GDP per capita rose by 0.67 log points over the same period, and the juxtaposition of this income growth with roughly flat happiness trend appears to provide useful support for the Easterlin Paradox. Indeed, an happiness-income gradient of 0.2-0.4 would have led one to expect the happiness index to have grown by 0.13-0.26 points. Translating this to the individual happiness categories, U.S. GDP growth from 1972-2006 was large enough to suggest that by the end of the sample another 8% of the population should be “very happy”, and the proportions “not too happy” and “fairly happy” should each be about 4% lower.

The U.S. time series pattern is suggestive of the Easterlin paradox—aggregate income rose, yet the population did not become happier. Moreover, there is clear evidence of the absence of a happiness-income relationship here—the data clearly reject growth in happiness predicted by the happiness-income gradient estimated within the U.S. or across countries. While the U.S. time series is a data point supporting the Easterlin paradox, it is an interesting exception and warrants further scrutiny.

To better understand trends in happiness within the U.S. and its relationship to recent income growth it is necessary to look closer at the patterns of income growth. In particular, the fruits of this economic growth through this period were quite unequally distributed. Indeed, from 1972-2005, CPS data suggest average real family income grew by only 15-20% in each of the three bottom quintiles; the fourth quintile experienced growth of 30% and only the top quintile realized income growth of 59% (DeNavas-Walt, Proctor and Lee 2006). In turn, the top two quintiles of the family income distribution experienced mild growth in happiness, while happiness actually declined for the bottom three quintiles.

²⁷ We have corrected these data for the biases due to changes in question ordering noted by Smith (1979 and 1990). Specifically, he noted that happiness among married respondents tends to be higher when preceded by a question about marital happiness (as was the norm, except in 1972), and when preceded by a five-item satisfaction scale (as was the norm except 1972 and 1985). In order to estimate the extent of these biases, we regress happiness on a dummy variable equal to one for those affected by each sampling change or a subsequent split ballot experiment (the first including married people in 1972, and married form 3 respondents in 1980 and 1987; the second including all 1972 and 1985 respondents, 1986 form 2 respondents and 1987 form 2 and 3 respondents), controlling for year fixed effects, entered separately for both married and unmarried respondents. Thus the ballot experiments identify the effect of changing questionnaire order, separate from background trends in happiness by marital status. The aggregate happiness time series is then estimated as simply the unadjusted annual happiness aggregates, less the estimated question order effects (for those subject to the varied question order).

²⁸ Stevenson and Wolfers (2008) show that this aggregate mild decline in happiness over time represents a combination of little change in male happiness and a decrease in female happiness.

The family income data recorded in the GSS suggest roughly similar real income growth—an average increase of about 32% over the full sample that was quite unequally distributed—with real declines reported in the bottom quintile. While 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 micro income-happiness link 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 the average level of log income and happiness aggregates. However, the sharp rise in inequality over recent decades puts 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 have computed the rise in income inequality in both the CPS and GSS samples. From 1972-2006, the CPS measure the log of average real household income rose by 41 log points, while inequality—measured by the mean log deviation—rose by 19 log points.²⁹ Together, these numbers imply that the average level 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, while the measured rise in inequality (again measured as the mean log deviation) is 15 log points.

Thus, within the GSS sample, the average level of the log of family income has risen by only around 17 log points since 1972 (an annual rate of growth of only around 0.5% per year).³⁰ Based on a happiness-income gradient of around 0.2-0.4, it seems reasonable to expect happiness in the United States to have been basically flat over the past 35 years. Thus, by re-focusing our attention on the appropriate macroeconomic aggregate (in the bottom panel of Figure 18), it can be seen that the U.S. experience is still roughly consistent with the evidence of a robust happiness-income link.

²⁹ It is worth being a bit more explicit about how reasonably robust economic growth translates into weaker growth in the average log income. From 1972-2006, real GDP per capita grew by 93%, or 66 log points, and disposable personal income per capita rose by a similar amount. Beyond these aggregate BEA data, the Census Bureau also calculates income per household from the March CPS. These alternative data suggest that personal income per capita (in 2005 dollars) rose by 65.8%, or that the log of personal income rose by 51 log points. Much of the gap between the BEA and CPS measures reflects differences in deflators. (From 1972-2006 the CPI-U-RS (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) were we to deflate 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.

³⁰ We have deflated the GSS income data using the CPI-U-RS, rather than the CPI-U-X. If instead we used the official deflator, we would register barely any growth at all in the average log of family income.

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 relative to a “ladder” with the best and worst possible lives bounding the scale. Yet this still leaves a lot unsaid about the subjectively-experienced lives of the rich and the poor, and major advances in recent cross-national 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 (1969) Affect Balance Scale which were included in the first two waves of the World Values Survey. The scale is intended to separately assess both positive and negative affect and does this by probing direct reports of whether various pleasurable and unpleasurable feelings were experienced recently. This battery of questions asked “During the past few weeks, did you ever feel...” each of five positive experiences (“particularly excited or interested in something”; “proud because someone had complimented you on something you had done”; “pleased about having accomplished something”; “on top of the world / feeling that life is wonderful”) and each of five negative experiences (“so restless you couldn’t sit long in a chair”; “very lonely or remote from other people”; “bored”; “depressed or very unhappy”; or “upset because somebody criticized you”).³¹

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. In order to separately isolate the between-country and within-country variation, we run one regression where the independent variable is log GDP per capita, and then another analyzing log household income, conditioning on country fixed effects. In order 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 contrasts of both rich and poor people in a country (*within-country*) and of rich and poor countries (*between-country*), measures of positive affect are typically positively associated with income, while measures of negative affect are negatively associated with income. While not all of the 10 measures of affect are individually statistically significantly related to income, the summary measures of negative affect (the sum of the five individual negative questions), and positive affect (the sum of the five positive questions) are both related to income, and our within- and

³¹ Bradburn (1969, chapter four) 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.

between-country contrasts yield roughly similar magnitudes. Putting these together into a measure of net affect (the average number of positive experiences less negative experiences) yields a measure that is strongly related to both log household income, and log GDP, and these estimates reflect the impact of income on positive and negative affect in roughly equal measure.³²

Figure 21 presents the cross-country comparisons graphically, with the top row revealing 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”). The bottom row of the figure reveals that a smaller proportion of the population in richer countries typically report negative experiences. Interestingly, as with self-assessed “happiness” Nigeria is an outlier for all of the measures of positive affect, with Nigerians reporting a much higher likelihood of experience positive feelings compared to other low-income countries. The bottom panel shows the relationship between each of the 5 measures of negative feelings and GDP. For all of the measures the negative affect-GDP gradient is negative with a higher proportion of people in poor countries experiencing negative feelings. (Note, that these measures of negative affect suggest that Nigerians have a more typical experience for their GDP level.)

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?” including enjoyment, physical pain, worry, sadness, boredom, depression, anger, and love. The middle panel of Table 6 shows that, among the positive emotions, the enjoyment-income gradient is positive and similar for both the between- and 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 anger all fall with rises in income, at both the national and individual level. Figure 22 shows how each of these feelings vary across countries.

The final regressions analyze the relationship between income and some more specific experiences in people’s lives, such as feeling respected, smiling, doing interesting activities, feeling proud, and learning. Most of these assessments are related to one’s income in the within-country estimates. Fewer show signs of a similar sized effect when examining the relationship between countries. Indeed, Figure 23 shows that there is tremendous variation across countries in how people report on their daily experiences, and these are not well explained by GDP levels. However, there are some notable exceptions. Wealthier people are more likely to say that felt that they were treated with respect yesterday

³² The summary measure of net affect is computed by adding up the positive and negative measures, with 0 reflecting an equal number of positive and negative experiences.

and people in wealthier countries are more likely to feel respected. People in wealthier countries also report smiling more. 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 higher levels of life satisfaction. Finally, people in wealthier countries are more likely to report having been able to eat good tasting food the previous day and the magnitude of the relationship is similar to that seen within countries.

All told, these alternative measures of well-being point 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 rich new cross-national 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

Our task in this paper has been to revisit—and revise—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 frame 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 analyze separately the relationship between income and happiness that one obtains from contrasting rich and poor members of a society, with that obtained from contrasting rich and poor countries, with that obtained from observing the path of average happiness as the average income of a country changes. Our measurement framework allows us to assess the extent to which these relationships may differ.

Thus, our key contribution is to note that the relationship between subjective well-being and income within a country (that is, contrasting the happiness of rich and poor members of a society) is similar to that seen across countries (contrasting rich and poor countries), which in turn is similar to the time series relationship (comparing the happiness of a country as it gets richer or poorer). In multiple datasets from several decades, and covering various populations, we consistently estimate a well-being–income gradient of around 0.2 to 0.4. We estimate slightly steeper gradients *between* countries, although reading across datasets and taking account of sampling error, we can neither reject that the gradients are the same within and between countries nor can we reject small differences between the two. Our

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.

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. The United States stands out as a notable exception, with virtually no increase in happiness over the past 35 years (and indeed, a decrease in happiness among U.S. women). 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; moreover, 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 international panel data, 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 similar role for income as that seen in our within- and between-country contrasts. Moreover, taken as a whole, the time series evidence is difficult to reconcile with earlier claims that economic growth yields no boost to happiness.

Thus, we conclude that the accumulation of new data (and our re-analysis of earlier data) has not been kind to the Easterlin Paradox. Easterlin and others have argued that comparisons of rich and poor people 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 (2008) 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 of income matter to well-being”. In its strong form, this hypothesis suggests that people (and public policy) are powerless to deliver lasting gains in happiness, as individual happiness returns inexorably to one’s set point of happiness. Our findings clearly falsify this strong form of adaptation—those enjoying materially better circumstances also enjoy greater 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 lie in the range of 0.2-0.4 and do not have enough 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 had suggested a *prima facie* case for relative income playing a dominant role, our updated re-analysis yields no such case. Equally, if 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 due to 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 0.35 (with the time series is a bit weaker still). This is consistent with both absolute and relative income impacting well-being, with the former having a weight about three times larger than the latter. Thus, our findings should not be interpreted as a falsification 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, as in Luttmer (2005).

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.

VIII. References

- Atkinson, A.B., L. Rainwater and T.M. Smeeding. *Income Distribution in OECD Countries*. OECD Social Policy Studies, Paris: OECD, 1995.
- 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, NJ: 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. *Guidelines for National Indicators of Subjective Well-Being and Ill-Being*. mimeo, University of Illinois, 2005.
- 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.
- 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 Veenhoven." *Social Indicators Research* 74, no. 3 (2005): 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.
- 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.
- Hout, Michael. "Getting the Most Out of the GSS Income Measures." GSS Methodological Report, 2004.
- 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.
- Krueger, Alan. "Are We Having More Fun Yet? Categorizing and Evaluating Changes in Time Allocation." *Brookings Papers on Economic Activity*, forthcoming, 2007.
- Layard, Richard. *Happiness: Lessons from a New Science*. London: Penguin, 2005.
- 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, 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. *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.
- 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, March 25, 2007.
- 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*year (or country*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 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” is given a value of 2; “quite happy” is given a value of 3, and “very happy” takes 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 existing 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 asked for three, four, five, seven, ten or eleven-point scales (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, the metric would have an economic interpretation, scaling differences in well-being relative by its cross-sectional standard deviation. (As we shall see, this approach yields results very close to our approach.)

Even so, the limitation of this approach is it imposes a linear structure, implying 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”. While psychologists often seem willing to accept that the subjective distances between successive points on category scales are similar, data on the proportions of the population who report themselves as being in each category allow us to relax (or test) this assumption.

In order to make use of these population proportions, the ordered probit 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 a standard deviation of one. The country or country*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 cutpoints. van Praag and Ferrer-i-Carbonell (2004) describe this as “probit-adapted OLS”. The table below provides the mapping between the underlying categorical responses, standardized categorical responses, and our scaling derived from these ordered probits,.

	WVS: Happiness			WVS: Satisfaction			Gallup World Poll: Satisfaction		
	<i>Simple</i>	<i>Std.</i>	<i>S-W</i>	<i>Simple</i>	<i>Std.</i>	<i>S-W</i>	<i>Simple</i>	<i>Std.</i>	<i>S-W</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	-0.55	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
Mean	3.01	0	6.60	5.32	0	0.00
SD	0.74	1	3.01	2.25	1	1.14

Notes: Table shows three alternative cardinalizations:

1. “Simple” is the rank of each ordinal category
2. “Std.” is “Simple”, standardized (subtract the mean, divide by standard deviation)
3. “S-W” is the expected value of the latent well-being variable, conditional on being in that response category, assuming that the latent well-being index (conditional on country*wave fixed effects) is $\sim N(0,1)$. Note that the standard deviation of this variable is not exactly one because: a) discretizing the variable into its expected value within each category reduces its variance; and b) adding back in the country*wave fixed effects add variance.

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 is quite useful, because it provides a useful approximation: to map results in other studies to ours, 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 below. In particular, the figure 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 comparing the indices assesses this approach relative to four alternative metrics, of which three are typically used in the literature, and the fourth is an interesting extension of our approach.

1. 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 WVS life satisfaction question which asks for a response on a scale of 1-10).
2. Population proportions: An alternative involves reporting, say, the proportion of the population reporting themselves as “quite happy” or “very happy”. This approach has the advantage that it yields a natural scaling (from 0-1), 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. In order 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.
3. Ordered logits: The ordered logit is similar to our ordered probit approach, although it 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 the ordered probit.
4. Heteroscedastic ordered probit: The ordered probit imposes an equal variance in residual happiness, while 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 cutpoints for each country and year, allowing proportional shifts in these cutpoints, by country-year.

The following three charts compare these alternative aggregators to our ordered probit approach, analyzing separately the the satisfaction ladder from the Gallup World Poll, and life satisfaction and happiness data, by country and wave, in the World Values Survey.

Alternative Estimates of Life Satisfaction

Gallup World Poll, 2006




Life Satisfaction: Ordered Probit Index

Dashed line shows linear regression; dotted line shows lowest fit.

Alternative Estimates of Average Happiness

World Values Survey, Waves 1-4




Happiness: Ordered Probit Index

Dashed line shows linear regression; dotted line shows lowest fit.

Alternative Estimates of Average Life Satisfaction

World Values Survey, Waves 1-4



Life Satisfaction: Ordered Probit Index

Dashed line shows linear regression; dotted line shows lowess fit.

Table 1: Cross-Country Relationship between Subjective Well-Being and GDP

Each cell reports coefficient from a regression: $Well\text{-}being = \beta \log(GDP\text{ per capita})$						
	Micro Data: Ordered Probit		Macro Data: OLS regression			Sample Respondents (clusters)
	Dep. Var: Indiv. well-being		Dep. Var: National Well-being			
	No controls	Controls	All countries	Only GDP>\$15k	Only GDP<\$15k	
Gallup World Poll <i>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?</i>						
2006	0.396*** (0.023)	0.422*** (0.023)	0.418*** (0.026)	1.076*** (0.194)	0.348*** (0.040)	139,051 (131 countries)
World Values Survey: Life Sat <i>All things considered, how satisfied are you with your life as a whole these days? [1] Dissatisfied – [10] Satisfied</i>						
1981-1984 wave	0.538** (0.241)	0.363 (0.301)	0.508** (0.230)	1.674** (0.701)	0.691 (0.523)	24,617 (20 countries)
1989-1993 wave	0.200*** (0.073)	0.200*** (0.073)	0.208*** (0.073)	0.504 (0.467)	-0.030 (0.070)	61,867 (42 countries)
1994-1999 wave	0.290*** (0.060)	0.305*** (0.059)	0.320*** (0.068)	0.324 (0.421)	0.195* (0.105)	75,911 (52 countries)
1999-2004 wave	0.320*** (0.040)	0.327*** (0.040)	0.347*** (0.045)	0.451* (0.223)	0.221** (0.088)	99,095 (69 countries)
Combined (with wave FE)	0.296*** (0.048)	0.302*** (0.047)	0.316*** (0.052)	0.474** (0.198)	0.165** (0.081)	261,490 (82 countries)
World Values Survey: Happiness <i>Taking all things together, would you say you are: [4] Very happy; [3] Quite happy; [2] Not very happy; [1] Not at all happy.</i>						
1981-1984 wave	0.676*** (0.210)	0.596*** (0.218)	0.594*** (0.193)	1.658 (0.987)	0.577 (0.644)	23,380 (19 countries)
1989-1993 wave	0.238** (0.096)	0.255*** (0.095)	0.259** (0.098)	0.328 (0.475)	-0.129** (0.059)	60,533 (42 countries)
1994-1999 wave	0.195*** (0.062)	0.217*** (0.060)	0.213*** (0.069)	0.246 (0.235)	0.038 (0.093)	74,716 (54 countries)
1999-2004 wave	0.117* (0.061)	0.137* (0.060)	0.124* (0.072)	0.767*** (0.217)	-0.132 (0.112)	97,654 (68 countries)
Combined (with wave FE)	0.169*** (0.056)	0.188*** (0.055)	0.181*** (0.063)	0.612*** (0.170)	-0.068 (0.083)	256,283 (82 countries)
Pew Global Attitudes Survey <i>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? [0-10 steps].</i>						
2002	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: ***, ** and * denote statistically significant at 1%, 5% and 10%, respectively.

(Robust standard errors in parentheses, clustered by country.)

- Micro data: Ordered probit regression of subjective well-being measure against log(GDP). Controls include a quartic in age, interacted with gender, an indicators for missing age or sex values.
- Macro data: OLS regression of national well-being index against log(GDP). The well-being index is calculated in a previous ordered probit regression of well-being on country*wave fixed effects.

Table 2: Within-Country Variation in Subjective Well-Being and Income

Each cell reports coefficient on log Household Income from a regression:
Ordered Probit: Well-being = β log (Household Income)
+ Country*wave fixed effects

	No controls	Age*gender	IV: country-specific education
Gallup World Poll			
	0.321*** (0.005)	0.318*** (0.005)	0.592*** (0.014)
World Values Survey: Life Satisfaction			
1981-1984 wave	0.167*** (0.019)	0.199*** (0.022)	n.a.
1989-1993 wave	0.125*** (0.009)	0.144*** (0.010)	0.057 (0.038)
1994-1999 wave	0.236*** (0.012)	0.251*** (0.012)	0.240*** (0.020)
1999-2004 wave	0.266*** (0.007)	0.275*** (0.007)	0.296*** (0.017)
Combined (w/ cty*wave FE)	0.211*** (0.005)	0.226*** (0.005)	0.257*** (0.013)
World Values Survey: Happiness			
1981-1984 wave	0.324*** (0.021)	0.281*** (0.023)	n.a.
1989-1993 wave	0.200*** (0.010)	0.191*** (0.011)	0.086** (0.042)
1994-1999 wave	0.217*** (0.013)	0.217*** (0.013)	0.264*** (0.022)
1999-2004 wave	0.255*** (0.007)	0.244*** (0.008)	0.298*** (0.020)
Combined (w/ cty*wave FE)	0.222*** (0.005)	0.237*** (0.005)	0.270*** (0.014)
Pew Global Attitudes Survey, 2003	0.320*** (0.008)	0.324*** (0.008)	0.451*** (0.016)

Notes: ***, ** and * denote statistically significant at 1%, 5% and 10%, respectively.

(Robust standard errors in parentheses, clustered by country.)

Column 1: An ordered probit regression of well-being on log household income, and country fixed-effects

Column 2: Adds gender, a quartic in age, and their interaction as controls

Column 3: Instruments for log household income using indicator variables for levels of education, entered separately for each country. Second stage is an ordered probit regression of well-being on the predicted values, the residuals, and country fixed-effects.

Table 3: Estimates of the Effect of Economic Growth on Well-being
World Values Survey

Each cell reports coefficient from a regression: $Well-being_{ct} = \beta \log (GDP \text{ per capita}_{ct})$

	Microdata Estimates (Respondent observations)	Macro Estimates (Country-wave observations)	Sample
Dep Var: <i>All things considered, how satisfied are you with your life as a whole these days?</i> Life Satisfaction <i>[1] Dissatisfied – [10] Satisfied</i>			
Levels	0.315*** (0.032)	0.338*** (0.031)	261,490 (183 country-waves)
Levels Country FE	0.153** (0.078)	0.150 (0.105)	261,490 (183 country-waves)
Levels Country and Wave FE	0.266** (0.109)	0.250* (0.146)	261,490 (182 country-waves)
First Differences	n.a.	0.404*** (0.133)	101 diffs
Long Differences	0.171 (0.114)	0.180 (0.118)	150,952 (110 country-years = 55 diffs)
Happiness <i>Taking all things together, would you say you are:</i> <i>[4] Very happy; [3] Quite happy; [2] Not very happy; [1] Not at all happy.</i>			
Levels	0.166*** (0.040)	0.180*** (0.035)	256,283 (183 country-waves)
Levels Country FE	0.294*** (0.075)	0.283*** (0.101)	256,283 (183 country-waves)
Levels Country and Wave FE	0.036 (0.101)	0.009 (0.133)	256,283 (183 country-waves)
First Differences	n.a.	0.081 (0.120)	101 diffs
Long Differences	-0.014 (0.109)	-0.024 (0.117)	152,452 (112 country-waves = 56 diffs)

Notes: ***, ** and * denote statistically significant at 1%, 5% and 10%, respectively.
(Robust standard errors in parentheses, clustered by country.)

Sample: All four waves of the World Values Survey, 1981-84, 1989-93, 1994-99 and 1999-2004.

Macro estimates: The dependent variable is a well-being index, estimated by an ordered probit of well-being on country*wave fixed effects. The table reports estimated coefficients on the independent variable, ln(GDP per capita).

Micro estimates: Ordered probit regression in which the dependent variable is individual well-being, regressed against ln(GDP per capita) for that country, clustering standard errors by country-wave.

**Table 4: Estimates of the Effect of Economic Growth on Well-being
Eurobarometer Survey**

Each cell reports coefficient from a regression: <i>Well-being_{c,t} = β log (GDP per capita_{c,t})</i>		
	Estimated Well-being-GDP slope	Sample
Dep Var: Life Satisfaction (1973-2007)	<i>On the whole, are you [4] very satisfied, [3] fairly satisfied, [2] not very satisfied, or [1] not at all satisfied with the life you lead?</i>	
Levels	0.758*** (0.177)	759 country-waves 32 countries
Levels Country fixed effects	0.190*** (0.060)	759 country-waves 32 countries
Levels Country & wave fixed effects	0.202** (0.095)	759 country-waves 32 countries
Dep Var: Happiness (1975-1986)	<i>Taking all things together, how would you say things are these days – would you that you're [3] Very happy; [2] Fairly happy; or [1] Not too happy these days?</i>	
Levels	0.445 (0.489)	139 country-waves 12 countries
Levels Country fixed effects	0.625* (0.346)	139 country-waves 12 countries
Levels Country & wave fixed effects	1.263 (0.905)	139 country-waves 12 countries

Notes: ***, ** and * denote statistically significant at 1%, 5% and 10%, respectively.

(Robust standard errors in parentheses, clustered by country.)

Sample period: 1973-2007

OLS regression of national well-being index against log(GDP). The well-being index is calculated in a previous ordered probit regression of well-being on country*wave fixed effects.

Table 5: Happiness in Japan – Life in Nation Survey

	Very satisfied	Fairly satisfied	Not very satisfied	Not at all satisfied	Unsure	DK/NA	Sample
Japanese <i>Literal</i> Idiomatic	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>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> How do you feel about your circumstances at home? Please choose one of the following.						
	上みればきりがなが、大体 において今の生活に満足してい る <i>Although I am not innumerablely satisfied, I am generally satisfied with life now</i> Satisfied	満足とはいえないが、今後の 生活が続ければ、まあまあ だと思う <i>Although I can't say that I am satisfied, if life continues in this way, it will be okay.</i> Not satisfied, not dissatisfied	今の生活ではまだまだ 不満だ <i>Somewhat dissatisfied with life now</i> Somewhat dissatisfied	今のままの生活ではとても やりきれない <i>Life now is very unbearable</i> Extremely dissatisfied	不明 <i>Unclear</i> 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 <i>Literal</i> Idiomatic	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ちに一番近いものを選んで下さい。 <i>How do feel about your life at home? Please choose the thing that is closest to how you feel.</i> 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 <i>Completely satisfied</i> Completely satisfied	充分十分とはいえないが、一応満 足している <i>Although I can't say I am completely satisfied, I am satisfied</i> Satisfied	まだまだ不満だ <i>Somewhat dissatisfied</i> Somewhat dissatisfied	きわめて不満だ <i>Completely dissatisfied</i> Completely dissatisfied	不明 <i>Unclear</i> Not sure		
Jan 1964 ^(a)	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 Idiomatic	How do you feel about your life now? Which of the following? How do you feel about your life now? Please choose one of the following.					
	十分満足している Completely satisfied Completely satisfied	十分満足しているとはいえないが、一応満足している Although I can't say I am completely satisfied, I am satisfied. Satisfied	まだまだ不満だ Somewhat dissatisfied Somewhat dissatisfied	きわめて不満だ Completely dissatisfied Completely dissatisfied	不明 Unclear Not sure	
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、全として、現在の生活どの程度満足していますか。この中でどうでしょうか。					
Literal	Overall, to what degree are you satisfied with your life now? Which of the following?					
Idiomatic	Overall, to what degree are you satisfied with your life now? Please choose one of the following.					
	満足している Satisfied	まあ満足している You might say I'm satisfied	やや不満だ Somewhat dissatisfied	不満だ Dissatisfied	どちらともいえない I can't say any of the above	わからない Don't know

	Satisfied	Somewhat satisfied	Somewhat dissatisfied	Dissatisfied	None of the above	Don't know	
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

Notes: We provide questions and alternative responses in three forms: the original Kanji (from the survey instruments or documentation), a literal translation (in italics), emphasizing the grammatical and lexical form of the source text, and idiomatic, emphasizing the natural English translation.

Note two very small deviations in the second part of the question wording:

(a) January 1964: The question is slightly different than the rest of 1965-69.

Kanji: 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ちに一番近いものを1つだけ選んで下さい。

Literal: How do you feel about your life at home? Please choose one answer that is closest to how you feel.

Yielding an identical idiomatic translation: How do you feel about your life at home? Please choose one of the following.

(b) January 1970: The question is slightly different than the rest of 1971-91.

あなたは、現在の暮らしについてどう思っていますか.....この中であなたの気持ちに一番近いものを選んで下さい。

How do feel about your life now? Please choose the answer that is closest to how you feel.

Yielding an identical idiomatic translation: How do you feel about your life now? Please choose one of the following.

Source: “Life in nation” surveys, 1958-2007, available (in Japanese) at: <http://www8.cao.go.jp/survey/index-ko.html>.

Table 6: Relationship Between Income and Affect

	Mean (Dependent Var).	Between- Country Estimates	Within-Country Estimates
World Values Survey: Bradburn Affect Balance Scale <i>During the past few weeks, did you ever feel...</i>			
Net Affect Balance	1.38 [sd=1.99]	0.324*** (0.111)	0.457*** (0.014)
Positive affect	2.49 [sd=1.53]	0.177 (0.114)	0.286*** (0.010)
Pleased	69.4%	0.177** (0.077)	0.177** (0.010)
Proud	41.2%	0.135* (0.077)	0.112*** (0.010)
Excited or interested	54.3%	-0.038 (0.054)	0.223*** (0.010)
On top of the world	35.4%	0.120 (0.087)	0.084*** (0.010)
Things going your way	49.3%	0.071 (0.065)	0.209*** (0.009)
Negative affect	1.13 [sd=1.30]	-0.116* (0.058)	-0.171*** (0.010)
Bored	24.5%	-0.134*** (0.044)	-0.118*** (0.010)
Upset / Criticized	18.1%	-0.072* (0.038)	-0.014 (0.011)
Restless	30.5%	-0.030 (0.047)	-0.044*** (0.010)
Lonely	18.1%	-0.102* (0.054)	-0.222*** (0.011)
Depressed	21.5%	-0.101* (0.053)	-0.183*** (0.011)
Gallup World Poll, 2006 <i>Did you experience the following feelings during a lot of the day yesterday?</i>			
Enjoyment	72.30%	0.128*** (0.021)	0.198*** (0.007)
Physical Pain	27.3%	-0.081*** (0.014)	-0.178*** (0.007)
Worry	35.0%	0.023 (0.020)	-0.128*** (0.006)
Sadness	21.9%	-0.019 (0.016)	-0.204*** (0.007)
Boredom	23.1%	-0.047*** (0.019)	-0.114*** (0.007)
Depression	14.6%	-0.075*** (0.023)	-0.195*** (0.008)
Anger	19.0%	-0.031* (0.017)	-0.056*** (0.007)
Love	66.2%	0.037 (0.029)	0.143*** (0.007)

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?	66.3%	0.028* (0.015)	0.115*** (0.007)
Did you feel well rested?	64.9%	0.023 (0.014)	0.055*** (0.006)
Were you treated with respect?	84.1%	0.154*** (0.029)	0.122*** (0.008)
Were you able to choose how you spent your time all day?	69.4%	0.045** (0.018)	0.005 (0.007)
Did you smile or laugh a lot yesterday?	70.2%	0.070*** (0.017)	0.161*** (0.007)
Were you proud of something you did?	59.8%	0.001 (0.023)	0.132*** (0.007)
Did you learn or do something interesting?	52.4%	-0.001 (0.022)	0.171*** (0.006)
Did you eat good tasting food?	73.4%	0.185*** (0.021)	0.220*** (0.007)

Notes: ***, ** and * denote statistically significant at 1%, 5% and 10%, respectively.

(Robust standard errors in parentheses, clustered by level of aggregation of income variable.)


Mean: First column shows the mean of each variable (%responding yes)

Between-country estimates: Second column shows the coefficient from a probit regression of the binary variable dependent variable on log(GDP), clustering standard errors by country-wave. The Gallup World Poll yielded 130 valid country observations; World Values Survey yielded 56 country-wave observations from 40 countries).

Within-country estimates: Third column shows the coefficient from a probit regression of the binary dependent variable on log household income, controlling for country fixed effects (and hence exploiting only within-country income comparisons). Sample size for the Gallup World Poll is 134,954; WVS yields a sample size of 77,038-78,498..

Exceptions: We report OLS coefficients for “Positive affect” which refers to the sum of the five positive measures in the World Values Survey; “Negative affect” refers to the sum of the five negative measures. “Net affect balance”=Positive Affect – Negative Affect.

Figure 1: Early Cross-National Surveys of Well-Being



Sources: Income data in all panels comes from Maddison (2007), and we match income data to the year the happiness survey commenced in each country.

1946 data were extracted from (H. Cantril 1951), reporting on polls by four Gallup affiliates (AIPO, BIPO, CIPO and FIPO), asking: “In general, how happy would you say you are—very happy, fairly happy, or not very happy?”.


1948 data were extracted from (Buchanan and Cantril 1953) reporting on a UNESCO study of “Tensions Affecting International Understanding”. This 9-nation study (AUS, GBR, FRA, DEU, ITA, NLD, NOR and MEX) asked: “How satisfied are you with the way you are getting on now?—Very, All right, or Dissatisfied.”

1949 data were drawn from Strunk (1950) and asked 7 Gallup affiliates (AUS, GBR, NLD, CAN, NOR, USA and FRA) “In general, how happy would you say you are—very happy, fairly happy, or not very happy?”

1960 data were extracted from tabulations of (H. Cantril 1965) reported in (Veenhoven, World Database of Happiness, Trend in Nations 2007), missing data from Philippines. Data for USA tabulated from ICPSR #7023. These surveys were run from 1957-63, and Cantril’s “Self-Anchoring Striving Scale” begins by probing about the best and worst possible futures, then 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? Step number [0-10]”

1965 data were extracted from Easterlin (1974, Table 7), who reported crosstabs for 7 countries from the USIA-funded World Survey III (GBR, DEU, THA, PHL, MYS, FRA and ITA), and added data for the USA from the October 1966 AIPO poll, and for Japan from the 1958 survey of Japanese national character. Easterlin reports only the proportion “not very happy” for Japan, and hence we infer the well-being index based only on the lower cutpoint of the ordered probit regression run on all 8 other countries. These surveys asked “In general, how happy would you say you are? Very happy, fairly happy, not very happy”.


Figure 2: Life Satisfaction and Real GDP per Capita



Source: World Values Survey

Notes: This survey asks: “All things considered, how satisfied are you with your life as a whole these days?” asking respondents to choose a number from 1 “Dissatisfied” to 10 “Satisfied”.


Figure 3: Pew Global Attitudes Survey: Satisfaction and GDP, 2002



Source: Pew Global Attitudes Survey

Notes: Survey asks: “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? [0-10 steps].”


Figure 4: Gallup World Poll: Satisfaction and GDP, 2006



Source: Gallup World Poll, 2006

Notes: Survey asks: “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?”


Figure 5: Comparing Income with Life Satisfaction and Happiness



Source: World Values Survey, 1999-2004 Wave.

Notes: Life satisfaction question asks: “All things considered, how satisfied are you with your life as a whole these days?” asking respondents to choose a number from 1 ‘Dissatisfied’ to 10 ‘Satisfied’. Happiness question asks: “Taking all things together, would you say you are: ‘Very happy’; ‘Quite happy’; ‘Not very happy’; ‘Not at all happy’.”

Figure 6: Happiness v. GDP and Life Satisfaction v. GDP




Source: Eurobarometer: Happiness data drawn from Eurobarometer 66.3; Life satisfaction from Eurobarometer 66.1.

“Euro_QoL” refers to the First European Quality of Life Survey, 2003.


“Kettering” refers to the Kettering-Gallup survey conducted in 1975. Data were extracted from Veenhoven (2007).

Figure 7: Assessing the Functional Form of the GDP-Satisfaction Gradient



Source: Gallup World Poll, 2006.

Figure 8: Cross-Sectional Relationship between Income and Happiness in the U.S.



Notes: Data is from the General Social Survey (1972-2006). The figure plots the coefficients from an ordered probit regression of happiness on income category by year fixed effects where income is real family income, deflated by the CPI-U-RS. Each circle represents an income category in a particular year with the diameter proportional to the population of the income category that year.

Figure 9: Relative Happiness and Relative Income






Figure 10: Within-Country Life Satisfaction-Income Gradient



Source: Gallup World Poll, 2006

Notes: Figure shows a histogram, displaying the within-country satisfaction-income gradient estimated for each country. That is, each data point shows the coefficient from a country-specific ordered probit regression of satisfaction on log household income.


Figure 11: Comparing Within-Country and Between-Country Estimates—Gallup World Poll



Source: Gallup World Poll, 2006

Notes: Each dot shows the level of GDP, and the average level of well-being in each country. The slope of the arrow for each country corresponds with the estimated income-well-being gradient estimated from a country-specific ordered probit of well-being on log household income.


Figure 12: Comparing Within-Country and Between-Country Estimates—World Values Survey



Source: World Values Survey, 1999-2004 wave.

Notes: Each dot shows the level of GDP, and the average level of well-being in each country. The slope of the arrow for each country corresponds with the estimated income-well-being gradient estimated from a country-specific ordered probit of well-being on log household income.

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



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

Figure 14: Economic Growth and Subjective Well-Being




Figure 15: Analyzing Economic Growth and Changes in Life Satisfaction





Figure 16: Eurobarometer data on Growth in GDP and Life Satisfaction



Notes: Newey-West standard errors are reported, accounting for first-order autocorrelation.

Figure 17: Well-Being Trends in Europe




Figure 18: Happiness in the United States—General Social Survey

Income and Happiness Trends in the U.S.






Figure 19: Subjective Well-Being in Japan




Notes: Note that each separate series—shown separately in each of four panels—has a different mean, and hence comparisons should only be made within each panel.

Figure 20: Economic Progress and Growing Subjective Well-Being in Japan



Notes: The four dotted series are the measures of subjective well-being obtained from four separate questions, each covering different periods, and hence their levels are not immediately comparable. We pool these four series, and run a regression of well-being on series fixed effects, the unemployment rate and log GDP per capita. The coefficient estimates—shown in the text—were used to create the continuous series shown, adjusted for breaks and unemployment, respectively.


Figure 21: Bradburn Affect Scale—World Values Survey



Source: World Values Survey, 1981-84 and 1989-93 waves.

Notes: Top row shows responses to five questions relating to “positive affect.” Bottom row shows responses to five questions regarding “negative affect.”

Figure 22: Feelings and GDP—Gallup World Poll




Source: Gallup World Poll, 2006

Notes: Regression equations listed on chart are ordinary least squares regressions in which the dependent variable is the proportion of each country's population agreeing with each statement, and the independent variable is log GDP per capita.

Figure 23: Daily Experiences and GDP—Gallup World Poll

Daily Experiences and GDP

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.



Source: Gallup World Poll, 2006

Solid line shows regression line; dotted line is lowess fit.

Source: Gallup World Poll, 2006

Notes: Regression equations listed on chart are ordinary least squares regressions in which the dependent variable is the proportion of each country's population agreeing with each statement, and the independent variable is log GDP per capita.