Definitely incomplete: Please do not cite, post or distribute

Inequality and Subjective Well-Being

Betsey Stevenson
Wharton, University of Pennsylvania
CESifo and NBER
betsey.stevenson@wharton.upenn.edu
http://bpp.wharton.upenn.edu/betseys

Justin Wolfers
Wharton, University of Pennsylvania
CEPR, CESifo, IZA and NBER
jwolfers@nber.org
www.nber.org/~jwolfers

Abstract

Write your abstract here. It will turn out to be indented from the left, and use 1.5 line spacing. If you run out of room, the easiest thing to do is to simply make this single-line spaced. But for the moment, it is a useful format

This draft: February 1, 2010

First draft: 4/26/2009

Keywords: .

JEL codes: D03, D31, D6, E01, I3, O1, O4, O57, P0

I. Introduction

Does income inequality undermine subjective well-being? Existing research comparing average levels of happiness, or life satisfaction, with the level of income inequality in a country has yielded surprisingly mixed results. All told, subjective well-being data have yielded little reliable evidence to support the utilitarian notion that unequal incomes undermine average levels of well-being. This is all the more the puzzling because the pre-conditions of the utilitarian argument—that higher income yields successfully smaller marginal increases in well-being—finds strong support in the data.

Our initial contribution in this paper is to re-examine available data, so as to be more precise about the strength of any association between inequality and well-being. As with previous authors, we find mixed results. But whereas such findings have typically been interpreted as suggesting that well-being is unaffected by inequality, we provide a different interpretation. Previous authors have failed to be precise about the quantitative strength of the inequality–well-being link one should expect to find in the data, leading them to mis-label a statistically imprecise finding as a falsification of the null that inequality matters. By providing a more precise framework, we are able to test somewhat more precisely the utilitarian intuition that average levels of well-being in a society can be raised by distributing income more equally. This framework is useful in allowing us to distinguish the claim that inequality is unrelated to average levels of subjective well-being, from the reality that there simply is not enough variation in inequality to find this relationship.

Yet despite this fairly uncontroversial foundation, researchers have been surprisingly unsuccessful at linking higher income inequality with lower average levels of well-being. Armed with data from the first three waves of the world's largest cross-national study of well-being we revisit this earlier literature. As with previous research, we fail to locate a robust relationship between inequality and well-being. However, as our framework makes clear, this finding that reflects statistical imprecision, and the data should not be interpreted as convincingly showing the absence of an effect of inequality on happiness.

The key intuition of this result is as follows. The data in Figure 2 suggest that measured satisfaction is a function of the log of individual income. Consequently, if income were equally distributed, the average happiness of a nation should be proportional to the log of it's average income. Instead, average happiness is a function of the average of log(income). The wedge between the log of average income and the average of log income, which is often called mean log deviation, represents the proportionate reduction in average incomes that could accompany an elimination of inequality, and still

keep the population just as well off. To test this, we compare average income across nations with measures of both the log of average income, and this measure of income inequality. As it turns out, there is surprisingly little variation in the mean log deviation, and hence this sort of cross-country exercise has very little power to reject the usual utilitarian

II. Diminishing Marginal Benefit of Income and the Cost of Inequality

The utilitarian case against income inequality is simple: concentrations of income among the rich are also typically concentrations of income among those receiving low marginal benefit from that extra income. As such, redistributing from rich to poor will reduce the well-being of the rich by less than it increases the well-being of the poor. That is, the gains from reducing income inequality derive from the diminishing marginal benefit of income. Figure 1, below illustrates the basic mechanism: if well-being is a convex function of income, then redistribution can raise average well-being. The example shown considers redistributing $50,000 from a high-income household (initially earning $250,000) to a low-income household (initially earning $25,000). This more equal distribution of income yields only a small decline in measured well-being for the rich, which is more than offset by a large increase in well-being among the poor. Consequently, with diminishing marginal benefits declining inequality yields higher average levels of well-being. We should note that this cost of inequality comes directly from standard neoclassical assumptions about self-interested preferences, subject to diminishing returns; the possibility of other-regarding preferences or a direct distaste (or preference) for inequality can also affect the inequality–well-being link.

Figure 1: Diminishing Marginal Benefit and the Cost of Inequality

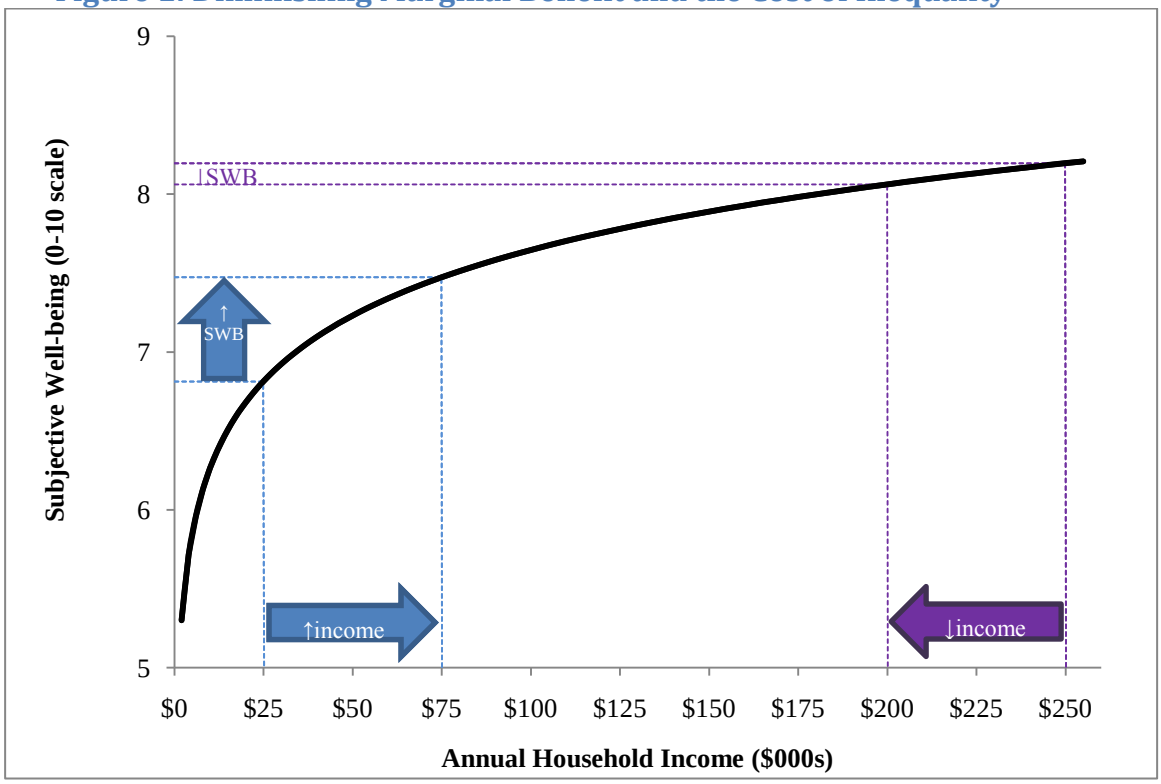

Before we turn to the data to check for evidence of diminishing marginal returns, it is worth clarifying the limits of what this exercise can reveal. Oswald (2008) notes that a person's response to a question asking about his or her subjective well-being may not be a linear function of one's true level of well-being. Consequently, a finding that reported well-being is a convex function of income may reflect either that true well-being is a convex function of income, or that reported well-being is a convex function of true well-being. Even so, this Oswald's observation doesn't negate the logic behind Figure 1: If reported well-being is a convex function of income, then greater income inequality should yield lower average levels of reported well-being. Instead, it reveals the limits of interpreting such a result: It may be that a finding that reported well-being is not convex function of income does not necessarily imply that true levels of well-being would rise if income inequality fell. For related reasons, we cannot simply rely on revealed preferences in order to make inferences about whether well-being is a convex function of income. In particular, one might be tempted to infer that robust evidence that people make risk-averse choices implies that subjective well-being is a convex function of income. We resist making such an argument simply because choices reveal the convexity of one's utility function (or indeed, whatever function it is that people maximize when they make choices), and it is by no means clear that one can or

should equate utility with reported subjective well-being. Our interest, for now, is to clarify the measurable relationship between reported well-being and income

Thus, in Figure 2, we show some simple evidence of the diminishing marginal benefit of income, graphing a non-parametric lowess fit of average levels of satisfaction against household income, for the world's 25 most populous nations. These data collect responses from the first three waves of the Gallup World Poll, and the specific measure of well-being is the Cantril Self-Anchoring Striving Scale, which 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. On which step of the ladder would you say you personally feel you stand at this time, assuming that the higher the step the better you feel about your life, and the lower the step the worse you feel about it? Which step comes closest to the way you feel?" Rungs on this ladder are numbered from zero, representing the worst possible life, through to ten, representing the best possible life.

The left panel of Figure 2 shows quite clearly that measured subjective well-being is a convex function of income. That is, the well-being–income gradient, which we interpret as the marginal benefit of an extra dollar, is particularly high for the poor, and much lower for the rich, suggesting a substantial cost of income inequality. Indeed, the similarity between these data and the hypothetical textbook relationship posited in Figure 1 is striking. In the right panel we re-estimate these lowess fits as a function of log household income, and the approximate linearity of this relationship for each of these 25 countires strongly suggests that the well-being–income relationship is approximately linear-log relationship. That is, this measure of subjective well-being well-being is roughly a linear function of the log of household income. While we focus on these data, alternative, alternative datasets such as the World Values Surveys, the Pew Global Attitudes Survey, the International Social Survey Program or the U.S.-based General Social Survey, and other measures of subjective well-being such as life satisfaction, or happiness, show a similar pattern.

Each of these comparisons was based on comparing rich and poor people within the same country. Consequently in Figure 3 we turn to cross-national comparisons, showing that that the same striking linear-log relationship also appears when contrasting national averages of satisfaction and income (measured as real GDP per capita, at purchasing power parity). Moreover, as Stevenson and Wolfers (2008) note, the relationship between subjective well-being and income derived from these cross-national comparisons is very similar to that derived from the earlier within-country cross-sectional comparisons.

All told, the left panels of both Figure 2 and Figure 3 suggest that there are large differences in the marginal benefit of income, suggesting a substantial cost to income inequality. Given the finding of a linear-log relationship, this suggests that 10 percent higher income is associated with roughly similar gains in income for the poor as for the rich. That is, the rise in well-being associated with the marginal dollar for a household earning $100,000 (roughly the 80th percentile of the U.S. distribution), is associated with one-fifth the increase in well-being as for a household earning $20,000 (which is both at around the 80th percentile in the distribution, and close to the poverty line for a family of four). These substantial differences in marginal returns suggest that income inequality may be quite an important factor undermining well-being. We now turn to assessing its quantitative importance.

III. A Simple Framework for Assessing the Costs of Income Inequality

At this point, it is worth formalizing our intuitions about the extent to which we expect income inequality to affect average levels of well-being. In line with the evidence in Figure 2, we begin by positing that the subjective well-being of an individual i, in country c is function of the log of their household income, y_{ic}, and other idiosyncratic factors, captured by the error term ϵ_{ic}:

$$Well-being_{ic} = \alpha + \beta \log(y_{ic}) + \epsilon_{ic} \quad (1)$$

If equation (1) accurately portrays individual income, then aggregating up to country-level averages yields:

$$\overline{Well-being}_c = \alpha + \beta \overline{\log(y)}_c + \epsilon_c \quad (2)$$

That is, average well-being in a country-year observation is a function of the average level of log income in a country. By contrast, previous research tracing the relationship between well-being and GDP (such as Deaton, 1997; Stevenson and Wolfers, 2008; and indeed, Figure 3, above) regresses average levels of happiness on the log of a measure of average income, such as GDP per capita. However, Income inequality drives a wedge between the log of average income $\log(\overline{y}_c)$, and the average of log income, $\overline{log(y)}_c$. Thus, we can rewrite equation (2) as:

$$\overline{Well-being}_c = \alpha + \beta \log(\overline{y}_c) - \beta \underbrace{\left(\log(\overline{y}_c) - \overline{\log(y)}_c\right)}_{\text{Mean log deviation}} + \epsilon_c \quad (3)$$

That is, average well-being in a country is positively related to the log of a measure of average income (such as GDP per capita), but negatively related to the gap between the average of the log of income and the log of average income. This latter term (in parentheses) is a standard measure of income

inequality, sometimes called the mean log deviation, or alternatively, Theil's Index. To be more precise, we estimate a mean log deviation for the United States of around 0.3, suggesting that—if Americans have log utility—aggregate well-being would be just as high if inequality were eliminated, even if it came at the cost of reducing average income by 0.3 log points, or around 26 percent. Thus, the mean log deviation can be considered a measure of compensating (or equivalent) variation. This intuitive scaling is particularly helpful for interpreting the results of cross-country well-being equations: if individual well-being really is a function of log income (as in equation (1)), then the coefficient on this measure of inequality should be equal to, but oppositely from the conventionally measured linear-log well-being–income gradient in a cross-national comparison (when we estimate regressions of the form of equation (3)). This precise quantitative prediction will allow us to discern whether the extent to which statistically insignificant effects of inequality on average happiness simply reflect a lack of statistical power, rather than a falsification of the hypothesis that the effects of log average income, and the mean log deviation are equal and opposite.

Before taking this to the data, it is also worth considering the possibility that well-being is a function not only of absolute income, but also of one's income relative to the national average:

$$\begin{aligned} Well-being_{ic} &= \alpha + \beta \log(y_{ic}) + \gamma \log\left(\frac{y_{ic}}{\overline{y}_c}\right) + \epsilon_{ic} \\ &= \alpha + (\beta+\gamma)\log(y_{ic}) - \gamma \log\left(\overline{y}_c\right) + \epsilon_{ic} \end{aligned} \tag{4}$$

This formulation nests both the simpler case in which well-being is only a function of absolute income (as in equation (1), when $\beta > 0$ and $\gamma = 0$), as well as the extreme case of well-being depending only on relative income as suggested by Easterlin (1974) (when $\beta = 0$ and $\gamma > 0$), as well as intermediate weights on these two objectives. Aggregating up to country averages yields:

$$\begin{aligned} \overline{Well-being}_c &= \alpha + (\beta+\gamma)\,\overline{\log(y)}_c - \gamma \log\left(\overline{y}_c\right) + \epsilon_c \\ &= \alpha + \beta \log(\overline{y}_c) - (\beta+\gamma)\underbrace{\left(\log\left(\overline{y}_c\right) - \overline{\log(y)}_c\right)}_{\text{Mean log deviation}} + \epsilon_{\epsilon} \end{aligned} \tag{5}$$

That is, the coefficient on the log of average income and the mean log deviation will still be oppositely signed, but to the extent that relative income comparisons are important, the estimated coefficient on income inequality will be of a larger magnitude. , This provides a way to test for relative income effects using aggregated country-level data.[1] (Indeed, in the extreme case where only relative

[1] Equally, we should recognize that this implication derives from an assumption about functional form—that well-being is a convex function of relative income, which in turn implies it is a convex function of individual income. The evidence in Figure 2 strongly suggests that this functional form assumption is warranted.

income comparisons matter ($\beta = 0$), and average well-being is positively related to income inequality, but is unrelated to the average level of income.) If the relevant comparison group is a sub-national reference group, this "extra" effect of inequality will be somewhat smaller.[2] To the extent that the relevant reference point is not another group, but one's own income in a previous period τ (as in theories X the invoking the "hedonic treadmill," then there is no extra effect of inequality on steady-state levels of average well-being. Instead, these intertemporal linkages will lead to a decline in average well-being during periods in which income inequality is widening, and a rise in average well-being when inequality is declining.[3] Finally, if comparisons are made on the basis of relative levels of well-being, rather than one's relative level of income, or rankings of relative income, then the conclusions in our central case—equation (3)—continue to hold, because neither of these comparisons can change, on average.

To summarize: this simple framework suggests that the mean log deviation is an interpretible measure of the likely consequences of income inequality. In fact, in our central case, the quantitative prediction is that the coefficient on the mean log deviation should be equal and opposite to the coefficient on the log of average income (or perhaps even a bit larger if local income comparisons are important).

[It is also worth thinking about identification when countries are choosing redistributional policies to maximize well-being.]

IV. Average Well-Being and Income Inequality

Our primary source of data on subjective well-being is the Gallup World Poll. These data are ideal for our purposes, because they contain observations on subjective well-being for 149 countries,

---

[2] That is, if we revise equation (4), so that comparisons are made relative to a contemporaneous sub-national reference group, r, whose average income is denote $\overline{y}_{rc}$: $Well-being_{irc} = \alpha + \beta \log(y_{irc}) + \gamma \log\left(\frac{y_{irc}}{\overline{y}_{rc}}\right) + \epsilon_{irc}$ then averaging across individuals within a country yields: $\overline{Well-being}_c = \alpha + \beta \log(\overline{y}_c) - (\beta + \gamma)\left(\log(\overline{y}_c) - \overline{\log(y)}_c\right) + \gamma\left(\log(\overline{y}_c) - \overline{log(y)}_{rc}\right)$. The final term in parentheses suggests that there is a partial offset to the cost of inequality equal to the mean log deviation, calculated as if each person's income as at their group-specific mean. That is, relative income comparisons contribute to the cost of inequality only to the extent that individual income differs from the group-specific norms.

[3] That is, if we revise equation (4), so that comparisons are made relative to one's own income τ periods ago: $Well-being_{ict} = \alpha + \beta \log(y_{ict}) + \gamma \log\left(\frac{y_{ict}}{y_{ic,t-\tau}}\right) + \epsilon_{ict}$ and averaging across individuals within a country yields: $\overline{Well-being}_{ct} = \alpha + \beta \log(\overline{y}_{ct}) - \beta\left(\log(\overline{y}_{ct}) - \overline{\log(y)}_{ct}\right) + \gamma\Delta \log(\overline{y}_{ct}) - \gamma\Delta\left(\log(\overline{y}_{ct}) - \overline{log(y)}_{c,t}\right) + \epsilon_{ict}$. Comparing this expression to our central case in equation (3), the only new element added by taking account of this historic reference point is that it adds interesting dynamics, with the first differences of both average income and inequality now also important.

which account for over 95% of the world's population. While our sample contains three waves of data, corresponding to those data from the 2006, 2007 and 2008 waves that had been processed by April 20, 2009, we combine them to form a single observation for each country.

We begin by analyzing responses to the "ladder of life" question, which 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. On which step of the ladder would you say you personally feel you stand at this time, assuming that the higher the step the better you feel about your life, and the lower the step the worse you feel about it? Which step comes closest to the way you feel?"*

These same data also contain a more standard measure of life satisfaction, which (following the World Values Survey), asks for a zero (dissatisfied) to ten (satisfied) response to:

> *"All things considered, how satisfied are you with your life as a whole these days?"*

While we have individual-level responses to these questions, our analysis will focus on national aggregates, which requires an appropriate normalization and aggregation of individual responses to qualitative questions. Our approach throughout this paper will be to simply take the zero to ten responses as given, and to standardize, by subtracting the mean, and dividing by the standard deviation of the relevant sample. The advantage of this approach is that it yields measures that are transparent, easy to calculate, and at least somewhat comparable when comparing the estimated well-being–income or well-being–inequality gradient measured from responses to a 3-item scale as when measured from responses to a 4-item, 10, or 11-point scale. This normalization also ensures the estimated well-being–income gradient is directly interpretable, as the well-being measure is scaled by the unconditional cross-sectional dispersion in happiness.[4] The disadvantage of this approach is that it is clearly ad hoc, as it assumes, for instance, that the difference between adjacent rungs on the response ladder are equally spaced. Fortunately, these scaling issues turn out to be more troubling in theory than in practice as Stevenson and Wolfers (2008) show that this standardized measure yields estimates of national happiness averages that are extremely highly correlated ($\rho > 0.99$) with alternative approaches.

[4] In Stevenson and Wolfers (2008), we estimated well-being aggregates as the coefficients from an ordered probit of well-being on country fixed effects, which yielded very similar estimates. The most important difference is that the ordered probit scales differences relative to the standard deviation of well-being conditional on country dummies, while the simpler normalization in this paper scales differences relative to the (larger) unconditional standard deviation of well-being. Given that country fixed effects account for about 20% of the variation in well-being (that is, $R^2 \approx 0.2$ in an OLS regression of satisfaction on country fixed effects), this simpler normalization will tend to yield estimates of the well-being–income gradient that are about nine-tenths as large ($\sqrt{1 - R^2} \approx 0.9$).

Our primary source of income inequality data is the World Bank's World Development Indicators database, which uses data on the distribution of income or consumption from household surveys provided by country teams, and consistent methods to estimate Gini coefficients from grouped data.

□□□□□□□□□□□□□□□□□□□□□□□□□□
income and expenditure data shows household surveys that collect income and expenditure data. Names and detailed information on household surveys can be found on the website of the International Household Survey Network (www.surveynetwork.org). Core Welfare Indicator Questionnaire Surveys (CWIQ), developed by the World Bank, measure changes in key social indicators for different population groups—specifically indicators of access, utilization, and satisfaction with core social and economic services. Expenditure survey/budget surveys (ES/BS) collect detailed information on household consumption as well as on general demographic, social, and economic characteristics. Integrated household surveys (IHS) collect detailed information on a wide variety of topics, including health, education, economic activities, housing, and utilities. Income surveys (IS) collect information on the income and wealth of households as well as various social and economic characteristics. Labor force surveys (LFS) collect information on employment, unemployment, hours of work, income, and wages. Living Standards Measurement Studies (LSMS), developed by the World Bank, provide a comprehensive picture of household welfare and the factors that affect it; they typically incorporate data collection at the individual, household, and community levels. Priority surveys (PS) are a light monitoring survey, designed by the World Bank, for collecting data from a large number of households cost-effectively and quickly. Income tax registers (ITR) provide information on a population's income and allowance, suchas gross income, taxable income, and taxes by socioeconomic group. 1-2-3 surveys (1-2-3) are implemented in three phases and collect sociodemographic and employment data, data on the informal sector, and information on living conditions and household consumption.

are themselves an update of the original Deininger and Squire database, and important sources include the Luxembourg Income Study, and Transmonee.

Bannerjee & Duflo: We also follow previous studies in adding 6.6 to the gini when it was constructed from expenditure instead of income.

These data have the virtue of being collected by country experts, and compiled into a cross-national dataset by , and of having been made as comparable as possible. Even so we should emphasize that there are vast methodological differences in the dozens of individual country data collections behind each datapointThese cross-national comparisons have been cobbled together from national sources, which are based on different surveys, population concepts, equivalence scales, weighting procedures, whether net or gross, and the use of income versus consumption.

The measure of average income we analyze is Gross Domestic Product per capita, measured at Purchasing Power Parity. The primary data that we analyze come from the same survey, and Gallup have carefully constructed measures of real household income, measured at purchasing power parity. We

have identified several anomalies in these data, and so we have dropped data from Sweden and X from our analysis.

Our empirical approach follows two strategies. In the top panel of Table X, I estimate equation (3)

V. Errors in Measuring Income Inequality

VI. Analysis of Gallup World Poll

Results still to come. Some are shown in the attached PowerPoint document.

VII. Conclusions

Your text goes here…

VIII. Works Cited

Blanchflower, David G. *International Evidence on Well-being.* NBER Working Paper, Cambridge, MA: National Bureau of Economic Research, 2008.

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.

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, aging, health and wellbeing around the world: Evidence from the Gallup World Poll." Working Paper, 2007.

Di Tella, Rafael, and Robert MacCulloch. *Happiness Adaptation to Income beyond "Basic Needs".* NBER Working Paper, National Bureau of Economic Research, 2008.

Diener, Ed. "Subjective well-being: The Science of Happiness and a Proposal for a National Index." *American Psychologist* 55, no. 1 (2000): 34-43.

Diener, Ed, and Martin E.P. Seligman. "Beyond money: Toward an economy of well-being." *Psychological Science in the Public Interest* 5 (2004): 1-31.

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.

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. "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., and Onnicha Sawangfa. *Happiness and Economic Growth: Does the Cross Section Predict Time Trends? Evidence from Developing Countries.* mimeo, University of Southern California, 2009.

Frey, Bruno S., and Alois Stutzer. "What Can Economists Learn from Happiness Research?" *Journal of Economic Literature* 40 (2002): 402-435.

Layard, Richard. "Happiness: Has Social Science a Clue." *Lionel Robbins Memorial Lectures 2002/3.* London School of Economics, 2003.

—. *Happiness: Lessons from a New Science*. London: Penguin, 2005.

Oswald, Andrew J. "On the Curvature of the Reporting Function from Objective Reality to Subjective Feelings." *Economics Letters*, 2008.

Stevenson, Betsey, and Justin Wolfers. "Economic Growth and Happiness: Reassessing the Easterlin Paradox." *Brookings Papers on Economic Activity*, Spring 2008: 1-87.

Stevenson, Betsey, and Justin Wolfers. "Marriage and Divorce: Changes and Their Driving Forces." *Journal of Economic Perspectives*, 2007: 27-52.

Stevenson, Betsey, and Justin Wolfers. *The Paradox of Declining Female Happiness.* mimeo, University of Pennsylvania, 2007.

Wolfers, Justin. "Is Business Cycle Volatility Costly? Evidence from Surveys of Subjective Well-being." *International Finance* 6, no. 1 (2003): 1-26.

Table 1: Cross-Country Comparisons—Gallup World Poll

Dependent Variable: Average level of satisfaction in each country (ladder score, 0-10)	
	Gallup_MLD
Inequality measure	
Panel A: Bivariate relationship: $Well-being_c = \alpha + \beta_{inequality} Mean\ log\ deviation_c$ **(n=?)**	
$Inequality$	
$\overline{R}^2$	
Panel B: Controlling for Log(GDP)	
Panel B: Bivariate relationship: $Well-being_c = \alpha + \beta_y Log(GDP_c) + \beta_{inequality} Mean\ log\ deviation_c$ **(n=?)**	
Test: $\beta_y + \beta_{inequality} = 0$	
Panel C: Controlling for country characteristics:	
Test: $\beta_y + \beta_{inequality} = 0$	

Figure 2: Within-Country Comparisons of Subjective Well-Being and Income

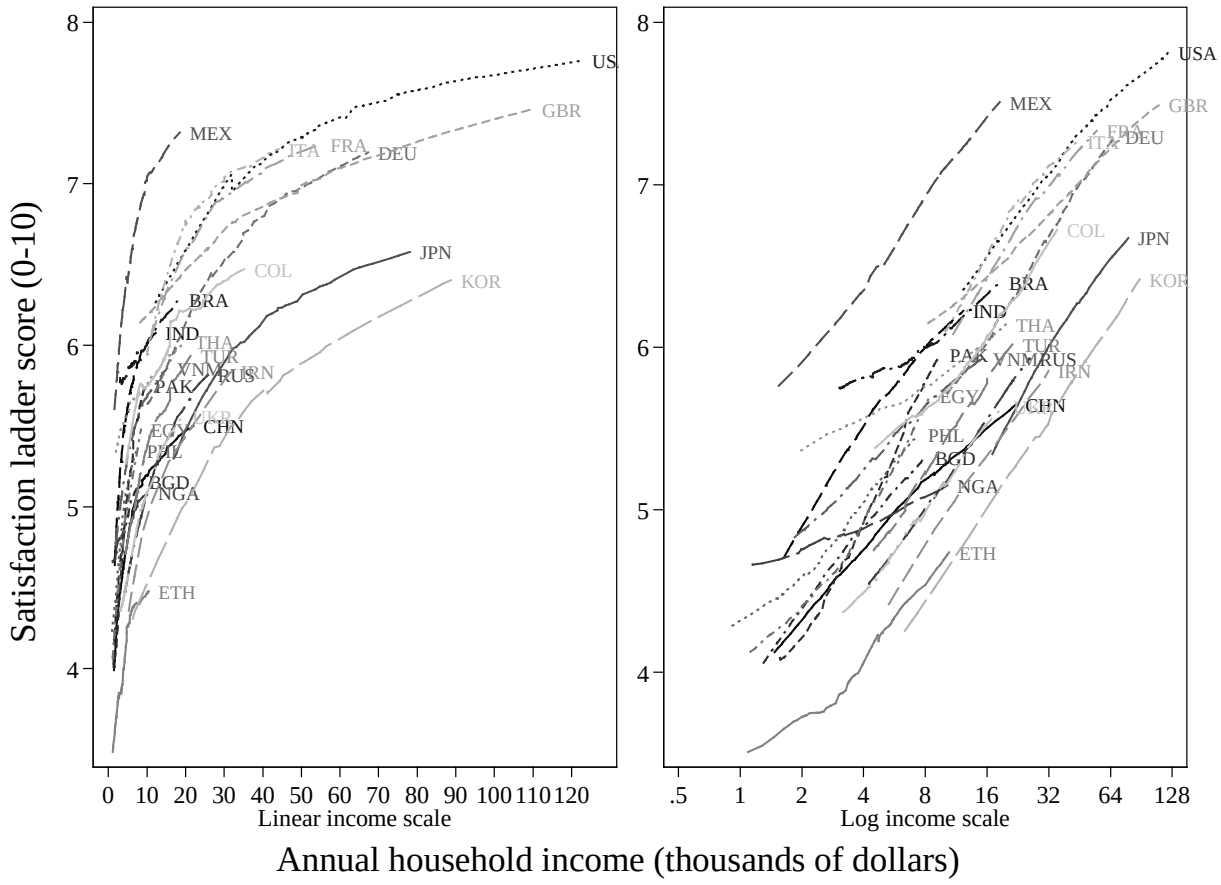

Source: Gallup World Poll, waves 1-3.

Notes: Figure shows the relationship between satisfaction and household income, estimated separately for each of the 25 most populous countries. This bivariate relationship was estimated separately for each country, using local linear (lowess) regressions with a bandwidth of 0.8. The results are shown over the range of incomes ranging from the 10th to the 90th percentiles of that country's income distribution.

Figure 3: Cross-Country Comparisons of Average Satisfaction and Income