

Inequality and Subjective Well-Being

Betsy Stevenson
University of Michigan

betseys@umich.edu
www.nber.org/~bstevens

Justin Wolfers
University of Michigan
R

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

Abstract

Diminishing marginal well-being from income suggests that redistributing income from the rich to the poor will raise average levels of well-being within a society. This utilitarian logic suggests that—conditional on average income—countries with greater income inequality will experience lower levels of average well-being. Yet existing research has failed to find clear evidence that inequality undermines average levels of subjective well-being, and many have concluded that therefore no such relationship exists. We develop a model that quantifies the utilitarian hypothesis quite precisely and find that existing data cannot reject the utilitarian intuition that economic inequality undermines average levels of well-being. However, equally, we are unable to reject the hypothesis that the inequality does not impact well-being. In short, we show that there is so little variation in economic inequality around the world that existing results reflect imprecise estimates, rather than important insight into the link between inequality and well-being.

This draft: October 12, 2018

Keywords: Income, subjective well-being, inequality

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

I. Introduction

Diminishing marginal well-being from income suggests that an equal distribution of income within a society will result in the greatest total well-being for the society. If some individuals are wealthier than others, their well-being gains from each extra dollar are smaller than the losses incurred by removing those dollars from poorer households. This utilitarian logic suggests that—conditional on average income—societies with more economic inequality will experience lower levels of subjective well-being.

The pre-conditions of the utilitarian argument—that each extra dollar of income yields successfully smaller marginal increases in well-being—finds strong support in the data. However, research comparing average levels of subjective well-being with the level of income inequality in a country has yielded mixed results. The literature to date has yielded little reliable evidence to support the utilitarian notion that unequal incomes undermine average levels of well-being.

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. Examining bivariate correlations across countries, we find that average subjective well-being is associated with lower economic inequality, but much of this relationship is due to a negative correlation between inequality and per capita GDP. When controls for GDP per capita are included, inequality is no longer statistically significantly related to subjective well-being.

Earlier papers that yielded similar findings have typically been interpreted as suggesting that well-being is unaffected by inequality; however 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 economic inequality undermines average well-being. By providing a clear theoretical framework, we are able to test more precisely the utilitarian intuition that average levels of well-being in a society are higher when income is distributed more equally. Our 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 around the world to be able to estimate the link between inequality and well-being with much precision.

Armed with data from the first four waves of the world’s largest cross-national study of well-being we revisit this earlier literature. We also bring new data to bear on the question, providing new internationally-comparable indices of income inequality. Even so, we fail to isolate a clearly statistically

significant relationship between income inequality and subjective well-being—a finding consistent with the mixed results found in previous research. However, as our framework makes clear, this finding 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. Our data suggest that measured satisfaction is a function of the log of individual income. Aggregating to the national level, this suggests that average happiness in a country is proportional to the average of log income. More often, cross-national studies associate average happiness with the log of a measure of average income (like GDP per capita). This substitution of the log of average income for the average of log income would be irrelevant were income equally distributed, since these two measures are equal when there is no income inequality. But when income is not equally distributed there is a wedge between the log of average income and the average of log income. This wedge is a measure of income inequality called the mean log deviation. Under the usual utilitarian logic, this wedge also represents the proportionate reduction in average incomes that could accompany an elimination of inequality, while keeping the population just as well off. That is, the mean log deviation is a compensating variation measure of the costs of income inequality.

Consequently the usual utilitarian logic suggests that average well-being bears the same proportionate relationship to $\log(\text{GDP})$ as it does to this measure of income inequality. To test this, we compare measures of average well-being across nations with measures of both the log of average income (such as $\log(\text{GDP})$), and this particular measure of income inequality. As it turns out, there is surprisingly little variation in the mean log deviation around the world, and hence this sort of cross-country exercise has very little power to reject the utilitarian null that income inequality undermines average well-being. We use power calculations to illustrate this point. In addition to the utilitarian logic there may also be direct consequences of inequality on well-being if people's preferences include the well-being of others. We expand our model to include this possibility and will show that this does little to solve the power problem.

We proceed by developing this logic as follows. In section II we articulate the textbook case that the diminishing marginal benefit of income means that income inequality undermines average levels of well-being. We also collect evidence showing that subjective well-being does in fact exhibit diminishing marginal benefit from income. Section III provides our theoretical innovation, developing precise quantitative predictions from this textbook model. Section IV turns to assembling the relevant data, describing first the normalization of ordinal well-being data into cardinal country-aggregates, and next,

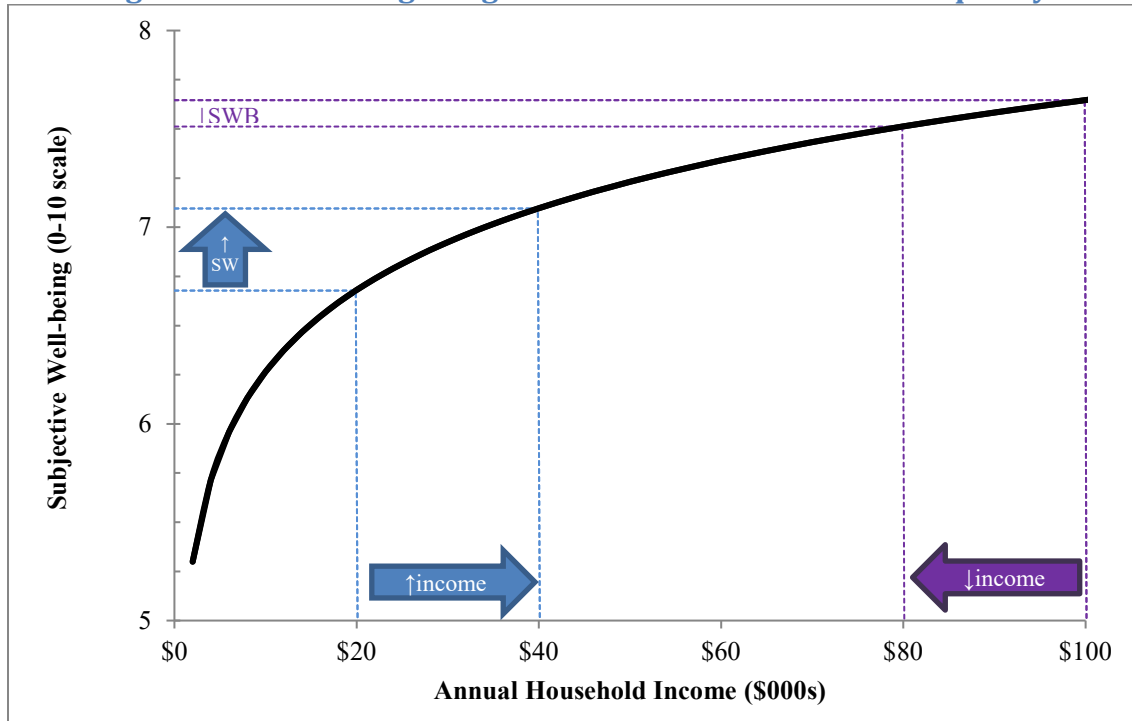
our sources of income inequality data. Section V provides the heart of the analysis; section VI adds some robustness checking, and section VII concludes.

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 an extra unit of 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.¹ Figure 1, below illustrates the basic mechanism: if well-being is a concave function of income, then redistribution can raise average well-being. The example shown considers redistributing \$50,000, reducing the income of the rich household from \$250,000 to \$200,000, and increasing the income of the poor, from \$25,000 to \$75,000. 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, as long as income is subject to diminishing marginal well-being benefits, lower levels of income inequality yield higher levels of average 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.

¹ While subjective well-being and utility are likely related, we shall resist confounding the two terms. Consequently we refer to diminishing marginal benefits of income, rather than diminishing marginal utility.

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 concave function of income may reflect that true well-being is a concave function of income. Alternatively, true well-being may not be a concave function of income, but if true well-being rises with income and reported well-being is a concave function of true well-being, then reported well-being will be a concave function of income. Even in this latter case, the logic behind Figure 1 holds. However, the interpretation is restricted to the relationship between reported well-being and inequality. If reported well-being is a concave function of income, then greater income inequality should yield lower average levels of reported well-being. Yet this need not imply that true levels of average well-being would rise if income inequality fell. While this distinction is important for interpretation, our concern right now is with the measured relationship and, under either interpretation, measured levels of average well-being would rise if income inequality fell.

A different concern regarding interpretation is the link between subjective well-being and utility. While the two are undoubtedly related, it is quite possible that subjective well-being is simply one component of utility and should not be interpreted as representing utility directly (Becker and Rayos, 2007). As such, we cannot simply rely on revealed preferences in order to make inferences about whether subjective well-being is a concave 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 concave function of income. We resist making such an argument simply because choices reveal the concavity of one's utility function, and it is by no means clear that one can or should equate their utility with reported subjective well-being.

While understanding the link between reported well-being and true well-being or that between reported well-being and utility is obviously crucial for evaluating the usefulness of well-being data, our focus in this paper is to clarify the relationship that is the focus of the existing data: the relationship between average levels of reported well-being, income, and income inequality.

We now turn to the data and in Figure 2 we show some simple evidence of the diminishing marginal benefit of income. This figure analyzes data from both poor and rich households, plotting non-parametric lowess fits of levels of satisfaction against household income, in each the world's 25 most populous nations. These data collect responses from the first four waves of the Gallup World Poll. 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 concave 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 the relationships replacing household income with the log of household income. The approximate linearity of this relationship for each of these 25 countries strongly suggests that the well-being–income relationship is an approximately linear-log relationship. That is, this measure of subjective well-being is roughly a linear function of the log of household income. While we focus on data assessing scores on the satisfaction ladder in the Gallup World Poll, 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, which include 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-country 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-country comparisons is very similar to that derived from the earlier within-country cross-sectional comparisons.

The left panels of both Figure 2 and Figure 3 show large differences in the marginal benefit of income, thereby 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 is one-fifth the increase in well-being that would occur if that marginal dollar went instead to a household earning \$20,000. These substantial differences in marginal returns when comparing the 80th and 20th percentile of US household income 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 impact average levels of well-being. In line with the evidence in Figure 2, we begin by positing that the subjective well-being of individual i , in country c is a function of the log of their household income, y_{ic} , and other idiosyncratic factors, captured by the error term ϵ_{ic} :

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

If equation (1) accurately portrays the individual relationship between well-being and income, then aggregating up to country-level averages yields:

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

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) regressed average levels of happiness on the log of a measure of average income, such as GDP per capita. However, this simplification is inaccurate in the presence of income inequality. Income inequality creates a wedge between the log of average income $\log(\bar{y}_c)$, and the average of log income, $\overline{\log(y)}_c$, and this wedge is known as the mean log deviation. Thus, we can rewrite equation (2) as:

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

When there is no income inequality the two measures—the log of average income $\log(\bar{y}_c)$, and the average of log income, $\overline{\log(y)}_c$ —are equal. If this is the case, then it is equivalent to estimate average levels of happiness in a country as a function of either the average level of log income in a country, or as a function of the log of average income. However, in the presence of inequality, average well-being in a country is positively related to the log of a measure of average income (such as GDP per capita) and 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 L.

Writing out the relationship clarifies the cost of inequality in terms of subjective well-being in the utilitarian framework. If the mean log deviation rises by x points, it will decrease average well-being by βx , while an increase in average income of x log points raises well-being by βx . Thus, the mean log deviation can be considered a measure of compensating (or equivalent) variation. To give an example, we estimate a mean log deviation for the United States of around 0.3, suggesting that—since the well-being of Americans rises with the log of their income—aggregate well-being would be just as high if inequality were eliminated and average income was lower by 0.3 log points, or around 26 percent.

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 in a cross-country regression of the form shown in equation (3), the coefficient on this measure of inequality should be equal to, but oppositely signed from, the coefficient on the log of average income. This precise quantitative prediction yields clearer hypotheses to test regarding the relationship between subjective well-being and inequality. Previous studies have simply asked whether the coefficient on inequality in a regression of subjective well-being on income inequality was statistically significantly different from zero. The utilitarian framework yields a new null: the coefficient on inequality is equal and oppositely signed to the coefficient on the log of average income. This hypothesis will allow us to discern the extent to which coefficient estimates that are not statistically significantly different from zero provide evidence that inequality does not affect subjective well-being from their simply lacking statistical power.

Thus far we have described the simple utilitarian framework in which subjective well-being is related to one's income. We can also extend this analysis to the case in which subjective well-being is a function of relative income. Following Easterlin (1995) we can re-write the well-being income relation

such that well-being is a not only of absolute income, but also of one's income relative to the national average:

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

This formulation nests both the simpler case in which well-being is only a function of absolute income—when $\beta > 0$ and $\gamma = 0$, equation (4) reduces to equation (1)—as well as the extreme case of well-being depending only on relative income— $\beta = 0$ and $\gamma > 0$ —as suggested by Easterlin (1974). Intermediate weights on β and γ allow for both absolute and relative income to impact well-being.

Aggregating equation (4) up to country averages yields:

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

The coefficient on the log of average income and the mean log deviation are still 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.² A hypothesis that relative income and not absolute income matters for well-being implies that $\beta = 0$ and $\beta + \gamma > 0$. Failure to find such a relationship is a failure to confirm the relative income hypothesis. However, a finding that the coefficient on inequality is larger than that on the log of average income suggests either that both relative income and absolute income both matter, or that beyond the utilitarian calculus represented in equation (4), there should be an additional term reflecting a direct individual aversion to income inequality.

Equations (4) and (5) illustrate the relationship when the country is the relevant reference group. If instead the relevant comparison group is a smaller subset of the country then the “extra” effect of inequality will be somewhat smaller.³ To the extent that the relevant reference point is not another group,

² This implication derives from an assumption about functional form—that well-being is a concave function of relative income, which in turn implies it is a concave function of individual income. The evidence in Figure 2 strongly suggests that this functional form assumption is warranted.

³ 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 \bar{y}_{rc} : $Well - being_{irc} = \alpha + \beta \log(y_{irc}) + \gamma \log\left(\frac{y_{irc}}{\bar{y}_{rc}}\right) + \epsilon_{irc}$ then averaging across individuals within a country yields: $\overline{Well - being}_c = \alpha + \beta \log(\bar{y}_c) - (\beta + \gamma) \left(\log(\bar{y}_c) - \overline{\log(y)}_c\right) + \epsilon_c$

but one's own income in a previous period τ 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.⁴ 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 a directly interpretable measure of the likely costs 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 income comparisons are important).

IV. Measuring Subjective Well-Being

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 154 countries, which account for over 95% of the world's population. While our sample contains four waves of data, corresponding to those data from the 2006-2009 waves that had been processed by October 16, 2009, we combine them to form a single well-being index 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

$\overline{\log(y)_c} + \gamma (\log(\bar{y}_c) - \overline{\log(y)_{rc}})$. 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.

⁴ 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:

$$\overline{Well - being}_{ct} = \alpha + \beta \log(\bar{y}_{ct}) - \beta (\log(\bar{y}_{ct}) - \overline{\log(y)_{ct}}) + \gamma \Delta \log(\bar{y}_{ct}) - \gamma \Delta (\log(\bar{y}_{ct}) - \overline{\log(y)_{c,t}}) + \epsilon_{ict}.$$

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.

your life, and the lower the step the worse you feel about it? Which step comes closest to the way you feel?”

The same survey poses a binary happiness question, asking:

“Did you experience the following feelings during A LOT OF THE DAY yesterday? How about Happiness?”

This question was asked in 131 countries in the third and fourth waves. Similar questions also ask about daily experiences of: enjoyment, physical pain, worry, sadness, stress, boredom, depression, anger, love and fear. A related question asks:

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. Did you smile or laugh a lot yesterday?”

In 122 of these countries Gallup also collects data on an alternative life satisfaction question which tracks that posed in the World Values Survey, asking for a zero (dissatisfied) to ten (satisfied) response to:

“All things considered, how satisfied are you with your life as a whole these days?”

We will also show some results when we focus instead on happiness data from the World Values Survey, which asks:

Taking all things together, would you say you are very happy, rather happy, not very happy, not at all happy?”

While we have individual-level responses to all of 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 (across all respondents in all countries), and dividing by the standard deviation of the relevant sample. The advantage of this approach is that it yields measures that are transparent and easy to calculate; in our graphical analysis, we will also use the secondary axis to label each measure on its original scale. Applying the same approach to alternative qualitative indicators ensures that the estimated coefficients will be at least somewhat comparable when analyzing the estimated well-being–inequality gradient measured from responses to a binary, four-item, or 10, or 11-point scale. This normalization also ensures the estimated well-being–income gradient is directly interpretable, as this index is scaled by the unconditional cross-sectional dispersion in each relevant well-being measure.⁵ The disadvantage of this approach is that it is clearly ad hoc, as it assumes,

⁵ 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

for instance, that the differences 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 well-being averages that are extremely highly correlated ($\rho > 0.99$) with alternative approaches.

V. Estimates of Income Inequality

Surprisingly, the more difficult task is in obtaining comparable international data on income inequality. As such, we pursue two avenues. First, we exploit newly-available data from the Gallup World Poll to generate new estimates of income inequality by country. The coverage of the Gallup World Poll is so large that we are able to calculate income inequality for a number of countries for which previously there existed no useful data. The other real strength of this approach is that it provides the first truly comprehensive global survey of income inequality that uses a consistent unit of observation, income concept, survey methodology and secondary processing. Our data come from an individual-level survey, which asks respondents about their household income. Consequently the unit of observation is an individual. The underlying income concept—which is the same in all countries—is real household income, which is a gross (or pre-tax) measure, rather than a net (or post-tax) measure. Gallup has adjusted for international differences in purchasing power parity, using the PPP adjustments from the 2005 International Comparison Program and these GDP-based PPP ratios are adjusted for subsequent inflation relative to the United States. For our purposes, the PPP adjustments aren't important since the mean log deviation of income in a country remains the same, no matter what currency it is measured in. However the inflation adjustments are important since they make household incomes measured in different years comparable.

The shortcoming of our income concept is that it does not adjust for differences in household size, because we currently have only incomplete data on the household register. But the existing literature gives reason to believe that may not be an important concern. Deininger and Squire (1996) analyze the “sixty-seven cases... in which information on both households and individuals is available from reasonably reputable sources,” finding that “the mean difference between person-based and household-based Gini coefficients is 1.69 [on a 0-100 scale],” leading them to conclude that measures of inequality

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 ($\sqrt{1 - R^2} \approx 0.9$).

based on either definition are both acceptable and roughly comparable. Of course, our measure is somewhere between the two types of measures they compare—our income concept is household income, but the sampling frame is individuals.

We are able to assess empirically the effect of variation in household size on our estimates of inequality for a subset of countries for which we have data on the presence of children in the household. We estimate the relationship between subjective well-being and the log of household income separately for those with and without children present. We find that this relationship is extremely similar in both cases, suggesting that adjustments for household size may yield only second-order changes when considering the well-being-equivalent level of income for any household size. (We are expecting to be able to make these adjustments in future drafts.)

We have also had to significantly clean these data. In particular, we do not use any income from the first wave in which the income questions were experimental. We only use data from the second wave if no income information is available for a country in either the third or fourth waves. If income data are available from both the third and fourth waves, we use both sets of data. In some country-wave surveys, there are an unusually high proportion of respondents reporting zero income, and Gallup have informed us that this may reflect a mis-coding of respondents who actually refused to answer the income question. We simply dropped country-waves in which more than 5 percent of all respondents reported zero income, a rule which still allowed us to retain all but a handful of poor-country samples. Many country samples contain implausible outliers, and so we “winsorized” the top and bottom five percent of the income distribution,⁶ so that no observation has undue influence on our measures of inequality. If a country sample yielded a measured income distribution that was so coarse that the measured 5th and 25th percentiles (or 75th and 95th percentiles) were equal, we simply dropped it (this only affects Mali and Hong Kong). We also drop two particularly problematic observations—those from Sweden in wave 2 and Sierra Leone in wave 3—each of which yielded results suggesting inconsistent coding of the units for measuring income. Stevenson and Wolfers (2010) provide further detail on the steps taken to clean these data.

⁶ Typically, “winsorizing” means replacing observations below the 5th, or above the 95th percentiles with the values of the 5th, or 95th percentile, respectively. But this would lead to downward biased of the extent of income inequality. Consequently, we replaced these tail observations with an estimate of the geometric mean income of the bottom 5% or top 5% of incomes. To generate the left tail estimate, we ran an interval regression of log income on a constant, setting the incomes of those in the bottom 5% of the distribution as unbounded below the 5th percentile; in order to ensure that the distribution we estimate reflects only the shape of the left tail, incomes above the 25th percentile were set as unbounded above that percentile. A symmetric approach (using the 75th-95th percentiles of the income distribution) was used to estimate the average log income of those in the top 5% of the distribution.

Our alternative inequality indicators come from the most reliable public data source—the World Bank’s World Development Indicators database. This database is an updated and refined version of the original Deininger and Squire (1996) database, and important sources include the Luxembourg Income Study, and Transmonee. These data include the Gini coefficient (and often also quintile income shares) from nationally-representative household surveys which country-teams have reported are “acceptable.” When only income-share data were available, the POVCAL procedure is used to impute a Gini index from grouped data. These data have the virtue of being collected by country experts, and compiled into a cross-national dataset that is easy to access and use. Even so we should emphasize that there are vast methodological differences in the dozens of individual country data collections behind each datapoint. These cross-national comparisons have been cobbled together from national sources, which are based on different surveys, units of observation (household or individual), population concepts, equivalence scales, weighting procedures, income concepts (both net and gross income are acceptable), and the use of either income or expenditure measures. Moreover, there exists scant documentation on many national surveys, and few adjustments are made to ensure comparability. Deininger and Squire (1996) argue that a particularly important distinction is between surveys of expenditure (which are particularly common in poor countries), and surveys of income. Based on 47 observations in which they observe both income- and expenditure-based Gini coefficients, Deininger and Squire recommend adjusting expenditure-based measures up by 6.6 points, to make them comparable with income-based measures. While we follow Deininger and Squire’s recommendation, we don’t make any further adjustments to the published data.

Another concern is that the availability of these data depends on when a particular country most recently produced a usable household survey. In order to maximize the coverage of these data, we simply take the most recent estimate of the Gini coefficient for any country. This yields a total of 128 countries, with estimates taken from as far back as 1992 (the average year of the relevant survey is 2002.6).

Finally, the World Bank estimates the Gini coefficient, while our theoretical framework highlights the importance of an alternative measure of income inequality, the mean log deviation. Nonetheless, we use a simple approximation, noting that if income is log-normally distributed, then:

$$\text{Mean log deviation}_c = \left(F_N^{-1} \left(\frac{1 + gini_c}{2} \right) \right)^2 \quad (6)$$

where $F_N^{-1}(\cdot)$ is the inverse cumulative normal distribution. We can use the underlying microdata from the Gallup World Poll to calculate both the Mean Log Deviation and the Gini Coefficient, and check whether this formula yields a reasonable mapping between the two. Figure 4 shows a very close relationship, suggesting that the assumption of log-normally distributed income is quite reasonable. In

fact, for the Gallup data, the correlation between the mean log deviation and the approximation suggested by equation (6) is 0.47.

Throughout this paper, we will report on the well-being–inequality link based first on our new estimates of inequality by country, and then based on the above transformation of the World Bank estimates. In order to get some sense of why these results may differ, Figure 5 shows the relationship between these two alternative measures of the mean log deviation.

VI. Analysis of Gallup World Poll

We begin by examining the bivariate relationship between subjective well-being and inequality. Figure 6 shows there is a clear negative correlation between subjective well-being and inequality. The countries with the most inequality have the lowest average life satisfaction and this relationship holds equally in the Gallup World Poll data (shown in the left panel) and the World Bank data (shown in the right panel). In both datasets average life satisfaction is negatively correlated with the mean log deviation, with a correlation coefficient of -0.4. The regression estimates at the bottom of each panel shows that the relationship is of a similar magnitude in each dataset and is highly statistically significant.

This relationship does not, however, take into account differences in GDP per capita in the various countries. Indeed, it appears as if the most unequal countries—for example, Zimbabwe, Namibia, Angola, and Tunisia—are also some of the poorest countries. Figure 7 illustrates that this is indeed the case. Log GDP per capita is negatively correlated with the mean log deviation, with a correlation coefficient of -0.4. Previous research has shown that average happiness rises with GDP per capita (Stevenson and Wolfers, 2008; Deaton, 2007). Moreover, the utilitarian framework that we’ve laid out suggests that well-being should be negatively related to inequality, *conditional on the log of average income*. Thus, we next turn to examining the relationship between subjective well-being and inequality controlling for log GDP per capita.

Figure 8 shows that the clear negative correlation between life satisfaction and inequality does not hold once we have conditioned on the log of GDP per capita. The partial correlation is approximately -0.1 in both datasets. While there still appears to be a negative relationship between well-being and inequality, the relationship is no longer precisely estimated and in both the Gallup World Poll data and the World Bank data the coefficient on inequality is statistically insignificantly different from zero.

Thus, in the simple bivariate relationship we are able to reject a null hypothesis that the coefficient on inequality is equal to zero. However, once controls are added for log GDP per capita, we

can no longer reject the null that the coefficient on inequality is zero. Recall that the utilitarian framework provides another null—that the coefficient on inequality should be equal in magnitude and opposite in sign to the coefficient on the log of GDP per capita. An F-test reveals that in neither the Gallup nor World Bank data can we reject this hypothesis. Therefore while we cannot reject that the coefficient is zero—that inequality does not matter for subjective well-being—neither can we reject that the coefficient is that which is predicted by the utilitarian framework—that inequality reduces well-being due to the diminishing marginal well-being of income. This alternative null illustrates the imprecision of the estimates since we are neither able to reject a hypothesis that inequality has no impact on well-being nor are we able to reject a hypothesis that inequality impacts well-being.

In order to attempt to distinguish between these two hypotheses we next turn to adding additional controls to the regressions. In Table 2 we use data from the Gallup World Poll and add controls. The first two columns show the results of the bivariate relationship and that controlling for the log of GDP per capita. The third column adds a control for continents, the fourth column adds controls for each of the sub-indices that comprise the Human Development Index, and the fifth adds a battery of controls: percent urban, population age, the log of inflation, and the agricultural sector, services, government spending, private consumption, exports, and imports each as a share of GDP. The addition of these controls reduces the coefficient estimate on the mean log deviation however it remains in all cases both statistically insignificantly different from zero and statistically insignificantly different from the negative of the coefficient on log GDP per capita. That is, we are unable to reject the utilitarian hypothesis in all cases.

We next turn to alternative data sets and alternative measures of subjective well-being. In table 3 we show that we get similarly imprecise results using both measures of life satisfaction in the Gallup World Poll, using both life satisfaction and happiness in the World Values Survey, using a satisfaction ladder question in the Pew Global Attitudes survey, and a happiness question in the European Social Survey. Using all of these measures of subjective well-being we remain unable to reject a hypothesis that inequality has no impact on well-being and unable to reject a hypothesis that the impact of inequality on well-being is the opposite of the impact of log GDP on well-being. It should be noted that in all of the regressions, the log of GDP per capita has a statistically significant positively relationship with well-being.

The utilitarian framework is useful in understanding the variation in inequality around the world. Recall that the mean log deviation represents the proportionate reduction in average incomes that could accompany an elimination of inequality, while keeping the population just as well off. That is, the utilitarian approach suggests that average well-being bears the same proportionate relationship to

$\log(\text{GDP})$ as it does to this measure of income inequality. Therefore we can examine the distribution of inequality around the world in terms of the amount of GDP that would be required to eliminate it. Figure 9 shows the variation around the mean in log points of the log of GDP per capita and the mean log deviation. This comparison highlights the fact that there is very little variation in inequality relative to the variation in income around the world. The bottom panel shows that the variation in mean log deviation is similar whether it is calculated conditional on the log of GDP per capita or unconditionally. It is this lack of variation in mean log deviation relative to GDP per capita that means that the cross-country exercise has very little power to reject the utilitarian null that income inequality undermines average well-being.

VII. Conclusions

Reported subjective well-being is linearly related to the log of income and therefore clearly exhibits diminishing marginal sensitivity to income. This relationship provides a framework for forming a hypothesis about the relationship between reported subjective well-being and inequality. The coefficient on inequality as measured by the mean log deviation should be equal and opposite to that on the log of GDP per capita. Taking relative income comparisons into account leads to a refined hypothesis that the coefficient on inequality should be oppositely signed to and larger in absolute value to that on the log of GDP per capita.

We examine cross-country data and find that inequality is negatively correlated with subjective well-being, but that this relationship does not hold once controls for GDP are added. However, while this relationship is statistically insignificantly different from zero, it is still quantitatively important since the alternative hypothesis generated by the utilitarian framework also cannot be rejected. The problem that the utilitarian framework makes clear is that there is too little variation in inequality across countries or through time to be able to accurately discern the relationship. Data cannot convincingly falsify most reasonable views about the quantitative link between happiness and inequality. This is not a problem stemming from data collection but rather reflects the fact that most of the inequality in the world occurs between countries rather than within countries. Therefore it is unlikely that it is possible to accumulate sufficient data to resolve the happiness-inequality link.

VIII. Works Cited

- Blanchflower, David G. 2008. *International Evidence on Well-being*. NBER Working Paper, Cambridge, MA: National Bureau of Economic Research.
- Blanchflower, David, and Andrew Oswald. 2004. "Well-Being Over Time in Britain and the USA." *Journal of Public Economics* 88 (7-8): 1359-1386.
- Clark, Andrew E., Paul Frijters, and Michael A. Shields. 2008. "Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles." *Journal of Economic Literature* 46 (1): 95-144.
- Deaton, Angus. 2007. "Income, aging, health and wellbeing around the world: Evidence from the Gallup World Poll." Working Paper.
- Di Tella, Rafael, and Robert MacCulloch. 2008. *Happiness Adaptation to Income beyond "Basic Needs"*. NBER Working Paper, National Bureau of Economic Research.
- Diener, Ed. 2000. "Subjective well-being: The Science of Happiness and a Proposal for a National Index." *American Psychologist* 55 (1): 34-43.
- Diener, Ed, and Martin E.P. Seligman. 2004. "Beyond money: Toward an economy of well-being." *Psychological Science in the Public Interest* 5: 1-31.
- Easterlin, Richard A. 1995. "Will Raising the Incomes of All Increase the Happiness of All?" *Journal of Economic Behavior and Organization* 27 (1): 35-48.
- Easterlin, Richard A. 2005b. "Diminishing Marginal Utility of Income? Caveat Emptor." *Social Indicators REsearch* 70 (3): 243-255.
- Easterlin, Richard A. 1974. "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.
- Easterlin, Richard A. 2005. "Feeding the Illusion of Growth and Happiness: A Reply to Hagerty and Veenhoven." *Social Indicators Research* 74 (3): 429-443.
- Easterlin, Richard A. 2001. "Income and Happiness: Towards a Unified Theory." *The Economic Journal* 111 (473): 465-484.
- Easterlin, Richard A., and Onnicha Sawangfa. 2009. *Happiness and Economic Growth: Does the Cross Section Predict Time Trends? Evidence from Developing Countries*. mimeo, University of Southern California.
- Feenstra, R. C., Inklaar, R., Timmer, M. (2013). *The next generation of the Penn World Table*. NBER Working Paper 19255.

- Frey, Bruno S., and Alois Stutzer. 2002. "What Can Economists Learn from Happiness Research?" *Journal of Economic Literature* 40: 402-435.
- Layard, Richard. 2003. "Happiness: Has Social Science a Clue." *Lionel Robbins Memorial Lectures 2002/3*. London School of Economics.
- . 2005. *Happiness: Lessons from a New Science*. London: Penguin.
- Luttmer, Erzo F. P. 2005. "Neighbors as Negatives: Relative Earnings and Well-Being." *Quarterly Journal of Economics* 120 (3): 963-1002.
- Neves, P. C., Afonso, Ó., Tavares Silva, S. (2016). A meta-analytic reassessment of the effects of inequality on growth. *World Development*, 78, 386-400.
- Ostry, J., Berg, A., Tsangarides, C. G. (2014). Redistribution, inequality, and growth. IMF Staff Discussion Note.
- Oswald, Andrew J. 2008. "On the Curvature of the Reporting Function from Objective Reality to Subjective Feelings." *Economics Letters*.
- Stevenson, Betsey and Justin Wolfers. 2013 "Subjective Well-Being and Income: Is There Any Evidence of Satiation?" *American Economic Review*, 103(3), 598–604.
- . 2009. "The Paradox of Declining Female Happiness." *American Economic Journal: Economic Policy* 1 (2): 190-225.
- . 2008. "Economic Growth and Happiness: Reassessing the Easterlin Paradox." *Brookings Papers on Economic Activity* 1-87.
- . 2007. "Marriage and Divorce: Changes and Their Driving Forces." *Journal of Economic Perspectives* 27-52.
- Wolfers, Justin. 2003. "Is Business Cycle Volatility Costly? Evidence from Surveys of Subjective Well-being." *International Finance* 6 (1): 1-26.

Table 1: Inequality and Income

	Reference Group	Individual Well-being	Coefficient on log(GDP)	Coefficient on inequality (MLD)
Absolute Income	None	$\beta \log(y_{ic})$	β	$-\beta$
Relative Income	Country	$+\gamma \log\left(\frac{y_{ic}}{\bar{y}_c}\right)$	0	$-\gamma$
Relative Income (Local Reference)	Local	$+\delta \log\left(\frac{y_{ic}}{\bar{y}_{rc}}\right)$	0	$-\delta\left(1 - \frac{\sigma_{\bar{y}_r}^2}{\sigma_{\bar{y}}^2}\right)$
Hedonic Treadmill	Own past income	$+\zeta \log\left(\frac{y_{i,t}}{y_{i,t-\tau}}\right)$	0	0
Relative Well-being	Any	$+\eta \log\left(\frac{U_{ic}}{\bar{U}_c}\right)$	0	0
Income Rank	Any	$+\theta \text{rank}(y_{ic})$	0	0

Table 2: Subjective Well-Being and Inequality

	(1)	(2)	(3)	(4)	(5)
MLD (Gallup)	-1.35*** (0.25)	-0.20 (0.18)	-0.20 (0.19)	-0.13 (0.20)	-0.06 (0.21)
Log GDP		0.29 (0.02)	0.27 (0.02)	0.24*** (0.04)	0.30*** (0.08)
Controls	None	+Log GDP	+Continent	+HDI Indices	+% urban, %0-14; %0-65; %ag; %service; ln(π), C/Y; G/Y; M/Y; X/Y
N	140	140	140	137	127
Adjusted R ²	0.17	0.65	0.68	0.69	0.74
H ₀ : $\beta_{inequality} = 0$	Reject	Accept	Accept	Accept	Accept
H ₀ : $\beta_{inequality} = -\beta_{income}$		Accept	Accept	Accept	Accept
		F=0.22 (p=0.64)			

Note: Dependent Variable: Satisfaction ladder score (standardized) Numbers in parentheses are robust standard errors. Asterisks indicate statistical significance at the *10 percent, **5 percent, and *1 percent level.

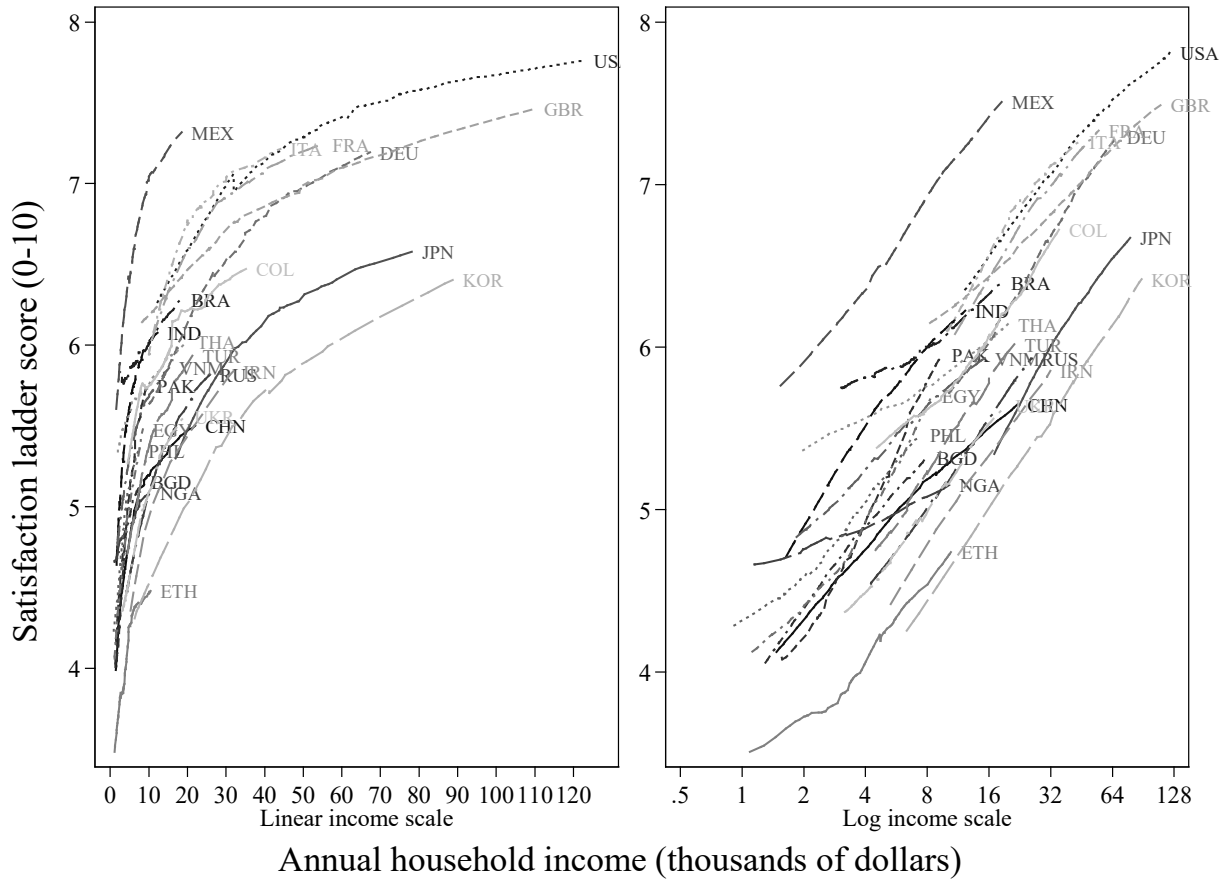
Table 3: Alternative Measures of Subjective Well-Being

Survey:	Gallup World Poll		World Values Survey (2004-08)		Pew Global Attitudes, 2007	European Social Survey, 2006-07
Dependent Variable:	Satisfaction Ladder	Life Satisfaction	Life Satisfaction	Happiness	Satisfaction Ladder	Happiness
Inequality (MLD, Gallup)	-0.20 (0.18)	-0.28 (0.25)	0.44 (0.52)	0.47 (0.45)	-0.06 (0.37)	-0.19 (0.62)
Log GDP	0.29*** (0.02)	0.33*** (0.03)	0.27*** (0.05)	0.13*** (0.04)	0.23*** (0.04)	0.68*** (0.07)
N	140	113	54	54	43	23
Adjusted R ²	0.65	0.63	0.36	0.12	0.48	0.79
H ₀ : $\beta_{inequality} = 0$	Accept	Accept	Accept	Accept	Accept	Accept
H ₀ : $\beta_{inequality} = -\beta_{income}$	Accept	Accept	Accept	Accept	Accept	Accept

Note: Dependent Variable: Subjective well-being (standardized)


Numbers in parentheses are robust standard errors. Asterisks indicate statistical significance at the *10 percent, **5 percent, and *1 percent level.

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



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. DATA ARE FROM THE GALLUP WORLD POLL.

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



Notes: Figure shows the relationship between average satisfaction and real GDP per capita in purchasing power parity in constant 2005 international dollars. Sample includes 131 developed and developing countries. In each panel the short- and long-dashed lines are fitted from regressions of satisfaction on GDP per capita and the log of GDP per capita, respectively. Real GDP per capita is at purchasing power parity in constant 2005 international dollars. Data are from the Gallup World Poll.

Figure 4: Comparing Mean Log Deviation with a Proxy Based on Adjusted Gini Coefficients

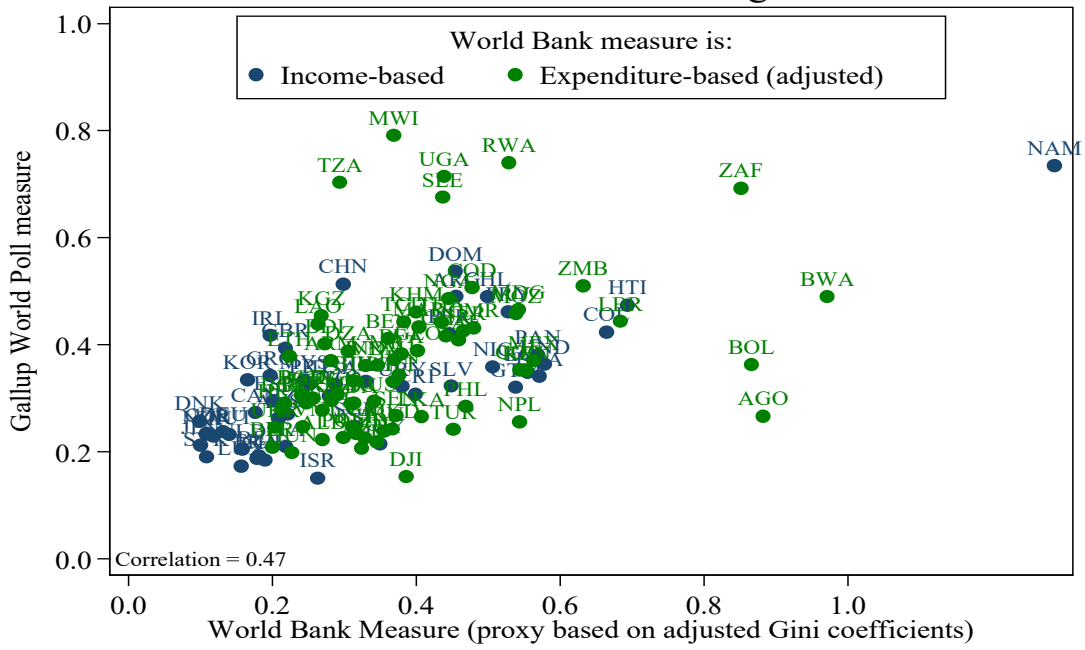


Figure 5: Alternative Measures of the Gini Coefficient

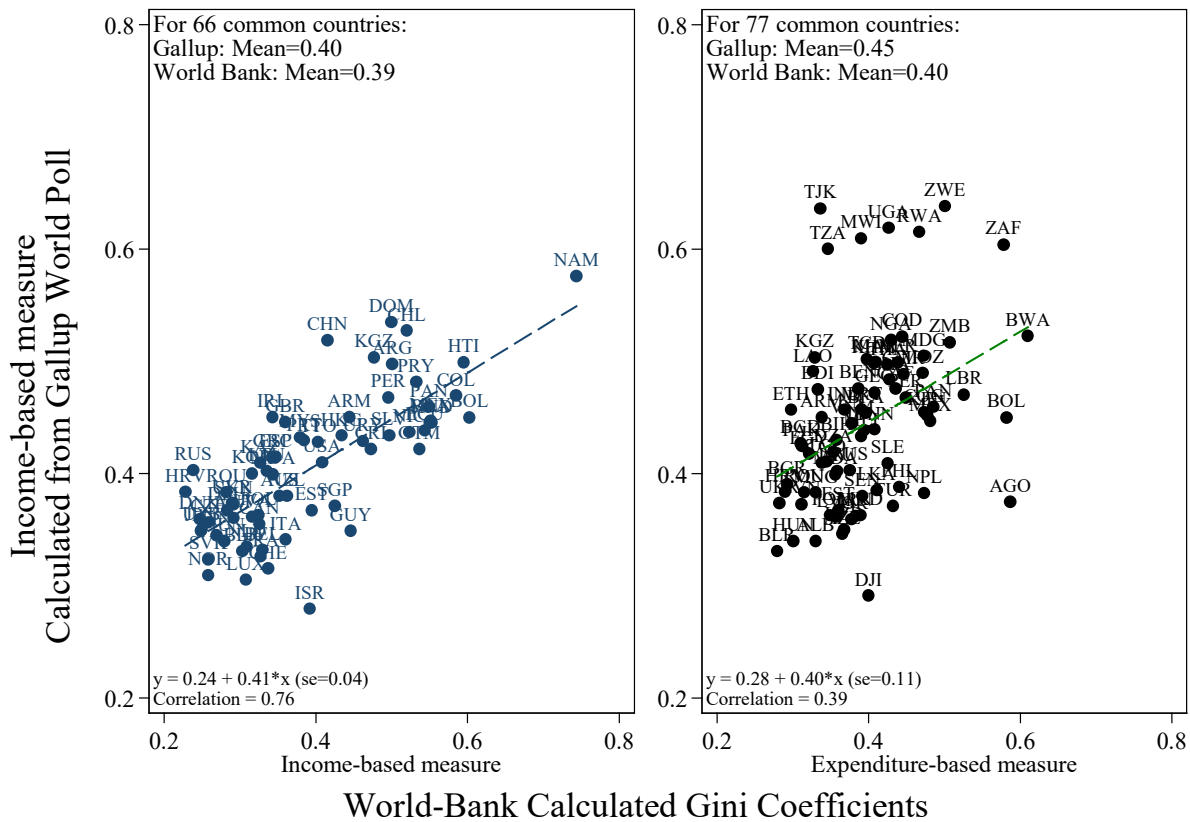


Figure 6: Raw Correlation Between Satisfaction and Inequality

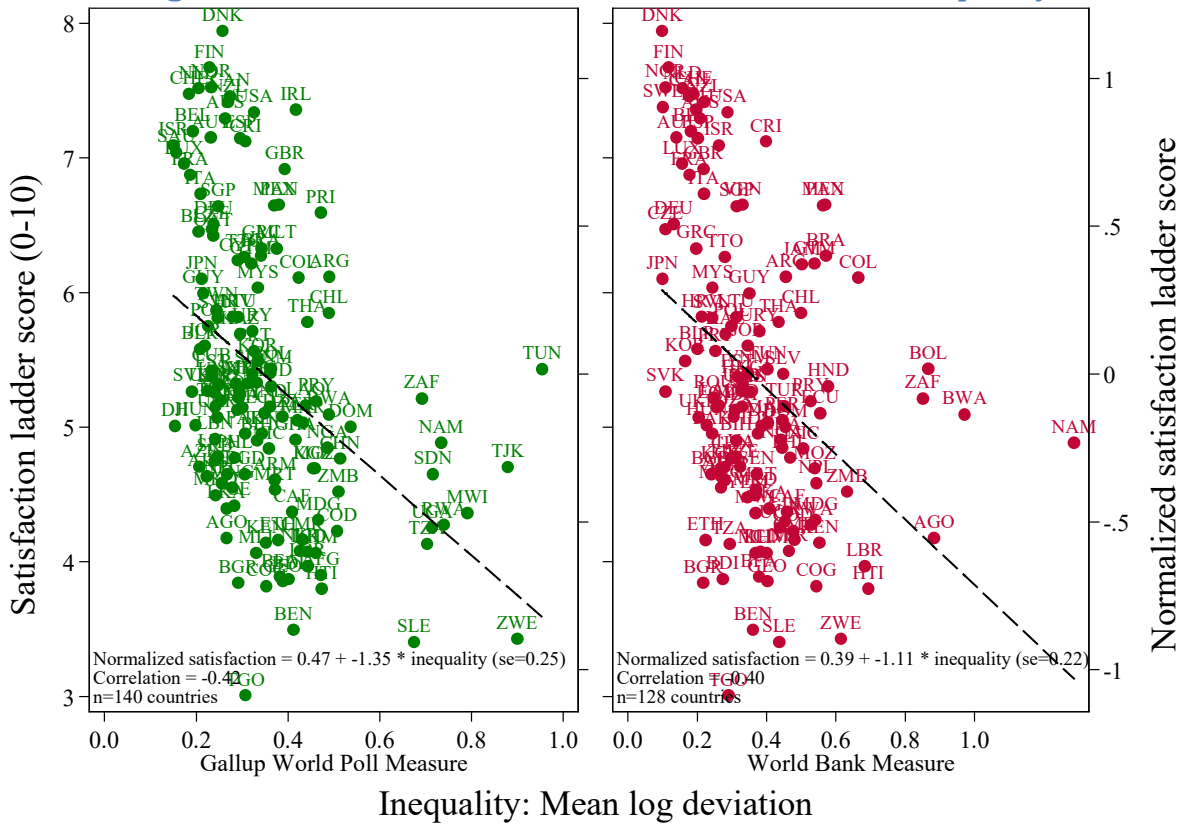


Figure 7: Raw Correlation Between Log GDP and Inequality

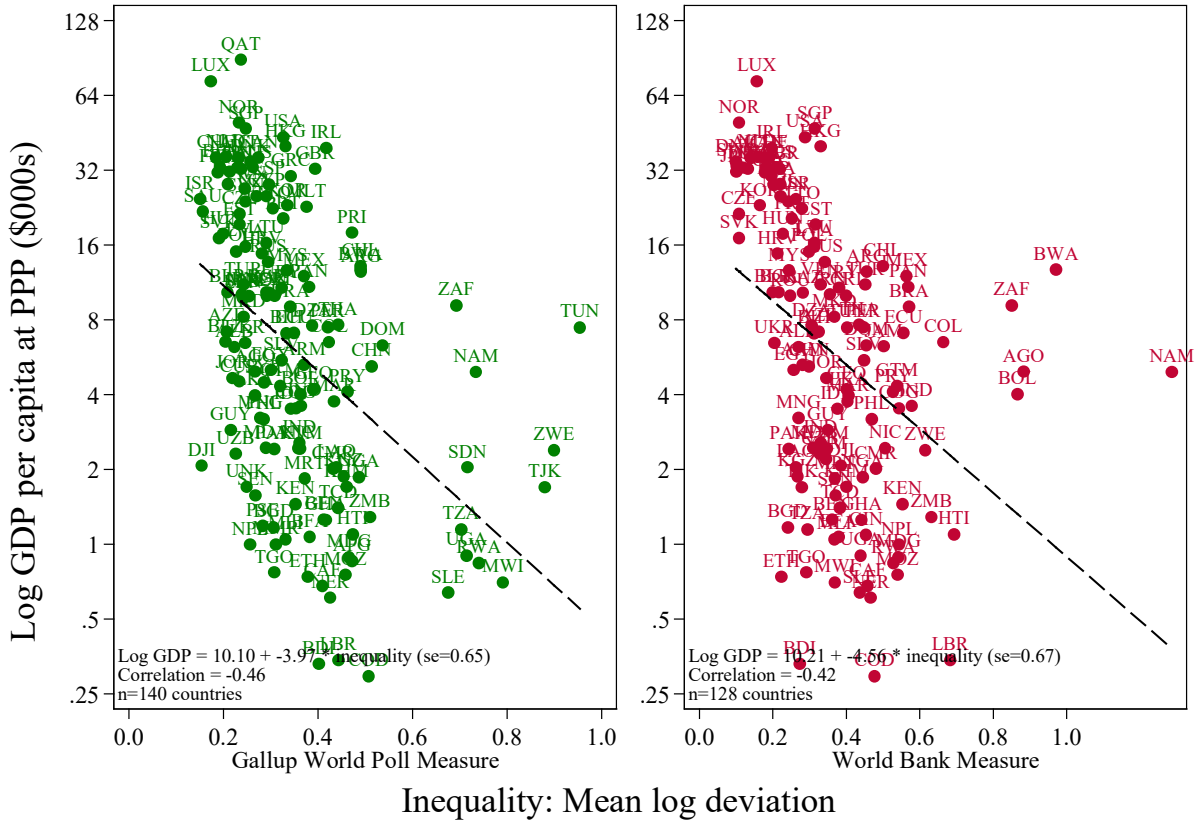


Figure 8: Relationship Between Satisfaction and Inequality Conditional on Log GDP

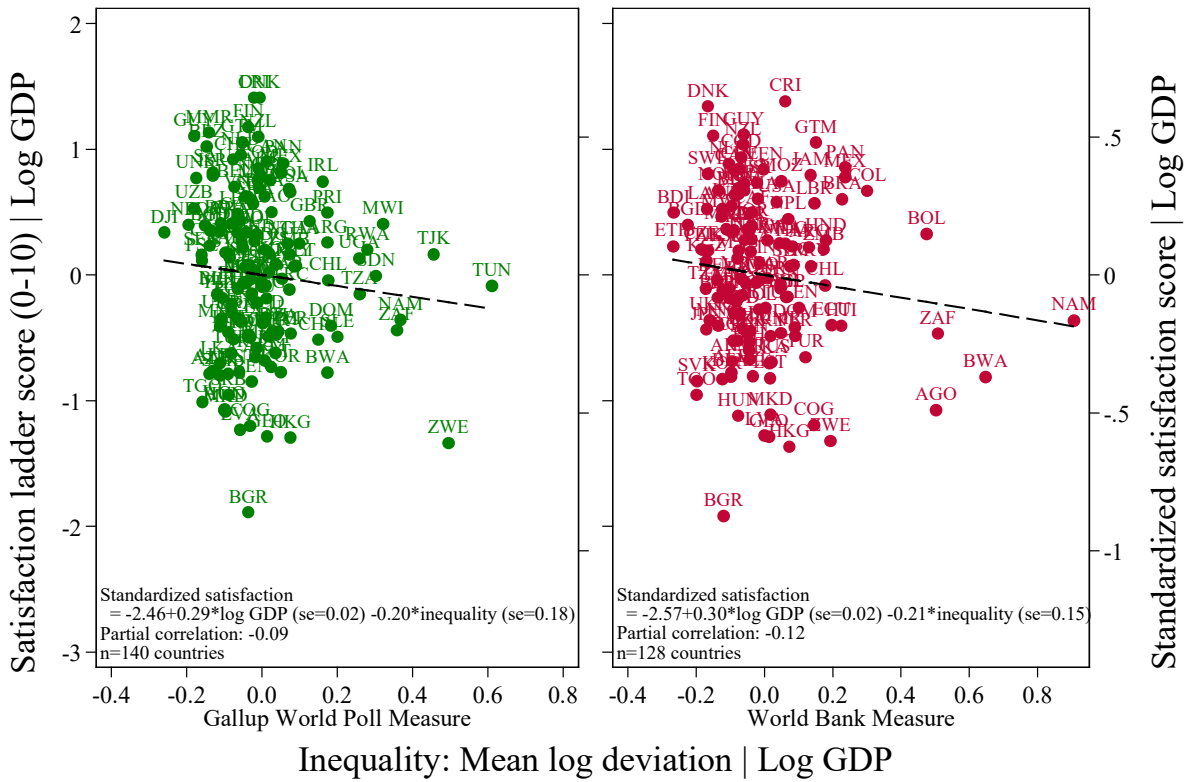


Figure 9: Variation in Log GDP and Mean Log Deviation

