

Motu Working Paper 20-09

The Relationship between Subjective Wellbeing and Subjective Wellbeing Inequality: Taking Ordinality and Skewness Seriously

Motu economic & public policy research

Arthur Grimes, Stephen P. Jenkins, Florencia Tranquilli

October 2020



Document information

Author contact details

Arthur Grimes

Motu Economic and Public Policy Research, and Victoria University of Wellington
arthur.grimes@motu.org.nz

Stephen P. Jenkins

London School of Economics and Political Science, and IZA
s.jenkins@lse.ac.uk

Florencia Tranquilli

Motu Economic and Public Policy Research, and Victoria University of Wellington
flortranquilli@hotmail.com

Acknowledgements

We thank Shine Wu for excellent technical assistance, and also David Maré and Phillip Morrison for insightful comments on an earlier draft.

Motu Economic and Public Policy Research

PO Box 24390 info@motu.org.nz +64 4 9394250

Wellington www.motu.org.nz

New Zealand

© 2020 Motu Economic and Public Policy Research Trust and the authors. Short extracts, not exceeding two paragraphs, may be quoted provided clear attribution is given. Motu Working Papers are research materials circulated by their authors for purposes of information and discussion. They have not necessarily undergone formal peer review or editorial treatment. ISSN 1176-2667 (Print), ISSN 1177-9047 (Online).

Abstract

We argue that the relationship between individual satisfaction with life (SWL) and SWL inequality is more complex than described by leading earlier research such as Goff, Helliwell, and Mayraz (Economic Inquiry, 2018). Using inequality indices appropriate for ordinal data, our analysis using the World Values Survey reveals that skewness of the SWL distribution, not only inequality, matters for individual SWL outcomes; so too does whether we look upwards or downwards at the (skewed) distribution. Our results are consistent with there being negative (positive) externalities for an individual's SWL from seeing people who are low (high) in the SWL distribution.

JEL codes

D31, D63, I31

Keywords

subjective wellbeing, ordinal data, inequality, skewness, WVS

Summary haiku

Inequality
and skewness of wellbeing
matter for people

1 Introduction

People have higher satisfaction with life (SWL) if they live in a country with low inequality of subjective wellbeing, *ceteris paribus*, according to an important recent paper by Goff, Helliwell, and Mayraz (GHM, 2018).¹ GHM also find that the association between individuals' SWL and their country's inequality of wellbeing is stronger than the association between their SWL and the country's inequality of income.² These findings are significant, not only for highlighting a link between personal wellbeing and inequality, but also for highlighting the potential usefulness of wellbeing inequality as a comprehensive measure that incorporates differences in aspects of utility beyond factors that relate to income alone.³ In this paper, we show that when more methodologically-appropriate measures of SWL inequality are employed, the link between an individual's SWL and their country's SWL distribution is more complex than found by earlier research.

The principal measure that GHM use to summarize wellbeing inequality is the standard deviation of SWL within a country. By using this measure, they are following much other research, from pioneering early analysis of wellbeing inequality (e.g. Veenhoven, 1990) through to the latest *World Happiness Report* (Helliwell et al., 2020).

However, recent research argues convincingly that the standard deviation is not an appropriate index of inequality for ordinal data, and SWL is an ordinal measure. Although GHM also utilize an ordinal measure of dispersion, it is not one of the inequality indices that is consistent with the recommendations of theoretical research about how to summarize wellbeing inequality. In this paper, we analyze whether – and in what form – the relationship between individuals' SWL and the wellbeing inequality of their country holds when a more appropriate set of inequality indices for ordinal data are employed. As a result, we provide not only a novel methodological contribution but also a substantive contribution to understanding of the relationship between individual-level SWL and SWL inequality.

Our paper unfolds as follows. In Sections II and III, we review GHM's analysis and findings and the modern literature on inequality measurement using ordinal data. After introducing the World Values Survey (WVS) data used by GHM and ourselves (Section IV), we replicate GHM's empirical analysis using the WVS data, as well as extend their model specification to better account for the time-series cross-section nature of the data. This analysis uses GHM's preferred SWL inequality

¹ We use the terms 'subjective wellbeing', 'wellbeing', and 'satisfaction with life' interchangeably.

² For reviews of the relationship between individual SWL and income inequality, see Schneider (2016) and Ngamaba et al. (2018).

³ Other studies using measures of SWL inequality to explain outcomes include Veenhoven (1990), Stevenson and Wolfers (2008), Ott (2015), and Grimes and Wesselbaum (2019).

measure, the standard deviation (*SD*). In Section V, we repeat the analysis substituting for the standard deviation each of eight SWL inequality indices derived from recent research.

Our regression results indicate that the negative association between individual SWL and country-level SWL inequality that GHM report is not robust to an extension of their specification, and nor does a relationship exist between wellbeing inequality and individual SWL for some of the alternative inequality indices we employ. However, when we use theoretically appropriate measures for summarizing SWL inequality – those due to Cowell and Flachaire (2017) – a strong relationship is re-established. Particularly interesting is that we find either a positive or a negative relationship between wellbeing inequality and individual SWL depending on which of the measures we use.

We therefore analyze the reasons behind these differing results and show that the skewness of the SWL distribution in a country may be at least as important as inequality *per se* in affecting individuals' SWL (Section VI). Section VII contains our summary and conclusions.

2 GHM'S analysis and findings

GHM analyze the relationship between individual SWL and wellbeing inequality within countries using a battery of tests and multiple time-series cross-section data sources. They examine whether both wellbeing inequality and income inequality (separately and together) are associated with individuals' SWL once one controls for differences in personal characteristics and region. GHM's core regression model specification is of the form:

$$(1) \quad SWL_{ijt} = \alpha + \beta_1 SD_{ijt} + \beta_2 G_{ijt} + \beta_3 Y_{ijt} + \beta_4' X_{ijkt} + \beta_5' R_{ijkt} + \varepsilon_{ijt},$$

where SWL_{ijt} denotes the SWL of person j in country i and wave t ,⁴ SD_{ijt} denotes the country-wave standard deviation of SWL, G_{ijt} is the country-wave Gini coefficient of income, Y_{ijt} is the country-wave logarithm of GDP per capita in purchasing power parity (PPP) terms, and X_{ijkt} is a set of individual-level controls for sex, age, education, employment, and marital status. GHM define 'clusters' as country-wave groups of observations. GHM's regressions include 5 'region' indicator variables (R): West (Europe, North America, Oceania); Latin America; Asia; Middle East and North Africa; and Sub-Saharan Africa. Country and wave dummy variables are not included in their core specification.⁵

GHM recognize several issues regarding their core specification. First, since SWL scales are bounded (with '10' as the maximum and '1' as the minimum categories), responses may be heaped

⁴ WVS data are collected in waves which can cover multiple adjacent years, e.g. wave 2 covered 1990–94. We match our non-WVS variables for each country to the specific year in which the country was surveyed.

⁵ Country dummy variables are included in the first of our robustness checks discussed below and are included in one of GHM's specifications.

at the top of the distribution (or at the bottom, though GHM do not look at this aspect). Second, use of the *SD* to summarize SWL inequality implicitly assumes that latent ('true') SWL as measured by the manifest data is a cardinal and linear measure (for instance, the change in actual wellbeing when moving from recording a '6' to a '7' is identical to that when moving from a '7' to an '8'). However, latent SWL might be better represented by a non-linear scale and it may be inappropriate to treat observed SWL responses as cardinal. Instead analysis should treat SWL as a purely ordinal measure.

Consequently, GHM undertake a set of robustness checks of their analysis based on the core specification. They examine the effects of: (i) adding as a regressor an interaction between SWL inequality (SD_{ijt}) and a variable summarizing whether an individual thinks inequality is too high; (ii) fitting the core specification using a non-linear model (assuming SWL follows a logistic distribution) rather than a linear model; and (iii) substituting the *SD* SWL inequality index with the 'ordinal variation ratio' index, which is the proportion of observations not equal to the mode.⁶

GHM apply these checks to three cross-country data sources – the European Social Survey (ESS), the World Values Survey (WVS), and the Gallup World Poll – plus the USA-specific Gallup-Healthways Well-Being Index, for which there are separate data by state. The analysis in the current paper uses only WVS data because the Gallup data are not easily accessible and the WVS have world-wide coverage unlike the ESS.

The SWL question asked in the WVS is: '*All things considered, how satisfied are you with your life as a whole these days?*'. Respondents provide an answer on an integer-valued scale running from 1 to 10, with 1 labelled 'completely dissatisfied' and 10 'completely satisfied'.

GHM's key findings using the WVS data include: (a) the SWL *SD* has a consistent negative association with individual-level SWL ($\widehat{\beta}_1 < 0$); (b) in the core specification this association nullifies any effect of income inequality on individual SWL ($\widehat{\beta}_2 \approx 0$); (c) the association between SWL inequality and individual-level SWL is greater for those who consider that inequality is too high (although in that specification – which is the only specification in which GHM include country fixed effects – income inequality becomes an additional significant determinant); (d) SWL inequality remains a significant predictor of social trust when mean SWL is included as a regressor in addition to the *SD*; (e) the association between SWL inequality and individual-level SWL is consistent across linear and logistic models; and (f) the 'ordinal variation ratio' remains a significant determinant of SWL when interacted with "thinks inequality is too high", although the main effect of the SWL dispersion variable is no longer significant (and income inequality is significant in that case). GHM's estimates derived from the core specification (1) are reproduced in our Appendix Table A1. GHM find similar conclusions using their other three datasets.

⁶ GHM also test whether addition of mean SWL impacts the effect of SWL dispersion on social trust.

Although GHM’s sensitivity checks (a)–(f) suggest that the negative association between individual SWL and their country’s SWL inequality is robust, there are reasons to question this conclusion. These are discussed further in the next section but, in short, they relate to the reservations that GHM expressed. Whether the SWL measure should be treated as being cardinal or ordinal is at the root of potential concerns over the GHM findings. If the SWL measure is ordinal, then better indices than GHM’s ordinal variation ratio are available. One family of indices, due to Cowell and Flachaire (CF, 2017), also enables us to investigate the role of skewness of the SWL distribution in affecting individual-level SWL outcomes. We also have some econometric concerns. Specifically, GHM exclude country fixed effects from all but one specification, and wave (i.e. time) fixed effects from all specifications. Without controlling for the effects of time-invariant unobservable factors within countries, or common shocks across countries, conclusions may not be robust.

3 Inequality indices for ordinal data

SWL measures based on integer-value Likert-like scales are inherently ordinal. A key characteristic of ordinal data is that the order of values is well-defined but the magnitudes of differences between scale levels are not. To compare SWL distributions, one must assume that there is a common and fixed ‘reporting function’, across time for the same individual and across individuals at a given time, so that the intrinsic meaning of the SWL scale levels is the same for all individuals and does not change over time. However, even assuming a common and fixed reporting function – as virtually all SWL researchers do, and we do too – the ordinal nature of the SWL scale means that there are no grounds for saying that the difference between ‘very satisfied’ and ‘fairly satisfied’ is of the same magnitude as the difference between ‘fairly satisfied’ and ‘unsatisfied’. To do this, and thereby make SWL responses cardinally measurable, we need to make additional assumptions about the scale.⁷

As noted by GHM, even if one assumes cardinality, the bounded nature of SWL scales may cause ‘true’ SWL inequality to be under-estimated by the *SD*. Delhey and Kohler (2011) address this issue, proposing an index that standardized the *SD* by the maximum possible value. DK note, assuming cardinality, that the maximum standard deviation possible for a given mean SWL (SD_{MAX}) is:

$$(2) \quad SD_{MAX} \equiv \sqrt{\frac{(U - \mu)(\mu - L)N}{N - 1}}$$

where U is the SWL scale’s maximum value, L is the minimum value, μ is mean SWL, and N is the number of observations in the analysis sample. SD_{MAX} reaches its greatest value when μ is midway

⁷ Bond and Lang (2019) discuss these issues in detail. For a more positive interpretation regarding what can be inferred, see Kaplan and Zhuo (2019).

between U and L , and decreases as μ approaches either U or L . DK's 'instrument-effect-corrected' index is:

$$(3) \quad SD^* = \frac{SD}{SD_{MAX}}$$

where SD is the standard deviation defined earlier. Hence SD^* is inflated relative to SD as μ approaches either U or L . We use SD^* as one of several alternatives to SD in our empirical work.

Using a calibrated example, Delhey and Kohler (2011) argue their measure 'works'. But their approach remains based on an assumption of cardinality and so only corrects for one potential issue regarding SD .

A related fundamental issue with the SD and other indices based on mean SWL is that rankings of distributions of ordinal variables based on them are not robust to changes in the scale ('scale dependent'). Scale dependence also implies that the mean is not a robust measure of central tendency for ordinal data such as these.⁸ On these issues, see Allison and Foster's (2004) pioneering article.

Allison and Foster (2004) argue instead that, for ordinal data applications, the median should be used to measure central tendency and be the fundamental reference point underpinning inequality indices. Allison and Foster's concept of S -dominance encapsulates this idea: for two SWL distributions X and Y with cumulative distribution functions F_X and F_Y respectively, and the same median k , distribution Y S -dominates distribution X if and only if: (a) for all categories below the median, $F_Y \leq F_X$; and (b) for all categories at or above the median category, $F_Y \geq F_X$. Building on this idea, Allison and Foster develop an inequality index (AF) which equals the difference between the mean SWL response for median-and-above categories minus the mean SWL response for below-median categories. For a linear integer scale, AF is the average across respondents of the number of scale points required to move from the observed SWL value to the median. Allison and Foster (2004) argue that AF has intuitive appeal but acknowledge that it is scale-dependent and so recommend checking the robustness of results to changes in the scale.⁹

Scale-independent inequality indices appropriate for ordinal data have been developed using axiomatic approaches by, for example, Abul Naga and Yalcin (2008), Apouey (2007), and Cowell and Flachaire (2017). All three classes of inequality index allow for different degrees of inequality aversion when summarizing dispersion across individuals. However, only the Cowell and Flachaire (CF, 2017) approach also allows other features to be accounted for straightforwardly, notably skewness, as we explain shortly. For this reason, we focus on a portfolio of CF indices in this paper (in addition to the SD , SD^* , and AF indices already mentioned).

⁸ For instance, if one applies a transformation to the numerical labels associated with the categories, e.g. using an exponential function, the mean changes. For further discussion of this point, see Bond and Lang (2019) and Jenkins (2019).

⁹ See Dutta and Foster (2013) for an application to US trends in the inequality of happiness.

Derivation of CF indices has three components: definition of each individual's 'status' (which is related to SWL), specification of the reference status value (inequality is conceptualized in terms of differences in individuals' status from the reference value), and assumptions about inequality aversion (how differences in status from the reference value are aggregated).

Status is the position of a person in the distribution of an ordinal variable (e.g. SWL) within society. CF distinguish between downward-looking status (i.e. measuring your position by the proportion of people with lower or the same SWL than you) and upward-looking status (measuring your position by the proportion of people with the same or higher SWL than you). For each of these concepts, status can be either peer-inclusive or peer-exclusive. The peer-inclusive status measures are as explained in the previous sentence; peer-exclusive measures exclude the people with the same SWL response as you when calculating the proportions. In this paper, we focus on peer-inclusive measures of status as they are what CF worked with, and so too have all other researchers to date.

In order to make inequality comparisons, we need both a definition of equality and a method of characterizing departures from equality. Equality is the situation in which all individuals have the same status; for peer-inclusive measures, CF argue persuasively that the natural reference point is the maximum value of status.

Having defined status and the reference point, CF summarize inequality by aggregating across individuals the 'distances' between each person's status and the status reference point. There are different ways of characterizing 'distance';¹⁰ CF do so using a one-parameter family of distance functions suitable for ordinal data, for which the parameter α represents the degree of inequality aversion – the extent to which large distances are given greater weight than small distances in the aggregation across individuals. CF's peer-inclusive downward-looking inequality index, I_α , is shown in (5) where s_i is person i 's status, being the proportion of people with the same or lower response for SWL as person i , and N is the number of people surveyed:

$$(5) \quad I_\alpha = \frac{1}{\alpha(\alpha - 1)} \left[\frac{1}{N} \sum_{i=1}^N s_i^\alpha - 1 \right], 0 < \alpha < 1.$$

$$I_0 = -\frac{1}{N} \sum_{i=1}^N \log(s_i).$$

CF's peer-inclusive upward-looking inequality index is defined analogously with s_i redefined to be the proportion of people with the same or higher response for SWL as person i . We denote the downward-looking versions of the indices I_α^D and the upward-looking versions I_α^U . All of CF's indices

¹⁰ For cardinal variables, standard inequality indices summarize inequality by aggregating across individuals the distances between each person's value and the mean. But, as explained earlier, the mean is an inappropriate reference point when measuring the inequality of ordinal data.

are scale-independent. For a given distribution, a larger α – greater inequality aversion – results in a higher index value.

A useful feature of CF's upward- and downward-looking indices is that they differ in their sensitivity to the skewness of the SWL distribution. For given α , $I_\alpha^D = I_\alpha^U$ if the distribution in question is symmetric. When the distribution is negatively skewed (i.e. has a long left tail), $I_\alpha^D > I_\alpha^U$, whereas $I_\alpha^D < I_\alpha^U$ when the distribution is positively skewed (a long right tail). Inequality aversion also plays a role: for asymmetric distributions, the difference between I_α^D and I_α^U is smaller, the larger that α is. The practical lesson, which we exploit below, is that downward- and upward-looking CF inequality indices used in combination allow one to take account of distributional skewness in addition to inequality per se.¹¹ Because findings are potentially sensitive to the choice of the inequality aversion parameter, we use values of α from across the full range (0, 0.5, and 0.9).

4 Data

We downloaded data on life satisfaction (SWL), employment status, education, marital status, gender and age from the WVS website (Inglehart et al., 2014), consistent with GHM's WVS application. There is information from 6 waves and 100 countries with a total of 348,532 individual-level observations.¹²

Like GHM, we obtained data for GDP per capita (real, PPP terms) and for the Gini coefficient of income from the World Bank's World Development Indicators dataset.¹³ We used linear interpolation to impute values that were missing (as GHM did). Our estimates from these data are very similar to those obtained by GHM for corresponding model specifications (see Section V).

CF also use WVS data (from wave 5) to illustrate their inequality indices in action. We use the data from this wave to calculate alternative inequality indices (using the *ineqord*

¹¹ We illustrate the points made in this paragraph in Appendix Table A2, reproducing and extending the examples presented in CF's Tables 2 and 3. CF focus on the case $\alpha = 0$; we also consider the cases of $\alpha = 0.5$ and $\alpha = 0.9$ (noting that $\alpha < 1$). In Table A2, Case 0 is negatively skewed, Case 3 is positively skewed, and Cases 1 and 2 are symmetric.

¹² To be consistent with GHM, we include countries which have data for only one wave. Identification of regression coefficients on country-level variables such as SWL inequality comes solely from countries with more than one wave of data, but all observations contribute to identifying coefficients on individual-level variables such as age. We lose 13% of WVS observations owing to incomplete data, so our sample size is 302,919.

¹³ We retrieved World Bank World Development Indicators GDP per capita in constant prices (Purchasing Power Parity, 2011 international dollars) from [https://datacatalog.worldbank.org/search?search_api_views_fulltext_op=AND&query=GDP%20Per%20Capita,%20PPP%20\(Constant%202011%20International%20\\$\)&nid=&sort_by=search_api_relevance&sort_order=DESC](https://datacatalog.worldbank.org/search?search_api_views_fulltext_op=AND&query=GDP%20Per%20Capita,%20PPP%20(Constant%202011%20International%20$)&nid=&sort_by=search_api_relevance&sort_order=DESC), and World Bank World Development Indicators Gini data from https://data.worldbank.org/indicator/si.pov.gini?fbclid=IwAR07jKjrl6Ual3tsnIY4d7dggJV4_3eRTfzj0a1BFBCBeZbMN1lp-x7mBeg. Data for the Gini coefficient in New Zealand were missing from World Development Indicators; we obtained estimates from the OECD database.

software of Jenkins, 2020), and to investigate the bivariate relationships between the nine SWL inequality indices discussed in Section III: SD , SD^* , AF , and six CF indices: I_0^D , I_0^U , $I_{0.5}^D$, $I_{0.5}^U$, $I_{0.9}^D$, $I_{0.9}^U$. (In subsequent sections we use all 6 waves of the WVS data.)

Table 1 presents Spearman rank correlation coefficients for the nine indices. All correlations are positive and in some cases are close to one. For instance, the AF ranks countries almost identically to SD , while SD^* is also closely correlated with SD .

As inequality aversion increases, the country rankings for CF indices become more similar to the rankings produced by the other three indices, possibly because skewness is de-emphasized in the CF indices as α increases (and the other indices are insensitive to skewness). One of the lowest rank correlations is between I_0^D and I_0^U , 0.32, indicating that the upward- and downward-looking indices measure different aspects reflecting skewness (recalling that $I_\alpha^D = I_\alpha^U$ for symmetric distributions). As α increases, the country rankings from the upward- and downward-looking indices become very similar (correlation 0.99 for $\alpha = 0.9$), again indicating that skewness is more relevant when inequality aversion is relatively low.

TABLE 1
Rank correlations between 9 SWL inequality indices

	SD	SD^*	AF	I_0^U	I_0^D	$I_{0.5}^U$	$I_{0.5}^D$	$I_{0.9}^U$	$I_{0.9}^D$
SD	1.00								
SD^*	0.88	1.00							
AF	0.93	0.71	1.00						
I_0^U	0.50	0.14	0.55	1.00					
I_0^D	0.53	0.44	0.42	0.32	1.00				
$I_{0.5}^U$	0.70	0.39	0.69	0.93	0.55	1.00			
$I_{0.5}^D$	0.78	0.57	0.70	0.70	0.83	0.88	1.00		
$I_{0.9}^U$	0.78	0.53	0.75	0.85	0.66	0.98	0.95	1.00	
$I_{0.9}^D$	0.80	0.56	0.75	0.81	0.71	0.96	0.97	0.99	1.00

Notes. Based on data for 58 countries from WVS wave 5. The inequality indices are explained in the main text.

5 The relationship between SWL and inequality

Table 2, columns (1)–(3) report WVS-based OLS estimates (with standardized beta coefficients) for the same model specifications as GHM report in columns (1)–(3) in their Table 3 (reproduced as Appendix Table A1). We refer to columns (1)–(3) as GHM’s base specifications since they do not

include the extra terms relating to preference for equality.¹⁴ Our dataset includes a slightly larger number of observations than GHM's (302,919 compared with their 243,875 to 271,677), reflecting extensions to the GDP and Gini series enabled by our use of updated data sources. Nevertheless, our results are very similar to those of GHM.

SD has an estimated coefficient of -0.15 using our sample compared with -0.17 in GHM's (significant at $p < 0.001$ in each case). The coefficients on per capita GDP are also almost identical for the two samples (both significant at $p < 0.001$) while the income Gini coefficient is insignificantly different from zero in both cases. Thus, the two samples produce almost identical results.

The GHM base specifications include region indicators (defined earlier). These control for time-invariant unobservable differences across regions but not within regions and there are many countries within each region. Given the wide variability of country conditions, cultures, and experiences within some of GHM's regions, the regional level of aggregation may be inadequate for controlling for country-specific characteristics.¹⁵

TABLE 2
Estimates from extended WVS specifications (standardised beta coefficients)

	Dependent variable: SWL (1–10 scale)				
	(1)	(2)	(3)	(4)	(5)
SWL <i>SD</i>	-0.147^{***} (0.027)		-0.151^{***} (0.027)	-0.048 (0.034)	0.002 (0.036)
Income Gini coefficient		-0.018 (0.033)	0.027 (0.030)	-0.189^{***} (0.061)	-0.132^{**} (0.054)
Log(GDP per capita)	0.198^{***} (0.033)	0.276^{***} (0.029)	0.196^{***} (0.033)	0.202^{***} (0.059)	0.045 (0.064)
Individual controls	Yes	Yes	Yes	Yes	Yes
Region dummies	Yes	Yes	Yes	No	No
Country dummies	No	No	No	Yes	Yes
Wave dummies	No	No	No	No	Yes
No. of observations	302,919	302,919	302,919	302,919	302,919
R^2	0.106	0.092	0.106	0.157	0.160

Notes. Estimates are weighted using the WVS-supplied sample weights. Countries are weighted equally. Cluster-robust standard errors in parentheses (clusters are country-wave combinations). Individual-level controls comprise sex, age, age squared, education (dummies), and marital status (dummies). Statistical significance indicators: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

¹⁴ When we fitted models that included "Thinks inequality is too high" as a regressor, our estimates differed somewhat from those presented by GHM. Hence, we restrict our focus here to model specifications that we replicate. This also enables us to compare more easily results based on different SWL inequality indices.

¹⁵ For instance, Chile and Venezuela are included in Latin America, and Albania and Canada in the West.

Table 2, column (4) displays estimates of the base specification (including both SWL and income inequality measures) with country dummies replacing region dummies. The inclusion of country dummies means that identification of the effects of the three focal variables presented in Table 2 comes from changes across time within countries. The coefficient on per capita GDP retains its size and significance but the relative importance of SWL inequality and income inequality switches: the coefficient on the Gini is now negative and significant, while *SD* is no longer significant (albeit remaining negative).

It is possible that SWL varies globally across time for reasons unrelated to the observable variables included in our specification. For instance, wars or pandemics may cause a global reduction in SWL. To investigate this potential effect, Table 2 column (5) displays estimates from a specification that adds wave indicators in addition to country indicators. The Gini coefficient remains negative and significant (at the 5% level) while per capita GDP is no longer significant (while remaining positive), potentially because the wave indicators capture the impact of globally-increasing incomes over time. The coefficient on *SD* is now virtually zero.

One interpretation of the estimates from these extended specifications is that the significant negative impact of *SD* that we find for specifications without country indicators reflects cross-national differences in SWL inequality – which reflect unobservable country circumstances rather than within-country SWL inequality *per se*. If this were the case, changes in SWL inequality within each country do not necessarily impact on individual SWL.

Another possibility is that OLS estimation of the relationship is inappropriate since the dependent variable (SWL) is ordinal rather than cardinal. Henceforth, we concentrate on ordered logit (OLogit) estimates because they respect the ordinality of the dependent variable. Tables 3 and 4 present estimates from OLogit models using regressor specifications corresponding to columns (3) and (5) respectively of Table 2. Each table includes estimates based on nine (standardized) SWL inequality indices. The Table 3 models include only region indicators (as in GHM), whereas Table 4 models include both country and wave indicators. Column (1) of each table refers to the same regressor variables as the corresponding specification in Table 2 but estimates are derived using OLogit rather than OLS. Each table reports log odds coefficients and standard errors are again clustered by country-wave combination.

The estimated coefficients for the income Gini and for per capita income in Table 3 (models with region indicators) are similar regardless of which index is used to summarize SWL inequality. The coefficient on the income Gini is always insignificant and the coefficient on per capita GDP is always significantly positive.

The results for six of the alternative SWL inequality indices are similar to those derived using *SD* (a significant negative relationship with individual SWL); *SD** also has a negative, but insignificant, coefficient. Perhaps counter-intuitively, however, we estimate the coefficient on I_0^D to be

(significantly) positive. We return to this finding when interpreting results of Tables 3 and 4 together. The explanatory power when I_0^U is the SWL inequality index is higher than when any of the other SWL inequality measures are used (with a pseudo- R^2 of around 0.07 rather than 0.06).

In Table 4 (models with country and wave indicators), we find that in each case the coefficient on the income Gini is negative and significant; thus an increase in income inequality within countries is associated with lower individual SWL. Per capita GDP remains positive but is insignificant in each case, with the wave dummies likely to be capturing the effects of globally increasing incomes over time.

We find no association between individual SWL and each of the SD , AF , $I_{0.5}^D$, $I_{0.9}^U$ or $I_{0.9}^D$ measures of SWL inequality. SD^* has a positive and statistically significant association with individual SWL but, since it is based on a cardinality assumption, we do not place much emphasis on this result other than noting it as a counter-example to that obtained from using the (unadjusted) SD .

I_0^U and I_0^D both yield significant estimates, as does $I_{0.5}^U$. The upward-looking indexes, I_0^U and $I_{0.5}^U$, have a similar relationship to that found by GHM (using SD): increased SWL inequality within a country is associated with a decline in individual SWL. As in Table 3, I_0^U delivers (marginally) the highest explanatory power of any of the SWL inequality measures. With country and wave fixed effects, we again find a positive relationship (significant at the 1% level) between individual SWL and SWL inequality within a country when the latter is measured by I_0^D . A similar pattern of oppositely signed coefficients occurs for $\alpha = 0.5$. We investigate why the upward and downward looking CF indices provide apparently conflicting results in the next section.

6 Upward- versus downward-looking inequality indices and a role for skewness

We now take advantage of the distinctive features of the CF approach to inequality measurement. As pointed out in Section III, upward-looking and downward-looking CF indices provide different estimates (for given α) unless the distribution of responses is symmetric. This observation helps us to interpret our findings in Tables 3 and 4 for I_0^U and I_0^D , noting that skewness has a greater differentiating impact on the CF indices for low α .

$I_0^U > I_0^D$ when positive skewness is present, i.e. the mass of individuals is concentrated towards lower categories coupled with a right tail of high-SWL individuals (as illustrated by Case 3 in Appendix Table A2). Interpreting this pattern in the manner of GHM and given the negative coefficient on I_0^U in Tables 3 and 4, one might conclude that individuals' SWL is lowered in a country in which there is an over-representation of low SWL people relative to a symmetric distribution. In effect, there is a negative externality from observing a preponderance of unhappy individuals in society.

$I_0^D > I_0^U$ when negative skewness is present, i.e. when the mass of individuals is concentrated towards higher categories coupled with a left tail of low-SWL individuals (as illustrated by Case 0 in Appendix Table A2). Given the positive coefficient on I_0^D in Tables 3 and 4, one might conclude that individuals' SWL is raised in a country in which there is an over-representation of high SWL people relative to a symmetric distribution: there is a positive externality from observing a preponderance of happy individuals in society. This interpretation is consistent with the 'tunnel effect' observed in transition countries in which envy of others' success is replaced by a positive demonstration effect that individuals have the scope to improve their lot when they see others succeed (Ravallion and Lokshin, 2000).

We investigate these conjectures further by including both I_0^U and I_0^D in the same regression equation (a specification that also includes country and wave indicators). For completeness, we also employ specifications including either $I_{0.5}^U$ and $I_{0.5}^D$ or $I_{0.9}^U$ and $I_{0.9}^D$. The estimates are presented in columns (1)–(3) of Table 5. In each case, the upward-looking index has a statistically significant negative coefficient and the downward-looking index has a statistically significant positive coefficient. Thus, the same patterns arise regardless of the degree of inequality aversion. Furthermore, the coefficient on the income Gini is significantly negative in each specification.

The externality explanation is appealing – and consistent with the hypothesis put forward by GHM – but another possibility is that the statistically significant (and contrasting) estimates for the pairs of CF indices are simply due to an arithmetic association between individual SWL and the indices in the presence of time-varying unobservable factors that affect SWL for some individuals. For instance, suppose some time-varying unobservable factor results in the SWL of some individuals being reduced so falling into the lower tail of the distribution. This will increase measured inequality, particularly as measured by I_α^U .¹⁶ The correlation of a higher I_α^U with reduced SWL for some individuals will result in an estimated negative coefficient on I_α^U in the SWL regression. In the opposite case, in which the SWL of some individuals increases so falling into the upper tail of the SWL distribution, we would again see an increase in measured inequality, particularly as measured by I_α^D .¹⁷ The correlation of a higher I_α^D with raised SWL for some individuals will result in an estimated positive coefficient on I_α^D in the SWL regression.

¹⁶ For instance, if the proportions for SWL levels 1 to 4 in Case 0 of Table A2 (0.25, 0.25, 0.5, 0) are replaced by (0.26, 0.26, 0.48, 0), representing a more pronounced left tail, I_0^U increases by 0.0121 whereas I_0^D increases by just 0.0004.

¹⁷ If the proportions for SWL levels 1 to 4 in Case 3 of Table A2 (0, 0.5, 0.25, 0.25) are replaced by (0, 0.48, 0.26, 0.26), representing a more pronounced right tail, I_0^D increases by 0.0121 whereas I_0^U increases by just 0.0004.

TABLE 3
The relationship between SWL and SWL inequality, by inequality index (ordered logit estimates, with region dummies)

SWL Inequality index:	Dependent variable: SWL (1–10 scale)								
	<i>SD</i>	<i>SD</i> *	<i>AF</i>	I_0^U	I_0^D	$I_{0.5}^U$	$I_{0.5}^D$	$I_{0.9}^U$	$I_{0.9}^D$
SWL Inequality	−0.253*** (0.0540)	−0.0586 (0.0553)	−0.319*** (0.0525)	−0.624*** (0.0970)	0.182*** (0.0568)	−0.457*** (0.0978)	−0.124* (0.0737)	−0.322*** (0.0858)	−0.266*** (0.0815)
Income Gini coefficient	0.0604 (0.0592)	0.00503 (0.0643)	0.0640 (0.0573)	−0.0416 (0.0404)	−0.0363 (0.0582)	0.00622 (0.0480)	0.00322 (0.0598)	0.0206 (0.0538)	0.0197 (0.0560)
Log(GDP per capita)	0.359*** (0.0629)	0.465*** (0.0597)	0.330*** (0.0630)	0.251*** (0.0531)	0.471*** (0.0512)	0.316*** (0.0636)	0.463*** (0.0603)	0.376*** (0.0643)	0.403*** (0.0637)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country dummies	No	No	No	No	No	No	No	No	No
Wave dummies	No	No	No	No	No	No	No	No	No
No. of observations	302,919	302,919	302,919	302,919	302,919	302,919	302,919	302,919	302,919
Pseudo- R^2	0.0597	0.0573	0.0607	0.0695	0.0584	0.0644	0.0577	0.0610	0.0598

Notes. See Table 2. Reported coefficients are for log odds.

TABLE 4

The relationship between SWL and SWL inequality, by inequality index (ordered logit estimates, with country and wave dummies)

SWL Inequality index:	Dependent variable: SWL (1–10 scale)								
	<i>SD</i>	<i>SD</i> *	<i>AF</i>	I_0^U	I_0^D	$I_{0.5}^U$	$I_{0.5}^D$	$I_{0.9}^U$	$I_{0.9}^D$
SWL Inequality	0.0573 (0.0706)	0.131** (0.0596)	-0.0100 (0.0543)	-0.369*** (0.0926)	0.220*** (0.0725)	-0.156** (0.0627)	0.100 (0.0743)	-0.0254 (0.0615)	0.0179 (0.0650)
Income Gini coefficient	-0.276** (0.109)	-0.286*** (0.104)	-0.262** (0.111)	-0.206** (0.0865)	-0.230** (0.0973)	-0.249** (0.105)	-0.261** (0.108)	-0.262** (0.110)	-0.264** (0.111)
Log(GDP per capita)	0.0731 (0.125)	0.0845 (0.126)	0.0806 (0.126)	0.0898 (0.101)	0.0247 (0.108)	0.0947 (0.121)	0.0623 (0.123)	0.0830 (0.126)	0.0771 (0.126)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region dummies	No	No	No	No	No	No	No	No	No
Country dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No. of observations	302,919	302,919	302,919	302,919	302,919	302,919	302,919	302,919	302,919
Pseudo- R^2	0.0747	0.0749	0.0747	0.0758	0.0755	0.0749	0.0748	0.0747	0.0747

Notes. As for Table 3.

TABLE 5

The relationship between SWL and SWL inequality including upward-looking and downward-looking CF measures with a measure of skewness (ordered logit estimates, with country and wave dummies)

	Dependent variable: SWL (1–10 scale)					
	(1) $\alpha = 0$	(2) $\alpha = 0.5$	(3) $\alpha = 0.9$	(4) $\alpha = 0$	(5) $\alpha = 0.5$	(6) $\alpha = 0.9$
I_{α}^U	-0.623*** (0.0285)	-1.091*** (0.0617)	-5.167*** (0.348)	-0.518*** (0.0438)	-0.757*** (0.104)	-2.488*** (0.572)
I_{α}^D	0.423*** (0.0280)	0.902*** (0.0585)	5.008*** (0.343)	0.351*** (0.0393)	0.625*** (0.0934)	2.412*** (0.559)
SWL skewness (tau-3)				-1.331*** (0.374)	-2.153*** (0.506)	-3.413*** (0.584)
Income Gini coefficient	-0.105** (0.0441)	-0.142*** (0.0537)	-0.161*** (0.0604)	-0.0875** (0.0429)	-0.107** (0.0511)	-0.100* (0.0542)
Log(GDP per capita)	-0.00867 (0.0620)	0.0294 (0.0712)	0.0501 (0.0782)	-0.00445 (0.0577)	0.0275 (0.0656)	0.0378 (0.0699)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Region indicators	No	No	No	No	No	No
Country indicators	Yes	Yes	Yes	Yes	Yes	Yes
Wave indicators	Yes	Yes	Yes	Yes	Yes	Yes
No. of observations	302,919	302,919	302,919	302,919	302,919	302,919
Pseudo- R^2	0.0782	0.0777	0.0774	0.0782	0.0778	0.0776

Notes. As for Table 3.

In each of the cases just outlined, the skewness of the SWL distribution alters as a result of the changes in unobservable factors. In order to control for this effect, the specifications shown in columns (4)–(6) of Table 5 add a measure of skewness as a regressor. The measure of skewness that we adopt is the tau-3 index which is based on an L-moments approach (Cox, 2013).¹⁸

In each case the coefficient on the skewness index is negative and statistically significant at the 1% level, indicating that an increase in skewness towards lower categories is associated, *ceteris paribus*, with a worsening in individual SWL consistent with GHM's intuition (though their study did not explicitly examine skewness). Even after controlling for skewness, we nevertheless continue to find a significant negative association of I_{α}^U with individual SWL and a significant positive association of I_{α}^D with individual SWL (albeit with a reduction in the point estimates once skewness is accounted for).

We tested the robustness of the estimated associations further by splitting the sample according to whether the country-wave median SWL > mean SWL, and vice versa. The results are unaltered: I_{α}^U is always significantly negative and I_{α}^D always significantly positive (at the 5% level) for each sample split, and with tau-3 always negative (significantly so in 5 of the 6 cases).

Our estimates indicate that the associations of individual SWL with both upward-looking and downward-looking SWL inequality hold after controlling for income inequality and the skewness of a country's SWL distribution. The implications are that having a proportionately large group of people which is low in the SWL distribution (i.e. high I_{α}^U relative to I_{α}^D) is associated, *ceteris paribus*, with lower individual SWL whereas having a large group of people which is high in the SWL distribution (high I_{α}^D relative to I_{α}^U) is associated with higher individual SWL; this association holds even after controlling for individuals' personal characteristics and other country characteristics. Our results are consistent with the idea that individuals experience negative externalities from seeing large groups of fellow citizens with low SWL and experience positive side-effects from seeing large groups of fellow citizens with high SWL. This explanation provides a richer interpretation of the effects of SWL inequality than is accorded by GHM's analysis. However, this interpretation remains a conjecture, and investigation into other potential explanations is warranted. So too is analysis based on other data sets in addition to the WVS.

7 Summary and conclusions

It is plausible that societal inequality impacts on individuals' wellbeing through negative externalities that arise from seeing others faring poorly. It is also plausible that positive externalities may arise from observing people faring well. Prior research has shown an association between income inequality and individual wellbeing. More recently, the contribution of GHM indicates an

¹⁸ Estimating the models using a measure of skewness equal to the median minus the mean led to very similar results.

association between individual wellbeing and SWL inequality. However, the measure of SWL inequality on which GHM focus (SD) is not an appropriate inequality measure for ordinal data, and SWL is an ordinal variable.

We analyze the link between individual wellbeing and SWL inequality using inequality measures that have been derived specifically for use with ordinal data. We also extend GHM's econometric specifications. The models that are extended to include country and wave fixed effects do not uphold GHM's findings of a negative association between SWL inequality and individual SWL when using their preferred measure of wellbeing inequality (the SD). Some of the alternative inequality measures that we employ, however, do provide support to GHM's finding that greater SWL inequality contributes negatively to individuals' satisfaction with life. These include the upward-looking peer-inclusive CF inequality indices incorporating low to moderate degrees of inequality aversion (I_0^U and $I_{0.5}^U$). Paradoxically, when a high degree of aversion to inequality ($\alpha = 0.9$) is assumed, the relationship between I_α^U and SWL evaporates; but that could be because such a high parameter value does not appropriately represent society's preferences. More paradoxical still is the finding that CF's downward-looking peer-inclusive inequality index indicates the opposite relationship: greater SWL inequality is associated with higher individual wellbeing.

We leverage a property of the CF family of indices to show that skewness of the SWL distribution matters for individual SWL in addition to inequality. Nevertheless, when we control for skewness, we still find that the upward- and downward-looking measures of inequality have opposing effects on individual SWL.

Our estimates are consistent with a hypothesis that individuals experience negative externalities when they see many of their fellow citizens having low SWL, whereas they experience positive side-effects when they see many citizens having high SWL. This explanation is certainly plausible. The negative externalities may arise from altruistic components to the utility function or from negative societal effects (e.g. crime) arising from people with predominantly low SWL. The positive externalities may arise from a tunnel effect when observing people doing well, so providing an aspirational example, and/or from their provision of beneficial amenities for the wider population.

Although these explanations are plausible, as are GHM's inequality findings, they need to be tested further lest other (possibly time-varying unobservable) factors are responsible for the associations that we find. In extending the research in this direction, a key requirement will be to adopt inequality (and skewness) measures, and estimation techniques, that are appropriate for ordinal data. Our analysis shows that the practice of applying an inequality index developed for cardinal data is inappropriate when dealing with ordinal data. We can learn a lot more by using the appropriate tools.

References

- Abdul Naga, R. and T. Yalcin. "Inequality Measurement for Ordered Response Health Data." *Journal of Health Economics*, 27, 2008, 1614–25.
- Allison, R.A. and J. E. Foster. "Measuring Health Inequality using Qualitative Data." *Journal of Health Economics*, 23, 2004, 505–24.
- Apouey, B. "Measuring Health Polarization with Self-Assessed Health Data." *Health Economics*, 16, 2007, 875–94.
- Bond, T. N. and K. Lang. "The Sad Truth about Happiness Scales." *Journal of Political Economy*, 127 (4), 2019, 1629–40.
- Cowell, F. A. and E. Flachaire. "Inequality with Ordinal Data." *Economica*, 84, 2017, 290–321.
- Cox, N. J. "Lmoments: Stata Module to Generate L-Moments and Derived Statistics." Statistical Software Components S341902, Boston College Department of Economics, 2013.
- Delhey, J. and U. Kohler. "Is Happiness Inequality Immune to Income Inequality? New Evidence through Instrument-Effect-Corrected Standard Deviations." *Social Science Research*, 40, 2011, 742–56.
- Dutta, I. and J. E. Foster. "Inequality of Happiness in the US: 1972–2010." *Review of Income and Wealth*, 59, 2013, 393–415.
- Goff, L., J. F. Helliwell, and J. Mayraz. "Inequality of Subjective Well-Being as a Comprehensive Measure of Inequality." *Economic Inquiry*, 56 (4), 2018, 2177–94.
- Grimes, A. and D. Wesselbaum. "Moving towards Happiness." *International Migration*, 57 (3), 2019, 20–40.
- Helliwell, J. F., R. Layard, J. D. Sachs, and J.-E. De Neve. *World Happiness Report 2020*. New York: Sustainable Development Solutions Network, 2020.
- Inglehart, R., C. Haerpfer, A. Moreno, C. Welzel, K. Kizilova, J. Diez-Medrano, M. Lagos, P. Norris, E. Ponarin & B. Puranen et al., eds. *World Values Survey: All Rounds - Country-Pooled Datafile Version*: www.worldvaluessurvey.org/WVSDocumentationWVL.jsp. Madrid: JD Systems Institute, 2014.
- Jenkins, S. P. "Better Off? Distributional Comparisons for Ordinal Data about Personal Well-Being." *New Zealand Economic Papers*, online ahead of print, 2019.
<https://www.tandfonline.com/doi/full/10.1080/00779954.2019.1697729>
- Jenkins, S. P. "Ineqord: Stata module to calculate indices of inequality and polarization for ordinal data." Statistical Software Components S458723, Boston College Department of Economics, 2020.
- Kaplan, D. M. and L. Zhuo. "Comparing Latent Inequality with Ordinal Data." Unpublished Paper, 2019, available at <https://faculty.missouri.edu/~kaplandm>.

- Ngamaba, K. A., M. Panagioti, and C. J. Armitage. "Income Inequality and Subjective Well-Being: a Systematic Review and Meta-Analysis." *Quality of Life Research*, 27, 2018, 577–96.
- Ott, J. "Level and Inequality of Happiness in Nations: Does Greater Happiness of a Greater Number imply Greater Inequality in Happiness?" *Journal of Happiness Studies*, 6 (4), 2015, 397–420.
- Ravallion, M. and M. Lokshin. "Who Wants to Redistribute? The Tunnel Effect in 1990s Russia." *Journal of Public Economics*, 76 (1), 2000, 87–104.
- Schneider, S. M. "Income Inequality and Subjective Wellbeing: Trends, Challenges, and Research Directions." *Journal of Happiness Studies*, 17, 2016, 1719–39.
- Stevenson, B. and J. Wolfers. "Happiness Inequality in the United States." *Journal of Legal Studies*, 37 (S2), 2008, S33–S79.
- Veenhoven, R. "Inequality in Happiness: Inequality in Countries Compared Between Countries." Paper presented at the 12th Work Congress of Sociology, Madrid, Spain, 1990.
<https://personal.eur.nl/veenhoven/Pub1990s/90c-full.pdf>

Appendix

TABLE A1
GHM (Table 3) estimates: WVS base specification

	Dependent variable: SWL (1–10 scale)		
	(1)	(2)	(3)
SWL standard deviation	−0.17*** (−5.87)		−0.17*** (−6.01)
Income Gini coefficient		−0.05 (−1.11)	−0.01 (−0.18)
Log(GDP per capita)	0.18*** (5.60)	0.28*** (8.34)	0.21*** (5.81)
Region indicators	Yes	Yes	Yes
No. observations	271,667	243,875	243,875

Notes. Standardized beta coefficients; cluster-robust *t*-statistics in parentheses (clusters defined by country-wave combination). Individual-level controls include sex, age, age squared, education (dummies), and marital status (dummies). Statistical significance indicators: ** $p < 0.01$; *** $p < 0.001$.

TABLE A2
SWL Inequality estimates for 4 hypothetical distributions

	Case 0	Case 1	Case 2	Case 3
SWL level		Population proportions		
1	0.25	0.25	0.0	0.0
2	0.25	0.25	0.5	0.5
3	0.5	0.25	0.5	0.25
4	0.0	0.25	0.0	0.25
Index		Inequality estimates		
I_0^D	0.5199	0.5918	0.3465	0.4185
I_0^U	0.4185	0.5918	0.3465	0.5199
$I_{0.5}^D$	0.7929	0.9269	0.5858	0.7198
$I_{0.5}^U$	0.7198	0.9269	0.5858	0.7929
$I_{0.9}^D$	3.2693	3.9029	2.5784	3.2120
$I_{0.9}^U$	3.2120	3.9029	2.5784	3.2693

Notes. Examples based on CF's Tables 2 and 3. SWL categories 1, 2, 3, 4 in the top panel correspond to CF's categories N, G, E, B respectively; population proportions in each case correspond to those in CF's Table 2. Case 0 is negatively (left) skewed, Cases 1 and 2 are symmetric, and Case 3 is positively (right) skewed. In the lower panel, values reported for I_0^D and I_0^U replicate those in CF's Table 3; $I_{0.5}^D$, $I_{0.5}^U$, $I_{0.9}^D$ and $I_{0.9}^U$ are calculated as indicated in Section III.

Motu



economic & public policy research

for other Motu working papers: www.motu.nz