

Federal Reserve Bank of New York
Staff Reports

World Welfare Is Rising: Estimation Using Nonparametric Bounds on Welfare Measures

Maxim L. Pinkovskiy

Staff Report No. 662
December 2013



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the author.

World Welfare Is Rising: Estimation Using Nonparametric Bounds on Welfare Measures

Maxim L. Pinkovskiy

Federal Reserve Bank of New York Staff Reports, no. 662

December 2013

JEL classification: I31, C02

Abstract

I take a new approach to measuring world inequality and welfare over time by constructing robust bounds for these series instead of imposing parametric assumptions to compute point estimates. I derive sharp bounds on the Atkinson inequality index that are valid for any underlying distribution of income conditional on given fractile shares and the Gini coefficient. While the bounds are too wide to reject the hypothesis that world inequality may have risen, I show that world welfare rose unambiguously between 1970 and 2006. This conclusion is valid for alternative methods of dealing with countries and years with missing surveys, alternative survey harmonization procedures, and alternative GDP series, or if the inequality surveys used systematically underreport the income of the very rich or suffer from nonresponse bias.

Key words: world income distribution, inequality and welfare measures, nonparametric bounds

Pinkovskiy: Federal Reserve Bank of New York (e-mail: maxim.pinkovskiy@ny.frb.org). The author thanks Daron Acemoglu, Tony Atkinson, David Autor, Arun Chandrasekhar, Pierre-André Chiappori, Victor Chernozhukov, Angus S. Deaton, Melissa Dell, Richard Eckaus, Susan Elmes, Jerry Hausman, Horacio Larreguy, James McDonald, Whitney Newey, Benjamin Olken, Adam Sacarny, Emmanuel Saez, Xavier Sala-i-Martin, Bernard Salanie, Paolo Siconolfi, James Snyder, Edward Vytlačil, Michael Woodford, and seminar participants at the MIT Development Lunch and the MIT Labor Lunch for very insightful comments on drafts and presentations of this paper. The author is very grateful to Thomas Piketty and two anonymous referees for extremely useful comments. He also thanks the Paul and Daisy Soros Fellowship for New Americans for intellectual stimulation and the National Science Foundation Graduate Research Fellowship Program and the Institute for Humane Studies for funding. The views expressed in this paper are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

There is a substantial literature on estimating the evolution of the global distribution of income over time to assess whether global poverty and inequality are rising or falling. An important strand of the literature argues that while inequality between countries treated as observations of equal weight is rising, inequality between all people on the globe is falling, as some of the fastest-growing countries (China and India) have initially been among the poorest. This claim has been advanced by, e.g. Schultz (1998), Bhalla (2002), Bourguignon and Morrisson (2002), and most recently Sala-i-Martin (2002a and b, 2006) and Chotikapanich et al. (2007). Similarly, Bhalla (2002), Chen and Ravallion [2001, 2010] and Sala-i-Martin (2002a and b, 2006) document that world poverty has been falling since 1990 (or 1970). There also is an alternative part of the literature (Dikhanov and Ward [2001], Milanovic [2002, 2005, 2012] that contends that global inequality, even if measured between individuals, has increased. The results derived in the literature are often cited to buttress or undermine the contention that the recent period of globalization has been good for the global poor, e.g. Bhalla (2002) or Milanovic (2005).

A common feature of this literature is that all its results are computed using grouped data on within-country inequality, typically obtained from a secondary dataset such as that of Deininger and Squire (1996) or its successor, the World Income Inequality Dataset (WIID). This has been done because data on income distributions, particularly for developing countries, is typically available only through tabulations, quintile shares, or Gini coefficients, with the microdata either not existing (as with very old surveys) or not being available to the public (as with many surveys administered by national statistical agencies, including those of critically important countries such as China).¹ While the World Bank has released some unit record data through its Living Standards Measurement Surveys (LSMS), it is extremely difficult to find, not to mention exploit, a sufficient number and variety of surveys to obtain

¹Reddy and Minoiu (2007) comment on the paucity of unit record data: "The analysis of unit data may be prohibitive in terms of time and manpower, and since unit data may be unavailable for numerous country-years...unit data from nationally representative household surveys for many countries....are not publicly available.

even one survey per country, let alone come close to the degree of coverage provided by a panel dataset like the WIID. The authors writing in the literature on estimating the world distribution of income implicitly or explicitly use parametric distributional assumptions to convert the grouped data into a distribution of income. Bourguignon and Morrisson (2002) and Milanovic (2002) assume that the distribution inside each quintile or decile² is egalitarian. Bhalla (2002) and Chen and Ravallion [2001, 2010] use parametrized Lorenz curves (for instance, Bhalla (2002) uses the World Bank’s Simple Accounting Procedure). Dikhanov and Ward (2001) approximate the income distribution using a polynomial approximation, and Chotikapanich et. al. (2007) fit rich parametric distributions, such as the four-parameter generalized beta distribution, to income data.

However, there is no consensus on what is a good parametric assumption for the distribution of income. The literature on functional forms for income distributions is large and largely inconclusive, going back to Pareto (1897) for the Pareto distribution; Gibrat (1931), Kalecki (1945), Aitchison and Brown (1957), and more recently Lopez and Servén (2006) for the lognormal distribution, Salem and Mount (1974) for the gamma distribution, Singh and Maddala (1976) for the Singh-Maddala distribution, and McDonald and Xu (1995) on the generalized beta family of distributions, nesting all the above. All of the above distributions are unimodal; Zhu (2005) has suggested that empirical income distributions may be multimodal, which opens the door to further candidate distributions.³ Given the number of distributions considered as plausible candidates, it is hard to be particularly confident about any given parametric assumption.

Parametric assumptions have a deeper methodological problem: since they yield point estimates of income distributions, parametric assumptions force the researcher to reach a conclusion on the evolution of the distribution of income. Yet, it may be the case that the

²Hereafter; quintiles, deciles, and other partitions of the income distributions will be referred to as fractiles.

³An alternative approach is to use nonparametric estimators to obtain the world distribution of income strictly from the data. Sala-i-Martin (2002a and b, 2006) uses kernel density estimation on fractile means to obtain estimates for country income distributions, and integrates them to obtain the world income distribution. Such an approach avoids the critique of making arbitrary distributional assumptions that the parametric approach is subject to, and succeeds in obtaining an estimate for the world income distribution as a whole, rather than just of some of its statistics. However, the approach estimates income distributions consistently only as the grouping of the data becomes arbitrarily fine, while in practice, income distribution data is presented in only a few groups (5 or 10). Therefore, there is need for bounds on poverty and inequality measures of the world distribution of income that are valid for any underlying functional form of the individual country distributions.

data are not fine enough to reach any conclusion; there may exist valid income distributions that generate the data, yet imply that the overall distribution of income has widened, and there may also exist equally valid distributions that also generate the data and imply the overall distribution has narrowed. A critical question is: can we know if this is ever the case? Yet more generally: if the functional forms that the literature has been assuming are wrong, it gives no guidance as to what alternative paths global poverty and inequality might have taken, and what paths they most certainly could not have taken. While there have been many proposed time paths of global poverty and inequality based on parametric assumptions, are there any paths of poverty and inequality measures that can be ruled out on the basis of the data we have?

Hence, it is interesting to ask whether we can dispense with parametric assumptions completely. In particular, to analyze a time series of the values of some functional of the world distribution of income (e.g. poverty or inequality), we do not want point estimates, but sharp upper and lower bounds on the value of this series at each point in time, so that there may exist income distributions compatible with the data such that these bounds are attained, but there exist no distributions compatible with the data that imply estimates for the series outside the bounds. Such a pair of sharp bounds would completely summarize what the data implies about the series in question: any series passing outside the bounds would be impossible, whereas any series contained in them would be conceivable.

Another feature of the literature on analyzing the world distribution of income is that it analyzes poverty and inequality separately, and typically reaches normative conclusions without a formal aggregation of the two in some theoretically justified manner. While different ways have been proposed to aggregate growth and inequality into a measure of welfare, no measure has been selected as definitive, and the papers looking at welfare measures are few (Pinkovskiy and Sala-i-Martin [2009], Atkinson and Brandolini [2010]). There has neither been an attempt to provide bounds for the evolution of such a measure based on weak and plausible assumptions. However, the need for such a measure in the normative analysis of the trends in the world distribution of income is crucial. It is clear that the assertion that

inequality has risen does not imply that welfare has fallen for most reasonable notions of welfare; the rise in inequality may have accompanied a Pareto improvement. Similarly, an assertion that poverty has risen need not automatically imply that welfare has decreased, since such a connection would be valid only if the welfare function was concerned exclusively about the poor. Otherwise, rapid growth for a large number of people in other parts of the distribution could (for a suitable welfare function) offset the negative effects of a rise in poverty. Without a well-specified welfare function that is derived from clear axiomatic normative principles, we cannot rigorously weigh the differential benefits of poverty and inequality reduction.

The contributions of this paper are threefold. First, I derive sharp nonparametric bounds for the Atkinson welfare measure (which is commonly calculated and theoretically justified) in terms of the inequality statistics typically made available by statistical agencies: fractile shares and Gini coefficients. The formula for a tight upper bound to any inequality measure when fractile means and boundaries are known is well-known and is presented in the review by Cowell (2000). Cowell (2000) also reviews tightening of the bounds under the assumption that the density of income is monotonic decreasing in a given fractile, while Cowell (1991) provides tight upper bounds in cases in which the fractile boundaries are known, but the fractile means are unknown. However, to my knowledge, there has been no work on deriving the bounds for the Atkinson welfare index using the Gini coefficient, with or without fractile shares.⁴ This problem is useful, since Gini coefficients are often reported by statistical agencies and used in the literature, and deriving bounds for the Atkinson welfare index based on the Gini coefficient allows the researcher to deduce something about the first measure from the second, both for empirical and theoretical purposes. The mathematical problem of deriving the bounds given the richest available data is nontrivial, but important to solve as taking advantage of all the available data substantially decreases the width of the bounds.

⁴The most related part of the literature derives tight upper bounds to the Gini coefficient based on fractile shares. Gastwirth (1972) provides a tight upper bound on the Gini coefficient given fractile means and boundaries, while Murray (1978) provides such a bound when fractile means are unknown, and Mehran (1975) calculates the bound when fractile boundaries are unknown. This paper also belongs to the more general econometrics literature on nonparametric bounds, e.g. Manski (1995).

Second, I compute the welfare bounds for all possible countries and years using traditional GDP and inequality data from the literature: the Penn World Tables (PWT) and the World Income Inequality Database (WIID) maintained by UNU-WIDER. Notwithstanding that the ways to compute some of the bounds have been previously known, to my knowledge, there has not been any work systematically applying these bounds to the problem of calculating world welfare, and *a fortiori*, no previous work has used the sharp bounds that I derive. Under very conservative imputation assumptions for countries and years without inequality data, I aggregate up the country welfare bounds to obtain bounds for world welfare. My first and main conclusion is that between 1970 and 2006, world welfare has risen unambiguously. The lower bound on world welfare growth in the baseline specification implies that welfare rose by 88%, with the effect being equivalent to an increase in the income of every person in the world of over \$2,600 in PPP-adjusted 2005 U.S. dollars. While the fact that world welfare measured in this way has increased may not be very surprising owing to the strong GDP growth enjoyed by the world over this period, the extent of growth is remarkable. Moreover, the sharper bounds that I derive allow a finer analysis of the consistency of welfare growth over the sample period: using these bounds, one can conclude that world welfare growth was always positive for any 10-year period under consideration, whereas under the bounds of Cowell (2000), such a conclusion could not be drawn. I subject my results to a battery of increasingly radical robustness checks that confirm that neither the finding that world welfare rose nor (to a lesser extent) the magnitude of this rise are sensitive to neither the necessary assumptions of my procedure nor the more substantive challenges to the validity of either the mean income or the survey data.

A related result is that the traditional data used in calculating the world distribution of income cannot reject the hypothesis that world inequality has risen. For almost all variations in the methodology of computing the bounds, my estimates are compatible with world inequality rising, or world welfare growing slower than world GDP per capita. The lower bound estimate for the baseline scenario suggests that welfare rose by only 93% of what it would have potentially risen by under uniform GDP per capita growth. However,

I can reject the hypothesis that rising inequality destroyed more than 50% of the welfare growth that would have obtained under uniform GDP per capita growth for most robustness checks for different methodologies of computing the bounds, and for many robustness checks I can reject the hypothesis that rising inequality destroyed more than 20% of potential welfare growth. In particular, no robustness check can yield a conclusion that inequality has unambiguously risen.

The paper is organized as follows: Section 2 presents the Atkinson measure of welfare and reviews its microfoundations. Section 3 presents the derivations of the nonparametric bounds. Section 4 discusses the data, the numerical implementation of the bounds, and the imputation assumptions made in order to construct bounds for world welfare. Section 5 presents the baseline results for world welfare and inequality, and discusses the gains from basing the bounds on more inequality statistics. Section 6 presents the robustness checks: 1) sampling error, 2) alternative imputation procedures, 3) alternative survey selections from the WIID dataset, 4) replacing the WIID dataset with the UTIP-UNIDO dataset pioneered by Galbraith and Kum (2005), 5) replacing the Penn World Tables GDP data by World Bank GDP data with different PPP, and 6) accounting for failure of survey coverage at the top. Section 7 concludes.

2 The Welfare Measure

An important aspect of the debate concerning the world distribution of income is the question of what are the appropriate metrics of poverty, inequality and welfare. In this paper, I will concentrate on a single family of welfare measures and its associated family of inequality measures, appealing to a well-known theoretical argument for the use of this family. This is the family of Atkinson equally-distributed income equivalents and inequality indices, introduced in Atkinson (1970). Atkinson treats the problem of assigning a welfare rating to an income distribution as equivalent to the problem of assigning utility ratings to lotteries. The Atkinson equally-distributed income (hereafter the welfare measure) is the certainty equivalent of the distribution of income treated as a lottery, whereas the Atkinson

inequality measure is the risk premium divided by the mean income. If it is assumed that utility is CRRA with risk aversion γ , which is a standard assumption in most empirical work, and is empirically supported in e.g. Chiappori and Paiella (2006), the relevant welfare index becomes

$$W(\gamma) = \left(\int_0^\infty x^{1-\gamma} dF(x) \right)^{\frac{1}{1-\gamma}}$$

and the relevant inequality index is

$$A(\gamma) = 1 - \frac{W(\gamma)}{\mu}$$

where μ is the mean income.

In terms of choice over lotteries, $A(\gamma)$ is the relative risk premium of the income distribution.

Atkinson's index resolves the chief problem in the construction (or even conceptualization) of a welfare index: the need to make interpersonal comparisons. Instead of assuming that the evaluator has some social preferences that allow her to trade off some utilities against others, the evaluator treats the income distribution as a lottery out of which she must draw a prize, and values this distribution accordingly. The choice of the (perfectly selfish) evaluator is then the choice she would make behind the Rawlsian veil of ignorance (although the welfare index would be the Rawlsian SWF only for $\gamma \rightarrow \infty$).

It is obvious that for $\gamma \geq 1$, any income distribution with an atom at zero income produces maximum inequality (an Atkinson index of 1, or an equally-distributed income of zero), and it is also immediate that no allocation of fractile shares nor any value of the Gini coefficient can rule out the income distribution having an atom at zero income, so for all $\gamma \geq 1$, it is impossible to construct a lower bound for welfare. Moreover, as $\gamma \rightarrow 1$ from below, the lower bound continuously drifts towards zero. Hence, bounds based on the fractile shares and the Gini coefficient can only be constructed for $\gamma \in (0, 1)$. (In particular, using the Rawlsian SWF or attempting to test for first-order stochastic dominance cannot be done

if one is to remain totally agnostic about the distribution as is done in this paper). I will use the central value in this interval, $\gamma = 0.5$, as the baseline for this paper. Statistics for the Atkinson inequality index for $\gamma = 0.5$ are routinely reported by developed countries (e.g. the United States Bureau of the Census reports Atkinson inequality indices for $\gamma = 0.25$, $\gamma = 0.5$ and $\gamma = 0.75$; the Luxembourg income study reports indices for $\gamma = 0.5$ and $\gamma = 1$), although I will show that the baseline result holds for γ as high as 0.9. I also perform a robustness check for values of γ higher than unity by assuming a lower bound for income equal to $1/5$ of the lowest fractile mean. Under this assumption, my baseline results are valid for γ as high as about 1.5.

3 Analytical Derivation of Uniform Bounds

In this section, I will first list a few facts about Lorenz curves and about the formulation of the optimization problems that yield the bounds in Lorenz curve space. I will then solve the optimization problems when the Gini coefficient is specified. Finally, I will then describe the solutions of the problems with fractile shares only, which is known by the literature on bounding inequality measures, and is reviewed in Cowell (2000).

3.1 General Remarks

I first review a number of basic facts about Lorenz curves, which can be found in, e.g. Gastwirth (1972) or derived by inspection.

1. The Lorenz curve $L(p)$ of a nonnegative random variable distributed according to $F(x)$ is $L(p) = \frac{1}{\mu} \int_0^{F^{-1}(p)} x dF(x)$, where $p \in [0, 1]$. This is an increasing and convex function, such that $L(0) = 0$ and $L(1) \leq 1$ (if a set of measure zero of the population holds a set of positive measure of income, then $L(1) < 1$).
2. Conversely, for any increasing and convex map $L(p)$ of $[0, 1]$ into itself, there exists a distribution function $F(x)$ such that $L(p)$ is the Lorenz curve of the nonnegative

random variable distributed according to $F(x)$. Hence, the set of Lorenz curves is

$$\mathcal{L} = \{L \in C[0, 1] : L \text{ is increasing, convex, } L(0) = 0, L(1) \leq 1\}$$

3. Any Lorenz curve is continuously differentiable at all but countably many points, and where it exists, $L'(p) = \frac{x}{\mu}$, where $x = F^{-1}(p)$.
4. The Gini coefficient of a distribution with Lorenz curve $L(p)$ is $G = 1 - 2 \int_0^1 L(p) dp$.
5. The Atkinson welfare index of a distribution with Lorenz curve $L(p)$ is $W(\gamma) = \mu \left(\int_0^1 (L'(p))^{1-\gamma} dp \right)^{\frac{1}{1-\gamma}}$.

Now, suppose we are given k fractile shares, or statements that individuals from the p_i th to the p_{i+1} st percentile (or in $[p_i, p_{i+1}]$) own fraction Q_i of the national income, with the cumulative share of national income owned by the lowest p_i earners being q_i . Suppose also that the mean of the income distribution is normalized to 1; then, the fractile boundaries, a_i^- and a_i^+ are defined as the (normalized) incomes of the p_i th and p_{i+1} st percentile of the income distribution respectively. The (normalized) mean income of fractile i is defined as $m_i = Q_i / (p_{i+1} - p_i)$.

6. By the definition of a Lorenz curve, an assignment of fractile shares equivalent to a set of constraints $L(p_i) = q_i$, $i = 1, \dots, k$, where p_i is the fraction of the population in or below fractile i , and q_i is the cumulative share of national income owned by this fraction of the population.
7. The statement that the boundaries of fractile i are $[a_i^-, a_i^+]$ is equivalent to the constraint $\lim_{p \downarrow p_i} L'(p) \geq a_i^-$ and $a_i^+ \geq \lim_{p \uparrow p_{i+1}} L'(p)$ (the inequalities are strict whenever there is no mass in the distribution of income around the p_i th or p_{i+1} st percentile).

Therefore, the main analytical problem of this paper can be formulated as follows:

$$\begin{aligned} \max_{L \in \mathcal{L}} \text{ or } \min_{L \in \mathcal{L}} \int_0^1 (L'(p))^\alpha dp \text{ st. } 1) \forall i = 1, \dots, k, L(p_i) = q_i, \\ 2) \int_0^1 L(p) dp = 0.5(1 - G) =: \bar{G} \end{aligned} \quad (1)$$

where $\alpha \in (0, 1)$

It will be useful to define the constraint set:

$$\mathcal{L}_c = \left\{ L \in \mathcal{L} : \bar{G} = \int_0^1 L(p) dp \text{ and } \forall i = 1, \dots, k, L(p_i) = q_i \right\}$$

3.2 Maxima

3.2.1 Maxima with Gini only

The optimization problem to maximize $W(1 - \alpha)$ for given Gini is

$$\max_{L \in \mathcal{L}} \int_0^1 (L'(p))^\alpha dp \text{ st. } \bar{G} = \int_0^1 L(p) dp$$

Since the objective function is concave in L , the maximum is unique and is attained in the interior, so this is a standard problem in the calculus of variations, since $L(p)$ is a.e.-twice differentiable. To solve this problem, I form the Lagrangian

$$\mathcal{L} = \int_0^1 (L'(p))^\alpha dp - \lambda \int_0^1 L(p) dp$$

and compute the Euler equation:

$$\frac{\partial \mathcal{L}}{\partial L} = \frac{d}{dp} \left(\frac{\partial \mathcal{L}}{\partial L'} \right) \Leftrightarrow \lambda = \alpha(1 - \alpha) (L'(p))^{\alpha-2} L''(p)$$

Proposition 1

The solution to the optimization problem is given by

$$L^*(p) = \frac{1 - \alpha}{\lambda} \left(c_1 - \frac{\lambda}{\alpha} p \right)^{-\frac{\alpha}{1-\alpha}} - c_2$$

where the constants c_1, c_2 and the Lagrange multiplier λ can be calculated from the equations

$$L^*(0) = 0, \quad L^*(1) = 1, \quad \bar{G} = \int_0^1 L^*(p) dp.^5$$

Proof. *In text.* ■

Note that the solution is a convex function, and thus a valid Lorenz curve. Hence, the value of the optimum attained corresponds to a sharp lower bound.

3.2.2 Maxima with Gini and fractile shares

The optimization problem to maximize $W(1 - \alpha)$ for given Gini and fractile shares is

$$\max_{L \in \mathcal{L}} \int_0^1 (L'(p))^\alpha dp \text{ st. } 1) \bar{G} = \int_0^1 L(p) dp \text{ and } 2) \forall i = 1, \dots, k, L(p_i) = q_i \quad (2)$$

The Euler equation is the same as in problem 1, however, the constants c_1 and c_2 are now allowed to vary by the interval i . An upper bound could be obtained by calculating the values of $c_{1,i}$ and $c_{2,i}$ from the conditions $L(p_i) = q_i$, however, unless the sequence $\{c_{2,i}\}$ is monotone increasing, the resulting solution is not convex, and hence, does not belong to \mathcal{L} . Hence, although it provides a greater lower bound than the solution when only the Gini constraint, or only the fractile shares are present, this bound is not sharp.

If in addition to the fractile shares, the fractile boundaries $\{a_i^-, a_i^+\}_{i=1}^k$ were given, a sharp upper bound could be obtained by solving the equivalent problem:

$$\max_{L \in \mathcal{L}} \int_0^1 (L'(p))^\alpha dp \text{ st. } 1) \bar{G} = \int_0^1 L(p) dp, 2) \forall i = 1, \dots, k, L(p_i) = q_i$$

$$3) \lim_{p \uparrow p_i} L'(p) \leq a_i^- \text{ and } \lim_{p \downarrow p_{i+1}} L'(p) \geq a_i^+$$

for some $\{a_i^-, a_i^+\}_{i=1}^k$ where a_i^- is the lower bound of interval i , while a_i^+ is the upper bound of interval i . Now, define

$$R(p) := \max_{i=1, \dots, k} \left\{ a_i^- (p - p_i) + q_i, a_i^+ (p - p_{i+1}) + q_{i+1} \right\}$$

⁵The implied CDF of this distribution is $F(x) = \frac{\alpha}{\lambda} \left(c_1 - \left(\frac{x}{\mu} \right)^{\sigma-1} \right)$ for $x \in \left[\mu (c_1)^{-\frac{1}{1-\alpha}}, \mu \left(c_1 - \frac{\lambda}{\alpha} \right)^{-\frac{1}{1-\alpha}} \right]$, $F(x) = 0$ for $x < \mu (c_1)^{-\frac{1}{1-\alpha}}$, and $F(x) = 1$ for $x > \mu \left(c_1 - \frac{\lambda}{\alpha} \right)^{-\frac{1}{1-\alpha}}$.

and consider the modified optimization problem:

$$\max_{L \in \mathcal{L}} \int_0^1 (L'(p))^\alpha dp \text{ st. } 1) \bar{G} = \int_0^1 L(p) dp, 2) \forall i = 1, \dots, k, L(p_i) = q_i, 3) L(p) \geq R(p) \quad (3)$$

Since the optimal solution must be convex, we have

$$\begin{aligned} & \left(\lim_{p \downarrow p_i} L'(p) \leq a_i^- \text{ and } \lim_{p \downarrow p_{i+1}} L'(p) \geq a_i^+ \forall i = 1, \dots, k, \right) \\ \Leftrightarrow & L(p) \geq \max_{i=1, \dots, k} \{a_i^- (p - p_i) + q_i, a_i^+ (p - p_{i+1}) + q_{i+1}\} =: R(p) \end{aligned}$$

so the two problems are equivalent.

By Kamien and Schwartz (1991), the solution to this latter problem is characterized as follows:

1. If $L(p) > R(p)$, then $L(p) = \frac{1-\alpha}{\lambda} (c_{1,i} - \frac{\lambda}{\alpha} p)^{-\frac{\alpha}{1-\alpha}} - c_{2,i}$ for some $c_{1,i}$ and $c_{2,i}$
2. If p is a "switching point" between $L(p)$ and $R(p)$ (so $\forall \varepsilon, \exists p_-, p_+ \in N_\varepsilon(p)$ st. $L(p) > R(p)$ on $[p_-, p]$ and $L(p) = R(p)$ on $[p, p_+]$, or vice versa), then, $L(p) = R(p)$ and $L'(p) = R'(p)$.

Hence, the solution $L^*(p, \mathbf{a})$ taking the vector of fractile boundaries $\mathbf{a} := \{a_i^-, a_i^+\}_{i=1}^k$ as given is characterized as follows:

Proposition 2 $\forall i = 1, \dots, k, \exists p_i^-, p_i^+ : p_i \leq p_i^- < p_i^+ \leq p_{i+1}$ st. $L^*(p, \mathbf{a}) = R(p)$ on $[p_i, p_i^-] \cup [p_i^+, p_{i+1}]$ and $L^*(p, \mathbf{a}) = \frac{1-\alpha}{\lambda} (c_{1,i} - \frac{\lambda}{\alpha} p)^{-\frac{\alpha}{1-\alpha}} - c_{2,i}$ on $[p_i^-, p_i^+]$, where λ solves $\bar{G} = \int_0^1 L^*(p, \mathbf{a}) dp$, while $p_i^-, p_i^+, c_{1,i}$ and $c_{2,i}$ solve

1. $\frac{1-\alpha}{\lambda} (c_{1,i} - \frac{\lambda}{\alpha} p_i^-)^{-\frac{\alpha}{1-\alpha}} - c_{2,i} = R(p_i^-) = a_i^- (p_i^- - p_i) + q_i$
2. $\frac{1-\alpha}{\lambda} (c_{1,i} - \frac{\lambda}{\alpha} p_i^+)^{-\frac{\alpha}{1-\alpha}} - c_{2,i} = R(p_i^+) = a_i^+ (p_i^+ - p_{i+1}) + q_{i+1}$
3. $(c_{1,i} - \frac{\lambda}{\alpha} p_i^-)^{-\frac{1}{1-\alpha}} \geq a_i^-$ with strict inequality iff $p_i^- = p_i$, and

4. $(c_{1,i} - \frac{\lambda}{\alpha} p_i^+)^{-\frac{1}{1-\alpha}} \leq a_i^+$ with strict inequality iff $p_i^+ = p_{i+1}$.

Proof. *In text.* ■

The problem reduces to the finite-dimensional problem of optimization along the finite sequence \mathbf{a} , which can be done with standard software. If the fractile boundaries are known, the formula can be used directly and no optimization is required.

3.3 Minima

The problem is

$$\inf_{L \in \mathfrak{L}} \int_0^1 (L'(p))^\alpha dp \text{ st. } 1) \bar{G} = \int_0^1 L(p) dp, 2) \forall i = 1, \dots, k, L(p_i) = q_i \quad (4)$$

It can be shown that the infimum of this problem is attained by a Lorenz curve that is a linear spline with corners (possibly) at the points $\{p_i\}$ and with no more than one corner in any interval (p_i, p_{i+1}) .

Proposition 3 *Consider the minimization problem (4). Then, the value of this problem is identical to the value of the following finite-dimensional problem:*

$$\min_{\{a_i^-, a_i^+\}_{i=1}^k} \left\{ \left(\sum_{i=1}^k \left(\lambda_i \left(\frac{a_i^-}{\mu} \right)^\alpha + (1 - \lambda_i) \left(\frac{a_i^+}{\mu} \right)^\alpha \right) \right)^{1/\alpha} \right\} \quad (5)$$

$$\text{st. } \forall i = 1, \dots, k, \lambda_i = \frac{a_i^+ - m_i}{a_i^+ - a_i^-}, a_i^- \leq m_i \leq a_i^+ \leq a_{i+1}^- \text{ and } \bar{G} = \int_0^1 L(p, \mathbf{a}) dp$$

where $L(\cdot, \mathbf{a})$ is the Lorenz curve defined by the sequence $\{a_i^-, a_i^+\}_{i=1}^k$.

Proof. See appendix. ■

The proof relies on the following lemma:

Lemma 4 (3-2) *Lemma: Suppose that $L \in \mathfrak{L}_c$ is piecewise linear with finitely many corners. Then, there exist numbers $\{a_i^-, a_i^+\}_{i=1}^k$ such that $\forall i = 1, \dots, k, a_i^- \leq m_i \leq a_i^+ \leq a_{i+1}^-$, and*

the Lorenz curve

$$\hat{L} = \max_{i=1,\dots,k} \left\{ \max \left\{ a_i^- (p - p_i) + q_i, a_i^+ (p - p_{i+1}) + q_{i+1} \right\} \right\}$$

satisfies $\hat{L} \in \mathfrak{L}_c$ and $\int_0^1 \left(\hat{L}'(p) \right)^\alpha dp \leq \int_0^1 \left(L'(p) \right)^\alpha dp$.

Proof. See appendix. ■

The 3-2 Lemma states that any piecewise linear Lorenz curve with finitely many kinks that are off the fractile constraints can be replaced with a modification with only $2k$ such kinks that satisfies the constraints in the minimization problem (4), and decreases the value of the program. Heuristically, the minimizing curve must have no more than one kink in each interval, or no more than $2k$ kinks overall, which makes the family of possible minimizing curves finite-dimensional, and allows them to be computed by standard numerical methods. If the fractile boundaries are known (or assumed), the formula can be used directly and no optimization is required.

In the special case that there are no fractile constraints, it is easy to derive an explicit formula for the maximum Atkinson given the Gini coefficient, as the optimal curve must have no more than one kink.

Proposition 5 Consider the maximization problem (4) and omit the fractile constraints. Then, the maximum value of the Atkinson index is given by

$$\max \left(G, 1 - (1 - G)^{\frac{1-\alpha}{\alpha}} \right)$$

Proof. See appendix. ■

3.4 Results with Fractiles Only

Cowell (1977) proves that if there are k fractiles, each with known mean m_i and fractile boundaries $[a_i^-, a_i^+]$, then the maximum welfare is attained at intrafractile egalitarianism: the distribution is concentrated at the fractile means, and the value of the problem is given

by $\left(\frac{1}{k} \sum_{i=1}^k \left(\frac{m_i}{\mu}\right)^\alpha\right)^{1/\alpha}$. The minimum welfare is attained by complete concentration on the fractile boundaries, or by the solution to the finite-dimensional optimization problem (5) omitting the Gini constraint. In particular, a crude (non-sharp) approximation to the minimizer of welfare with fractiles only can be computed in closed form by setting $a_i^- = m_{i-1}$ and $a_i^+ = m_{i+1}$ for each i .

Figures Ia) and Ib) presents plots of some welfare-maximizing and welfare-minimizing Lorenz curves (the quintile shares that they are based on are denoted by bold circles in the diagrams). It is clear that the welfare-minimizing curves are all piecewise linear, while the welfare-maximizing curves that involve the Gini are nonlinear. Note how the crude approximations to both the welfare-maximizing and welfare-minimizing curves are nonconvex.

4 Implementation

4.1 GDP Data: Penn World Table

The Penn World Table (hereafter PWT) is one of the most cited sources for purchasing-power-parity-adjusted GDP data. The latest edition (version 7.0) has nearly comprehensive coverage of 189 currently existing countries since 1970 to 2009. I reconstruct GDP for currently nonexisting countries (e.g. the Soviet Union, Czechoslovakia, East Germany) by applying the growth rates of Penn World Tables version 5.6 to the implied GDP for these countries in version 7.0. This procedure is discussed in detail in Pinkovskiy and Sala-i-Martin (2009).

A major controversy in the literature is whether estimates of GDP should come from national accounts or from household surveys. Ahluwalia et al. (1979) pioneered the combination of national accounts GDP and survey inequality data, which is the dominant approach today, and is used, e.g. by Bourguignon and Morrisson (2002), Sala-i-Martin (2002a and b, 2006) and Bhalla (2002). Proponents of using national accounts to estimate the mean of the distribution of income argue that survey means tend to understate mean income and sometimes yield implausible implications (see e.g. the discussion in Bhalla (2002)).

Moreover, national accounts estimates of GDP, and the Penn World Table in particular, are extensively used in cross-country research on growth and development: in particular, the seminal works of Barro and Sala-i-Martin (1992,1995); Barro (1999); Acemoglu, Johnson and Robinson (2001, 2002, 2005); and Banerjee and Duflo (2004) all use Penn World Table GDP, sometimes in conjunction with the Deininger-Squire dataset on inequality, whereas no such paper to my knowledge uses survey means. Other papers, such as Milanovic (2002) and Anand and Segal (2008), strongly criticize the use of national accounts on the grounds that it is inconsistent to take the distribution of income from one source and the mean of income from another. A practical consideration in favor of national accounts is that national accounts estimates are calculated using common methodology for virtually all countries and years, whereas survey means tend to be available for far fewer countries only in select years, so further assumptions are required in order to use a series of survey means. Since the focus of this paper is to present the methodology of the uniform bounds and to observe their implications for widely used data on the distribution of income, I will use national accounts as my source of GDP, but I will conduct a robustness check using survey means. Moreover, in the robustness check correcting for nonresponse, I will attempt to control for possible mechanisms that lead national accounts and survey means to diverge.

Figures IIa) and IIb) present a brief summary of the data. First, we see that GDP growth in 1970-2006 has been extraordinary – GDP has nearly doubled.⁶ Second, we see that between-country inequality – the value of the Atkinson inequality index with $\gamma = 0.5$ that would obtain if the income distribution in each country were egalitarian – fell significantly, with much of the fall taking place after 2000. These results are suggestive of the claim that world welfare has increased, but are not conclusive, since if within-country inequality has increased by a substantial amount, welfare could have actually fallen. In fact, these results could be perfectly consistent with a "nightmare scenario" of a global elite, tiny in number but evenly distributed across nations, capturing most of the gains to growth in the past several decades.

⁶Note that the sample period for this paper ends before the beginning of the global recession in 2008.

4.2 Inequality Data: The World Income Inequality Dataset

The World Income Inequality Dataset (WIID), maintained by UNU-WIDER, is a significantly improved and expanded version of the Deininger-Squire (DS) dataset pioneered in 1996. It is probably the most comprehensive, and the most cited source on income inequality around the world⁷, presenting over 5,200 surveys for over 150 countries and 79 years. Over 75% of the surveys listed took place after 1970. All the survey data reported include an estimate of the Gini coefficient, over 2,700 surveys contain quintile shares, and over 2,000 contain decile shares. Moreover, for nearly all the surveys, the database records the coverage of the survey of different parts of the country in question, the income concept asked for in the survey (income or consumption, gross income or net, whether in-kind income is included), the conversion factor used to obtain inequality between persons from household-level data, and the statistical agency conducting the survey and the researchers reporting its results.⁸ However, the WIID, as well as the DS dataset from which it was constructed, errs on the side of comprehensiveness of coverage rather than quality of surveys. Atkinson and Brandolini (2001) criticize the DS dataset for including poorly conducted and methodologically unclear surveys alongside well-organized ones, and criticize much research using the dataset for disregarding the noncomparabilities of surveys with different income concepts, different equivalence scales, and different underlying populations.⁹ Hence, an important problem for any researcher using the dataset is to provide a method of selecting which surveys to use that avoids these pitfalls.

⁷The paper introducing the dataset, Deininger and Squire (1996), has 1,884 citations on GoogleScholar.

⁸There may be concern that the Gini coefficients listed in the WIID are estimated from the presented fractile shares. I have checked that no Gini coefficient that the WIID presents is given by either the minimum or the maximum value of the Gini coefficient that is theoretically compatible with the fractile shares (formulas are given in Mehran (1975)). The documentation to WIID mentions explicitly for a few surveys (121 out of 2240 survey groups) that the Gini was constructed from the fractile shares. In results not presented, I have recomputed the baseline estimates excluding these Gini data, and the results are indistinguishable from the baseline results reported in this paper.

⁹However, it appears that the WIID has drawn lessons from the critique of the DS dataset; in personal communication, Tony Atkinson noted that "WIDER did a great deal to clean the original DS database; the WIID database is much less subject to the kind of criticisms that [Atkinson and Brandolini (2001)] made." (Tony Atkinson, personal communication, December 2009).

4.2.1 Choice of Surveys

I subdivide all surveys in the WIID into groups, hereafter surveygroups, within which all surveys 1) describe the same country, and the same geographic, demographic and socio-economic population within the country, 2) have the same income concept, 3) collect income data on the same unit, use the same unit of analysis, and use the same equivalence scale to convert between the two if they are distinct, and 4) have the same primary and secondary source. I identify 2240 such surveygroups in the WIID, which means that each contains on average a little over two surveys, but some have much wider coverage than others. I then, instead of selecting separate surveys, select entire surveygroups on a heuristic basis by weighting the following considerations in approximately the following lexicographical order:

1. I give preference to those surveygroups that provide decile shares over those that provide only quintile shares, and I give preference to surveygroups providing fractile shares over those that only provide Gini coefficients,
2. I attempt to ensure that the surveygroups cover the longest date range for each country, and be well-distributed over the sample period, preferring a few surveys in each decade to thorough coverage of some periods at the expense of others,
3. I attempt to ensure that the surveygroups selected for each country have the same or similar income concepts, the same or similar equivalence scales and geographical extent, and the same primary source.
4. I attempt to maintain homogeneity in the characteristics of surveygroups selected across countries, with surveys asking about disposable income and equivalizing on the basis of household per capita being preferred. However, there are gross exceptions to this (e.g. India offers quintile shares only for consumption surveys).
5. In view of the fact that I will need to interpolate and extrapolate to estimate inequality for years for which I do not have survey data, I choose some surveys outside the sample period (i.e. before 1970).

I also maintain some region-specific conventions aside from these four general principles. In particular, I use the Luxembourg Income Study surveys for OECD countries unless long and detailed series are available in the WIID from the countries' own statistical bureaus. For Latin American and Caribbean countries, I almost invariably use the surveys provided by the Socio-Economic Database for Latin America and the Caribbean (SEDLAC), following the recommendation of WIID. For the populous East Asian countries of China, India and Indonesia, I use survey data provided by the national statistical bureaus. Finally, owing to a dearth of surveys for Africa, particularly from the beginning and middle of the sampling period, I suspend many of the homogeneity requirements and often use consumption surveygroups when these offer more extensive coverage than income surveygroups do. Overall, I choose surveygroups containing 1094 surveys, of which 1011 lie in the sample period.

For some countries and years, the unit records from household surveys are publically available. Chen and Ravallion (2010) use a database of over 700 surveys, many of which were made available to them as microdata, whereas others were provided only in grouped data form. Their PovCal website contains a description of all the surveys, including whether their unit records were available, and the parametric estimates of their underlying income distribution obtained using the Kakwani-Podder method. Furthermore, the Luxembourg Income Study provides microdata for many household income surveys in the OECD. For all countries and years for which microdata is available, I use either the published inequality statistics (Atkinson inequality indices) directly, or (in the case of the Chen-Ravallion data) I compute Atkinson inequality indices from the parametric estimates obtained by Chen and Ravallion on the basis of the microdata they used, assuming that these estimates are probably very close to the actual values of the Atkinson inequality indices in the microdata. Using microdata decreases the width of my bounds slightly but noticeably in the period 1995-2005 (to which most of the available microdata corresponds), and does not affect them for the preceding period.¹⁰

¹⁰For the OECD countries with Luxembourg Income Study data, the Atkinson inequality indices are available only for $\gamma = 0.5$ and $\gamma = 1$. Whenever I compute Atkinson indices with other parameters, I use the $\gamma = 1$ index from the LIS. This does not affect the qualitative conclusions reached.

4.2.2 Breadth of Coverage

As is intuitive from Figures IIa) and IIb), inequality does not tend to vary much over short periods of time, especially when compared to variation in GDP, so interpolation (as opposed, possibly, to extrapolation) procedures to impute inequality measures for years without data should be relatively reliable. Hence, while one intuitive measure of the breadth of coverage is the percent of the world population in the given year who are covered by surveys, a potentially better measure is the percent of the world population who are either covered by a survey in that year, or whose inequality measures will be obtained by interpolation (rather than extrapolation). Hereafter, I define the *core* to be the set of individuals who are so covered. Figure IIc) presents these measures. While the direct coverage measure is highly erratic (depending significantly on whether China and / or India are covered in a given year), and tends to be below 60%, the size of the core as a percent of the world population is remarkably continuous, and tends to be above 70% until about 1998, and around 90% for most of the 1980s and 1990s. Hence, at least until the 2000s, coverage using the WIID is rather good.

4.3 Numerical Implementation

I have implemented numerically all the bounds described above except for the sharp bound for the maximum welfare given both the Gini coefficient and the fractile shares. In its place, I am reporting the crude upper bound for welfare based on Ginis and fractiles, which, while superior to both the Gini and the fractile upper bounds taken individually, may only be attained by curves that are not convex. All the other bounds for the maximum are very easy to implement, as they only rely on finding the root of a monotonic function in one variable. The sharp bounds for the minimum welfare given fractile shares (with or without the Gini) require numerical optimization over a long vector of arguments, the sequence $\{a_i^-, a_i^+\}_{i=1}^k$.¹¹

¹¹It is possible to implement this optimization straightforwardly using the Matlab program `fmincon` on a standard PC. While the solution does depend on the initial value chosen for the optimization, and while the program occasionally fails to

4.4 Assumptions for Interpolation and Extrapolation

The bounds I have computed given fractile shares and the Gini coefficient say absolutely nothing in theory about the behavior of inequality in countries and years for which we do not have data, so the only fully conservative bound for those country-years is the trivial bound $[0, 1]$. However, it is accepted in the area of inequality research that inequality tends to change very slowly and very continuously,¹² so in practice, inequality data in a given year should give a great deal of information about inequality in that country in nearby years. The average coefficient of variation of the Gini within a surveygroup is only 0.06, and it does not exceed 0.41 for any surveygroup.

A plausible and easy-to-implement interpolation assumption is that inequality in any given year for which data is missing is bounded above and below by the inequality in the closest preceding and following years (hereafter, closest available years) with data available. Then, the upper bound of inequality for that year is the maximum of the upper bounds of the inequality in the closest available years, while the lower bound is the minimum of the lower bounds. For country-years outside the core, this method is tantamount to horizontal extrapolation of the bounds, which is problematic as it may artificially truncate rising trends in inequality. Hence, I use a more conservative extrapolation procedure that interpolates the upper and lower bounds linearly if this would result in the bounds widening further apart, and horizontally otherwise. To avoid upper and lower bounds from reaching implausible values, I bound them by the maximum upper and minimum lower bounds obtained from the data within each World Bank region¹³. As it is very rare that inequality rises or falls at a linear rate for an extended interval of time, it is plausible that such extrapolation would account for the possible dynamics of inequality at the ends of the sample period.

A more difficult problem is to impute inequality for countries with no survey data at

converge, the variation in the result as a function of the initial value is extremely small (the bound on the inequality index varies by less than 1% of the maximum value of this index with $\gamma = 0.5$ for most observations). As a compromise between speed and accuracy, I run the program for each survey for no less than twenty randomly selected starting values, and stopping at the first subsequent time the program converges.

¹²See e.g. Bhalla (2002) or Galbraith and Kum (2005). The latter source considers changes of 5 Gini points or more per year to be "unlikely, except when they coincide with moments of major social upheaval."

¹³For a classification of countries into World Bank regions, see Sala-i-Martin (2006)

all. Reasoning that countries may tend to be like other countries around them, and following Sala-i-Martin (2006) and Pinkovskiy and Sala-i-Martin (2009), I impute these inequality measures on the basis of the inequality of the other countries in their World Bank region. However, to be conservative, I impute the upper bound in every year to be the maximum upper bound observed in the data for all countries and years in that region, and the lower bound similarly. I will investigate more and less conservative methodologies for interpolation in the robustness checks.

5 Baseline Results

5.1 A Simple Test for a Rising or Falling Series

The results for each time series will take the form of upper and lower bounds, rather than point estimates. Under the assumptions for imputation, as well as under the assumption that the data is valid, and that sampling error can be ignored, these bounds contain the true value of the measure of interest with probability 1. Any path of the measure that is contained within the bounds is therefore consistent with the data, whereas any path that violates the bounds at any point is inconsistent with it.

There is a simple procedure for drawing conclusions as to whether a series increased or decreased between two dates. If the lower bound of the series at the earlier date exceeds the upper bound at the later date, the series fell for sure (with the caveats expressed in the previous paragraph). If, on the contrary, the *upper* bound of the series at the earlier date is exceeded by the lower bound of the series at the later date, the series rose for sure. However, neither of these statements may be true; in which case, it is impossible to draw conclusions on whether the series rose or fell between the two dates without further assumptions or data.

5.2 Example for a Single Country: Chinese Inequality

Before presenting my baseline results for inequality and welfare in the world as a whole, I present upper and lower bounds on the Atkinson inequality index for China in order to demonstrate the interaction of my bounding technique, interpolation and extrapolation on a single consistent set of surveys. Moreover, estimates for China are interesting in their own right because microdata from Chinese official income surveys is not released to the public. Figure Ic) presents the series for the $\gamma = 0.5$ Atkinson inequality index for China. We see that inequality in China was between 0.07 and 0.24 in 1970 and 0.16 and 0.28 in 2006, which is consistent with Chinese inequality rising or falling over this time period (which involved a transition to capitalism and is believed to have witnessed a substantial rise in Chinese inequality). We also see that the interpolation and extrapolation procedures appear to be reasonable and to yield results that are not radically different from the observed values of the bounds.

5.3 Baseline results

All the baseline results and robustness checks are summarized in Table I. The table presents for all variations (except sampling error) 1) the minimum and maximum amounts by which Atkinson welfare (interpreted as the certainty equivalent of the world distribution of income) increased between 1970 and 2006, 2) the minimum and maximum percentage increases in Atkinson welfare since 1970, and 3) the minimum and maximum percentage increases in Atkinson welfare as a percentage of what they would have potentially been if all incomes grew at the same rate (uniform GDP growth). Thus, the lower bound in part 3) is informative as to how much less welfare growth there is because of the fact that growth in GDP per capita is distributed unequally. Note that we can reject the hypothesis that inequality rose if and only if the lower bound in part 3) is greater than 100%; welfare grew faster than did GDP because inequality shrank.

I present the time series of world welfare for implied risk aversion $\gamma = 0.5$ and $\gamma = 0.9$

in Figure IIIa) For both indices of risk aversion considered, world welfare rose between 1970 and 2006. For $\gamma = 0.5$, we can also reach the conclusion that world welfare rose between 1990 and 2006, and even between 2000 and 2006. These are very important findings, since they establish that for plausible levels of risk aversion, (and even for relatively high ones, such as $\gamma = 0.9$, for which fully nonparametric bounds may be expected to be difficult to construct), the only series consistent with the data imply that even accounting for its uneven distribution, growth was sufficiently high relative to any increase in inequality that overall welfare rose. This conclusion is also intuitive given the more primitive facts of the dataset we use: if per capita GDP grew by nearly a factor of two, and between-country inequality fell substantially, and within-country inequality as measured by the Gini varied very little, the only way that welfare could have fallen was if movements in the Gini coefficient and in the fractile shares were unrelated to movements in the Atkinson index. No less important is it to note by how much welfare rose: Table I indicates that for $\gamma = 0.5$, welfare rose by at least 88% between 1970 and 2006.

From the time series of inequality in Figure IIIb), I must remain agnostic about the direction of world inequality: it is impossible to tell whether inequality rose or fell without additional assumptions. While I can almost reject the hypothesis that inequality rose according to $\gamma = 0.5$ (the relative risk premium of the income distribution could have, at most, risen from 0.396 to 0.414), it is obvious by inspection that for $\gamma = 0.9$, the data is consistent with many possible rising or falling time paths of inequality. In particular, this finding indicates that the large drop in between-country inequality could have been more than overridden by a rise in within-country inequality. However, these bounds also display the relatively limited feasible variation in inequality. Table I shows that for the baseline specification, rising inequality could have eroded at most 7% of the welfare benefits of GDP growth, and for a risk aversion coefficient even as high as $\gamma = 0.9$, the largest possible inequality increase could have decreased the growth rate of welfare relative to uniform growth by at most 31%.

The benefit of using uniform bounds is that this failure to reject should not be

interpreted as a "null result," but rather as a criticism of the (amount and presentation of the) data. It indicates that, at least without stronger assumptions on the form and evolution of inequality within countries, it is impossible to tell whether inequality rose or fell. If we wish to reach a conclusion, what is required is more surveys, more finely presented. In the robustness checks, I will show that the failure to reject comes largely from the paucity of information in the surveys (from the width of the bounds when inequality data is given) rather than from the conservatism of my imputation assumptions.

5.4 Gain from fine bounds

It is useful to see how much we gain by basing our bounds on additional data, and how much we gain by using sharp rather than loose bounds. Figure IVa) presents welfare estimates for $\gamma = 0.5$ using four methods: 1) crude bounds based on fractiles, 2) sharp bounds based on fractiles, 3) sharp bounds based on the Gini, and 4) bounds (crude upper bound and sharp lower bound) based on both the Gini and fractiles. We see that 1) including the Gini coefficient as well as the fractile shares in estimating the bounds decreases the width of the intervals (more in the earlier than in the later part of the sample because there are substantially more surveys with unit records in the later part of the sample), 2) sharpening the fractile-based bounds does not appreciably decrease the width of the intervals, and 3) while the lower bound on welfare based on the Gini is very poor, the upper bound on welfare based on the Gini is quite good, and can be superior to the upper bound based on fractile shares (which, in practice, tend to be decile shares). In particular, the upper bounds on welfare are all extremely close together, which is consistent with the idea that there is little gain in further improving the upper bound, so our omission of the sharp upper bound with the Gini and fractile shares is not a large loss.

It is reasonable to ask whether there is anything that we gain from using finer bounds. From Figure IVa) , we see that if we could base our bounds only on the Gini coefficient, we could not reject the hypothesis that welfare fell between 2000 and 2006 even for $\gamma = 0.5$. We can also deduce from Figures IIIb) and IVa) that if we did not have a formula for

calculating the upper bound on inequality given both the Gini and fractile shares, the data would be consistent with a large, rather than a trivial, rise in inequality for risk aversion coefficient $\gamma = 0.5$, and would have been consistent with a rise in inequality for risk aversion coefficients much lower than 0.5. However, at least within our baseline results, there are no other immediately obvious hypotheses that critically depend on the use of the finer bounds.

The power of the finer bounds can, however, be seen if we turn to the analysis of the fine structure of the welfare time series. It is of interest to ask what we can say about the rate of welfare growth over the period 1970-2006. Given upper and lower bounds for welfare, bounds for welfare growth can be easily constructed without losing sharpness by computing the upper bound of growth as the growth between the lower bound at date 1 and the upper bound at date 2, and vice versa for the lower bound of growth. It is easy to see from Figure IVa) that the width of the bounds exceeds the typical 1-year growth rate, so it is useful to compute average growth rates over long periods, such as 10 years; however, this procedure prevents us from talking about growth trends at the ends of the sample period.¹⁴ Figure IVb) shows bounds on the growth rate using fractiles only, and using fractiles together with the Gini coefficient for averaging periods of 10 years. It is obvious that using the finer bounds connotes an important improvement; we can reject the hypothesis that average annual growth rates in any 10-year period in the sample were negative using the fine bounds, but not using the fractile-based bounds. Moreover, the width of the bounds shrinks considerably when the bounds are finer, and we can make some nontrivial statements about the level of growth in different time periods with the fine bounds, such as the average annual growth rates being bounded away from zero. Unfortunately, we cannot make any statements about welfare growth accelerating or decelerating during the sample period, a question of obviously great interest.

¹⁴This limitation may actually be appropriate in practice, since the growth dynamics at the ends of the sample period may be products of extrapolation. However, it is clear we can have too much of a good thing, as with 20-year average growth rates, we lose more than half of our 36-year long sample period.

5.5 Higher Atkinson Parameters

As mentioned in section 2, it is impossible to construct nontrivial bounds for Atkinson welfare indices with coefficient greater than unity because the lower bound is zero whenever a distribution of income with positive mass at zero is allowed. However, bounds for higher degree Atkinson indices can be constructed under the assumption of a minimum income. A flexible procedure for selecting such a minimum income is to assume that the minimum income is a fixed fraction of the lowest fractile mean, which allows poorer countries to have lower minimum incomes than richer countries. Figures Va) and Vb) show plots for Atkinson welfare indices with $\gamma = 1.25$ and $\gamma = 1.5$ under the assumption that the minimum income is one-fifth of the mean income of the lowest fractile. Since the magnitudes of the higher-parameter Atkinson indices are much lower than the magnitude of the $\gamma = 0.5$ Atkinson index, the series are plotted at different scales, but the growth dynamics of the higher-parameter Atkinson indices are clear. Even for these higher values of the Atkinson parameter, welfare rises unambiguously, although we see from the figures and from Table I that the bounds are much wider and are compatible with much smaller rises in welfare. For $\gamma = 1.25$, rising inequality could have destroyed 31% of potential welfare growth, and for $\gamma = 1.5$, it could have eroded as much as 69%.

6 Robustness Checks

The formulae for the uniform bounds given fractile shares and the Gini coefficient are derived analytically, and hence need to be checked for robustness only to the relaxing the assumptions underlying them. The substantive assumptions underlying the baseline results presented in the previous sections are as follows: 1) the GDP data in the PWT and the inequality data in the WIID selected as described in fact do describe accurately the true GDP and inequality measures of the countries in question, and 2) the interpolation and extrapolation method assumed in section 4 is a good approximation for the actual behavior of the time series in question. These assumptions will be scrutinized in what follows.

6.1 Sampling error in the fractile shares and the Gini coefficient

The derivations in section 3 took the fractile shares and Gini coefficient to be known without error; in fact, these quantities are survey estimates that depend on the sample collected, so there may be a nonzero probability that the true values of the Atkinson index are not contained in the bounds constructed from the empirical estimates. The idea that sharp bounds based on empirical estimates may fail to contain the population value for which they have been constructed is explored in McDonald and Ransom (1981) and is a serious problem. In the context of the paper, there is reason to believe that this problem is small, since the WIID provides information on the sample sizes of most of the listed household surveys, and these sample sizes are very large, with median sample size equal to 23,900. However, since the surveys are nonrandom samples, and in particular, probably have high degrees of clustering, the variances of the resulting estimates are higher than the corresponding variances would have been had the surveys been simple random samples.

I perform a robustness check for sampling error by the following procedure:

1. I assume all underlying country distributions of income to be lognormal with inequality parameter implied by the Gini coefficient.
2. I draw 100 simple random samples from each survey with sample size equal to 1/10th of the listed sample size in order to control conservatively for the variance-inflating effect of nonrandom sampling procedures.
3. I compute the decile share-based bounds (using the crude closed-form version of the bound for the minimum welfare) for each draw.
4. I compute sampling-error-adjusted confidence bounds for welfare as the upper bound plus 2 standard deviations, and the lower bound minus 2 standard deviations.
5. I aggregate these bounds to obtain bounds for the entire world using the baseline imputation assumptions.

This procedure is very conservative, as sampling error is likely to be independent (or very weakly correlated) across surveys, so aggregating all the lower and upper bounds considers the very unlikely case that sampling error (as opposed to systematic error) consistently was in a downward (or upward) direction for *all* surveys in the dataset. Hence, the resulting confidence bounds contain the true value of world welfare in a given year with a probability far higher than 95%. The graph of the resulting bounds, along with the bounds based on the lognormal fractiles without sampling error, are presented in Figure VIa). It is obvious that adding sampling error, even in a highly conservative fashion, does not substantially affect the bounds.

6.2 More conservative interpolation / extrapolation

One of the substantive assumptions that had to be made in the aggregation of country estimates to get the world welfare estimates concerned the imputation of bounds for country-years without any inequality data. Our baseline assumption is that the survey data gives us all the peaks and the troughs of the time series, so observations for the missing country-years should be contained between the outer envelope of the bounds of the closest available observations. An (extreme) alternative methodology would be to compute the highest upper bound and lowest lower bound observed in the data for the given country, and assume that inequality in this country never violates these bounds. Hence, we relax our assumption that all the peaks and troughs of the inequality series are observed to the assumption that we observe the highest peak and the lowest trough. Extrapolation is still performed linearly, so as to allow inequality to grow to values not observed in the sample.

I present the resulting bounds for welfare along with the baseline bounds for $\gamma = 0.5$ in Figure VIb) and Table I. It is clear that we can reject the hypothesis that welfare did not grow in favor of the hypothesis that it grew for most periods of interest, and over the course of the sampling period. The bounds do widen, and we see that rising inequality could have destroyed as much as 41% of potential welfare growth (although the bounds are compatible with inequality falling as well).

One may argue that I fail to reject the hypothesis that world inequality rose because my interpolation scheme is too conservative: in particular, 1) the bounds for countries without surveys are too wide since they capture uncertainty in the level of inequality in the country as well as uncertainty coming from functional form, 2) the outer envelope interpolation is too cautious, since inequality tends to rise smoothly, 3) the linear extrapolation is too conservative as it inflates uncertainty due to functional form. In particular, Sala-i-Martin (2006) and Pinkovskiy and Sala-i-Martin (2009) impute inequality for countries with no survey data using regional average inequality, while Milanovic (2002) and Chen and Ravallion (2001) (implicitly) interpolate and extrapolate horizontally by using surveys from nearby years to stand in for surveys in years of interest. Therefore, Figure VIII) considers what happens to the baseline inequality series ($\gamma = 0.5$) when these assumptions are relaxed. One modification replaces the bounds for countries without data by the average (rather than the envelope) of the bounds for countries in the same region with data, and the second modification also interpolates the bounds within the core (rather than taking their envelope) and uses horizontal extrapolation of the bounds rather than linear extrapolation. For the second modification, it is possible to barely reject the hypothesis that world inequality rose (the lower bound in 1970 is 0.4006 and the upper bound in 2006 is 0.3997), but this is entirely a result of using horizontal as opposed to linear extrapolation. Hence, in order to understand whether inequality has risen or fallen since 1970, it is necessary to collect more recent survey data (some of which has probably not been processed in the case of recent surveys) in finer categorizations.

6.3 Alternative inequality data: different procedures for choosing surveys¹⁵

As noted by, e.g. Atkinson and Brandolini (2001), in using the DS database or the WIID, it is crucial to select comparable surveys so as to avoid comparing the inequality of conceptually different distributions. While it is difficult to write a formula that can combine the various considerations that go into determining which surveys to select, the methodology

¹⁵In results not reported, I also consider replacing the WIID survey data with data on Gini coefficients from Galbraith and Kum (1999). The results are essentially identical to my baseline Gini results.

for selecting surveys that I have presented in section 4 can be justly criticized for being heuristic and difficult to replicate. Therefore, I provide two alternative methodologies; one that seeks to ensure comparability of the surveys selected within each country at the cost of a substantial loss of coverage, and another that attempts to control for the range of sampling and nonsampling error in the computation of the Gini coefficient at the cost of not being able to use fractile shares to reduce the width of the distribution-free bounds.

The first (hereafter homogeneous) methodology entails selecting the surveygroup with fractile shares with the largest number of surveys within the sample period for each country, and taking surveys for that country only from the selected surveygroup. Hence, all surveys for a given country must be identical along all dimensions that are held fixed within a surveygroup: source, underlying population, unit of analysis and equivalence scale, and income concept. However, this methodology does not attempt to ensure homogeneity across countries, and recognizes that while within-country trends in inequality will be measured using comparable data, the levels of inequality in different countries will not necessarily be comparable. (Trying to ensure homogeneity across countries by further excluding surveys from the WIID would do violence to the procedures, as either China, which has almost exclusively income surveys, or the Indian subcontinent, which has almost exclusively consumption surveys, would be excluded). From Figure IXa), we see that this methodology drastically restricts coverage; only for the 1980s is more than 70% of the world covered even indirectly (in the core), and (not shown) the inequality series for China stops in 1992.

The second methodology (hereafter the extreme Ginis methodology) involves ignoring the differences between all surveys in terms of income concept and unit of analysis (but acknowledging the differences in terms of the underlying population), and for each country-year, taking as the final bounds the outer envelope of the bounds based only on the Gini coefficient for each Gini coefficient presented in the WIID for that country-year. The extreme Ginis methodology conjectures that all the income concepts and equivalence scales in the surveys are imperfectly implemented, but the range of resulting estimates captures the Gini coefficient that would result from an ideal implementation of a consistent income concept and

equivalence scale. The average standard deviation of the Gini estimates is 0.042, which far exceeds the time standard deviation of the Gini coefficient within a given surveygroup across multiple years (whose mean and median are approximately 0.02). Hence, it is plausible that, given the wide range of the Gini estimates, this range contains the true value of the Gini. Obviously, this methodology expands the coverage of the surveys: Figure IXb) shows that more than 90% of the world population are in the core until 2000, and more than 80% until 2003, while about 60% are directly covered by surveys.

Figure VIc) shows bounds for world welfare using the homogeneous methodology, while Figure VIId) shows the nonparametric bounds for world welfare for $\gamma = 0.5$ for both the baseline estimates and the extreme Gini estimates. Since the extreme Ginis methodology can use only Gini-based bounds, I use the Gini-based baseline bounds for comparison. These estimates are remarkably close to the baseline estimates, and yield the same implications; for the homogeneous survey selection it is possible to conclude that world welfare rose for every decade for $\gamma = 0.5$, and we see from Table I that any rise in inequality could have destroyed no more than 17% of potential welfare growth. The fact that the extreme Gini estimates largely coincide with the baseline estimates is not unexpected when one considers the great width of the Gini bounds, which dwarfs most empirically plausible ranges of the Gini coefficient.¹⁶

6.4 Alternative GDP data: Treatment of Chinese GDP

In 2007, in the wake of concluding a series of price surveys in the developing world, the World Bank revised the prices it used in its purchasing-power-parity adjustments, which led to major changes in its GDP series in the World Development Indicators, in particular, the lowering of Chinese and Indian GDP by 40% and 35% respectively. This development has been reviewed in the popular press (The Economist: Nov. 29, 2007; Dec. 19, 2007). The revision has been criticized, in particular on the grounds that it considered prices in urban China only. Penn World Tables version 7 fully incorporates these PPP revisions, but, mindful

¹⁶In fact, the extreme Gini bounds are sometimes narrower than the baseline bounds. This is because using Gini coefficients increases the number of country-years with surveys, thus replacing very conservative imputation procedures with much narrower nonparametric bounds for these country-years.

of the controversy of the new Chinese PPP estimates, reports two estimates for China: one based exclusively on Chinese national income accounts (version 1) and the 2005 ICP price survey for PPP adjustment, and the other one including some further PPP adjustments to compensate for the potentially nonrepresentative geographical character of the Chinese price surveys in the 2005 ICP (version 2). In my analysis, I have used the version 2 China series from the Penn World Tables because it delivers more conservative results. Figure VIe) shows bounds for world welfare using the version 1 series for Chinese GDP. We see from Table I that using the version 1 series actually strengthens my conclusion: world welfare rises by at least 101% from 1970 to 2006, and in particular, rises by at least 107% of uniform growth, suggesting that if we use the version 1 series, we could actually reject the hypothesis that inequality rose. However, such a rejection would not be robust to alternative extrapolation and interpolation methodologies, so I interpret this result cautiously.

6.5 Alternative GDP data: World Bank GDP

To check for robustness to the source of GDP more radically, I re-estimate world welfare and inequality measures using World Bank estimates of GDP from the World Development Indicators (hereafter WB). Figure VI f) presents the WB welfare estimates for $\gamma = 0.5$ along the baseline results. The bounds are extremely close to each other. One may conjecture that the PWT GDP series might do a better job of describing GDP earlier in the sample period while the WB GDP should do a better job later in the sample period, which would mean that welfare rose by less than either set of bounds would imply separately. We see that even taking the outer envelope of the nonparametric bounds for PWT GDP and WB GDP, we can conclude that welfare rose over the sample period, and we can establish more restrictive hypotheses as well (e.g. welfare rose from 1990 to 2006). Hence the WB series does not substantially change our results.

6.6 Alternative GDP data: World Bank Survey Means

As discussed in Section 4.1, there is disagreement in the literature on whether to combine national accounts GDP with survey data on inequality, or to use the survey mean as a measure of the mean of the income distribution. In particular, Milanovic (2005) and Chen and Ravallion (2001) use survey means in their calculations of world poverty and inequality. In this paper, I have used national accounts data in order to 1) remain comparable to most of the literature on the evolution of the world distribution of income and to the growth literature, and 2) avoid problems relating to the unavailability of survey means for many countries and years, given that the coverage of national accounts is nearly universal. In this section, I investigate the robustness of my results to the use of survey means in place of national accounts GDP.

I use the data on survey means from the World Bank's poverty calculator, PovCalNet, which is the most complete and consistent panel of survey means that I am aware of. However, even this panel does not match the nearly complete coverage of the Penn World Tables. For 79 major countries (including China, India, Nigeria, Argentina, Mexico and the former Soviet Union) I can obtain survey mean data from 1990 to 2004, extrapolating and interpolating the survey means using the methods of section 4.4. I use the World Bank national accounts data to construct a comparison sample of the same 79 countries during the same time period. There are substantial differences between the national accounts and the survey means, which typically result in lower income and slower growth in the survey means than in the national accounts. For example, China's annual rate of growth is more than 1 percentage point smaller if computed using survey means than using national accounts. There are many explanations for these differences (Deaton 2005, 2010), such as intentional and unintentional survey misreporting, problems in monetizing in-kind income, and inappropriate national accounting.

Figure VII) presents the sharp upper and lower bounds for the welfare series computed for the 79-country composite using 1) the survey means, and 2) Penn World Table national accounts data for the period 1990-2004. The increase and the growth rate of welfare

in this sample is much lower than for the baseline because the sample covers a much shorter period of time – in fact, the average annualized growth rate in the 79-country sample is 1.6% per year, while the average annualized growth rate in the baseline sample is 1.77% a year. It is clear that whether one chooses to use national accounts or survey means, world welfare rises unambiguously during this period, and, in fact, would rise unambiguously if we restricted our analysis to some subperiods of this data, such as 1990-2000. Table I shows the minimum absolute rise in welfare and the minimum growth rates for these welfare series. In particular, it is clear that in this subsample of countries, and even using survey means, we can actually reject the hypothesis that inequality rose within the subsample of countries because the lower bound of the ratio of welfare growth to per capita GDP growth exceeds 100%.

6.7 Accounting for nonresponse and nonrepresentativeness

A major problem with the survey data is that the people who respond to the surveys might systematically differ from people who do not. This concern is raised in Deaton (2005), who argues that falling response rates to many household surveys, and the large discrepancy between the national accounts means and the survey means, are problems of the first magnitude for the validity of the surveys. A particularly worrisome problem is that rich people in developing countries are systematically not covered by surveys (e.g. they live in gated communities to which surveyors have no access, or they openly lie to the surveyors for fear that their truthful answers might be given to the government). Atkinson, Piketty and Saez (2009) make the general case for the importance of inequality at the very top of the distribution in calculations of global inequality, and Banerjee and Piketty (2005) argue that failure of the Indian NSS to cover the top of the Indian income distribution may account for as much as 20%-40% of the much-documented growing gap between Indian national accounts and survey means in the NSS.

Given that the Atkinson welfare measure is decomposable, a simple method to generate sharp bounds given systematic nonresponse is to divide the population into two parts,

respondents and nonrespondents, and combine the bounds for respondents with worst-case assumptions about nonrespondents. Specifically, suppose that the survey represents a fraction λ of the population, who have mean income μ_c . Note that λ is bounded below by the response rate, but may in fact be larger than the response rate if it is possible to adjust for nonresponse within the survey. Let $z := \mu_c/\mu$, and let $A_{LB}(\gamma)$ and $A_{UB}(\gamma)$ be the upper and lower bound for the Atkinson inequality index computed on the basis of the survey data without adjusting for nonrepresentativeness. Then, the sharp nonrepresentativeness-adjusted bounds for $A(\gamma)$ are as follows:

$$\text{lower bound: } \hat{A}_{LB}(1 - \alpha) = 1 - \left[\lambda (z)^\alpha (1 - A_{LB}(1 - \alpha))^\alpha + (1 - \lambda) \left(\frac{1 - \lambda z}{1 - \lambda} \right)^\alpha \right]^{1/\alpha}$$

$$\text{upper bound: } \hat{A}_{UB}(1 - \alpha) = 1 - \lambda^{1/\alpha} z (1 - A_{UB}(1 - \alpha))$$

with the restriction that $\lambda z < 1$. This restriction always holds if $z < 1$. It is apparent that $\hat{A}_{LB}(1 - \alpha)$ increases in (λ, z, α) , and $\hat{A}_{UB}(1 - \alpha)$ decreases in these variables, so as surveys become less representative (λ falls) and as survey mean income falls further below the national accounts GDP (z falls), the confidence intervals widen. Therefore, if we want to make the broader assumption that $z \geq \bar{z}$ and $\lambda \leq \bar{\lambda}$ for some \bar{z} and $\bar{\lambda}$, we just compute the upper bound, and set the lower bound at $\hat{A}_{LB}(1 - \alpha) = A_{LB}(1 - \alpha)$.

The variable λ has the intuitive meaning of the response rate. The variable z is the ratio of the mean income of respondents to true mean income. If λ is equal to unity, and $z < 1$, then z should be interpreted as the fraction of national income accounted for by the survey because $1 - z$ is the fraction of income owned by the small number of unsurveyed super-rich individuals.

Unfortunately, the WIID does not contain data on the response rates to the surveys (either in the database itself or in the documentation), and reports mean survey incomes very sporadically (only one survey mean is ever reported for China). Therefore, for the purposes of

this robustness check, I will assume common values λ and z for the entire world, and compute bounds on their basis. Such bounds will, in a sense, be more informative, as these bounds will be valid for all values of λ and z higher than the values chosen, whereas attempting to retrieve z from the WIID would introduce additional sources of error. I let λ take on values in $\{0.80, 0.90, 1\}$, following Korinek et. al. (2003), who give a reasonable range for nonresponse of 10%-30%. I let z take on values in $\{0.75, 0.875, 1\}$ following Banerjee and Piketty (2005), who argue that the top 1% of the Indian income distribution hold approximately 12% of Indian national income, and following Atkinson, Piketty and Saez (2009) who document that the top 1% of the US income distribution hold a share of national income approaching 25%.

First, suppose that nonresponse is negligible ($\lambda = 1$) and is coming from a small group of super-rich individuals who are not captured by the surveys. Figure Xa) plots the welfare bounds for the baseline case, for $z = 0.875$ and for $z = 0.75$. For either value of z , we can be confident that welfare rose. We can barely fail to reject that welfare fell for the outer envelope of all the bounds presented. Thus, under the assumption that the discrepancy between survey mean income and national accounts income can be explained by part of the growth enriching a very small and very rich minority, with full response otherwise, we can conclude that welfare rose even if the fraction of national income held by this minority rose, for instance, from 0 to 25% over the sample period. However, the bounds are consistent with rising measured inequality and a rising share of income held by the super-rich substantially eroding welfare gains relative to uniform growth: Table I shows that if $z = 0.875$, the realized welfare growth may have been only 68% of per capita GDP growth, and if $z = 0.75$, the realized welfare growth may have been only 43% of per capita GDP growth. These are of course much more demanding sets of bounds because they try to not only capture the uncertainty in the population measured by surveys, but also attempt to account for factors changing the composition and size of this population.

We now consider the more general case when only λ of the population is covered by any household survey, and its mean income is z of the true mean income. This situation

is very general; in particular, there are no assumptions at all on exactly how nonrandom the sampling is, and how large or small inequality is in the fraction of the population not surveyed. However, if survey means are low and nonresponse is high, the implication is that an increasingly large and economically significant part of the population is not being covered, which leads to wide nonparametric bounds and an inability to reach any conclusions without further assumptions. Figure Xb) presents the baseline bounds as well as the bounds for $(\lambda, z) = (0.9, 0.875)$, $(\lambda, z) = (0.9, 0.75)$ and $(\lambda, z) = (0.8, 0.875)$. It is clear that we are confident for each set of the bounds that world welfare rose, but only barely so in some cases. In particular, setting the nonresponse rate to 30% ($\lambda = 0.7$), which is deemed possible by Korinek et al. (2003), would prevent us from rejecting the hypothesis that welfare fell. While we can be confident that welfare rose if the nonresponse rate is sufficiently small and if the mean income of respondents is sufficiently close to total mean income, we cannot be confident that welfare rose essentially for any higher nonresponse rates, which might nevertheless be plausible.

7 Conclusion

In this paper, I presented formulae for sharp, nonparametric bounds for typical measures of inequality and welfare that can be computed from standard summary statistics of income distributions that are routinely provided to the public. These bounds are valid independently of the functional form of the underlying distribution of income. Hence, these bounds illustrate exactly the extent of knowledge about the inequality measures in question that we gain from our data; they render moot any questions about appropriate assumptions for the form of within-country or within-fractile income distributions and they help focus the debate over whether we have seen improvement in living standards onto issues of the validity of the data from which they were computed.

Using national accounts GDP estimates and WIID survey data, the nonparametric bounds imply that for any series compatible with these data, welfare must have risen, and risen substantially, but it is not possible to conclude whether world inequality rose or fell.

This claim is supported through sensitivity analysis over the exact use of the WIID data as well as over the methodology used to compute the national accounts GDP estimates. In particular, it appears reasonable to conclude that notwithstanding the incomparability problems of surveys in the WIID, and notwithstanding recent issues in adjusting national accounts for PPP, welfare has risen over the period 1970-2006, as well as over most shorter periods of interest within this time. It is also unlikely that sampling error or vagueness in the income concepts used in administering the surveys could change this result, so long as surveys are fully representative of whatever income concept they are measuring. For most of my robustness checks, I can rule out substantial rises in inequality, but if I use extremely conservative interpolation procedures, or if I attempt to account for the possible deterioration of survey representativeness of the population, substantial rises in inequality that destroy around 50% of potential welfare growth are compatible with the estimated bounds.

Hence, the major challenge that remains to the claim that welfare rose consistently over the period 1970-2006 is the concern that the WIID surveys are significantly nonrepresentative of the underlying population because of selective nonresponse. While under some assumptions, this concern does not overturn our conclusion, for some possible estimates of the extent of nonresponse, we obtain that further assumptions are needed to conclude anything about the path of welfare during the sample period. It is therefore of great importance to accurately gauge the reliability of the survey data we are using in order to reach conclusions about what happened to welfare in recent times.

References

- [1] Aitchison J. and J.A.C. Brown. *The Lognormal Distribution*. Cambridge, UK: Cambridge University Press. 1957
- [2] Ahluwalia, Montek S., Nicholas G. Carter, and Hollis B. Chenery. "Growth and Poverty in Developing Countries." *Journal of Development Economics* 6, no. 3: 299-341. 1979.

- [3] Atkinson, A. B. "On the Measurement of Inequality." *Journal of Economic Theory* 2, no. 3: 244-263. 1970.
- [4] Atkinson, Anthony B., and Andrea Brandolini. "Promise and Pitfalls in the Use of 'Secondary' Data-Sets: Income Inequality in OECD Countries As a Case Study." *Journal of Economic Literature* 39, no. 3: 771-799. 2001.
- [5] Atkinson, Anthony B., Thomas Piketty and Emmanuel Saez. "Top Incomes in the Long Run of History." NBER Working Paper #15408. 2009.
- [6] Atkinson, Anthony B., and Andrea Brandolini. "On Analysing the World Distribution of Income," forthcoming in the *World Bank Economic Review*. 2010.
- [7] Banerjee, Abhijit, and Thomas Piketty. "Top Indian Incomes, 1922-2000." *World Bank Economic Review* 19, no. 1: 1-20. 2005.
- [8] Bhalla, Surjit S. *Imagine there's no country: Poverty, inequality, and growth in the era of globalization*. Washington, D.C.:Penguin. 2002.
- [9] Bourguignon, Francois, and Christian Morrisson. "Inequality among World Citizens: 1820-1992." *American Economic Review* 92, no. 4: 727-744. 2002.
- [10] Chen, Shaohua, and Martin Ravallion. "How Did the World's Poorest Fare in the 1990s?." *Review of Income and Wealth* 47, no. 3: 283-300. 2001.
- [11] Chen, Shaohua, and Martin Ravallion. "The Developing World Is Poorer Than We Thought, but No Less Successful in the Fight against Poverty." *Quarterly Journal of Economics* 125, no. 4: 1577-1625. 2010.
- [12] Chiappori, Pierre-Andre, and Monica Paiella. "Relative Risk Aversion is Constant: Evidence from Panel Data," forthcoming *Journal of the European Economic Association*. 2006.

- [13] Chotikapanich, Duangkamon, William E. Griffiths, and D. S. Prasada Rao. "Estimating and Combining National Income Distributions Using Limited Data." *Journal of Business and Economic Statistics* 25, no. 1: 97-109. 2007.
- [14] Cowell, Frank A. *Measuring Inequality*, Oxford: Philip Allan. 1977
- [15] Cowell, Frank A. "Inequality Decomposition: Three Bad Measures." *Bulletin of Economic Research* 40, no. 4: 309-312. 1988.
- [16] Cowell, Frank A. "Grouping Bounds for Inequality Measures under Alternative Informational Assumptions." *Journal of Econometrics* 48, no. 1-2: 1-14. 1991.
- [17] Cowell, Frank. "Measuring Inequality" in Atkinson T. and Bourguignon F. eds. *Handbook of Income Distribution*, Elsevier. 2000
- [18] Deaton, Angus. "Measuring Poverty in a Growing World (or Measuring Growth in a Poor World)." *Review of Economics and Statistics* 87, no. 1: 1-19. 2005.
- [19] Deaton, Angus. "Price Indexes, Inequality, and the Measurement of World Poverty." *American Economic Review* 100, no. 1: 5-34. 2010.
- [20] Deininger, Klaus, and Lyn Squire. "A New Data Set Measuring Income Inequality," *World Bank Economic Review*, X, 565-591. 1996
- [21] Dikhanov, Yuri, and Michael Ward. "Evolution of the Global Distribution of Income, 1970-99," mimeo. 2001.
- [22] Galbraith, James K., and Hyunsub Kum. "Estimating the Inequality of Household Incomes: A Statistical Approach to the Creation of a Dense and Consistent Global Data Set." *Review of Income and Wealth* 51, no. 1: 115-143. 2005.
- [23] Gastwirth, Joseph L. "The Estimation of the Lorenz Curve and Gini Index." *Review of Economics and Statistics* 54, no. 3: 306-316. 1972.
- [24] Gibrat, R. *Les inégalités économiques*, Paris: Sirey. 1931.

- [25] Heston, Allan, Robert Summers, and Bettina Aten. Penn World Table Version 6.2, Center for International Comparisons at the University of Pennsylvania (CICUP). 2006
- [26] Kalecki, M. "On the Gibrat Distribution", *Econometrica*, 13, 161-170. 1945.
- [27] Kamien, Morton I. and Nancy L. Schwartz. *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*. New York: Elsevier. 1991.
- [28] Lopez, J. Humberto and Servén, Luis. "A Normal Relationship? Poverty, Growth and Inequality," World Bank Policy Research Working Paper 3814. 2006.
- [29] Maddison, Angus. *Historical Statistics of the World Economy: 1-2003 A.D.* <http://www.ggdc.net/maddison/> Date Accessed: June 1 2007.
- [30] Manski, Charles. *Identification Problems in the Social Sciences*. Cambridge, MA: Harvard University Press. 1995
- [31] McDonald, James B., and Michael R. Ransom. "An Analysis of the Bounds for the Gini Coefficient." *Journal of Econometrics* 17, no. 2: 177-188. 1981.
- [32] McDonald, James B., and Yexiao J. Xu. "A Generalization of the Beta Distribution with Applications." *Journal of Econometrics* 66, no. 1-2: 133-152. 1995.
- [33] Mehran, Farhad. "Bounds on the Gini Index Based on Observed Points of the Lorenz Curve." *Journal of the American Statistical Association* 70, no. 349: 64-66. 1975.
- [34] Milanovic, Branko. "True World Income Distribution, 1988 and 1993: First Calculation Based on Household Surveys Alone." *Economic Journal* 112, no. 476: 51-92. 2002.
- [35] Milanovic, Branko. *Worlds Apart: Measuring International and Global Inequality*. Princeton and Oxford: Princeton University Press. 2005.
- [36] Milanovic, Branko. "Global Inequality Recalculated and Updated: The Effect of New PPP Estimates on Global Inequality and 2005 Estimates." *Journal of Economic Inequality* 10, no. 1: 1-18. 2012.

- [37] Murray, David. "Extreme Values for the Gini Coefficients Calculated from Grouped Data." *Economics Letters* 1: 389-393. 1978.
- [38] Pareto, W. "*Cours d'Economie Politique*", F. Rouge, Lausanne. 1897.
- [39] Pinkovskiy, Maxim and Xavier Sala-i-Martin. "Parametric Estimations of the World Distribution of Income." NBER Working Paper #15433. 2009.
- [40] Ravallion, Martin, Johan A. Mistiaen, and Anton Korinek. "Survey nonresponse and the distribution of income." World Bank Policy Research Working Paper 3543. 2005.
- [41] Reddy, Sanjay G., and Camelia Minoiu. "Has World Poverty Really Fallen?" *Review of Income and Wealth* 53, no. 3: 484-502. 2007.
- [42] Sala-i-Martin, Xavier. "The Disturbing 'Rise' of Global Income Inequality." NBER Working Paper #8904. 2002a.
- [43] Sala-i-Martin, Xavier. "The World Distribution of Income (estimated from Individual Country Distributions)." NBER Working Paper #8933. 2002b.
- [44] Sala-i-Martin, Xavier. "The World Distribution of Income: Falling Poverty and . . . Convergence, Period." *Quarterly Journal of Economics* 121, no. 2: 351-397. 2006.
- [45] Salem, A. and T. Mount. "A Convenient Descriptive Model of Income Distribution: The Gamma Density", *Econometrica* 42, 1115-1128. 1974.
- [46] Schultz, T. Paul. "Inequality in the Distribution of Personal Income in the World: How It Is Changing and Why." *Journal of Population Economics* 11, no. 3: 307-344. 1998.
- [47] Sen, Amartya K. "Real National Income." *Review of Economic Studies* 43, no. 1: 19-39. 1976.
- [48] Singh, S. and G. Maddala. "A Function for Size Distribution of Incomes", *Econometrica* 44, 963-970. 1976.

- [49] Socio-Economic Database for Latin America and the Caribbean (CEDLAS and The World Bank). Date Accessed: September 1, 2009.
- [50] World Bank: World Development Indicators 2009. Accessed July 15, 2007 and June 1, 2008. <http://go.worldbank.org/U0FSM7AQ40>
- [51] Zhu, Feng. “A Nonparametric Analysis of the Shape Dynamics of the U.S. Personal Income Distribution: 1962-2000.” BIS Working Paper #184. 2005.

A Proofs of Various Propositions

A.1 Proof of Lemma 1 (3-2 Lemma):

Suppose that the increasing and convex curve $L(p)$ on $[0, p^*]$, where $p^* > 1$, is defined by three line segments, where the first segment is defined by $q = mp$, and the third segment is defined by $q = zp - (z - m)$, where $0 \leq m < z \leq \infty$. Let the second segment lie between the point $(\bar{p}, m\bar{p})$ on the first line and $(1 + \frac{\bar{q}-m}{z}, \bar{q})$ on the third line, so its equation is $M(p) := q = m\bar{p} + \frac{z(\bar{q}-m\bar{p})}{(z-m)+(\bar{q}-z\bar{p})}(p-\bar{p})$. Therefore, $L(p) = \max\left(mp, m\bar{p} + \frac{z(\bar{q}-m\bar{p})}{(z-m)+(\bar{q}-z\bar{p})}(p-\bar{p}), zp - (z - m)\right)$ on $[0, p^*]$. Let the value of the third line segment at p^* be defined as rp^* .

Then, the Gini constraint is

$$\int_0^{p^*} L(p) dp = K_1 + K_2 [(\bar{q} - m)(1 - \bar{p})] = \bar{G}$$

for some K_1 and K_2 that are constants in \bar{p} and \bar{q} , so the Gini constraint is equivalent to $\bar{q} = \frac{S}{1-\bar{p}} + m$, for some $S \in [0, \infty]$. Note that \bar{p} is then constrained to lie in the interval $[0, \tilde{p}]$ for some $\tilde{p} < 1$, since $\bar{q} \leq rp^*$

The objective is given by

$$O = \int_0^{p^*} (L'(p))^\alpha dp = \bar{p}m^\alpha + z^{-(1-\alpha)}(\bar{q} - m\bar{p})^\alpha (z(1 - \bar{p}) + \bar{q} - m)^{1-\alpha} + z^\alpha \left(K - \frac{\bar{q} - m}{z}\right)$$

If and only if it can be shown that O is minimized by some $\bar{p} \in \{0, \tilde{p}\}$, or that the problem of minimizing O in \bar{p} subject to the Gini constraint yields a corner solution, then the lemma is proved; since then the line $M(p)$ dominates either the first or the third line on $[0, p^*]$. Hence, the lemma is equivalent to the problem

$$\min_{\bar{p} \in [\tilde{p}, 1]} \left\{ (1 - p)m^\alpha + z^{-(1-\alpha)} \left(\frac{S}{p}\right) \left[\left(1 + \frac{z}{S}p^2\right)^\alpha \left(1 + \frac{m}{S}p^2\right)^{1-\alpha} - 1 \right] \right\}$$

having a corner solution (where we replace $p = 1 - \bar{p}$). Now, let $w = p^2/S \in (0, \infty)$. Then, we obtain.

$$D_p = z^\alpha \left[\left(\frac{1+mw}{1+zw}\right)^\alpha \left(2 \left((1-\alpha) + \alpha \frac{m}{z} \frac{1+zw}{1+mw}\right) - 1\right) + \frac{1}{zw} \left(1 - \left(\frac{1+mw}{1+zw}\right)^\alpha\right) - \left(\frac{m}{z}\right)^\alpha \right], \text{ and}$$

$$D_p^2 = -\frac{2}{p} \left[D_p + z^\alpha \left(\left(\frac{m}{z}\right)^\alpha - \left(\frac{1+mw}{1+zw}\right)^\alpha \left(1 - \frac{1}{z} \left(\frac{\alpha(z-m)}{1+mw}\right) \left(1 + \frac{2(1-\alpha)(z-m)w}{(1+mw)(1+zw)}\right)\right) \right) \right]$$

A necessary and sufficient condition for a corner solution to the optimization problem is that $D_p = 0 \Rightarrow D_p^2 < 0$, or that

$$\begin{aligned} \left(\frac{m}{z}\right)^\alpha &= \left(\frac{1+mw}{1+zw}\right)^\alpha \left(2 \left((1-\alpha) + \alpha \frac{m}{z} \frac{1+zw}{1+mw}\right) - 1\right) + \frac{1}{zw} \left(1 - \left(\frac{1+mw}{1+zw}\right)^\alpha\right) \\ &\Rightarrow \left(\frac{m}{z}\right)^\alpha \geq \left(\frac{1+mw}{1+zw}\right)^\alpha \left(1 - \frac{1}{z} \left(\frac{\alpha(z-m)}{1+mw}\right) \left(1 + \frac{2(1-\alpha)(z-m)w}{(1+mw)(1+zw)}\right)\right) \end{aligned}$$

which is trivially true.

Hence, the minimum of the problem is achieved at the boundary, and the optimizing segment of the curve has only one interior corner rather than two. It is obvious by induction that for any piecewise linear curve with z corners between two consecutive constrained points, there exists another piecewise linear curve that

has only one corner between these points, satisfies the constraints of the original curve, and attains a weakly smaller value of the program.

In particular, for any piecewise linear Lorenz curve L with finitely many kinks, there exists a sequence $\{a_i^-, a_i^+\}_{i=1}^k$ such that $a_i^- \leq m_i \leq a_i^+ \leq a_{i+1}^- \forall i = 1, \dots, k$, and the Lorenz curve given by

$$\hat{L} = \max_{i=1, \dots, k} \left\{ \max \left\{ a_i^- (p - p_i) + q_i, a_i^+ (p - p_{i+1}) + q_{i+1} \right\} \right\}$$

satisfies the constraints and attains a weakly smaller value of the objective than does the curve L .

A.2 Proof of Proposition 2:

The problem is

$$\inf_{L \in \mathfrak{L}} \int_0^1 (L'(p))^\alpha dp \text{ st. } 1) \bar{G} = \int_0^1 L(p) dp, 2) \forall i = 1, \dots, k, L(p_i) = q_i$$

Since the functional $\int_0^1 (L'(p))^\alpha dp$ is bounded below by zero, it must be the case that $S := \inf_{L \in \mathfrak{L}} \int_0^1 (L'(p))^\alpha dp \in \mathbb{R}$. Moreover, there must be a sequence $\{L_i\} \in \mathfrak{L}_c$ such that $\lim_{i \rightarrow \infty} \int_0^1 (L'_i(p))^\alpha dp = S$. Now, since the function $L'_i(p)$ is Riemann integrable, it must be the case that for any ε , and for any i , there exists $\tilde{L}_i \in \mathfrak{L}_c$ such that \tilde{L}_i is piecewise linear with finitely many corners, and $\left| \int_0^1 (L'_i(p))^\alpha dp - \int_0^1 (\tilde{L}'_i(p))^\alpha dp \right| \leq \varepsilon$.¹⁷ Hence, let $\{\varepsilon_i\} \in \mathbb{R}$ be a sequence such that $\lim_{i \rightarrow \infty} \varepsilon_i = 0$, and let $\{\tilde{L}_i\}$ be a sequence of piecewise linear functions with finitely many corners such that $\left| \int_0^1 (L'_i(p))^\alpha dp - \int_0^1 (\tilde{L}'_i(p))^\alpha dp \right| \leq \varepsilon_i$. Then, $\lim_{i \rightarrow \infty} \int_0^1 (\tilde{L}'_i(p))^\alpha dp = S$.

Now, by Lemma 1, for every piecewise linear Lorenz curve $\tilde{L}_i \in \mathfrak{L}_c$, there exists a piecewise linear Lorenz curve $\hat{L}(\cdot; \mathbf{a}_i) \in \mathfrak{L}_c$ given by

$$\hat{L}(p; \mathbf{a}_i) = \max_{s=1, \dots, k} \left\{ \max \left\{ a_{i,s}^- (p - p_s) + q_s, a_{i,s}^+ (p - p_{s+1}) + q_{s+1} \right\} \right\}$$

for some $\mathbf{a}_i = \{a_{i,s}^-, a_{i,s}^+\}_{s=1}^k$ such that $a_{i,s}^- \leq m_s \leq a_{i,s}^+ \leq a_{i,s+1}^-$, such that $\int_0^1 (\hat{L}'(p; \mathbf{a}_i))^\alpha dp \leq \int_0^1 (\tilde{L}'_i(p))^\alpha dp$.

Hence, $\lim_{i \rightarrow \infty} \int_0^1 (\hat{L}'(p; \mathbf{a}_i))^\alpha dp \leq S$, and by definition of infimum, $\lim_{i \rightarrow \infty} \int_0^1 (\hat{L}'(p; \mathbf{a}_i))^\alpha dp = S$. Now, let the set \mathbf{A} be the set of all \mathbf{a}_i satisfying the restriction $a_{i,s}^- \leq m_s \leq a_{i,s}^+ \leq a_{i,s+1}^- \forall i = 1, \dots, k$, and note that this is a closed subset of the compact set $\bar{\mathbb{R}}_+^{2k}$, and is therefore compact. Hence, the sequence $\{\mathbf{a}_i\}_{i=1}^\infty$ has a convergent subsequence, $\{\mathbf{a}_{i(k)}\}_{k=1}^\infty$, which converges to a limit \mathbf{a} . Finally, observe that by the definition of $\hat{L}(p; \mathbf{a}_i)$, the integral $\int_0^1 (\hat{L}'(p; \mathbf{a}_i))^\alpha dp$ is continuous in \mathbf{a}_i , so

$$S = \lim_{i \rightarrow \infty} \int_0^1 (\hat{L}'(p; \mathbf{a}_i))^\alpha dp = \lim_{k \rightarrow \infty} \int_0^1 (\hat{L}'(p; \mathbf{a}_{i(k)}))^\alpha dp = \int_0^1 (\hat{L}'(p; \mathbf{a}))^\alpha dp$$

and $L(p; \mathbf{a})$ is a Lorenz curve that attains the infimum value S . Since $L(p; \mathbf{a})$ is defined by $2k$ parameters, its coefficients $\{a_i^-, a_i^+\}_{i=1}^k$ can be solved for using standard numerical methods.

¹⁷I am very grateful to Paolo Siconolfi for help with this part of the proof.

A.3 Proof of Proposition 3:

The 3-2 Lemma implies that if there is only a Gini constraint, the Atkinson is maximized by a Lorenz curve $L(p)$ with only one interior corner (p, q) .

The Gini of this Lorenz curve is given by $G = p - q$, and the Atkinson is given by $1 - q^\alpha p^{1-\alpha} - (1 - q)^\alpha (1 - p)^{1-\alpha}$, so the parameters p and q of the optimal curve are given by

$$\bar{p} = \arg \min_{p \in [G, 1]} \left\{ (p - G)^\alpha p^{1-\alpha} + (1 - p + G)^\alpha (1 - p)^{1-\alpha} \right\}$$

and $\bar{q} = p - G$. The second derivative of the minimand is given by

$$D_p^2 = -\alpha(1 - \alpha) \bar{q}^{\alpha-1} \bar{p}^{-\alpha} \left[\left(\frac{\bar{q}}{\bar{p}} + \frac{\bar{p}}{\bar{q}} - 2 \right) + \left(\frac{1 - \bar{q}}{1 - \bar{p}} + \frac{1 - \bar{p}}{1 - \bar{q}} - 2 \right) \right]$$

which is negative, so any minimum must be a corner solution, and the maximized Atkinson is given by $\max \left(G, 1 - (1 - G)^{\frac{1-\alpha}{\alpha}} \right)$.

Table I: Bounds on Increase and Growth of World Welfare

	Diff. Lower Bound	Diff. Upper Bound	Growth Lower Bound	Growth Upper Bound	Growth, Lower Bound	Growth, Upper Bound
			Pct	Pct	Relative to Uniform Pct	Relative to Uniform Pct
IIIa) Baseline	2626	3361	88	127	93	134
IIIa) Gamma=0.9	1257	3234	65	330	69	350
IVa) Gini Bounds Only	1949	3937	65	190	69	201
IVa) Fractile Bounds, Sharp	2501	3579	83	147	88	156
IVa) Fractile Bounds, Crude	2484	3623	82	152	87	161
Va) Baseline, lowest income at 0.2 of lowest fractile mean	2610	3331	87	124	92	131
Va) Gamma=1.25, lowest income at 0.2 of lowest fractile mean	1052	2035	67	167	71	176
Vb) Gamma=1.5, lowest income at 0.2 of lowest fractile mean	382	1585	29	172	31	182
VIb) Extreme Interpolation	1731	4011	56	167	59	176
VIc) Homogeneous Survey Choice	2365	3541	78	132	83	140
VIc) Maximum Ginis	1510	3760	50	190	53	201
VIc) Alternative PWT Series for China	2804	3507	101	142	107	150
VIc) World Bank GDP	2548	3222	90	128	96	136
VII) Survey Means	187	305	26	47	107	193
VII) PWT GDP for Survey Means sample	780	1171	35	58	143	238
Xa) Super-Rich have <12.5 percent	1910	3695	64	160	68	169
Xa) Super-Rich have <25 percent	1214	4024	40	203	43	215
Xb) <10 percent Nonresponse, >87.5% Mean Income Ratio	984	4133	33	221	35	234
Xb) <20 percent Nonresponse, >75% Mean Income Ratio	154	4524	5	306	5	324
Xb) <10 percent Nonresponse, >87.5% Mean Income Ratio	419	4399	14	275	14	291

Note: Table I summarizes the results from all the graphs. The number for every row indicates the graph from which the relevant bounds are taken. The first two columns present bounds on the absolute increase of world welfare (the certainty equivalent of the income distribution) in dollars. The next two columns present bounds on the aggregate growth rate of world welfare in percent. The last two columns present bounds on the ratio of the aggregate growth rate of world welfare to the growth rate of world GDP per capita in percent (so 100 percent would correspond to welfare growing by as much as GDP per capita, or uniform growth).

The asterisked rows correspond to estimates from Gini-based bounds, which are much wider than the baseline bounds and should be compared to the series "Gini Based Baseline" rather than to the baseline series.

The variations are:

IIIa) Baseline: PWT 7 GDP, WIID surveys selected by the procedure described in section 4. Microdata used where available as described in section 4.

For country-years with no survey data: interpolate the bounds as the outer envelope of the bounds of the adjacent years with survey data;

extrapolate the bounds as the outer envelope of horizontal and linear extrapolation of each bound;

impute bounds for countries with one or no surveys as the outer envelope of the bounds for all countries in the given region at any time in the sample period.

Bounds based on Gini coefficient and fractile shares for all country-years with survey data. Atkinson parameter γ is 0.5

IIIa) $\gamma=0.9$: Same as Baseline, but $\gamma=0.9$

IVa) Gini Bounds Only: Same as Baseline, but bounds based on only Gini coefficient.

IVa) Fractile Bounds, Sharp: Same as Baseline, but bounds based on only fractile shares (exact computation)

IVa) Fractile Bounds, Crude: Same as Baseline, but bounds based on only fractile shares (approximate computation)

Va) Baseline, lowest income at 0.2 of lowest fractile mean: Same as Baseline, but bounds computed under assumption that the lowest income is 0.2 of the lowest fractile mean.

Va) $\gamma=1.25$, lowest income at 0.2 of lowest fractile mean: Same as Baseline, but bounds computed under assumption that the lowest income is 0.2 of the lowest fractile mean.

Vb) $\gamma=1.5$, lowest income at 0.2 of lowest fractile mean: Same as Baseline, but bounds computed under assumption that the lowest income is 0.2 of the lowest fractile mean.

VIb) Extreme Interpolation: Same as Baseline, except for country-years with no survey data, impute the outer envelope of the bounds over the country over the sample period (if country has more than one survey), or impute the outer envelope of the bounds over the entire region over the sample period (if country has one or no surveys).

VIc) Homogeneous Survey Choice: Same as Baseline, except choose longest survey series from the same source

and with same equalization and income concept for each country from WIID.

VIId) Maximum Ginis: Same as Baseline, except for each country-year with any WIID survey data on a national scale, use envelope of all Gini-based bounds implied by WIID surveys in country-year.

VIe) Alternative China GDP: Same as Baseline, except use the PWT 7 version 1 GDP series, which bases Chinese GDP exclusively on the 2005 PPP revision without further adjustments.

VIf) World Bank GDP: Same as Baseline, except use GDP data from World Development Indicators, 2008.

VII) Survey Means: Same as Baseline, except use survey means from the WIID instead of PWT GDP for all years in which they are available.

VII) PWT GDP for Survey Means sample: Same as Baseline, except restrict to the sample of countries with survey mean data.


Xa) Super-Rich have <12.5 percent: Assume WIID statistics exclude a set of people of measure zero who own at most 12.5% of national income respectively.

Xa) Super-Rich have <25 percent: Assume WIID statistics exclude a set of people of measure zero who own at most 25% of national income respectively.


Xb) <10 percent Nonresponse, >87.5% Mean Income Ratio: Assume WIID statistics exclude nonrespondents, with respondents' mean income at least 87.5% of the overall mean income but the distribution of this income among the nonrespondents is arbitrary.

Xb) <20 percent Nonresponse, >87.5% Mean Income Ratio: Assume WIID statistics exclude nonrespondents, with respondents' mean income at least 87.5% of the overall mean income but the distribution of this income among the nonrespondents is arbitrary.

Xb) <10 percent Nonresponse, >75% Mean Income Ratio: Assume WIID statistics exclude nonrespondents, with respondents' mean income at least 75% of the overall mean income but the distribution of this income among the nonrespondents is arbitrary.




Ia)




Ib)




Ic)



IIa)




IIb)




IIc)

Note: The source for Figures IIa and IIb is the Penn World Tables 6.2. The source for Figure IIc is the WIID inequality database. In Figure IIc, fraction covered directly is fraction of world population in given year in countries with surveys used from WIID. Fraction in core is fraction of world population in given year in countries with at least one earlier survey and at least one later survey.




IIIa)




IIIb)

Note: Figure IIIa (resp., Figure IIIb) shows the bounds on global Atkinson welfare indices (resp., global Atkinson inequality indices) for parameters $\gamma = 0.5$ and $\gamma = 0.9$. Each pair of identically formatted lines represents a pair of bounds. The bounds are sharp given the fractile shares and Gini coefficient for any household survey used, subject to interpolation, extrapolation and imputation for countries and years without inequality data. Mean incomes to construct Figure IIIa are taken from the Penn World Tables 6.2.




IVa)

Note: Figure IVa presents sharp global Atkinson welfare bounds based on different statistics from the household surveys in the WIID database. Note that most of the upper bounds coincide.




IVb)

Note: Figure IVb presents sharp bounds on global Atkinson welfare growth based on different statistics from the household surveys in the WIID database.




Va)




Vb)


Note: Figures Va) and Vb) presents sharp global Atkinson welfare bounds for higher values of the parameter γ . Different series have their own y-axes because of the difference in scales in the indices.




VIa)




VIb)




VIc)



VI d)




VIe)



VI f)


Note: Figure VIa checks robustness with respect to sampling error by assuming the underlying inequality distributions in WIID are lognormal and comparing the bounds based on the predicted quintile shares with bounds that incorporate sampling error in the quintile shares. Figure VIb checks for robustness to imputing country and regional extreme values of the Atkinson welfare index for countries and years with missing data instead of extrapolation. Figure VIc checks for robustness to constraining all surveys selected from WIID for a given country to come from the same source. Figure VI d checks for robustness to using the envelope of the bounds for the highest and lowest

Gini coefficient provided for every country-year with inequality data; the baseline series is also based only on Gini coefficients. Figure VIe checks for robustness to using GDP from the 2008 World Development Indicators. Figure VI f checks for robustness to using an alternative GDP series for China reported in the Penn World Tables.



VII)

Note: Figure VII checks robustness for the assumption that survey means provide a better measure for mean income than do national accounts GDP. I compute bounds for a sample of 79 countries, which includes the largest and most populous countries in the developing world. The series in red presents bounds computed using survey means from the World Bank’s PovCalNet website, while the series in blue presents bounds computed using PWT 7 GDP. The left vertical axis is to be used with the survey means series, while the right vertical axis is to be used with the national accounts series.




VIII)

Note: Figure VIII shows sharp bounds for global Atkinson inequality indices for three different assumptions about interpolation, extrapolation and imputation of the country Atkinson inequality index bounds when inequality data for the given country-years is not available. The three different assumptions are:


Baseline: interpolate the bounds as the outer envelope of the bounds of the adjacent years with survey data. Extrapolate the bounds as the outer envelope of horizontal and linear extrapolation of each bound. Impute bounds for countries with one or no surveys as the outer envelope of the bounds for all countries in the given region at any time in the sample period.

Regional Means: same as Baseline, but impute upper (lower) bound for countries with one or no surveys as the average of the upper (lower) bounds for all countries in the given region in that year.

Regional Means plus Interpolated Bounds: same as Regional Means, but for countries with more than one survey, interpolate each bound linearly and extrapolate the bounds horizontally.




IXa)




IXb)

Note: Figure IXa shows breadth of coverage when selecting only surveys from the same source from the WIID database; Figure IXb shows breadth of coverage when selecting all surveys with national coverage from WIID. Fraction covered directly is fraction of world population in given year in countries with surveys used from WIID according to the selection procedure. Fraction in core is fraction of world population in given year in countries with at least one earlier survey and at least one later survey selected by the procedure.



X.a)



X.b)

Note: Figure Xa checks robustness for the assumption that the WIID statistics exclude a set of people of measure zero who own <12.5% and <25% of GDP respectively. Figure Xb checks robustness for the assumption that the WIID statistics exclude nonrespondents, and that the ratio between respondents' mean income and true mean income is Z.