

Global inequality and global macroeconomics[☆]

James K. Galbraith^{*}

LBJ School of Public Affairs, The University of Texas at Austin, Austin, TX 78713-8925, USA

Received 1 August 2006; received in revised form 1 February 2007; accepted 1 April 2007

Available online 7 May 2007

Abstract

This paper presents evidence for a common global pattern in the movement of inequality in national structures of pay, over the years 1963 to 1999. We find a worldwide pattern of declining inequality from 1971 until 1980, followed by a long and sharp period of increasing inequality from 1981 through the end of the century. The existence of a global pattern suggests that the study of inequality, long associated with the disparate effects of technology, trade in local or national labor markets and with national policy choices, would be better treated as a branch of a global macroeconomics, associated with the breakdown of Bretton Woods in 1971–73 and with the onset of the global debt crisis in 1981–82. The work is based on data sets developed by the University of Texas Inequality Project.

© 2007 Society for Policy Modeling. Published by Elsevier Inc. All rights reserved.

JEL classification: F02; J31; E0; P5

Keywords: Inequality; Macroeconomic forces; Globalization; Development; Kuznets; Inverted “U” hypothesis

1. Introduction

Inevitably the study of economic inequality is conditioned by the preconceptions of the economist, and the taxonomic history of the discipline powerfully reinforces those preconceptions. The classical economics separated theories of value and distribution from theories of growth; much of the functional analysis of wages, profits and rent could be carried out with reference to a stationary or steady state. Though in Marx distribution and growth did merge, the marriage was brief, and in the neoclassical revival their divorce was complete.

[☆] Reprinted with permission from David Held and Ayse Kaya (Eds.), *Global Inequality*, Ralph Miliband Lecture Series on Inequalities at the London School of Economics. Cambridge: Polity Press, 2007, 148–175.

^{*} Tel.: +1 512 471 1244; fax: +1 512 471 1835.

E-mail address: galbraith@mail.utexas.edu.

Keynes and Kalecki rejected the determining power of the labor market over wages and employment, and their theory of investment and profits was bound up with prospects for growth. But Keynesians after Keynes did not read Keynes, or if they did they lacked the nerve to push his revolution to its full extent. And so they contented themselves with modeling aggregative variables such as income and the interest rate, unemployment and inflation. This allowed microeconomists to retain control over the theory of distribution, and the result was the post-war compromise division of introductory economics into a semester of micro and a semester of macro. These were ostensibly the same subject, but as any undergraduate could see they had essentially nothing in common.

After gathering strength for some decades, the microeconomists ventured forth to reclaim the entire subject—to oblige macroeconomics to accept a microeconomic foundation. They succeeded as the world knows. In most treatments nowadays, the same forces – supply and demand, relative supply and relative demand, aggregate supply and aggregate demand – determine all essential economic outcomes; it is only a matter of defining the appropriate framework for the market. Macroeconomics has practically been swallowed up. Even the term has been supplanted in some circles, and resold second-hand to a sub-discipline concerned with statistical analyses of econometric time-series.

In consequence we have today a wage doctrine that economists in the late 19th century would have found largely familiar. Relative wage rates are governed by marginal productivities, and these are determined by the relative supply and relative demand for labor of various grades. Supply and demand are largely governed by such factors as technology, trade, and within specific national labor markets by the various institutional arrangements that either impede or facilitate the adjustment of relative wage rates to the market-clearing equilibria. In a well-functioning economy, markets will clear, unemployment will be minimal, and the pattern of relative wages will reflect the efficient levels. In an economy beset by rigidities, unemployment will be visible, and relative wage rates will be distorted by the offending institutions. Given predominant views of “skill-bias” in technological change and the substitutability of foreign for domestic labor, it is expected that such distorted economies will show greater economic equality than the efficient ones. Having made a commitment to particular social structures, such economies face the trade-off between equity and efficiency in a particularly acute and intractable form.

Today’s well-bred economist has essentially no alternative to thinking along these lines. It is my purpose here to provide one. My argument is simple and empirical. The presence of *any* strong global pattern to changing inequalities in pay – whatever its source or precipitating causal factor – would refute the view that relative wages are determined wholly by the interactions of firms and workers in local or even national labor markets. It would show that they are, at the least, strongly conditioned by forces sweeping around the world. Such forces are, by definition, *macroeconomic*.

Suppose such a pattern exists. The question then becomes: what lies behind it? Once such a question is posed, the lines of causality are reversed, *ipso facto*. At that point, the study of macroeconomic change becomes an essential component of any effort to understand inequality outcomes. Depending on the size and relative importance of the pattern, compared to local variations, it may emerge as the *central* element in any such effort.

Fifty years ago, Simon Kuznets articulated the prototype of my view, holding that in general the processes of inter-sectoral transition would govern the evolution of pay inequality as economic development progressed. Thus, the famous hypothesis of the inverted “U”. Inequality would rise as the differential between city and countryside came to dominate the development landscape, and then decline as industrialization deepened and social democracy took hold. Kuznets’ hypothesis implies that most countries in the world eventually should surmount the agro-industrial transition

point, and be found on the downward-sloping portion of a Kuznets surface. Once that is the general case, then in general strong growth should reduce inequality, while recessions and economic crises, which reduce income levels, should increase it.

These are not the only relationships compatible with a general version of Kuznets' view. Kuznets surely understood that smaller countries, some specialized in commodities such as oil, would not follow the same inter-sectoral path as the United Kingdom or the United States. He would have appreciated that the emergence of sectors specializing in technology goods would conspire to lend a procyclical bias to inequality in the most advanced countries (such as the US, UK and Japan). Countries that sell into investment and export booms would tend to experience rising pay inequality as incomes rise and falling inequality when incomes fell. (In other work Pedro Conceição and I have called this the “augmented Kuznets hypothesis.” (Conceição & Galbraith, 2001)) Finally, Kuznets would surely have agreed that in a globalized economy, global forces are capable of imposing themselves over the purely national relationship between inequality and income.

The modern microeconomic position in these matters thus requires that the Kuznets hypothesis in its general form be rejected. There can be no consistent relationship, of any shape, between *levels* of income and *levels* of inequality, no global pattern to the movement of inequality. And this is the present mainstream position. The most widely subscribed theories hold that inequality is a matter of policy choice, idiosyncratic to each country and its political system.

The mainstream today links the choice of inequality level to the prospects for economic growth, but it divides into two rivulets on the direction of the effect. The first, associated with the authors of the 1994 *East Asian Miracle* report, holds that egalitarian policies (especially land reform and universal education) are preconditions for accelerated growth, on the ground that they improve work incentives and capacities (Birdsall, Ross, & Sabot, 1995). The second holds that increased *inequality* of income and wealth generates economies of concentration in leading sectors, and that this is a motor of accelerated growth (Bénabou, 1996; Galor & Tsiddon 1997a, 1997b). Both positions are supported, in the main, by theoretical hypothesis and casual empiricism, though an occasional article claims to find systematic support in the data for one or the other (e.g., Forbes, 2000). On all of this, more presently. The principal empirical support for the new mainstream view lies simply in the apparent absence of contradictory evidence, which would take the form of consistent support for the Kuznets hypothesis, in modern data. And for that the main modern source has been until recently the “high-quality” subset of the Deininger-Squire (DS) data set of world inequality measures, published by the World Bank (Deininger & Squire, 1996, 1998). But as a crucible for testing Kuznets, this source of information is defective in two critical respects.

First, DS is a very difficult data set from which to draw systematic conclusions. It combines measures of inequality of highly diverse types, including income and expenditure, household and personal income, income gross and net of tax (Atkinson & Brandolini, 2001). It combines these measures in an overall panel that remains sparse, with fewer than 800 country-year observations over nearly 50 years in the most widely used version. These measurements are unbalanced, with far more coming from advanced countries such as the United States, the UK, Japan and Taiwan, than from most of the countries of the developing world. Moving from this foundation to general statements about the evolution of inequality in the world economy requires steps to reconcile the differing data types, and steps to fill in the gaps, usually by interpolating over long intervals. Without additional sources of information, this is very difficult to do with any confidence in the result.

Second, even if DS were reliable, it is not a record of measures of inequality in *pay*. While many economists find measures of *household income inequality* attractive for welfare reasons,

Kuznets himself was concerned only about disparities in the rewards to labor; he considered other sources of income to be unnecessary complications. In this, Kuznets was entirely within the theoretical tradition of professional economics, according to which the distribution of labor income is governed by principles entirely separable from those governing the distribution of capital assets and political entitlements. Indeed economics has no theory governing the latter, which can take any form at all. But it is the distribution of pay, capital income and entitlements together that determine the overall distribution of *income*. Thus it is possible for a Kuznets relationship to exist in *pay* and yet be unobservable in *income*. If that is the case, the relationship nevertheless exists. The distribution of non-labor income (and the distribution of all income across household of varying sizes and compositions) may mask the Kuznets relationship, but it cannot render it invalid for the purposes Kuznets intended, of assessing whether inter-sectoral transitions and global macroeconomic forces are central determinants of changing inequality in the structure of remuneration.

Evaluating the Kuznets hypothesis broadly stated thus requires additional data. It would in principle require actual measurements of inequalities in pay, on a balanced, comparable and annual basis, for many if not most of the countries of the world. However, until recently no such data set has existed and indeed none was thought possible by those few analysts (if any) who considered the question.

The breakthrough – I will not minimize it – of the work of the University of Texas Inequality Project (UTIP, 2007) lies in seeing a way to construct precisely such a data set, and to do so from inexpensive and readily available raw materials. Our work is based on a sequence of assumptions, each of which has turned out, on close examination and checking against related evidence, to be remarkably robust.

We begin with the well-accepted argument of Theil (1972) that the between-groups component of a Theil inequality measure provides a lower-bound estimate of total inequality. We *infer*, first, that where a consistent group structure (sampling frame) exists through time, change of the between-groups component is usually a robust indicator of change in the dispersion of the entire distribution. We *discover*, second, that where a consistent group structure exists in measurements taken in different countries, the between-groups components are again reasonably robust measures of the relative degree of inequality in the different places. We *surmise*, finally, that observed changes in a consistently measured part of an overall dispersion is likely, though not certain, to be consonant with changes of the dispersion as a whole.

In practical terms, these inferences together imply that one can construct an entire worldwide data set for pay inequality from the between-industries component of a Theil statistic measured across manufacturing sectors for most economies. And the data permitting such a calculation are readily available from the UNIDO Industrial Statistics, a respected but generally neglected data source. First, data for individual countries are consistently available through time, generally on an annual basis from 1963 through 1999. Second, data are in a common classification scheme across countries, permitting between-groups components of the Theil statistic to be compared. Third, and most surprisingly, while the dispersion of manufacturing pay is more volatile than the dispersion of pay as a whole or the dispersion of incomes, the correlation across types of inequality measures is high, and one can therefore use the manufacturing pay dispersions as a robust indicator of the behavior of broader but often elusive economic distributions.

Galbraith and Kum (2003) present the UTIP-UNIDO data set of manufacturing pay dispersions, with about 3,200 country-year measurements over the period 1963 to 1999. Galbraith and Kum (2005) extend this work to show the correlation of these measures to a matched subset of about 500 observations from DS, and then show how pay inequality measures can be used as instruments,




Fig. 1. Inequality in China and Hong Kong.

alongside other variables, to fill out a balanced and dense data set of *estimated* measures of gross household income inequality. The result is called the Estimated Household Income Inequality data set, and along with UTIP-UNIDO it is freely available on the web-site of the University of Texas Inequality Project, at <http://utip.gov.utexas.edu>.

The emphasis here is on the patterns revealed by the UTIP-UNIDO measures, and their relationship to the Kuznets hypothesis. We proceed in three stages: patterns of change within individual countries, patterns of change across regions of the world through time; and global patterns.

2. Patterns within countries

Fig. 1 shows inequality in China and Hong Kong.¹ In both places, inequality remained nearly constant from the start of the reform period around 1979, through until the economic slowdown and inflation that led to the insurrection and repression of 1989, associated in the West with Tienanmen Square.² Following that period, our measurements show a steady increase in inequality in both China and Hong Kong.

Inequality declined slightly in the early reform period because the impetus to rapid growth came from agricultural reform, greatly increasing the real incomes of the peasantry. After Tienanmen, the Chinese government loosened central control and inaugurated the “open-door policy.” The loci of most rapid growth shifted to the coastal areas and the cities, and in particular to Guangdong province, in the South, and to Shanghai and to Beijing—the seat of government where the rebellion had most threatened the regime (Galbraith, Krytynskaia, & Wang, 2004). Here incomes grew dramatically in the years that followed. The spillover to Hong Kong can be accounted for partly as a question of migration, partly because Hong Kong handles most of Guangdong’s exports, and partly because the opportunities to become richer in certain sectors in China were then amplified in those same sectors (for instance, banking) in Hong Kong.

Fig. 2 shows the especially interesting cases of Iran and Iraq, neither of which are rich in survey data. Both countries experienced the oil boom of the early 1970s. But in Iraq this led to a marked decline in our measure of inequality, reflecting the transformation of the country (under

¹ The Chinese data in this diagram partly from UNIDO, and partly from information gleaned from the State Statistical Yearbook, which we believe to be consistent with the information provided by UNIDO.

² Most observers now believe that the major casualties of the June 4 incidents occurred in battles between the army and the city population of Beijing during the march into the square, and not in the Square itself, which most protesters had left by the time the Army arrived.




Fig. 2. Inequality in Iran and Iraq.

the Baathist regime, in the early years of Saddam Hussein's power) into an urban, middle class nation with the strongest health and welfare services in the Middle East at that time. In Iran, no similar transformation occurred. There then came the abrupt revolution of 1979, which deposed the Shah and installed the Islamic Republic. At that point, a radical equalization was imposed (with considerable violence) and it shows up with striking clarity in the figure. The two countries then went to war, maintaining popular mobilization and low inequality until the war ended in 1988. After that, both liberalized to some degree, and with striking, simultaneous effects. Inequality went up.

The Southern Cone of Latin America provides a third case; Fig. 3 presents the examples of Chile and Argentina, as well as Brazil. Notice that in Chile in 1973 and Argentina in 1976, turning points in measured inequality occurred. These were the years of military *coups d'état* against Salvador Allende and Isabel Peron, and the start, in both countries, of repression against leftists and the "dirty war." The experience of the two countries then diverge in 1982. This was a year of war (the Falklands) and the collapse of military government in Argentina, but of banking crisis in Chile, whose military remained in power for another 8 years. Chile began recovering from its extreme crisis in the mid 1980s, and inequality declined. Argentina entered a new period of crisis under the radical government of Raul Alfonsín, which deteriorated into hyper-inflation.




Fig. 3. Inequality in the Southern Cone.




Fig. 4. Inequality in Scandinavia.

By the time that Alfonsin resigned prematurely in 1989, cumulative increases in inequality in the two countries had again converged. Notice that throughout this period, Brazil too experienced rising inequality.

A striking counter-example to the general pattern of rising inequality occurs in the Scandinavian countries, where our measure of inequality remained largely stable in most cases, and actually fell substantially in the case of Denmark. Fig. 4 illustrates the pattern. No iron law dictated that inequality had to rise, even in the 1980s.

A few hundred miles away, in Central Europe, we find yet another pattern in the data. Measured pay inequality in the communist countries of Poland, Czechoslovakia and Hungary was stable through the 1970s and 1980s, reflecting the planned and command nature of the economic regime. When the regime collapsed, in 1989, inequality immediately sky-rocketed, as wages rose in some sectors but fell sharply for the less-skilled (many of whom also lost their jobs). Fig. 5 illustrates.

Before 1989, inequality in Central Europe was the lowest in the world—no doubt, so low as to interfere with economic efficiency and work incentives. After 1989, inequality rose sharply. But while this is true, these figures merely show the evolution of inequality within each country. We have not yet established that these same measures can be used to produce valid comparisons of inequality from one country to the next.




Fig. 5. Inequality in Central Europe.

3. Comparing inequality between countries

It is not obvious that one can use the between-groups component of Theil's T statistic to compare inequality meaningfully from one country to the next. Theil's T is not bounded between zero and one, as the Gini coefficient is³; measures taken from populations of different sizes are normally different on that ground alone. Moreover it is obvious that if one moves from a coarse to a fine classification scheme (for instance, from a two-digit to a four-digit industrial classification, or from a between-states to a between-counties measure with U.S. geographic data), a larger part of the total inequality will be observed as lying "between-groups" in the finer scheme. The between-groups component of Theil's T statistic will thus necessarily be larger in that case, even though population inequality is unchanged.

But in fact, the existence of a common classification scheme takes care of this problem. We chose to work with the UNIDO's Industrial Statistics initially because it is an inexpensive, reliable source of harmonized data covering many countries. We then found that the between-groups Theil's T statistics computed across industrial categories for the ISIC categories yielded measures that provide remarkably plausible cross-country comparisons. This came as a surprise. There is no compelling mathematical reason for it.

Yet we know the comparisons are plausible for three reasons. First, the measures are consistent across frontiers: countries that are economically integrated, with low barriers to trade and capital movement (as in Europe) tend to show similar inequality measures, even where those countries are markedly different in overall size (for instance, Denmark and Germany). Second, the measures correspond broadly to those of the Luxembourg Income Studies, a "gold standard" in measures of household income inequality (Atkinson, Rainwater, & Smeeding, 1995). The LIS achieves a high level of comparability across countries by careful comparison of large micro-data sets, but with only a handful of annual observations for each country and at a great cost of effort. Third, our measures correspond to casual expectation. They are low in socialist and social-democratic countries, higher in capitalist democracies, higher in middle-income developing countries, and generally speaking much higher in the dualistic economies of the Third World.

Table 1 gives partial evidence on UTIP-UNIDO rankings. The table presents three columns of countries for which at least ten observations are available, the lowest or most equal, those in the middle of the range, and those at the top. There is a clear consistency in the measurements. (Note that data for China in this particular version terminate in 1986.)

We infer that our measures of inequality across industrial pay, even though measured quite crudely across broad categories from a standard international data set, are in many ways reliable proxies for measures that others had spent far more time and money working to produce. So far as we know, this is merely a practical reality, a feature of economic life. It must reflect the fact that differences in industrial structure are highly influential in determining the larger differences in inequality characteristics of different countries. It must reflect the fact the ISIC categories, while crude, are sufficiently fine to capture large parts of the actual inequalities in the pay and also the income structure of most countries, along with the fact that the standardization of the group structure at 29 industrial categories holds the computed statistic within a definite range.

³ The bounds for Theil's T are zero and $\ln(n)$, where n is the size of the population, and the upper bound is approached when one person has all of the income. In principle the upper bound for the between groups component of Theil's T is no different, since one can imagine a two-group division of any "totally unequal" population into one group with population $1/n$ and all of the income, and the other with population $(n-1)/n$ and no income. But in practice an industrial group structure never approaches this bound.

Table 1
Low, middle and high inequality countries ranked by UTIP-UNIDO

Low inequality	Medium inequality	High inequality
China	Israel	Swaziland
Germany, East	Pakistan	Yemen
Cuba	Uruguay	Uganda
Czechoslovakia	Myanmar	Ghana
Denmark	Ecuador	Oman
Sweden	Somalia	Jamaica
Seychelles	Argentina	Cameroon
Romania	Venezuela	Congo
Macao	El Salvador	Trinidad and Tobago
Norway	Haiti	Mozambique
Netherlands	Zimbabwe	Kuwait

Table 2
UTIP-UNIDO inequality measures: distribution of observations by region and time

Continent	Before 1965	1966–1970	1971–1975	1976–1980	1981–1985	1986–1990	1991–1995	1996–1999
Africa	28	91	111	122	116	87	97	40
Asia	36	78	92	104	109	102	82	33
Europe	55	104	110	115	120	122	103	47
South America	11	21	27	35	41	46	43	17
Central and North America	24	48	62	58	67	55	49	20
Oceania	6	12	15	15	19	20	16	5

Table 2 gives the comparative coverage of the UTIP-UNIDO measures of inequality in manufacturing pay, for regions and 5-year intervals around the world. As the table makes clear, the coverage is consistent and reasonably uniform around the world – including many hard-to-survey countries of Africa, Asia and Central America – though it does taper off in the earlier years (before many countries had consistent reporting systems, or perhaps simply in some cases for failure to go back and fill in blanks in the historical record), and in the most recent period – no doubt due to lags in the reporting of recent data to UNIDO or in the preparation by the latter of the finished data set. Still, we have, for the first time, a snapshot of changing patterns of inequality within countries across the entire world, on a nearly annual basis over nearly 40 years.

Overall, the UNIDO source permits calculation of inequality measures for nearly 3200 country-year observations, covering over 150 countries during the period 1963 to 1999. We then match this data to real gross domestic product (GDP) per capita, from the Penn World Tables version 5.6. Including only countries with four or more observations on both variables, this matching reduces our data to 2836 country-year observations. The coverage of observations in region and time is tabulated in Table 2. Observations are annual for virtually all of the Americas, Europe, and Asia; only in Africa and for small island countries do we face significant gaps in coverage.

Fig. 6 gives a picture of the distribution of the UTIP-UNIDO inequality measures. Note that the raw values of the T statistic are grouped heavily to the left and diminish sharply as one moves to the right: there are many low observations and a handful of larger ones. Since the T statistic has a lower bound at zero, this distribution approximates what would be most likely if our measure were sampling from some random, larger universe of inequality numbers, generated by a process of random percentage change, over time, in the inequality statistic. This is reassuring. It suggests,




Fig. 6. Distribution of the UTIP-UNIDO inequality measures: raw and log form.

among other things, that there is nothing particularly strange or distorted about our collection of numbers. We find that the logarithm of our statistic is “nearly normal” in its distribution, which is what the underlying numerical properties of the index would lead one, under a simple rule and in the absence of any particular contradictory information, to expect. Fig. 6 gives both the distribution of measures and their log transformation.

Now let us look at the distribution of inequality measures over the entire world. The left panel of Fig. 7 presents a simple series of the average of our (log) UTIP-UNIDO pay inequality measures, annually for developed (OECD) and less developed (non-OECD) countries, together with bands indicating the standard error (variation) of the series. From this, we can see that in general, within country inequality measures are higher for developing countries, and that the gap in pay inequality between developed and developing countries remains nearly steady over four decades. The stability of the error bands tells us that our series is comparatively regular across countries from year to year; there may be fluctuations from country to country, but on the whole the range across countries remains remarkably constant.

Looking at the trend of these measures, we arrive directly at a central finding, which is that inequality in manufacturing pay in both OECD and non-OECD countries, on average, rose sharply from the early 1980s. The pattern is unambiguous. It is clear. It is consistent. It is uncomfortable. It is unmistakable. It is flatly contradictory to the idea, expressed by Sala-i-Martin and taken up by




Fig. 7. Trends in global inequality: contrasting views.

Stanley Fischer, that the years since 1975 have seen general declines in world inequality—outside the effect of rising average incomes in China and India on the global distribution. This is the first critical generalization of our work.

Why is this pattern clear to us and not to others? The right-hand panel provides an answer. When the same procedure is applied to the DS data, great fluctuations both within and between groups are found. In 1964, 1966 and 1982, but not in other years, non-OECD countries appear to enjoy less income inequality on average than OECD countries. And since the early 1980s, while the poorer countries appear to have experienced increased income inequality, the rich countries appear to have not. This is despite the fact that pay inequalities increased in both groups of countries, despite the fact that the industrial sectors are generally much larger in the OECD countries, despite the fact that from the early 1980s onward the developed countries were hit by policies, under Reagan in the United States and Thatcher in Great Britain, that explicitly aimed to reduce labor income and raise the share and the power of the rich.

The DS results can occur in a data set, but not in real life. As it turns out, they are due mainly to large changes in the composition of the data set from 1 year to the next. In a low inequality-year DS may be reporting on a handful of East European countries, while in a high-inequality year the sample might be weighted to Africa or Latin America. So the comparison is, in a sense, unfair. It does not necessarily mean that any particular number in the DS collection is wrong. But it is by no means obvious how one can effectively overcome the gaps in coverage, in order to make the DS data into a meaningful measure of world developments. As Milanovic (2002) showed, the one attempt to do so (Sala-i-Martin, 2001) was shot through with approximations and interpolations. Sala-i-Martin's method, in short, was designed so that it could not, in general, have picked up the influence of the year-to-year trends apparent in the left-hand-panel. But are those trends in fact significant for income inequality? We think so. But that fact, let us remind ourselves, has yet to be established.

4. Regional and global trends

The next step is to examine patterns of *change* in the global pay inequality data over a series of 6-year windows. To permit as complete a picture as possible in a world of incomplete data, we use an average value for the 3 years immediately surrounding the start and the finishing dates of each window. This permits us to capture the major patterns of change in the data and to see, at a glance, where and when rising inequality began, and how it moved around the world. This approach will also permit us to identify significant exceptions to the trend.

Fig. 8 presents the pattern as we have it for the 1960s, specifically from 1963 to 1969. The map is spotty, owing to the relatively fragmentary character of the records for the earliest years of the data set. In any event, the important characteristic of the map is that there is no clearly discernible pattern. Our measure of inequality rises in some countries and falls in others, and that is about as much as one can reasonably and reliably say.

In the 1970s the plot thickens, the data grow more comprehensive, and a distinct pattern emerges. This is of rising inequality in the richest industrial countries, and indeed across the world of countries that rely mainly on imported oil. Meanwhile among oil exporters, especially across North Africa and the Middle East, the case is different; our measure of inequality in these regions undergoes a systematic decline (Fig. 9 illustrates).

Is it too much to infer the obvious? The oil shocks of the early 1970s created boom conditions in the producing countries, recession among the major consumers. Where economies boomed, pay inequalities declined: low wage workers gained more rapidly, in proportionate terms, than




Fig. 8. Changing inequality in the UTIP-UNIDO data, 1963–1969.




Fig. 9. Changing inequality in the UTIP-UNIDO data, 1970–1976.




Fig. 10. Changing inequality in the UTIP-UNIDO data. 1981–1987.

those at the top of the scale. This may be simply a demand effect. With rapid growth, the hours worked of the most contingent members of the workforce tend to grow, while those who already had secure employment are less favorably affected. With recession, conversely, those at the bottom of the pay scale are more vulnerable to reduced hours and intermittent layoffs than those at the top, and inequalities widen. In any event, for our purposes the precise mechanism is not important. As Fig. 10 shows, there is a pattern in the data by this point which is very difficult to assign to any other general cause, than the world-wide shift in relative purchasing power that occurred in the favor of exporters of oil and other commodities during the inflationary 1970s. Meanwhile in certain developing countries (notably, Brazil and Argentina), rapidly expanding bank debts kept up a momentum of economic growth despite the higher bill for imported oil.

The commodities boom of the 1970s gave way, of course, to the debt crisis and economic collapse of the early 1980s, and this pattern is clearly reflected in Fig. 11. Inequality rises most sharply in Latin America, notably Brazil, Chile, Peru and Mexico: the epicenters of the debt crisis. Though our information is less good, inequality also appears to be rising sharply in Africa at this time, and in parts of Asia. But there are notable exceptions. Iran and Iraq, in the grip of revolution and war, defy the trend. So do India and China. The common distinguishing feature of all four countries is their prior insulation from the developing world's rush to dependence on borrowing from commercial banks, and their consequent protection from the financial upheavals that followed. China in particular had remained entirely autarkic through 1979, while India's external debt was concentrated on long-term, concessional loans from the World Bank's International Development Agency. Both countries maintained strict capital control through this period which had prevented debt entanglement by entities outside the government. For the rest of the world, we may term this period the Age of Debt.




Fig. 11. Changing inequality in the UTIP-UNIDO data, 1988–1994.

Finally, in the late 1980s, a wave of reform swept the socialist world, leading ultimately to the collapse of communist governments in Europe and then of the Soviet Union itself. In this period, rising inequality is almost everywhere. But it is greatest in those countries which had previously registered the most exaggerated degree of equality in their pay structures: the former Soviet Union and its neighbors. Fig. 12 illustrates the pattern observed through 1994.

While a few countries show declining inequality in this period (some, like Mozambique, with data that are surely open to question) only one region shows a pattern of increasing equalization consistently across countries at this time. That was Southeast Asia – notably Malaysia, Thailand, Indonesia, and Singapore – then in the grip of a boom driven by foreign direct investment. This boom came to an end in the crisis of 1997. Consistent worldwide data are not yet available for the years since the crisis, so we are as yet unable to measure the effects of those events on global inequality. Nevertheless, what we have so far is – by far – the most complete and consistent global pattern of changing inequalities within countries yet measured.

The general patterns of change observed in this data point to the predominant influence of regional and of global economic forces on conditions within countries. The effect of these forces depends, of course, on the circumstances into which they are projected. High oil prices lower inequality in producing countries and raise them among the consumers. High interest rates – the proximate cause of the debt crisis in 1981–1982 – raise inequality in debtor countries but leave those who have sheltered themselves from the whirlwind of global finance comparatively unscathed. Globalization – falling trade barriers and the freeing of the capital account – increases inequality as the barriers that once shielded an industrial middle class in developing and socialist countries come down. Against this, the force of FDI can keep inequality at bay for a time—though only for a time, and only in countries small enough, and tied closely enough to rapidly expanding export markets for FDI to make a large effect on national conditions.




Fig. 12. A downward-sloping Kuznets surface in pay data.

These patterns lead to another important question. What is the relationship between inequality and the *level* of national income? Is there a systematic relationship between the level of one and the level of the other, as Simon Kuznets argued five decades ago? Or, alternatively as modern theorists prefer to believe, is the correct relationship between the *level* of inequality and the *rate* of economic growth?

5. Reexamining the Kuznets hypothesis

Kuznets' famous argument was based on straightforward consideration of the central processes of development, and in particular of the *inter-sectoral transition* from agriculture to industry. In the northern states of the United States, notably, the starting point had been a relatively egalitarian society of agrarian freeholders; these were prosperous by then-prevailing world standards but they still had much lower real incomes, on average, than could be had in the emerging industrial cities. Thus as the industrial revolution took hold, capitalism opened a vast urban-rural income gap, and overall inequality in the society rose primarily on this account. Thus Kuznets argued that the initial transition toward industrial development would be accompanied, in general, by rising inequality in the distribution of income.

Later on – especially with the mechanization of agriculture – the center of economic gravity would shift entirely to the cities. And with a declining relative rural population the entire evolution of inequality would come to be determined by relationships inside the cities (and later, the suburbs). In these relationships, countervailing power, modern industrial relations, democratization and the rise of the welfare state would assure, past a certain point, declining inequality in the overall structures of pay. From this, Kuznets inferred that the relationship between income and inequality

would follow an inverted “U” shape: first rising, and then falling, as the ordinary processes of industrialization unfolded.

Whether Kuznets’ view of the initial conditions of agrarian life were a realistic portrayal of the general case is very much open to doubt. They surely did not characterize either feudal Europe or the American South. There is also no special reason to assume that the income level at which the transition from agriculture to industry begins would necessarily correspond, over long historical intervals, from one country to another. Thus the inverted “U” may or may not be found in any particular set of modern data. For that matter, it may or may not be found in the historical data, if we were to obtain superior measures of the relative incomes prevailing between farm and cities, and within cities themselves, in the long process of industrialization.

But this is not the true test of Kuznets’ procedure. What Kuznets offered was a *general* method for coming to some state of expectation concerning the pattern of inequalities that one might reasonably expect. That method consists of assessing inequality primarily as a matter of an appropriate pattern of inter-sectoral transitions. The key to the analysis lies in identifying the principal sectoral structures relevant to the problem, and the key characteristics of each. Kuznets’ particular case was of the transition from egalitarian agriculture to a mixed industrial economy, followed by the decline of the rural population share and the democratization of industry. But other patterns can easily exist, and these too will yield “Kuznets relations” of varying forms. If Kuznets is right in general terms, such patterns – if correctly identified – are the key to understanding why inequality sometimes rises and sometimes declines. They will, as a general rule, involve some consistent pattern of evolution of inequality as income levels change.

Fifty years after Kuznets, the development literature continues to make reference to his inverted “U” hypothesis, but his larger founding of the causal force in inter-sectoral transitions has been for the most part forgotten. Rather, the test of the inverted “U” curve has become a statistical exercise, largely free of historical and institutional context. And in the meanwhile, the available evidence – mostly, Deininger and Squire – has been deployed to argue that no Kuznets relationship, indeed no systematic relationship between income levels and inequality of any kind, exists. Various studies have, to be sure, found curves of one form or another in various subsets of the DS data. But the evidence is weak, and no general pattern has been widely accepted.

In 1994, important work associated with the World Bank and published as a report entitled *The East Asian Miracle* advanced a very different claim. EAM argued that reductions in inequality were *preconditions* to the growth surge experienced especially in Korea, China, and Taiwan. The underlying theoretical relationship tied the expected return on education, economic engagement (and entrepreneurial activity) to the distribution of current returns. In a more equal society (it was said) people feel a closer connection between work and reward and are therefore more strongly motivated to effort. Thus, policies that reduced inequality – land reform and universal public education, in particular – would lead toward more rapid growth.

The EAM argued for redistribution within a market-oriented policy framework emphasizing personal incentives and micro-structural causal factors. In this, the EAM study embraced neo-classical economics without entirely going over to neoliberal policy ideas. The essence of the argument, however, was to neglect institutional forces and inter-sectoral transitions in favor of an argument about personal behavior in the labor market. Unfortunately, once the venue is conceded, the decision in the case may be determined, and in ways that the plaintiff’s attorney may not be able to control. So it happened here.

In the late 1990s, a contrary position emerged. This one held that *rising* inequality could set the preconditions for more rapid growth. This argument was essentially similar to the view of Victorian political economists, who believed in the concentration of wealth as the motor of capital

accumulation, and therefore the indispensable social role of the upper class. In the new version, the relevant theory of growth was based on the idea that concentrations of knowledge, savings, capital and entrepreneurship provide the seed-bed for the transition to advanced development. In an influential article in the *American Economic Review* in 2000, Kristen Forbes claimed to have found evidence for this position in (where else?) the high-quality subset of the DS data.

The EAM and the Forbes positions were, of course, diametrically opposed. One held that the relationship between initial inequality and later growth slopes downward; the other held that the relationship slopes upward. Yet in their view of the character of the causal relation between inequality and economic growth, they shared an important element of common ground. Both held that there exists an important causal link between the *level* of inequality and a subsequent *rate* of economic growth, which is to say, between an initial *state* and a later *rate of change*. Whichever way this connection ran, it is plainly incompatible with Kuznets' idea, which connects the level of inequality to the level of income, and the evolution of inequality to inter-sectoral changes in the pattern or stage of economic development.

To see this point, consider two countries perched side-by-side in income-inequality space. Let us say, for instance, they are both in the midst of a rural-to-urban transition, and both are therefore experiencing high inequality at the peak of the inverted "U". Under the original Kuznets formulation, therefore, both would expect to enjoy declining inequality as development progresses and incomes increase. Both would therefore descend the downward-sloping surface of the Kuznets inverted "U", and divergences in future inequality would depend largely on divergences in future income performance.

Under either the EAM or the Forbes model, however, the Kuznets relationship – even if it existed at some one point in time – would quickly evaporate. Let us suppose the EAM hypothesis were correct. Let us further suppose one of the two countries were to embark on a program of land reform and public education. At that point, income inequality would fall. This would happen first, *before* there was any growth in income. The socially progressive country would therefore fall off the Kuznets curve. Only with time, as growth accelerated and income levels grew, would the inverted "U" relationship be even partly restored (and that assumes that the growth of incomes would not, itself, reverse the earlier decline in inequality). At any moment of time, moreover, *any number* of countries might be in the off-curve position, that of embarking (or renewing) egalitarian commitments as a precondition to more rapid growth. There is therefore no reason to believe that a stable, general downward-sloping relationship between inequality and income levels would persist in the data.

Under the Forbes hypothesis, contrariwise, a country that enacted neoliberal reform would first experience rising inequality, and then rising incomes. The further evolution of that country would be entirely off the Kuznets curve. And since, again, at any moment any number of countries might be in the process of implementing reform, there is no reason to expect an enduring relationship between income levels and inequality measures. As general matter, therefore, truth of either proposition relating levels of inequality and later rates of growth precludes finding a Kuznets-type relationship between levels of inequality and levels of income. Put another way, the finding of a Kuznets relationship of any kind – whether sloping downward or upward or curved in any predictable respect – between the levels of income and the levels of inequality would constitute strong evidence against either the East Asian Miracle or the Forbes propositions.

But when we correlate levels of manufacturing pay inequality and levels of national per capita income, what do we find? We find, in fact, that a stable and persistent relationship exists!

There is a strong negative correlation in the UTIP-UNIDO data: higher incomes are associated with lower inequality, generally speaking. This is by itself decisive evidence against either thesis

linking an initial state of inequality to a later rate of growth—though it is, to be sure, somewhat more decisive against the Forbes version than against that of the EAM report.

The finding of a negative relationship between income and inequality in a dense global data set based on manufacturing pay is broadly consistent with Kuznets general method and also loosely consistent with his specific historical thesis. It is important to remember that the UTIP-UNIDO data measure only inequalities in the structure of manufacturing pay. And the measures are for a specific historical period – 1963–1999 – which was associated with deepening industrialization in much of the world. We therefore should not be expected to observe many countries in the initial phases of industrialization, undergoing a large rural–urban transition (here, China is a major exception). Instead, most of the information we observe is of a type that Kuznets would have expected to display a downward-sloping relationship between income and inequality, and most of it does.

An important qualification concerns the behavior of a relatively small number of advanced industrial countries – the US, the UK, and Japan – and a somewhat larger number of small oil producers, mainly in the Persian Gulf. In these cases, rising incomes have led in recent years to rising inequality. The effect is to give our version of the Kuznets curve a slight upward loop on the right-hand side of the scale (Conceição & Galbraith, 2001). We believe that in the case of the most advanced industrial countries, a principal reason for the upward-slope lies in the fact that the most advanced sectors supply capital goods and therefore have strongly pro-cyclical patterns of employment and pay. Thus when growth accelerates, pay differentials widen in such countries.

We find that a broad relationship between the level of income and the level inequality – when measured, as Kuznets intended, in structures of pay – exists, contrary to the implications of recent theoretical models. But on closer examination, something else important emerges from this data. And that is, that the Kuznets relationship is not stable over time. Rather, taken over the entire period of observation, it has tended to shift outward, so that more recent years are characterized by higher levels of inequality, for a given level of real per capita national income.

Fig. 13 illustrates this relationship with the visual device of a regression plane. Inequality is on the vertical axis, per capita GDP on the left horizontal, and a “time” variable on




Fig. 13. The global macroeconomic movement of pay inequality.




Fig. 14. Actual change in pay inequality and simulation with global component removed.

the right horizontal. The downward slope of the plane as the figure recedes into the distance captures the core relationship between inequality and income. But it is well to notice that the plane is not horizontal in the third dimension. Instead, it is tilted upward, indicating that inequality has tended to rise, year on year, even when changes in the income level are controlled for. This is something new. It represents a common element, or *global tendency*, in the data.

It is possible, moreover, to gain a considerably more precise view of the character of this global tendency, by estimating a regression model with inequality as the dependent variable, per capita income as the dependent variable, and vectors of dummy variables or categorical effects, one for each country and one for each year. This is a two-way fixed effects model. The resulting vector of time effects shows the pattern of the common tendency, relative to the final year in the sample. Fig. 14 illustrates, and the arrows provide a schematic of the major movements. We find that, over all, there was no global trend in inequality during the 1960s. From 1973 through 1980, on the other hand, inequality declined: the years of the oil shocks, the commodities boom, and the global buildup of commercial debt sponsored strong economic growth and relatively rapid growth of pay for less-well-paid workers.

Then came the U-turn of the 1980s. From 1981 through the end of the millennium – the years of debt crisis, communist collapse, and neoliberal globalization – inequality within countries rose relentlessly *as a global pattern*. And what is more, this pattern is essentially identical to that observed when one measures the dispersion of incomes *between* countries. The inset gives this pattern, as calculated by Milanovic (2007) from the Penn World Tables, using only average GDP per capita.

What is the interpretation of this remarkable matching of trends? The simplest one, and by all odds the most appealing, is *that they measure the same events*. The UTIP-UNIDO measures capture a global pattern in the evolution of inequalities within countries, and in particular a sharp turnaround and increase from 1981 forward. These data measure the gap, across industries, between the well- and the poorly paid. The Milanovic/Penn World Tables measure captures the gap between rich and poor across countries. But these developments could, and probably do, have a common cause. Surely a force that widens the income difference between the rich and the poor within any given country should also widen the income difference between a country comprised mainly of rich people and a country comprised mainly of poor people. The fact that this appears to have been the case is merely a confirmation that UTIP-UNIDO and the Penn World Tables are reliable sources of information.

How important is the global component of inequality change? Fig. 14 provides the answer. The left panel, as before, reflects the data as they are. The right-hand panel is a simulation of what the UTIP-UNIDO inequality measures would show if the global component of rising inequality were removed. The evidence is clear: the whole of the worldwide rise in inequality in the years of globalization and neoliberalism is attributable to the common worldwide pattern. It was, in other words, a global macroeconomic event.

The fact that these patterns are common across countries establishes an important forensic fact. It cannot be the case, as most writing on the subject and in particular the “new theories” relating inequality to subsequent growth persistently allege that inequality is merely the reflection of policy choices taken within countries. If that were true, we should expect no global pattern: countries have differing preferences and should be expected to act on reformist impulses at differing times. This is not, however, what we observe. To the contrary, we find that there is a consistent global pattern to the movement of inequality, and this pattern transcends domestic political choice. Presumably, few countries would deliberately choose to lower their growth rate or to slip further behind in average per capita income, but this occurred, as the divergence in the Milanovic measure reveals.

There must be a common external cause. And this brings us to the critical question: what did cause the decline in global inequality, both within countries and between them, from 1973 to 1980, and the long, sordid increase in the two decades that followed?

As it happened, changes in global financial governance occurred on precisely the dates now revealed to be the critical turning points in the inequality data. Up until 1973, the world lived under the globally stabilizing financial order of the Bretton Woods system, created in 1944 to provide a framework within which countries could pursue reconstruction and development, with short-term financial assistance from the International Monetary Fund and long-term development aid from the World Bank. Capital movements between countries were generally controlled, and international commercial banking played a minor role in global finance. This system began to break down when Richard Nixon ceased to exchange dollars for gold in central bank settlements in 1971, and it fell apart altogether in 1973.

From 1973 through the end of the decade, commercial banks financed global development on exceptionally favorable terms. In a time of rising commodity prices (notably oil, but also cotton, coffee, copper and other mainstays of third world agriculture and mining), real interest rates facing the developing world were effectively less than zero. It was a great moment to borrow and spend, and the developing countries did so with abandon. As a result they grew rapidly—and, moving down the inverted “U” of the Kuznets curve, inequalities within those countries generally fell. But it was also to prove the last hurrah.

From 1979 forward, with the arrival of Margaret Thatcher (and economic monetarism) in power in Britain, and then the appointment of Paul A. Volcker to chair the Federal Reserve Board in the United States, and then again the election of Ronald Reagan in 1980, the global economic climate changed. Interest rates soared, the pound and the dollar rose, commodities slumped, and the currencies of Latin America and Africa collapsed. Imports were slashed, and what was experienced as recession in the global North was full-throated depression in the global South. Now, with declining per capita incomes, the Kuznets process went into reverse, and inequality rose sharply. But it was not only a Kuznets process. Everywhere, with high interest rates, creditors gained on debtors. Everywhere, with a rising dollar, those who held dollar assets gained on those who did not. Everywhere, with collapsing local currencies, those who could sell to the outside world gained on those whose market was largely internal. Everywhere – and continually for 20 years, with just the limited exceptions already mentioned – global inequalities rose.

Acknowledgements

The author holds the Lloyd M. Bentsen, jr. Chair in Government/Business Relations at the Lyndon B. Johnson School of Public Affairs, The University of Texas at Austin, where he directs the University of Texas Inequality Project. He is a Senior Scholar with the Levy Economics Institute and Non-Resident Fellow, Center for Global Development. This article is based partly on joint work theoretical work with Pedro Conceição and joint empirical work with Hyunsub Kum, and on many conversations with the UTIP team.

References

- Atkinson, A., & Brandolini, A. (2001). Promise and pitfalls in the use of secondary data-sets: Income inequality in OECD countries as a case study. *Journal of Economic Literature*, 34, 771–799.
- Atkinson, A., Rainwater, L., & Smeeding, T. (1995). *Income distribution in OECD countries: Evidence from the Luxembourg income study*. Paris: OECD.
- Bénabou, R. (1996). Inequality and growth. In B. S. Bernanke & J. J. Rotemberg (Eds.), *NBER macroeconomics annual 1996* (pp. 11–74). Cambridge, MA: MIT Press.
- Birdsall, N., Ross, D. R., & Sabot, R. (1995). Inequality and growth reconsidered: Lessons from East Asia. *World Bank Economic Review*, 9(3), 477–508.
- Conceição, P., & Galbraith, J. K. (2001). Towards a new Kuznets hypothesis: Theory and evidence on growth and inequality. In J. K. Galbraith & M. Berner (Eds.), *Inequality and industrial change: A global view* (pp. 139–160). Cambridge: Cambridge University Press.
- Deininger, K., & Squire, L. (1996). A new data set measuring income inequality. *World Bank Economic Review*, 10, 565–591. Available at <http://www.worldbank.org/research/growth/dddeisqu.htm>
- Deininger, K., & Squire, L. (1998). New ways of looking at old issues: Inequality and growth. *Journal of Development Economics*, 57, 259–287.
- Forbes, K. (2000). A reassessment of the relationship between inequality and growth. *American Economic Review*, 90, 869–887.
- Galbraith, J. K., & Kum, H. (2003). Inequality and Economic Growth: A Global View Based on Measures of Pay. *CESifo Economic Studies*, 49(4), 527–556.
- Galbraith, J. K., & Kum, H. (2005, March). *Estimating the Inequality of Household Incomes: Toward a Dense and Consistent Global Data Set*, *Review of Income and Wealth*, 51(1), 115–143.
- Galbraith, J., Krytynskaia, L., & Wang, Q. (2004). The Experience of Rising Inequality in Russia and China during the Transition. *European Journal of Comparative Economics*, 1(1), 87–106.
- Galor, O., & Tsiddon, D. (1997a). The distribution of human capital and economic growth. *Journal of Economic Growth*, 2, 93–124.
- Galor, O., & Tsiddon, D. (1997b). Technological progress, mobility and economic growth. *American Economic Review*, 87(3), 363–382.
- Milanovic, B. (2002, September). The Ricardian Vice: Why Sala-i-Martin's calculation of world income inequality cannot be right, mimeo.
- Milanovic, B. (2007). *Worlds Apart: Measuring International and Global Inequality*. Princeton: Princeton University Press.
- Sala-i-Martin, X. (2001, March). The disturbing 'rise' in global income inequality, mimeo.
- Theil, H. (1972). *Statistical decomposition analysis: With application to the social and administrative science*. Amsterdam and London: North Holland.
- UTIP (University of Texas Inequality Project). (2007), <http://utip.gov.utexas.edu>.