NBER WORKING PAPER SERIES

EXPLAINING INEQUALITY THE WORLD
ROUND: COHORT SIZE, KUZNETS
CURVES, AND OPENNESS

Matthew Higgins
Jeffrey G. Williamson

Working Paper 7224
http://www.nber.org/papers/w7224

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1999

is given to the source.

Explaining Inequality the World Round:
Cohort Size, Kuznets Curves, and Openness
Matthew Higgins and Jeffrey G. Williamson

ABSTRACT

Klaus Deininger and Lyn Squire have recently produced an inequality data base for a panel of countries from the 1960s to the 1990s. We use these data to decompose the sources of inequality into three central parts: the demographic or cohort size effect; the so-called Kuznets Curve or demand effects; and the commitment to globalization or policy effects. We also control for education supply, the so-called natural resource curse and other variables suggested by the literature. While the Kuznets Curve comes out of hiding when the inequality relationship is conditioned by the other two, cohort size seems to be the most important force at work. We resolve the apparent conflict between this macro finding on cohort size and the contrary implications of recent research based on micro data.

Matthew Higgins
International Economics Function
Federal Reserve Bank of New York
Federal Reserve P.O. Station
New York, New York 10045
matthew.higgins@ny.frb.org

Jeffrey G. Williamson
Department of Economics
216 Littauer Center
Harvard University
Cambridge, MA 02138
and NBER
jwilliam@kuznets.fas.harvard.edu

1. Introduction

The empirical results presented in this paper provide strong support for cohort size effects on inequality the world round: large mature working age cohorts are associated with lower aggregate inequality, and large young adult cohorts are associated with higher aggregate inequality. This finding is consistent with the writings of Richard Easterlin and others regarding the fallout from America's previous baby boom. It is also of interest because standard theoretical models associated with Angus Deaton and others argue the opposite. In addition, the paper reports compelling evidence that inequality follows the inverted-U pattern described by Simon Kuznets, tending to rise as a country passes through the early stages of development, and tending to fall as a country passes through the later stages. This is a littered academic battlefield, but our work differs from most previous studies of the Kuznets hypothesis by examining the inequality-development relationship conditional on other variables. In particular, and as we have noted, the analysis stresses a country's position in the demographic transition, as measured by the mature adult share of the labor force, and on a country's degree of economic openness. However, and consistent with so much of recent inequality debate about rising wage inequality in the US and in other OECD economies in the 1980s, we find only limited support for the hypothesis that a policy commitment to globalization has an impact on inequality.

Section 2 surveys the three main hypotheses upon which this papers dwells: cohort size, Kuznets Curves and openness. Section 3 describes patterns in inequality, openness and cohort size, across regions and since the 1950s. Section 4 presents pooled and fixed-effects estimates of the relationship between inequality and cohort size, Kuznets Curve effects, openness and other variables. It also explores the quantitative significance of the estimated effects. Section 5 conducts simulation exercises to evaluate

1

potential sources of the negative link between cohort size and inequality. Section 6 concludes.

2. Reviewing the Three Hypotheses

2.1 Inequality and Cohort Size

The cohort size hypothesis is simple enough: fat cohorts tend to get low rewards. When those fat cohorts lie in the middle of the age-earnings curve where life-cycle income is highest, this labor market glut lowers income in the middle, thus tending to flatten the age-earnings curve. Earnings inequality is moderated. When instead the fat cohorts are young or old adults, this kind of labor market glut lowers incomes at the two tails of the age-earnings curve thus tending to heighten the slope of the upside and the downside of the age-earnings curve. Earnings inequality is augmented. This demographic hypothesis has a long tradition in the United States starting with the entry of the baby boomers into the labor market when they faced such poor prospects (Easterlin 1980; Freeman 1979; Welch 1979), and it was surveyed recently by David Lam (1997: pp. 1023-4 and 1044-52). Kevin Murphy and Finis Welch (1992) and Murphy and Lawrence Katz (1992) have now extended this work to include the 1980s. All of these studies have shown that relative cohort size has had an adverse supply effect on the relative wages of the fat cohort in the United States since the 1950s. This tradition ignores the potential endogeneity of hours and weeks worked, educational attainment and labor-force participation rates with respect to cohort size. We shall do the same in this paper, but it should be noted that one effort to endogenize those effects for the United States has concluded that:

> "almost all of the change in the experience premium over the past 30 years (younger and older relative to prime-age workers) and a significant portion of the change in the college wage premium can be explained solely as a function of changing age structure." (Macunovich 1998: 263)

If the cohort size hypothesis helps explain United States post-war experience with wage inequality, it might do even better world wide. After all, there is far greater variance in the age distribution of populations between regions and countries than there has been over time in the United States. Furthermore, the demographic transition in the Third World has generated much more dramatic changes in relative cohort size than did the baby boom in the OECD. The higher demographic variance between countries at any point in time versus within countries over time can also be illustrated by a pair of summary statistics from the data set used in this paper. Define the variable MATURE as the proportion of the adult population 15-69 who are ages 40-59. When the standard deviation of MATURE is calculated between countries in the sample we get a figure, 5.10, that far exceeds the standard deviation over time within countries for the sample, 1.66. Thus, the variance in cohort size across countries and regions is more than nine times the variance for countries over time.

All of this suggests that cohort size is likely to matter in explaining inequality the world around since the 1950s, fat young-adult cohorts creating inequality while fat prime-age cohorts doing just the opposite. Interestingly, a recent and influential paper by Angus Deaton and Christina Paxson (1997) suggests that faster population growth (and, thus, fat young and thin prime-age cohorts) should *reduce* inequality. The resolution of the apparent conflict is, we think, straightforward, but is reserved for section 5.

Two caveats are in order before we proceed. First, we have relied on the micro cohort size literature to motivate the discussion of demographic effects on inequality. This literature assumes that cohort size effects reflect the competitive market-clearing equilibrium, driven by imperfect substitutability in production between workers of different experience levels. We are unable to test this assumption, but the validity of our empirical results does not rest on it. It is also possible, for example, that more mature workers are better at "gaming" the economic system, and thus, in extracting rents from other age groups.

4

The fact of cohort-size-effects-on-income requires only that the total income accruing to a cohort rises less-than-proportionately with cohort size, whatever the causal mechanism. Second, as a related matter, the micro cohort size literature focuses on *earnings*; the international macro inequality data pertains to total *income*, and sometimes consumption. We know of no way to address this mismatch without abandoning the attempt to link international demographic variation with international variation in inequality. Given the much greater demographic variation in the international data, we hold that this would be throwing out the baby with the bathwater. In effect, we assume that what holds true for earnings holds true for income as well. The true links between demography and income inequality are no doubt more complex, depending on the links among demography, savings rates, the transmission of wealth across generations, and the mean and variability of returns to accumulated assets.

2.2 Inequality and Openness[1]

After 1973 and especially in the 1980's, the US experienced a dismal real wage performance for the less skilled, mostly due to declining productivity growth coupled with increasing wage inequality between skills. The ratio of weekly wages of the top decile to the bottom decile increased from 2.9 in 1963 to 4.4 in 1989 (Kosters 1994; Freeman 1996). This inequality was manifested primarily by an increasing wage premiums for workers with advanced schooling and age-related skills. While the same inequality trends were apparent elsewhere in the OECD in the 1980s, the increase was typically far smaller (Kosters 1994). The debate over explanations for these inequality facts started with the observation that rising inequality coincided with rising globalization in the form of trade and immigration. The latter underwent increasing quantity and declining quality (Borjas 1994). Trade shares in the US increased from 12 to 25% of GNP between 1970 and 1990 (Lawrence and Slaughter 1993), while World

---

[1] Much of this section is taken from Williamson 1997, pp. 119-121.

Bank figures document that the share of output exported from low-income countries rose from 8 to 18%

between 1965 and 1990 (Richardson 1995, p. 34). These inequality developments also coincided with a

shift in US spending patterns which resulted in large trade deficits. Thus, economists quite naturally

explored the linkages between trade and immigration, on the one hand, and wage inequality, on the other.

The standard Heckscher-Ohlin two-factor, two-good trade model makes an unambiguous

prediction. Every country exports those products which use intensively abundant and cheap factors of

production. Thus, a trade boom induced by either declining tariffs or transport costs will cause exports and

the demand for the cheap factor to boom too. Globalization in poor countries should favor unskilled labor

and dis-favor skilled labor; globalization in rich countries should favor skilled labor and dis-favor unskilled

labor. Robert Lawrence and Matthew Slaughter (1993) used the standard Heckscher-Ohlin trade model

to explore wage inequality and concluded that there was little evidence to support it. Instead, the authors

concluded that technological change was the more important source of rising wage inequality. Hot debate

ensued.

This strand of the debate stressed the evolution of labor demand by skill, ignoring the potential

influence of supply. George Borjas (1994) and his collaborators (Borjas, Freeman and Katz 1992) took a

different approach, emphasizing instead how trade and immigration served to augment US labor supply. In

order to do this, they first estimated the implicit labor supply embodied in trade flows. Imports embody

labor thus serving to augment effective domestic labor supply. Likewise, exports imply a decrease in the

effective domestic labor supply. In this way, the huge US trade deficit of the 1980s implied a 1.5%

increase in the US labor supply and, since most of the imports were in goods which used unskilled labor

relatively intensively, it also implied an increasing ratio of unskilled to skilled effective labor supplies. In

addition, there was a shift in national origin of immigrants from the 1960s to the 1980s so that an

increasing proportion of immigrants were from the less developed nations (e.g., Mexico and Asia) and

thus more unskilled, which in turn meant a far higher fraction of immigrants were relatively unskilled just when there were more of them.

These relative supply shifts gave economists the desired qualitative result—wage inequality between skill types. The quantitative result, at least in Borjas's hands, also seemed big. Borjas estimated that 15 to 25% of the relative wage decline of high school to college graduates was due to trade and immigration. He also estimated that 30 to 50% of the decline in relative wage of high school dropouts to all other workers was due these same globalization forces, one-third of which was due to trade and two-thirds to immigration.

Thus far, the discussion has focused mainly on the United States, perhaps because this is where rising inequality and immigration have been greatest. But the question is not simply why the United States and even Europe experienced a depressed relative demand for low-skilled labor in the 1980s and 1990s (Freeman 1995, p. 19), but whether the same factors were <u>stimulating</u> the relative demand for low-skill labor in the poor Third World. This is where Adrian Wood (1994, Chp. 6; 1995) entered the debate. Basing his results on insights derived from classical Heckscher-Ohlin theory (extended by Stolper-Samuelson), Wood concluded that trade globalization could account for rising inequality in the rich North <u>and</u> falling inequality in the poor South. Wood's research has been met with stiff critical resistance. Since his book appeared, we have learned more about the inequality and globalization connection in the Third World. The standard Stolper-Samuelson prediction would be that unskilled labor abundant poor countries should undergo egalitarian trends in the face of globalization forces, unless they are overwhelmed by industrial revolutionary labor-saving events on the upswing of the Kuznets Curve (Kuznets 1955), or by young adult gluts generated by the demographic transition (Bloom and Williamson 1997, 1998). A recent review by Donald Davis (1996) reports the contrary, and a study of seven countries in Latin America and East Asia shows that wage inequality typically did not fall after trade liberalization, but rather <u>rose</u>

(Robbins 1996). This apparent anomaly has been strengthened by other studies, some of which have been rediscovered since Adrian Wood's book appeared. Of course, none of these studies are very attentive to the simultaneous role of emigration from these developing countries.

As detailed below, we have designed our empirical specification with an eye to the possibility of non-standard Stolper-Samuelson effects. Here, Davis's study is of particular interest. Davis shows that, given partial specialization, the textbook SS propositions linking external prices hold only *within* a given cone of specialization (for example, Mexico might be the capital-rich country within its cone, even if it is capital-poor relative to the U.S.). The rough empirical analogue of this observation is that greater openness might raise returns to capital or skilled labor (and thus raise inequality) only for the richest countries, and might lower returns to capital or skilled labor only for the poorest countries. As a result, we interact our measures of openness with dummy variables capturing the top and bottom thirds of the world national income per capita ranking.

As with our discussion of demographic effects, two caveats are in order before we proceed. First, the standard Stolper-Samuleson predictions can fail for reasons other than partial specialization, and the possible violation of these standard assumptions should be kept in mind in interpreting our empirical results. Second, the Stolper-Samuleson predictions apply to relative factor rewards, e.g., those of land vs. labor, of capital vs. labor, or of skilled vs. unskilled labor. Relative factor rewards have a clear intuitive connection with aggregate inequality measures, but the actual correspondence between factor rewards and inequality is no doubt fairly rough.

2.3 Strong vs Weak Versions of the Kuznets Curve Hypothesis

Simon Kuznets (1955) noted that inequality had declined in several nations across the mid-20th century, and supposed that it probably had risen earlier. Furthermore, Kuznets thought it was demand-side

forces that could explain his Curve: that is, technological and structural change tended to favor the demand for capital and skills, while saving on unskilled labor. These labor-saving conditions eventually moderated as the rate of technological change (catching up) and the rate of structural change (urbanization and industrialization) both slowed down. Eventually, the labor-saving stopped, and other, more egalitarian forces were allowed to have their impact. This is what might be called the strong version of the Kuznets Curve hypothesis, that income inequality first rises and then declines with development. The strong version of the hypothesis is strong because it is unconditioned by any other effects. Demand does it all.

The weak version of the Kuznets Curve hypothesis is more sophisticated. It argues that these demand forces can be offset or reinforced by any other force if it is sufficiently powerful. The forces of some demographic transition at home may glut the labor market with the young and impecunious early in development, reinforcing the rise in inequality. Or emigration to labor-scarce OECD or oil-rich economies may have the opposite effect, making the young and impecunious who stay home more scarce (while the old receive remittances). It depends on the size of the demographic transition and whether the world economy accommodates mass migration. A public policy committed to high enrollment rates and to the eradication of illiteracy may greatly augment the supply of skilled and literate labor, eroding the premium on skills and wage inequality. Or public policy might not take this liberal stance, allowing instead the skill premium to soar, and wage inequality with it. A commitment to liberal trade policies may allow an invasion of labor-intensive goods in labor-scarce economies, thus injuring the unskilled at the bottom of the distribution. Or, trade policies may protect those interests. And a commitment to liberal trade policies in industrializing labor-abundant countries may allow an invasion of labor-intensive goods in OECD markets, the export boom raising the demand for unskilled labor and thus augmenting incomes of common labor at the bottom. Or, trade policies may instead protect the interests of the skilled in the import-

9

competing industries. Finally, natural resource endowment may matter since an export boom in such

economies will raise the rents on those resources and thus augment incomes of those at the top who own

those resources.

The strong version of the Kuznets Curve has received most of the attention since 1955, while the

weak version has received very little. A phalanx of economists, led by Hollis Chenery and Montek

Ahluwalia at the World Bank (Chenery et al., 1974; Ahluwalia, 1976), looked for unconditional Kuznets

Curves in a large sample of countries, and the results are illustrated in Figure 1. The inequality statistic

used by Ahluwalia was simply the income share of the top 20%. Based on his 60-country cross-section

from the 1960s and 1970s, it looked very much like there was a Kuznets Curve out there. True, the more

robust portion of the Curve lay to the right; income inequality clearly fell with the development of

economically mature economies. The left tail of the Curve appeared to be less robust; there was

enormous variance in inequality experience during earlier stages of development. This strong version of

the Kuznets Curve also seemed to be supported by what historical data was available at that time, some

of it reported in Figure 2.

Oddly enough, the attack on the Kuznets Curve continued to take aim at the strong and

unconditional version long after the 1970s. Even as late as 1993, Sudhir Anand and S. Kanbur published a

paper critical of the Kuznets Curve which contained no other explanatory variable but GDP. As is by

now well known, it turned out that the Kuznets Curve disappeared from Figure 1 when dummy variables

for Asia and Latin America were added. The Latin countries tend to have higher inequality, and in the

1960s, before the Asian miracle, they were located closer to the middle of the income per capita ranking.

The Asian countries tend to have lower inequality, and were located closer to the bottom of the income

per capita ranking in the 1960s.

It seems to us that the more effective attacks on the Kuznets Curve (including that by Kuznets himself) have always been based on the quality of the income distribution data. The World Bank data was poor: there was simply very little consistency as to how income was measured, how the recipient unit was defined, and how comprehensive was the coverage of the units. Thanks to Klaus Deininger and Lyn Squire (1996), we now have an excellent inequality data base which this paper exploits. Even with this new data base, however, Deininger and Squire were unable to find any evidence supporting the Kuznets Curve that Ahluwalia saw 25 years ago in Figure 1. Once again, the strong version of the Kuznets Curve hypothesis fails. While some countries may conform to the Kuznets Curve in the late 20th century, just as many do not.

But for which countries does the strong version of the hypothesis fail, and why? When it does fail is it because demand is being overwhelmed by some combination of other forces, including cohort size and openness?

3. Inequality, Cohort Size and Openness: The Data

Deininger and Squire subject their inequality data to various quality and consistency checks. In order to be included in their "high quality" data set, an observation must be drawn from a published household survey, provide comprehensive coverage of the population, and be based on a comprehensive measure of income or expenditure. The resulting data set covers 111 countries and four decades (the 1960s through the 1990s), yielding 682 annual observations. We exclude from our analysis here a number of countries with insufficient economic data, yielding a data set covering 85 countries, and including a total of 600 annual observations. Although many countries contribute only one or two annual observations, 19 countries contribute 10 or more, permitting the analysis of inequality trends over time.

11

We focus on two measures of inequality, the Gini coefficient (GINI) and the ratio of income earned by the top income quartile to income earned by the bottom quartile (GAP). To highlight inequality patterns across regions and over time, Table 1 reports unweighted averages of these inequality measures by region and decade. Inequality follows the expected regional patterns. It is quite high in Latin America and sub-Saharan Africa, with Gini coefficients in the 1990s of 50 and 46.4, respectively. Inequality is much lower among OECD countries and along the Pacific Rim, with Gini coefficients in the 1990s of 33.0 and 39.2, respectively. T. Paul Schultz (1998) has also used this data to decompose statistically the sources of world inequality into its within and between components, concluding that two-thirds of world inequality is due to between country variation. Two-thirds is a big number, and it justifies all the recent attention of the new growth theory on country growth performance since the 1960s. Yet, it is the within country variance that motivates this paper. The within country inequality data summarized in Table 1 also confirm a point already noted by Deininger and Squire (1996) and Li, Squire and Zou (1998): inequality displays little apparent variation over time within regions. The OECD s Gini coefficient, for example, moves from 33.6 to 33.0 between the 1970s and the 1990s; and the Ginis for Latin America and the Pacific Rim are also quite stable over the past four decades, despite impressive growth, policy regime switches and demographic transitions. However, and this deserves stress, data limitations make it almost impossible to draw firm conclusions about regional inequality trends across the four recent decades. For example, the Gini coefficient for Latin America in the 1970s is based on 12 countries, while the Gini for the 1990s is based on 10; only 6 Latin countries, not necessarily representative, can be observed during both decades. Data limitations are even more severe for the GAP variable, which, it turns out, is even more easily distorted by changes in sample membership.

To study Kuznets effects, we rely on real GDP per worker, measured at purchasing power parity. Some earlier studies have relied on real GDP per capita rather than per worker, but we feel that

labor productivity is more closely connected to Kuznets's notion of stages of development. GDP per worker is viewed as a proxy for a constellation of variables which have unequal derived demand impact on factor markets, an impact which Kuznets himself summarized as (unskilled) labor-saving in early stages of development. Following many earlier studies, the possibility of this inequality turning point appearing at later stages of development is captured by adding a quadratic GDP per worker term to the model. Table 2 reveals the expected labor productivity growth patterns: real GDP per worker grows rapidly along the Pacific Rim, moderately in the OECD, and stagnates in sub-Saharan Africa and Latin America.

Our openness measure comes from Jeffrey Sachs and Andrew Warner (1995), who classify an economy as closed (dummy = 0) if it is characterized by any of the following four conditions: (i) a black market premium of 20 percent or more for foreign exchange; (ii) an export marketing board which appropriates most foreign exchange earnings; (iii) a socialist economic system; (iv) extensive non-tariff barriers on imports of intermediate and capital goods. The black market premium is generally the most decisive criteria of the four, by itself identifying the vast majority of countries considered closed. According to the Sachs-Warner index, the OECD region has been quite open since the 1960s. The Pacific Rim became open in the 1970s. Latin America waited until the first half of the 1990s to make a significant switch toward economic openness, while sub-Saharan Africa still remains closed. Since there is no generally accepted metric for assessing an country's degree of economic openness (Anderson and Neary 1994), we experiment with alternative measures of openness to test the robustness of our results based on the Sachs-Warner index.

To capture the effects of cohort size, we rely on the fraction of the labor force in its peak earning years (MATURE). Because data concerning age-specific labor-force participation rates are unavailable, we approximate this by the fraction of the adult population aged 40-59. This cohort size measure has

been relatively stable within regions over the past three decades, but it varies substantially across regions, standing far higher in the developed world than elsewhere (Table 2). Evidently, the mature adult share of the labor force rises substantially only during latter stages of the demographic transition.

4. Empirical Results

Our benchmark empirical model treats the data as decadal averages by country, following Deininger and Squire (1998). We first estimate the standard unconditional Kuznets Curve, with only real output per worker and its square as explanatory variables. We then add measures of openness and cohort size to the conditional Kuznets Curve. To assess the robustness of our results, we consider the stability of the estimated relationships over time; add to the model several additional variables identified in the literature as potential inequality determinants; experiment with alternative measures of economic openness; and explore alternative demographic variables for which our cohort size measure might act as a proxy. We then turn from our benchmark (pooled) specification to a fixed-effects specification. While the benchmark estimates are driven by the dominant cross-country variation in the data, fixed-effects estimate only the within-country variation over time. Both provide considerable support for the hypotheses that inequality follows an inverted-U as an economy's aggregate labor productivity rises, and that inequality falls as an economy's population matures. However, both the benchmark and fixed-effects models provide only limited support for the hypothesis that economic openness brings increased inequality. Cohort size has a consistent and powerful effect throughout.

4.1. Pooled Estimates

14

Since the benchmark model relies on decadal averages, each country contributes between one and four observations. The average number of observations per country in our largest sample is 2.4, or about two and a half decades. All specifications include three dummy variables describing whether an inequality observation is (i) measured at the personal or household level, (ii) based on income or expenditure, or (iii) based on gross or net income.[2] All specifications also include a dummy variable for the presence of a socialist government as well as decade dummies, the latter ensuring that the estimates are driven entirely by cross-sectional variation. The standard errors used to generate our test statistics are robust to heteroskedasticity of an unknown form.

We begin by estimating the unconditional Kuznets Curve, that is, a model containing only real output per worker and its square as explanatory variables (RGDPW and RGDPW2), along with the various dummy variables. These initial results point to a relationship between inequality and labor productivity, significant at the 1 percent level, but the relationship involving GINI does not follow the expected inverted-U (Table 3, columns 1 and 2). The estimated coefficients for RGDPW and RGDPW2 are both negative in the two GINI regressions, implying that inequality declines monotonically with the level of economic development. When inequality is measured instead by GAP (=Q5/Q1), the inverted-U does appear, but the individual coefficients are very imprecisely estimated, reflecting a high degree of collinearity between the two variables (Table 3, columns 3 and 4). Much the same holds true when the model is estimated for the four decades in our sample (not reported): for both the GINI and GAP variables, RGDPW and RGDPW2 are always jointly significant at the 1 percent level, but the estimated sign pattern is often perverse. Adding regional dummy variables for sub-Saharan Africa and Latin America changes these results but little (columns 2 and 4).

---

[2] Deininger and Squire (1996) note that measured inequality levels vary systematically along these dimensions, making it important to control for them in empirical work.

It is, of course, possible that the inverted-U posited by Kuznets is masked by other forces, such as cohort size and economic openness. After all, economic relationships are seldom expected to hold without controlling for other relevant influences.[3] In this spirit, we add to the model the measures of openness and cohort size discussed earlier, and when we do the Kuznets Curve emerges from hiding (Table 4, columns 1 and 2, 3 and 4). RGDPW and RGDPW2 are jointly and individually significant at the 1 percent level, and they display the expected sign pattern. However, it is worth noting that the estimated inequality turning point is quite high, at about $15,000 evaluated at purchasing power parity in 1985 prices.[4] For comparison, as of 1990, real output per worker stood at $36,800 in the U.S. (which passed the estimated turning point well before 1950), $16,000 in Korea, and $6,800 in Thailand. According to Kuznets, the transition from a traditional, agricultural economy to a modern, industrial economy should be essentially complete at the estimated turning-point, or at least the economy should undergo a pronounced slowdown in the rate of structural change at the turning point. Thus, it is difficult to interpret these results in the manner Kuznets would have preferred, as showing the path of inequality over the course of the agricultural-industrial transition. We return to this issue below.

Next, note that Table 4 reports emphatic support for a link between cohort size and aggregate inequality. The estimated coefficient for MATURE is negative and easily statistically significant at the 1 percent level for both the GINI and GAP variables, indicating that a more experienced labor force is associated with reduced inequality, regardless of schooling levels or its distribution. The estimated quantitative impact is also large. According to the estimated coefficients, a one-standard deviation

[3] The distinction between unconditional and conditional convergence in country income levels provides an apt analogy (Williamson 1998). Numerous studies fail to find support for unconditional convergence, but find powerful evidence of convergence after controlling for determinants of steady-state income levels.
[4] Recall that these estimates are based on output per worker, which is generally about twice as high as output per capita. Also, developing country productivity levels evaluated at purchasing power parity are often more than twice as high as productivity levels evaluated at current prices and exchange rates (Summers and Heston, 1992).

increase in this variable would lower a country's Gini coefficient by 6.5, and reduce the value of its GAP

variable by 2.8. We will return to the quantitative impact of these cohort size effects, as well as of the

other two explanatory variables, but these cohort size effects appear to be very big.

Finally, note that Table 4 does not support the view that economic openness is closely connected

with higher inequality. Nor does Table 4 support the more complex predictions of standard trade theory,

namely that poor countries who go open should become less unequal while rich countries who go open

should become more unequal. There are two specifications each under GINI and GAP. The first

specification interacts OPEN with an indicator variable which equals 1 if a country is in the top third of

the labor productivity distribution in 1975-79; this new variable is called RICH. The second specification

interacts OPEN with an indicator variable which equals 1 if a country is in the bottom third of the labor

productivity distribution in 1975-79; this new variable is called POOR. As Table 4 shows (columns 1

through 6), OPENxRICH and OPENxPOOR are always small and insignificant, indicating that the impact

of openness (as measured here) does not vary with income, productivity and human capital endowment.

Standard trade theory (Hecksher-Ohlin and Stolper-Samuelson) does not survive in these data.

As noted earlier, the standard trade theoretical predictions rest on several ancillary assumptions;

the failure of our empirical result to support those predictions may mean that one or more of the

assumptions is violated. Perhaps more important, our tests may simply lack statistical power against the

null hypothesis that inequality is unrelated to openness. Remember, we interact the Sachs-Warner

openness measure with a dummy variable which selects members of (depending on the specification) the

top or the bottom third of the world income distribution. It turns our that, by this measure, almost all

countries in the top third of the world income distribution are rated as open, and almost all countries in the

bottom third as closed. Because the available data may not permit a sharp test of the hypothesis that the

openness-inequality relationship should vary with the level of development—and in light of the negative

openness results reported above—the remainder of this paper treats the openness-inequality relationship as independent of level of development.

Turning to the direct effect of openness, the coefficient on the Sachs-Warner OPEN variable is negative and statistically significant at the 1 percent level for the GINI variable (columns 1 and 2); and negative but significant at the 10 percent level in only one of the two specifications for the GAP variable. According to these estimated coefficients, an economy rated as fully open (dummy = 1) would have a Gini coefficient of 3.5 below that of an economy rated as fully closed (dummy = 0). Given that the cross-country standard deviation for Gini coefficients is close to 10, the maximum quantitative impact of 3.5 does not appear to be very large (and only 7 percent of the Latin American Gini in the 1990s). Similarly, according to the estimated coefficients, the GAP variable is only 14 percent higher for a closed than for an open economy, a reduction of only about 1.3 percent evaluated at the sample average for the 1990s.[5]

4.2. Checking Robustness

To evaluate the robustness of these results, we experiment with a number of alternative specifications. We begin by adding dummy variables for sub-Saharan Africa and Latin America to control for unobserved factors peculiar to these regions (Table 4, columns 3 and 6). Now how do our three main hypotheses perform? First, and most important for this paper, the link running from older working age populations to lower inequality remains significant at the 1 percent level. Second, the Kuznets Curve persists. Deininger and Squire (1998) found that the Kuznets Curve disappeared when African and Latin American dummies were introduced, a finding consistent with those writing in the 1970s and 1980s in the wake of Montek Ahluwalia's (1976) work for the World Bank. In contrast, the addition of these regional dummies to our conditional model makes only modest changes in the evidence supporting the Kuznets

---

[5] The cross-country standard deviation is close to 5.0.

Curve. For the GINI variable, RGDPW and RGDPW2 are easily significant at the 1 percent level, while the estimated productivity turning point falls slightly. For the GAP variable, the statistical significance of the productivity variables falls from the 1 percent level, but still retains significance at the 5 percent level. Third, the evidence of any link between economic openness and inequality essentially disappears. The coefficient for OPEN retains its negative sign, but is far from significant statistically.

We next explore the stability of the empirical relationships over time, estimating the models separately for each decade. (The estimates will also be influenced by decadal differences in the availability of the inequality data.) The results lead to some softening of the evidence supporting the Kuznets Curve (Table 5). For the GINI variable, the coefficients for RGDPW and RGDPW2 are of the expected signs and jointly statistically significant at or close to the 1 percent level for the 1970s and 1980s, and are significant at the 10 percent level for the 1960s. However, there is no evidence of a Kuznets Curve in the 1990s. Similarly, for the GAP variable, coefficients for RGDPW and RGDPW2 are of the expected signs and jointly statistically significant at the 5 percent level for the 1970s and 1980s, but switch signs and fall well short of statistical significance for the 1990s. In short, it seems wise to be tentative even about the emergence of a conditional Kuznets Curve in these data. After all, while the poor results for the 1960s might reflect the small sample size (in particular, there are few inequality observations for Africa or Latin America), the results for the 1990s are just plain negative.

Splitting the sample by decade tends to increase the already strong support for cohort size effects on inequality. The MATURE variable attains 5 percent significance levels everywhere but once (for the GAP variable in the 1960s, a period for which the sample size is small). In contrast, the Sachs-Warner openness measure—treated here as the simple additive variable OPEN since Table 4 rejected complex interactions—attains a conventional significance level for only one specification, that for the GINI variable in the 1960s.

19

The large theoretical and empirical literature on inequality has identified many other potentially important inequality determinants. We further examine the robustness of our empirical results by adding a number of these other determinants to our benchmark equations (Table 6). Francois Bourguignon and Christian Morisson (1998) focus on the role of relative labor productivity in agriculture and non-agriculture to capture Kuznets's notion that the differential development of these sectors plays a key role in explaining inequality. These authors also include arable land per capita to capture a potential link between natural resource endowment and inequality, and the secondary school enrollment ratio, to capture the intuitive notion that broader access to education reduces inequality.

Table 6 confirms the importance of the Bourguignon-Morisson agricultural variables in explaining inequality. The productivity ratio between industry and agriculture is significant at the 1 percent level, bigger productivity gaps contributing to greater inequality. The estimated coefficient implies that a reduction in the productivity ratio from 7.0 to 1.5 (the values, respectively, for Peru and the U.S. in the early 1990s) would lower a country's Gini coefficient by 2.2, compared with a cross-sectional standard deviation of about 9.7. Similarly, a more abundant agricultural endowment is associated with higher inequality, supporting the view that abundant resources can be a social "curse" as well as a drag on growth (Sachs and Warner 1995).[6] The secondary school enrollment ratio has the expected sign, but it is significant at the 10 percent level for only the GINI inequality measure. For both the GINI and GAP variables, however, the Kuznets Curve and cohort size effects remain significant at the 1 percent level, with little change in the coefficient estimates.

Note that Table 6 also adds a measure of financial depth (M3/GDP) and political freedom

---

[6] We experimented by measuring natural resource abundance as the share natural resource exports in GDP, rather than as agricultural land per capita. The alternative variable was statistically insignificant. (Natural resource exports include fuels, minerals and primary agricultural products.)

(Freedom),[7] both of which were suggested by Squire and two collaborators (Li, Squire and Zou 1998). Regarding the former, some inequality theories argue that countries with poorly developed financial systems will have higher inequality since the poor, lacking collateral, will be unable to make profitable investments. In any case, neither variable is significant in our data. A final specification drops variables which are insignificant at the 10 percent level, and adds dummy variables for Latin America and Africa, with little effect on the results.

The largely negative results described above concerning the relationship between inequality and economic openness could reflect the choice of a poor or misleading index of the latter. Similarly, the positive results concerning the relationship between inequality and our measure of cohort size could reflect a proxy relationship between this variable and some relevant, omitted demographic variable. To explore these possibilities, we experimented with several alternative measures of openness, and added several alternative demographic variables to the model.

As alternative measures of openness, we used measures of the presence of capital controls,[8] quantitative and tariff restrictions on imports, the share of imports plus exports in GDP, and the portion of this variable orthogonal to variables designed to capture a country's "natural" level of openness: the logs of country size, population, per capita income, per capita crude proven oil reserves, the average distance from trading partners, and two dummy variables describing, respectively, whether a country is an island or is landlocked.[9] None of the alternative openness measures was significant at the 10 percent level when

---

[7] FREEDOM is taken from the Barro-Lee data set, and it is a geometric average of two indices, one measuring civil liberties and one measuring political rights.

[8] The IMF records four policies restricting capital flows: (1) separate exchange rates for capital account transactions, (2) payment restrictions for current transactions, (3) payment restrictions for capital transactions, (4) mandatory surrender of export proceeds. For each of the four possible restrictions, we define a dummy variable equal to 1 when the restriction is in place, and 0 otherwise. We then take the sum of the four dummy variables as our measure of the presence of capital controls.

[9] Exports plus imports as a share of GDP is often used as a measure of openness—indeed, Summers and Heston (1995) simply label the variable as OPEN—although it has no clear connection with openness in an economically relevant sense. Standard trade models imply that a country's product and factor prices might be determined entirely

used in place of the Sachs-Warner OPEN index. [10] The cross-country data, it appears, do not support the

hypothesis that more open economies will suffer from higher inequality. It should be stressed, however,

that the evidence supporting a Kuznets Curve was unaffected during these experiments, remaining

significant at the 1 percent level for both the GAP and GINI variables. The same held true for the cohort

size impact on inequality.

To check the robustness of the cohort size effect and our choice of MATURE, we added the

following demographic variables to the model, one at a time: the total fertility rate, the population growth

rate, the labor-force growth rate, the infant mortality rate and life expectancy at birth. Our preferred

cohort size measure, of course, depends on the behavior of age-specific fertility and mortality rates over

several previous decades. Even so, MATURE could serve as an excellent point-in-time proxy for such

demographic variables: for the 1990-94 period, the cross-country correlation of MATURE with labor

force growth and the total fertility rate is -0.88 and -0.74, respectively. This point is important because

some models of fertility choice imply that fertility will fall as income inequality declines (Perotti 1996).

According to this reasoning, the negative estimated coefficient for MATURE could be capturing the

endogenous response of fertility to inequality, rather than a cohort size effect, as we have inferred.

Our robustness tests suggest that our inference is correct: our principal cohort size findings are unaffected

by adding the alternative demographic variables to the model. Of the new variables, the total fertility rate

and life expectancy at birth are statistically significant at the 5 percent level, but only when the model

in the world market even with a low trade share, or diverge substantially from their free-trade values even with a high trade share. Moreover, much of the variation in the trade/GDP ratio is explained by country size and population, although these variables should be unrelated to a country's trade policy stance. We take the residual of OPEN from the variables listed in the text as a crude attempt to capture the variation in the trade/GDP ratio potentially explained by economic policy.

[10] For brevity, we do not report these results here. The specifications correspond to Table 4, columns 1 and 4, but with only a simple, non-interacted measure of openness.

does not include dummy variables for sub-Saharan Africa and Latin America. [11] In contrast, MATURE is always easily statistically significant at the 1 percent level, with little change in the estimated coefficient. RGDPW and RGDPW remain jointly significant at the 1 percent level, with little change in the estimated inequality turning point.

The results described above provide emphatic support for the link between inequality and cohort size. They also offer strong, even if not unequivocal, support for a Kuznets Curve. Even so, our empirical models are not without their flaws. First, the estimates suffer from possible simultaneity bias, as is true of most other work in this area. The dearth of variables correlated with the relevant explanatory variables, and clearly uncorrelated with disturbances to inequality, makes it difficult to address this issue in a satisfactory way. Equally important, the estimates are likely to suffer from omitted-variable bias. Our strategy thus has been to address this issue by testing the robustness of our principal results to the inclusion of other variables identified in the literature as potential inequality determinants. An alternative strategy is to rely on fixed-effects.

4.3. Fixed-Effects Estimates

Fixed-effects are estimated by adding country-specific dummy variables to the model, and it removes one potential source of bias: unobserved country-characteristics that (i) affect inequality and (ii)

---

[11] Again, for brevity, we do not report these results here. The specifications correspond to Table 4, columns 1 and 4, but with only a simple, non-interacted measure of openness, and including both MATURE and the alternative demographic variable.

are correlated with included explanatory variables. [12] It is important to keep in mind, however, that any

reduction in bias comes at a significant cost since it removes all cross-sectional variation from the data,

potentially reducing the efficiency of parameter estimates. The loss in efficiency could be sizeable in our

data since most of the variation in inequality and in the principal explanatory variables is across countries,

rather than within countries over time. Table 7 makes this point by showing the fraction of the variance in

GINI, GAP and the explanatory variables accounted for by a set of country dummies alone. The

explained fraction exceeds 85% for both dependent variables and for all five explanatory variables listed

in Table 6 but one, and for that case, OPEN, it still exceeds 75%.

Given the evidence in Table 7 just summarized, it is quite clear that fixed-effects estimates face a

daunting hurdle. Nonetheless, we press bravely on with the fixed-effects specifications in Table 8. As

---

[12] Roughly speaking, a fixed-effects model will exhibit smaller expected omitted-variable bias if the fixed-by-country components of the included and omitted variables are more closely related than the time-varying components of the included and omitted variables. Consider a simple bivariate example. Suppose we have a panel data for a variable,

Y, generated according to the process: $Y_{i,t} = b_1 X_{1,i,t} + b_2 X_{2,i,t} + m_{t,t}$, where $m_{t,t}$ is i.i.d. X_1 is observed, but

X_2 is unobserved. Let X_1 be a random variable generated according to the process: $X_{1,i,t} = h_{1,i} + e_{1,i,t}$. Here,

$h_{1,i}$ is a country-specific disturbance, constant over time, and $e_{1,i,t}$ is an orthogonal i.i.d. observation-specific

disturbance. Finally, let X_2 be generated according to a similar process. The simple algebra of omitted variable bias

implies that, for a pure pooled model: $E(\hat{b}_1) = b_1 + b_2 \dfrac{s_{12} + n_{12}}{s_1^2 + n_1^2}$. Here, $s_{12} = Cov(h_1, h_2)$,

$s_1^2 = Var(h_1)$, $n_{12} = Cov(e_{1,t}, e_{2,t})$, and $n_1^2 = Var(e_1)$, with all relationships holding for every individual i.

For a simple fixed-effects model: $E(\hat{b}_1) = b_1 + b_2 \dfrac{n_{12}}{n_1^2}$. As a result, the bias is smaller for the fixed effects model

when: $\dfrac{s_{12}}{s_1^2} > \dfrac{n_{12}}{n_1^2}$.

24

before, we find no "unconditional" Kuznets Curve: however, a Kuznets Curve does emerge when controls

for cohort size, educational supply and openness are added to the model. The estimated inequality turning

point remains high, in fact, somewhat higher than we found for the pooled model. Cohort size effects

remain powerful, with the mature adult populations having lower inequality, and the same is true of high

levels of educational supply (proxied by secondary enrollment rates).

The residuals from these fixed-effects estimates, however, display pronounced serial correlation.

At best, this suggests that the underlying structural disturbances are serially correlated, leaving estimated

standard errors biased, and invalidating statistical inference. Worse, residual serial correlation often points

to the influence of omitted, serially dependent explanatory variables, raising the prospect that coefficient

estimates as well as standard errors are biased (Davidson and MacKinnon 1993, p. 364). Serial

correlation can be addressed by relying on an autoregressive (AR) or lagged dependent variable (LDV)

specification. The LDV specification is more robust to the presence of omitted, serially dependent

variables, and is generally preferred in modern empirical work. Davidson and MacKinnon (1993, pp. 346-

366) describe how an AR model of a given order can be nested within a corresponding LDV model via

common factor restrictions.[13] Standard diagnostic tests favor the LDV model, and the results reported in

Table 8 (columns 3 and 7) are based on it. The LDV specification sharpens our empirical results

considerably. The coefficients for real output per worker and its square display the expected

positive/negative sign pattern in both the GINI and GAP equations, and are jointly significant at the 1

percent level. The estimated inequality turning points are somewhat lower than found using the non-LDV

specification, at about $17,000 for both inequality measures. Our cohort size measure is also significant at

---

[13] As consecutive annual inequality observations are rare in our data set, we consider only an AR(1) and the corresponding LDV(1) model. We now consider countries with as few as three complete observations for all variables (including the lagged dependent variable) in order to avoid restricting the sample too severely. Even so, the number of countries in the sample falls from 44 to 23 for the GINI equation, and from 37 to 21 for the GAP equation.

the 1 percent level in both inequality equations. The estimated quantitative impact of the variable, however, is just over half that found under the earlier, pooled specification. [14] The results for secondary school enrollment ratios are much stronger than for the corresponding pooled models: the estimated coefficients are negative and significant at the 1 percent level for both inequality equations, and the estimated quantitative impacts are more than twice as large. Finally, the new results again fail to confirm a link between economic openness and inequality.

Although the LDV estimates represent an improvement on the naive fixed-effects estimates generated earlier, they will not satisfy readers with delicate econometric scruples. In particular, the current disturbance is correlated with the mean-deviation for the lagged dependent variable (Nickell 1981). As a result, the estimated coefficient for the lagged dependent variable is biased downward, and coefficient estimates for other RHS variables are also potentially biased. [15] The econometrician's toolkit now contains techniques for deriving unbiased panel estimates in an LDV setting (for a review, see Judson and Owen 1997), but those techniques require a richer data set than that available here. [16]

---

[14] Letting $\hat{\boldsymbol{b}}_1$ represent the estimated coefficient for the lagged dependent variable, and $\hat{\boldsymbol{b}}_2$ the estimated coefficient for MATURE, the estimated long-run impact of a change in MATURE is given by $\dfrac{\hat{\boldsymbol{b}}_2}{1-\hat{\boldsymbol{b}}_1}$. The coefficient estimates imply that an increase in of one percentage point in MATURE (say, from 35 to 36) lowers a country's GINI coefficient by 0.50 (say, from 40 to 39.5). A similar expression holds for the long-run impact of changes in the other RHS variables.

[15] Fixed effects transforms all variables into deviations from country-specific means. Consider an LDV model with a single RHS variable, with a typical transformed observation:

$$y_{i,t} - \bar{y}_i = \boldsymbol{b}_1(y_{i,t-1} - \bar{y}_{i,-1}) + \boldsymbol{b}_2(x_{i,t-1} - \bar{x}_{i,-1}) + \boldsymbol{e}_{i,t}.$$ Given a total of T observations for each country, we have

$$E(\bar{y}_{i,-1}, \boldsymbol{e}_{i,t}) = \frac{1}{T-1}\boldsymbol{s}_e^2,$$ leading to biased parameter estimates.

[16] The estimators proposed by Anderson and Hsiao (1981) and Arellano and Bond (1991) instrument for the lagged dependent variable using deeper lags. The estimator proposed by Arellano and Bover (1995) transforms the data into deviations from forward-looking means. These estimators are infeasible given an unbalanced panel with few complete consecutive observations.

However, standard fixed-effects LDV estimates are consistent in *T*: any bias fades as the time series dimension of the panel grows large. We exploit this fact by restricting the sample to countries with at least seven complete observations (that is, current inequality *and* the first lag of inequality must be observed at least seven times). The only effect of this is to increase the quantitative and statistical significance of the estimated Kuznets Curve, the cohort size and the education effects (Table 8, columns 4 and 8).

The cross-section data, exploited using the pooled model, and the time series data, exploited using the fixed-effects LDV model, point to powerful Kuznets Curve and cohort size effects. Both approaches suffer from econometric shortcomings. However, as these shortcomings differ in character, the pooled and fixed-effects results can be seen as mutually supporting.

4.4. Quantitative Implications

Tables 9 and 10 explore the impact on inequality of demand (proxied by the Kuznets Curve), openness and cohort size. The figures in Table 9 show how inequality would be affected were the regional values of the three explanatory variables replaced by OECD values (columns 1 and 2) or by Pacific Rim values (columns 3 and 4). The biggest effects coming from this exercise are those associated with cohort size. Compared with the OECD, both Africa and Latin America had much greater inequality, the Gini Coefficient being 13.4 points higher in the 1990s in Africa and 17 points higher in Latin America (Table 1). Table 8 shows that if Africa had the same demographic mix as the OECD, inequality (measured by the Gini coefficient) would have been lower by 8.58 points, cohort size accounting for almost two-thirds of the difference between the two regions. If Latin America had the same demographic mix as the OECD, inequality would have been lower by 8.07 points, cohort size accounting for almost half of the difference between the two regions. Compared with the Pacific Rim countries, inequality (again measured by the Gini coefficient) in Africa and Latin America in the 1990s was much higher, bigger by

27

7.2 points in Africa and by 10.8 points in Latin America (Table 1). Table 8 shows that if Africa had the same demographic mix as the Pacific Rim, inequality would have been lower by 3.57 points, cohort size accounting for about half of the difference between the two regions. If Latin America had the same demographic mix as the Pacific Rim, inequality would have been lower by 3.06 points, cohort size accounting for almost a third of the difference between the two regions. Openness (OPEN) also helps account for the inequality differences between regions in Table 8 (whether GINI or GAP), but its contribution is tiny compared with cohort size. The Kuznets demand effects (RGDP per worker) are also smaller than cohort size, and they account for none of the differences between Africa and the Pacific Rim or Latin America and the Pacific Rim.

While Table 9 explores the impact of the three explanatory variables on between-region inequality differences in the 1990s, Table 10 explores their impact on within-region inequality changes from the 1970s to the 1990s. Table 1 shows that within-region inequality change over these two decades was small, and that cohort size changes were serving to raise inequality in Africa, lower it in the OECD and the Pacific Rim, and to change it not at all in Latin America.

4.5 The Future

The estimation results can also be used to assess the effect of anticipated demographic change on inequality over the next half century. As is well known, the currently developed world is greyer than the currently developing world. The contrast is starkest for the OECD and sub-Saharan Africa. MATURE, the proportion of the adult population 15-69 who are aged 40-59, stood at 34.7 among OECD countries in the early 1990s, but at only 23.4 in Africa (Table 11). Even among Pacific Rim countries, the mature adult share was only 27.7.

The coming decades will witness substantial convergence among regional age distributions, as

birth rates and adult mortality in the currently developing world continue to fall.[17] In Latin America and the Pacific Rim, MATURE is expected to rise by about 9 percentage points between the early 1990s and 2025, 33.4 and 36.9, respectively. A further, more modest increase is expected for the years between 2025 and the middle of the century. In Africa, the expected sequence is somewhat different: MATURE shows a moderate increase between the early 1990s and 2025, but a much larger increase between 2025 and 2050.   Among OECD countries, a moderate increase in the mature adult share is expected between 1995 and 2025, with a slight decline in the subsequent decades.

Our empirical results suggest that these demographic changes will be a powerful force promoting reduced inequality the world round. The impact should be strongest in the currently developing world, where the rise in MATURE will be most pronounced. According to our estimates, the rise in the mature adult share of the labor force, taken by itself, will reduce Latin America's Gini coefficient from 50 to 42.8 by 2025, with a further, more modest decline between 2025 and 2025. The Gini coefficient for Pacific Rim countries is estimated to fall from a relatively low 39.2 to a still lower 31.4 by 2025, before stabilizing. Demography is estimated to bring only a modest decline before 2025 in African inequality, with the Gini coefficient falling to 43.5, from 46.4 in the early 1990s. However, the rapid rise in MATURE during 2025-2050 would push Africa's Gini coefficient down sharply to 37.8.   The OECD, for its part, would see a moderate decline in inequality up to 2025, followed by a modest rise. Note that these demographic changes would still leave inequality in Latin America and Africa well above OECD or Pacific Rim levels, although the gap would be reduced.

5. Explaining Cohort Size Effects on Inequality

---

[17] The figures cited here come from the U.N.'s "medium variant" population projection.

This section attempts to place the cohort size effects estimated above in context, by drawing on earlier theoretical and empirical work linking demography and inequality. We find that our estimated cohort size effects are roughly twice as large as typical estimates from the U.S. micro literature. The effect of steady-state changes in population growth on aggregate inequality can be broken down into three channels. First, slower population growth increases the share of older, high-earning workers at the expense of younger, low-earning workers. Thus, the contribution of age structure to aggregate inequality is altered, even without any change in the age-earnings profile. Deaton and Paxson (1997) show that slower steady-state population growth *raises* aggregate earnings inequality, so long as the age-earnings profile slopes upward throughout the lifecycle.[18] Second, different age groups might be characterized by different inequality levels. Deaton and Paxson (1994, 1997) present evidence that income inequality tends to increase with age for several countries examined.[19] Slower population growth, by raising the average age of the population, should *raise* aggregate inequality through this channel. Finally, slower population growth tilts the population age distribution toward older, more experienced cohorts, possibly reducing the experience premium, and *lowering* aggregate inequality. As noted above, the consistent empirical finding is that smaller youth cohorts enjoy higher mean earnings, although estimates of the magnitude of this effect vary widely .

The first two channels identified above work through changes in the relative population weights of age groups which differ in the mean or variance of earnings, treating the age-income profile as fixed (in both first and second moments). There is no attempt to assess the impact of these two demographic events on labor markets. The third channel works through the effect of cohort size on the age-income

---

[18] The effect of slower population growth on inequality, operating through this channel, is ambiguous if labor earnings tend to decline for during the final years of working life. The ambiguity is compounded if labor-force participation declines for older adults.

[19] The authors present evidence that within-cohort inequality in consumption, income and earnings tends to rise with age in the U.S., UK, Taiwan and Thailand.

profile itself; this channel works entirely through labor-market effects. Notably, the first two channels

work against the empirical results found here, implying that a higher share of mature adults in the labor

force should be associated with higher aggregate inequality, while the third channel supports those results.

Which dominates: composition effects or labor-market effects? To our knowledge, nowhere is there an

attempt in the existing literature to assess how these three channels, working together, might affect

aggregate inequality.

We rely on simulations to answer this question. The simulation results depend on three key sets

of parameters: the age profile of labor productivity over the lifecycle; the age profile of the variance of

earnings over the lifecycle; and the elasticity of substitution in the aggregate production function between

different age groups or experience levels. A high elasticity of substitution implies of course small cohort

size effects. To fix ideas, assume that there are only two age groups, the mature and the young. The

ratio of expected earnings for old and young individuals is then given by: $\dfrac{\overline{w}_m}{\overline{w}_y} = \dfrac{\boldsymbol{g}_m}{\boldsymbol{g}_y}\left(\dfrac{L_y}{L_m}\right)^{1/\boldsymbol{e}}$, where

$\boldsymbol{g}$ is an age specific productivity parameter, and $\boldsymbol{e}$ is the elasticity of substitution in production between

young and mature workers. Mature adults enjoy higher expected incomes both because they are more

productive ($\boldsymbol{g}_m > \boldsymbol{g}_y$) and because (given positive population growth) they are relatively scarce

($L_m < L_y$).

For the age profile of the mean and variance of log income, we draw on estimates for the U.S.

from Deaton and Paxson (1994, 1997). Importantly, we treat the estimated mean age-income as

representing the age profile of labor productivity.[20] We select various values for the elasticity of

substitution across age groups. We then evaluate the inequality indexes associated with various steady-

state population growth rates (and the corresponding labor-force age distributions). The Appendix

contains a more complete description of the simulation experiments.

Several simulation details deserve note. First, the age profiles for the mean and variance of log

income refer to total, rather than simply labor, income. This choice corresponds to our country inequality

data, which also refer to total income. Second, we apply the assumed cohort size effects to total, rather

than simply labor, income. We make this simplifying assumption due to lack of information concerning the

evolution of the mix between labor and non-labor income over the course of the lifecycle. If non-labor

income rises to a sizeable fraction of labor income during the later years of working life, the simulations

will overstate the effect of relative cohort size on the age-income profile. Third, in deriving cohort size

effects, we assume that all surviving, non-elderly adults are active in the labor force. We make this

simplifying assumption to avoid having to specify the potential endogenous response of relative labor-force

participation rates to relative cohort size. To the extent that labor-force participation is lower among more

mature adults (boosting their relative scarcity), the simulations will understate the effect of cohort size on

the age-income profile.[21] Finally, estimated age effects on the mean and variance of log earnings are

based on household rather than personal income, with households identified by age of household head. It

is possible, of course, that sustained changes in population growth might have systematic effects on

changes in household composition, but it is beyond the scope of this exercise to evaluate the effect of such

---

[20] Deaton and Paxson (1994) divide household survey data into age x cohort (year of birth) cells, and calculate the mean and variance of log income for each cell. The cell observations are then regressed on a set of age and cohort dummies to derived estimated age effects.

[21] Lower labor-force participation among older adults would raise the *level* of the age premium. The derivative of total labor income with respect to cohort size depends on whether labor force participation responds positively or negatively to higher wages (that is, on whether the substitution effect outweighs the income effect). If the substitution effect is the stronger, the impact of relative cohort size on relative labor income will be magnified.

changes on aggregate inequality.

The first three sets of simulations provide a point of reference by assuming perfect substitutability in production across age groups (Table 12). The first set of simulations considers the effect of population growth rates on the mix between older, high-earning workers and younger low-earning workers; the variance of log earnings over the lifecycle is held constant. The second set of simulations considers the effect of population growth rates on the mix between older, more unequal workers and younger, more equal workers; the mean of log earnings over the lifecycle is held constant. The third set of simulations considers these two channels working together. We show the Gini coefficient and the Q5/Q1 income ratio at population growth rates of 0, 2 and 4% per annum, along with the associated values for MATURE.

The most striking result is the small magnitude of changes in inequality working through changes in the mix between older, high-wage and younger, low-wage workers (row 1). Moving from steady-state population growth of 0 to 4% indeed lowers inequality, as suggested by Deaton and Paxson, but only from 32.5 to 32.1 for the Gini coefficient and from 5.3 to 5.1 for the Q5/Q1 income ratio. (The low aggregate inequality statistics are due to the fact that we have held within-cohort inequality constant at the estimated value for the 20-24 age group.) Additional simulations (not reported here) show that any decline in inequality would be quite small even if the age-income profile sloped upward throughout the lifecycle, rather than declining gently after age 50-54.

The effect of changes in the mix between younger, low-variance and older, high-variance workers is evidently more powerful (row 2). Moving from 0 to 4% steady-state population growth lowers inequality appreciably, from 43.1 to 39.7 for the Gini coefficient, and from 9.6 to 7.8 for the Q5/Q1 income ratio. Taking the mean-earnings and variance effects together results in an inequality reduction of similar magnitude (row 3).

33

Could cohort size effects be powerful enough to reverse the conclusion that slower population growth (and a higher mature adult population share) brings greater inequality? The answer to this question depends on the elasticity of substitution between older and younger workers. We take an elasticity of substitution of 3.0 as representative of the estimates from the U.S. micro literature (see the Appendix). Under that assumption, the addition of cohort size effects is enough to reverse the presumption that faster population growth reduces aggregate inequality (row 4); inequality now remains essentially unchanged in moving from 0 to 4% population growth, as measured by both the Gini coefficient and the Q5/Q1 income ratio.

Our estimates concerning the effects of cohort size evidently imply a lower elasticity of substitution across age groups than is typically found in the U.S. micro literature. We have already observed that such work usually ignores the potential endogeneity of hours and weeks worked, educational attainment and labor-force participation rates with respect to cohort size, suggesting that estimates based on total cohort population and income might yield larger elasticities. It is also possible, of course, that substitutability across age groups is higher in the U.S. than elsewhere, or that the variance of log income rises more steeply with age in the U.S. than elsewhere. We can only raise these possibilities here. For now, we merely ask whether our macro results might correspond to a lower—but still plausible—elasticity of substitution.

Accordingly, the next simulation considers an elasticity 1.5 (row 5). Cohort size effects now overwhelm the pure population weight effects. As the steady state population growth rate falls from 4 to 0%, inequality falls substantially, from 49.1 to 44.2 for the Gini coefficient, and from 12.5 to 10.1 for the Q5/Q1 income ratio. Notably, the bulk of the inequality decline occurs in moving from 4 to 2% population. Because 4% is an extremely fast population growth rate, and 2% is still considerable, it might be wondered whether the simulation results are informative about actual country experiences.

34

It turns out, however, that the steady state assumption used in generating the simulation results dramatically understates the typical variation in relative cohort size. For example, in the simulations, the ratio of the 20-29 age group to the 45-54 age group is 2.84 at a 4% steady-state population growth rate, and 1.75 at a 2% steady-state growth rate. Yet in 1985, fully 75% of 133 countries had 20-29/45-54 age ratios above 1.87; 50% were above 2.39; 25% were above 2.69; and 10% were above 2.87. The typical demographic transition, which features rapid and then slowing population growth, evidently results in cohort size ratios corresponding to very fast steady-state population growth rates. Thus, the simulation experiments comparing 4 and 2% steady-state population growth should be quite informative about actual country experiences.

Of course, the cohort size would have an even more powerful effect on inequality if the variance of log income rises less rapidly with age than suggested by Deaton and Paxson's estimates. The authors estimate the age-variance profile after controlling for cohort or year-of-birth effects (footnote 19). It is impossible also to accommodate unrestricted year effects: for any observation, the current year is collinear with cohort and age.[22] As Deaton and Paxson (1997) note, the result of this normalization is that aggregate "trends ... are attributed to age and cohort, not to time [year]" (p. 103).[23] This observation is important in light of the well-known upward trend in inequality in the U.S. and, to a lesser extent, other

---

[22] As Deaton and Paxson note, year effects can be included by restricting them to be orthogonal to a time trend.

[23] In particular, suppose the cross-sectional variance of log income evolves according to:

$$Var(Y_t) = a + b*TIME + e_t,$$ where *TIME* represents a time trend (1, 2, etc.) and e_t is an *iid* random

disturbance. Suppose also that $b > 0$. It is then easy to show that fitting an age-cohort model to this series will generate an upward tilt in estimated age effects, with positive cohort effects (higher cohort-specific inequality) for younger cohorts. The reason is that the earliest cohorts are observed only in the first or first several surveys; they then drop out. Similarly, the latest cohorts are observed only in the last or last several surveys. As a result, earlier cohorts are "older" on average, allowing age to function as an excellent proxy for the *TIME* variable. In contrast, age and year are approximately orthogonal under the alternative age-year model. Were this model fit to the DGP just described, the upward trend in aggregate inequality would be reflected in an upward tilt in estimated year effects, while the estimated age effects would be centered around the (assumed true) value of zero.

developed countries, since the 1970s. Preliminary work by one of the authors (Higgins 1999) considers these issues in more detail, and shows using the CPS data for the U.S. that the estimated upward tilt with age in variance of log income is much less pronounced under the age-year as opposed to age-cohort model. Moreover, whereas the age-year model finds an upward trend in aggregate inequality dating from the late 1970s (reflected in positive, increasing year effects), the age-cohort model finds a strong trend in cohort effects, with large positive effects for younger cohorts.[24]

It should be stressed that the age-cohort and alternative age-year models are algebraic transformations of each other, and thus, give the same fitted values. This seems to be a rare case in which a seemingly innocuous normalization has substantive implications. There is no clear metric for choosing between the two models, but we would suggest considering the plausibility of the corresponding data generating processes. The estimated year effects from the age-year model appear open to a simple, plausible interpretation, as summarizing institutional and technological changes affecting inequality. The estimated cohort effects from the cohort-year model appear to have no ready interpretation, as they imply that more recent cohorts are (conditional on age) intrinsically more diverse than earlier cohorts.

We illustrate the importance of this issue by relying on the estimated age effects from the CPS data for the variance of log income. Slower population growth is now even more strongly linked with lower inequality: in moving from 4 to 0% population growth, the simulated Gini coefficient falls from 45.9 to 38.2.

6. Conclusion

---

[24] Indeed, under the age-cohort model, the standard deviation of log income, conditional on age, is estimated to be 0.37 log points higher for the 1975-79 cohort than for the 1915-19 cohort.

The empirical results presented in this paper provide strong support for cohort size effects on inequality the world round: large mature working age cohorts are associated with lower aggregate inequality, and large young adult cohorts are associated with higher aggregate inequality. In addition, the paper reports strong, even if not unequivocal, evidence that inequality follows the inverted-U pattern described by Simon Kuznets, tending to rise as a country passes through the early stages of development, and tending to fall as a country passes through the later stages. In contrast with most previous work on the subject, our evidence in favor of the Kuznets hypothesis comes from examining the inequality-development relationship conditional on other variables. Finally, consistent with much of recent inequality debate about rising wage inequality in the US and in other OECD economies in the 1980s, we find only limited support for the hypothesis that a policy commitment to globalization has an impact on inequality.

This paper confirms a powerful world-wide connection between cohort size and inequality. Yet, that finding must be accompanied by an important caveat, which, in turn, suggests an agenda. The paper explores the connection using aggregate or economy-wide inequality data. The cohort-size hypothesis, however, concerns the relationship between relative size and the slope of the age-earnings profile. Aggregate inequality data provides only an indirect window to look in on such cohort size effects. A definitive analysis of cohort size effects awaits the appearance of internationally comparable age-earnings profile data. [25]

---

[25] In this connection, the data developed under the Luxembourg Income Study and the World Bank's LSMS study represent rich potential resources.

References

Anand, S. and S. Kanbur (1993), "Inequality and Development: A Critique," <u>Journal of Development Economics</u> 41, 19-43.

Anderson, J. E. and J. P. Neary (1994), "Measuring the Restrictiveness of Trade Policy," <u>The World Bank Economic Review</u> 8, 151-69.

Anderson, T.W. and C. Hsiao (1981), "Estimation of Dynamic Models with Error Components," <u>Journal of the American Statistical Association</u> 76, 589-606.

Ahluwalia, M. (1976), "Inequality, Poverty and Development," <u>Journal of Development Economics</u> 3, 307-42.

Arellano, M. and S. Bond (1991), "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Data," <u>Review of Economic Studies</u> 58, 277-297.

Arellano, M. and O. Bover (1995), "Another Look at the Instrumental Variable Estimation of Error-Components Models," <u>Journal of Econometrics</u> 68, 29-51.

Barro, Robert and Jong-Wha Lee (1994), <u>Data Set for a Panel of 138 Countries</u>, available from the National Bureau of Economic Research at: <u>http://www.nber.org/pub/barro.lee/zip</u>.

Bloom, D. and J. G. Williamson (1997), "Demographic Change and Human Resource Development," Ch. 3 in <u>Emerging Asia: Change and Challenges</u> (Manila: Asian Development Bank).

--- (1998), "Demographic Transitions and Economic Miracles in Emerging Asia," <u>World Bank Economic Review</u> 12, 419-55.

Borjas, G. (1994), "The Economics of Immigration," <u>Journal of Economic Literature</u> 32, 1667-717.

Borjas, G., R. Freeman and L. Katz (1992), "On the Labor Market Impacts of Immigration and Trade," in G. Borjas and R. B. Freeman (eds.), <u>Immigration and the Work Force: Economic Consequences for the United States and Source Areas</u> (Chicago: University of Chicago Press).

Bourguignon, F. and C. Morison (1998), "Inequality and Development: The Role of Dualism," <u>Journal of Development Economics</u> 57(2), 233-258.

Chenery, H. B. et al. (1974), <u>Redistribution with Growth</u> (London: Oxford University Press).

Davidson, R. and J. Mackinnon (1993), Estimation and Inference in Econometrics (Oxford: Oxford University Press).

Davis, D. R. (1996), "Trade Liberalization and Income Distribution," NBER Working Paper 5693, National Bureau of Economic Research, Cambridge, Mass. (August).

Deaton, A. S. and C. H. Paxson (1994), Intertemporal Choice and Inequality," Journal of Political Economy 102(3), 437-467.

--- (1995), "Saving, Inequality, and Aging: An East Asian Perspective," Asia-Pacific Economic Review 1, 7-19.

--- (1997) The Effects of Economic and Population Growth on National Saving and Inequality, Demography 34(1), 97-114.

Deininger, K. and L. Squire (1996), "A New Data Set Measuring Income Inequality," World Bank Economic Review 10, 565-91.

--- (1998), "New Ways of Looking at Old Issues: Inequality and Growth," Journal of Development Economics 57(2), 259-288.

Easterlin, R. A. (1980), Birth and Fortune: The Impact of Numbers on Personal Welfare (New York: Basic Books).

Freeman, R. (1979), The Effects of Demographic Factors on Age-Earnings Profiles, Journal of Human Resources 14, 289-318.

--- (1995), "Are Your Wages Set in Beijing?" Journal of Economic Perspectives 9, 15-32.

Gastil, Raymond and Lindsay Wright (1988, etc.), Freedom in the World (Westport, CT: Greenwood Publishing Group).

Judson, R. and A. Owen (1997), "Estimating Dynamic Panel Data Models: A Practical Guide for Macroeconomists," Federal Reserve Board Finance and Economics Discussion Series mimeo.

Katz, L. and K. Murphy (1992), "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," Quarterly Journal of Economics, 107, 35-78.

Kosters, M. H. (1994), "An Overview of Changing Wage Patterns in the Labor Market," in J. Bhagwati and M. H. Kosters (eds.), Trade and Wages: Leveling Wages Down? (Washington, D.C.: AEI Press).

Kuznets, S. (1955), "Economic Growth and Income Inequality," American Economic Review 45, 1-28.

Lawrence, L. and M. Slaughter (1993), Lawrence, L. and M. Slaughter [1993], "International Trade and American Wages in the 1980s: Giant Sucking Sound or Small Hiccup?," Brookings Papers on Economic Activity, Microeconomics 2: 161-226.

Li, H., L. Squire and H. Zou (1998), Explaining International and Intertemporal Variations in Income Inequality, Economic Journal 108, 26-43.

Lindert, P. H. (1997), "Three Centuries of Inequality in Britain and America," in A. B. Atkinson and F. Bourguignon (eds.), Handbook of Income Distribution (Amsterdam: North Holland).

Lindert, P. H. and J. G. Williamson (1985), Growth, Equality and History, Explorations in Economic History 22, 341-77.

Macunovich, D. J. (1998), Relative Cohort Size and Inequality in the United States, American Economic Review 88, 259-64.

Macunovich, D. J. (1999), " The Fortune of One's Birth: Relative Cohort Size and the Youth Labor Market in the U.S.," Journal of Population Economics, forthcoming.

Murphy, K. and F. Welch (1992), The Structure of Wages, Quarterly Journal of Economics 107, 285-326.

Nickell, S. (1981), "Biases in Dynamic Models with Fixed Effects," Econometrica 49, 1417-1426.

Perotti, R. (1996), "Growth, Income Distribution and Democracy," Journal of Economic Growth 1, 149-87.

Richardson, J. D. (1995), "Income Inequality and Trade: How to Think, What to Conclude," Journal of Economic Perspectives 9, 33-55.

Robbins, D. (1996), "Trade, Trade Liberalization and Inequality in Latin America and East Asia—Synthesis of Seven Countries," Harvard Institute for International Development, mimeo. (March).

Sachs, J. and A. Warner (1995), "Economic Reform and the Process of Global Integration," Brookings Papers on Economic Activity 1, 1-68.

Schultz, T. P. (1998), "Inequality in the Distribution of Personal Income in the World: How It Is Changing and Why," Journal of Population Economics 11, 307-44.

Summers and Heston (1991), "The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950-1988," Quarterly Journal of Economics 106(2), 327-368.

--- (1995), <u>Penn World Table (Mark 5.6)</u>, available on diskette from the National Bureau of Economic Research, Cambridge, Massachusetts.

United Nations (1996), <u>Age and Sex Quinquennial, 1950-2050</u> (New York: United Nations).

United Nations (1996), <u>Demographic Indicators, 1950-2050</u> (New York: United Nations).

Welch, F. (1979), "Effects of Cohort Size on Earnings: The Baby Boom Babies Financial Bust," <u>Journal of Political Economy</u> 87, S65-97.

Williamson, J. G. (1997), "Globalization and Inequality, Past and Present," <u>World Bank Research Observer</u> 12, 117-35.

--- (1998), "Growth, Distribution, and Demography: Some Lessons from History," <u>Explorations in Economic History</u> 35, 241-71.

Wood, A. (1994), <u>North-South Trade, Employment and Inequality: Changing Fortunes in a Skill-Driven World</u> (Oxford: Clarendon Press).

--- (1995), "How Trade Hurt Unskilled Workers," <u>Journal of Economic Perspectives</u> 9, 57-80.

World Bank (1998), <u>World Development Indicators on CD-ROM</u> (Washington, D.C.: The World Bank).

</ref_section>

Appendix

Data Sources. Inequality data come from Deininger and Squire (1996). The data can be downloaded from the World Bank web site: http://www.worldbank.org/growth/dddeisqu.htm. Demographic data are taken from the United Nations diskettes, Sex and Age Quinquennial, 1950-2050 and Demographic Indicators, 1950-2050. Data concerning real output per worker and exports plus imports as a share of GDP come from the data diskette Penn World Tables (Mark 5.6), available from the National Bureau of Economic Research.   Our principal measure of openness comes from Sachs and Warner (1995). Data concerning the incidence of capital controls were developed by the International Monetary Fund, compiled by Leonardo Bartolini and Alan Drazen, and obtained from these authors via personal communication. Data concerning political rights and civil liberties were taken from Barro and Lee (1994). The Barro-Lee data set is available from the NBER web site at: http://www.nber.org/pub/barro.lee/zip. The original source of the political rights and civil liberties data is Gastil and Wright (1988, etc.). All other data come from the World Bank CD-ROM, World Development Indicators: 1998.

Simulation Details. The simulation experiments concern a population ranging in age from 20 to 79. The parameters describing the age profile for the mean and variance of log income are taken from Deaton and Paxson's (1994, 1997) estimates for the U.S. The parameters are taken from 1994, Table 1, and from various graphs in 1994, 1997 via visual approximation. Deaton and Paxson's estimates are quite similar to the authors' own estimates using the CPS data for the U.S. As noted in the test, we use the estimated age profile of mean log income as a baseline; and then alter this profile to reflect different experimental assumptions about the age distribution of the labor force and the elasticity of substitution in production between different age groups. The key exception here is that we assume that persons aged 65-79 are no longer in the labor force. For this age group, we begin with mean log income for 64-year-olds, and adjust it downward using the appropriate age factors estimated by Deaton and Paxson.

An assumed steady-state population growth rate fixes the population age distribution at zero mortality. We then apply a Metropolitan Life Insurance Company (1996) mortality table to find the surviving population for each age group. Finally, we calculate age-specific probabilities of household headship using the CPS data for the U.S., and apply these probabilities to the surviving population to generate experimental survey samples. (Note that this procedure affects the number of *observations* by age group, not the total *population* by age group; the latter is

relevant for assessing cohort size effects.) We adopt this procedure because the Deininger-Squire dataset generally reports inequality at the household rather than individual level. Sampling the entire surviving population has little effect on our results.

The final simulation experiment relies on age-year rather than age-cohort model to assess the age profile of the variance of log income. We begin by estimating age-year and age-cohort models for the variance of log income using the 1967-1997 CPS data for the U.S. We break age groups and cohorts into five-year periods. As noted earlier, our estimates for the age-cohort model appear very close to those reported by Deaton and Paxson (1994, 1997). To ensure comparability with the earlier experiments, we then adjust the Deaton-Paxson age effects to reflect the difference we find in age effects from the age-year and age-cohort models.

Cohort Effects in the Micro Literature. Finis Welch (1979), in a seminal study on the subject, takes as his measure of cohort size the percentage of all workers belonging to a given age x education group. For new entrants to the labor force, he finds that the elasticity of annual earnings with respect to cohort size ranges from -.240 for high school dropouts to -.907 for college graduates (Table 9, p. S90). However, he finds that the effects of cohort size diminish over the lifecycle: the permanent effect for high school dropouts is in fact the smallest, at -.252; the effect for high school graduates (no college) is the smallest, at –0.08.

Welch's estimates do not correspond directly with the elasticity of substitution framework used in the simulations. In particular, the dependent variable is actual rather than relative wages. Moreover, in assessing the elasticity of substitution across age groups, we must remember that an increase in the young adult age share implies a decrease in other age shares. We proceed as follows to translate Welch's results into our framework. First, we calculate the labor force age shares associated with population growth of 0, 1, 2, 3 and 4% per annum, focusing on the 20-24 and 50-54 age groups. (For simplicity, we assume zero mortality and 100% labor force participation.) At successive population growth rates (and the associated labor force shares) we apply the average *entry* elasticity across education classes to the wages of the 20-24 age group, and the average *permanent* elasticity across age groups to the wages of the 50-54 age group. We then compare the change in the log wage gap with the change in the log labor force ratio to calculate implicit elasticities of substitution. The implicit elasticities range from 2.6, in moving from 0 to 1% population, to 2.9, in moving from 3 to 4% population growth.

43

Murphy and Welch (1992) estimate elasticities of complementarity across various age and education groups. Using these estimates, the authors assess the labor market effects of increasing the relative size of younger cohorts by 20%. They find that the wages of younger high school graduates would fall by 6% relative to older graduates, implying an elasticity of substitution of 3.3. They also find that the wages of younger college graduates would fall by between 9 and 15% relative to older graduates, implying an elasticity of substitution of between 2.2 and 1.3.

Katz and Murphy (1992) directly estimate the effect of changes in relative cohort size (measured by hours worked) and relative hourly wages. Aggregating across education categories, the authors find an elasticity of substitution of 2.9 (p. 76, footnote 24).

Macunovich (1998,1999) relies on the gross fertility rate during a cohort's year of birth as a measure of cohort size. (The gross fertility rate is the number of births per female population aged 15-44.) This measure has no natural interpretation in terms of *relative* steady-state cohort size. Holding mortality constant, a high steady-state gross fertility rate implies a high steady-state population growth rate, making older workers will be relatively scarce. Yet the gross fertility rate at birth would be the same for both older and younger workers. As a result, we are unable to interpret Macunovich's estimates in an elasticity of substitution framework. It should be noted, however, that her estimates imply quite large cohort size effects. For the period since 1960, the estimated differential impact on the demographically most favored and least favored cohorts is estimated to fall between 27% and 15% of labor earnings.[26]

---

[26] Note that First, Macunovich's coefficient estimates (1998, Table 1) pertain to normalized variables, that is, variables transformed to be $N(0,1)$. A one standard deviation increase in the gross fertility rate at year of birth is estimated to reduce log cohort earnings by between -.047 and -.074, depending on the specification. The gross fertility rate at year of birth for adults in the labor force (for any part of) the period since 1960 ranges between 134 and 67, a gap of approximately 3.8 standard deviations. The maximum differential impact on log wages then ranges between 0.18 and 0.28, for an impact in levels of between 20 and 32 percentage points. Also, note that some of Macunovich's specifications includes the first and/or second differences of the gross fertility rate in addition to the level. These additional variables also have no natural interpretation in terms of *relative* cohort size.

Table 1
Inequality: Patterns by Region and Decade

	1960s	1970s	1980s	1990s
Full Sample				
Gini Coefficient	37.7	38.8	37.6	39.7
	(10.3)	(9.71)	(9.20)	(9.68)
Q5 / Q1 Ratio	9.25	9.74	8.2	8.86
	(7.68)	(6.41)	(4.95)	(5.86)
Countries	37	61	73	63
OECD				
Gini Coefficient	34.7	33.6	32.6	33.0
	(7.86)	(5.72)	(4.30)	(4.86)
Q5 / Q1 Ratio	6.94	6.64	6.20	6.49
	(3.73)	(2.60)	(1.79)	(2.28)
Countries	12	19	20	13
Africa				
Gini Coefficient	45.3	49.8	41.6	46.4
	(10.5)	(8.39)	(7.74)	(9.35)
Q5 / Q1 Ratio	12.2	17.5	9.63	12.88
	(9.01)	(3.17)	(5.81)	(8.91)
Countries	4	4	11	15
Latin America				
Gini Coefficient	53.6	50.4	50.1	50
	(5.26)	(4.94)	(5.47)	(5.35)
Q5 / Q1 Ratio	21.2	17.0	16.2	13.3
	(10.9)	(6.54)	(5.26)	(3.30)
Countries	6	12	12	10
Pacific Rim				
Gini Coefficient	37.4	39.0	38.5	39.2
	(7.05)	(7.03)	(6.76)	(7.45)
Q5 / Q1 Ratio	8.28	8.96	7.88	8.14
	(3.89)	(3.98)	(3.10)	(4.25)
Countries	6	9	10	7

Mean values, with standard deviations in parentheses. See the *Appendix* for data
sources. For each decade-region pair, the number of countries with available
inequality data is indicated under that line item. Note that apparent trends in inequality
may reflect changes in data availability.

Table 2
Income, Openess and Cohort Size: Patterns by Region and Decade

	1960s	1970s	1980s	1990s
Full Sample				
RGDP per Worker	7,425	10,063	11,237	12,265
	(6,580)	(8,222)	(9,074)	(9.965)
Open	0.329	0.383	0.425	0.648
	(0.422)	(0.481)	(0.462)	(0.456)
Mature	28.4	27.5	26.9	27.1
	(4.66)	(4.10)	(4.56)	(5.02)
OECD				
RGDP per Worker	16,194	21,734	24,860	28,083
	(5,836)	(5,999)	(6,052)	(6,835)
Open	0.825	0.900	0.925	1.0
	(0.337)	(0.308)	(0.236)	(0.0)
Mature	34.3	32.9	32.4	33.8
	(2.92)	(2.14)	(2.93)	(3.04)
Africa				
RGDP per Worker	2,398	3,272	3,490	3,380
	(1,765)	(2,584)	(2,755)	(3,056)
Open	0.032	0.045	0.141	0.318
	(0.113)	(0.213)	(0.305)	(0.454)
Mature	25.5	25.3	24.4	23.7
	(2.20)	(1.99)	(2.07)	(1.95)
Latin America				
RGDP per Worker	8,059	10,413	10,364	9,334
	(5,109)	(5,565)	(5,173)	(4,217)
Open	0.320	0.227	0.273	0.822
	(0.407)	(0.413)	(0.349)	(0.278)
Mature	25.2	24.3	23.8	24.3
	(1.47)	(1.20)	(1.92)	(2.24)
Pacific Rim				
RGDP per Worker	3,995	6,995	10,472	14,612
	(2,071)	(4,166)	(6,341)	(9,046)
Open	0.490	0.900	0.900	0.900
	(0.375)	(0.316)	(0.316)	(0.316)
Mature	27.4	26.8	26.5	27.9
	(2.47)	(3.05)	(3.91)	(4.42)

Mean values, with standard deviations in parentheses. See the *Appendix* for data sources.
All available data are used, even if no corresponding inequality data are available for some
country-decade pairs.

Table 3
The Unconditional Kuznets Curve

	Gini Coefficient		Q5/Q1 Income Ratio	
RGDPW	-7.14E-02	-2.55E-02	4.77E-03	1.290E-03
	(0.31)	(0.13)	(0.31)	(0.10)
RGDPW^2	-1.34E-02	-9.52E-03	-8.08E-04	-4.22E-04
	(2.01)	(1.67)	(1.87)	(1.16)
Joint significance	<.0001	<.0001	<.0001	.0013
Turning Point	NA	NA	$2,952	$1,528
Africa dummy		10.64		0.614
		(6.22)		(5.45)
Latin dummy		12.63		0.751
		(10.99)		(8.94)
R^2 adj.	0.373	0.624	0.336	0.587
Observations	223	223	196	196

The Q5/Q1 income ratio is measured in logs. Absolute t-statistics, in parentheses, are based on heteroskedasticity-corrected standard errors. Data are pooled by decade, with countries contributing between one and four observations. All specifications include the following dummy variables: (i) inequality data based on expenditure rather than income; (ii) inequality measured at household rather than personal level; (iii) inequality data based on gross rather than net income; (iv) socialist government; and (v) - (vii) decade. See the Appendix for data sources and definitions.

Table 4
The Kuznets Curve, Openness and Cohort Size

	Gini Coefficient			Q5/Q1 Income Ratio		
RGDPW	0.739	0.801	0.580	4.61E-02	5.14E-02	3.07E-02
	(3.22)	(3.63)	(2.77)	(2.90)	(3.51)	(2.29)
RGDPW^2	-2.57E-02	-2.65E-02	-2.01E-02	-1.38E-03	-1.49E-03	-9.34E-04
	(4.16)	(4.23)	(2.74)	(3.34)	(3.72)	(2.02)
Joint significance	<.0001	<.0001	.0002	.0030	.0010	.0441
Turning Point	*$14,377*	*$15,113*	*$14,428*	*$16,703*	*$17,248*	*$16,435*
Open	-3.74	-3.71	-1.14	-0.152	-0.179	-2.04E-02
	(2.30)	(2.47)	(0.92)	(1.50)	(1.93)	(0.24)
Open x Rich	1.10			2.08E-02		
	(0.54)			(0.16)		
Open x Poor		1.58			0.177	
		(0.39)			(0.61)	
Mature	-1.15	-1.13	-0.852	-6.57E-2	-6.52E-2	-4.44E-2
	(7.65)	(7.95)	(6.89)	(6.69)	(7.39)	(4.98)
Africa dummy			9.71			0.555
			(5.81)			(4.95)
Latin dummy			9.02			0.550
			(6.92)			(5.39)
R^2 adj	0.554	0.554	0.688	0.494	0.496	0.627
Observations	219	219	219	193	193	193

The Q5/Q1 income ratio is measured in logs. Absolute t-statistics, in parentheses, are based on heteroskedasticity-
corrected standard errors. Data are pooled by decade, with countries contributing between one and four observations.
All specifications include the following dummy variables: (i) inequality data based on expenditure rather than income;
(ii) inequality measured at household rather than personal level; (iii) inequality data based on gross rather than net
income; (iv) socialist government; and (v) - (vii) decade. See the Appendix for data sources and definitions.

Table 5
Stability of Estimates Over Time

	Dependent Variable							
	Gini Coefficient				Q5/Q1 Income Ratio			
	1960s	1970s	1980s	1990s	1960s	1970s	1980s	1990s
RGDPW	1.20	1.29	0.565	0.175	6.32E-02	9.810E-02	4.05E-02	-1.88E-02
	(2.10)	(2.74)	(2.07)	(0.29)	(0.86)	(2.33)	(2.26)	(0.39)
$RGDPW^2$	-3.16E-02	-4.32E-02	-2.13E-02	-9.91E-03	-6.95E-04	-2.92E-03	-1.21E-03	9.55E-05
	(2.01)	(1.67)	(4.32)	(0.71)	(0.29)	(2.44)	(2.45)	(0.09)
Joint significance	*.0997*	*.0124*	*.0029*	*.4023*	*.2799*	*.0498*	*.0480*	*.4894*
Turning Point	*$18,987*	*$14,931*	*$13,263*	*$8,829*	*$45,468*	*$16,798*	*$16,736*	*NA*
Open	-9.63	-4.55	-0.348	-1.23	-0.699	-0.178	1.23E-02	-1.02E-02
	(2.17)	(1.48)	(0.16)	(0.31)	(1.69)	(0.88)	(0.10)	(0.03)
Mature	-1.22	-1.09	-0.734	-1.39	-8.73E-2	-7.09E-2	-4.76E-2	-8.74E-2
	(2.08)	(2.77)	(2.93)	(4.43)	(1.34)	(2.30)	(3.24)	(3.85)
R^2 adj.	0.539	0.620	0.629	0.399	0.367	0.553	0.581	0.357
Observations	34	56	69	60	28	49	64	52

The Q5/Q1 income ratio is measured in logs. Absolute t-statistics, in parentheses, are based on heteroskedasticity-corrected standard errors. Data are pooled by decade, with countries contributing between one and four observations. All specifications include the following dummy variables: (i) inequality data based on expenditure rather than income; (ii) inequality measured at household rather than personal level; (iii) inequality data based on gross rather than net income; and (iv) socialist government. See the Appendix for data sources and definitions.

Table 6
Extending the Basic Model

	Dependent Variable					
	Gini Coefficient			Q5/Q1 Income Ratio		
RGDPW	1.04	1.00	0.600	5.22E-02	4.85E-02	3.03E-05
	(4.63)	(4.22)	(2.54)	(3.43)	(2.92)	(1.74)
$RGDPW^2$	-3.02E-02	-2.95E-02	-1.94E-02	-1.40E-03	-1.33E-03	-8.52E-10
	(4.54)	(4.20)	(2.88)	(3.19)	(2.84)	(1.79)
Joint significance	<.0001	<.0001	.0126	.0028	.0130	.1893
Turning Point	$17,219	$16,949	$15,464	$18,643	$18,233	$17,782
Mature	-1.15	-1.09	-0.945	-8.95E-2	-8.94E-02	-6.34E-02
	(6.01)	(5.06)	(6.03)	(6.74)	(6.19)	(5.26)
Secondary Enroll.	-6.61E-2	-4.92E-2		-5.39E-4	6.67E-04	
	(1.74)	(1.17)		(0.22)	(0.25)	
Ind. / Agr. Labor Prod.	0.398	0.370	0.300	1.80E-2	1.61E-02	8.56E-03
	(2.61)	(2.33)	(2.12)	(2.16)	(1.72)	(1.05)
Arable Land / Pop.	1.22	1.52	0.657	9.37E-2	0.114	5.76E-02
	(3.16)	(3.52)	(1.82)	(3.57)	(3.66)	(2.52)
M3 / GDP		-1.02E-02			-1.03E-03	
		(0.32)			(0.48)	
Freedom		0.430			2.04E-02	
		(1.03)			(0.67)	
Africa dummy			8.50			0.50
			(4.83)			(3.93)
Latin dummy			7.76			0.41
			(5.50)			(3.49)
R^2 adj.	0.561	0.526	0.643	0.541	0.502	0.586
Observations	162	153	164	141	132	143

The Q5/Q1 income ratio is measured in logs. Absolute t-statistics, in parentheses, are based on heteroskedasticity-corrected standard errors. Data are pooled by decade, with countries contributing between one and four observations. All specifications include the following dummy variables: (i) inequality data based on expenditure rather than income; (ii) inequality measured at household rather than personal level; (iii) inequality data based on gross rather than net income; (iv) socialist government; and (v) - (vii) decade. See the Appendix for data sources and definitions.

Table 7
How Large are Country Fixed-Effects?

Fraction of Sample Variance Explained by Country Dummies

Dependent Variables	
Gini	0.874
Gap	0.878
Explanatory Variables	
RGDP	0.932
$RGDP^2$	0.895
Open	0.753
Mature	0.874
Secondary Enroll.	0.868

The figures above refer to the R^2 obtained when each variable is regressed on a
constant and a set of country dummies. The calculations are based on annual data.
Only countries which have four or more complete observations for the Gini coefficient
and all five explanatory variables are included, to correspond with the data used in the
fixed-effects estimates reported in Table 4. This left 449 observations, covering 44
countries. For our Gap variable, there were 387 observations, covering 37 countries.

Table 8
Fixed-Effects Specifications

	Dependent Variable							
	Gini Coefficient				Q5/Q1 Income Ratio			
LDV			0.372	0.453			0.296	0.389
			(4.40)	(5.74)			(4.11)	(6.01)
RGDPW	3.98E-02	0.405	0.349	0.956	4.27E-03	1.76E-02	2.78E-02	5.06E-02
	(0.30)	(2.33)	(1.43)	(4.08)	(0.42)	(1.49)	(2.12)	(4.04)
$RGDPW^2$	-3.09E-03	-8.82E-03	-1.04E-02	-1.81E-02	-1.93E-04	-4.56E-04	-8.04E-04	-1.13E-03
	(1.12)	(2.69)	(2.52)	(3.97)	(1.03)	(2.11)	(3.43)	(4.22)
Joint significance	0.3014	.0227	.0029	.0002	0.4493	0.0902	<.0001	.0001
Turning point	$6,440	$22,959	$16,779	$26,409	$11,062	$19,298	$17,289	$22,389
Open		1.03	0.114	-0.385		-5.76E-3	2.13E-02	1.81E-02
		(1.47)	(.196)	(.458)		(0.111)	(0.526)	(0.432)
Mature		-0.183	-0.309	-0.462		-6.99E-3	-2.33E-02	-2.79E-02
		(2.21)	(2.96)	(4.84)		(1.01)	(3.75)	(4.43)
Secondary Enroll.		-6.56E-2	-.161	-0.202		-2.24E-3	-8.78E-03	-1.17E-02
		(3.00)	(3.89)	(4.41)		(1.47)	(3.29)	(4.33)
R^2 adj	0.909	0.915	0.957	0.921	0.872	0.880	0.946	0.939
DW Statistic	1.46	1.47	NA	NA	1.54	1.53	NA	NA
Countries	44	44	23	10	37	37	21	10
Observations	459	449	216	162	394	387	202	156

The Q5/Q1 income ratio is measured in logs. Absolute t-statistics, in parentheses, are based on heteroskedasticity-corrected standard errors. For the model without a lagged dependent variable (LDV), the panel was restricted to countries with at least four complete observations. For the LDV model, the panel was restricted to countries with at least three complete observations in columns three and six, and seven complete observations in columns four and eight. For the LDV model, note that coefficient estimates and absolute t-statistics (other than for the LDV itself) pertain to the *sum* of the estimated coefficients for the current variable value and its first lag. Similarly, the test for the joint significance of RGDPW and $RGDPW^2$ refers to the joint restriction that the coefficients for RGDPW and its first lag sum to zero, and that the coefficients for $RGDPW^2$ sum to zero. In addition to the fixed country effects, all models also included a set of year dummies.

Table 9
Regional Counterfactuals

	OECD Variable Values		Pacific Rim Variable Values	
	Gini Coefficient	Q5 / Q1 Ratio	Gini Coefficient	Q5 / Q1 Ratio
Full Sample				
RGDP per Worker	-5.11	-1.40	0.09	0.07
Open	-1.23	-0.47	-0.73	-0.28
Mature	-8.71	-3.56	-0.61	-0.31
OECD				
RGDP per Worker	NA	NA	5.19	1.28
Open	NA	NA	0.49	0.14
Mature	NA	NA	8.11	3.99
Africa				
RGDP per Worker	-1.89	0.71	3.30	3.39
Open	-2.15	-1.18	-1.66	-0.92
Mature	-13.55	-7.10	-5.45	-3.54
Latin America				
RGDP per Worker	-4.59	-1.62	0.61	0.69
Open	-0.81	-0.47	-0.31	-0.18
Mature	-13.19	-7.20	-5.08	-3.45
Pacific Rim				
RGDP per Worker	-5.19	-1.34	NA	NA
Open	-0.49	-0.18	NA	NA
Mature	-8.11	-3.10	NA	NA

The figures above show how inequality would be affected were regional variable values replaced by the values for, respectively, the OECD and the Pacific Rim. Real GDP per Worker, Open and Mature are averages for the 1990-94 period, as reported in *Table 2*. The calculations are based on the pooled regression estimates, reported in *Table 3*, columns 3 and 6. See the *Appendix* for details as to data sources and variable definitions.

Table 10
The Impact of Demand, Globalization and
Cohort Size on Inequality: 1970s to 1990s

	Change, 1970s to 1990s	
	Gini Coefficient	Q5 / Q1 Gap
Full Sample		
RGDP per Worker	0.76	0.71
Open	0.39	0.24
Mature	0.26	0.15
OECD		
RGDP per Worker	-1.91	-0.39
Open	0.15	0.06
Mature	-0.58	-0.23
Africa		
RGDP per Worker	0.08	0.10
Open	0.40	0.40
Mature	1.02	0.98
Latin America		
RGDP per Worker	-0.44	-0.60
Open	0.88	0.88
Mature	0.00	0.00
Pacific Rim		
RGDP per Worker	2.77	2.57
Open	0.00	0.00
Mature	-0.70	-0.36

The figures above show the estimated impact on regional inequality of changes in RGDPW,
Open and Mature, comparing 1970-79 with 1990-94. The figures rely on the coefficient
estimates reported in Table 7, columns 3 and 6, and the regional data reported in Table 2.
Note that these "fitted value" inequality changes are based on all available data for the three
explanatory variables; and cannot be directly compared with measured regional inequality
changes (see Table 1), which are based on shifting sample of fewer countries.

Table 11
The Future: Cohort Size Effects on Inequality

	1990s	2025	2050
Full Sample			
Mature	27.1	33.8	35.8
Gini Coefficient	39.7	34.5	32.9
Q5 / Q1 Ratio	8.9	6.7	6.1
OECD			
Mature	33.8	38.7	36.7
Gini Coefficient	33.0	29.6	31.3
Q5 / Q1 Ratio	6.5	5.4	5.9
Africa			
Mature	23.7	26.8	33.6
Gini Coefficient	46.4	43.5	37.8
Q5 / Q1 Ratio	12.9	11.0	8.1
Latin America			
Mature	24.3	33.4	36.4
Gini Coefficient	50	42.8	40.3
Q5 / Q1 Ratio	13.3	9.0	7.9
Pacific Rim			
Mature	27.9	36.9	36.9
Gini Coefficient	39.2	31.4	31.4
Q5 / Q1 Ratio	8.1	5.3	5.3

21st-century age distributions are taken from the U.N.'s "medium variant" population projection. The estimated effects of expected demographic change on inequality are based on our fixed-effects estimation results (Table 8, columns 4 and 8). Given the LDV specification, the estimated effects correspond (approximately) to the long-run impact described in footnote 14. The inequality figures for the early 1990s are based on the available data for 1990-94, and repeat Table 1, column 4.

Table 12
Population Growth and Inequality: Population Weight and Cohort Size Effects

	Inequality Measure					
	Gini	Q5 / Q1	Gini	Q5 / Q1	Gini	Q5 / Q1
Population Growth Rate	4%	4%	2%	2%	0%	0%
Population Weight Effects Only						
Fixed age tilt: mean log earnings	32.1	5.1	32.3	5.2	32.5	5.3
Fixed age tilt: variance log earnings	39.7	7.8	41.4	8.7	43.1	9.6
Fixed age tilt, mean and variance	41.0	8.3	42.5	9.2	43.9	10.1
Adding Cohort Size Effects						
Elas. of substitution = 3.0	44.3	9.6	43.9	9.7	44.0	10.1
Elas. of substitution = 1.5	49.1	12.5	45.9	10.7	44.2	10.1
Elas. of substitution = 1.5, lower age tilt in variance	45.9	10.7	40.8	8.3	38.2	7.2
Mature	0.289	0.289	0.350	0.350	0.400	0.400
Pop4554 / Pop2029	2.84	2.84	1.75	1.75	1.07	1.07

Population growth rates refer to the steady state. The surviving population, given any birth cohort size, is based on current U.S. age-specific mortality rates. Given the size of the surviving cohort, the pseudo-survey "sample" incorporates age-specific probabilities of household headship, computed using average values from the the U.S. CPS from 1960-94. Importantly, however, simulated cohort size effects are based on the en*tire s*urviving cohort, no*t t*he population of household heads, assuming 100% labor-force participation. However, basing the pseudo-survey sample on the entire surviving population has little effect on our results. The simulation age profiles for the mean and variance of log earnings are based on Deaton and Paxson's (1994) estimates for the U.S. The final simulation consider an upward slope in the age profile of log earning of half that estimated by Deaton and Paxson.