

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 database 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 offer a resolution to the apparent conflict between this macro finding on cohort size and the contrary implications of recent research based on micro data.

Keywords: inequality, demography, Kuznets Curve, openness

The empirical results presented in this article 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 point in the opposite direction. In addition, the article 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

* We thank Abhijit Banerjee, Angus Deaton, Michael Kremer, Sumner J. La Croix, Andrew Mason, Andrew Warner, and David Weil for helpful comments and suggestions. The views expressed in this article are those of the authors, and not necessarily those of Merrill Lynch or Harvard University.

** Merrill Lynch, North Tower, World Financial Center, New York, NY 10281 U.S.A., e-mail: matthew.higgins@ml.org

*** Department of Economics, Harvard University, Cambridge, MA 02138 U.S.A., e-mail: jwilliam@harvard.edu

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 United States 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 I surveys the three main hypotheses upon which this article dwells: cohort size, Kuznets Curves and openness. Section II describes patterns in inequality, openness, and cohort size across regions and since the 1950s. Section III presents pooled and fixed-effects estimates of the relationships among inequality and cohort size, Kuznets Curve effects, openness, and other variables. It also explores the quantitative significance of the estimated effects. Section IV conducts simulation exercises to evaluate potential sources of the negative link between cohort size and inequality. Section V presents our conclusions.

I Reviewing the Three Hypotheses

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 their income, 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: 1023–1024, 1044–1052]. Murphy and Welch [1992] and Katz and Murphy [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 article, 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 U.S. postwar experience with wage inequality, it might do even better worldwide. 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 post-World War II demographic transition in the Third World has generated much more dramatic changes in relative cohort size than did the baby boom in the OECD countries. 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 analysis. Define the variable MATURE as the proportion of the adult population (taken to be persons in the age range 15–69) who are 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 whereas fat prime-age cohorts doing just the opposite. Interestingly, a recent and influential paper by Deaton and Paxson [1997] identifies forces linking faster population growth (and thus fat young and thin prime-age cohorts) with *reduced* inequality. The resolution of the apparent conflict is, we think, straightforward, but is reserved for section IV.

Two caveats are in order before we proceed. First, we have relied on the micro-economics literature on cohort size 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, and 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. Cohort-size effects on income require 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 pertain 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.

Inequality and Openness

After 1973 and especially in the 1980s, the United States experienced a dismal real wage performance for the less skilled, due mostly to declining productivity growth coupled with increasing wage inequality between skills.¹⁾ 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 1995]. This inequality was manifested primarily by an increasing wage premium for workers with advanced schooling and age-related skills. While the same inequality trends were apparent elsewhere in the OECD countries in the 1980s, the increase was typically far smaller [Kosters 1994]. Most of the current debate has focused on explaining these inequality facts, and it started with the observation that rising inequality coincided with rising globalization in the form of rising trade and immigration. The latter underwent rising rates and a decline in “quality” [Borjas 1994]. Trade shares in the United States increased from 12 percent of GNP in 1970 to 25 percent in 1990 [Lawrence and Slaughter 1993], while World Bank figures document that the share of output exported from low-income countries rose from 8 percent in 1965 to 18 percent in 1990 [Richardson 1995: 34]. These inequality developments also coincided with a shift in U. S. spending patterns, which resulted in large trade deficits. Thus economists have 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 unambiguous predictions. Every country exports those products that 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 disfavor skilled labor; globalization in rich countries should favor skilled labor and disfavor unskilled labor. Lawrence and Slaughter [1993] used the standard Heckscher-Ohlin trade model to explore wage inequality and concluded that there is little evidence to support it. Instead, they 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. Borjas [1994] and his collaborators [Borjas, Freeman and Katz 1992] took a different approach, emphasizing instead how trade and immigration served to augment the U. S. 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

1) This subsection is taken from Williamson [1997: 119–121].

effective domestic labor supply. In this way, the huge U.S. trade deficit of the 1980s implied a 1.5 percent increase in the U.S. 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 from the 1960s to the 1980s in the national origin of immigrants: an increasing proportion was from less developed areas (*e. g.*, Mexico and Asia) and thus less skilled. This in turn meant that a far higher fraction of immigrants were relatively unskilled just when there were more of them.

These shifts in relative supply 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 percent of the wage decline of high school graduates in relation to that of college graduates was due to trade and immigration. He also estimated that 30 to 50 percent of the decline in the relative wage of high school dropouts *vis-à-vis* all other workers was due to these same globalization forces, one-third of which was due to trade and two-thirds to immigration. Migration was the more important globalization force producing U.S. inequality trends in the 1980s, according to Borjas.

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: 19], but whether the same factors were *stimulating* the relative demand for low-skill labor in the poor Third World. This is where Wood [1994: Ch. 6; 1995] entered the debate. Wood was one of the first economists to examine systematically inequality trends across rich industrial countries in the North and poor developing countries in the South.

Basing his results on insights derived from classical Heckscher-Ohlin theory extended by Stolper-Samuelson (hereafter cited as SS), Wood [1994] concluded that trade globalization could account for rising inequality in the rich North and 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 SS prediction is that unskilled labor-abundant poor countries should undergo egalitarian trends in the face of globalization forces, unless those forces 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 Davis [1996] reports the contrary, and a study by Robbins [1996] of seven countries in Latin America and East Asia shows that wage inequality typically did not fall after trade liberalization, but rather rose.²⁾ This apparent anomaly has been strengthened by other

2) An even more recent survey by Lindert and Williamson [2002] suggests some reasons why the SS result was not forthcoming at the time Robbins was writing but now is [Robertson 2001].

studies, some of which have been rediscovered since Wood's book appeared. Of course, none of these studies is 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 nonstandard SS 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 in relation to the United States. The rough empirical analogue of this observation is that greater openness may raise the returns to capital or skilled labor (and thus raise inequality) only for the poorest countries, and may lower the returns to capital or skilled labor only for the richest countries. As a result, we interact our measures of openness with dummy variables capturing the top and bottom thirds of the world's national income distribution.

As with our discussion of demographic effects, two caveats are in order before we proceed. First, the standard SS predictions can fail for reasons other than partial specialization. The possible violation of these standard assumptions should be kept in mind in interpreting our empirical results. Second, the SS predictions apply to relative factor rewards, *e. g.*, capital versus labor or skilled versus 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.

Strong versus Weak Versions of the Kuznets Curve Hypothesis

Simon Kuznets [1955] noted that inequality had declined in several nations across the mid-twentieth 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 laborsaving 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 laborsaving 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. Factor 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 forces if they are sufficiently powerful. The forces of a 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 scarcer (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 may 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-competing industries. Finally, natural-resource endowment may matter since an export boom in economies having one will raise the rents on those resources and thus augment the 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, whereas 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; the results are illustrated in Fig. 1. The inequality statistic used by Ahluwalia was simply the income share of the top 20 percent. Based on his 60-country cross-section from the 1960s and 1970s, it looked very much as though there was a Kuznets Curve out there. True, the




Fig. 1 The Kuznets Curve: International 60-country Cross-section from the 1960s and 1970s

Source: [Ahluwalia 1976: Table 8, 340-341]




Fig. 2 The Kuznets Curve: Historical Time Series from Five European Countries and America

Source: See Appendix

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 the historical data available at that time, some of it reported in Fig. 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 that contained no other explanatory variable but GDP. As is by now well known, it turned out that the Kuznets Curve disappeared from Fig. 1 when dummy variables for Asia and Latin America were added. The Latin American 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 were poor: there was simply very little consistency as to how income was measured, how the recipient unit was defined, or how comprehensive was the coverage of the units. Thanks to Deininger and Squire [1996], we now have an excellent inequality database, which this article exploits. Even with this new database, however, Deininger and Squire were unable to find any evidence supporting the Kuznets Curve that Ahluwalia saw 25 years ago in Fig. 1. Once again, the strong version of the Kuznets Curve hypothesis fails. While some countries may have conformed to the Kuznets Curve in the late twentieth century, just as many did not.

But for which countries does the strong version of the hypothesis fail, and why?

When it does fail, is it because some combination of other forces, including cohort size and openness, is overwhelming demand?

II 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 19 countries with insufficient economic data, yielding a data set covering 92 countries and including a total of 600 annual observations. Although many countries contribute only one or two annual observations, 19 countries contribute ten or more, permitting the analysis of inequality trends over time.

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 (Q5/Q1). To highlight inequality patterns across regions and over time, Table 1 reports unweighted averages of these inequality measures by region and decade.³⁾ 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.

Schultz [1998] has also used these data to decompose statistically the sources of world inequality into its within and between components, concluding that two-thirds of world inequality are due to between-country variation. Two-thirds represent a big number, one that 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 analysis. The within-country inequality data summarized in Table 1 also confirm a point already noted by Deininger and Squire [1996] and Li, Squire, and Zhou [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 Gini coefficients for Latin America and the Pacific Rim are also quite stable over the past four decades, despite impressive growth, switches in policy regimes, and demographic transitions.

3) Note that, for each period, the total number of observations is greater than the sum of the observations in the four regional aggregates. We consider the remaining, miscellaneous countries as too heterogeneous to merit reporting as a separate category. See the Appendix for details as to regional-group membership.

Table 1 Inequality: Patterns by Region and Decade

Region and Measure	1960s	1970s	1980s	1990s
Full sample				
Gini coefficient	37.7 (10.3)	38.8 (9.71)	37.6 (9.20)	39.7 (9.68)
Q5/Q1 ratio	9.25 (7.68)	9.74 (6.41)	8.2 (4.95)	8.86 (5.86)
No. of countries	37	61	73	63
OECD				
Gini coefficient	34.7 (7.86)	33.6 (5.72)	32.6 (4.30)	33.0 (4.86)
Q5/Q1 ratio	6.94 (3.73)	6.64 (2.60)	6.20 (1.79)	6.49 (2.28)
No. of countries	12	19	20	13
Africa				
Gini coefficient	45.3 (10.5)	49.8 (8.39)	41.6 (7.74)	46.4 (9.35)
Q5/Q1 ratio	12.2 (9.01)	17.5 (3.17)	9.63 (5.81)	12.88 (8.91)
No. of countries	4	4	11	15
Latin America				
Gini coefficient	53.6 (5.26)	50.4 (4.94)	50.1 (5.47)	50 (5.35)
Q5/Q1 ratio	21.2 (10.9)	17.0 (6.54)	16.2 (5.26)	13.3 (3.30)
No. of countries	6	12	12	10
Pacific Rim				
Gini coefficient	37.4 (7.05)	39.0 (7.03)	38.5 (6.76)	39.2 (7.45)
Q5/Q1 ratio	8.28 (3.89)	8.96 (3.98)	7.88 (3.10)	8.14 (4.25)
No. of countries	6	9	10	7

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

However—and this point deserves stress—data limitations make it almost impossible to draw firm conclusions about regional-inequality trends over the four recent decades. For example, the Gini coefficient for Latin America in the 1970s is based on 12 countries, whereas the Gini for the 1990s is based on 10 countries; only 6 Latin American countries, not necessarily representative, can be observed during both decades. Data limitations are even more severe for the Q5/Q1 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 are persuaded that labor productivity is more closely connected to the

Kuznets notion of stages of development. GDP per worker is viewed as a proxy for a constellation of variables that have unequal derived-demand impact on factor markets, an impact that Kuznets himself summarized as (unskilled) laborsaving in early stages of development. Following many earlier studies, adding a quadratic GDP per worker term to the model captures the possibility that this inequality turning point appears at later stages of development. Table 2 reveals the expected labor-productivity growth pat-

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

Region and Measure	1960s	1970s	1980s	1990s
Full sample				
RGDP per worker	7,425 (6,580)	10,063 (8,222)	11,237 (9,074)	12,265 (9,965)
Open	0.329 (0.422)	0.383 (0.481)	0.425 (0.462)	0.648 (0.456)
Mature	28.4 (4.66)	27.5 (4.10)	26.9 (4.56)	27.1 (5.02)
OECD				
RGDP per worker	16,194 (5,836)	21,734 (5,999)	24,860 (6,052)	28,083 (6,835)
Open	0.825 (0.337)	0.900 (0.308)	0.925 (0.236)	1.0 (0.0)
Mature	34.3 (2.92)	32.9 (2.14)	32.4 (2.93)	33.8 (3.04)
Africa				
RGDP per worker	2,398 (1,765)	3,272 (2,584)	3,490 (2,755)	3,380 (3,056)
Open	0.032 (0.113)	0.045 (0.213)	0.141 (0.305)	0.318 (0.454)
Mature	25.5 (2.20)	25.3 (1.99)	24.4 (2.07)	23.7 (1.95)
Latin America				
RGDP per worker	8,059 (5,109)	10,413 (5,565)	10,364 (5,173)	9,334 (4,217)
Open	0.320 (0.407)	0.227 (0.413)	0.273 (0.349)	0.822 (0.278)
Mature	25.2 (1.47)	24.3 (1.20)	23.8 (1.92)	24.3 (2.24)
Pacific Rim				
GDP per worker	3,995 (2,071)	6,995 (4,166)	10,472 (6,341)	14,612 (9,046)
Open	0.490 (0.375)	0.900 (0.316)	0.900 (0.316)	0.900 (0.316)
Mature	27.4 (2.47)	26.8 (3.05)	26.5 (3.91)	27.9 (4.42)

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

terns: real GDP per worker grows rapidly along the Pacific Rim, grows moderately in the OECD economies, and stagnates in sub-Saharan Africa and Latin America.

Our openness measure comes from Sachs and Warner [1995], who classify an economy as closed (dummy = 0) if it is characterized by any of the following four conditions: (1) a black market premium of 20 percent or more for foreign exchange, (2) an export-marketing board that appropriates most foreign-exchange earnings, (3) a socialist economic system, or (4) extensive nontariff barriers on imports of intermediate and capital goods. The black market premium is generally the most decisive criterion 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, whereas sub-Saharan Africa still remains closed. Since there is no generally accepted metric for assessing a country's degree of economic openness [Anderson and Neary 1994; Rodriguez and Rodrik 2001], 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 later stages of the demographic transition.

III 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. Our results 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. We find only limited support, however, for the hypothesis that economic openness brings increased inequality. Cohort size has a consistent and powerful effect throughout.

Pooled Estimates

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 (a) measured at the personal or household level, (b) based on income or expenditure, or (c) based on gross or net income.⁴⁾ 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 (GINI or Q5/Q1) and labor productivity, significant at the 1 percent level, but the relationship does not follow the expected inverted U (Table 3,

Table 3 The Unconditional Kuznets Curve

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

Note: 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. NA—Not applicable.

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

columns 1 and 3). The estimated coefficients in GINI for RGDPW and RGDPW² are both negative, implying that inequality declines monotonically with the level of economic development. When inequality is measured instead by 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. Much the same holds true when the model is estimated for the four decades in our sample (not reported); for both the GINI and Q5/Q1 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).

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 unless other relevant influences are controlled.⁵⁾ In this spirit, we add to the model the measures of openness and cohort size discussed earlier, and when we do so the Kuznets Curve emerges (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. It is worth noting, however, that the estimated inequality turning point is quite high, at about \$15,000 evaluated at purchasing-power parity in 1985 prices.⁶⁾ For comparison, as of 1990, real output per worker stood at \$36,800 in the United States, \$16,000 in South 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.

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 Q5/Q1 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 increase in

5) 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 they do find powerful evidence of convergence after controlling for determinants of steady-state income levels.

6) 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 1991].

Table 4 The Kuznets Curve, Openness, and Cohort Size

	Dependent Variable					
	Gini Coefficient			Q5/Q1 Income Ratio		
RGDPW	0.739 (3.22)	0.801 (3.63)	0.580 (2.77)	4.61 E-02 (2.90)	5.14 E-02 (3.51)	3.07 E-02 (2.29)
RGDPW ²	- 2.57 E-02 (4.16)	- 2.65 E-02 (4.23)	- 2.01 E-02 (2.74)	- 1.38 E-03 (3.34)	- 1.49 E-03 (3.72)	- 9.34 E-04 (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 (2.30)	- 3.71 (2.47)	- 1.14 (0.92)	- 0.152 (1.50)	- 0.179 (1.93)	- 2.04 E-02 (0.24)
Open × Rich	1.10 (0.54)			2.08 E-02 (0.16)		
Open × Poor		1.58 (0.39)			0.177 (0.61)	
Mature	- 1.15 (7.65)	- 1.13 (7.95)	- 0.852 (6.89)	- 6.57 E-2 (6.69)	- 6.52 E-2 (7.39)	- 4.44 E-2 (4.98)
Africa dummy			9.71 (5.81)			0.555 (4.95)
Latin dummy			9.02 (6.92)			0.550 (5.39)
R ² adj.	0.554	0.554	0.688	0.494	0.496	0.627
Observations	219	219	219	193	193	193

Note: 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.

this variable would lower a country's Gini coefficient by 6.5 and reduce the value of its Q5/Q1 variable by 2.8. We return below 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 that go open should become less unequal whereas rich countries that go open should become more unequal. There are two specifications each under GINI and Q5/Q1. The first specification interacts OPEN (here, the Sachs-Warner measure) with an indicator variable that equals 1 if a country was 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 that equals 1 if a country was in the bottom third of the labor-productivity

distribution in 1975–79; this new variable is called POOR. As Table 4 shows, $OPEN \times RICH$ and $OPEN \times POOR$ are always small and insignificant, indicating that the impact of openness (as measured here) does not vary with income, productivity, or human-capital endowment. Standard Stolper-Samuelson trade theory does not survive in these data.

As noted earlier, the theoretical predictions of standard trade rest on several ancillary assumptions; the failure of our empirical results to support those predictions may mean that one or more of the assumptions are 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 that selects members of (depending on the specification) the top or the bottom third of the world-income distribution. It turns out 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 article treats the openness-inequality relationship as independent of the level of development.

Turning to the direct effect of openness, the coefficient on the Sachs-Warner 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 Q5/Q1 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 Q5/Q1 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.⁷⁾

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).⁸⁾ Now how do our three main hypotheses perform? First, and most important,

7) The cross-country standard deviation is close to 5.0.

8) We experimented with adding additional regional dummies for OECD and Pacific Rim economies. These dummy variables were statistically insignificant, and coefficient estimates for other variables remained essentially unchanged.

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 of observers writing in the 1970s and 1980s in the wake of 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 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 Q5/Q1 variable, the statistical significance of the productivity variable 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.⁹⁾ The results lead to some softening of the evidence supporting the Kuznets Curve (Table 5). For the GINI variable, the coefficients for

Table 5 Stability of Regression Estimates over Time

	Dependent Variable							
	Gini Coefficient				Q5/Q1 Income Ratio			
	1960s	1970s	1980s	1990s	1960s	1970s	1980s	1990s
RGDPW	1.20 (2.10)	1.29 (2.74)	0.565 (2.07)	0.175 (0.29)	6.32 E-02 (0.86)	9.810 E-02 (2.33)	4.05 E-02 (2.26)	- 1.88 E-02 (0.39)
RGDPW ²	- 3.16 E-02 (2.01)	- 4.32 E-02 (1.67)	- 2.13 E-02 (4.32)	- 9.91 E-03 (0.71)	- 6.95 E-04 (0.29)	- 2.92 E-03 (2.44)	- 1.21 E-03 (2.45)	9.55 E-05 (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 (2.17)	- 4.55 (1.48)	- 0.348 (0.16)	- 1.23 (0.31)	- 0.699 (1.69)	- 0.178 (0.88)	1.23 E-02 (0.10)	- 1.02 E-02 (0.03)
Mature	- 1.22 (2.08)	- 1.09 (2.77)	- 0.734 (2.93)	- 1.39 (4.43)	- 8.73 E-2 (1.34)	- 7.09 E-2 (2.30)	- 4.76 E-2 (3.24)	- 8.74 E-2 (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

Note: The Q5/Q1 income ratio is measured in logs. Absolute t-statistics, in parentheses, are based on heteroskedasticity-corrected standard errors. 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.

NA—Not applicable.

9) The estimates will also be influenced by decadal differences in the availability of the inequality data.

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; they are also significant at the 10 percent level for the 1960s. However, there is no evidence of a Kuznets Curve in the 1990s. Similarly, for the Q5/Q1 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 they 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 may 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 in all cases but one—for the Q5/Q1 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 because Table 4 rejected complex interactions—attains a conventional statistical significance level for only one specification, that for the GINI variable in the 1960s.

The extensive 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). Bourguignon and Morrisson [1998] focus on the role of relative labor productivity in agriculture and nonagriculture 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-Morrisson agricultural variables in explaining inequality. The productivity ratio between industry and agriculture is statistically 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 United States 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].¹⁰⁾ The secondary-

10) We experimented by measuring natural resource abundance as the share of 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.

Table 6 Extending the Basic Regression Model

	Dependent Variable					
	Gini Coefficient			Q5/Q1 Income Ratio		
RGDPW	1.04 (4.63)	1.00 (4.22)	0.600 (2.54)	5.22 E-02 (3.43)	4.85 E-02 (2.92)	3.03 E-05 (1.74)
RGDPW ²	- 3.02 E-02 (4.54)	- 2.95 E-02 (4.20)	- 1.94 E-02 (2.88)	- 1.40 E-03 (3.19)	- 1.33 E-03 (2.84)	- 8.52 E-10 (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 (6.01)	- 1.09 (5.06)	- 0.945 (6.03)	- 8.95 E-2 (6.74)	- 8.94 E-02 (6.19)	- 6.34 E-02 (5.26)
Secondary enroll.	- 6.61 E-2 (1.74)	- 4.92 E-2 (1.17)		- 5.39 E-4 (0.22)	6.67 E-04 (0.25)	
Ind./agr. labor prod.	0.398 (2.61)	0.370 (2.33)	0.300 (2.12)	1.80 E-2 (2.16)	1.61 E-02 (1.72)	8.56 E-03 (1.05)
Arable land/pop.	1.22 (3.16)	1.52 (3.52)	0.657 (1.82)	9.37 E-2 (3.57)	0.114 (3.66)	5.76 E-02 (2.52)
M3 /GDP		- 1.02 E-02 (0.32)			- 1.03 E-03 (0.48)	
Freedom		0.430 (1.03)			2.04 E-02 (0.67)	
Africa dummy			8.50 (4.83)			0.50 (3.93)
Latin America dummy			7.76 (5.50)			0.41 (3.49)
R ² adj.	0.561	0.526	0.643	0.541	0.502	0.586
Observations	162	153	164	141	132	143

Note: 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.

school enrollment ratio has the expected sign, but it is statistically significant at the 10 percent level for only the GINI inequality measure. For both the GINI and Q5/Q1 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 (FREEDOM),¹¹⁾ both of which were suggested by Squire and two collaborators

11) FREEDOM is taken from the Barro-Lee data set and is a geometric average of two indices, one measuring civil liberties and the other measuring political rights.

[Li, Squire and Zhou 1998]. Regarding the former, some inequality theories argue that countries with poorly developed financial systems will have higher inequality because 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 that 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,¹²⁾ 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.¹³⁾ None of the alternative openness measures was significant at the 10 percent level when used in place of the Sachs-Warner OPEN index.¹⁴⁾ 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 support-

12) 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, and (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. We thank Leonardo Bartolini and Alan Drazen for providing a tabulation of the IMF data.

13) Exports plus imports as a share of GDP are 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 may be determined entirely 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, country size and population size explain much of the variation in the trade/GDP ratio, although these variables should be unrelated to a country's trade policy. 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.

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

ing a Kuznets Curve was unaffected during these experiments, remaining significant at the 1 percent level for both the Q5/Q1 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 does not include dummy variables for sub-Saharan Africa and Latin America.¹⁵⁾ In contrast, MATURE is always statistically significant at the 1 percent level, with little change in the estimated coefficient. RGDPW and RGDPW² 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 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.¹⁶⁾

15) 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.

16) An alternative strategy, explored in Higgins and Williamson [1999], is to rely on a fixed-effects model by adding country-specific dummy variables to the regression specification. A fixed-effects estimator eliminates bias arising from unobserved country-specific characteristics that (a) affect inequality and (b) are correlated with included explanatory variables. ↗

Quantitative Implications

Tables 7 and 8 explore the impact on inequality of demand (proxied by the Kuznets Curve), openness, and cohort size. The figures in Table 7 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 economies, 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 (see Table 1). Table 7 shows that if Africa had the same demographic mix as the OECD, inequality (measured by the Gini coefficient) would have been lower by 8.61 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 8.09 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 7.2 points in Africa and by 10.8 points in Latin America (Table 1). Table 7 shows that if Africa had the same demographic mix as the Pacific Rim, inequality would have been lower by 3.58 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 3.07 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

↙ The authors' results, based on annual data, support the presence of a conditional Kuznets Curve, as well as a strong negative link between inequality and the MATURE cohort share. However, there is good reason to take these results with a grain of salt. First, the fixed-effect procedure removes the dominant cross-sectional variation from the data: the authors find that more than 85 percent of the variation in inequality and the principal explanatory variables is across countries, rather than within countries over time. Thus, any reduction in estimation bias comes at a substantial potential cost in estimation efficiency. Second, the regression residuals displayed significant serial correlation, implying that the estimated standard errors are biased (or worse yet, that relevant, serially dependent explanatory variables have been omitted from the model, leaving both coefficient and standard-error estimates biased [Davidson and MacKinnon 1993: 364]). However, extant techniques for controlling for serial dependence require a fairly large number of consecutive observations—a standard that the available inequality data do not accommodate. In our data set, only 10 countries have as many as 7 adjacent annual observations for inequality and the relevant macroeconomic variables. Judson and Owen [1997] show that lagged dependent-variable models with fixed effects are subject to substantial bias even with as many as 30 time-series observations. Similarly, the panel data estimators proposed by Anderson and Hsiao [1981] and Arellano and Bond [1991] instrument for the lagged dependent variable using deeper lags; while the panel-data estimator proposed by Arellano and Bover [1995] transforms the data into deviations from forward-looking means. These estimators are also infeasible given an unbalanced panel with few complete consecutive observations.

Table 7 Regional Counterfactuals

	Changes in OECD Values		Changes in Pacific Rim Values	
	Gini Coefficient	Q5/Q1 Ratio	Gini Coefficient	Q5/Q1 Ratio
Full sample				
RGDP per worker	− 3.65	− 0.92	0.09	0.12
Open	− 0.40	− 0.06	− 0.29	− 0.05
Mature	− 5.71	− 2.28	− 0.68	− 0.31
OECD				
RGDP per worker	NA	NA	3.75	0.85
Open	NA	NA	0.11	0.01
Mature	NA	NA	5.03	1.94
Africa				
RGDP per worker	− 1.29	0.43	2.45	2.18
Open	− 0.78	− 0.18	− 0.66	− 0.15
Mature	− 8.61	− 4.65	− 3.58	− 2.19
Latin America				
RGDP per worker	− 3.23	− 1.01	0.52	− 0.35
Open	− 0.20	− 0.05	− 0.02	− 0.02
Mature	− 8.09	− 4.58	− 3.07	− 1.96
Pacific Rim				
RGDP per worker	− 3.75	− 0.94	NA	NA
Open	− 0.11	− 0.02	NA	NA
Mature	− 5.03	− 1.88	NA	NA

Note: 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 4, columns 3 and 6. See the Appendix for details as to data sources and variable definitions.
NA—Not applicable.

Table 7 (whether GINI or Q5/Q1), but its contribution is tiny compared with cohort size. The Kuznets-factor 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 7 explores the impact of the three explanatory variables on between-region inequality differences in the 1990s, Table 8 explores their impact on within-region inequality changes from the 1970s to the 1990s. It shows that within-region inequality change over the 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 change it not at all in Latin America.

The Future

The estimation results can also be used to assess the effect of anticipated demographic change on inequality. As is well known, the currently developed world is grayer than the currently developing world. The contrast is starkest for the OECD region and sub-

Table 8 The Impact of Demand, Globalization, and Cohort Size on Inequality: Changes, 1970s to 1990s

Region and Measure	Gini Coefficient	Q5/Q1 Gap
Full sample		
RGDP per worker	0.29	0.21
Open	– 0.30	– 0.05
Mature	0.34	0.17
OECD		
RGDP per worker	– 2.67	– 0.63
Open	– 0.11	– 0.01
Mature	– 0.77	– 0.26
Africa		
RGDP per worker	0.05	0.05
Open	– 0.31	– 0.10
Mature	1.36	1.29
Latin America		
RGDP per worker	– 0.20	– 0.22
Open	– 0.68	– 0.21
Mature	0.00	0.00
Pacific Rim		
RGDP per worker	1.11	0.75
Open	0.00	0.00
Mature	– 0.94	– 0.43

Note: 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 4, columns 3 and 6, and the regional data reported in Table 2. 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.

Saharan Africa. MATURE, the share of the 40–59 age group in the adult population 15–69, stood at 33.8 percent among OECD countries in the early 1990s, but at only 23.7 percent in Africa (Table 9). Even among Pacific Rim countries, the mature-adult share was only 27.9 percent.

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, to 33.4 and 36.9, respectively. For

17) The figures cited here come from the United Nations’ “medium variant” population projection.

Table 9 The Future: Cohort-Size Effects on Inequality

Region and Measure	1990s	2025	2050
Full sample			
Mature	27.1	33.8	35.8
Gini coefficient	39.7	34.0	32.3
Q5/Q1 ratio	8.9	6.6	6.0
OECD			
Mature	33.8	38.7	36.7
Gini coefficient	33.0	28.8	30.5
Q5/Q1 ratio	6.5	5.2	5.7
Africa			
Mature	23.7	26.8	33.6
Gini coefficient	46.4	43.8	38.0
Q5/Q1 ratio	12.9	11.2	8.3
Latin America			
Mature	24.3	33.4	36.4
Gini coefficient	50	42.2	39.7
Q5/Q1 ratio	13.3	8.9	7.8
Pacific rim			
Mature	27.9	36.9	36.9
Gini coefficient	39.2	31.5	31.5
Q5/Q1 ratio	8.1	5.5	5.5

Note: 21st-century age distributions are taken from the United Nations' "medium variant" population projection. The estimated effects of expected demographic change on inequality are based on the pooled estimation results (Table 4, columns 3 and 6). The inequality figures for the early 1990s are based on the available data for 1990–94, and repeat Table 1, column 4.

Latin America, a further, more modest increase is expected for the years between 2025 and the middle of the century. In Africa, the expected sequences is the opposite: 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 throughout the world. 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 42.2 by 2025, with a further, more modest decline between 2025 and 2050. The Gini coefficient for Pacific Rim countries is estimated to fall from a relatively low 39.2 to a still lower 31.5 by 2025 before stabilizing. Population aging is estimated to bring only a modest decline before 2025 in African inequality, with the Gini coefficient falling to 43.8 from 46.4 in the

early 1990s. However, the rapid rise in MATURE during 2025–50 would push the region's Gini coefficient down to 38.0. The OECD, for its part, would see a moderate decline in inequality until 2025, followed by a modest rise. Note that these demographic changes would leave inequality in Latin America and Africa well above OECD or Pacific Rim levels, although the gap would be reduced.

Before concluding this section, it is worth emphasizing the obvious: this analysis considers only the potential effect of demography on inequality. It ignores the many other factors that drive it.

IV Explaining Cohort-Size Effects on Inequality

This section attempts to place the cohort-size effects estimated above in context, by drawing on earlier theoretical and empirical work linking demographic variables 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 may be characterized by different inequality levels. Deaton and Paxson [1994; 1997] present evidence that income inequality has tended 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 that 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

18) The effect of slower population growth on inequality, operating through this channel, is ambiguous if labor earnings tend to decline 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 has tended to rise with age in the United States, the United Kingdom, Taiwan, and Thailand.

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 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: $\frac{\bar{w}_m}{\bar{w}_y} \parallel \frac{\gamma_m}{\gamma_y} \left(\frac{L_y}{L_m} \right)^{1/\varepsilon}$, where γ is an age-specific productivity parameter, and ε is the elasticity of substitution in production between young and mature workers. Mature adults enjoy higher expected incomes both because they are more productive ($\gamma_m > \gamma_y$) 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 United States from Deaton and Paxson [1994; 1997]. It is important that 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 to simply labor, income. We make this simplifying assumption for lack of information concerning the evolution of the mix between labor and nonlabor income over the course of the lifecycle. If nonlabor 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, nonelderly adults are

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 derive estimated age effects.

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 may have systematic effects on changes in household composition, but it is beyond the scope of this exercise to evaluate the effect of such 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 10). 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 life-

Table 10 Population Growth and Inequality: Population Weight and Cohort-Size Effects

Population Growth Rate	Inequality Measure					
	Gini	Q5/Q1	Gini	Q5/Q1	Gini	Q5/Q1
	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
Mature	0.289	0.289	0.350	0.350	0.400	0.400
Pop 45–54/Pop 20–29	2.84	2.84	1.75	1.75	1.07	1.07

Note: 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 by using average values from the U. S. CPS for 1960–94. Importantly, however, simulated cohort-size effects are based on the entire surviving cohort, not the population of household heads, assuming 100 percent labor force participation for those aged 20–64. Basing the pseudo-survey sample on the entire surviving population has little effect on our results, however. The simulation age profiles for the mean and variance of log earnings are based on Deaton and Paxson’s [1994] estimates for the United States.

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.

cycle 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 percent 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 workers and younger, low-wage workers (row 1). Moving from steady-state population growth of 0 to 4 percent 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 ages 50–54.

The effect of changes in the mix between younger, low-variance workers and older, high-variance workers is evidently more powerful (row 2). Moving from 0 to 4 percent 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).

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 microeconomics literature on the U. S. baby boom (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 percent 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 microeconomics literature on the U. S. baby boom. 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 may yield larger elasticities. It is also possible, of course, that substitutability across age groups is higher in the United States than elsewhere, or that the variance of log income rises more steeply with age in the United States 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 of 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 percent, 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 percent population growth. Because 4 percent is an extremely fast population growth rate, and 2 percent is still considerable, it might be wondered whether the simulation results are informative about actual country experiences.

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 percent steady-state population growth rate, and 1.75 at a 2 percent steady-state growth rate. Yet in 1985, fully 75 percent of 133 countries had 20–29/45–54 age ratios above 1.87; 50 percent were above 2.39; 25 percent were above 2.69; and 10 percent 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 percent and 2 percent steady-state population growth should be quite informative about actual country experiences.

V Conclusion

The empirical results presented in this article 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 analysis 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. Our work differs from most previous studies of the Kuznets hypothesis by examining the inequality-development relationship conditional on other variables. Finally, in accordance with much of the recent debate about rising wage inequality in the United States and other OECD economies in the 1980s, we find little support for the hypothesis that a policy commitment to globalization has an impact on inequality.

Our results concerning cohort size and inequality should be accompanied by an important caveat. Throughout our analysis, we have worked with data concerning aggregate or economy-wide income inequality. The cohort-size hypothesis, however,

concerns the relationship between relative size and the slope of the age-earnings profile. Aggregate inequality data can provide only an indirect window on such cohort-size effects. A definitive analysis of cohort-size effects awaits the development of internationally comparable data concerning age-earnings profiles.

References

- Ahluwalia, M. 1976. Inequality, Poverty and Development. *Journal of Development Economics* 3: 307-342.
- Anand, S.; and Kanbur, S. 1993. Inequality and Development: A Critique. *Journal of Development Economics* 41: 19-43.
- Anderson, J. E.; and Neary, J. P. 1994. Measuring the Restrictiveness of Trade Policy. *The World Bank Economic Review* 8: 151-169.
- Anderson, T. W.; and Hsiao, C. 1981. Estimation of Dynamic Models with Error Components. *Journal of the American Statistical Association* 76: 589-606.
- Arellano, M.; and Bond, S. 1991. Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Data. *Review of Economic Studies* 58: 277-297.
- Arellano, M.; and Bover, O. 1995. Another Look at the Instrumental Variable Estimation of Error-Components Models. *Journal of Econometrics* 68: 29-51.
- Barro, R.; and Lee, J-W. 1994. *Data Set for a Panel of 138 Countries*. Available from the National Bureau of Economic Research at <http://www.nber.org/pub/barro.lee/zip>
- Bloom, D.; and Williamson, J. G. 1997. Demographic Change and Human Resource Development. In *Emerging Asia: Change and Challenges*, Ch. 3. Manila: Asian Development Bank.
- . 1998. Demographic Transitions and Economic Miracles in Emerging Asia. *World Bank Economic Review* 12: 419-455.
- Borjas, G. 1994. The Economics of Immigration. *Journal of Economic Literature* 32: 1667-1717.
- Borjas, G.; Freeman, R.; and Katz, L. 1992. On the Labor Market Impacts of Immigration and Trade. In *Immigration and the Work Force: Economic Consequences for the United States and Source Areas*, edited by G. Borjas and R. B. Freeman. Chicago: University of Chicago Press.
- Bourguignon, F.; and Morrisson, C. 1998. Inequality and Development: The Role of Dualism. *Journal of Development Economics* 57 (2): 233-258.
- Chenery, H. B.; Ahluwalia, M. S.; Bell, C. L. G.; Duloy, J. H.; and Jolly, R. 1974. *Redistribution with Growth*. London: Oxford University Press.
- Davidson, Russell; and MacKinnon, James G. 1993. *Estimation and Inference in Econometrics*. Oxford: Oxford University Press.
- Davis, D. R. 1996. *Trade Liberalization and Income Distribution*. NBER Working Paper No. 5693 (August). Cambridge, MA: National Bureau of Economic Research.
- Deaton, A. S.; and Paxson, C. H. 1994. Intertemporal Choice and Inequality. *Journal of Political Economy* 102 (3): 437-467.
- . 1997. The Effects of Economic and Population Growth on National Saving and Inequality. *Demography* 34 (1): 97-114.
- Deininger, K.; and Squire, L. 1996. A New Data Set Measuring Income Inequality. *World Bank Economic Review* 10: 565-591.
- . 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. B. 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, R.; and Wright, L. 1988 -. *Freedom in the World*. Westport, CT: Greenwood Publishing Group.
- Higgins, M.; and Williamson, J. G. 1999. *Explaining Inequality the World Round: Cohort Size, Kuznets Curves and Openness*. NBER Working Paper No. 7224 (June). Cambridge, MA: National Bureau of Economic Research.
- Judson, R.; and Owen, A. 1997. Estimating Dynamic Panel Data Models: A Practical Guide for Macroeconomists. Federal Reserve Board Finance and Economics Discussion Series. Washington, D. C. (Mimeographed)
- Katz, L.; and Murphy, K. 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 *Trade and Wages: Leveling Wages Down?* edited by J. Bhagwati and M. H. Kosters. Washington, D. C.: AEI Press.
- Kuznets, S. 1955. Economic Growth and Income Inequality. *American Economic Review* 45: 1–28.
- Lam, David. 1997. Demographic Variables and Income Inequality. In *Handbook of Population and Family Economics*, edited by Mark R. Rosenzweig and Oded Stark, pp. 1015–1059. Amsterdam: Elsevier Science B. V.
- Lawrence, L.; and Slaughter, M. 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.; Squire, L.; and Zhou, H. 1998. Explaining International and Intertemporal Variations in Income Inequality. *Economic Journal* 108: 26–43.
- Lindert, P. H.; and Williamson, J. G. 2002. Does Globalization Make the World More Unequal? In *Globalization in Historical Perspective*, edited by M. Bordo, A. M. Taylor, and J. G. Williamson. Chicago: University of Chicago Press.
- Macunovich, D. J. 1998. Relative Cohort Size and Inequality in the United States. *American Economic Review* 88: 259–264.
- . 1999. The Fortune of One's Birth: Relative Cohort Size and the Youth Labor Market in the U. S. *Journal of Population Economics* 12: 215–272.
- Murphy, K.; and Welch, F. 1992. The Structure of Wages. *Quarterly Journal of Economics* 107: 285–326.
- Perotti, R. 1996. Growth, Income Distribution and Democracy. *Journal of Economic Growth* 1 : 149–187.
- 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, Cambridge, MA. (Mimeographed)
- Robertson, R. 2001. Relative Prices and Wage Inequality: Evidence from Mexico. Unpublished paper (October). Macalester College, St. Paul, MN. 53p.
- Rodriguez, F.; and Rodrik, D. 2001. Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence. In *NBER Macroeconomics Annual 2000*, Vol. 15, edited by B. S. Bernanke and K. Rogoff, pp. 261–325. Cambridge, MA: MIT Press.
- Sachs, J.; and Warner, A. 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: 307–344.
- Summers, R.; and Heston, A. 1991. The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950–1988. *Quarterly Journal of Economics* 106: 327–368.
- . 1995. *Penn World Table (Mark 5.6)*. Available on diskette from the National Bureau of Economic Research, Cambridge, MA.
- United Nations, Department of International Economic and Social Affairs. 1996a. *Age and Sex*

- Quinquennial, 1950–2050*. New York: United Nations.
- . 1996b. *Demographic Indicators, 1950–2050*. New York: United Nations.
- Welch, F. 1979. Effects of Cohort Size on Earnings: The Baby Boom Babies Financial Bust. *Journal of Political Economy* 87: S65–S97.
- Williamson, J. G. 1997. Globalization and Inequality, Past and Present. *World Bank Research Observer* 12 (2): 117–135.
- . 1998. Growth, Distribution, and Demography: Some Lessons from History. *Explorations in Economic History* 35 (3): 241–271.
- Wood, A. 1994. *North-South Trade, Employment and Inequality: Changing Fortunes in a Skill-Driven World*. Oxford: Clarendon Press.
- . 1995. How Trade Hurt Unskilled Workers. *Journal of Economic Perspectives* 9: 57–80.
- World Bank. 1998. *World Development Indicators: 1998*. CD-ROM. Washington, D. C.

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 *Age and Sex Quinquennial, 1950–2050* [1996a] and *Demographic Indicators, 1950–2050* [1996b]. Data documenting 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 (NBER) in Cambridge, MA. 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 those authors via personal communication. Data concerning political rights and civil liberties were taken from Barro and Lee [1994]. The complete 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–]. All other data come from the World Bank's [1998] CD-ROM *World Development Indicators: 1998*.

Regional Aggregates

Sub-Saharan Africa: Botswana, Cameroon, Central African Republic, Cote D'Ivoire, Gabon, Ghana, Guinea-Bissau, Kenya, Lesotho, Madagascar, Mauritania, Mauritius, Niger, Nigeria, Rwanda, Senegal, Sierra Leone, South Africa, Tanzania, Uganda, Zambia, Zimbabwe.

Latin America: Bolivia, Brazil, Chile, Columbia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Honduras, Mexico, Nicaragua, Panama, Peru, Puerto Rico, Venezuela.

OECD: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Turkey, United Kingdom, United States.

Pacific Rim: China, Hong Kong, Indonesia, Japan, South Korea, Malaysia, Philippines, Singapore, Taiwan, Thailand.

Other: Algeria, Bangladesh, Barbados, Bulgaria, Czechoslovakia, Egypt, Fiji, Guyana, Hungary, India, Iran, Jamaica, Jordan, Laos, Morocco, Nepal, Pakistan, Poland, Romania, Soviet Union, Sri Lanka, Trinidad and Tobago, Tunisia, Yugoslavia.

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 United States. 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 our own estimates using the Current Population Survey (CPS) data for the United States. As noted in the text, 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 a 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 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 United States 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 data set generally reports inequality at the household level rather than the individual level. Sampling the entire surviving population has little effect on our results.

The final simulation experiment relies on an age-year rather than an 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 1967–97 CPS data for the United States. 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 $\times$ education group. For new entrants to the labor force, he finds that the elasticity of annual earnings with respect to cohort size ranges from -0.240 for high school dropouts to -0.907 for college graduates [*ibid.*: Table 9, S90]. He finds, however, that the effects of cohort size diminish over the lifecycle: the permanent effect for high school dropouts is in fact the smallest, at -0.252 ; the effect for high school graduates (with 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 percent per annum, focusing on the 20–24 and 50–54 age groups. For simplicity, we assume zero mortality and 100 percent 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 percent population growth, to 2.9, in moving from 3 to 4 percent population growth.

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 percent. They find that the wages of younger high school graduates would fall by 6 percent 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

percent 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 [*ibid.*: 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. With mortality held constant, a high steady-state gross fertility rate implies a high steady-state population growth rate, making older workers 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.