On the Impact of Inequality on Productivity Growth in the Short and Long Term: A Synthesis

Huw Lloyd-Ellis
Department of Economics
Queen's University
Kingston, Ontario

Au début des années 1990, la traditionnelle relation de compromis «équité-efficacité» a été mise en question par des preuves empiriques suggérant que l'inégalité des revenus avait un impact négatif sur la croissance. Cela a encouragé le développement de plusieurs théories alternatives fondées sur des marchés incomplets, et de travaux empiriques destinés à découvrir quels mécanismes sont les plus importants. Cependant, des découvertes empiriques récentes, aussi bien au niveau micro qu'au niveau macro, ont mis en question, sur plusieurs fronts, cette théorie des marchés incomplets. Cet article examine la littérature théorique et empirique récente entourant ce débat, et pose la question de savoir si, une fois qu'on fait la distinction entre relation à court terme et relation à long terme, les opinions conflictuelles sont effectivement aussi inconsistantes qu'elles apparaissent à première vue.

In the early 1990s, the traditional "equity-efficiency" trade-off was challenged by empirical evidence suggesting a negative impact of inequality on growth. This has spurred the development of several alternative theories based on incomplete markets and empirical work aimed at uncovering the most important mechanisms. However, recent empirical findings, at both the micro and macro level, have challenged this incomplete markets view on several fronts. This paper surveys the recent theoretical and empirical literature surrounding this debate, and asks whether the opposing views are really as inconsistent as they at first appear once we distinguish between short- and long-run relationships.

Introduction

A common view amongst economists is that a trade-off exists between equity and efficiency — greater equality of income can only be bought at the cost of lower productivity. However, in the early 1990s this view was challenged by mounting evidence of a statistically significant negative impact of income inequality on subsequent per capita income growth across countries. This evidence spurred on a theoretical literature that attempts to understand why this relationship might arise due to various types of market incompleteness, and an empirical literature aimed at uncovering which of these various mechanisms are the most important. Recently, however, this "incomplete markets view" has been challenged by empirical evidence at both the micro and macro level. This paper surveys the recent theoretical and empirical literature surrounding this ongoing debate and asks whether these apparently conflicting views are as inconsistent with each other as they first appear.

The incomplete markets view draws on two branches of economic theory whose development has led to an increased acceptance of the importance of certain market failures in the growth process and, hence, for there to be a potential role for inequality in that process. Organizational Economics emphasizes informational problems that result in fundamental breakdowns in the market mechanism which cannot easily be solved.[1] This branch of economics has made significant ground in understanding the interaction of various non-market institutions (e.g., contracts and modes of organization) with the process of economic development. Endogenous growth theory emphasizes the role of positive externalities from private investments made by individual firms or households on the productivity of investments made by others (either contemporaneously or in the future), which are not taken into account by the investor. This branch of economics has made important advances in identifying the key determinants of long-run productivity growth and in providing a framework for thinking about the role of policy in influencing this growth.[2]

A key implication of the theoretical mechanisms surveyed in the third section is that greater inequality reduces aggregate productivity when two conditions are satisfied: (i) there is a market failure or some constraint on the tradeability of a key input to production (e.g., human capital), and (ii) the impact of that input on an individual's contribution to aggregate productivity or productivity growth exhibits diminishing returns. In addition to the direct impact of inequality on productivity, the theoretical literature also emphasizes factors affecting the persistence of inequality through time and between generations. If wealth inequality is a temporary phase in the development process, then its implications for long-run productivity may be less important. However, if inequality persists through time and across generations, the overall effects are magnified.

As I discuss in the fourth section, this incomplete markets view has been challenged recently on three main fronts.

- First, although there is some evidence suggesting that borrowing constraints may affect entrepreneurial activity and schooling in less-developed countries, there is little evidence that short-run borrowing constraints have economically significant effects on human capital investment in developed countries.

- Second, in contrast to early estimates of large external effects of human capital on productivity, more recent studies that attempt to control for supply-side effects and simultaneity biases suggest that these effects are small. The absence of short-run human capital externalities implies that, along an economy's balanced growth path, the direct impact of human capital dispersion on aggregate productivity is likely to be small.

- Finally, using panel data rather than just cross-country regressions, Barro (1999) and Forbes (2000) have found that a *positive* relationship rather than a negative relationship may be observed between inequality and growth.

In the section following, I argue that much of this recent empirical evidence does not necessarily directly contradict the incomplete markets view and, in fact, is consistent with much of the theoretical literature. In particular, the absence of short-run borrowing constraints need not imply that family income and/or human capital have no impact on a child's educational outcomes. There may also exist important implicit long-term borrowing constraints that determine how well children do early on in the school system. Moreover, the absence of short-run positive externalities of human capital in production does not rule out the existence of long-run positive *dynamic* externalities, which are important in the process of education itself and in the adoption and implementation of new technologies. Finally, the empirical evidence coming from panel regressions is consistent with the view that the induced relationship between inequality and growth is positive in the short run, but does not imply that inequality has a positive impact on growth in the long run.

The evidence as yet is far from conclusive and there is much to be done in identifying the true mechanisms, if any, that link inequality and growth.[3] Nevertheless, section six outlines some of the implications of distinguishing between the short-run (induced) versus long-run (causal) relationships between inequality and growth for human capital policy in Canada. In particular, it implies that although reduced inequality and higher productivity need not be conflicting objectives, a balance must be struck between the short-run disincentive effects and long-run average investment effects of reduced inequality. This perspective has important implications for the choice between private and public financing of education, the allocation of public spending across different stages of education, regional disparities in the quality of education and the optimal response to rapid skill-biased technological change.

Beyond the Equity-Efficiency Trade-Off

In a world in which perfectly functioning markets exist for *all* commodities and in which there are no impediments to trade, there would be no fundamental impact of inequality on productivity or productivity growth.[4] In such a world, the only relationship would come from policy attempts to affect the distribution of income that also distort incentives. For example, a progressive income tax that reduces inequality while distorting the labour-leisure choice would reduce productivity.[5] It is this view of the world that underlies the so-called "equity-efficiency" trade-off that was the predominant view in the 1970s.[6]

The policy implications of this perspective are clear. Any attempt to move away from the market allocation via redistributive policies will result in lower productivity and/or productivity growth by distorting incentives. Thus there would be an equity-efficiency trade-off induced by redistributive policy. For example, a progressive tax system might reduce the after-tax returns to acquiring higher education, which would remove the incentives to go to college/university. The desire to reduce inequality might be reasonable from a welfare perspective, but it could only be achieved at the expense of lower growth.

In the early 1990s this traditional view was challenged with the emergence of new, cross-country evidence suggesting a negative long-run relationship between income inequality and subsequent growth in per capita income. These empirical studies use an index of income inequality to proxy subsequent inequality in human capital or wealth. The most common measures are Gini coefficients and the income shares of the top 20 percent of the income distribution. Other studies have used inequality indices of land ownership arguing that, in less-developed economies at least, these may be a better proxy for the true distribution of wealth.

Persson and Tabellini (1994), for example, employ two different datasets, one with historical observations for nine developed countries and one with postwar observations for 56 countries. Their results suggest that an increase of 0.07 (one standard deviation) in the share of income held by the top 20 percent of the population lowers average annual growth rates by just less than one-half of one percent. Alesina and Rodrik (1994) employ different data for the period 1960–85 from up to 70 countries. They report that an increase of one standard deviation for their Gini coefficient of land distribution would lower average per capita growth rates by 0.8 percentage points per year. Both studies employ a similar methodology, obtaining a measure of inequality from at or near the beginning of a long sub-period of the data (Persson and Tabellini use 20- and 15-year sub-periods, Alesina and Rodrik use 1960–85 and 1970–85) and observing the influence of this measure on subsequent growth rates. Clarke (1995) and Perotti (1996) provide additional evidence on the robustness of this long-run negative relationship to different measures of inequality and to different regression equation specifications.

Many of these earlier studies used income inequality data that were not of the highest quality.

Deininger and Squire (1996) develop a more consistent dataset that meets certain minimum standards — the data must be based on household surveys, the sample must be representative of the population, and the measure of income (or expenditure) must be comprehensive. Using these improved measures of income inequality, Deininger and Squire (1998) confirm that the coefficient on inequality in their baseline growth regression is negative and significant. Moreover, they find evidence suggesting that much of the impact of inequality on growth appears to come through investment in human capital (schooling) rather than physical capital.

This cross-country evidence has led to the development of a number of theories outlining potential mechanisms by which such a relationship might arise. Before discussing these mechanisms, two caveats should be noted: first, the theoretical literature is concerned with labour productivity — measured real output per worker. In general this is not the same as the measures of real income per capita used in empirical studies. Second, a serious drawback of most of these empirical studies is their use of after-tax income inequality as a proxy for subsequent wealth inequality or human capital dispersion. Typically, income distributions are thought to be much less skewed than distributions of wealth or human capital. Moreover, if the relationships between them vary considerably across countries, it is difficult to know how to interpret results based on cross-country growth regressions.

The Incomplete Markets View

A fundamental (i.e., not induced) impact of inequality on productivity growth can only arise when there exists some kind of market failure. That is, for some reason a "commodity" is not being priced correctly by the market, so that the marginal benefit of the commodity to society does not reflect the marginal cost of providing it. In this section, I discuss the incomplete markets view, which emphasizes the role of endogenous borrowing constraints and externalities in generating the observed impact of inequality on growth.

Decreasing Returns and Credit Constraints

As Stiglitz (1969) first pointed out, when there are decreasing returns to capital and there is credit-rationing, individual wealth need not converge to a common level and the aggregate level of output may be affected by its distribution. Credit rationing arises when, at the going rate of interest, there exist individuals who could profitably invest borrowed funds and repay with interest, but lenders are unwilling to lend to them in full. When this particular market failure arises, it typically drives less wealthy, but potentially productive borrowers out of the loan market, leading to an inefficient allocation of resources, underinvestment and reduced productivity.

To understand this idea, suppose that the world consists of just two investors, Agatha and Bart, who have access to identical production opportunities illustrated by the stylized production function in Figure 1. This production function exhibits decreasing returns to wealth at high enough levels due to diminishing returns in production. Suppose Agatha has an initial wealth of 4 which could be invested in her own project or loaned. Bart has initial wealth of 2, but could borrow to finance his investment. If Agatha loaned one unit of wealth to Bart, this would raise Bart's output by more than it would reduce Agatha's ($BC>CD$). It follows that there are "gains from trade": Bart is willing to pay a rate of interest on the loan that exceeds the opportunity cost to Agatha of taking it out of production. If capital markets worked perfectly, Agatha would lend 1 unit of wealth and Bart would borrow 1 unit of wealth, and each would invest 3 units. Aggregate output would equal 2 x OC. If capital markets are imperfect, Agatha invests 4 units in production and Bart 2 units. Aggregate output in this case is $OB+OD$, which is less than 2 x OC. If Agatha and Bart had started out with the same initial wealth, the efficient level of output would have been attained with no trade. In the presence of market imperfections and diminishing returns to privately owned capital

FIGURE 1
Returns to Scale and Credit Rationing

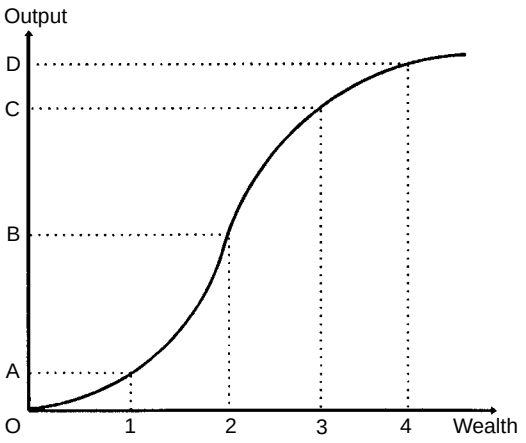

(physical or human), increased equality results in greater production efficiency.

Several recent theoretical analyses have built on Stiglitz's idea to study the interaction between economic development and the evolution of inequality in the presence of endogenous borrowing constraints and decreasing private returns to wealth. Galor and Zeira (1993) consider an overlapping generations models with intergenerational altruism in which individual lifetimes are divided into two periods: young and old. The young inherit heterogeneous levels of wealth and must decide whether to invest in a fixed and indivisible level of human capital. This determines their incomes when old and, hence, the size of the bequest they leave to their children.[7] Credit market imperfections result in the lending rate on capital being lower than the borrowing rate (due to monitoring costs). Thus, education is limited to those with sufficient wealth to purchase it outright or to pay a high rate of interest on loans. As a result of this mechanism, the initial distribution of wealth determines the aggregate amount of human capital investment and long-run per capita income. If initial inequality is sufficiently low, an egalitarian steady-state arises where all workers receive the same high-skilled wage and per capita income reaches a maximum. Otherwise, a low level one emerges where a fraction of the workforce earns disproportionately low wages. Banerjee and Newman (1993) also show the potential for long-run outcomes to be determined by initial levels of inequality, providing examples in which the economy either prospers or stagnates depending upon initial distributions of wealth.

In Aghion and Bolton (1997), agents either invest in a fixed-size, risky project, lending any remaining wealth or borrowing if necessary; or they earn a safe, low income, and lend. Limited liability and the dependence of the success probability on non-contractible effort induces credit-rationing based on inherited wealth. Equilibrium between borrowers and lenders determines a market interest rate that varies with the distribution of wealth. As wealth accumulates, demand for credit declines and supply rises, so that interest rates fall and, although it may initially rise, wealth inequality eventually falls. An important feature of their model is that the presence of idiosyncratic shocks to income implies that over time, even the wealthiest lineage could eventually become poor and the poorest lineage could become rich. This "ergodicity" implies that the initial distribution of wealth does not affect either the degree of inequality or per capita income in the long run. It follows that any positive effects of redistribution do not persist, so that perpetual redistribution is always necessary to achieve the maximum per capita income.

Lloyd-Ellis and Bernhardt (2000) develop a more general model than Banerjee and Newman in which individuals differ in their entrepreneurial efficiencies as well as their inherited wealth levels. They characterize the entire evolution of the distributions of wealth and income from a low-level state to an advanced economy and study its interaction with the development process. The impacts of wealth and ability are distinct and vary as the economy

develops. While in initial stages, wealth is the primary determinant of occupation because wealthy agents can invest in capital and profitably exploit cheap labour on a grander scale; in later stages, entrepreneurial efficiency matters more both because fewer agents are wealth constrained and because higher wages reduce the profitability of large-scale production. The consequence for the dynamics of income and wealth inequality is that they first rise and persist along family lineages, and then fall and are less persistent along lineages. That is, a Kuznet's curve arises endogenously and social mobility increases over time.

Although these analyses are suggestive of the impact of inequality on the speed of economic development, diminishing returns to human or physical capital imply that growth eventually peters out, so it is not possible to draw any implications for the impact of inequality on the long-run growth rate. To address this issue requires a framework in which growth is endogenously sustainable. Recently, several authors have studied the role of inequality in endogenous growth models in which human capital is the engine growth, and it is to these theories that I now turn.

Short-Run Human Capital Externalities, Inequality, and Sustained Growth

Although the importance of human capital in determining aggregate productivity has long been recognized (see, e.g., Shultz 1961), its role became part of mainstream macroeconomic thinking following the work of Lucas (1988). In his formulation, output is a function of physical capital and human capital devoted to production. Increments to human capital are a function of the fraction of the current stock of human capital devoted to learning (e.g., education and job training). The technologies for production and human capital accumulation both feature constant returns to scale. In the Lucas framework, productive knowledge is embodied in workers' skills that are, in turn, accumulated through *endogenous*, utility-maximizing investment decisions (schooling, training, and learning-by-doing) that sacrifice present consumption in order to raise future productivity and income. The assumption of constant returns at the aggregate level is crucial. If investments were subject to diminishing returns then sustained growth would be impossible.

The fact that education is almost always publicly financed, to a large degree, suggests that individual decisions to acquire human capital create external benefits for others. Moreover, in the absence of such externalities, it is difficult to reconcile observed pressures for migration from poor to rich countries with the absence of large capital flows in the other direction. Lucas (1988) therefore analyzes a version of his model in which the output of each firm depends on the human capital of its own workers as well as the average value of human capital per worker in the economy — *a short-run production externality*. With this technology, decentralized decision-making yields too little investment in human capital, as individual decisions to invest do not take into account the productivity gains from that investment that are realized by others. Steady-state output is too low relative to the social optimum, and growth is too slow.

The existence of a short-run externality in production, like that studied by Lucas (1988), opens the door for the dispersion of human capital to impact upon productivity and productivity growth. The existence of a positive externality coupled with the necessity for there to be constant returns to human capital at the aggregate level, implies that there must be diminishing returns to individual human capital. For heterogeneity in human capital not to affect the growth rate would require that embodied human capital is a perfectly tradeable input to production. While it is possible to think of quasi-examples of people trading pieces of their human capital (e.g., specialized consulting services), it is difficult to imagine someone selling analytical power to one firm and creative power to another at the same time in two different cities. In general, such markets are likely to be thin or non-existent.

Benabou (1996) illustrates how a static human capital externality in production can arise when

workers with different skills are complements in production and human capital is non-tradeable. The complementarity captures the idea that "poorly educated, insufficiently skilled production or clerical workers will drag down the productivity of engineers, managers, doctors and so on. Conversely, lagging advances in knowledge by scientists, engineers and other professionals will mean lagging wages for basic workers." In Benabou's production set up, this complementarity implies that an individual worker's wages depend positively on the current average level and equality of human capital in the economy, as well as his/her own schooling. Moreover, aggregate productivity is a decreasing function of the dispersion of human capital. Because household investments in human capital accumulation is subject to idiosyncratic shocks, inequality persists over time and creates a drag on long-run productivity growth.

An alternative micro-foundation to explain why a short-run externality from average human capital might arise is considered by Acemoglu (1996). Suppose there are two periods. In the first, firms make irreversible investments in physical capital and households invest in human capital. Workers and firms come together in the second period. The labour market is not competitive: instead, firms and workers are matched randomly. The only decision workers and firms make after matching is whether to produce together or not to produce at all. Firms base their investment decisions on the expected human capital of the workers they hire. Thus, although a worker's wages will depend on his/her own human capital, it will also depend positively on average human capital via the investment decision. With decreasing returns to individual human capital in production, returns also depend negatively on the variance of human capital. Thus, inequality again has a negative impact on aggregate productivity.

Socio-Political Mechanisms

Another strand of the positive theoretical literature on growth focuses on the relationship between inequality, the political process, and government policy. In this literature political outcomes determining government policy are endogenous to the distribution of wealth or income in the economy. Rational economic agents vote for or against tax policies which have redistributive consequences. Greater inequality (i.e., a poorer median voter) tends to result in higher equilibrium tax rates since a larger proportion of voters will favour redistributive policies. In Bertola (1991), Alesina and Rodrik (1994), and Persson and Tabellini (1992, 1994) redistributive policies of this sort reduce the private, after-tax marginal product of capital and create a disincentive to investment that leads to lower growth. These models thus predict that inequality and growth will be indirectly negatively related through the political process. Perotti (1993) examines similar political mechanisms except in a model where growth is driven by human capital accumulation and aggregate learning spillovers where redistribution is directly growth enhancing. Of particular interest is the possibility that both rich and poor would vote for redistributive policies (either in income or in terms of publicly-provided education) if the external benefits to all classes of having a better educated workforce are sufficiently large.

Although it accounts for the negative correlation between inequality and growth found by reduced-form equations, the political economy approach does not appear to be supported by the data. It implies that greater inequality increases the extent of redistribution, which in turn has a direct negative effect on economic growth. A corollary to this is that such a relation should be exclusive to democratic countries. However, Alesina and Rodrik (1994) and Clarke (1995) differentiate between democratic and non-democratic countries in examining the relationship between inequality and growth, and fail to find such evidence. Moreover, when measures of redistribution such as tax rates or the extent of social spending are regressed on measures of inequality, the coefficients are either insignificant or have an opposite sign to what the theory predicts (see Perotti 1996). It would seem that channels other than the political process must account for the influence of inequality on growth.[8]

There is also a literature that emphasizes the impacts of inequality of wealth and income on "disruptive" activities such as property crime, riots, and armed insurrection (e.g., Gupta 1990 and Benhabib and Rustichini 1996). It is argued that absolute and/or relative poverty may motivate people to participate in property crime and that this distorts the allocation of labour effort away from productive activities. Moreover, defensive efforts by potential victims represents a further loss of resources and threats to property rights deter investment. Through these various mechanisms, more inequality may tend to reduce the productivity of an economy. The negative effect of weak property rights on productivity and productivity growth is confirmed in the cross-country empirical work of Hall and Jones (1999) and Barro (1999). However, the empirical relationship between inequality and criminal behaviour and between crime and productivity is less clear.

Recent Challenges to the Incomplete Markets View

Evidence on Credit Constraints

Financial constraints have been found to play a crucial role in the entrepreneurial process in both developed and less-developed countries (LDCs). Evans and Jovanovic (1989) and Holtz-Eakin, Joulfaian and Rosen (1994), for example, find that capital is essential for starting a business in the United States and that borrowing constraints tend to exclude those with insufficient funds at their disposal. Similarly, Blanchflower and Oswald (1998) find that those individuals who inherit significant amounts of wealth are much more likely to start a business in the United Kingdom. Given that these two economies have the most developed capital markets, borrowing constraints are likely to be even more prevalent in other countries. In particular, Levy (1993) and Fidler and Webster (1996) find considerable evidence that entrepreneurs in LDCs are borrowing-constrained.

However, the evidence that short-run credit constraints are important for private human capital investment is far less clear, especially for developed economies. It is well-documented that children from low-income families complete fewer years of schooling than other individuals, despite high rates of return to schooling (see Jimenez 1986; Kane 1994). However, it is not clear that this is because they do not have sufficient access to credit to pay for higher education. For example, Cameron and Heckman (1998) find that after they account for other background characteristics and scores on ability tests, measured family income plays only a minor role in explaining schooling attainment in the US. Similarly, Shea (2000) finds that "exogenous'" changes in parents' income due to "luck" have a negligible impact on children's human capital.[9] Heckman and Klenow argue that "Long term factors, like ability, family structure, neighbourhood effects and the quality of the primary and secondary schools an individual attends may be more important than short term credit constraints in determining who goes to college" (1997, p. 23).

Rising tuition costs in Canada have generated concern that students may be deterred from entering postsecondary education. Finnie (2000) finds that student borrowing has increased since 1982, but only 10 to 15 percent of all postsecondary students in 1995 reported difficulties with the repayment of their student loans. However, the data do not "tell us how many worthy and interested students have chosen not to pursue (or continue) their postsecondary studies because they were unwilling to take on the required debt loads" (Finnie 2000, p. 9). Thus, although at present the evidence on the impact of short-term credit constraints on educational attainment in Canada is inconclusive, the US evidence provides little support. As I discuss later, this does not imply that family income and human capital is an unimportant determinant of a child's human capital.

Evidence on Short-Run Human Capital Externalities

A key implication of the link between inequality and productivity stemming from short-run production externalities is that, controlling for own human capital (i.e., schooling and experience), the productivity (i.e., the wage) of an individual worker is higher the higher the human capital of other workers in the economy. Is there any evidence to support this link?

One approach to measuring the social returns to schooling is to compare the aggregate output effect of schooling across countries with the individual micro returns. Early studies (e.g., Barro and Sala-i-Martin 1995) find that across countries, each additional year of average schooling is associated with about 30 percent higher gross domestic product (GDP) per capita. In contrast, across individuals within a country, each additional year of schooling is associated with roughly 7–10 percent higher wages (Psacharopoulos 1994). This large discrepancy between macro and micro rates of return has been interpreted by many as evidence that there exist huge positive externalities to schooling. However, as several recent studies (e.g., Bils and Klenow 1999; Heckman and Klenow 1997; Howitt 2000) point out, the macro estimates attribute too large an output effect to schooling. Causality may run from technology to income/life-span, or from future anticipated growth to schooling. When Heckman and Klenow (1997) include life expectancy in the macro regression to proxy for cross-country differences in technology, the average schooling coefficient falls to about 10 percent, which is in the same ballpark as the micro estimates. Similarly, Bils and Klenow (1999) argue that only a fraction of the correlation between schooling and per capita GDP growth, estimated by Barro and Sala-i-Martin (1995), comes from the impact of schooling on GDP growth. The remainder appears to be a result of reverse causation — the anticipation of future growth induces greater investment in human capital.[10]

Another approach is to include a measure of average schooling in a given area in regressions of individual's wages on their own schooling and other characteristics. Using data from US Standard Metropolitan Areas (SMSAs), Rauch (1993) finds that controlling for a worker's own education and experience levels, the worker's wages are higher the higher the average level of education in the worker's SMSA. Rauch finds that a worker's wages are 3.1 percent higher for each additional year of SMSA average education. But differences in average years of schooling across cities are also likely to be associated with differences in the relative supplies of skilled and unskilled labour which generate wage premia for average schooling even in the absence of externalities. When Ciccone, Perl and Almond (1999) build in these effects they cannot reject the hypothesis that no external effect is present.

Mare (1995), Peri (1998), and Moretti (1999) also estimate the effect of average schooling in US cities on individual wages, while Acemoglu and Angrist (2000) estimate the effect across US states. A key problem in estimating the returns to both own and average schooling is the potential endogeneity of schooling to wages. In order to avoid upward bias in their estimates, researchers must identify an exogenous source of variation in schooling (a "natural experiment") and use that to estimate the relevant coefficient. Moretti (1999) instruments for average schooling with changes in city age structure, tuition costs, and the presence of a land-grant college, but treats individual schooling as exogenous. He finds that the social returns to schooling exceed the private returns. Acemoglu and Angrist (2000) identify sources of exogenous variation for both own schooling (birth quarter) and average schooling (compulsory schooling laws). They find that workers' wages are typically less than 1 percent higher for each additional year of state average education.

On balance, recent evidence offers little support for sizeable social returns to education operating through a *short-run* externality in production. If there are close-to-constant returns to scale in production at the aggregate level, this implies that the

private returns to human capital cannot exhibit strongly decreasing returns, so that the impact of inequality in human capital is likely to be small. Note, however, that this evidence relates only to static externalities. As we discuss in the next section, a more important externality may be dynamic, operating through the education process itself or through the process of technology adoption.

Panel Regressions

Recently, the original cross-country evidence on the impact of inequality on growth has been challenged. Using panel data for a cross-section of countries at five-year intervals, Forbes (2000) finds a *positive* relationship between inequality (measured using Gini coefficients from Deininger and Squire 1996) and per capita income growth. Her regressions control for lagged per capita income, male and female education, market distortions (the price level of investment), and country-fixed effects. Her methodology and results have come under significant criticism (see Aghion, Caroli and García-Peñalosa 1999). In particular, because her results are based on fixed-effects estimates which have relatively few observations, they are particularly sensitive to measurement-error problems. However, also using panel data over ten-year intervals, Barro (1999) finds that the overall impact of inequality on growth over the subsequent decade to be weakly positive, once one controls for a broad set of other key determinants of growth.[11]

These results pose an important challenge to what had until recently become almost conventional wisdom. They suggest that perhaps an equity efficiency trade-off exists after all, or at least the basic incomplete markets plus diminishing private returns story is incomplete. Moreover, there are a number of theoretical mechanisms through which a positive relationship might manifest itself. For example, Greenwood and Jovanovic (1990) study the dynamic interaction between "financial superstructure," inequality, and economic growth. In their model, market imperfections arise because it is costly to engage in financial intermediation which, through risk-pooling, can allow entrepreneurs to earn a higher and safer return on their investments. If the higher returns available through financial intermediation justify the costs of forming such "syndicates," then these structures will tend to arise endogenously. The fixed cost associated with the formation of a financial intermediary structure ensures that the extent of financial intermediation, and thus the overall level of investment efficiency and growth, will be a function of the distribution of wealth. In the early stages of development, growth is slow, but as wealthy investors organize, their investments are made more efficiently and they become proportionally more wealthy. Thus, rising growth is associated with rising inequality.

Inequality and Growth in the Short and Long Run

In this section, I ask whether the recent evidence surveyed above is really much at odds with the incomplete markets view as it would first appear. In particular, I argue that, once we distinguish between short- and long-run relationships, the two bodies of evidence are quite consistent with each other, and with the associated theoretical literature which often predicts that the relationships between inequality and growth depend crucially on the time frame.

Short- versus Long-Term Credit Constraints

The absence of short-term borrowing constraints on investment in human capital does not imply that family income is unimportant for schooling attainment. Family income affects the kind of community that children grow up in, the schools they attend and important complementary inputs to human capital formation such as nutrition, location, books, family holidays, etc. All of these factors in turn affect how well they do early on in the schooling system (and hence their performance on tests), which determines the feasibility and optimality of continuing on to higher levels of education and training. Thus, estimating the marginal impact of parental income on schooling attainment while controlling for these other variables only picks up the short-run impacts.

In a recent paper, Acemoglu and Pischke (2000) attempt to sidestep these potential endogeneity problems by exploiting the fact that families at the bottom of the income distribution became much poorer and families at the top became much richer in the 1990s than they were in the 1970s. They argue that these movements are unlikely to be correlated with other characteristics affecting educational choices. Using this source of exogenous variation, they find that a 10 percent increase in family income is associated with a 1.4 percent increase in the probability of attending a four-year college. However, they are unable to determine whether the impact exhibits diminishing returns. Relatedly, Duflo (2000) finds a positive effect of family resources on child health in South Africa, using the expansion of old age pensions as an instrument. In a Canadian sample, Beach and Finnie (1988) find that much of the impact of family background variables on income appears to operate through education.

Thus, while short-term credit constraints on attending college or university may not play a pivotal role in human capital investment, there may be a crucial *long-term credit constraint*: parents provide key complementary inputs to their child's human capital, but poor parents cannot borrow against their children's anticipated earnings in order to finance them. For example, it seems unlikely that a bank would be willing to lend to a low-income parent enough money for them buy a house in a well-off neighbourhood with high quality schools and positive social interactions, using the increase in their child's expected earnings as collateral.

Dynamic Externalities

Tamura (1991) considers an alternative externality to that considered by Lucas (1988) whereby individuals *learn more*, the higher the human capital of others in the (local) economy. That is, he supposes that an aggregate spillover arises in the human capital accumulation process rather than production. This *dynamic externality* implies that controlling for own schooling, the greater is the average stock of human capital in the economy the larger will be the *increments* to an individual's productivity. If there are decreasing returns to individual human capital and human capital is non-tradeable, the inequality has a negative effect on the average growth of human capital.

Although Tamura's model is suggestive, it is rather stylized and does not provide an explicit account of how or why average human capital impacts upon the individual learning process. Recently, however, several authors have developed alternative theories which explicitly model the way in which such dynamic externalities operate and which provide specific implications for policy. In particular, dynamic externalities may arise through the public education system and through the innovation process.

Dynamic Externalities in Education

In an overlapping generations model, Glomm and Ravikumar (1992) illustrate the role of formal education in determining the interactions between growth and inequality.[12] In their framework, parents care about the quality of their children's education which, in conjunction with their own level of human capital, determines the human-capital acquisition by their child. Under a public-education regime all students receive an equal quality of education determined by average income via taxation. Although the private marginal returns to parental human capital are diminishing the overall returns to public and private investment are constant, so that growth is sustained. Greater initial inequality leads to lower productivity growth. Moreover, with diminishing private returns (and public education) inequality declines over time and the growth rate rises.

Rather than thinking of dynamic human capital externalities as being economywide, it may be more useful to think of them as being local in nature; for example, if the financing of public primary and secondary schools has a large local component and it will be a function of community income.[13] If households cannot borrow against the future earnings of

their children to finance current expenditures, inefficiently low investment in human capital among children from poor neighbourhoods may result. Alternatively, such externalities may be social, arising from the interaction of the aspirations and performance of students in affluent and less-affluent neighbourhoods. In either case, there is a tendency for stratification to imply self-replicating neighbourhoods and the segregation of the poor (see Durlauf 1994). Benabou (1996) analyzes the effect of schooling on growth when students of heterogenous abilities can either be segregated or mixed together.[14] In the short run, segregation may increase growth because talented people are complements in producing new human capital. In the long run, however, segregation leaves intact the overall heterogeneity of skills in the economy, perpetuating inequality in the long run, and creating a drag on productivity growth.

In these models, the human capital acquired by each generation depends on parental inputs (income and/or human capital) and public education expenditures financed by taxing the previous generation. As a result, average investments in human capital (and therefore wages) are increasing in the mean and equality of the human capital of the *previous* generation. Williamson (1993) and Lloyd-Ellis (2000) provide cross-country evidence suggesting that greater parental inequality is indeed associated with lower secondary school enrolment (controlling for per capita income). Another piece of supportive evidence is that the quality of schools attended by workers seems to matter for their wages later in life and that the quality of schools is a function of local human capital. For example, Card and Krueger (1992) find that American men who were educated in states with higher quality schools (measured by pupil-teacher ratios, average term length, and relative teacher pay) have a higher rate of return to additional years of schooling. Rates of return are also higher for individuals from states with better-educated teachers. Indeed, Hanushek finds that "the difference in student performance in a single academic year from having a good as opposed to a bad teacher can be more than one full year of standardized achievement" (1992, p. 84). While it does not prove their existence, this empirical evidence is consistent with the potential importance of dynamic externalities operating through education.

Dynamic Externalities in Technology Adoption

Romer (1990) develops a paradigm of endogenously sustainable growth based on the accumulation of *disembodied knowledge* in the form of new technologies and ideas.[15] In his model, as in Lucas (1988), output is a function of physical capital and the portion of human capital devoted to production. However, physical capital is made up of heterogeneous intermediate inputs, whose overall productivity is a function of the state of applied knowledge. Increments to knowledge are a function of past knowledge and the labour effort devoted to research and development. The incentives to invest in research and development (R&D) come from the monopoly profits from new innovations which are protected, at least temporarily, by patents. Although the technologies themselves cannot be used by others, the knowledge generated by R&D is "non-rival" and feeds into future innovations, thereby making growth endogenously sustainable. Thus, in contrast to the Lucas model, positive *dynamic* externalities are an inherent part of a growth process driven by disembodied knowledge accumulation. Because investors do not take into account the positive effects of their R&D on future innovations, economic growth is inefficiently low and there is a role for government intervention in the process (e.g., by subsidizing R&D or strengthening intellectual property rights).

If innovation responds endogenously to incentives and different technologies have different skill requirements, the nature or the rate of innovation may be affected by the distribution of skills in the economy. For example, Acemoglu (1998) argues that if the *nature* of innovations depends upon the distribution of skills, changes in the distribution of skills may have effects on the skill-bias of new technologies. He supposes that researchers target their

effort to innovations that complement either skilled or unskilled labour. Since research is a fixed cost, the returns to R&D depend on the number of workers that will be able to use the new technology. He argues that the expansion of education since the 1960s made it profitable to invent machinery to be used by skilled rather than unskilled workers, so that technical change *became* skill-biased.

An alternative hypothesis that explicitly links the dispersion of skills to the *rate* of productivity growth is explored by Lloyd-Ellis (1999). In that model minimum skill levels are required to implement new ideas and technologies. Workers are distinguished by the range of ideas and technologies that they are capable of implementing and it takes time to acquire the necessary skills. The existence of endogenous skill-biased technological change provides another mechanism through which the distribution of skills may impact upon productivity growth. If appropriately skilled workers are scarce relative to less-skilled workers, this may result in an allocation of resources toward current production and away from the introduction and implementation of new technologies, to the detriment of long-run productivity growth. However, for such a situation to persist for so long implies that there must be some reason why individuals are unable or are not choosing to acquire the necessary skills, despite the apparent high returns to doing so. There are several possible explanations: public institutions are not offering an appropriate match between skills and technology, there are borrowing constraints or there is under-investment due to some kind of externality.

Interpreting the Growth-Inequality Regressions

The panel regression results of Forbes (2000) and Barro (1999) do not directly contradict those of the earlier work by Persson and Tabellini (1994), Alesina and Rodrick (1994), Perotti (1996) and Deininger and Squire (1998). Forbes' (2000) main estimates are based on panel regressions over five-year intervals with country-fixed effects. It follows that the results she obtains can be viewed as an estimate of the average *short-run* relationship between growth and inequality *within* countries. In contrast, the earlier negative estimates relate to the *long-run* impact of initial inequality on growth over the subsequent 25–30 years *across* countries.[16] When Forbes estimates the regressions over ten-year intervals, the relationship becomes insignificant, a fact that is also consistent with Barro's results. Unfortunately, due to the limitations on the time span of the Deininger and Squire (Gini coefficient) data that they use, it is not possible to obtain within-country estimates over longer time intervals.

These results are therefore not inconsistent with the view that the positive short-run relationship could become reversed over significantly longer periods. Indeed as the discussion above suggests there are good reasons to believe that any negative effects of inequality should be stronger over longer periods of time. For example, if higher levels of inequality limit investment in education, the negative impact would be greater in the long term. Although it is not possible to test this hypothesis using the broad and (relatively) high quality datasets of Forbes and Barro, there is one analysis that provides some evidence in this regard. In fact, one of the very first studies in this literature, that of Persson and Tabellini (1994), estimates the within-country relationship between initially inequality and subsequent growth over 20-year intervals from 1830 to 1985 for nine countries (Austria, Denmark, Finland, Germany, the Netherlands, Sweden, the United Kingdom and the United States). They measure inequality using the after-tax income share of the top 20 percent of the population. They find a significant negative impact in regressions with fixed country effects, which control for several key variables (e.g., initial income and schooling).[17]

A reasonable (though not unique) interpretation of these regression results is that the long-run impact of inequality on growth is negative, but the short-run, induced relationship between the two variables is positive. Given that the microeconomic evidence discussed above suggests that the short-

run negative impacts of inequality are likely to be small, factors which induce a positive short-run relationship may very well dominate. Indeed there are several forces that provide reasonable candidates, but which have not been controlled for in the panel regressions discussed above.

Time Variation in Policies

If the "equity-efficiency trade-off" dominates in the short run, changes in the progressivity of the tax system should be expected to induce a positive relationship. Greater progressivity would lower both inequality and growth in the short run, while a movement toward less progressivity might be expected to raise both variables. Similarly, increased taxation that is used to raise the quality of public education might have disincentive effects that reduce productivity in the short run but raise long-run productivity growth. Even in an incomplete markets model, negative incentive effects may dominate in the short run. One problem with the credit constraint models of Galor and Zeira (1993), Aghion and Bolton (1997), etc. is that they assume that individuals simply inherit their wealth so that, for example, redistribution does not affect incentives. Ghatak, Morelli and Sjöstrom (1997) develop an alternative model in which individual's must first work and save before they can acquire enough wealth to invest in their own enterprise. Their wages depend on their own effort and the existence of a credit constraint generates incentives to work harder (they call this the "American Dream" effect). In their model, redistribution (via income taxation, say) can destroy these incentives and will tend to offset the productivity gains from redistribution arising from decreasing returns to wealth in production. Hence, a trade-off arises between the short-run costs and long-run gains of redistribution.

Skill-Biased Technological Change

Alternative hypotheses for the rise in returns to skill in some OECD countries since the mid-1970s are that it stems from changes in government policies, increased international competition (Wood 1994), or skill-biased technological change (SBTC). Murphy and Welch (1993) argue that the rise in returns to skill appears to have progressed in a smooth fashion for a relatively long time in the US and thus it is unlikely that they are a short-run consequence of governmental policies. A consensus is also emerging that trade effects can at best explain a small portion of the changes (see Berman, Bound and Machin 1998; Wood 1994).[18] Thus, the leading hypothesis appears to be that the rise in skill-premia is associated with the skill bias of newly introduced technologies.[19] This hypothesis finds support in the work of several authors including Berman, Bound and Griliches (1994) and Autor, Krueger and Katz (1998) for the US, and Machin and Van Reenen (1998) for other OECD countries. Berman, Bound and Machin (1998, 2000) go further in suggesting that the effects of skill-biased technological change on the relative wages of skilled and unskilled workers can be seen throughout the world in both developed and less-developed countries.

It should be recognized, however, that while SBTC is the leading hypothesis for rising inequality in the US, it is by no means incontrovertible. The econometric evidence on the impact of technological change on relative wages is indirect, and subject to the usual criticisms, most importantly omitted variable biases. Mishel (this volume) argues that the SBTC hypothesis is inconsistent with various aspects of the rise in US wage inequality and that other factors may have been more important. Moreover, even if the hypothesis were correct for the US, it is not clear that it provides a good explanation for events in other countries. Beaudry and Green (1997), for example, find no evidence of increased within-cohort dispersion of earnings in Canada. Rather, they document that much of the increase in Canadian wage dispersion is mainly due to a deterioration in the wages of more recent cohorts. Nevertheless, to the extent that technological change contributes both to increased growth and to rising skill premia, it is possible that a short-run positive relationship between inequality and growth would emerge, even if the long-run effects of inequality are negative.

Changing Age Structure

Measures of inequality such as the Gini coefficient do not allow for a decomposition of incomes by age. If there are returns to experience say, then an increase in the proportion of new entrants into the workforce might lead to an increase in inequality by boosting the fraction of the population on low incomes and possibly an increase in per capita income growth. The key difference in Barro's (1999) panel regression results relative to those of previous authors who find a significant negative impact, appears to be the inclusion of a fertility-rate variable. Once Barro drops this variable from his regressions, he gets similar results to Perotti (1996). Thus, the interpretation of these results depends on one's view of why fertility and income inequality are contemporaneously correlated. Foot and Gomez (this volume) find a similar result using the dependency ratio. At ten years, including this ratio causes the relationship to become insignificant. However, in estimating the impact of initial inequality on growth over the subsequent 20 years they find that the significant negative effect of inequality is robust to the inclusion of the dependency ratio.

Implications for Human Capital Policy

Assuming that the removal of the underlying market failure is not feasible, then long-run productivity growth may be enhanced by policies that ultimately reduce inequality in the distribution of wealth or human capital. However, the short-run impact of inequality on productivity in Canada is likely to be small, so that any negative incentive effects of redistributive policies (e.g., high taxation to pay for greater investment in public education) will tend to dominate. In contrast, the long-run impact of inequality on productivity growth is likely to be negative and more substantial. This perspective suggests the "hump-shaped" relationship between the progressivity of policy and growth illustrated in Figure 2. At low levels of progressivity (and, hence, high inequality), the long-term efficiency gains from greater progressivity, emphasized by the "incomplete markets" view, tend to dominate, so that greater progressivity enhances growth. However, at sufficiently high levels of progressivity (and, hence, low inequality), the short-run (but continual) efficiency losses, emphasized by the "equity-efficiency trade-off" view, start to dominate so that greater progressivity inhibits growth. In this section I outline some of the policy implications of this perspective.

Balancing the Short- and Long-Run Impacts of Public Education

Sufficient direct public spending on education ensures a minimum standard of quality for all and makes it possible for students from disadvantaged backgrounds to advance through the system. If there are decreasing returns to private parental inputs into human capital, transferring resources from private individuals into public education will reduce inequality and raise average investment in human capital in the long run. This in turn will enhance long-run productivity gains as these students enter the workforce. However, excessively high taxation of parents to finance this spending may also have negative incentive effects on work effort and investment in the short run. Ultimately, an optimal government policy must achieve a balance between these effects.

Figure 2
Impact of Progressivity on Growth

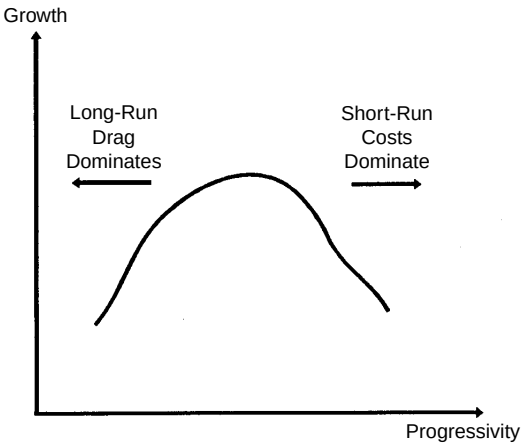

Public versus Private Funding

Glomm and Ravikumar (1992) compare the economy's growth path and the evolution of the distribution of income under private and public mechanisms for the provision of education. With private provision, the quality of a child's education is determined by his/her parent's income. As a result the child's human capital is a function only of his/her parent's human capital, so that inequality persists. In contrast, under a public education regime all students receive an equal quality of education determined by average income. Since the private marginal returns to parental human capital are diminishing, inequality declines relatively rapidly over time. Since, under private provision, all returns are appropriated by the individual, more effort is allocated to schooling than under public provision, so that the growth rate tends to be higher, for a given degree of inequality. However, if initial inequality is sufficiently high, growth under public provision of education will eventually surpass that under private provision, because of the faster reduction in inequality. If one introduces *idiosyncratic shocks* to income (due to differences in innate ability, say) into the Glomm and Ravikumar model, it turns out that public provision *always* leads to higher long-run growth, because inequality persists.

One implication of this kind of model is that economies with more emphasis on public education systems should exhibit more social mobility than those with less. However, when comparing Italy and the US, Checchi, Ichino and Rustichini (1999) find that Italy, while displaying less income inequality, also displays less intergenerational upward mobility between occupations and between education levels. They argue that when family background is important for labour market success, an overly centralized and egalitarian *tertiary* education system need not help poor children and may take away from them a fundamental tool to prove their talent and to compete with rich children. In contrast, at the primary and secondary levels the quality of schools matters more than individual effort, so that a state-run school system generates larger human capital investment by poor families.

The Distribution of Quality

Parental incomes and/or parental human capital are complementary to public expenditures in the production of children's human capital. Although the evidence suggests that short-term borrowing constraints may not be that important for investment in higher education, lack of parental resources are crucially important in the primary and secondary stages of education. This in turn makes it more difficult for children from poorer backgrounds to get to more advanced stages of education and/or training. In effect, this reflects a *long-term borrowing constraint*: it is not possible for parents to borrow against their children's anticipated earnings.

Benabou (1996) shows that if schools are financed locally, in communities that are sorted on talent or resources, then expenditures on education will tend to perpetuate inequality and reduce long-run growth. Greater funding equality (through centralized taxation) and reduced segregation on talent leads to lower long-run inequality and higher growth. In this model, centralized financing and a national curriculum may provide a long-run advantage relative to a decentralized system. Over the last two decades several US states introduced legislation which effectively increased their role in the provision of education relative to local districts, effectively equalizing expenditures per student across districts. Murray, Evans and Schwab (1998) find that court-ordered finance reforms between 1971 and 1996 reduced within-state inequality in spending by 19 to 34 percent. Fernandez and Rogerson (1998) examine the consequences of such reforms in the context of a dynamic equilibrium model of public education provision, calibrated using US data. They find that the policy increases both average income and the share of income spent on education, and significantly increases welfare.

Thus the distribution in the quality of education can have profound implications for both the rate of accumulation of human capital and the persistence of inequalities across generations, both of which have implications for productivity growth. Egalitar-

ian systems of primary and secondary education are crucial for raising the aggregate efficiency of human capital investments. In particular, the quality of schools should not reflect the average incomes and social conditions of the local community but, if anything, should compensate for them.

Relatedly, given that the quality of schools largely reflects the quality of their teachers, the incentives faced by teachers should not vary across localities. Recent policy discussions have suggested that teacher's pay should be linked to how well their students do on standardized tests. If the social backgrounds of students did not vary across schools this could provide good incentives and attract the best teachers. However, it is clear that average social backgrounds do vary considerably across schools, due to average incomes of local communities, the concentration of recent immigrants in particular areas, etc. In this context, linking pay to outcomes alone could have adverse implications for the persistence of inequality and long-term productivity. An effective compensation scheme must take into account the correlation between students' backgrounds and their academic outcomes.[20]

Public Funding at Different Stages of Human Capital Acquisition

Given a fixed total budget, what factors determine the optimal allocation of expenditures between primary, secondary, and tertiary education? As usual the literature implies a trade-off. In Lloyd-Ellis (2000), for example, basic education is compulsory, but increments to higher education are voluntary. Parental incomes affect their children's human capital accumulation (e.g., nutrition, books, family holidays, etc.), and there are decreasing returns to this input to human capital. Since the acquisition of higher education is costly in terms of time and effort, students who anticipate employment in low-skilled occupations have little incentive to acquire education beyond basic levels. As a result, the distribution of income among parents and the relative quality of different education levels affects the incentives of students to acquire higher education.

In this context, greater expenditures on higher education increase the human capital of those who make it that far through the educational system. This can have important "trickle-down effects" by making them more effective managers, engineers, etc. and perhaps most importantly, creative innovators, which will benefit society in general. On the other hand, if such improvements come at the expense of primary and secondary education, it may reduce the incentives for students from disadvantaged backgrounds to get through the system. Thus, higher quality university education may be concentrated amongst fewer students. The resulting decrease in enrolments and long-term effects of the (persistent) increase in inequality on average human capital accumulation could more than offset the gains.

The Distribution of Skills and Innovation

Although there are important exceptions, there is growing evidence that many recent technological innovations are skill-biased. To the extent that the rate of innovation responds endogenously to profit incentives this implies that the distribution of skills (as well as its average level) is an important determinant of R&D investment and other forms of innovation. Given that the innovation or adoption of new ideas and technologies often requires skilled workers, it is crucial for universities, colleges and other institutions of higher and technical learning to be adequately responsive to the technological frontier. Rosenberg (2000), for example, argues that a crucial factor in postwar US growth has been the responsiveness of the higher education system to the needs of industry. However, focusing resources on this stage of the educational process at the expense of others need not be the appropriate policy response. Although the innovative process requires skilled labour, implementation of new innovations in production also requires sufficient skills on the part of production and managerial workers. While raising the human capital of those at the top of the skill distribution may reduce the unit costs of innovation, doing so at the expense of those lower down the distribution reduces the ultimate profitability of implementing new technologies. Thus, the maintenance of incentive to

innovate requires a balance between the skills of those involved in the innovation or initial adoption of new technologies and those who ultimately implement it.[21]

On-the-Job Training: Matching Skills with Technology

Many argue that in the current climate of rapid skill-biased technological change, training workers on the job to be able to adapt new technologies should be a primary objective for public policy. In the short run, it may be true that well-targeted expenditures may be effective in this regard. It is likely that the optimal allocation of resources would be skewed toward lower income workers, thereby compensating for previous disadvantages in the acquisition of skills via the public education system. In the long run, however, a better allocation might be to direct public resources at raising the quality of primary and secondary education. This would enhance students learning abilities and creativity from the beginning, allowing them to adapt to new ideas and technologies more easily. Moreover, the evidence discussed by Heckman and Klenow (1997), for example, suggests that it is preferable to leave on-the-job training to private firms that can ensure a better match between their own needs and the skills learned by workers. Unless there is some evidence of a knowledge externality between firms, it is not clear why such training should be subsidized by the public sector.

Concluding Remarks

It should be recognized that both the microeconomic and macroeconomic evidence that I have discussed here is far from conclusive and is subject to a number of criticisms, some of which I have noted. Considerable research needs to be done in order to substantiate or reject some of the claims that are made in the literature. This is especially true outside the US, where the relative dearth of longitudinal datasets has made it difficult to assess many of the microeconomic hypotheses that labour economists would like to. While I have tried to demonstrate the direction in which, on balance, the currently available evidence appears to be pointing, the conclusions I reach ultimately involve a substantial degree of speculation on my part.

Nevertheless, I believe that this paper has made three contributions. First, I assess whether recent empirical evidence on the relationship between inequality and growth is really as contradictory to the incomplete markets view as it first appears. I argue that it is not, but that it does illustrate the importance of distinguishing between the long-run, causal effects of persistent inequality on growth, and the short-run, induced relationship between them. Second, I argue that the potential implied trade-off between the short-run negative effects and long-run positive effects of progressive policies on growth is completely consistent with much of the recent theoretical literature. Finally, I illustrate several specific implications for human capital policy in Canada that may be drawn from this perspective.

Notes

This survey was prepared for the CSLS-IRPP conference on the Linkages between Inequality and Growth. I am grateful to Benoit Delage, Curtis Eaton, David Love, Richard Roy and Andrew Sharpe for useful comments and suggestions. Funding from the Social Sciences and Humanities Research Council of Canada and the Canadian Institute for Advanced Research is gratefully acknowledged.

[1]Organizational economics is often referred to as "new institutional economics" by development economists (see Hoff, Braverman and Stiglitz 1993).

[2]For an excellent introduction to endogenous growth theory, see Aghion and Howitt (1998).

[3]There are, however, serious data constraints. For example, there is a complete absence in Canada of data that would allow for a convincing estimation of the impact of parental income on children's educational outcomes.

[4]The existence of a market for *all* commodities is a rather strong requirement. This is especially true when

we think about the market for human capital. For example, one implication of such "complete markets" would be that poor parents could finance a mortgage in a rich neighbourhood with a high quality local school and low crime, etc., by borrowing against the resulting expected increase in their child's expected future earnings.

[5]Here and throughout the paper, I define labour productivity as output per worker, not output per hour worked.

[6]See Osberg (1995) for a discussion. This trade-off presumes an attempt to generate lower inequality relative to the *laissez-faire* efficient economy. It is worth noting that distortionary policies which *increase* income inequality relative to the *laissez-faire* economy would also reduce productivity.

[7]To the extent that the preferences of parents and their children differ, this may actually be a more sensible representation.

[8]Note, however, that Alesina and Perotti (1994) provide empirical support for the more general hypothesis that high inequality lowers growth because of the political instability it causes.

[9]He defines "luck" as income variation due to union status, industry wage premia and job loss due to company closure or relocation. It is not clear, however, how correlated these variables are with parental attitudes toward education.

[10]This is because the benefit to having human capital is proportional to aggregate productivity while working, whereas the opportunity cost is proportional to aggregate productivity while in school.

[11]If anything, he finds a negative relationship for poorer countries and a positive relationship for richer ones.

[12]Given the long-run nature of dynamic externalities operating through the education system, overlapping generations models seem like a reasonable framework within which to think about these issues.

[13]A large portion of the funding of Canadian schools is at the provincial level, but there are still significant disparities within provinces. Moreover, the local environment is still an important factor in determining schooling outcomes.

[14]Benabou's model actually features *both* static and dynamic externalities. In his model, the impact of human capital inequality arises through a static externality due to the complementarity in aggregate productivity. However, this need not be so, one can also introduce it in the human capital accumulation process and obtain the same qualitative results (see Love and Lloyd-Ellis 1997).

[15]Grossman and Helpman (1991) and Aghion and Howitt (1992) develop somewhat different models of growth through disembodied technological change.

[16]These results are confirmed in Forbes' dataset.

[17]Although they are able to partially rule out problems associated with reverse causation and measurement error, the possibility of omitted variable bias remains.

[18]The case against the importance of international trade is really a case against the importance of trade in final goods. It is possible that changes in the trade of intermediate goods associated with outsourcing is a major source of rising skill premia (see Aghion, Caroli and Garcia-Penalosa 1999).

[19]Beaudry and Green (1997) find no evidence of increased within-cohort dispersion of earnings in Canada. Rather they document that much of the increase in Canadian wage dispersion is due to a deterioration in the wages of more recent cohorts.

[20]Of course, there are many other potential problems in devising such a compensation scheme, not least of which would be how to measure student outcomes in the first place.

[21]This effect is in addition to that described above, where skewing resources toward higher levels of education to raise the skills of those who reach that stage, make it more difficult for others to get this far.

References

Acemoglu, D. 1996. "A Microfoundation for Social Increasing Returns in Human Capital Accumulation," *Quarterly Journal of Economics* 111(3):779-804.

_____ "Why do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics* 113(4):1055-89.

Acemoglu, D. and J. Angrist. 2000. "How Large are the Social Returns to Education? Evidence from Compulsory Schooling Laws." Unpublished manuscript.

Acemoglu, D. and J.S. Pischke. 2000. "Changes in the Wage Structure, Family Income and Children's Education," NBER Working Paper No. 7986. Cambridge, MA: National Bureau of Economic Research.

Aghion, P. and P. Bolton. 1997. "A Trickle-Down Theory of Growth and Development with Debt Overhang," *Review of Economic Studies* 64(2):151-62.

Aghion, P. and P. Howitt. 1992. "A Model of Growth through Creative Destruction," *Econometrica* 60:323-51.

_____ *Endogenous Growth Theory.* Cambridge, MA: MIT Press.

Aghion, P., E. Caroli and C. García-Peñalosa. 1999. "Inequality and Economic Growth: The Perspective of the New Growth Theories," *Journal of Economic Literature* 37(4):1615-60.

Alesina, A. and D. Rodrik. 1994. "Distributive Politics and Economic Growth," *Quarterly Journal of Economics* 109(2):465-90.

Alesina, A. and R. Perotti. 1994. "The Political Economy of Growth: A Critical Survey of the Recent Literature," *The World Bank Economic Review* 8(3):351-72.

Autor, D., A. Krueger and L. Katz. 1998. "Computing Inequality: Have Computers Changed the Labour Market?" *Quarterly Journal of Economics* 113(4): 1169-1213.

Banerjee, A. and A.F. Newman. 1993. "Occupational Choice and the Process of Development," *Journal of Political Economy* 101(2):274-98.

Barro, R. 1999. "Inequality, Growth and Investment," NBER Working Paper No. 7038. Cambridge, MA: National Bureau of Economic Research.

Barro, R. and X. Sala-i-Martin. 1995. *Economic Growth*. New York: McGraw Hill.

Beach, C.M. and R.E. Finnie. 1988. "Family Background in an Extended Earnings-Generation Model: Further Evidence," *Eastern Economic Journal* 14(1):39-49.

Beaudry, P. and D. Green. 1997. "Cohort Patterns in Canadian Earnings: Assessing the Role of Skill Premia in Inequality Trends," Discussion Paper No. 9703. Vancouver: University of British Columbia.

Benabou, R. 1996. "Heterogeneity, Stratification and Growth: Macroeconomic Implications of Community Structure and School Finance," *American Economic Review* 86:584-609.

Benhabib, J. and A. Rustichini. 1996. "Social Conflict and Growth," *Journal of Economic Growth* 1(1):129-46.

Berman, E. and S. Machin. 2000. "Skill-Biased Technology Transfer: Evidence of Factor Biased Technological Change in Developing Countries." Mimeo.

Berman, E., J. Bound and S. Machin. 1998. "Implications of Skill-Biased Technological Change: International Evidence," *Quarterly Journal of Economics* 113(4):1245-80.

Berman E., J. Bound and Z. Griliches. 1994. "Changes in Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures," *Quarterly Journal of Economics* 109(2):367-97.

Bertola, G. 1991. "Market Structure and Income Distribution in Endogenous Growth Models," NBER Working Paper No. 3851. Cambridge, MA: National Bureau of Economic Research.

Bils, M. and P. Klenow. 2000. "Does Schooling Cause Growth or the Other Way Around?" *American Economic Review* 90(5):1160-83.

Blanchflower, D.G. and A.J. Oswald. 1998. "What Makes an Entrepreneur?" *Journal of Labor Economics* 16(1):26-60.

Cameron, S. and J. Heckman. 1998. "Life Cycle Schooling and Dynamic Selection Bias: Models and Evidence for Five Cohorts of American Males," *Journal of Political Economy* 106(2):262-333.

Card, D. and A.B. Krueger. 1992. "Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States," *Journal of Political Economy* 100(1):1-40.

Checchi, D., A. Ichino and A Rustichini. 1999. "More Equal but Less Mobile? Education Financing and Intergenerational Mobility in Italy and in the US," *Journal of Public Economics* 74(3):351-93.

Ciccone, A., G. Peri and D. Almond. 1999. "Capital, Wages and Growth: Theory and Evidence," Discussion Paper No. 2199. Paris: Centre for Economic Policy Research.

Clarke, G.R.G. 1995. "More Evidence on Income Distribution and Growth," *Journal of Development Economics* 47(2):403-27.

Deininger, K. and L. Squire. 1996. "A New Data Set Measuring Income Inequality," *World Bank Economic Review* 10(3):565-91.

_____ 1998. "New Ways of Looking at Old Issues: Inequality and Growth," *Journal of Development Economics* 57(2):259-87.

Duflo, E. 2000. "Grandmother and Granddaughters: The Effects of Old Age Pension on Child Health in South Africa," Working Paper No. 00/05. Massachusetts Institute of Technology, Department of Economics.

Durlauf, S.N. 1994. "Spillovers, Stratification and Inequality," *European Economic Review* 38(April): 836-45.

Evans, D.S. and B. Jovanovic. 1989. "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints," *Journal of Political Economy* 97(4):808-27.

Fernandez, R. and R. Rogerson. 1998. "Public Education and Income Distribution: A Dynamic Quantitative Evaluation of Education-Finance Reform," *American Economic Review* 88(4):813-33.

Fidler, L. and P. Webster, eds. 1996. *The Informal Sector and Microfinance Institutions in West Africa*, World Bank Regional and Sectoral Studies. New York: World Bank.

Forbes, K. 2000. "A Reassessment of the Relationship Between Inequality and Growth," *American Economic Review* 90(4):868-87.

Finnie, R. 2000. "Student Loans: Empirical Evidence and Policy Implications," Working Paper No. 7. Kingston: School of Policy Studies, Queen's University.

Foster, A.D. and M.R. Rosenzweig. 1996. "Technical Change and Human Capital Returns and Investment: Evidence from the Green Revolution," *American Economic Review* 86:931-53.

Galor, O. and J. Zeira. 1993. "Income Distribution and Macroeconomics," *Review of Economic Studies* 60(1):35-52.

Ghatak, M., M. Morelli and T. Sjostrom. 1997. "Dynamic Incentives, Occupational Choice and the American Dream," CORE Discussion Paper No. 9728.

Glomm, G. and B. Ravikumar. 1992. "Public vs. Private Investment in Human Capital: Endogenous Growth and Income Inequality," *Journal of Political Economy* 100:818-34.

Greenwood, J. and B. Jovanovic. 1990. "Financial Development, Growth and the Distribution of Income," *Journal of Political Economy* 98(5):1076-1107.

Grossman, G. and E. Helpman. 1991. *Innovation and Growth in the Global Economy.* Cambridge, MA: MIT Press.

Gupta, D. 1990. *The Economics of Political Violence.* New York: Praeger.

Hall, R.E. and C.I. Jones. 1999. "Why Do Some Countries Produce so much more Output per Worker than Others?" *Quarterly Journal of Economics* 114(February):83-116.

Hanushek, E.A. 1992. "The Trade-Off between Child Quantity and Quality," *Journal of Political Economy* 100(1)84-117.

Heckman, J. and P. Klenow. 1997. "Human Capital Policy," Working Paper. Chicago: University of Chicago.

Hoff, K., A. Braverman and J. Stiglitz. 1993. *The Economics of Rural Organization: Theory, Practice and Policy.* New York: World Bank Books, Oxford University Press.

Holtz-Eakin, D., D. Joulfaian and H. Rosen. 1994. "Sticking it Out: Entrepreneurial Survival and Liquidity Constraints," *Journal of Political Economy* 102(1):53-75.

Howitt, P. 2000. "Endogenous Growth and Cross-Country Income Differences," *American Economic Review* 90:829-46.

Jimenez, E. 1996. "The Public Subsidization of Education and Health in Developing Countries: A Review of Equity and Efficiency," *Research Observer* 1:141-59.

Juhn, C., K.M. Murphy and B. Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill," *Journal of Political Economy* 101:410-22.

Kane, T. 1994. "College Entry by Blacks Since 1970: The Role of College Costs, Family Background, and the Returns to Education," *Journal of Political Economy* 102(5):878-911.

Levy, B. 1993. "Obstacles to Developing Indigenous Small and Medium Enterprises: An Empirical Assessment," *The World Bank Economic Review* 7(1):65-83.

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

Lloyd-Ellis, H. 1999. "Endogenous Technological Change and Wage Inequality," *American Economic Review* 89(1):47-77.

_____ 2000. "Public Education, Occupational Choice and the Growth-Inequality Relationship," *International Economic Review* 41(1):171-201.

Lloyd-Ellis, H. and D. Bernhardt. 2000. "Enterprise, Inequality and Economic Development," *Review of Economic Studies* 67(January):147-68.

Love, D. and H. Lloyd-Ellis. 1997. "Inequality, Growth and the Role of Government" in *Equality and Prosperity: Finding Common Ground*, ed. W. Robson and W. Scarth. Policy Study No. 30. Toronto: C.D. Howe Institute.

Lucas, R. 1988. "On the Mechanics of Economic Development," *Journal of Monetary Economics* 22:3-42.

Machin, S. and J. Van Reenen. 1998. "Technology and Changes in Skill Structure: Evidence from Seven

OECD Countries." Berkeley, CA: Berkeley Center for Labor Economics.

Mare, D. 1995. "Living with Graduates: Human Capital Interactions in Metropolitan Areas." New Zealand Department of Labor. Mimeo.

Moretti, E. 1999. "Estimating the External Return to Education: Evidence from Repeated Cross-Sectional and Longitudinal Data." Berkeley, CA: Department of Economics, University of California. Mimeo.

Murphy, K.M. and F. Welch. 1993. "Inequality and Relative Wages," *American Economic Review* 83(2):104-10.

Murray, S.E., W.N. Evans and R.M. Schwab. 1998. "Education-Finance Reform and the Distribution of Education Resources," *American Economic Review* 88(4):789-812.

Osberg, L. 1995. "The Equity/Efficiency Trade-off in Retrospect," *Canadian Business Economics* (Spring):5-19.

Peri, G. 1998. "Human Capital Externalities in U.S. Cities." Berkeley, CA: Department of Economics, University of California. Mimeo.

Perotti, R. 1993. "Political Equilibrium, Income Distribution and Growth," *Review of Economic Studies* 60:755-76.

_____ 1996. "Growth, Income Distribution and Democracy: What the Data Say," *Journal of Economic Growth* 1(2):149-87.

Persson, T. and G. Tabellini. 1992. "Growth, Distribution and Politics," *European Economic Review* 36:593-602.

_____ 1994. "Is Inequality Harmful for Growth," *American Economic Review* 84:600-21.

Psacharopoulos, G. 1994. "Returns to Investment in Education: A Global Update," *World Development* 22:1325-43.

Rauch, J. 1993. "Productivity Gains from Geographic Concentration of Human Capital: Evidence from the Cities," *Journal of Urban Economics* 34:380-400.

Romer, P. 1990. "Endogenous Technological Change," *Journal of Political Economy* 98:S71-S102.

Rosenberg, N. 2000. "America's University/Industry Interfaces: 1945-2000." Stanford, CA: Stanford University. Mimeo.

Shea, J. 2000. "Does Parents' Money Matter?" *Journal of Public Economics* 77:155-84.

Shultz, T.W. 1961. "Investment in Human Capital," *American Economic Review* 51:1-17.

Stiglitz, J.E. 1969. "The Distribution of Income and Wealth among Individuals," *Econometrica* 37(3):382-97.

Tamura, R. 1991. "Income Convergence in an Endogenous Growth Model," *Journal of Political Economy* 99:522-40.

Williamson, J.G. 1993. "Human Capital Deepening, Inequality and Demographic Events along the Asia-Pacific Rim," in *Human Resources in Development Along the Asia-Pacific Rim*, ed. N. Ogawa, G.W. Jones and J. Williamson. Singapore: Oxford University Press.

Wood, A. 1994. *North-South Trade, Employment and Inequality.* Oxford: Clarendon Press.