CHAPTER 6

Does Poverty Matter for Growth?

There is ample evidence that growth reduces poverty. This justifies having a pro-growth package at the heart of any poverty reduction strategy. However, is it also the case that poverty reduction is good for growth? Is there a possibility of entering a virtuous circle by which growth lowers poverty and in turn lower poverty results in faster growth?

THE PREVIOUS CHAPTERS HAVE EXPLORED the link between growth and poverty by focusing on the poverty-reducing effect of growth and the factors that shape it. It was argued that in poorer and more equal countries, development strategies aimed at poverty reduction should emphasize growth. As countries become richer or more unequal, however, policy makers should try to balance growth and distribution concerns because in those circumstances poverty may be much more sensitive to changes in relative incomes than to changes in mean income.[1] We also addressed whether the pattern of growth associated with specific policies and sectors is more pro-poor in some circumstances than in others. We concluded that even though over long-run horizons most pro-growth policies will also be pro-poor (in the sense that the poor receive some benefit from the particular policy), in principle one can expect that growth will have differing effects on poverty in the short run depending on the policies with which it is associated.

A debate on the pro-poorness of a particular pattern of growth can be very appealing from an intellectual viewpoint but of little practical relevance if there is no growth—of any type—to start with or if growth is too low to make a dent in poverty. This in fact may be the root problem because as some development practitioners argue, existing global poverty levels are probably more related to the insufficient growth experienced by developing countries over the past decades than to particularly anomalous patterns of growth. Today the annual median per capita income in developing countries is $3,000, a figure that indicates only modest progress since 1975, when the median income level was about $2,500. Over this same time period, median per capita income in developed countries increased from about $15,000 to more than $25,000.

Against this background and given that the achievement of growth—any type of growth—is a big challenge in itself, should a discussion on growth and poverty reduction, or pro-poor growth, focus first on how to achieve growth and only then consider how to ensure that its pattern is pro-poor? This chapter argues that, on the contrary, the disappointing growth performance of developing countries makes the growth-poverty link even more critical. Not only does low growth mean that even small deteriorations in income inequality may lead to higher poverty (see Cord, Lopez, and Page 2005 for a discussion). It also means that poverty per se may be a barrier to growth, as suggested by several theoretical models developed in the economics

This chapter relies heavily on the background paper "Too Poor to Grow," prepared for this report by H. Lopez and L. Servén (2005b).

literature. In other words, countries do not grow fast because they are too poor to grow. This direction of causality from poverty to growth in turn opens the door to the existence of poverty traps, in which poverty and growth interact in a vicious circle where high poverty leads to low growth and low growth in turns leads to high poverty.

The theoretical appeal of poverty-traps models is clear: these models explain a number of stylized facts on the growth-poverty link (such as the disappointing growth performance of developing countries relative to the developed world or the existence of convergence clubs[2]) for which the traditional neoclassical growth model is inappropriate. Beyond the theoretical appeal, however, several aspects related to the poverty-traps view of the development process are likely to have important policy implications. First, at a strategic level, the existence of poverty traps should mitigate the debate on whether development strategies should rely more on pro-growth or pro-poor policies, because strategies that do not take into account the bidirectional relation between poverty and growth will likely lead to disappointing results: poverty will not decline without growth, but growth will be difficult unless the constraints affecting the poor are also addressed. Second, if a country is trapped in a bad equilibrium, then market policies may not be enough to break the vicious circle between poverty and growth, and policies that change the state of development may be needed. In this regard, country-specific analytical work that blends growth and poverty analyses into a single entity and tries to uncover the potential complex set of interactions operating in a given country would be a first step toward determining exactly which policies are needed to break the poverty trap. Third, at a more operational level, one implication of the potential existence of poverty traps is that the biggest payoff to growth (and hence to poverty reduction) would likely result from policies that not only promote growth but also exert an independent, direct impact on poverty—thereby reducing the drag of poverty on growth.

This chapter elaborates on these issues. It motivates the discussion by briefly reviewing arguments put forward in the literature suggesting how poverty can become self-reinforcing and potentially lead to multiple equilibriums. The chapter then examines the empirical evidence on the dynamics of per capita income. First, it reviews the recent growth experience in the developed and developing worlds, concluding that the developing world has underperformed systematically relative to the developed countries. In fact, the evidence presented here suggests a bimodal income distribution, with countries showing a tendency to cluster around either a high-level efficient equilibrium or a low-level inefficient equilibrium. This clustering is consistent with one of the predictions of poverty-traps models.

Of particular interest here are the findings for the cross-country distribution of incomes in the Latin American region. In contrast to the global data, this distribution appears to be roughly unimodal, implying that most Latin American countries belong to the same convergence club and thus share the same dynamics of the development process in the region. When we also ask to which country cluster the region belongs—the rich or the poor—the results are mixed. On the one hand, it is difficult to argue that the region is stuck in the low, inefficient equilibrium (although admittedly some weak evidence suggests that a few countries in the region—namely, Bolivia, Honduras, and Nicaragua—could be trapped in the poor-countries club).[3] On the other hand, the region does not seem to belong to the rich-countries club either. On the whole, the region would be better described as in an intermediate state somewhere between the very poor and the very rich.

Finally, the chapter presents new empirical evidence suggesting that poverty deters investment and growth, especially where the degree of financial development is limited. This result appears consistent with stylized theoretical models in which financial market imperfections prevent the poor from taking advantage of their investment opportunities, and it suggests a particular mechanism through which poverty affects growth. Admittedly, this mechanism is not necessarily exclusive; moreover, there are other channels, such as education, health, and innovation, through which high poverty can potentially feed back into lower growth rates. In any case, we emphasize here that this chapter, and more generally this report, does not aim at setting the debate on the existence of poverty traps (defined as the existence of multiple steady states); admittedly the empirical evidence on this question is mixed at best. Instead, its main concern is whether the empirical evidence supports a weaker version of the predictions derived from poverty-traps models, namely, that poverty tends to hold back growth.

A poverty-traps view of the development process

The past few years have witnessed the emergence of a booming theoretical literature aimed at explaining why poverty may be self-reinforcing and therefore why countries that start out being poor continue to be persistently

FIGURE 6.1
Traditional view of the growth-poverty relationship

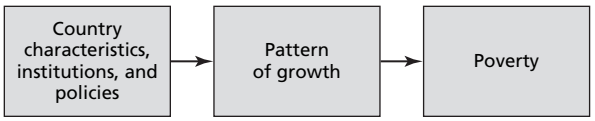

FIGURE 6.2
Poverty-traps view of the growth-poverty relationship

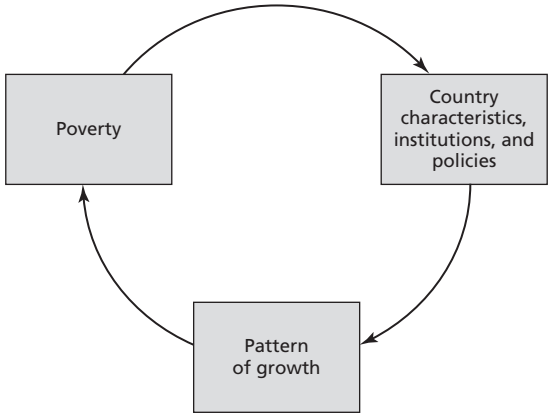

poor over the long run (see Azariadis and Stachurski 2005 for a survey). In the traditional view of development (presented schematically in figure 6.1), country constraints (institutions, policies, internal and external shocks, and the like) are considered to be largely exogenous (that is, they are not determined within the system). In contrast, the poverty-traps literature stresses the possibility that poverty has feedback effects on growth, a dynamic that has the potential to create poverty traps and that results in a very different picture of the development process (figure 6.2).

One critical difference between the two development views is that in the poverty-traps view, different equilibriums may exist that are stable and self-reinforcing so that the initially poor may stay poor and the initially rich stay rich. Figure 6.3 illustrates this point, comparing the results of the standard neoclassical growth model with decreasing returns to scale (panel A) with a model that exhibits increasing returns to scale (panel B). In the case of the standard neoclassical growth model, the equilibrium is uniquely determined by the intersection of per capita savings and investment ($s \times y$) with the rate of depreciation of the per capita capital

FIGURE 6.3
Multiple equilibriums in the presence of increasing returns to scale

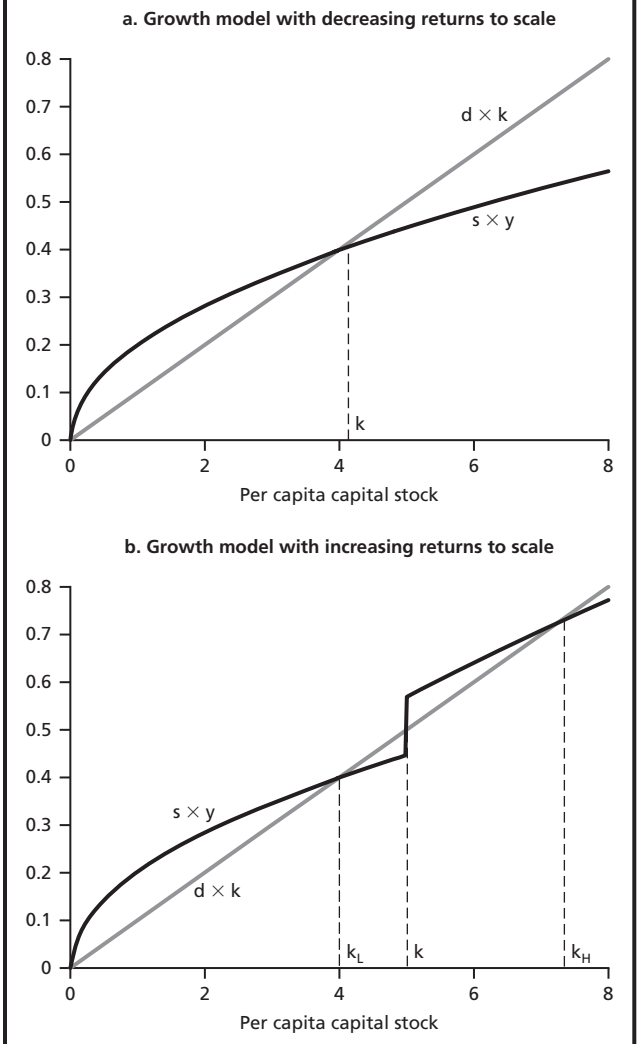

Source: Authors.

stock ($d \times k$). If, however, the production function experiences a technological jump (discussed in more detail later), there would be two steady states, and countries would tend toward one or the other equilibrium depending on their initial position. The lower equilibrium could be thought of as a poverty trap. Countries with capital below k_L would initially grow and converge toward the steady-state k_L. Countries between k_H and k would converge toward k_H. Thus initially poor countries would converge toward the low, inefficient equilibrium whereas initially rich countries would tend toward the high, efficient equilibrium, producing a bimodal cross-country distribution of income.

In these circumstances policies aimed at eliminating market distortions that prevent the economy from moving toward its equilibrium may be highly effective at achieving their objective. The problem is that the economy may be headed toward an inefficient equilibrium. Thus poverty-traps models have the ability to explain both why poor economies may have a tendency to underperform richer economies and why the benefits of good policies fail to materialize. What are the mechanisms that lead to this type of feedback from poverty to growth? Several channels, typically in the form of departures from the basic neoclassical model, have been explored in the literature.[4] We briefly discuss three of those channels here.

Increasing returns to scale and poverty traps

As suggested earlier, one mechanism that may potentially lead to poverty traps is the existence of increasing returns to scale (this is the issue illustrated in panel B of figure 6.3). Increasing returns may appear when the adoption of newer and more efficient technologies has an associated fixed cost. For example, Murphy, Shleifer, and Vishny (1989) argue that even if modern technologies are freely available to poor countries, when the size of the domestic market is small relative to the fixed costs required to adopt the new, more efficient technology, firms may not have the right incentive to do so. As a result, initially richer economies may enter a virtuous circle, whereas initially poorer economies may end up stuck with less-efficient technologies and lower income levels. Increasing returns may also appear in the presence of complementary production processes that act as an incentive for firms to match workers of similar skills, in which case the incentive to educate increases as the initial pool of skilled workers increases (Kremer 1993).

Market failures and poverty traps

A second mechanism that may generate poverty traps is related to the existence of potential market imperfections in credit and insurance markets. With perfect capital markets, investment decisions in physical or human capital depend on the expected returns (probably adjusted by risk) of the investment and on the associated cost. When the returns are higher than the cost of capital, an individual would have the same incentive to invest regardless of his or her initial income level: theoretically, poor people could always borrow the capital they need to make the investment and then repay the loan out of the returns of the investment.

However, in real life—and especially in developing countries—capital and financial markets are plagued with imperfections. In many economies large segments of the population may not have access to credit at all. In some cases, access to credit is denied because the poor do not have the necessary collateral. In other cases, financial operators may find it difficult to enforce contracts, and an individual's access to credit will likely be constrained by his or her initial wealth; those with low or no initial wealth may be excluded from capital markets. Moreover, even those with access to credit may encounter significant constraints. Since deposit rates tend to be much lower than borrowing rates (figure 6.4), the opportunity cost of capital is lower for those who need to borrow less. For example, the average interest rate spread (lending minus deposit) for 2003 in the

FIGURE 6.4
Interest rate spreads in Latin America, 2003

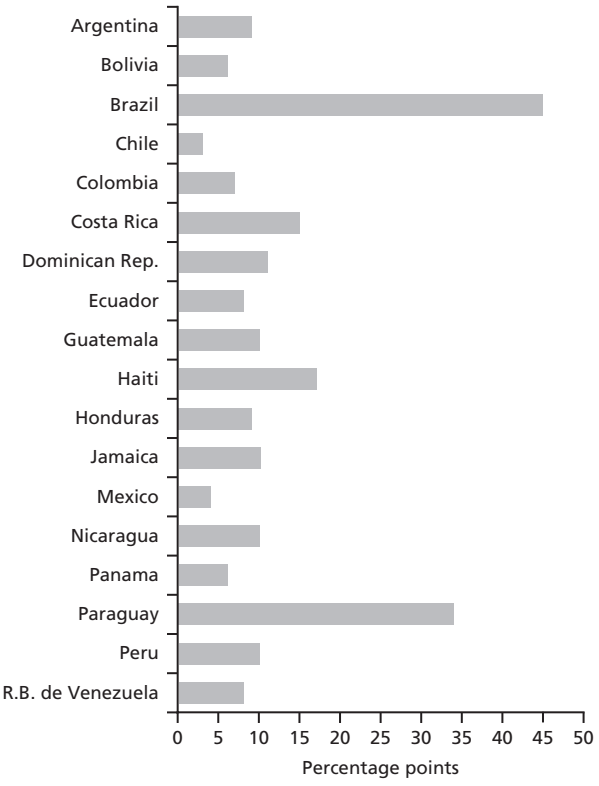

Source: WDI database.
Note: The figure reports lending minus deposit rates.

sample of Latin American countries included in figure 6.4 is about 10 percentage points, but in specific countries (Brazil and Paraguay), it is more than 30 points. Thus, if both a rich and a poor person face a similar rate of return on a project, it is likely that the rich person will invest much more than the poor person. In other words, the opportunities and costs of borrowing can be very different for rich and poor people and play against the latter group.

Imperfect capital markets coupled with fixed costs imply that important segments of the population are excluded from investment opportunities. For example, Banerjee and Newman (1994) stress the effect that an individual's initial wealth has on the level of physical investment when there are credit constraints. Thus high poverty rates might result in low investment rates and hence in lower growth.

Galor and Zeira (1993) make a similar argument. They note that people at the bottom of the income distribution may not be able to cover the expense of education or access the financial sector to borrow for that purpose. In this case high poverty rates result in low educational outcomes because poor individuals likely opt out of the education sector and work at unskilled, low-return labor. Note that this effect goes beyond the lower supply of education possibilities in poorer countries and focuses on the demand side. As argued in de Ferranti and others (2003), education levels are a vital complement for technological advance and are thus a critical element in understanding growth rates (box 6.1).

Institutional mechanisms and poverty traps

The theoretical literature also stresses the role played by the institutional framework in generating poverty traps. For example, Engerman and Sokoloff (forthcoming) argue that institutions that place economic opportunities beyond the reach of broad segments of society are likely to result in reduced growth rates because modern economies require broad participation in entrepreneurship and innovation. In addition, a natural tendency for those who hold power to try to perpetuate that power results in path dependence and persistence for the institutional framework. These two elements together help explain the tendency for poverty and bad institutional arrangements to coexist and persist over time.

Similarly, Mauro (2002) considers low economic growth in countries with persistent corruption and notes that some countries appear to be stuck in a bad equilibrium characterized by pervasive corruption with no sign of improvement. He argues that one reason why rooting out widespread corruption is so difficult may be that it just does not make sense for individuals to attempt to fight it, even if everybody would be better off if corruption were to be eliminated. For example, if corruption is widespread in an administration, civil servants might find it difficult to decline bribes in exchange for favors because their superiors may expect a portion of the bribe for themselves. In contrast, in bureaucracies that are generally honest, a real threat of punishment deters individual civil servants from behaving dishonestly. This is an example of a strategic complementarity, whereby if one agent does something it becomes more profitable for another agent to do the same thing. The tendency of corruption to persist, together with the negative impact of corruption on growth (Mauro 1995), would then explain why some countries may be caught in inefficient equilibriums.

BOX 6.1

Education and technology

Productivity differences between countries and between firms within countries are profoundly affected by differences in skills and technology. It is therefore no surprise that the East Asian tigers—Hong Kong (China), Republic of Korea, Singapore, and Taiwan (China)—which exhibit well-above-average rates of total factor productivity growth, also outperform Latin America on measures of technology and skills. The same is true for some of the successful natural resource–based economies. Within Latin America, the best-performing country, Chile, concurrently had positive increases in productivity, substantial skill upgrading, and increases in all indicators associated with technology transfer and innovation.

Source: de Ferranti and others 2003.

In summary, a variety of mechanisms that typically do not fit the assumptions underlying the neoclassical model may both cause poverty and perpetuate it over the long run. Moreover, many of these mechanisms may well interact with and reinforce each other. For example, corruption may exacerbate credit access problems if the public sector subsidizes or guarantees credit to some privileged groups in society at the expense of poorer segments of the population. Similarly, institutional frameworks with weak enforcement of the rule of law may discourage investment in sectors where intellectual property rights have a high value for the firm. That in turn can lower the demand for skilled workers and hence the incentives for individual workers to invest in skill acquisition. The next section reviews some existing empirical evidence on the practical relevance of these models.

Empirical evidence on poverty traps

Over the last decades, the world has become increasingly divided into two clubs—one of rich countries, the other of poor countries. Figure 6.5 plots median per capita growth rates for industrial and developing countries between 1963 and 2003.[5] It also plots median per capita growth rates for Latin America. The figure indicates that, apart from one short period in the second half of the 1970s and another in the early 2000s, the typical developing country has experienced lower growth rates than the typical rich country. Over the 1963–2003 period, median per capita growth in industrial countries has outpaced median growth in developing countries by an average of more than 1 percent a year.[6] Moreover, there are two extended periods of time—the 1960s and early 1970s, and the mid- to late 1980s—where the differences are consistently in the range of 2 percent a year.

Latin America does not seem to be an exception among developing countries; the growth performance of the region over the 40-year period was fairly consistent with the performance observed in other developing countries. The differences between Latin America and all developing countries were notable for three periods: the early 1980s, when Latin America was badly hit by the debt crisis and recorded median growth rates below −1 percent; the early 1990s, when the region did much better than the rest of the developing countries; and the late 1990s, when once again Latin America experienced a significant deceleration following the financial crises in East Asia in 1997 and in Russia in 1998.

The underperformance of the developing world relative to the developed world appears even more dramatic when one looks at the evolution of median per capita income levels over time (figure 6.6). Because developing countries have been experiencing lower growth rates for prolonged periods of time, the gap between the per capita income levels of rich and poor countries has been steadily increasing. In the early 1960s, the income level of the median

FIGURE 6.5
Growth in developed (OECD) and developing countries, 1963–2000

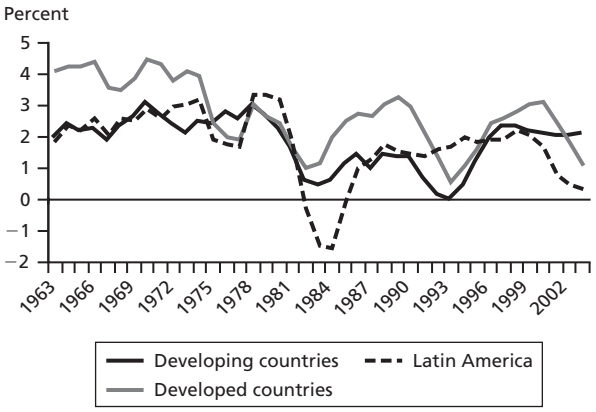

Source: WDI database.
Note: The chart reports the 3-year moving average of the median per capita growth for each group of countries.

FIGURE 6.6
Income in Latin America relative to the OECD countries, 1960–2002

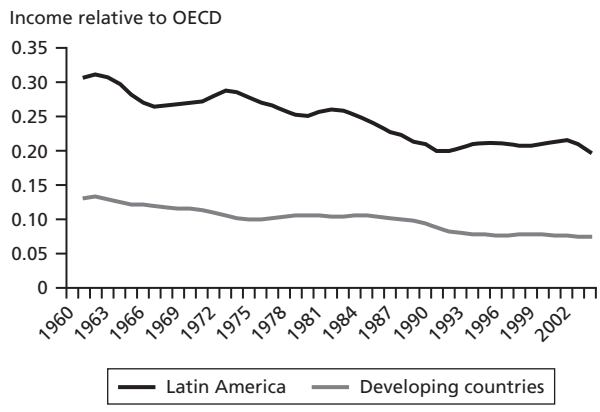

Source: Authors' calculations.

developed country was six times greater than the income level of the median developing country; today income in the median developed country is close to nine times greater (representing a 50 percent increase in the gap). More dramatically, in 1960 the income of the richest country at the time, Switzerland, was about 50 times the income of the poorest country, Malawi. Today, the richest country is Luxembourg, which has a per capita income level that in purchasing power parity is almost 120 times that of Sierra Leone, now the poorest country.

The use of the median as a summary statistic is somewhat limited because it does not show the significant heterogeneity that exists at the country level. Yet, even if we focus on the evolution of income on a country-by-country basis (table 6.1), the majority of the Latin American countries (the exception is the Dominican Republic) have income levels today that are lower than they were in 1960 relative to the income of OECD countries. Not only have the majority of Latin American countries lost ground over the past 25 years but in some cases the decline has been very significant. Take the case of Argentina, the richest country of the region in 1960 with a per capita income level that was close to the level of industrial countries (85 percent). Forty years later Argentina's relative income has declined to 43 percent of the industrial countries' level. Similarly, the relative per capita income of Nicaragua has declined from 49 percent in 1960 to about 12 percent in 2000. Today three countries in Latin America (Bolivia, Haiti, and Honduras) have PPP-adjusted per capita GDP levels that are less than 10 percent of the income of the developed countries. In 1960 no country in the region had a relative income level below 20 percent.

On the whole, this evidence is at odds with the convergence predictions of the simple neoclassical model and instead is more consistent with what World Bank economist Lant Pritchett (1997) refers to as "divergence big time": "Whichever way the debate about whether there has been some 'conditional' convergence in the recent period is settled, the fact remains that one overwhelming feature of the period of modem economic growth is massive

TABLE 6.1

Median income in Latin America and the Caribbean relative to the industrial countries

Country	1960	1970	1980	1990	1998	2003
Argentina	0.85	0.72	0.64	0.40	0.52	0.43
Bolivia	0.22	0.15	0.14	0.10	0.10	0.09
Brazil	0.30	0.28	0.38	0.29	0.29	0.27
Chile	0.37	0.30	0.26	0.26	0.36	0.36
Colombia	0.32	0.27	0.27	0.27	0.25	0.24
Costa Rica	0.46	0.37	0.38	0.29	0.32	0.34
Dominican Republic	0.21	0.18	0.22	0.19	0.22	0.24
Ecuador	0.22	0.16	0.19	0.16	0.14	0.13
El Salvador	0.38	0.32	0.24	0.17	0.18	0.17
Guatemala	0.26	0.22	0.23	0.16	0.15	0.15
Guyana	0.30	0.23	0.20	0.16	0.17	0.15
Haiti	0.31	0.18	0.18	0.11	0.07	0.06
Honduras	0.20	0.15	0.15	0.11	0.10	0.09
Jamaica	0.29	0.28	0.19	0.18	0.15	0.14
Mexico	0.42	0.39	0.44	0.34	0.33	0.32
Nicaragua	0.49	0.46	0.26	0.13	0.12	0.12
Panama	0.26	0.28	0.29	0.21	0.24	0.24
Paraguay	0.25	0.21	0.27	0.21	0.19	0.17
Peru	0.41	0.35	0.30	0.18	0.19	0.19
Trinidad and Tobago	0.49	0.46	0.53	0.32	0.32	0.38
Uruguay	0.62	0.43	0.43	0.33	0.37	0.29
Venezuela, R.B. de	0.69	0.54	0.38	0.26	0.25	0.17
Latin America	0.31	0.28	0.26	0.19	0.21	0.19

Source: Authors' calculations using GDP per capita ($2,000 PPP) from the *World Development Indicators* for various years. Data before 1975 has been computed using available per capita growth rates for the period 1960–75 and the per capita GDP level of 1975.

divergence of absolute and relative incomes across countries, a fact which must be grappled with in a fully satisfactory model of economic growth and development."

Convergence clubs

What explains this apparent divergence between developed and developing countries? Could it be attributable to the existence of multiple states of development toward which different countries converge, creating convergence clubs? If so, where is the Latin American region in this picture? Are there also regional convergence clubs that will result in regional clusters of development or, instead, can the region be viewed as a single entity? We now address these questions in turn.

Convergence clubs in absolute income levels

The first question concerns the dynamics of cross-national per capita income levels and the existence of convergence clubs. Panel a of figure 6.7 presents the histograms of per capita income for 1960 and 1999 computed for 102 countries using data from the Penn World Table (PWT6.1). The histograms suggest that whereas in the early 1960s the distribution of income appeared to be unimodal in the early 1960s, by the late 1990s it had become trimodal,

FIGURE 6.7

Histograms for per capita income, 1960s versus the 1990s

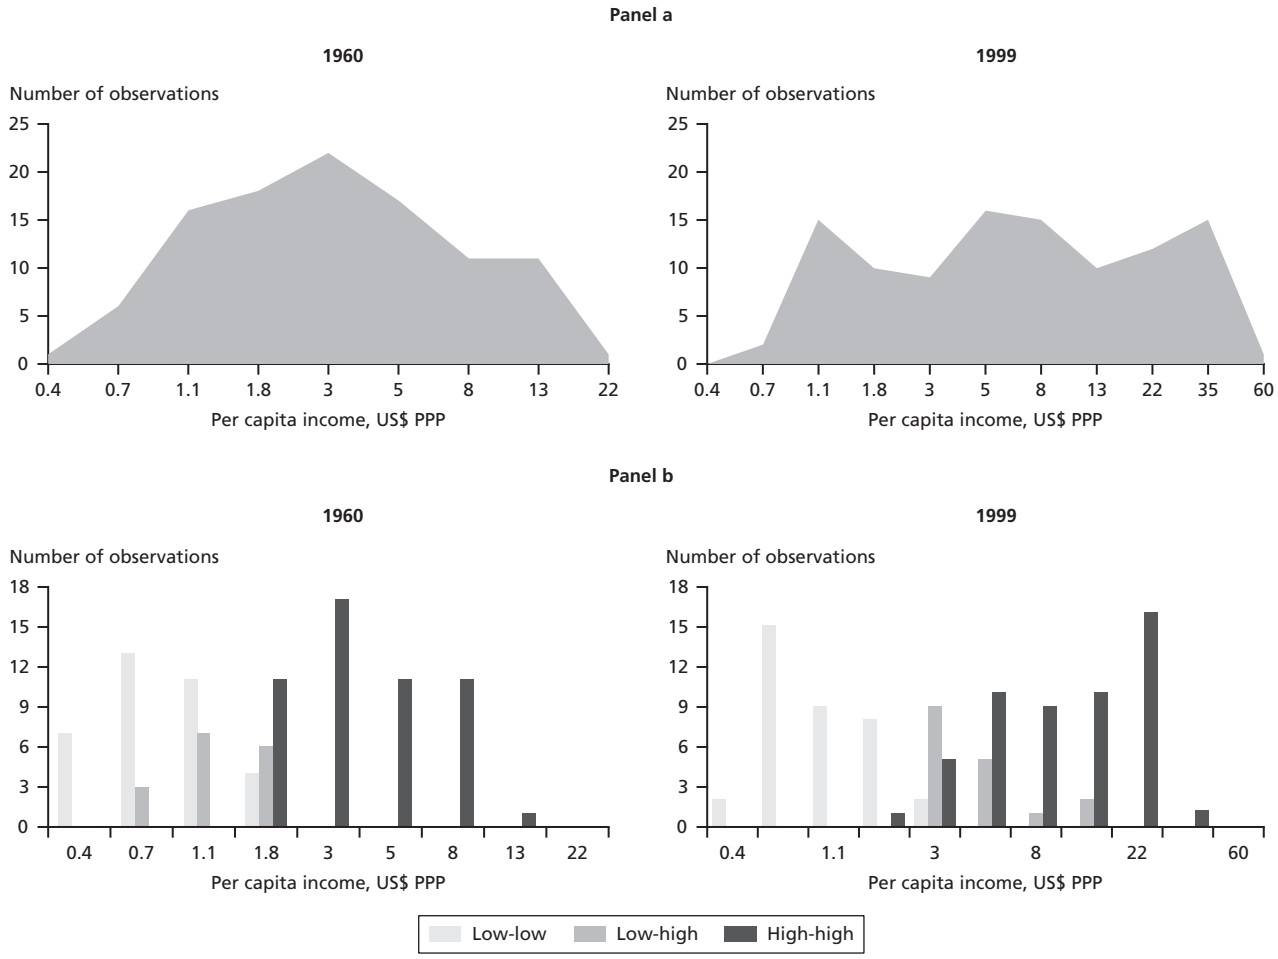

Source: Penn World Tables (PWT) 6.1.

Note: The top panel reports the histograms of the cross country per capita income distribution (102 countries) in 1960 and 1999. The bottom panel presents the transitions of three groups of countries: low-low shows countries that in both 1960 and 1999 had per capita income levels below $3,400; high-high shows countries that in both 1960 and in 1999 had per capita income levels above $3,400; low-high shows countries that in 1960 were below $3,400 and in 1999 were above $3,400.

with a low peak at $1,100; a second peak between $5,000 and $8,000, and a third peak around $35,000.[7]

In panel b we attempt to discriminate between convergence clubs and present the histograms for three groups of countries. Here we follow an approach similar to the one used by Mayer-Foulkes (2003) in his study of convergence clubs in life expectancy and divide the sample into four groups. The first group includes those countries whose per capita income levels were below $3,400 in both 1960 and 1999. This is the per capita income level of the poorest industrial country in 1960 (Portugal) and is very close to the observed peak in 1960. We refer to that group as low-low. The second group includes countries with per capita income levels above $3,400 in both 1960 and 1999. This is the high-high group. The third group (low-high) comprises countries with per capita income levels below $3,400 in 1960 and above $3,400 in 1999. No country falls in the fourth group, which notionally corresponds to a high-low group, and the numbers of countries in each of the other three groups are quite balanced.

Panel b shows three markedly different behaviors. The initially rich countries present the highest per capita growth rates. The median income of the high-high club increased from about $7,500 in 1960 to about $22,000 in 1999 (table 6.2). The transition countries (the low-high group) also show considerable growth (from a median income of about $2,400 in 1960 to about $5,400 in 1999), but the average annual growth rate is lower than in the high-high group by almost 0.7 percentage point. Finally, the low-low group shows very low growth. The median income for the 37 countries in this group increased from about $1,050 in 1960 to just $1,300 in 1999, which implies an average annual increase of about half a percent.

Clearly, the peaks in the histogram for 1999 may not correspond to the equilibriums for the different groups, especially if the groups are in a transition toward a steady state. Where, then, is each of these groups heading? The annex to this chapter discusses a simple procedure that can be used to estimate the steady state for each group. Implementation of this procedure suggests convergence but to three dramatically different steady states. For the low-low group, the estimated equilibrium for per capita income is around $1,700. For the low-high group, the equilibrium is around $11,000, and for the high-high group, the point estimates suggest an equilibrium well above current levels.

How does the Latin American region fare in this context? Is the apparent bi- or trimodality of the world distribution also observed in the region, or do all the countries in the region belong to a single cluster? To answer these questions, figure 6.7 plots a histogram similar to the one in panel A of figure 6.6 but restricts the sample to Latin American countries. In contrast to the full sample, the estimated cross-country distributions of per capita income for Latin America appear to be unimodal for both the early 1960s and the late 1990s. The peak in 1960 is around $3,000, which is fully consistent with the global data. The peak in 1999 is around $8,000, which implies average annual growth in the 2.5 percent range, approximately halfway between the growth rates for the global high-high and low-high groups.

How do we interpret these results? Well, it depends on whether we see the glass as half full or half empty. As a half-full glass, it seems difficult to argue that the region is stuck in the low, inefficient equilibrium (the one corresponding to the equilibrium around $1,700). More likely, taking into account the initial starting point and the evolution of income levels over the 1960–99 period, the region is better characterized as belonging to the low-high transition group (for which the estimated equilibrium for income per capita is in the $11,000 range). As a half-empty glass, the region does not seem to belong to the high-high equilibrium either. On the whole, the region would be better described by an intermediate state somewhere in between the very poor and the very rich.

One issue needs to be highlighted before we continue, however. Careful observation of figure 6.8 indicates that the dispersion of regional income in 1999 is significantly higher than it was in 1960. This results from the relatively good performance of some of the economies that were richer to begin with (Chile, Mexico, and Uruguay) and the modest performance of some of the poorer economies (Bolivia, Honduras, and Nicaragua), which initially

TABLE 6.2

Median income of convergence clubs

Club	Countries	Median income		Annual increase (%)
		1960	1999	
Low-low	37	1,046	1,277	0.51
Low-high	33	2,395	5,442	2.13
High-high	32	7,417	21,632	2.78

Source: Authors' calculations.

FIGURE 6.8
Histograms for per capita income in Latin America, 1960s versus the 1990s

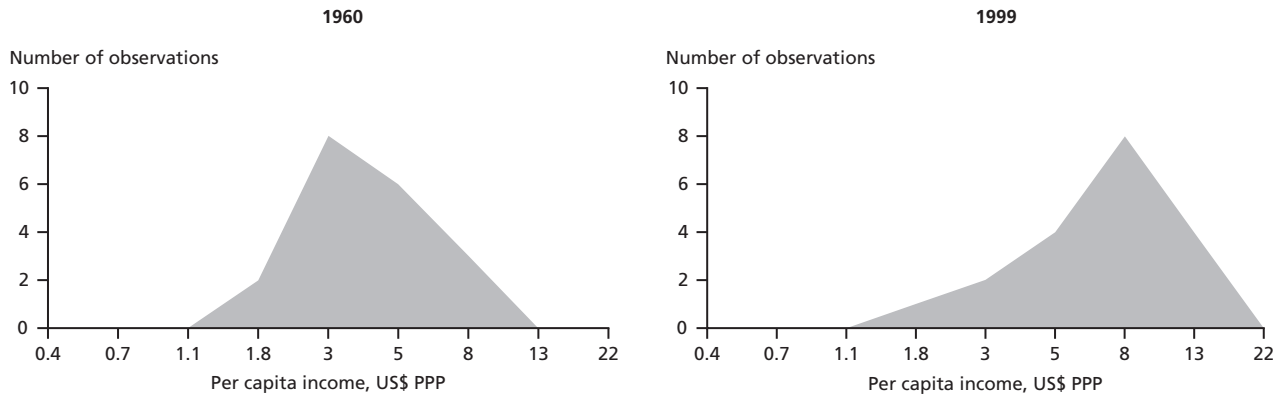

Source: Penn World Tables (PWT) 6.1.
Note: The figure reports the histograms of the cross-country per capita income distribution (18 countries) for the Latin American region in 1960 and 1999.

experienced average annual growth rates below 0.5 percent (Nicaragua's average annual growth rate was in negative territory). At least three countries in the region appear to have a performance that is more consistent with that observed for the low-low group in figure 6.7, and these countries could potentially be trapped in the low equilibrium. In other words, behind figure 6.8 there could be a bimodal distribution, with a second steady state toward the lower end of the distribution that is not apparent because the associated probability mass is very low (that is, because only a few countries belong to that group).

Convergence clubs in relative incomes

An alternative way to look at the cross-national distribution of income is based on an analysis of relative income levels and on the probability that a country moves between states of development. In the technical annex to this chapter, we review some methodological details and present some empirical results that can be used to estimate equilibrium values for the distribution of income. Figure 6.9 reports results for five states of relative development. In state 1 are the poorest countries of the world: those with per capita income levels below 25 percent of average world per capita income. In state 2 is a group of richer but still relatively poor countries: those with per capita income levels between 25 and 50 percent of average world per capita income. State 3 includes economies that have income levels between 50 percent and the world average. States 4 and 5 cover the richest countries: those with per capita incomes between the world average and twice the average, and those with incomes above twice the average, respectively.

FIGURE 6.9
Twin peaks

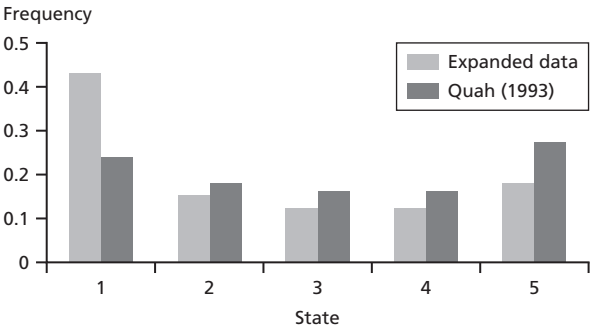

Source: Quah (1993) and authors' calculations.

Figure 6.9 plots the equilibrium as computed by Quah (1993) on the basis of data spanning 1962–84, and it also plots the equilibrium that results when the analysis is based on an expanded sample covering 1960–99. A number of interesting points are revealed in this figure. First, both samples suggest the presence of convergence clubs at either end of the income distribution: there is a cluster of poor countries around a low per capita income equilibrium and a second cluster around the high per capita income equilibrium (that is, the poor tend to stay poor and the rich tend to stay rich). However, while the 1962–84 sample results in a picture of the world that is divided almost

FIGURE 6.10
Equilibrium and distribution in 1999

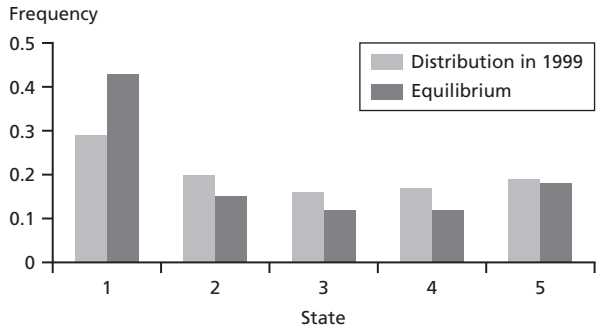

Source: Authors' calculations.

FIGURE 6.11
Latin American states: One peak?

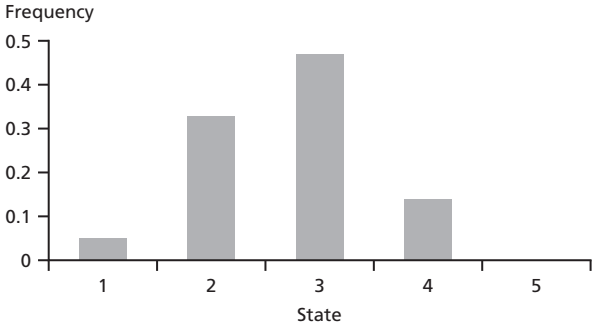

Source: Authors' calculations.

symmetrically, the 1960–99 sample produces a distribution that is clearly skewed toward the lower equilibrium (that is, the cluster of poor countries has more members).

In other words, while evidence of some type of bimodality still exists, the expected long-run frequency of countries in the first state increases by almost 20 percentage points (from 0.24 to 0.43) by expanding the sample. This finding implies that the updated estimates predict more countries falling behind (at least relative to the world average). This is further explored in figure 6.10, which compares the distribution in 1999 to the estimated equilibrium. The figure suggests that unless there are changes in the transitional dynamics of the growth process, the number of countries in the first state can be expected to increase.

Unlike our previous analysis where the empirical evidence pointed toward a three-club characterization, this body of evidence is more consistent with the existence of two convergence clubs. One is composed of very poor countries, apparently with loose rules of admission; on the basis of the data to 1999, more than 40 percent of the countries belong to this club. The second club—the rich-countries club—is much more exclusive, and our estimates suggest that only about 20 percent of the countries belong to it. (The remaining 40 percent of the countries lie somewhere in between these two convergence clubs.)

The difference between having two or three clubs is key for Latin America, given our earlier conclusion that the region fell somewhere between the low and the high equilibrium. To explore this issue, we replicate the previous exercise but use data only for Latin America. The results suggest that there are important differences in the estimated long equilibrium (figure 6.11). As in figure 6.8, the obtained results for the region do not show evidence of bimodality. Instead, there seems to be a long-run equilibrium around state 3. The cross-country distribution of income, however, is not symmetrical, and long-run equilibrium computed on the basis of the estimated transition matrix places 80 percent of Latin American countries in states 2 and 3; these are countries whose relative income ranges from 25 percent of the world average to the world average.

These results are largely consistent with those of the previous analysis and show the region on an equilibrium that is well below the world average. The estimates also show a disturbing tendency for Latin American countries to cluster around the lower tail of the equilibrium. Here the only thing we can do is to speculate that a relatively small group of countries in the region do not belong to the state 3 equilibrium and instead converge around state 2.

Convergence clubs in other dimensions of poverty

So far we have focused on the cross-national distribution of per capita income. However, there is no reason to constrain the analysis to the income dimension of welfare. Convergence clubs may also involve specific health phenomena. For example, the theory of efficiency wages in Dasgupta and Ray (1986) implies the possibility of a low-productivity, low-nutrition trap. Mayer-Foulkes (2003) argues that the existence of convergence clubs is also apparent in life-expectancy dynamics. Figure 6.12 presents cross-national life-expectancy histograms for 1960 and 2002. These histograms indicate the presence of a two-peaked pattern in

FIGURE 6.12
Convergence clubs in life expectancy

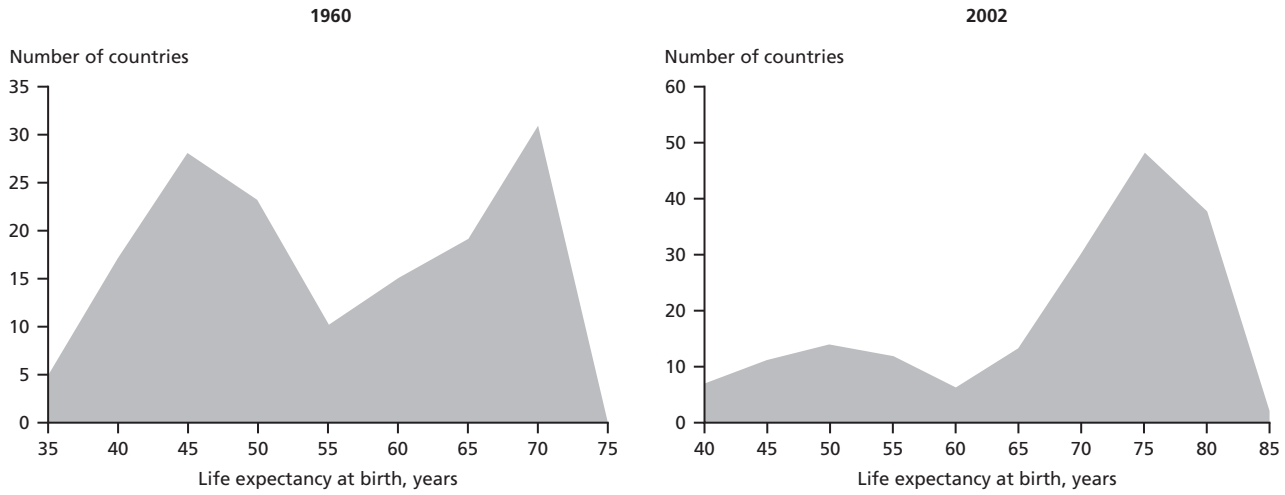

Source: Authors' calculations.

both periods. It is also evident that the mass of the low peak declines between 1960 and 2002, whereas the mass of the high peak increases. These figures are basically a replica of those in Mayer-Foulkes (2003) and indicate that the cross-country data on life expectancy are consistent with the presence of three convergence clubs (with a different number of members): one for the low equilibrium, one for the high equilibrium, and one for a third transitional group.

Formal tests of the poverty-traps hypothesis

The empirical evidence discussed here is supportive of a multimodal distribution in cross-national per capita income levels, which is consistent with the predictions of poverty-traps models. However, as Azariadis and Stachurski (2006) argue, one has to be extremely careful to avoid taking these empirical findings as evidence of poverty-traps phenomena. In fact, in a recent study, Bloom, Canning, and Sevilla (2003) stress that a multimodal distribution in cross-country income levels is also consistent with the existence of fundamental differences between countries that result in different but unique equilibriums for each country. Thus, in principle one has to be able to determine whether bimodality results from two "similar" countries having completely "different" states of development (that is, poverty traps) or from fundamental differences between the two countries.

With these ideas in mind, Bloom, Canning, and Sevilla (2003) move beyond the pure description of the cross-national income distribution and find that the existence of twin peaks in the data is more likely attributable to multiple equilibriums than to fundamental forces. This, in turn, supports the hypothesis that poverty traps with low and high equilibriums underlie the dynamics of per capita income.

An alternative way to determine the existence of poverty traps is to investigate specific sources of multiple equilibriums. One such approach is the calibration of models consistent with the poverty-trap hypothesis. Once a model has been calibrated, its empirical relevance can be assessed. For example, Graham and Temple (2004) calibrate a two-sector general equilibrium model and then explore the extent to which this model is able to explain the real data. The model considers a traditional agricultural sector with diminishing returns and a nonagricultural sector with increasing returns (in the vein of our earlier discussion about poverty traps in the presence of increasing returns to scale). As it turns out, the degree of increasing returns is one of the key parameters underlying the simulations, and depending on its assumed value, the model can explain between 15 and 60 percent of the variance of incomes. The Graham and Temple analysis has the same limitations in the Latin American context, however. In particular, as

the authors recognize, the model appears to explain the existing income differences between the low- and middle-income countries better than it explains the differences between middle-income and developed countries. Thus while the results they obtain offer some ideas of why African countries are so poor, they have much less to say about the current positions of Latin America relative to the industrial countries.

Kraay and Raddatz (2005) also calibrate simple aggregate models capable of generating poverty traps through low savings or low technology at low levels of development.[8] The basic idea behind these models is that if either the saving rate or productivity increases above a certain threshold of development, it would then be possible to find poverty-trap-like features in the data. To assess the empirical relevance of these models, Kraay and Raddatz explore whether saving rates exhibit the sort of nonlinear relationship implied by the model for the existence of poverty traps, and whether scale effects on productivity are of a magnitude consistent with the theoretical model. Unlike Graham and Temple's findings, their results do not lend much support to the existence of poverty traps based on these mechanisms. In particular, their technology-based model suggests that for a poverty trap to exist, the estimated returns to scale would have to be in the 1.4 to 2.5 range. This interval is much higher than is typically found in the literature, where most studies report constant to moderate increasing returns.

Another strand of the empirical poverty-traps literature has explored the existence of nonconvexities by exploiting existing microeconometric evidence. For example, McKenzie and Woodruff (2004) examine the empirical relevance of the assumptions that minimum start-up costs are high relative to wealth and that returns to capital are low at low investment levels (see Banerjee and Newman 1993). Using microenterprise data for Mexico, McKenzie and Woodruff show that the median investment levels of new firms in some sectors are very low (about US$100, or less than half of the monthly earnings of even a low-wage worker). They also show that the marginal return to capital is quite high even for low levels of invested capital (in the $200 range), concluding that the Mexican evidence does not support this particular mechanism as a candidate to justify the existence of poverty traps.

Similarly, Lokshin and Ravallion (2004) test for the existence of a threshold effect in household incomes using data for Russia and Hungary. They find no evidence to support the poverty-traps hypothesis (although they do find that the adjustment of income to shocks is nonlinear). Their results indicate that households tend to bounce back from transient shocks, although the adjustment process is slower for poorer individuals. Jalan and Ravallion (2002) use household panel data from China, however, and find that aggregate physical and human capital endowments play a significant role in household consumption growth, a finding that they argue is consistent with the existence of regional poverty traps.

On the whole, it must be acknowledged that the empirical evidence on the existence of poverty traps is, at best, mixed. How then do we explain the existence of convergence clubs alongside the relative lack of evidence on the existence of poverty traps? One possibility is that poverty traps do exist and that the econometric models used to test such hypotheses are unable to capture the dynamics behind the data. An alternative possibility is that poverty traps do not exist in the strict theoretical sense (multiple equilibriums created, for example, by increasing returns to scale or any other mechanism), but that poverty is still a barrier to growth by which poorer countries find it more difficult to grow than richer countries. In this regard, Azariadis and Stachurski (2006) use a much more general definition and classify any self-reinforcing mechanism that causes poverty to persist as a poverty trap. Note that with this alternative definition in mind, the important question is not whether the development process is characterized by the existence of multiple equilibriums but rather how persistent and self-reinforcing the mechanisms are that lock in poverty over time frames that matter from a policy perspective. But is there any empirical evidence suggesting that poverty may represent a barrier to growth? The next sections explore this issue.

What is the empirical evidence on poverty's impact on growth?

The past few years have witnessed a renewed interest in both the theoretical and the empirical relationship between inequality and growth. At the theoretical level, two main types of arguments have been put forward: sociopolitical economy arguments and credit constraint–factor accumulation arguments.

The sociopolitical economy arguments stress the role that high inequality may play in the decisions of various

agents and how these decisions may influence growth. For example, Alesina and Rodrik (1994) suggest that high inequality may lead to lower growth if the level of taxation has a negative impact on capital accumulation, if taxes are proportional to income but the benefits of public expenditure accrue equally to all individuals (implying that an individual's preferred levels of taxation and expenditure are inversely related to her income), and if the tax rate selected by the government is the one preferred by the median voter. Similarly, Alesina and Perotti (1996) argue that highly unequal societies create incentives for individuals to engage in activities, such as crime, that are outside normal markets and that sociopolitical instability discourages accumulation because of current disruptions and future uncertainty. In both cases, high levels of inequality may lead to lower future growth.

The credit constraint–factor accumulation argument emphasizes the possibility that some individuals will be excluded from the economic process because they have neither the resources nor the means to borrow them to engage in potentially profitable economic activities. For example, as discussed earlier, Galor and Zeira (1993) argue that the process of development is characterized by complementarity between physical and human capital so that growth increases as investment in human capital increases. However, credit constraints may prevent poorer individuals from investing in education and thus affect growth prospects by reducing the number of individuals who are able to invest in human capital. Similarly Aghion, Caroli, and García-Peñalosa (1999) show that if there are decreasing returns to individual capital investments and if credit imperfections mean that individual investments are an increased function of initial endowments, then the concentration of investment in fewer richer people will negatively affect growth.

Admittedly for a given level of income, higher inequality will lead to higher poverty. But note that the credit constraint–factor accumulation argument is more a poverty argument than an inequality argument. Yet, to the best of our knowledge, the hypothesis that countries suffering from higher levels of poverty grow less rapidly than those countries with less poverty has remained untested. To fill that gap, in a background paper for this report, Lopez and Servén (2005b) make a first attempt to provide a direct empirical assessment of the impact of poverty on growth (see the technical appendix).

The main results of that work are the following:

- Poverty has a consistently negative impact on growth that is significant both statistically and economically.
- This negative growth effect seems to work through investment in the sense that high poverty deters investment, which in turn lowers growth.
- The data suggest that this mechanism operates only at low levels of financial development, consistent with the predictions of theoretical models that underscore financial market imperfections as a key mechanism of poverty traps.

We now review each of these findings in some detail.

Poverty is bad for growth

Lopez and Servén (2005b) begin with the observation that if poverty hampers growth, then countries with higher initial poverty should grow less rapidly than comparable countries with lower initial poverty, all else being equal. This hypothesis is a weaker version of the predictions derived from the analytical models on poverty traps, in that to support it one does not need to find evidence of multiple equilibriums but simply empirical proof that poverty tends to hold back growth. Using a standard growth model augmented to include a suitable poverty measure among the explanatory variables, the authors find that after controlling for other factors, poverty has a negative and strongly significant impact on growth, which is also economically significant. On average, a 10 percent increase in poverty reduces annual growth by 1 percentage point. This finding is robust to a number of basic departures from the basic specification in Lopez and Servén (2005b),[9] including:

- *The use of alternative poverty lines.* The estimated impact on growth of a change in headcount poverty is very similar regardless of the poverty line ($2, $3, or $4 a day) used in the computation of the poverty index. Changes to the poverty line have an impact on the estimated coefficient of poverty of around 0.01.
- *The use of different sets of control variables.* Changing controls seems to have only a moderate effect on the estimated negative impact of poverty on growth. Depending on the control set used, a 10 percent increase in headcount poverty reduces growth prospects by between 0.7 and 1.3 percent.

- *The use of different poverty measures (headcount, poverty gap, squared poverty gap).* Changing the definition of poverty does affect the estimated coefficients of poverty, which are not comparable across definitions. However, the coefficients continue to be negative and statistically significant; in absolute value, the coefficients of the poverty gap and square poverty gap tend to be larger than the coefficient corresponding to the headcount definition.
- *The use of alternative estimation methods.* One of the problems dealing with highly persistent endogenous data is that the standard GMM estimation method based on internal instruments may not fully eliminate the potential reverse causality bias. To control for this problem, Lopez and Servén (2005b) also present results based on cross-sections that should not suffer from reverse causality (although admittedly they may suffer from fixed-effects bias). The results also confirm the negative impact of poverty on growth.
- *Adding inequality to the regression models.* When inequality is added to the empirical models, the sign, significance, and magnitude of the poverty effect decline somewhat in absolute value, and the estimate is less accurate. It remains highly significant, however, suggesting that the poverty variable does capture a true poverty effect rather than an inequality effect. This result is also robust to adding inequality and squared inequality to control for the likely nonlinear relation between poverty and inequality.

In principle, the finding that poverty lowers growth does not necessarily rule out the convergence of cross-national incomes (conditional convergence in this case) predicted by the neoclassical model, but the empirical estimates in Lopez and Servén (2005b) do imply the existence of a threshold poverty level beyond which divergence would occur. For example, with the baseline estimates in Lopez and Servén, there would be divergence for levels of the poverty headcount (with a $2-per-day poverty line) above 10 percent.

BOX 6.2

Is Latin America different?

Although the Lopez and Servén (2005b) results do not explicitly consider whether the impact of poverty on growth varies by geographic region, extending the model to test this possibility is relatively simple. In fact, we have reestimated their basic models to allow Latin American poverty levels to have an impact on growth that is different from the average for the group (that is, we are allowing the Latin American region to be "different"). The table below reports the results of this exercise.

This table suggests that Latin America may indeed be different. In particular, the estimates of the coefficients for Latin America are always negative and significant (in other words, poverty would reduce growth more in Latin America than in the typical country of the world). The magnitude of the Latin American dummy declines significantly in absolute value as the poverty line used in the computation of headcount poverty increases, from -0.23 under a $2-a-day poverty line to about -0.10 under a $4-a-day poverty line (although admittedly the standard error in the former case is also much larger than in the latter).

Poverty and growth: Is Latin America different?

	Poverty line					
	$2 a day		$3 a day		$4 a day	
	All	LAC dummy	All	LAC dummy	All	LAC dummy
Parameter	–0.114	–0.237	–0.128	–0.165	–0.140	–0.098
	(0.02)	(0.08)	(0.02)	(0.05)	(0.02)	(0.03)

Source: Authors' calculations.

Transmissions channels from poverty to growth

What are the channels through which poverty might influence growth? A quick review of the theoretical literature suggests a number of potential channels including investment, human capital (both education and health), innovation and mobility, and risk.

Poverty and investment

Several theoretical models on poverty traps exploit the result that high poverty levels (typically coupled with credit constraints) are likely to affect the investment rate negatively. But what do we actually know about the relationship between poverty and investment? Although the literature has paid significant attention to the impact of income levels on the investment rate (see, for example, Ben-David, 1995), little is known about the impact of poverty on investment. As a first pass at the issue, we ranked 99 countries for which we have income, poverty, and investment data according to their per capita income in the mid-1990s.[10] Then we partitioned these countries into 10 groups of 10 countries each (the last group has only 9 countries). The poorest countries in the sample are in the first group, the next poorest 10 countries are in the second group, and so on; thus the 9 richest countries form the tenth group.

For each group, figure 6.13 plots median (log) income in panel A, poverty ($2 poverty line) in panel B, and gross fixed capital formation relative to GDP in panel C.[11] Inspection of this figure reveals a clear nonlinear pattern in the relationship between income, poverty, and investment. For example, headcount poverty falls dramatically between the first and fourth groups—from about 66 percent to less than 8 percent, but after that it declines much more modestly as one moves up the income-group classification. Similarly, investment increases from 14 to about 22 percent of GDP between the first and fourth groups, and then remains virtually constant between the fourth and tenth groups. Note that these nonlinearities are not driven by the underlying income data (panel A), whose association with investment seems to be well described by a linear pattern.

The figure suggests a closer association between poverty and investment than between income levels and investment. In fact, the correlation coefficient between the income series in panel A and the investment series in panel C is about 0.55 (that is, investment tends to be higher in richer countries), whereas the correlation coefficient between the investment series and the poverty series shown in panel B is −0.77.

FIGURE 6.13
Income, poverty, and investment

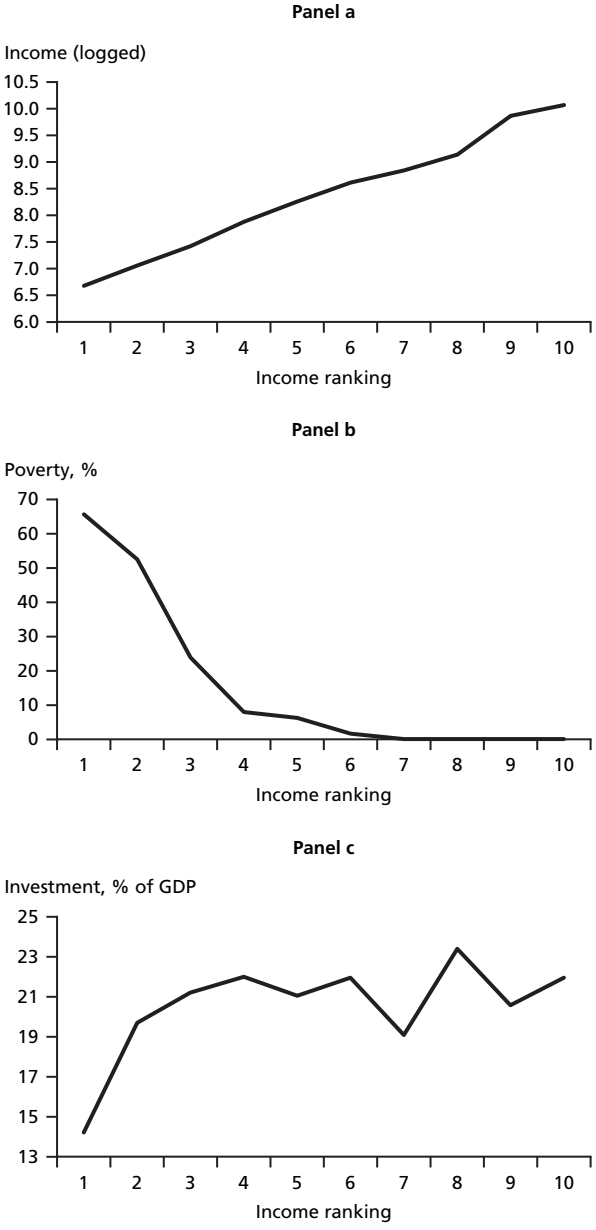

Source: López and Servén (2005b).

Note: The picture plots median income, headcount poverty ($2 poverty line), and investment (gross fixed capital formation as a percentage of GDP) by group of countries. Countries have been ranked by their income in the 1990s and then grouped in 10 groups of 10 countries each (except for the last group, which has 9 countries) for a total of 99 countries. The poorest countries would be in group 1 and the richest in group 10.

Does this apparent close association between poverty and investment withstand econometric scrutiny? Apparently

yes. Lopez and Servén (2005b) estimate the impact of poverty on investment using a simple accelerator model and find that, all else being equal, a 10 percent increase in poverty is likely to be associated with a decline in investment of about 6–8 percentage points. This result is robust to the use of different poverty lines and alternative measures of investment.

This finding suggests a potential explanation for poverty's negative effect on growth: a higher poverty rate leads to a lower investment rate, which leads to lower growth. In fact, when one econometrically explores the impact of poverty on growth controlling for investment, the investment rate turns out to belong to the growth equation, but poverty does not enter significantly in the various specifications (that is, the impact of poverty on growth is captured by the investment variable).[12]

The role of the financial sector

As noted above, theoretical models on poverty traps tend to exploit the joint impact of high poverty and credit constraints on growth. The basic idea is that poverty is likely to have a greater effect on investment when financial sector development is limited. Thus, one would expect to find that the impact of poverty on investment is affected by the degree of financial sector development.

Table 6.3 reports the results of estimating an empirical investment equation (based on the simple accelerator model) augmented with two variables aimed at capturing any potential difference in the effect of poverty on investment in countries with a highly developed financial sector ($Poverty^{HFD}$) and in those with a less developed financial sector ($Poverty^{LFD}$).[13] The results of this exercise indicate that, as expected, investment rates tend to be highly persistent, to

TABLE 6.3

Does financial sector development play a role in the poverty-investment interaction?

	GFCF			GCF		
Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Investment ($t-1$)	0.721	0.716	0.735	0.653	0.656	0.674
	(0.04)	(0.04)	(0.05)	(0.03)	(0.03)	(0.03)
Income (in logs) ($t-1$)	–0.005	–0.011	–0.010	–0.005	–0.006	–0.002
	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)
Growth (t)	0.524	0.507	0.498	0.620	0.616	0.612
	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)
PPP ($t-1$)	–0.004	0.001	–0.001	0.000	0.000	–0.001
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Terms of Trade (t)	0.079	0.089	0.100	0.071	0.078	0.079
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
$Poverty^{HFD}$ ($2) ($t-1$)	0.031			0.016		
	(0.03)			(0.03)		
$Poverty^{LFD}$ ($2) ($t-1$)	–0.055			–0.057		
	(0.03)			(0.02)		
$Poverty^{HFD}$ ($3) ($t-1$)		–0.002			0.011	
		(0.03)			(0.03)	
$Poverty^{LFD}$ ($3) ($t-1$)		–0.059			–0.038	
		(0.02)			(0.02)	
$Poverty^{HFD}$ ($4) ($t-1$)			0.003			0.025
			(0.03)			(0.03)
$Poverty^{LFD}$ ($4) ($t-1$)			–0.039			–0.010
			(0.03)			(0.03)

Source: Lopez and Servén (2005b), table 9.
Note: Numbers in parentheses are standard errors. The table reports the results of regressing investment on the variables in the first column. In columns 1, 2, and 3, we use the ratio of gross fixed capital formation (GFCF) to GDP as the measure of investment. In columns 4, 5, and 6, we use the ratio of gross capital formation (GCF) to GDP. PPP is a measure of the price of capital goods, and $Poverty^{LFD}$ and $Poverty^{HFD}$ are the poverty headcounts of countries with low and high financial sector development, respectively. The poverty line used for each variable is given in US$.

be procyclical, and to negatively depend on the cost of capital goods. Moreover, the impact of poverty on investment is more adverse in countries with less developed financial sectors. In fact, poverty does not seem to affect investment at high levels of financial sector development when credit constraints for the poor may not be so relevant.

These findings are consistent with those in Giuliano and Ruiz-Arranz (2005) who analyze the impact on investment and growth of foreign workers' remittances. Giuliano and Ruiz-Arranz find that remittances typically have a positive impact on investment but that this impact declines with the level of financial sector development. In other words, remittances seem to alleviate the credit constraints on the poor and through that channel contribute to capital accumulation and growth.

Poverty and education

There is a clear relationship between education and poverty reduction. Education has a very strong impact on earning potential, expands labor mobility, promotes the health of parents and children, and reduces fertility and child mortality. For example, the World Bank's 2005 poverty assessment for El Salvador (World Bank 2005) estimated that the per capita income of a household whose head had a primary education was 13 percent higher, on average, than that of a household with an uneducated head. The gain from a household head with a secondary school education was about 26 percent relative to a head with a primary school education, whereas the average gain from a household head having a university education was about 38 percent.

Similarly, the Bank's poverty assessment for Honduras in 2001 (World Bank 2001a) reported that in urban areas during the 1990s, workers with 7 years of school increased their labor income by 9 percent over workers with 6 years of school, whereas an increase from 15 to 16 years resulted in additional income of 14 percent. The income gains in rural areas from comparable improvements in schooling were estimated at 11 and 18 percent, respectively.

Education is also crucial to achieve sustained economic growth and hence sustained poverty reduction. As noted in chapter 5, human capital plays a central role in long-run growth. Education directly contributes to worker productivity and to more rapid technological adaptation and innovation. This point is particularly relevant for growth in Latin America because most Latin American countries have massive deficits in secondary enrollment (de Ferranti and others 2003). For the region the deficit is estimated at about 19 percent, but in some countries it is much higher. In Brazil, for example, the secondary school enrollment deficit is estimated at 36 percent and in República Bolivariana de Venezuela at 42 percent.

However, as discussed in detail in chapter 9, poverty may also affect education levels so that the relationship between poverty reduction and education is one of double causality. In table 6.4 we present the results of estimating a simple econometric model for the years of secondary schooling using cross-country data.[14] In addition to the lagged dependent variable, it includes among the explanatory variables the following indicators: per capita income to control for the country's level of development, the pupil-to-teacher ratio to capture quantity and quality efforts at the country level, and poverty (as measured by the headcount index using the $2-, $3-, and $4-a day poverty lines).

Table 6.4 shows that, as expected, secondary education is highly persistent. It also indicates that richer countries (as measured by per capita income levels) have more-educated populations, and that a lower quality of education (as measured by a higher pupil-to-teacher ratio) is associated with less-educated populations. Finally, higher poverty levels typically result in lower average years of secondary education.

On the whole, this discussion highlights the possibility that poverty and growth interact through the education

TABLE 6.4

Does poverty lead to lower secondary education?

Dependent variable is average years of secondary education			
Secondary education ($t-1$)	0.95	0.94	0.94
	(0.00)	(0.00)	(0.00)
Income	0.12	0.11	0.08
	(0.01)	(0.01)	(0.01)
Pupil/teacher ratio	–0.01	–0.01	–0.01
	(0.00)	(0.00)	(0.01)
Poverty ($2 a day)	–0.08		
	(0.04)		
Poverty ($3 a day)		–0.09	
		(0.03)	
Poverty ($4 a day)			–0.16
			(0.03)

Source: Authors' calculations.
Note: Numbers in parentheses are standard errors. The table reports the results of regressing the average years of secondary education on the variables. Although not reported here, the standard specification tests do not indicate any particular problem with the estimated model or the instruments used.

channel. As the literatures on both growth and microeconomic determinants of poverty stress, higher education levels result in higher growth and higher household income levels and therefore in lower poverty. At the same time, lower poverty levels feed back into the system and result in higher education, creating the potential for a virtuous circle between growth and poverty.

Poverty and health

Poorer countries have much worse health indicators than richer countries, most likely because of the bidirectional causality between income and health. On the one hand, empirical evidence indicates that higher income levels lead to better health indicators For example, Pritchett and Summers (1996) estimate that the long-run income elasticity of infant and child mortality in developing countries lies between 0.2 and 0.4. On the basis of those estimates, they calculate that more than 500,000 child deaths in the developing world in 1990 alone could be attributed to the poor economic performance in the 1980s.

On the other hand, there are a number of channels through which health can affect growth and income levels.

- *Productive efficiency.* Healthier workers are more productive. When health improves, more output can be produced with any given combination of skills, physical capital, and technological knowledge. One way to think about this effect is to take health as another component of human capital, analogous to the skill component.
- *Learning capacity.* Health plays an important role in determining the rate of return to education. Children who are well nourished and alert gain more from a given amount of education.
- *Creativity.* Just as a healthier person is more efficient in producing goods and services, so is the person likely to be more efficient in producing new ideas and hence in his or her ability to innovate (see also below).
- *Life expectancy.* Increases in life expectancy have a direct effect on the average skill level of the population. This is a consequence of two forces. When the probability of dying young is high, the discount rate is also high, making it optimal for people to start working early in their life and not to stay at school too long. Similarly, when life expectancy is short, the depreciation rate of human capital is high, making its accumulation less profitable.

For example, Fogel (1994) argues that nutrition and health have a significant influence on labor productivity and estimates that when labor is adjusted for intensity (measured by calories), improved gross nutrition explains about one-third of economic growth in the United Kingdom since 1800. Similarly, Boucekkine, de la Croix, and Licandro (2003) estimate that the observed improvements in adult mortality since the 18th century account for 70 percent of the growth acceleration that occurred before the industrial age. They argue that exogenous improvements in adult mortality between 1600 and 1800 increased individual incentives to build human capital and, as a consequence, investment in education rose, which in turn exerted a positive effect on economic growth.

Mayer-Foulkes (2001) has studied the long-term impact of health on economic growth in Latin America. Although he is unable to disentangle the relative contribution of such factors as nutrition and adult mortality, his results indicate that typical health improvements for adults may be associated with a permanent incremental increase in annual growth of between 0.8 and 1.5 percent. Thus poverty can also affect growth through the health channel. High poverty may result in worse health, which feeds back into lower growth, creating the possibility of a vicious circle.

Poverty and innovation

The discussion so far has suggested that poverty can hamper economic growth by choking an economy's ability to accumulate various forms of productive capital. Another potential link between poverty and growth exists, however, one that concerns an economy's ability to innovate and thus improve the productivity or efficiency of capital, labor, and other factors of production. Moreover, poverty's negative effect on capital accumulation can itself hamper innovation when capital investments are required to cover the costs of innovation. For instance, introducing new export products can require investments to understand market regulations and product standards, or simply to experiment with various business plans to achieve an efficient production process. Similarly, more sophisticated innovations with commercial value can be achieved only through investments in research and development. And both types of innovations can require at least a minimum amount of education. Consequently poverty, which is associated with low levels of human and physical capital, can be associated with lower levels of innovation at the national level (for a given level of national income per capita). In other words, poverty can effectively

limit the number of potential innovators, not because community members are not talented, but because poverty prevents them from undertaking the necessary investments to bring about economically meaningful innovations.

While the links between poverty and innovation remain understudied, and our understanding of the drivers of innovation and technical progress in general is quite modest, recent research by Klinger and Lederman (2005) sheds some light on this important issue. These authors studied the determinants of two types of innovations, namely, the introduction of new export products and patenting activity across countries and over time. This study reports the so-called marginal effects of population and poverty, and their interaction on the number of new products exported by a sample of 70 countries during 1994–2003. It also presents the same marginal effects, but for patenting activity during the 1980s and the 1990s. It is worth highlighting that these analyses controlled for numerous other variables that might also affect innovative activity.[15]

In any case, Klinger and Lederman find that the median (or typical) effect of poverty on export "discoveries" is about −0.02; for patenting activity, it is about −0.06. In other words, for each 1 percent increase in a country's poverty rate, the number of export innovations falls by 0.02 percent and the number of patents falls by 0.06 percent. Since the monetary value of exports and patents can be quite high, the economic consequences of poverty through these innovation channels should be worrisome. Perhaps more interesting, the empirical evidence also suggests that poverty affects innovation by affecting the number of potential innovators within a country. For both export discoveries and patenting, the effect of population size on innovation activity declines with poverty. A plausible explanation for this result is that poverty reduces the number of people with sufficient human and physical capital needed to produce innovation.

Poverty, mobility, and risk

According to de Ferranti and others (2000), volatility is considerably higher in all developing regions than in industrial economies. The less-diversified economies in lower-income countries, as well as limited access to external financing, expose these countries to higher risk and thus greater volatility. This then translates into higher volatility in aggregate wage measures and unemployment rates. Thus poverty seems to lead to higher risk.

At the same time, mobility through the income distribution may have impacts that promote growth. Hart (1981, 9), for example, argues that "it is mobility which provides the sticks for those who do not wish to move down the distribution and the carrots for those who wish to move up." More generally, the accumulation of human capital that is so critical to intergenerational mobility has effects on growth; a greater possibility for moving up the income ladder stimulates greater investment, which in turns leads to higher growth.

Mobility is also seen as an indicator of efficiency: high levels of income fluctuations may be seen as evidence that individuals are moving fluidly from one position to another, responding to changes in supply and demand for labor. Labor legislation that leads to segmented labor markets where certain classes of workers are therefore rationed out of good jobs, liquidity constraints that prevent individuals from migrating to more prosperous regions, or deficient financial markets that deny good entrepreneurs the resources they need to grow both restrict mobility and lead to poor allocation of resources. They can also be elements of poverty traps, which are explicitly about the inability of low-income groups to move up in the distribution.

However, chapter 2 argued that the unpredictable element of mobility constitutes risk that adversely affects welfare. For this reason, advanced societies have developed insurance and other mechanisms to reduce the risk that individuals and families face. Simulations that measure how risk-averse people are suggest that these welfare effects are large. In addition, a recent strand of the literature (Krebs 2003) argues that risk also has negative impacts on growth. As chapter 9 discusses, individuals' decisions to invest in education are strongly dependent on the perceived long-run gains in income. But like any other investment, the riskier the expected return, the less attractive it becomes. Cunha, Heckman, and Navarro (2005) argue that college attendance is lower than expected given the relatively high average return to education because roughly 40 percent of the observed variability in postcollege incomes is unpredictable: if individuals could make their decisions based on their actual incomes, 25 percent of high school graduates would rather be college graduates and 31 percent of college graduates would have stopped at high school. Hence, "uncertainty about future outcomes greatly affects schooling choices, and there is plenty of scope for ex-post regret," the three write (54). In countries where

workers face large shocks to their labor incomes, because of either frequent bouts of unemployment or high earnings volatility caused, perhaps, by inflation, or where frequent illness prevents working, the incentive to invest in education may fall even more. The resulting lower levels of education in turn dampen growth.

Here, then, is another example where two dimensions of poverty—health and risk—undercut growth, and the magnitudes appear large. Krebs, Krishna, and Maloney (2005) make an attempt to assess empirically the effect on human capital accumulation and growth of declines in the level of income risk of Argentina and Mexico to the U.S. levels. Their findings indicate that if Mexico could lower its labor market risk to Argentine levels, it could potentially increase its growth rate permanently by almost half a percentage point (table 6.5). The amount that growth would have to increase to increase the total welfare measure by an equivalent amount has two components. The first is the direct loss that is attributable to workers' and families' dislike of risk; this effect is worth the equivalent of a 0.59 percent permanent loss in yearly growth. The second component is the additional effect that arises because risk also makes workers and their families invest less in human capital; this has a direct impact on welfare of 0.48 percent. On the whole, the effect of lowering Mexico's risk to Argentine levels is equivalent to increasing growth by slightly more than 1 percent, a huge amount in a country where growth rates hover around 2 percent. If Argentina could reduce its risk to U.S. levels the effect would be less dramatic—growth would increase only about 0.2 percent—but still important over the long run.

These are only ballpark estimates. Clearly, the Mexican and Argentine economies are not identical to the U.S economy, and, more fundamentally, simple algebraic models cannot capture all the very subtle effects. Nonetheless, the exercise suggests that the magnitudes of effects arising from the presence of high risk in Latin America are large and that risk thus needs to be treated as an important dimension of an effective poverty reduction and growth strategy. Not only are policies to ameliorate risk beneficial from a pure vulnerability point of view, they may also be central to growth.

TABLE 6.5

The impact of risk on growth

Factor	United States	Argentina	Mexico
Income risk	0.15	0.18	0.21
Growth rate (%)	2.00	1.81	1.33
In education (%)	28.12	25.8	21.8
Direct loss due to risk (%)			0.59
Loss due to lower growth (%)			0.48
Total welfare loss (%)			1.07

Source: Krebs, Krishna, and Maloney (2005b) for Argentina and Mexico; Meghir and Pistaferri (2004) for United States.

Concluding remarks

This chapter explored the possible existence of links between growth and poverty reduction by which growth lowers poverty and lower poverty in turn contributes to faster growth. We reviewed several possible theoretical arguments that support the existence of such links. Among the most prominent are those arguments in the poverty-traps literature that suggest that the countries of the world are increasingly divided into two convergence clubs—the rich and the poor. Membership in the poor club is considered a huge handicap for growth and hence for poverty reduction.

The chapter then assessed the empirical evidence on this front and found mixed results. On the one hand, we presented evidence of convergence clubs in both absolute and relative income levels: richer countries converging toward the rich-club equilibrium, and poorer countries toward the poor-club equilibrium. By these measures, Latin America seems to be a homogeneous entity that is converging toward an equilibrium somewhere between the rich and the poor clubs. On the other hand, we also reviewed several empirical works that have formally tested whether the bimodality in the cross-national distribution of income is driven by poverty traps. In this regard, most, although not all, of the studies tend to reject the poverty-traps hypothesis.

Finally, we posed one simple question. Even if there is no evidence of poverty traps in the strict sense, is it still possible that poverty is a barrier to growth? We addressed this question from two different directions. First, we reviewed the empirical evidence contained in a background paper for this report, which found that countries with higher poverty levels tend to grow less than countries with lower poverty levels. The estimates presented in this chapter suggest that an additional 10 percentage points in the headcount poverty index cut growth prospects by about 1 percentage point. Second, we explored a number of potential channels through

which poverty might lead to lower growth. This evidence indicated that in countries with higher poverty rates, accumulation of both physical and human capital (education and health) is lower. Evidence also suggests that countries with higher poverty levels have lower rates of innovation (a critical contributor to growth) and higher risk.

It must be noted that in many of these channels the financial sector may play a very significant role, either by imposing a binding financial constraint on the poor that may prevent them from undertaking investments in human and physical capital or by preventing them from hedging against risk. Thus, the development and operation of the financial sector also appear to matter for the potential feedback effect from poverty to growth.

Overall, the results of this chapter suggest two main messages. First, the focus of the growth-poverty discussion needs to be shifted from the possible effects of growth on the poor (on which ample evidence has already been collected) to the relationships between growth and poverty. That shift in focus should mitigate the debate on whether development strategies should rely more on pro-growth or pro-poor policies, because strategies that do not take into account the bidirectional relation between poverty and growth will likely lead to disappointing results: poverty will not decline without growth, but growth will be difficult unless the constraints affecting the poor are also addressed. Second, at a more operational level, considering poverty and growth as part of the same problem suggests that the biggest payoff to growth (and hence to poverty reduction) is likely to result from policies that not only promote growth, but also exert an independent, direct impact on poverty—hence reducing the drag of poverty on growth.

Annex 6A

Convergence clubs and long-run equilibriums

One way to estimate the long-run per capita income equilibrium for each convergence club is based on the concept of β-convergence (see Barro and Sala-i-Martin 1995). This concept relies on the estimation of the following simple model:

$$(6A.1) \quad [\ln(Y_{1999}) - \ln(Y_{1960})]/39 = \mu + \beta \ln(Y_{1960}),$$

where Y denotes per capita income and the subscript refers to the year in question. Values of $\beta < 0$ would indicate convergence (β-convergence, to be more precise), and one could expect the countries in the group to cluster around the equilibrium values over the long run. In contrast, values of $\beta > 0$ would indicate divergence, and one would expect to observe that the dispersion in the cross-country distribution of per capita income increases as time goes by. Finally, for $\beta = 0$ there is neither convergence nor divergence. This simple model can be used to estimate the expected value of income over time when $\beta < 0$, which is given by $-\mu/\beta$.

The table below reports the results of estimating the previous model for the full sample of countries and for the three clubs discussed in the text (low-low, low-high, and high-high). The first noteworthy point is that, not surprisingly, in view of figure 6.7, the full sample presents divergence ($\beta > 0$). However, when we reestimate the model for each of the three clubs we obtain convergence, the point estimates of β are always negative (although admittedly for the high-high group, the estimate is not significant, which in turn may suggest that although there is no divergence, there may not be convergence either).

Convergence clubs

	Parameter		Equilibrium
Club	β	μ	US$
All	0.0033*	–0.007	Divergence
	(0.0017)	(0.014)	
Low-low	–0.0117*	0.087*	1,717
	(0.003)	(0.024)	
Low-high	–0.0178*	0.165*	10,600
	(0.0069)	(0.053)	
High-high	–0.006	0.07*	120,000
	(0.004)	(0.036)	

Source: Authors' calculations.
*Significant at the 5 percent level.

Convergence clubs and country transitions

To explore the distribution of income levels across countries, Quah (1993) takes each country's income level relative to the world average; allocates each observation to one of five states: 0–0.25, 0.25–0.5, 0.5–1, 1–2, and 2 and above (that is, the first state includes the poorest countries and the fifth state the richest); computes a transition matrix measuring the probability that a country in one state changes state by averaging the observed one-year transitions over every year from 1962 to 1984; and evaluates the long equilibrium consistent with the stationary distribution.

When we replicate all these calculations but use data for 1960–99, we obtain the following transition matrix:

$$M = \begin{Bmatrix} 0.987 & 0.013 & 0 & 0 & 0 \\ 0.038 & 0.935 & 0.026 & 0 & 0 \\ 0 & 0.033 & 0.936 & 0.031 & 0 \\ 0 & 0 & 0.032 & 0.954 & 0.014 \\ 0 & 0 & 0 & 0.009 & 0.991 \end{Bmatrix},$$

where a typical element m_{ij} measures the probability that a country in state i shifts to state j. So, for example, the probability that a country in the first state remains in its state is almost 99 percent, whereas the probability that it moves to the second state is about 1 percent. Similarly, the probabilities that a country in the second state remains in the same state, progresses to the third, and returns to the first state are 93 percent, 2.6 percent, and 3.8 percent, respectively; thus suggesting that the probability that an economy in state 2 falls behind is slightly larger than the probability of the same economy going ahead. This type of asymmetric behavior also applies to countries in state 3 and more markedly to those in state 4 where the probability of falling behind is more than double the probability of advancing.

Using the transition matrix M, it is now possible to compute the associated long-run equilibrium for the distribution of income levels by allowing the time horizon of the iterations to expand. This exercise results in the following equilibrium values for each of the five states under consideration: 0.43, 0.15, 0.12, 0.12, and 0.18.

Convergence clubs and country transitions in Latin America

The previous exercise can be replicated using data only for Latin America. The resulting transition matrix in this case is as follows:

$$M_{LAC} = \begin{Bmatrix} 0.875 & 0.125 & 0 & 0 & 0 \\ 0.02 & 0.928 & 0.052 & 0 & 0 \\ 0 & 0.036 & 0.948 & 0.016 & 0 \\ 0 & 0 & 0.055 & 0.945 & 0 \\ 0 & 0 & 0 & 0.154 & 0.846 \end{Bmatrix}.$$

There are at least two important differences between M_{LAC} and M. First, M displays more persistency in the first and fifth states than M_{LAC} does (the estimated persistency of states 2, 3, and 4 is very similar in both cases). Whereas the estimated probability that an economy in either state 1 or state 5 of the global sample continues in the same state is about 99 percent, the same probability for Latin America is estimated at 87 percent and 85 percent, respectively. Thus the Latin American region seems to display more mobility at the extremes of the distribution than does the global distribution: both getting out of extreme poverty and getting out of extreme richness seems easier in Latin America than in the rest of the world.

The second difference regards the probability of moving ahead for a Latin American country in state 3 or 4; that probability appears to be lower than it is in the rest of the world. In particular, a Latin American country in state 3 has about half the probability of moving to state 4 as do state 3 countries in the global sample (1.6 percent and 3.1 percent, respectively). More dramatically, the estimated probability of moving from state 4 to state 5 is nil in Latin America. These differences would result in a regional equilibrium given by 0.052, 0.33, 0.47, 0.14, and 0.

Estimating the impact of poverty on growth

The empirical strategy that Lopez and Servén (2005b) use to explore the links between poverty and growth in the data is based on the addition of a suitable measure of poverty to an otherwise standard empirical cross-nation growth regression:

$$(6A.2) \quad (y_{it} - y_{it-1}) = \delta y_{it-1} + \omega' x_{it} + \beta p_{it-1} + \nu_i + \upsilon_{it},$$

where y is the log of per capita income, p is a measure of poverty, x represents a set of control variables other than lagged income (discussed shortly), ν_i is a country-specific effect, and υ_{it} is an i.i.d. (independent and identically distributed) error term. However, several aspects of this empirical strategy require attention.

Estimation issues

Estimation of the previous equation poses two main challenges, namely, the presence of country-specific effects and the possible simultaneity of some of the explanatory variables with growth. These problems are addressed by using a GMM estimator (Arellano and Bover 1995 system estimator) that relies on internal instruments. Admittedly, with highly persistent instruments, that estimation method may not fully eliminate the potential bias related to reverse causality. To control for this problem, Lopez and Servén (2005b) also present results based on cross-sections, which should not suffer from reverse causality. In this regard, changing the estimation method does not dramatically affect the results.

Control variables

The empirical growth literature has experimented with so many alternative sets of explanatory variables that according to Durlauf and Quah (1999), by 1998 the number of individual regressors that had been considered as potential explanatory variables in growth regressions exceeded the number of countries in the standard growth data set. Rather than adding to the already huge variety of growth models, Lopez and Servén (2005b) use a baseline specification that relies on the controls used by Perotti (1996), Forbes (2000), Banerjee and Duflo (2003), and Knowles (2005). However, Lopez and Servén also experiment with two alternative sets to check whether the results are sensitive to changes in the controls. The basic finding is that changing controls does not significantly affect the estimated impact of poverty on growth.

Missing variables

The problem of missing variables is quite standard in this type of analysis. However, one variable in this context—inequality—needs particular attention. A relatively extensive literature already relates inequality and growth. For example, Alesina and Rodrik (1994) and Perotti (1996) find a negative relationship between inequality and growth on the basis of cross-section data, but Li and Zou (1998) and Forbes (2000) obtain the opposite result using aggregate panel data. Barro (2000) finds that inequality may affect growth in different directions depending on the country's level of income, while Banerjee and Duflo (2003) conclude that the response of growth to inequality changes has an inverted U-shape. Given the relation between inequality and poverty, excluding inequality from the equation could lead to the poverty variable capturing a pure inequality effect rather than a poverty effect. The empirical findings in this regard confirm that the estimated impact of poverty on growth does not result from poverty acting as a proxy for inequality either in a linear or in a nonlinear fashion.

Notes

1. Clearly, given the aversion of societies to high income inequality levels (see de Ferranti and others 2004), one could also justify the need to pay attention to distributional issues on the basis of political economy arguments.

2. By convergence club, we refer to a tendency of countries to converge to different equilibriums for per capita income levels. For example, Quah (1993), among others, finds evidence suggesting that the cross-country distribution of income may be well characterized by a twin-peak structure with poor and rich countries clustering around two different equilibriums.

3. Although not included in the sample, it is likely that Haiti also belongs to this group.

4. See Azariadis and Stachurski (2005) for a complete survey, and Lustig, Arias, and Rigolini (2003) for a nontechnical review.

5. For the purposes of this report, the industrial, or developed, countries group covers the OECD economies that are not eligible for lending from the International Bank for Reconstruction and Development. Figure 6.5 was constructed as follows. First, for each year we compute the median growth rate for all the countries in the relevant group for which the annual *World Development Indicators* report data. Then we apply a three-year, backward-moving average filter to smooth the series.

6. Admittedly, if the analysis were to take into account population weights, the story for the 1990s would be different: per capita growth would be approximately the same in both the developing and developed worlds. China and India account for much of this evening out, not only because they had almost 40 percent of the world's population during the 1990s, but also because India and especially China had excellent growth records. These differences are a reflection of the different ways in which economic performance can be measured. If individuals are the preferred unit of analysis, then weighted averages are probably more useful. If, instead, the unit of analysis is the country (as is the case when one focuses on country policies and country performance), then medians seem more appropriate.

7. Admittedly, it would be possible to argue that the 1960s distribution has two peaks: one around $3,000 and the other around $13,000.

8. For savings, Kraay and Radatz (2005) use a representative agent framework, something that rules out the possibility of credit market failure. In the Solow framework they use, the roles of jumps in saving and jumps in technology are more or less interchangeable.

9. Overall the results are backed by almost 90 robustness checks.

10. This approach is similar to that of Ben-David (1995) who focuses on the impact of income levels on investment. We pick the 1990s because it is the period over which more poverty observations are available.

11. The results remain virtually unchanged if one uses gross capital formation (GFC) as the investment measure.

12. This result is robust to the use of different measures of the investment rate.

13. Poverty^{HFD} is equal to the poverty headcount when the stock of credit to the private sector in the country/year in question is larger than the sample median and zero otherwise. Poverty^{LFD} equals the poverty headcount when the stock of credit to the private sector in the country/year in question is smaller than the sample median and 0 otherwise. Clearly, $\text{Poverty}^{HFD} + \text{Poverty}^{LFD} = \text{Poverty}$.

14. Estimation is performed using the GMM system estimator with internal instruments. This estimator therefore controls for unobserved fixed effects and potential endogeneity of the explanatory variables. The data are the same as in Lopez and Servén (2005b), except for the pupil-to-teacher ratio and expenditure in education, which come from the *World Development Indicators*.

15. Klinger and Lederman (2005) control for GDP per capita, export growth, population size, the sectoral concentration of innovation, past innovation activity, expenditures in research and development (in the case of patents granted by the U.S. Patent and Trademark Office), and exports to the United States (in the case of patents granted by the U.S. Patent and Trademark Office). These authors obtained similar results when using the share of the population with less than a high school education, but they were unable to differentiate between the effects of poverty on both human capital and physical capital reducing the effective share of the population capable of undertaking productive innovations.