attained, further expansion reduces growth. A possible interpretation is that, in extreme dictatorships, an increase in political rights tends to raise growth because the limitation on governmental authority is critical. However, in nations that have already achieved some political rights, further democratization may retard growth because of the heightened concern with social programs and income redistribution.

In contrast to the weak effect of democracy on growth, there is a strong positive linkage from prosperity to the propensity to experience democracy, a relation called the Lipset (1959) hypothesis. Various measures of the standard of living—real per capita GDP, life expectancy, and a smaller gap between male and female educational attainment—are found to predict democracy. Additional effects considered include urbanization, natural resources, country size, inequality, colonial history, and religious affiliation.

The final essay details the link between inflation–monetary policy and economic growth. The basic finding is that higher inflation goes along with a lower rate of economic growth. Moreover, the adverse effect of higher inflation on economic outcomes is quantitatively important. This pattern shows up clearly for inflation rates in excess of 15 to 20 percent annually but cannot be isolated statistically for the more moderate experiences. However, there is no evidence in any range of a positive relation between inflation and growth. The analysis also suggests that the estimates isolate the direction of causation from inflation to growth, rather than the reverse.

1 Economic Growth and Convergence

Neoclassical and Endogenous Growth Theories

In the 1960s, growth theory consisted mainly of the neoclassical model, as developed by Ramsey (1928), Solow (1956), Swan (1956), Cass (1965), and Koopmans (1965). One feature of this model, which has been exploited seriously as an empirical hypothesis only in recent years, is the convergence property: the lower the starting level of real per capita gross domestic product (GDP), the higher is the predicted growth rate.

If all economies were intrinsically the same except for their starting capital intensities, then convergence would apply in an absolute sense; that is, poor places would tend to grow faster per capita than rich ones. However, if economies differ in various respects—including propensities to save and have children, willingness to work, access to technology, and government policies—then the convergence force applies only in a conditional sense. The growth rate tends to be high if the starting per capita GDP is low in relation to its long-run or steady-state position, that is, if an economy begins far below its own target position. For example, a poor country that also has a low long-term position, possibly because its public policies are harmful or its saving rate is low, would not tend to grow rapidly.

The convergence property derives in the neoclassical model from the diminishing returns to capital. Economies that have less capital per worker (relative to their long-run capital per worker) tend to have higher rates of return and higher growth rates. The convergence is conditional because the steady-state levels of capital and output per worker depend in the neoclassical model on the propensity to save, the growth rate of population, and the position of the production function—characteristics that may vary across economies. Recent extensions of the model suggest the inclusion of additional sources of cross-country variation, especially government policies with respect to levels of consumption spending, protection of property rights, and distortions of domestic and international markets.

The concept of capital in the neoclassical model can be usefully broadened from physical goods to include human capital in the forms of education, experience, and health. (See Lucas 1988, Rebelo 1991, Caballe and Santos 1993, Mulligan and Sala-i-Martin 1993, and Barro and Sala-i-Martin 1995, chap. 5.) The economy tends toward a steady-state ratio of human to physical capital, but the ratio may depart from its long-run value in an initial state. The extent of this departure generally affects the rate at which per capita output approaches its steady-state value. For example, a country that starts with a high ratio of human to physical capital (perhaps because of a war that destroyed mainly physical capital) tends to grow rapidly because physical capital is more amenable than human capital to rapid expansion. A supporting force is that the adaptation of foreign technologies is facilitated by a large endowment of human capital (see Nelson and Phelps 1966, and Benhabib and Spiegel 1994). This element implies an interaction effect whereby a country's growth rate is more sensitive to its starting level of per capita output the greater is its initial stock of human capital.

Another prediction of the neoclassical model, even when extended to include human capital, is that in the absence of continuing improvements in technology, per capita growth eventually must cease. This prediction, which resembles those of Malthus (1798) and Ricardo (1817), comes from the assumption of diminishing returns to a broad concept of capital. The long-run data for many countries indicate, however, that positive rates of per capita growth can persist over a century or more and that these growth rates have no clear tendency to decline.

Growth theorists of the 1950s and 1960s recognized this modeling deficiency and usually patched it up by assuming that technological progress occurred in an unexplained (exogenous) manner. This device can reconcile the theory with a positive, possibly constant per capita growth rate in the long run, while retaining the prediction of conditional convergence. The obvious shortcoming, however, is that the long-run per capita growth rate is determined entirely by an element—the rate of technological progress—that comes from outside the model. (The long-run growth rate of the level of output depends also on the growth rate of population, another element that is exogenous in the standard theory.) Thus, we end up with a model of growth that explains everything but long-run growth, an obviously unsatisfactory situation.

Recent work on endogenous growth theory has sought to supply the missing explanation of long-run growth. In the main, this approach provides a theory of technical progress, one of the central missing elements of the neoclassical model. The inclusion of a theory of technological change in the neoclassical framework is difficult, however, because the standard competitive assumptions cannot be maintained. (These assumptions work fine in the framework of Frank Ramsey, David Cass, and Tjalling Koopmans.)

Technological advance involves the creation of new ideas, which are partially nonrival and therefore have aspects of public goods. For a given technology—that is, a given state of knowledge—it is reasonable to assume constant returns to scale in the standard, rival factors of production, such as raw labor, broad capital, and land. But then the returns to scale tend to be increasing if the nonrival ideas are included as factors of production. These increasing returns conflict with perfect competition. Moreover, the compensation of nonrival old ideas in accordance with their current marginal

cost of production—zero—will not provide the appropriate reward for the research effort that underlies the creation of new ideas.

Arrow (1962) and Sheshinski (1967) constructed models in which ideas were unintended by-products of production or investment, a mechanism described as learning by doing. In these models, each person's discoveries immediately spilled over to the entire economy, an instantaneous diffusion process that might be technically feasible because knowledge is nonrival. Romer (1986) showed later that the competitive framework can be retained in this case to determine an equilibrium rate of technological advance, but the resulting growth rate typically would not be Pareto optimal. More generally, the competitive framework breaks down if discoveries depend in part on purposive research and development (R&D) effort and if an individual's innovations spread only gradually to other producers. In this realistic setting, a decentralized theory of technological progress requires basic changes in the framework to incorporate elements of imperfect competition. These additions to the theory did not come until Romer's (1987, 1990) research in the late 1980s.

The initial wave of the new research—Romer (1986), Lucas (1988), and Rebelo (1991)—built on the work of Arrow (1962), Sheshinski (1967), and Uzawa (1965) and did not really introduce a theory of technological change. In these models, growth may go on indefinitely because the returns to investment in a broad class of capital goods, which includes human capital, do not necessarily diminish as economies develop. (This idea goes back to Knight 1944.) Spillovers of knowledge across producers and external benefits from human capital are parts of this process, but only because they help to avoid the tendency for diminishing returns to capital.

The incorporation of R&D theories and imperfect competition into the growth framework began with Romer (1987, 1990) and includes significant contributions by Aghion and Howitt (1992) and Grossman and Helpman (1991, chaps. 3, 4). Barro and Sala-i-Martin (1995, chaps. 6, 7) provide expositions and extensions of these models. In these settings, technological advance results from purposive R&D activity, and this activity is rewarded, along the lines of Schumpeter (1934), by some form of ex post monopoly power. If there is no tendency to run out of ideas, then growth rates can remain positive in the long run. The rate of growth and the underlying amount of inventive activity tend, however, not to be Pareto optimal because of distortions related to the creation of the new goods and methods of production. In these frameworks, the long-term growth rate depends on governmental actions, such as taxation, maintenance of law and order, provision of infrastructure services, protection of intellectual property rights, and regulation of international trade, financial markets, and other aspects of the economy. The government therefore has great potential for good or ill through its influence on the long-term rate of growth.

One shortcoming of the early versions of endogenous growth theories is that they no longer predicted conditional convergence. Since this behavior is a strong empirical regularity in the data for countries and regions, it was important to extend the new theories to restore the convergence property. One such extension involves the diffusion of technology (see Barro and Sala-i-Martin 1997). Whereas the analysis of discovery relates to the rate of technological progress in leading-edge economies, the study of diffusion pertains to the manner in which follower economies share by imitation in these advances. Since imitation tends to be cheaper than innovation, the diffusion models predict a form of conditional convergence that resembles the predictions of the neoclassical growth model. Therefore, this framework combines the long-run growth of the endogenous growth theories (from the discovery of ideas in the leading-edge economies) with the convergence behavior of the neoclassical growth model (from the gradual imitation by followers).

Endogenous growth theories that include the discovery of new ideas and methods of production are important for providing possible explanations for long-term growth. Yet the recent cross-country empirical work on growth has received more inspiration from the older, neoclassical model, as extended to include government policies, human capital, and the diffusion of technology. Theories of basic technological change seem most important for understanding why the world as a whole can continue to grow indefinitely in per capita terms. But these theories have less to do with the determination of relative rates of growth across countries, the key element studied in cross-country statistical analyses. The remainder of this chapter deals with the findings from this kind of cross-country empirical work.

Framework for the Analysis of Growth Across Countries

The framework for the determination of growth follows the extended version of the neoclassical model as already described. In equation form, the model can be represented as

$$Dy = f(y, y^*), \tag{1.1}$$

where Dy is the growth rate of per capita output, y is the current level of per capita output, and y^* is the long-run or steady-state level of per capita output.[1] The growth rate, Dy, is diminishing in y for given y^* and rising in y^* for given y. The target value y^* depends on an array of choice and environmental variables. The private sector's choices include saving rates, labor supply, and fertility rates, each of which depends on preferences and costs. The government's choices involve spending in various categories, tax rates, the extent of distortions of markets and business decisions, maintenance of the rule of law and property rights, and the degree of political freedom. Also relevant for an open economy is the terms of trade, typically given to a small country by external conditions.

For a given initial level of per capita output, y, an increase in the steady-state level, y^*, raises the per capita growth rate

over a transition interval. For example, if the government improves the climate for business activity—say, by reducing burdens from regulation, corruption, and taxation or by enhancing property rights—the growth rate increases for awhile. Similar effects arise if people decide to have fewer children or (at least in a closed economy) to save a larger fraction of their incomes.

In these cases, the increase in the target, y^*, translates into a transitional increase in the economy's growth rate. As output, y, rises, the workings of diminishing returns eventually restore the growth rate, Dy, to a value determined by the rate of technological progress. Since the transitions tend to be lengthy, the growth effects from shifts in government policy or private behavior persist for a long time.

For given values of the choice and environmental variables—and, hence, y^*—a higher starting level of per capita output, y, implies a lower per capita growth rate. This effect corresponds to conditional convergence. Note, however, that poor countries would not grow rapidly on average if they tend also to have low steady-state positions, y^*. In fact, a low level of y^* explains why a country would typically have a low observed value of y in some arbitrarily chosen initial period.

The last result shows that the framework can be reconciled with the now familiar lack of correlation between the growth rate and initial level of real per capita GDP across a large number of countries over the period 1960 to 1990.

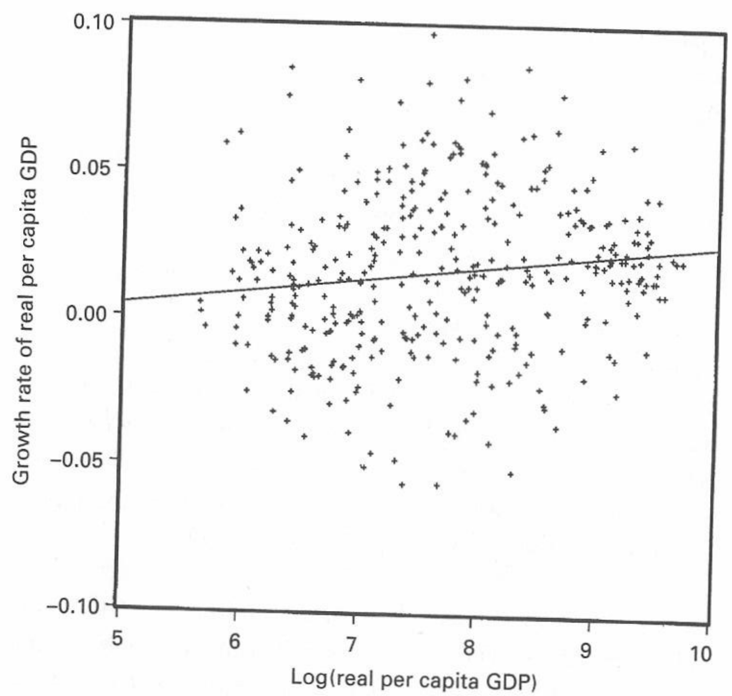

Figure 1.1
Simple correlation between growth and level of GDP

Figure 1.1 shows that this relationship is virtually nil.[2] (The slope actually has the wrong sign—slightly positive—but is not statistically significant.) The interpretation from the standpoint of the neoclassical model is that the initially poor countries, which show up closer to the origin along the horizontal axis, are not systematically far below their steady-state positions and therefore do not tend to grow relatively fast. The isolation of the convergence force requires

a conditioning on the determinants of the steady state, as in the cross-country empirical analysis discussed in the next section.

Even if convergence held in an absolute sense—that is, if y^* were identical across economies and the poorer places thus tended to grow faster—the dispersion of per capita product would not necessarily narrow over time. The evolution of dispersion, or inequality, depends on a weighing of the convergence force against effects from the shocks that impinge on each economy. These shocks, if independent across economies, tend to create dispersion and therefore work against the equalizing pressure from convergence.

The idea that the tendency for the poor to grow faster than the rich implies a negative trend in inequality is a fallacy; in fact, it is Galton's fallacy, as discussed in the growth context by Quah (1993) and Hart (1995). Galton's (1886, 1889, chap. 7) research indicated that the deviation of children's heights (and other physical and mental characteristics) from the mean of the population was positively correlated with the parents' deviation, but the amount of deviation tended to regress—or converge—toward zero. Nevertheless, the population's distribution of heights did not tend systematically to narrow over time.

The explanation of these facts is that a measure of the population's dispersion—say, the standard deviation of the log of height (or GDP)—would tend to adjust toward a long-run value that depends on the rapidity of the reversion to the mean (the rate of convergence) and the variance of random shocks to height (or GDP). If the determinants of the long-run distribution do not change, then dispersion would tend to rise or fall depending on whether it happened to start below or above its long-run value. Moreover, if the underlying determinants stay constant for a long time, then the observed distribution for a large population would remain fixed (despite the presence of the convergence tendency).

Empirically, for 114 countries with data, the standard deviation of the log of real per capita GDP rose from 0.89 in 1960 to 1.14 in 1990. This observation of increased inequality does not reject the convergence implications of the neoclassical growth model, partly because the predicted convergence is only conditional and partly because the poor tending to grow faster than the rich is not the same as a declining trend in inequality.

Empirical Findings on Growth Across Countries

Table 1.1 shows results from regressions that use the general framework of equation 1.1. The regressions apply to a panel of roughly one hundred countries observed from 1960 to 1990.[3] The dependent variables are the growth rates of real per capita GDP over three periods: 1965–1975, 1975–1985, and 1985–1990.[4] (The first period begins in 1965 rather than 1960, so that the 1960 value of real per capita GDP can be used as an instrument.) Henceforth, the term GDP will be used as a shorthand to refer to real per capita GDP.

Table 1.1
Regressions for per capita growth rate

Independent variable	(1)	(2)
Log(GDP)	−.0254 (.0031)	−.0225 (.0032)
Male secondary and higher schooling	.0118 (.0025)	.0098 (.0025)
Log(life expectancy)	.0423 (.0137)	.0418 (.0139)
Log(GDP) * male schooling	−.0062 (.0017)	−.0052 (.0017)
Log(fertility rate)	−.0161 (.0053)	−.0135 (.0053)
Government consumption ratio	−.136 (.026)	−.115 (.027)
Rule of law index	.0293 (.0054)	.0262 (.0055)
Terms of trade change	.137 (.030)	.127 (.030)
Democracy index	.090[a] (.027)	.094 (.027)
Democracy index squared	−.088 (.024)	−.091 (.024)
Inflation rate	−.043 (.008)	−.039 (.008)
Sub-Saharan Africa dummy		−.0042[b] (.0043)
Latin America dummy		−.0054 (.0032)
East Asia dummy		.0050 (.0041)
R^2	.58, .52, .42	.60, .52, .47
Number of observations	80, 87, 84	80, 87, 84

Table 1.1 (continued)

Notes: The system has three equations, where the dependent variables are the growth rate of real per capita GDP for 1965–1975, 1975–1985, and 1985–1990. The variables GDP (real per capita gross domestic product) and male schooling (years of attainment for the population aged twenty-five and over at the secondary and higher levels) refer to 1965, 1975, and 1985. Life expectancy at birth is for 1960–1964, 1970–1974, and 1980–1984. The variable log(GDP) ∗ male schooling is the product of log(GDP) (expressed as a deviation from the sample mean) and the male upper-level schooling variable (also expressed as a deviation from the sample mean). The rule of law index applies to the early 1980s (one observation for each country). The terms of trade variable is the growth rate over each period of the ratio of export to import prices. The inflation rate is the growth rate over each period of a consumer price index (or of the GDP deflator in a few cases). The other variables are measured as averages over each period. These variables are the log of the total fertility rate, the ratio of government consumption (exclusive of defense and education) to GDP, and the democracy index. Column 2 includes dummy variables for sub-Saharan Africa, Latin America, and East Asia. Individual constants (not shown) are also estimated for each period.

Estimation is by three-stage least-squares (with different instrumental variables used for each equation). The instruments include the five-year earlier value of log(GDP) (for example, for 1960 in the 1965–1975 equation); the actual values of the schooling, life expectancy, rule of law, and terms of trade variables; and, in column 2, the three area dummy variables. Additional instruments are earlier values of the other variables except the inflation rate. For example, the 1965–1975 equation uses the averages of the fertility rate and the government spending ratio for 1960–1964. Dummies for former colonies of Spain or Portugal and for former colonies of other countries aside from Britain and France are included as instruments. The instrument list also includes the cross product of the lagged value of log(GDP) (expressed as a deviation from the sample mean) with the male schooling variable (expressed as a deviation from the sample mean).

The estimation weights countries equally but allows for different error variances in each period and for correlation of these errors over time. The estimated correlation of the errors for column 1 is −0.13 between the 1965–1975 and 1975–1985 equations, 0.05 between the 1965–1975 and 1985–1990

Table 1.1 (continued)

(Notes, continued) equations, and 0.04 between the 1975–1985 and 1985–1990 equations. The pattern is similar for column 2. The estimates are virtually the same if the errors are assumed to be independent over the time periods. Standard errors of the coefficient estimates are shown in parentheses. The R^2 values and numbers of observations apply to each period individually.

[a] *p* value for joint significance of two democracy variables is 0.0006 in column 1 and 0.0004 in column 2.

[b] *p* value for joint significance of three dummy variables is 0.11.

Some previous analysis, such as Barro (1991), used a cross-sectional framework; that is, the growth rate and the explanatory variables were observed only once per country. The main reason to extend to a panel setup is to expand the sample information. Although the main evidence turns out to come from the cross-sectional (between-country) variation, the time-series (within-country) dimension provides some additional information. This information is greatest for variables that have varied a good deal over time within countries, such as the terms of trade and inflation.

The underlying theory relates to long-term growth, and the precise timing between growth and its determinants is not well specified at the high frequencies characteristic of business cycles. For example, relationships at the annual frequency would likely be dominated by mistiming and, hence, effectively by measurement error. In addition, many of the variables considered—such as fertility rates, life expectancy, and educational attainment—are not actually measured for many countries at periods finer than five or ten years. These considerations suggest a focus on the determination of growth rates over fairly long intervals. As a compromise with the quest for additional information, I settled on periods of five or ten years; specifically, growth rates were considered for 1965–1975 and 1975–1985 and for a final five-year period, 1985–1990. When the data through 1995 become available, the third period will be lengthened to 1985–1995.

The estimation uses an instrumental-variable technique, where some of the instruments are earlier values of the regressors. (The method is three-stage least squares, except that each equation contains a different set of instruments; see the notes to table 1.1 for details.) This approach may be satisfactory because the residuals from the growth rate equations are essentially uncorrelated across the periods. In any event, the regressions describe the relation between growth rates and prior values of the explanatory variables.

The regression shown in column 1 in table 1.1 includes explanatory variables that can be interpreted as initial values of state variables or as choice and environmental variables. The state variables include the initial level of GDP and measures of human capital in the forms of schooling and health. The GDP level reflects endowments of physical capital and natural resources (and also depends on effort and the unobserved level of technology). The choice and environmental variables are the fertility rate, government consumption spending, an index of the maintenance of the rule of law, the change in the terms of trade, an index of democracy (political rights), and the inflation rate. The roles of democracy and inflation will be discussed in the subsequent chapters.

Initial Level of GDP

For given values of the other explanatory variables, the neoclassical model predicts a negative coefficient on initial GDP, which enters in the system in logarithmic form.[5] The coefficient on the log of initial GDP has the interpretation of a conditional rate of convergence. If the other explanatory variables are held constant, then the economy tends to approach its long-run position at the rate indicated by the magnitude of the coefficient.[6] The estimated coefficient of −0.025 (s.e. = 0.003) is highly significant and implies a conditional rate of convergence of 2.5 percent per year.[7] The rate of convergence is slow in the sense that it would take the economy twenty-seven years to get halfway toward the steady-state level of output and eighty-nine years to get 90 percent of the way. Similarly slow rates of convergence have been found for regional data, such as the U.S. states, Canadian provinces, Japanese prefectures, and regions of the main Western European countries (see Barro and Sala-i-Martin 1995, chap. 11).

Figure 1.2 shows the partial relation between growth and the starting level of GDP, as implied by the regression from column 1 of table 1.1. The horizontal axis plots log(GDP) for 1965, 1975, and 1985 for the observations in the regression sample. The vertical axis shows the corresponding growth rate of GDP after filtering out the parts explained by all

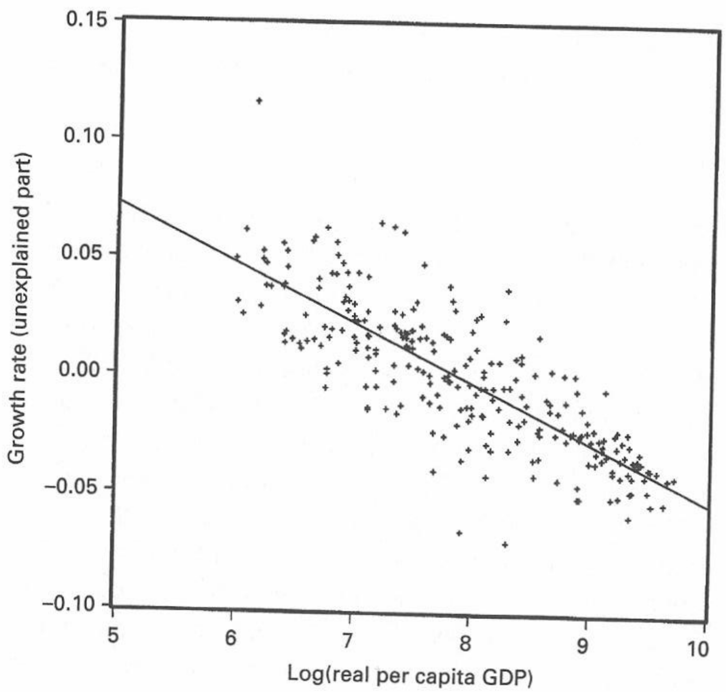

Figure 1.2
Growth rate versus level of GDP

explanatory variables other than log(GDP).[8] Thus, the negative slope shows the conditional convergence relation, that is, the effect of log(GDP) on the growth rate for given values of the other independent variables. In contrast to the lack of a simple correlation in figure 1.1, the conditional convergence relation in figure 1.2 is clearly defined in the graph. Also, the graph indicates that the relation is not driven by a few outliers and does not appear to be nonlinear.

Initial Level of Human Capital

Initial human capital appears in three variables in the system: average years of attainment for males aged twenty-five and over in secondary and higher schools at the start of each period, the log of life expectancy at birth at the start of each period (an indicator of health status),[9] and an interaction between the log of initial GDP and the years of male secondary and higher schooling. The data on years of schooling are updated and improved versions of the figures reported in Barro and Lee (1993).

The results show a significantly positive effect on growth from the years of schooling at the secondary and higher level for males aged twenty-five and over (0.0118 [0.0025]).[10] On impact, an extra year of male upper-level schooling is therefore estimated to raise the growth rate by a substantial 1.2 percentage points per year. (In 1990, the mean of the schooling variable was 1.9 years, with a standard deviation of 1.3 years.) The partial relation between the growth rate and the schooling variable—constructed analogously to the method described for log(GDP) in note 8—is shown in figure 1.3.

Male primary schooling (of persons aged twenty-five and over) has an insignificant effect if added to the system—the estimated coefficient is −0.0005 (0.0011)—whereas that on upper-level schooling remains similar to that found before (0.0119 [0.0025]). Thus, growth is predicted by male schooling at the upper levels but not at the primary level. Primary

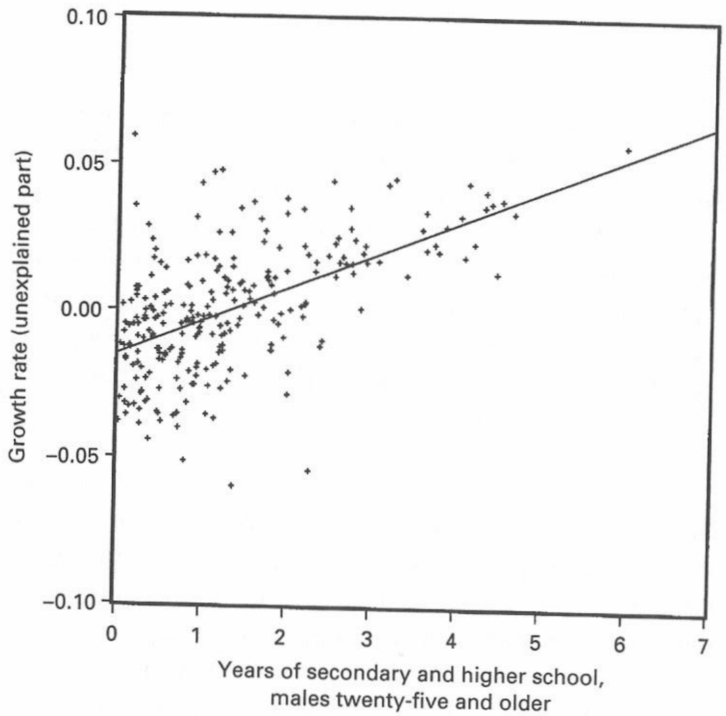

Figure 1.3
Growth rate versus male schooling

schooling, nevertheless, is indirectly growth enhancing because it is a prerequisite for training at the secondary and higher levels.

More surprising, female education at various levels is not significantly related to subsequent growth. For example, if years of schooling at the secondary and higher levels for females aged twenty-five and over is added to the system

shown in column 1 of table 1.1, then the estimated coefficient of this variable is −0.0023 (0.0046), whereas that for males remains significantly positive at 0.0132 (0.0036). For primary schooling of women aged twenty-five and over, the estimated coefficient is −0.0001 (0.0012), whereas that for men (twenty-five and over for secondary and higher schools) is 0.0118 (0.0025). Thus, these findings do not support the hypothesis that education of women is a key to economic growth.[11]

Some additional results indicate that female schooling is important for other indicators of economic development, such as fertility, infant mortality, and political freedom (see the next chapter). Specifically, female primary education has a strong negative relation with the fertility rate (see Schultz 1989, Behrman 1990, and Barro and Lee 1994). A reasonable inference from this relation is that female education would spur economic growth by lowering fertility, and this effect is not captured in the regressions shown in table 1.1 because the fertility rate is already held constant. If the fertility rate is omitted from the system, then the estimated coefficient on female primary schooling (the level of female schooling that affects fertility inversely) is 0.0012 (0.0012), which is positive but not significantly different from zero. Thus, there is only slight evidence that female education enhances economic growth through this indirect channel.

Returning to column 1 of table 1.1, the significantly negative estimated coefficient of the interaction term between male schooling and log(GDP), −0.0062 (0.0017), implies that more years of school raise the sensitivity of growth to the starting level of GDP. Starting from a position at the sample mean, an extra year of male upper-level schooling is estimated to raise the magnitude of the convergence coefficient from 0.026 to 0.032. This result supports theories that stress the positive effect of education on an economy's ability to absorb new technologies. The partial relation between the growth rate and the interaction variable appears in figure 1.4. (The points at the far right of the diagram are for the most developed countries, such as the United States, Canada, and Sweden, which have high values of GDP and schooling.)

The regression in column 1 of table 1.1 also reveals a significantly positive effect on growth from initial human capital in the form of health. The coefficient on the log of life expectancy is 0.042 (0.014). As an interpretation, it may be that life expectancy proxies not only for health status but more broadly for the quality of human capital. The partial relation between growth and life expectancy is shown in figure 1.5.

Fertility Rate

If the population is growing, then a portion of the economy's investment is used to provide capital for new workers rather than to raise capital per worker. For this reason, a higher rate of population growth has a negative effect on y^*, the steady-state level of output per effective worker in

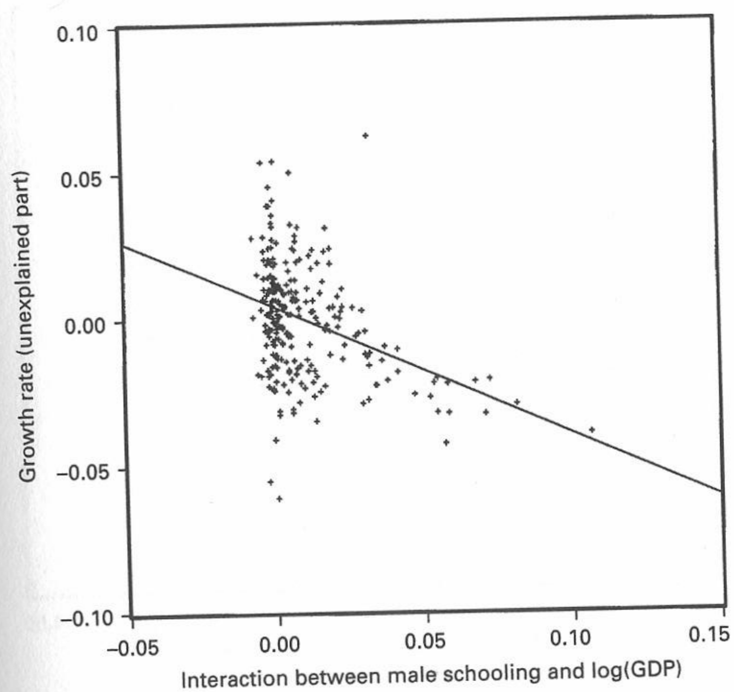

Figure 1.4
Growth rate versus interaction between schooling and level of GDP

the neoclassical growth model. A reinforcing effect is that a higher fertility rate means that increased resources must be devoted to child rearing rather than to production of goods (see Becker and Barro 1988). The regression in column 1 of table 1.1 shows a significantly negative coefficient, −0.016 (0.005), on the log of the total fertility rate. The partial relation between growth and fertility is in figure 1.6.

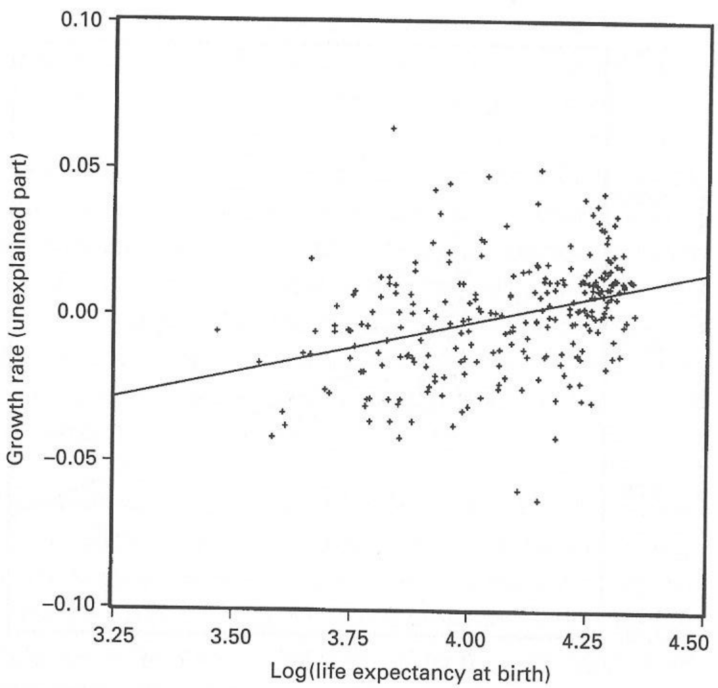

Figure 1.5
Growth rate versus life expectancy

Fertility decisions are surely endogenous; previous research has shown that fertility typically declines with measures of prosperity, especially female primary education (see Schultz 1989, Behrman 1990, and Barro and Lee 1994). The estimated coefficient of the fertility rate in the growth regression shows the response to higher fertility for given values of male schooling, life expectancy, GDP, and so on. Since the average of the fertility rate over the preceding five years

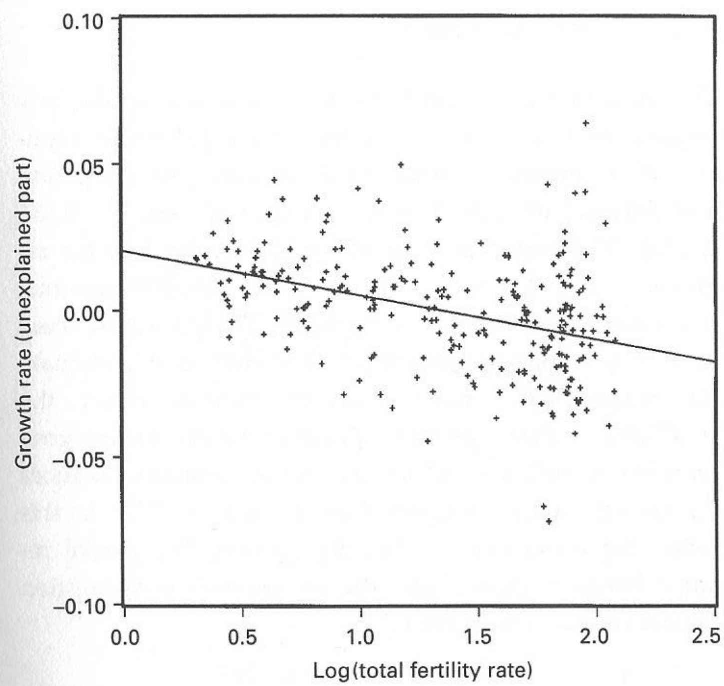

Figure 1.6
Growth rate versus fertility rate

is used as an instrument, the coefficient likely reflects the impact of fertility on growth, rather than vice versa. (In any event, the reverse effect would involve the level of GDP rather than its growth rate.) Thus, although population growth cannot be characterized as the most important element in economic progress, the results do suggest that an exogenous drop in birthrates would raise the growth rate of per capita output.

Government Consumption

The regression in column 1 of table 1.1 shows a significantly negative effect on growth from the ratio of government consumption (measured exclusive of spending on education and defense) to GDP. The estimated coefficient is −0.136 (0.026). (The period average of the ratio enters into the regression, and the average of the ratio over the previous five years is used as an instrument.) The particular measure of government spending is intended to approximate the outlays that do not improve productivity. Hence, the conclusion is that a greater volume of nonproductive government spending—and the associated taxation—reduces the growth rate for a given starting value of GDP. In this sense, big government is bad for growth. The partial relation between growth and the government consumption variable appears in figure 1.7.

The Rule of Law Index

Knack and Keefer (1995) discuss a variety of subjective country indexes prepared for fee-paying international investors and distributed as the *International Country Risk Guide*. (The various time series cover 1982 to 1995 and are available from Political Risk Services of Syracuse, New York.) The concepts covered include quality of the bureaucracy, political corruption, likelihood of government repudiation of contracts, risk of government expropriation, and overall maintenance of the rule of law. The general idea is

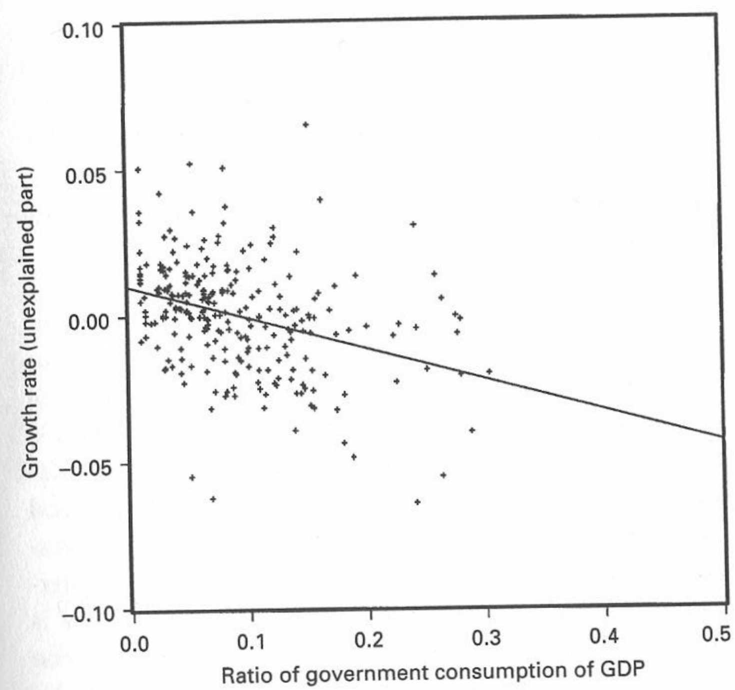

Figure 1.7
Growth rate versus government consumption ratio

to gauge the attractiveness of a country's investment climate by considering the effectiveness of law enforcement, the sanctity of contracts, and the state of other influences on the security of property rights. Although these data are subjective, they have the virtue of being prepared contemporaneously by local experts. Moreover, the willingness of customers to pay substantial fees for this information is perhaps some testament to their validity.

Among the various series available, the indicator for overall maintenance of the rule of law seemed a priori to be most relevant for investment and growth. This indicator was initially measured in seven categories on a 0 to 6 scale, with 6 the most favorable. The scale has been revised here to 0 to 1, with 0 indicating the worst maintenance of the rule of law and 1 the best.

The rule of law variable (observed, because of lack of earlier data, only once for each country in the early 1980s) was included in the regression system reported in column 1 of table 1.1 and has a significantly positive coefficient, 0.0293 (0.0054). (The other measures of investment risk, including political corruption and various indicators of political instability, are insignificant in these kinds of growth regressions if the rule of law index is also included.) The interpretation is that greater maintenance of the rule of law is favorable to growth. Specifically, an improvement by one rank in the underlying index (corresponding to a rise by 0.167 in the rule of law variable) is estimated to raise the growth rate on impact by 0.5 percentage point. The partial relation between growth and the rule of law index is in figure 1.8. (Note that only seven values for the index are observed.)

Terms of Trade

Changes in the terms of trade have often been stressed as important influences on developing countries, which typically specialize their exports in a few primary products. The

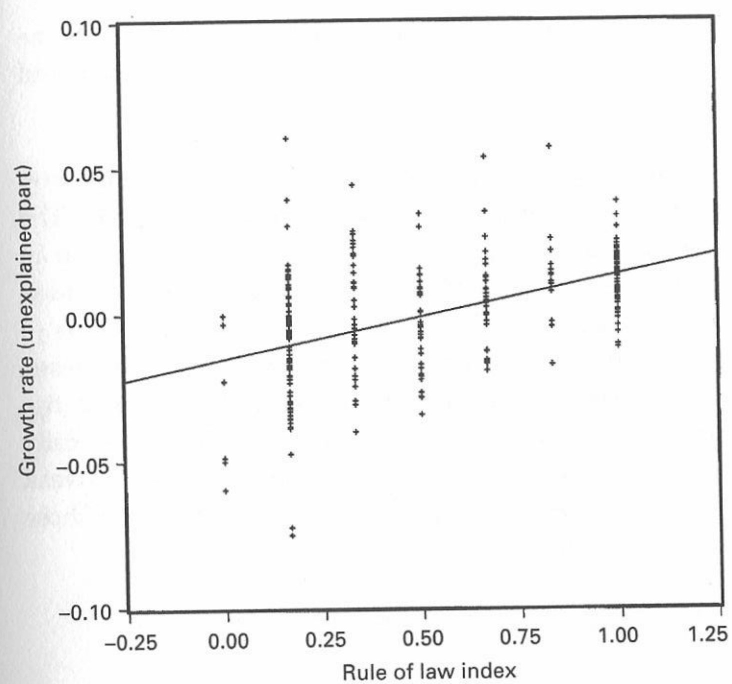

Figure 1.8
Growth rate versus rule of law index

effect of a change in the terms of trade, measured as the ratio of export to import prices, on GDP is, however, not mechanical. If the physical quantities of goods produced domestically do not change, then an improvement in the terms of trade raises real domestic income and probably consumption but would not affect real GDP. Movements in real GDP occur only if the shift in the terms of trade stimulates a change in domestic employment and output. For example, an oil-importing country might react to an increase in the relative price of oil by cutting back on its employment and production.

The result in column 1 of table 1.1 shows a significantly positive coefficient on the terms of trade: 0.14 (0.03). (The change in the terms of trade is regarded as exogenous to an individual country's growth rate and is therefore included as an instrument.) Thus, an improvement in the terms of trade apparently does stimulate an expansion of domestic output. The partial relation with growth appears in figure 1.9. Although the terms of trade variable is statistically significant, it turns out not to be the key element in the weak growth performance of many poor countries, such as those in sub-Saharan Africa.

Regional Variables

It has often been observed that recent rates of economic growth have been surprisingly low in sub-Saharan Africa and Latin America and surprisingly high in East Asia. For 1975–1985, the mean per capita growth rate for all 124 countries with data was 1.0 percent, compared with −0.3 percent in 43 sub-Saharan African countries, −0.1 percent in 24 Latin American countries, and 3.7 percent in 12 East Asian countries. For 1985–1990, the average growth rate was again 1.0 percent (for 129 places), compared with 0.1 percent in 40 sub-Saharan African countries, 0.4 percent in 29 Latin American countries, and 4.0 percent in 15 East Asian countries. An important question is whether these regions

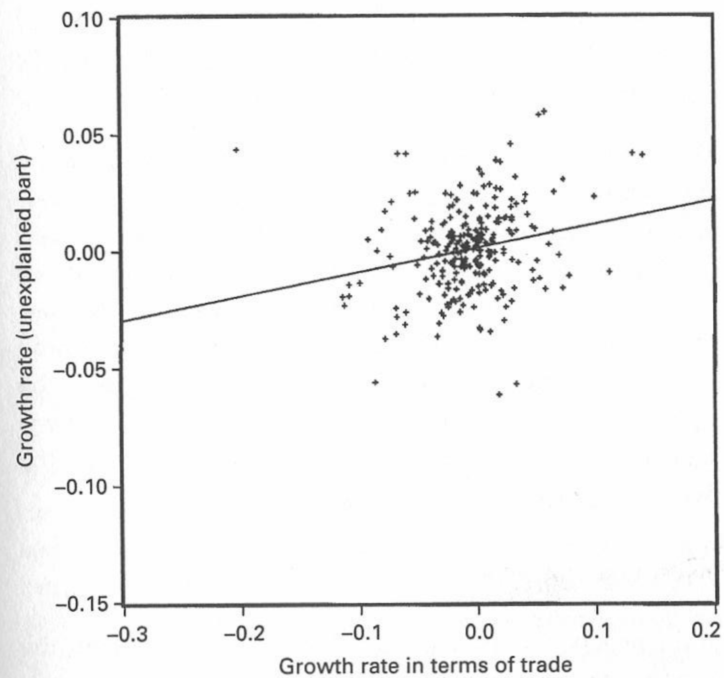

Figure 1.9
Growth rate versus change in terms of trade

continue to look like outliers once the explanatory variables considered in table 1.1 have been taken into account.

In some previous cross-country regression studies, such as Barro (1991), dummy variables for sub-Saharan Africa and Latin America were found to enter negatively and significantly into growth regressions. However, column 2 of table 1.1 shows that dummies for these two areas and also for East Asia are individually insignificant. (The p value for joint significance of the three dummy variables is 0.11.) Thus, the unusual growth experiences of these three regions are mostly accounted for by the explanatory variables.

The inclusion of the inflation rate is critical for eliminating the significance of the Latin America dummy (this interaction is discussed in the next chapter). The Latin America dummy also becomes significant if the fertility rate or the government consumption ratio is omitted. In the case of sub-Saharan Africa, the government consumption ratio is the only individual variable whose omission causes the dummy to become significant. For East Asia, the dummy is significant if male schooling, the rule of law indicator, or the democracy variables are deleted.

Investment Ratio

In the neoclassical growth model for a closed economy, the saving rate is exogenous and equal to the ratio of investment to output. A higher saving rate raises the steady-state level of output per effective worker and thereby raises the growth rate for a given starting value of GDP. Some empirical studies of cross-country growth have also reported an important positive role for the investment ratio (see, for example, DeLong and Summers 1991 and Mankiw, Romer, and Weil 1992).

Reverse causation is, however, likely to be important here. A positive coefficient on the contemporaneous investment ratio in a growth regression may reflect the positive rela-

tion between growth opportunities and investment rather than the positive effect of an exogenously higher investment ratio on the growth rate. This reverse effect is especially likely to apply for open economies. Even if cross-country differences in saving ratios are exogenous with respect to growth, the decision to invest domestically rather than abroad would reflect the domestic prospects for returns on investment, which would relate to the domestic opportunities for growth.

The system from column 1 of table 1.1 has been expanded to include the period-average investment ratio as an explanatory variable. If the instrument list includes the investment ratio over the previous five years but not the contemporaneous value, then the estimated coefficient on the investment variable is positive but not statistically significant: 0.027 (0.021). In contrast, the estimated coefficient is almost twice as high and statistically significant if the contemporaneous investment ratio is included as an instrument: 0.043 (0.018). These findings suggest that much of the positive estimated effect of the investment ratio on growth in typical cross-country regressions reflects the reverse relation between growth prospects and investment. Blömstrom, Lipsey, and Zejan (1993) reach similar conclusions in their study of investment and growth.

To interpret these results further, table 1.2 shows regression systems in which the dependent variables are the average ratios of investment to GDP for 1965–1974, 1975–1984, and 1985–1989. The independent variables (aside from the

Table 1.2
Regressions for investment ratio

Independent variable	(1)	(2)
Log(GDP)	−.010 (.011)	−.005 (.011)
Male secondary and higher schooling	−.0032 (.0088)	−.0064 (.0085)
Log(life expectancy)	.259 (.050)	.274 (.051)
Log(GDP) * male schooling	−.0004 (.0057)	.0009 (.0055)
Log(fertility rate)	−.0028 (.0192)	.0056 (.0186)
Government consumption ratio	−.264 (.089)	−.216 (.087)
Rule of law index	.092 (.023)	.074 (.024)
Terms of trade change	.074 (.068)	.070 (.064)
Democracy index	.148 (.069)	.168 (.070)
Democracy index squared	−.142 (.061)	−.153 (.062)
Inflation rate	−.053 (.022)	−.036 (.021)
Sub-Saharan Africa dummy		−.013 [a] (.019)
Latin America dummy		−.038 (.014)
East Asia dummy		.010 (.017)
R^2	.59, .62, .61	.60, .65, .67
Number of observations	80, 87, 84	80, 87, 84

Table 1.2 (continued)

Notes: The systems correspond to those described in table 1.1, except that the dependent variables are now the average ratios of real investment (private plus public) to real GDP over the periods 1965–1974, 1975–1984, and 1985–1989. The correlation of the errors across the equations is substantial in the systems for investment. For example, for column 1, the correlation between the first and second periods is 0.53, that between the first and third periods is 0.35, and that between the second and third periods is 0.62.

[a] p value for joint significance of three dummy variables is 0.03.

investment ratio) are the same as those used in table 1.1. The key finding in column 1 of table 1.2 is that a number of the variables that are found to enhance the growth rate in table 1.1 also appear as stimulants to investment. In particular, the investment ratio is positively related to life expectancy (a proxy for the quality of human capital) and the rule of law index and negatively related to the government consumption ratio and the inflation rate. The investment ratio also follows the same sort of quadratic relation with democracy that showed up for the growth rate. The effects of democracy are explored in the next chapter.

A reasonable interpretation of the results is that some policy variables—such as better maintenance of the rule of law, lower government consumption, and price stability—encourage economic growth partly by stimulating investment. However, if investment is higher for given values of the policy instruments—perhaps because of variations in thriftiness across economies that lack perfect capital mobility—then the positive effect on growth is weak, as indicated by the estimated coefficient of 0.027 (0.021) on the investment ratio.

Cross-Country Regressions and Country Fixed Effects

A comparison of figures 1.1 and 1.2 shows that it is critical to hold fixed the determinants of the long-run target value, y^*, in equation 1.1 to isolate the conditional convergence force—that is, the effect of initial GDP, y, on the growth rate, Dy, for a given y^*. Since y and y^* tend to be positively correlated, the estimated coefficient on y would be biased upward if y^* were not held constant. Since the true coefficient on y is negative, the omission of y^* tends to generate an underestimate of the rate of convergence, possibly even to the extent of estimating divergence (a positive coefficient on y) rather than convergence. Hence, the omission of y^* can account for the incorrect (positive) sign in the simple relation between Dy and y shown in figure 1.1.

One remaining problem in figure 1.2 is that the estimated rate of convergence would still tend to be underestimated if the measures used to hold fixed y^* were imperfect (as they must be). Specifically, underestimation of the convergence rate would tend to apply if the omitted determinants of y^* were still positively correlated with y after holding fixed the variables included to measure y^*. It is hard to get a direct assessment of the magnitude of this problem, although the isolation of y^*—like variables that have a lot of explanatory power for growth—as highlighted in the previous discussion—should lessen the error.

Some researchers prefer to handle this type of estimation problem by allowing for an unobserved fixed effect for each

country (see Knight, Loayza, and Villanueva 1993; Islam 1995; and Caselli, Esquivel, and Lefort 1996). Usually this treatment is applied by first differencing all variables in order to eliminate the fixed effect. This procedure works if the underlying determinants of y^*, such as government policies and preferences about saving and fertility, do not vary over time within a country. In practice, problems would still exist because unobserved shifts in y^* could still be correlated with the movements in y.

The main drawback of the fixed-effects technique is that it relies on time-series information within countries; that is, it eliminates the cross-sectional information, which is the principal strength of the broad cross-country data. Aside from losing information and, hence, precision, first differencing of the data tends to emphasize measurement error over signal. In particular, the estimation becomes more sensitive to incorrect timing in the relation between growth and its determinants.

If first differences of the determinants of y^* are retained in the estimation, then measurement error tends to bias toward zero the estimated coefficients of these variables. For the estimated coefficient of Dy on y, one should consider a regression of y (log[GDP]) on its own lag. Measurement error tends to bias this value toward zero and leads accordingly to an overestimate of the rate of convergence.

Column 1 of table 1.3 shows the results from estimation of a first-differenced version of the system from column 1 of

Table 1.3
Results from first differences and a cross-section

Independent variable	(1) First difference	(2) Cross-section	(3) Panel	(4) p value
Log(GDP)	−.0444 (.0066)	−.0220 (.0041)	−.0242 (.0028)	.000
Male schooling	−.0032 (.0045)	.0141 (.0030)	.0123 (.0023)	.68
Log(life expectancy)	−.0820 (.0381)	.0172 (.0184)	.0388 (.0124)	.002
Log(GDP) ∗ male schooling	.0052 (.0035)	−.0077 (.0019)	−.0070 (.0015)	.18
Log(fertility rate)	−.0396 (.0116)	−.0206 (.0066)	−.0156 (.0049)	.11
Government consumption ratio	.000 (.048)	−.114 (.026)	−.110 (.021)	.024
Rule of law index		.0294 (.0066)	.0300 (.0051)	
Terms of trade change	.102 (.027)	.078 (.078)	.129 (.029)	.92
Democracy index	.019 (.029)	.071 (.026)	.048 (.019)	.51
Democracy index squared	−.014 (.026)	−.074 (.023)	−.051 (.016)	.28
Inflation rate	−.032 (.005)	−.030 (.006)	−.028 (.004)	.12
R^2	.29, .44	.76	.56, .53, .49	
Number of observations	88, 91	80	83, 88, 84	

Notes: The systems are variants of the one shown in column 1 of table 1.1. Column 1 uses first differences of all variables and is estimated by the seemingly unrelated (SUR) technique, which allows for a different error variance for the two periods and correlation of the errors across the

Table 1.3 (continued)

(Notes, continued) periods. Column 2 uses means of all variables and is estimated by ordinary least squares (OLS). Column 3 is the same as column 1 of table 1.1, except that estimation is by SUR rather than three-stage least squares. The p values in column 4 refer to Wald tests of equality of the coefficients from columns 1 and 2.

table 1.1. This setup includes two equations. In the first, the dependent variable is the growth rate of GDP from 1975 to 1985, less that from 1965 to 1975. In the second, the dependent variable is the growth rate from 1985 to 1990, less that from 1975 to 1985. Similarly, the independent variables are first differences of the variables that appear in column 1 of table 1.1, for example, the first equation contains log(GDP) for 1975 less log(GDP) for 1965. The system is estimated in a seemingly unrelated (SUR) framework, which allows for correlation of the errors across the two equations. (Since the residuals from the growth rate equations in table 1.1 were essentially uncorrelated across the time periods, the residuals for the two equations in column 1 of table 1.3 have a strong negative correlation.)

Column 2 of table 1.3 shows the results from ordinary least squares (OLS) estimation of a pure cross-section, which contains one observation for each country. In this case, the dependent and independent variables are means over the three time periods of the variables used in column 1 of table 1.1.

Finally, column 3 of table 1.3 is the same as column 1 of table 1.1, except that the estimation is by the SUR technique instead of instrumental variables. This setup is basically a weighted combination of the time-series information from column 1 of table 1.3 with the cross-sectional information from column 2 of the table. In the main, these estimates are close to those shown in column 1 of table 1.1. The principal differences from the use of instruments show up in the estimated coefficients of the democracy and inflation variables.

If one compares the estimated coefficients from the first-difference specification with those from the cross section, then the biggest discrepancy is in the estimated convergence rate: −0.044 (0.007) in column 1 versus −0.022 (0.004) in column 2. The hypothesis of equality for these coefficients is rejected by a Wald test with a p value of 0.000 (see column 4 of the table). For the other independent variables, the only cases in which the estimated coefficients from the two specifications differ significantly at the 5 percent level (when variables are considered one at a time) are those for life expectancy and government consumption. However, a joint test of equality for all ten pairs of coefficients rejects decisively.

The standard errors of the coefficients in columns 1 and 2 indicate the information available from the time-series and cross-sectional dimensions of the panel data. For many of the variables—log(GDP), male schooling, log(life expectancy), the interaction between log(GDP) and male schooling, log(fertility rate), and the government consumption ratio—the standard errors are much smaller in column 2 than in column 1. This pattern indicates that the

cross-country (between) variation in these independent variables is much more informative than the time-series (within-country) variation. The extreme situation is for the rule of law variable, which has no time-series dimension (as currently measured) and therefore effectively has an infinite standard error in the first-difference form. The only case in which the standard error is noticeably smaller in column 1 is for the terms of trade; the variations here relate more to changes over time than to differences across countries. For democracy and inflation, the standard errors are similar in the two contexts.

Many researchers seem to prefer the results from variants of first-difference specifications, as in column 1 of table 1.3, because of their concern with the possible bias from correlated fixed effects. The high estimated convergence coefficient from this column—4.4 percent per year—is similar to that reported from more sophisticated but related techniques by Knight, Loayza, and Villanueva (1993, p. 529), Islam (1995, tables 3, 4); and Caselli, Esquivel, and Lefort (1996, tables 3, 4). However, the higher magnitude of these convergence coefficients, relative to those found from the panel estimation in column 3 of table 1.3 or column 1 of table 1.1, may reflect an increase in the relative amount of measurement error from the exclusion of the cross-sectional information. That is, instead of eliminating the fixed-effects bias (which tends to underestimate the convergence rate), the first-difference procedure may mainly exaggerate the measurement error bias (which tends to overestimate the convergence rate).

The results in column 1 also show that it is hard to isolate effects from the explanatory variables other than lagged GDP in a pure time-series context. The only estimated coefficients that are significant at the 5 percent critical level are those for the fertility rate, the terms of trade, and the inflation rate. Life expectancy is marginally significant with the wrong sign. One reason for these findings is that the time series offers little variation in many of the variables. In addition, the model likely misspecifies the timing between growth and its determinants, and this error is much more important for time-series estimation than in a cross section.

Undoubtedly the confidence in the results would be greater if the estimated coefficients from first-difference and cross-sectional forms did not differ significantly. Improvements in specification—for example, with regard to the lag structure between growth and its determinants—may produce more uniform results, but at this stage, there seems to be no basis for preferring the first-difference estimates to the cross-sectional ones. I have focused on panel results—column 3 of table 1.3 or column 1 of table 1.1—as a weighing of these two imperfect sources of information, where the weights are determined (by means of the SUR or three-stage least-squares procedures) from the relative informativeness of the two sources.

Growth Projections

The results from column 1 of table 1.1 can be used to construct long-term forecasts of economic growth for individ-

ual countries. These predictions have been constructed by using recent observations of the explanatory variables: GDP in 1994 (or sometimes earlier), schooling in 1990, life expectancy and fertility in 1993, consumer price index (CPI) inflation through 1993 or 1994, the rule of law indicator for 1995, the democracy index for 1994, and government consumption in the late 1980s.[12] Table 1.4 shows the twenty predicted best and worst performers from 1996 to 2000 out of the eighty-six countries that have the necessary data to make these projections.[13] There is, however, a substantial margin of error (as much as two percentage points) in the prediction for an individual country.

For all eighty-six countries, the average forecast of per capita growth is 2.4 percent per year. The breakdown by region is 3.7 percent for 18 Asian countries, 2.9 percent for twenty-two Latin American countries, 2.4 percent for twenty-one Organization for Economic Cooperation and Development (OECD) countries (not including Japan, Turkey, and Mexico), and 0.5 percent for eighteen sub-Saharan African countries.

It is no surprise that many old and new tigers of East Asia are forecasted to grow rapidly; South Korea, Malaysia, Singapore, Thailand, Hong Kong, and Taiwan are on the high-growth list. (Japan falls short with 3.2 percent growth.) The unexpected finding is the presence in the high-growth group of Asian laggards of the past: the Philippines, India, Sri Lanka, and Pakistan. (China and Vietnam would likely also appear but are excluded because of missing data.)

Table 1.4
Winners and losers for prospective economic growth

Top 20 prospects		Bottom 20 prospects	
Country	Predicted growth rate of real per capita GDP 1996–2000 (% per year)	Country	Predicted growth rate of real per capita GDP 1996–2000 (% per year)
South Korea	6.2	Sierra Leone	−3.6
Philippines	5.6	Sudan	−2.7
Dominican Republic	5.4	Malawi	−0.2
		Bangladesh	−0.2
India	5.3	Niger	−0.1
Poland	5.2	Zaire	−0.1
Peru	5.2	Gambia	0.1
Sri Lanka	5.0	Botswana	0.1
Malaysia	5.0	Senegal	0.2
Argentina	4.7	Papua New Guinea	0.2
Singapore	4.6		
Thailand	4.6	Brazil	0.2
Greece	4.6	Congo	0.3
Chile	4.3	Algeria	0.3
Paraguay	4.2	Zambia	0.5
Hong Kong	4.2	Mali	0.8
Guyana	4.2	Nicaragua	0.8
Pakistan	3.9	Cameroon	1.1
Taiwan	3.8	Trinidad	1.2
Ecuador	3.8	Costa Rica	1.3
Egypt	3.8	Uganda	1.3

Regional patterns:

All countries (86)	2.4
Sub-Saharan Africa (18)	0.5
Latin America (22)	2.9
Asia (18)	3.7
OECD (21)	2.4

South Korea places at the top with 6.2 percent growth because it has high educational attainment, strong rule of law, low government spending, low fertility, high investment, and low inflation. Although their underlying growth determinants are less favorable, the Philippines, India, and Sri Lanka place nearly as high in projected growth rates because their levels of per capita GDP are only one-eighth to one-quarter as large as South Korea's. These are cases in which the convergence force generates rapid growth.

The high-growth list also has substantial representation in South America: Peru, Argentina, Chile, Paraguay, Guyana, and Ecuador. A key assumption here is that the recently achieved macroeconomic stability, as reflected in relatively low inflation rates, will be maintained. As a contrast, Brazil appears on the low-growth list with roughly zero per capita growth. Aside from low school attainment, a major element is projected inflation of around 50 percent.

In Central Europe, posttransformation Poland appears as a prospective fast grower, and Hungary (with 3.5 percent projected growth) just misses the list. Other countries, such as the Czech Republic, would likely have appeared but are excluded because of lack of data.

On the low-growth list, thirteen of the twenty countries are in sub-Saharan Africa. (Other countries, such as Nigeria, Rwanda, and Somalia, would likely have been included if not for their missing data.) Sierra Leone, as a prototype, has weak rule of law, low school attainment, high fertility, low life expectancy, no political freedom, high government consumption, moderately high inflation, and virtually no investment. Being poor, which Sierra Leone and the other African countries surely are, is not enough to generate high growth.

Among OECD countries, the only place on the high-growth list is Greece. (Spain comes close with 3.8 percent growth.) Many of the advanced economies nearly made the low-growth list: Denmark at 1.3 percent, Norway at 1.4 percent, the United States at 1.4 percent, Sweden at 1.7 percent, Finland at 1.9 percent, the United Kingdom at 2.0 percent, Canada at 2.0 percent, Germany at 2.1 percent, Italy at 2.2 percent, and France at 2.4 percent. (Note that the growth rate of the level of GDP adds the growth rate of population—roughly 1 percent per year for the United States and smaller amounts in Western Europe.)

One can also use the results to ask, somewhat more speculatively, whether some changes in institutions or policies could move the United States, the United Kingdom, or another advanced country to the high-growth list, that is, raise the long-term per capita growth rate from $1\frac{1}{2}$ to 2 percent to around 4 percent. Unfortunately, the answer seems to be no. The institutions and policies in the advanced countries are already reasonably good (despite possible excesses of transfer programs and regulations), and long-term per capita growth much above 2 percent seems to be incompatible with the prosperity that has already been attained.

It would probably be feasible to raise the long-term growth rate by a few tenths of a percentage point by cutting tax rates and nonproductive government spending or by eliminating harmful regulations. (Some of these variables may be important but could not be measured in the cross-country empirical work discussed above.) Moreover, increases in growth rates by a few tenths of a percentage point matter a lot in the long run and are surely worth the trouble. On the negative side, it would be possible to lower the growth rate by a few tenths of a percentage point by moving away from price stability or interfering further with free markets. There is no evidence that increases in infrastructure investment, research subsidies, or educational spending would help a lot. Basically, 2 percent per capita growth seems to be about as good as it gets in the long run for a country that is already rich.