

CHAPTER 4

Convergence and Divergence

4.1 Introduction

We now come to a central prediction of the Solow growth model, one that is a direct corollary of the assumption of diminishing returns. It states that two countries that are the same in all their parameters — savings rates, population growth rates, rates of technical progress, and so on — must ultimately exhibit similar levels of per capita income. For capital per efficiency unit of labor must converge to a steady state value that's common to both countries. Indeed, this will happen *irrespective of the initial state of each of these economies*, as measured by their starting levels of per capita income (or equivalently, their per capita capital stock).

Does this sound trivial to you, or totally wild? The trivial camp would say: after all, we are assuming that all long-run parameters are similar. Under such an assumption, how could we expect anything other than convergence? The wild camp would counter: the assertion is far from obvious. What are assumed to be the same are exogenous parameters of the model, but *not* the initial level of the capital stock or per capita income. The claim of convergence is then based on the analysis that we conducted for the Solow model: its content is that in the face of similar parameters governing the evolution of the economy, “history” in the sense of different initial conditions does not matter.

For what it's worth, such a prediction is *not* made by the Harrod-Domar model. It is easy to verify (and you should take the time to do so) that in that scenario, two countries that have some initial discrepancy in per capita incomes will maintain that relative discrepancy for all time, as long as they have the same underlying parameters. Constant or increasing returns to capital will not imply convergence — a theme that we will return to — but




FIGURE 4.1. Unconditional Convergence.

diminishing returns does. The attraction of the convergence prediction, apart from its rosy optimistic outlook on life, is that it relies on diminishing returns to capital, which is not at all a crazy hypothesis, given that labor is an equally fundamental factor of production.

The convergence prediction relies on both diminishing returns and the assumption that all parameters are the same across countries. It is an uncompromising prediction, built on strong assumptions. We therefore give it an equally uncompromising name: *unconditional convergence*. Soon we shall introduce a more accommodating sibling, *conditional convergence*, which will permit the cross-country parameters to be different.

Figure 4.1 illustrates unconditional convergence. It plots the logarithm of income against time, so that a constant rate of growth appears as a straight line. The line *AB* plots the time path of (log) per capita income in steady state, where income per efficiency unit of labor is precisely at the level generated by \hat{k}^* . The path *CD* represents a country that starts below the steady-state level per efficiency unit. According to the Solow model, this country will initially display a rate of growth that *exceeds* the steady-state level, and its time path of (log) per capita income will move asymptotically toward the *AB* line as shown. Over time, its growth rate will decelerate to the steady-state level. Likewise, a country that starts off above the steady state, say at *E*, will experience a *lower* rate of growth, because its time path

EF of (log) income flattens out to converge to the line AB from above. At any rate, that is what the hypothesis has to say.

Unconditional convergence, then, is indicated by a strong negative relationship between the initial value of per capita income and subsequent growth rates of per capita income. Is that prediction borne out by the data?

4.2 Unconditional Convergence: Evidence or Lack Thereof

4.2.1 Time Horizons. The first problem that arises when testing a hypothesis of this sort is the issue of time horizons. Ideally, we'd like to go back a century or more in history, but the systematic collection of data in developing economies is a modern phenomenon. There are just two choices: cover a relatively small number of countries over a large period of time or cover a large number of countries over a short period of time. That we can even contemplate the former is due to Angus Maddison, who has attempted to compile data on several countries back to the mid-19th century (and earlier). The Maddison Project — see Maddison (1982, 1991, 2007) — is extended and continued today at the Maddison Project webpage (see <http://http://www.ggdc.net/maddison/>) by his colleagues. The problem is, of course, that the number of countries with data stretching back into the nineteenth century continues to be small.

4.2.2 The Baumol Study. That last sentence held true *a fortiori* in 1986, when William Baumol published one of the first studies of long-run convergence. At the time, there were just sixteen countries in Maddison's database for which "reliable" estimates of per-capita income existed (and I put "reliable" in quotes because this sort of historical detective work must always be taken with a large pinch of salt). These were, in order of poorest to richest in 1870: Japan, Finland, Sweden, Norway, Germany, Italy, Austria, France, Canada, Denmark, the United States, the Netherlands, Switzerland, Belgium, the United Kingdom, and Australia. They are among the richest countries in the world today.

Figure 4.2 illustrates the exercise that Baumol conducted. It plots 1870 per capita income for these sixteen countries on the horizontal axis, and the *growth rate* of that income over 1870–1979 (measured by the difference in the logs of per capita income over this period) on the vertical axis. The convergence of these sixteen countries to one another, starting from widely different levels of per capita income in 1870, is unmistakable. It appears,




FIGURE 4.2. Growth and Per Capita Income for Baumol's Sixteen Countries.

then, that Baumol's finding supports the unconditional convergence hypothesis quite strongly.¹

4.2.3 Selection Bias. Unfortunately, there's a classic statistical pitfall lurking both in the picture and in the study. The sixteen countries are the first to have historical records for good reason: they are rich countries today! Yet in 1870 they were all over the economic map. Japan is a perfect case in point. It is there precisely because of hindsight: Japan is rich today, but in 1870, it was probably midway in the world's hierarchy of nations arranged by per capita income. If Japan, why not Argentina or Chile or Portugal?²

¹Baumol regressed the log difference in per capita income between 1870 and 1979 on the logarithm of 1870 per capita income and a constant. A slope coefficient of -0.995 and an R^2 of 0.88 was obtained. A slope close to -1 means that by 1979, almost all the initial gaps in per capita income had been erased.

²Suppose you were to look at today's successful basketball stars. They would come from a variety of backgrounds; some poor and some rich. You could say that they "converged" to success, and in fact the rags-to-riches stories that we see often in the media bolsters such unwarranted perceptions. However, from this you cannot predict that a randomly chosen sample of kids who *aspire* to be basketball players will all "converge"! Hindsight is no substitute for prediction. One of my favorite documentaries, *Hoop Dreams*, indeed takes

Therefore, the “convergence” that Baumol found is a result of a statistical bias rather than any underlying real tendency of convergence.³ A *true* test of convergence would have to look at a set of countries with no *ex ante* selection at all. Or, to be a bit more relaxed about it, one would choose a set of countries that, *ex ante*, seemed *likely* to converge to the high per capita GDP levels that came to characterize the richest nations several decades later.⁴ And indeed, a host of countries (many in Europe, and some in Latin America)—and left out of Baumol’s study—were very much in the same position as many of the countries included in Maddison’s set of sixteen, in terms of income levels and economic promise (as perceived *then*). Does the evidence on convergence hold up in a rigorous statistical test if we broaden the set of countries in the manner suggested above?

De Long (1988) addressed this question by adding, to Maddison’s sixteen, seven other countries, which in 1870 had as much claim to membership in the “convergence club” as many included in Baumol’s original set. These additional countries are Argentina, Chile, East Germany, Ireland, New Zealand, Portugal, and Spain. Three of these countries (New Zealand, Argentina, and Chile) figure in the list of top ten recipients of British and French overseas investment (in per capita terms) as late as 1913. Investors’ faith in these economies reflects a very favorable perception regarding their growth prospects at that point in time. *All* the new countries included had per capita GDP levels in 1870 higher than Finland, which was the second lowest in Baumol’s sample.⁵

Figure 4.3 shows the modified pictures after De Long’s countries are added and Japan is dropped from the initial sixteen. The earlier observations from Figure 4.2 appear as unlabeled dots. Now matters don’t look so good for convergence, and indeed, De Long’s statistical analysis confirms this gloomier story. Baumol’s original regression can now be repeated on this new data set: regress the log difference in per capita income between 1870 and 1979 on the logarithm of 1870 per capita income and a constant.⁶ The slope coefficient of the regression is still appreciably negative, but the

this bolder step, studying a group of aspiring basketball professionals *ex ante* rather than *ex post*.

³Although we use this study as an illustrative warning, Baumol himself did see the problem immediately after a number of scholars, notably De Long (1988), pointed out the error.

⁴This likelihood may be determined on the basis of the extent of their existing integration into the world economy, as well as their per capita income levels then.

⁵The lowest is that of Japan and it is significantly below the rest. De Long dropped the data on Japan in his analysis, because its inclusion would necessitate the inclusion of several other countries at comparable income levels in 1870. But accurate 1870 data for such countries do not exist.

⁶For a discussion of regression and associated concepts, see Appendix 2.




FIGURE 4.3. Growth and per capita income for the twenty-two countries studied by De Long. New countries are shown as squares; original countries as diamonds. Source: De Long [1988].

“goodness-of-fit” is very bad, as indicated by the fact that the residual disturbance term is very large.

De Long also argued, and correctly so, that the 1870 data are likely to contain large measurement errors (relative to those in 1979), which make the various observations more scattered than they actually should be and makes any measurement of convergence more inflated than the case actually merits.⁷ De Long repeated his regression exercise, assuming a stipulated degree of measurement error in the 1870 data and making necessary amendments to his estimation technique to allow for this, and found that the slope coefficient comes out to be very close to zero—indicating that there is very little systematic relationship between a country’s growth rate and its per capita GDP, at least in the cross section of the twenty-two countries studied.

A few decades have passed since De Long wrote his critique, and now we have many more countries — over 60 of them — including several Latin

⁷Imagine that a set of families all have the same incomes in 1960 as well as in 1990. They are surveyed in both these years. However, the survey in 1960 is inaccurate, so there are errors of measurement, whereas the survey in 1990 is accurate, showing that they all have the same income. Then these families will (statistically) appear to have “converged” to the same income from different starting levels.




FIGURE 4.4. Growth and per capita income, 1870–2010, for 60 countries from the Maddison Project. Source: the extended dataset described in Bolt and Van Zanden (2013). Original Maddison 16 shown by filled diamonds.

American countries and Asian countries such as India or China, for which estimates of per-capita income in 1870 are available. The data continue to be overwhelmingly absent for Africa, so the selection problem hasn't gone away by any means. Yet, despite that, there is no sign of convergence. Figure 4.4 makes this clear by running the same exercise for the updated Maddison dataset. The filled data points represent Baumol's original 16, with their now-familiar downward sloping signature of convergence. But as you can see, that convergence is nowhere to be found among the larger database depicted in the figure. We must conclude that the convergence in the Baumol study is an artificial consequence of the following points: (a) the original 16 are all rich today, so (b) that largely explains why they are in the original dataset to begin with, but (c) in 1870 they were all over the economic map. "Convergence" is a consequence of (a)–(c), but it's an artifact.

4.2.4 A Shorter Time Horizon. A second option is to include a very large set of countries to test for unconditional convergence. This approach has the advantage of "smoothing out" possible statistical irregularities in looking at a small sample. The disadvantage is that the time span of analysis must be shortened to a few decades, which is the span over which reliable data are available for a larger group of countries.

In Chapter 2, we looked at the world distribution of income over the period 1970–2010. If you turn back to that chapter and reread the discussion there, it should be clear that unconditional convergence sounds like a pretty long shot. At the very least, the gap between the richest and the poorest countries does not seem to have appreciably narrowed. The richest 10% of countries have a GDP per capita that's approximately four times the world average, while the poorest 10% earn about 7% of that same average. This discrepancy does not seem to have altered in any significant way over 1970–2010; if anything, it's grown. (This is certainly not to say that the poorest countries have not moved up in some absolute sense.)

Of course, a good sample should be broad-based; that is, focused not merely on the richest and the poorest in the sample. Indeed, in Chapter 2, we did note how diverse the experiences of different countries have been over this period. We observed, moreover, that if we group countries into different clusters and then construct a mobility matrix to track their movements from one cluster to another, there is little tendency for countries to move toward a common cluster.

Several studies support a similar finding. For instance, Parente and Prescott [1993] study 102 countries over the period 1960–85. In this study, each country's per capita real GDP is expressed in relative terms: as a fraction of U.S. per capita GDP for the same year. The authors then calculate the standard deviation⁸ of these values separately for each year. Whereas the convergence hypothesis says that countries move closer to each other in income levels, we expect the standard deviation of their relative incomes to fall over time. In Parente and Prescott's study, however, it actually *increased* by 18.5% over the twenty-six year period, and the increase was fairly uniform from year to year. However, there is some variation here if we look at geographical subgroupings. The standard deviation in relative incomes for Western European countries shows a clear decline. In fact this decline persists through the period 1913–85. On the other hand, the same measure applied to Asian countries displays a significant and pronounced increase, and the divergence in this region is consistent with data going all the way back to 1900.

To put the data together in yet another way, suppose that we compare average per capita growth between 1970 and 2011 with per capita GDP in 1970. Just as in the Baumol study, the tendency toward level convergence

⁸The standard deviation of a set of observations is a statistical measure that indicates how closely bunched or how spread out in value the set of observations is. A higher measure of the standard deviation means a higher "spread" or dispersion around the average. See Chapter ?? and Appendix 2 for more detail.

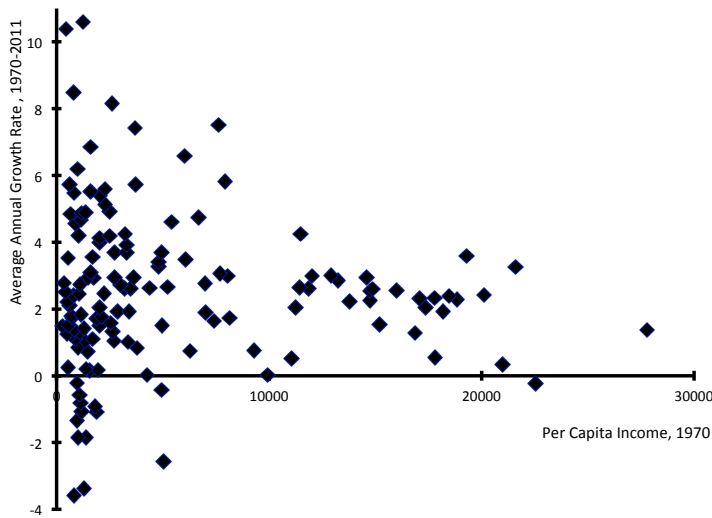


FIGURE 4.5. Per capita GDP (1970) Versus Subsequent Annual Growth for 140 Countries, 1970–2011. Source: Penn World Tables v8.0.

would show up in a negative relationship between these two variables.⁹ But in line with the discussion in the rest of this section, there is no such tendency at all. Indeed, the correlation coefficient¹⁰ between these two variables is just -0.07 , which is pretty much zero (Barro (1991) obtains a similarly small number, 0.09 , for the period 1960–1985). Figure 4.5 uses the Penn World Tables to plot per capita income in 1970 (in 2005 USD) against the average annual growth rate between 1970 and 2011. The lack of a pattern in the data needs no comment.

4.2.5 Unconditional Convergence: A Summary. Our understanding so far has been negative, but it is useful all the same. Recall that the Harrod–Domar model, in its simplest form, predicts the “neutrality” of growth rates with respect to per capita income. Because of the assumption that returns to capital are *constant*, parameters such as the savings rate have effects on the growth rate. The Solow model, on the other hand, demotes the savings rate to a parameter that only has effects on the *level* of per capita income. An ongoing increase in per capita capital (and therefore in per capita output) contains the seeds of a subsequent slowdown, because the economy runs into diminishing returns to capital. Thus, provided

⁹Note that in considering the entire data set, we avoid the ex post bias discussed earlier.

¹⁰For a definition of this term, see Appendix 2.

that all the parameters of the Solow model are constant across countries, convergence across countries is an implication of that model.

But this prediction of the Solow model in its simplest, strongest form appears to be clearly rejected by the data. Our task, then, is to examine several directions of reconciliation between the theory and the data, and, in so doing, to learn new ways to look at the data.

In closing this section, let's take care of a natural objection. We could just say that the Harrod-Domar model was right after all. It predicts the neutrality of growth rates with respect to per capita income. By and large, that is what we seem to get, so why study the Solow model? The main rejoinder to this point of view is that the assumption of constant returns to *physical* capital alone is simplistic in a way that does not seem to fit the facts. Physical capital does need labor to operate: we are far from a world of pure automation. There is little doubt that in the absence of other inputs, or technical progress, a straightforward accumulation of machines would fail to lead to corresponding increases in output. Capital and labor do go hand in hand. Thus we may think of the Solow model as being absolutely correct in postulating diminishing returns to each input separately, but at the same time lacking in some important dimension that enriches the story and does not predict convergence.

xx Insert box on Divergence Big Time?

4.3 Conditional Convergence

4.3.1 Countries Have Different Parameters. The prediction of unconditional convergence takes us out on a rather long limb. It makes the assumption that across all countries, the level of technical knowledge (and its change), the rate of savings, the rate of population growth, and the rate of depreciation are all the same. This notion certainly flies in the face of the facts: countries differ in many, if not all, these features. Although this has no effect on the Solow prediction that countries must converge to *their* steady states, the steady states can now be different from country to country, so that there is no need for two countries to converge to *each other*. This weaker hypothesis leads to the notion of *conditional convergence*.

To discuss this concept, we retain the assumption that knowledge flows freely across countries, so that technological know-how is the same for all countries. That's not an entirely crazy assumption, given that much of what we think of as pure "productivity" differences can often be traced to just different endowments of human and physical capital. For instance, a




FIGURE 4.6. Convergence in growth rates.

country might have a lower endowment of educated labor, but that is not a difference of technological knowledge. Relatedly, a poor country may use techniques of production that are relatively labor-intensive, leading to a lower marginal product of labor, but that isn't the same as asserting that overall productivity is lower.¹¹ But we allow other parameters, such as savings rates and population growth rates, to differ.

4.3.2 Convergence, But to Different Steady States. The discussion so far leads to a weaker notion of convergence. Recall that in the Solow model, parameters such as savings and population growth rates only have “level effects” on per capita income, not on its growth. The growth rate of per capita income in the long run is determined entirely by the rate of technical progress, which we've assumed to be the same for all countries.

Figure 4.6 illustrates this situation. Now there is no single line that depicts the steady-state time path of (log) per capita income for all countries. Instead, because the parameters vary, different countries have their own steady-state paths, as illustrated by the lines AB and $A'B'$. We've assumed,

¹¹For instance, very different techniques are used in dairy farming in India and in the United States. But that doesn't mean that the overall *set* of known blueprints for dairy farming differs across the two countries. The different techniques in dairy farming more likely are due to the different relative availabilities of labor and capital. In India, where labor is plentiful and capital is scarce, it would be absurd to adopt the capital-intensive methods of dairy farming used in the United States. This does not mean that Indians are unaware of U.S. farming methods (or vice versa). The assumptions of an identical *technology* and an identical *technique* are distinct.

however, that these paths are all *parallel* to one another, given that the hypothesis maintains the same rate of technical progress (and therefore the same rates of steady-state growth) across all countries.

Now imagine that the country with steady-state path AB starts at a point C above its steady-state path. The Solow model predicts that over time, this country will exhibit a slower rate of growth than the steady-state path as it slopes in to its steady-state path AB . This path is given by the curve CD . Likewise, a country that starts at point C' below its steady-state path $A'B'$ will exhibit a rate of growth higher than that of the steady state, with the resulting path $C'D'$ converging upward to its steady-state path.

4.3.3 Conditioning. In terms of a practical test that we can apply to data, what exactly is implied by this form of convergence? Recall how unconditional convergence asserted that the change in income of a country must be negatively related to its starting level of per capita income. Does our weaker notion of convergence assert something similar—that poorer countries have a tendency to grow faster? The answer is no, and Figure 4.6 shows us why. The country that starts at point C is actually poorer than the country that starts at point C' , but the former country actually grows *slower* than the latter. A country that is below its own steady state indeed grows faster than its steady-state growth rate, but to test this we have to also use the data to identify *where* these steady states are. For this reason, unconditional convergence is really an unconditional hypothesis: asserting that all the steady states are in the same place obviates the need to condition for the positions of different steady states. On the other hand, the weaker hypothesis needs to be appropriately “conditioned” on the position of steady states.

The idea of controlling for the position of steady states amounts to factoring out the effects of parameters that might differ across countries and *then* examining whether convergence occurs. Such a concept is called *conditional convergence*, because we are “conditioning” on possible intercountry differences before we examine the possibility of convergence.

To apply notions of conditioning properly to the data, we return to the material in Chapter 3, Section 3.4.4. We’re going to recall equation (3.10) from that section, which gives us a formula for steady states per *effective* unit of labor:

$$(4.1) \quad \frac{\hat{k}^*}{\hat{y}^*} \simeq \frac{s}{n + \pi + \delta}.$$

Our first task is to relate \hat{k}^* to \hat{y}^* using a production function, and to do so we will use the (now-familiar) Cobb-Douglas specification. We recall equation (3.4) from Chapter 3, which we rewrite here for convenience:

$$(4.2) \quad Y = K^\alpha (eL)^\beta$$

where we will interpret e as the efficiency of labor.¹² This represents the level of technology, and as in Chapter 3, we presume that it is growing exogenously: $e(t + 1) = (1 + \pi)e(t)$ for all t , where π is the rate of technical progress.

Dividing through by effective labor eL and defining (as before) $\hat{y} = Y/eL$ and $\hat{k} = K/eL$, we can write the production function in “per effective labor” form,

$$\hat{y} = \hat{k}^\alpha,$$

and by moving the variables around as we’re by now in the habit of doing, we see that

$$\frac{\hat{k}}{\hat{y}} = \hat{y}^{(1-\alpha)/\alpha}.$$

Substitute this formula, evaluated at the steady state, into (4.1) to obtain

$$(4.3) \quad \hat{y}^* \simeq \left[\frac{s}{n + \pi + \delta} \right]^{\alpha/(1-\alpha)}.$$

A word of reassurance: take a look at equation (3.13) in Chapter 3. We did *exactly* the same thing there; nothing new is going on except that we’re redoing things with effective units of labor. But now we go a step further.

4.3.4 The Calibration Game. We’re going to use equation (4.3) to get a sense of how much variation in steady states we can explain. Does the model “calibrate” well, in the sense that if we throw in reasonable variations in the parameters, does it “explain” the diversity in incomes that we observe?

It turns out that the parameter α is crucial in answering this question. So we need an estimate of this object. We already know that the smaller α is, the greater the extent of diminishing returns to the capital input. At the other end, as α converges to 1, there is no diminishing returns and we are in the world of Harrod and Domar. How do we pin down what a reasonable value of α is?

¹²If you are awake, you will see that the equation is missing the technical level term A . But that doesn’t matter. The term A can be folded with no loss of generality into the efficiency of labor.

There's an interesting way to do this. In competitive markets, we know that factors of production are paid their marginal products. A single line of calculus will tell you that the marginal product of capital is given by

$$\frac{\partial Y}{\partial K} = \alpha AK^{\alpha-1}(eL)^{1-\alpha} = \alpha \frac{Y}{K},$$

so that the *share of capital income in total income* is just

$$\frac{\partial Y}{\partial K} \cdot \frac{K}{Y} = \alpha!$$

The nice thing is that we do have estimates of the share of capital in national income. It varies, of course, as all estimates do, but a good approximate range for the United States is between a quarter (Parente and Prescott, 2000) and two-fifths (Lucas, 1990). A third is not a bad compromise at all.

Take $\alpha = 1/3$ to equation (4.3), to see that $\alpha/(1 - \alpha)$ is $1/2$ (or between $1/3$ and $2/3$, if you want to use the whole range above). But now we are in trouble. Suppose that two countries, sat 1 and 2, both adhere to equation (4.3), and have the same parameters except for different savings rates s_1 and s_2 . Then the ratio of their "per-capita" incomes in steady state is

$$(4.4) \quad \frac{y_1^*}{y_2^*} = \frac{\hat{y}_1^*}{\hat{y}_2^*} = \left(\frac{s_1}{s_2}\right)^{\alpha/(1-\alpha)} \simeq \left(\frac{s_1}{s_2}\right)^{1/2}.$$

That means that if country 1 has *double* the savings rate of country 2, it is predicted by our calibration to have per-capita income that is *only* $2^{1/2}$, or 1.41 times higher; just 40% higher in other words! That gets us nowhere close to the diversity that we see in the world today. Even if we take the highest estimate of $\alpha = 2/5$, so that $\alpha/(1 - \alpha)$ equals $2/3$, that gives us an explanatory magnitude of $2^{2/3}$ from doubling the savings rate, or a (relatively) paltry 60% increase in incomes. This is orders of magnitude less than the differences we do observe.

4.3.5 Conditional Convergence and Cross-Country Regressions. The same problem reappears if we try regress per-capita incomes for different countries on the different parameters. What follows is a more holistic attempt to keep track of different parameters, rather than just the savings rate differences in the previous section. To implement a cross-country regression, we need to set the appropriate equation up first. To this end, express equation (4.3) in logarithmic form to see that

$$(4.5) \quad \ln \hat{y}^* \simeq \frac{\alpha}{1 - \alpha} \ln s - \frac{\alpha}{1 - \alpha} \ln(n + \pi + \delta).$$

Then unwrap \hat{y}^* to recover per capita incomes, which we can observe directly from the data. Note that

$$\hat{y}^* = \frac{Y(t)}{L(t)e(t)} = \frac{Y(t)}{L(t)e(0)(1 + \pi)^t} = \frac{y(t)}{e(0)(1 + \pi)^t},$$

where $y(t)$ is per capita income at date t , $e(0)$ is technical knowledge at some baseline date (say 1960 or 1985), and t is counted in terms of years elapsed from the baseline date. Taking logarithms of both sides of this equation, we see that

$$\ln \hat{y}^* = \ln y(t) - \ln e(0) - t \ln(1 + \pi).$$

Substituting this expression into (4.5) and moving terms around, we obtain

$$(4.6) \quad \ln y(t) \simeq A + \frac{\alpha}{1 - \alpha} \ln s - \frac{\alpha}{1 - \alpha} \ln(n + \pi + \delta),$$

where A is just the collection of terms $\ln e(0) + t \ln(1 + \pi)$.

The plan now is to regress $y(t)$ on the parameters exactly along the lines suggested by (4.6). Observe that the intercept term A is an “unknown” from the point of view of the regression exercise; it will be estimated after the best possible fit is found to the data.

Likewise, as far as the empirical exercise goes, the coefficients on $\ln s$ as well as on $\ln(n + \pi + \delta)$ are unknowns and will be estimated by the best fit to the data. However — and this is the power of a theoretical prediction — the theory suggests that after we are done estimating the equation, the coefficients will be close in value to each other (they are *both* supposed to equal $\alpha/(1 - \alpha)$), but of opposite sign. In fact, if we take the estimate of α from the previous section, $\alpha/(1 - \alpha)$ should be around 0.5.

Therefore we enter the empirical study with the following expectations:

- (1) The coefficients on the term $\ln s$ are positive and the coefficient on the term $\ln(n + \pi + \delta)$ is negative. This captures the Solow prediction that savings has a positive (level) effect on per capita income and population growth has a negative (level) effect on per capita income.
- (2) The estimated coefficients have the same approximate magnitude, and this magnitude is around 0.5.

Mankiw, Romer, and Weil [1992] tested these predictions using the Heston–Summers data set. They took $\pi + \delta$ to be approximately 0.05, or around 5% per year, and used the average of investment–GDP ratios over the period 1965–85 to form an estimate of the savings rate. The variable y is given by per capita GDP in the year 1985. The resulting regression shows the following features:

(1) More than half the worldwide variation in per capita GDP in 1985 can be explained by the two variables s and n . The correlation coefficient of the regression is 0.59. This is a powerful finding indeed.

(2) As predicted by the Solow model, the coefficient of $\ln s$ is significant and positive, whereas that of $\ln(n + \pi + \delta)$ is significant and negative. In qualitative terms, as long as we do not stick to the absurd assumption of equal savings and population growth rates (and therefore the prediction of unconditional convergence), the Solow model predicts broad relationships that do show up in worldwide data. However, there is a bug:

(3) The coefficients are too large to be anywhere close to 0.5: the coefficient on savings is 1.42 and that on population is -1.97 . Moreover, the coefficients are far from being of similar magnitude. Population growth rates seem to have a larger depressive effect on per capita incomes than the upward kick from savings rates.