

Proximate and fundamental causes of growth

Now let's talk a little bit about what the data say regarding economic growth. There is a very long line of research trying to empirically assess the determinants of growth – an area that is still very vibrant. In order to organise what this literature has to say, it is useful to start by distinguishing between what Acemoglu (2009) calls *proximate* and *fundamental* causes of economic growth. If we think of any generic production function $Y = F(\mathbf{X}, A)$, where \mathbf{X} is a vector of inputs (capital, labour, human capital) and A captures productivity, we can attribute any increase in output to an increase in \mathbf{X} or A . In that sense, the accumulation of physical capital, human capital, or technological progress generates growth, but we still want to learn why different societies choose different accumulation paths. We can thus think of these as *proximate* causes, but we want to be able to say something about the *fundamental* causes that determine those choices. Our survey of the empirical literature will address what economists have been able to say about each of those sets of causes.

7.1 | The proximate causes of economic growth

There are three basic empirical tools to assess the importance of proximate causes of growth (factor accumulation, productivity): growth accounting, regression-based approaches, and calibration. We briefly go over the advantages and pitfalls, and the message they deliver. Factor accumulation has significant explanatory power, but in the end productivity matters a lot.

The natural starting point for this investigation is our workhorse, the Neoclassical Growth Model (NGM). The basic question, to which we have already alluded, is: how well does the NGM do in explaining differences in income levels and in growth rates?¹

Several methods have been devised and used to assess this question, and they can be broadly grouped into three classes: *growth accounting*, *growth regressions*, and *calibration*. Let us address each of these.

How to cite this book chapter:

Campante, F., Sturzenegger, F. and Velasco, A. 2021. *Advanced Macroeconomics: An Easy Guide*.

Ch. 7. 'Proximate and fundamental causes of growth', pp. 87–112. London: LSE Press.

DOI: <https://doi.org/10.31389/lsepress.ame.g> License: CC-BY-NC 4.0.

7.1.1 | Growth accounting

This is another founding contribution of Robert Solow to the study of economic growth. Right after publishing his “Contribution to the Theory of Economic Growth” in 1956, he published another article in 1957 (Solow 1957) noting that an aggregate production function such as

$$Y(t) = A(t)F(K_t, L_t), \quad (7.1)$$

when combined with competitive factor markets, immediately yields a framework that lets us account for the (proximate) sources of economic growth. Take the derivative of the log of the production function with respect to time,

$$\begin{aligned} \frac{\dot{Y}}{Y} &= \frac{\dot{A}}{A} + \frac{AF_K}{Y} \dot{K} + \frac{AF_L}{Y} \dot{L} \Rightarrow \\ \frac{\dot{Y}}{Y} &= \frac{\dot{A}}{A} + \frac{AF_K K}{Y} \frac{\dot{K}}{K} + \frac{AF_L L}{Y} \frac{\dot{L}}{L} \Rightarrow \\ g_Y &= g_A + \alpha_K g_K + \alpha_L g_L, \end{aligned} \quad (7.2)$$

where g_X is the growth rate of variable X , and $\alpha_X \equiv \frac{AF_X X}{Y}$ is the elasticity of output with respect to factor X . This is an identity, but adding the assumption of competitive factor markets (i.e. factors are paid their marginal productivity) means that α_X is also the share of output that factor X obtains as payment for its services. Equation (7.2) then enables us to estimate the contributions of factor accumulation and technological progress (often referred to as total factor productivity (TFP)) to economic growth.

This is how it works in practice: from national accounts and other data sources, one can estimate the values of g_Y, g_K, g_L, α_K , and α_L ; from (7.2) one can then back out the estimate for g_A .² (For this reason, g_A is widely referred to as the Solow residual.) Solow actually computed this for the U.S. economy, and reached the conclusion that the bulk of economic growth, about 2/3, could be attributed to the residual. Technological progress, and not factor accumulation, seems to be the key to economic growth.

Now, here is where a caveat is needed: g_A is calculated as a residual, not directly from measures of technological progress. It is the measure of our ignorance!³ More precisely, any underestimate of the increase in K or L (say, because it is hard to adjust for the increased quality of labour input), will result in an overestimate of g_A . As a result, a lot of effort has been devoted to better measure the contribution of the different factors of production.

In any event, this approach has been used over and over again. A particularly famous example was Alwyn Young’s research in the early 1990s (1995), where he tried to understand the sources of the fantastic growth performance of the East Asian “tigers”, Hong Kong, Singapore, South Korea, and Taiwan.⁴ Most observers thought that this meant that they must have achieved amazing rates of technological progress, but Young showed that their pace of factor accumulation had been astonishing. Rising rates of labour force participation (increasing L), skyrocketing rises in investment rates (from 10% of GDP in 1960 to 47% of GDP in 1984, in Singapore, for instance!) (increasing K), and increasing educational achievement (increasing H). Once all of this is accounted for, their Solow residuals were not particularly outliers compared to the rest of the world. (This was particularly the case for Singapore, and not so much for Hong Kong.) Why is this important? Well, we know from the NGM that factor accumulation cannot sustain growth in the long run! This seemed to predict that the tigers’ performance would soon hit the snag of decreasing returns. Paul Krugman started to become famous beyond the circles of economics by explicitly predicting as much in a famous article in 1994 (Krugman 1994), which was interpreted by many as having predicted the 1997 East Asian crisis.

Of course, the tigers resumed growing fast soon after that crisis – have they since then picked up with productivity growth?

7.1.2 | Using calibration to explain income differences

We have seen in Chapter 2 that a major issue in growth empirics is to assess the relative importance of factor accumulation and productivity in explaining differences in growth rates and income levels. A different empirical approach to this question is *calibration*, in which differences in productivity are calculated using imputed parameter values that come from microeconomic evidence. As it is closely related to the methodology of growth accounting, we discuss it here. (We will see later, when discussing business cycle fluctuations, that calibration is one of the main tools of macroeconomics, when it comes to evaluating models empirically.)

One of the main contributions in this line of work is a paper by Hall and Jones (1999). In their approach, they consider a Cobb-Douglas production function for country i ,

$$Y_i = K_i^\alpha (A_i H_i)^{1-\alpha}, \quad (7.3)$$

where K_i is the stock of physical capital, H_i is the amount of human capital-augmented labour and A_i is a labour-augmenting measure of productivity. If we know α , K_i and H_i , and given that we can observe Y , we can back out productivity A_i :

$$A_i = \frac{Y_i^{\frac{1}{1-\alpha}}}{K_i^{\frac{\alpha}{1-\alpha}} H_i}. \quad (7.4)$$

But how are we to know those?

For human capital-augmented labour, we start by assuming that labour L_i is homogeneous within a country, and each unit of it has been trained with E_i years of schooling. Human capital-augmented labour is given by

$$H_i = e^{\phi(E_i)} L_i. \quad (7.5)$$

The function $\phi(E)$ reflects the efficiency of a unit of labour with E years of schooling relative to one with no schooling ($\phi(0) = 0$). $\phi'(E)$ is the return to schooling estimated in a Mincerian wage regression (i.e. a regression of log wages on schooling and demographic controls, at the individual level). As such, we can run a Mincerian regression to obtain H_i . (Hall and Jones do so assuming that different types of schooling affect productivity differently.)

How about physical capital? We can compute it from data on past investment, using what is called the perpetual inventory method. If we have a depreciation rate δ , it follows that

$$K_{i,t} = (1 - \delta)K_{i,t-1} + I_{i,t-1}. \quad (7.6)$$

It also follows that


$$K_{i,t} = (1 - \delta)^t K_{i,0} + \sum_{s=0}^{t-1} I_{i,s} (1 - \delta)^{t-s-1}. \quad (7.7)$$

If we have a complete series of investment, we can calculate this for any point in time. (We assume $\delta = 0.06$ for all countries). Since we don't, we assume that, before the start of our data series, investment had been growing at the same rate that we observe in the sample. By doing that, we can compute the $K_{i,0}$ and obtain our value for the capital stock.

How about α ? Well, we go for our usual assumption of $\alpha = 1/3$, which is thought of as a reasonable value given the share of capital returns in output as measured by national accounts. This is subject to the caveats we have already discussed, but it is a good starting point.

Since we are interested in cross-country comparisons, we benchmark the data with comparisons to the U.S. series. This comparison can be seen in Figure 7.1, from Acemoglu (2009).

Figure 7.1 Productivity differences, from Acemoglu (2009)



If all countries had the same productivity, and all differences in income were due to differences in factor accumulation, we would see all countries bunched around a value of 1 in the y -axis. This is clearly not the case! Note also that the pattern seems to become stronger over time: we were farther from that benchmark in 2000 than in 1980.

To summarise the message quantitatively, we can do the following exercise. Output per worker in the five countries with the highest levels of output per worker was 31.7 times higher than output per worker in the five lowest countries. Relatively little of this difference was due to physical and human capital:

- Capital intensity per worker contributed a factor of 1.8
- Human capital per worker contributed a factor of 2.2
- Productivity contributed a factor of 8.3!

Hall and Jones associate this big impact of productivity to the role of social capital: the ability of societies to organise their economic activity with more or less costs. For example, a society where theft is prevalent will imply the need to spend resources to protect property; a society full of red tape would require lots of energy in counteracting it, and so on. In short, productivity seems a much bigger concept than just technological efficiency.

However, just as in the regression approaches, calibration also relies on important assumptions. Now, functional forms make a huge difference, both in the production function and in the human capital equation. If we lift the Cobb-Douglas production function or change the technological assumptions in the production of human capital (e.g. assuming externalities), things can change a lot.

7.1.3 | Growth regressions

Another approach to the empirics of economic growth is that of growth regressions – namely, estimating regressions with growth rates as dependent variables. The original contribution was an extremely influential paper by Robert Barro (1991), that established a canonical specification. Generally speaking, the equation to be estimated looks like this:

$$g_{i,t} = \mathbf{X}'_{i,t} \beta + \alpha \log(y_{i,t-1}) + \epsilon_{i,t}, \quad (7.8)$$

where $g_{i,t}$ is the growth rate of country i from period $t - 1$ to period t , $\mathbf{X}'_{i,t}$ is a vector of variables that one thinks can affect a country's growth rate, both in steady state (i.e. productivity) and along the transition path, β is a vector of coefficients, $y_{i,t-1}$ is country i 's output in the previous period $t - 1$, α is a coefficient capturing convergence, and $\epsilon_{i,t}$ is a random term that captures all other factors omitted from the specification.

Following this seminal contribution, innumerable papers were written over the subsequent few years, with a wide range of results. In some one variable was significant; in others, it was not. Eventually, the results were challenged on the basis of their robustness. Levine and Renelt (1991), for example, published a paper in which they argued *no* results were robust. The counterattack was done by a former student and colleague of Barro, Sala-i-Martin (1997), that applied a similar robustness check to all variables used by any author in growth regressions, in his amusingly titled paper, "I Just Ran Two Million Regressions". He concluded that, out of the 59 variables that had shown up as significant somewhere in his survey of the literature, some 22 seem to be robust according to his more lax, or less extreme, criteria (compared to Levine and Renelt's). These include region and religion dummies, political variables (e.g. rule of law), market distortions (e.g. black market premium), investment, and openness.

Leaving aside the issues of robustness, the approach, at least in its basic form, faces other severe challenges, which are of two types, roughly speaking.

1. Causality (aka Identification; aka Endogeneity): The variables in $\mathbf{X}_{i,t}$ are typically endogenous, i.e. jointly determined with the growth rate. As you have seen in your courses on econometrics, this introduces bias in our estimates of β , which in turn makes it unlikely that we identify the causal effect of any of those variables (at the end of this chapter, when discussing institutions, we will discuss the solution of this problem suggested by Acemoglu et al. (2001), one of the most creative and influential proposed solutions to this endogeneity problem).
2. Interpretation: The economic interpretation of the results might be quite difficult. Suppose that openness truly affects economic growth, but it does so (at least partly) by increasing investment rates; if our specification includes both variables in $\mathbf{X}_{i,t}$, the coefficient on openness will not capture the full extent of its true effect.

Both of these are particularly problematic if we want to investigate the relationship between policies and growth, a point that is illustrated by Dani Rodrik's (2012) critique. Rodrik's point is that if policies are endogenous (and who could argue they are not?) we definitely have a problem. The intuition is as follows. Imagine you want to test whether public banks are associated with higher or lower growth. If you run a regression of growth on, say, share of the financial sector that is run by public banks, you may find a negative coefficient. But is that because public banks are bad for growth? Or is it because politicians resort to public banks when the economy faces a lot of constraints (and thus its growth is bound to be relatively low)?

To see the issue more clearly, consider a setup, from a modified AK model, in which

$$g = (1 - \theta)A - \rho, \quad (7.9)$$

where θ is a distortion. Now consider a policy intervention s , which reduces the distortion, but that has a cost of its own. Then,

$$g(s, \theta, \phi) = (1 - \theta(1 - s))A - \phi\alpha(s) - \rho. \quad (7.10)$$

The optimal intervention delivers growth as defined by the implicit equation

$$g_s(s^{**}, \theta, \phi) = 0. \quad (7.11)$$

In addition, there is a diversion function of the policy maker $\pi(s)$, with $\pi'(s) > 0$, $\pi''(s) < 0$, and $\pi'(s^P) = 0$ with $s^P > s^{**}$. This means that the politicians will use the intervention more than is actually desirable from a social perspective. The politician will want to maximise growth and their own benefit, placing a weight λ on growth. This means solving

$$\max_s u(s, \theta, \phi) = \lambda g(s, \theta, \phi) + \pi(s), \quad (7.12)$$

which from simple optimisation yields the FOC

$$\lambda g_s(s^*, \theta, \phi) + \pi'(s^*) = 0. \quad (7.13)$$

Because we have assumed that $\pi'(s) > 0$, it follows from (7.13) that $g_s(s^*, \theta, \phi) < 0$, and this implies that a reduction in s will increase growth. Does this imply that we should reduce s ? Marginally, yes, but not to zero, which is the conclusion that people typically read from growth regressions.

Now, what if we were to run a regression to study the links between policy s and growth? We need to take into account the changes in s that happen when the parameters vary across countries. Consider the effect of changes in the level of distortions θ . Recall that, from (7.10):

$$g_s(s, \theta, \phi) = \theta A - \phi \alpha'(s). \quad (7.14)$$

Replacing in (7.13) and totally differentiating yields

$$d\theta \lambda A + [-\lambda \phi \alpha''(s) + \pi''(s^*)] ds^* = 0 \quad (7.15)$$

$$\frac{ds^*}{d\theta} = \frac{\overset{(+)}{\lambda A}}{\underbrace{\lambda \phi \alpha''(s) - \pi''(s^*)}_{(+)}} > 0. \quad (7.16)$$

This implies that in an economy with greater inefficiencies we will see a higher level of the policy intervention, as long as politicians care about growth. But growth will suffer with the greater inefficiencies: differentiating (7.10) with respect to θ we have

$$\frac{dg}{d\theta} = -A(1 - s^*) + g_s(s^*, \theta, \phi) \frac{ds^*}{d\theta} < 0 \Rightarrow \frac{\frac{dg}{d\theta}}{\frac{ds^*}{d\theta}} < 0. \quad (7.17)$$

The fact that this coefficient is negative means nothing, at least from a policy perspective (remember that it is optimal to increase the policy intervention if the distortion increases).

Because of challenges like these, people later moved to analyse panel growth regressions, which rearrange (7.8) as

$$g_{i,t} = \mathbf{X}'_{i,t} \beta + \alpha \log(y_{i,t-1}) + \delta_i + \mu_t + \epsilon_{i,t}, \quad (7.18)$$

where δ_i and μ_t are country and time fixed effects, respectively. By including country fixed effects, this removes fixed country characteristics that might affect both growth and other independent variables of interest, and thus identifies the effects of such variables out of within-country variation. However, in so doing they might be getting rid of most of the interesting variation, which is across countries, while also increasing the potential bias due to measurement error. Finally, these regressions do not account for time-varying country-specific factors. In sum, they are no panacea.

Convergence

Another vast empirical debate that has taken place within the framework of growth regressions is one to which we have already alluded when discussing the Solow model: convergence. We have talked, very briefly, about what the evidence looks like; let us now get into more detail.

Absolute convergence

As you recall, this is the idea that poorer countries grow faster than richer countries, unconditionally. Convergence is a stark prediction of the NGM, and economists started off by taking this (on second

thought, naive) version of it to the data. Baumol (1986) took Maddison's core sample of 16 rich countries over the long run and found

$$\text{growth} = 5.251 - 0.749 \text{ initial income} \\ (0.075)$$

with $R^2 = 0.87$. He thus concluded that there was strong convergence!

However, De Long (1988) suggested a reason why this result was spurious: only successful countries took the effort to construct long historical data series. So the result may be a simple fluke of sample selection bias (another problem is measurement error in initial income that also biases the results in favour of the convergence hypothesis). In fact, broadening the sample of countries beyond Maddison's sixteen leads us immediately to reject the hypothesis of convergence. By the way, there has been extensive work on convergence, within countries and there is fairly consistent evidence of absolute convergence for different regions of a country.⁵

Conditional convergence

The literature then moved to discuss the possibility of conditional convergence. This means including in a regression a term for the initial level of GDP, and checking whether the coefficient is negative when controlling for the other factors that determine the steady state of each economy. In other words, we want to look at the coefficient α in (7.8), which we obtain by including the control variables in \mathbf{X} . By including those factors in the regression, we partial out the steady state from initial income and measure deviations from this steady state. This, of course, is the convergence that is actually predicted by the NGM.

Barro (1991) and Barro and Sala-i-Martin (1992) found evidence of a negative α coefficient, and we can say that in general the evidence is favourable to conditional convergence. Nevertheless, the same issues that apply to growth regressions in general will be present here as well.

7.1.4 | Explaining cross-country income differences, again

Another regression-based approach to investigate how the NGM fares in explaining the data was pioneered by Mankiw et al. (1992) (MWR hence). Their starting point is playfully announced in the very first sentence: "This paper takes Robert Solow seriously" (p. 407).⁶ This means that they focus simply on the factor accumulation determinants that are directly identified by the Solow model as the key proximate factors to explain cross-country income differences, leaving aside the productivity differences. They claim that the NGM (augmented with human capital) does a good job of explaining the existing cross-country differences.

Basic Solow model

There are two inputs, capital and labour, which are paid their marginal products. A Cobb-Douglas production function is assumed

$$Y_t = K_t^\alpha (A_t L_t)^{1-\alpha} \quad 0 < \alpha < 1. \quad (7.19)$$

L and A are assumed to grow exogenously at rates n and g :

$$\frac{\dot{L}}{L} = n \quad (7.20)$$

$$\frac{\dot{A}}{A} = g. \quad (7.21)$$

The number of effective units of labour $A(t)L(t)$ grows at rate $n + g$.

As usual, we define k as the stock of capital per effective unit of labour $k = \frac{K}{AL}$ and $y = \frac{Y}{AL}$ as the level of output per effective unit of labour.

The model assumes that a constant fraction of output s is invested. The evolution of k is

$$\dot{k}_t = sy_t - (n + g + \delta) k_t \quad (7.22)$$

or

$$\dot{k}_t = sk_t^\alpha - (n + g + \delta) k_t, \quad (7.23)$$

where δ is the rate of depreciation. The steady state value k^* is

$$k^* = \left[\frac{s}{(n + g + \delta)} \right]^{\frac{1}{1-\alpha}}. \quad (7.24)$$

Output per capita is

$$\left(\frac{Y_t}{L_t} \right) = K_t^\alpha A_t^{1-\alpha} L_t^{-\alpha} = k_t^\alpha A_t. \quad (7.25)$$

Substituting (7.24) into (7.25)

$$\left(\frac{Y_t}{L_t} \right) = \left[\frac{s}{(n + g + \delta)} \right]^{\frac{\alpha}{1-\alpha}} A_t \quad (7.26)$$

and taking logs

$$\log \left(\frac{Y_t}{L_t} \right) = \frac{\alpha}{1-\alpha} \log(s) - \frac{\alpha}{1-\alpha} \log(n + g + \delta) + \log A(0) + gt. \quad (7.27)$$

MRW assume that g (representing advancement of knowledge) and δ do not vary across countries, but A reflects not only technology but also resource endowments. It thus differs across countries as in

$$\log A(0) = a + \epsilon, \quad (7.28)$$

where a is a constant and ϵ is a country-specific shock. So we have

$$\log \left(\frac{Y}{L} \right) = a + \frac{\alpha}{1-\alpha} \log(s) - \frac{\alpha}{1-\alpha} \log(n + g + \delta) + \epsilon \quad (7.29)$$

We assume s and n are not correlated with ϵ . (What do you think of this assumption?) Since it is usually assumed that the capital share is $\alpha \cong \frac{1}{3}$, the model predicts an elasticity of income per capita with respect to the saving rate $\frac{\alpha}{1-\alpha} \cong \frac{1}{2}$ and an elasticity with respect to $n+g+\delta$ of approximately -0.5 .

Table 7.1 Estimates of the basic Solow model

	Update		
	Log GDP per Capita		
	MRW 1985	Acemoglu 2000	Update 2017
$\log(s_k)$	1.42*** (.14)	1.22*** (.13)	.96* (.48)
$\log(n+g+\delta)$	-1.97*** (.56)	-1.59*** (.36)	-1.48*** (.21)
Implied α	.59	.55	.49
Adjusted R ²	.59	.49	.49

Note: *p<0.1; **p<0.05; ***p<0.01

What do the data say?

With data from the real national accounts constructed by Summers and Heston (1988) for the period 1960-1985, they run (7.29), using ordinary least squares (OLS) for all countries for which data are available minus countries where oil production is the dominant industry.

We reproduce their results in Table 7.1, to which we add an update by Acemoglu (2009), and one of our own more than 30 years after the original contribution. In all three cases, aspects of the results support the Solow model:

1. Signs of the coefficients on saving and population growth are OK.
2. Equality of the coefficients for $\log(s)$ and $-\log(n+g+\delta)$ is not rejected.
3. A high percentage of the variance is explained (see R^2 in the table).

But the estimate for α contradicts the prediction that $\alpha = 1/3$. While the implicit value of α seems to be falling, in each update it is still around or above .5. Some would have said it is OK (remember our discussion in Chapter 2), but for MRW it was not.

Introducing human capital

MRW go on to consider the implications of considering the role of human capital. Let us now recall the augmented Solow model that we saw in Chapter 5. The production function is now

$$Y_t = K_t^\alpha H_t^\beta (A_t L_t)^{1-\alpha-\beta}, \quad (7.30)$$

where H is the stock of human capital. If s_k is the fraction of income invested in physical capital and s_h the fraction invested in human capital, the evolution of k and h are determined by

$$\dot{k}_t = s_k y_t - (n + g + \delta) k_t \quad (7.31)$$

$$\dot{h}_t = s_h y_t - (n + g + \delta) h_t, \quad (7.32)$$

where k , h and y are quantities per effective unit of labour.

It is assumed that $\alpha + \beta < 1$, so that there are decreasing returns to all capital and we have a steady state for the model. The steady-state level for k and h are

$$k^* = \left[\frac{s_k^{1-\beta} s_h^\beta}{(n + g + \delta)} \right]^{\frac{1}{1-\alpha-\beta}} \tag{7.33}$$

$$h^* = \left[\frac{s_k^\alpha s_h^{1-\alpha}}{(n + g + \delta)} \right]^{\frac{1}{1-\alpha-\beta}} \tag{7.34}$$

Substituting (7.33) and (7.34) into the production function and taking logs, income per capita is

$$\log \left(\frac{Y_t}{L_t} \right) = \frac{\alpha}{1 - \alpha - \beta} \log(s_k) + \frac{\beta}{1 - \alpha - \beta} \log(s_h) - \frac{\alpha + \beta}{1 - \alpha - \beta} \log(n + g + \delta) + \log A(0) + gt. \tag{7.35}$$

To implement the model, investment in human capital is restricted to education. They construct a SCHOOL variable that measures the percentage of the working age population that is in secondary school, and use it as a proxy for human capital accumulation s_h .

The results are shown in Table 7.2. It turns out that now 78% of the variation is explained, and the numbers seem to match: $\hat{\alpha} \cong 0.3, \hat{\beta} \cong 0.3$. (For the updated data we have a slightly lower R^2 and a higher $\hat{\beta}$ indicating an increasing role of human capital, in line with what we found in the previous section.)

Table 7.2 Estimates of the augmented Solow model

	Update		
	Log GDP per Capita		
	MRW 1985	Acemoglu 2000	Update 2017
$\log(s_k)$.69*** (.13)	.96*** (.13)	.71 (.44)
$\log(n+g+\delta)$	-1.73*** (.41)	-1.06*** (.33)	-1.43*** (.19)
$\log(s_h)$.66*** (.07)	.70*** (.13)	1.69*** (.43)
Implied α	.30	.36	.28
Implied β	.28	.26	.33
Adjusted R ²	.78	.60	.59

Note: *p<0.1; **p<0.05; ***p<0.01

Challenges

The first difficulty with this approach is: is it really OK to use OLS? Consistency of OLS estimates requires that the error term be orthogonal to the other variables. But that error term includes technology differences, and are these really uncorrelated with the accumulation of physical and human capital? If not, the omitted variable bias (and reverse causality) would mean that the estimates for the effects of physical and human capital accumulation (and for the R^2) are biased upwards, and the NGM doesn't do as good a job as MRW think, when it comes to explaining cross-country income differences. This is the very same difficulty that arises from the growth regressions approach – not surprising, since the econometric underpinnings are very much similar.

A second difficulty has to do with the measure of human capital: is it really a good one? The microeconomic evidence suggests that the variation in average years of schooling across countries that we see in the data is not compatible with the estimate $\hat{\beta}$ obtained by MRW.

7.1.5 | Summing up

We have seen many different empirical approaches, and their limitations. Both in terms of explaining differences in growth and in income levels at the cross-country level, there is a lot of debate on the extent to which the NGM can do the job.

It does seem that the consensus in the literature today is that productivity differences are crucial for understanding cross-country differences in economic performance. (A paper by Acemoglu and Dell (2010) makes the point that productivity differences are crucial for *within*-country differences as well.) This means that the endogenous growth models that try to understand technological progress are a central part of understanding those differences.

In the previous chapter we talked about some of the questions surrounding those models, such as the effects of competition and scale, but these models focused on productive technology, that is, how to build a new blueprint or a better variety for a good. The empirical research, as we mentioned above, suggests that productivity differences don't necessarily mean technology in a narrow sense. A country can be less productive because of market or organisational failures, even for a given technology. The reasons for this lower productivity may be manifold, but they lead us into the next set of questions: what explains them? What explains differences in factor accumulation? In other words, what are the *fundamental* causes of economic performance? We turn to this question now.

7.2 | The fundamental causes of economic growth

We go over four types of fundamental explanations for differences in economic performance: luck (multiple equilibria), geography, culture, and institutions.

As North (1990) point out, things like technological progress and factor accumulation “are not causes of growth; *they are growth*” (p.2). The big question is, what in turn causes them? Following Acemoglu (2009), we can classify the main hypotheses into four major groups:

1. *Luck*: Countries that are identical in principle may diverge because small factors lead them to select different equilibria, assuming that multiple equilibria exist.

2. *Geography*: Productivity can be affected by things that are determined by the physical, geographical, and ecological environment: soil quality, presence of natural resources, disease environment, inhospitable climate, etc.
3. *Culture*: Beliefs, values, and preferences affect economic behaviour, and may lead to different patterns of factor accumulation and productivity: work ethic, thrift, attitudes towards profit, trust, etc.
4. *Institutions*: Rules, regulations, laws, and policies that affect economic incentives to invest in technology, physical, and human capital. The crucial aspect is that institutions are *choices* made by society.

Let us discuss each one of them.

7.2.1 | Luck

This is essentially a catchier way of talking about multiple equilibria. If we go back to our models of poverty traps, we will recall that, in many cases, the same set of parameters is consistent with more than one equilibrium. Moreover, these equilibria can be ranked in welfare terms. As a result, it is possible (at least theoretically) that identical countries will end up in very different places.

But is the theoretical possibility that important empirically? Do we really believe that Switzerland is rich and Malawi is poor essentially because of luck? It seems a little hard to believe. Even if we go back in time, it seems that initial conditions were very different in very relevant dimensions. In other words, multiple equilibria might explain relatively small and short-lived divergence, but not the bulk of the mind-boggling cross-country differences we see – at least not in isolation.

In addition, from a conceptual standpoint, a drawback is that we need to explain the coordination failures and how people fail to coordinate even when they are trapped in a demonstrably bad equilibrium. This pushes back the explanatory challenge by another degree.

In sum, it seems that multiple equilibria and luck might be relevant, but in conjunction with other explanations. For instance, it may be that a country happened to be ruled by a growth-friendly dictator, while another was stuck with a growth-destroying one. Jones and Olken (2005) use random deaths of country leaders to show that there does seem to be an impact on subsequent performance. The question then becomes why the effects of these different rulers would matter over the long run, and for this we would have to consider some of the other classes of explanations.⁷

7.2.2 | Geography


This is somewhat related to the luck hypothesis, but certainly distinctive: perhaps the deepest source of heterogeneity between countries is the natural environment they happened to be endowed with. From a very big picture perspective, geographical happenstance of this sort is a very plausible candidate for a determinant of broad development paths, as argued for instance by Jarred Diamond in his 1999 Pulitzer-Prize-winning book *Guns, Germs and Steel*⁸. As an example, Diamond suggests that one key reason Europe conquered America, and not the other way around, was that Europe had an endowment of big animal species that were relatively easy to domesticate, which in turn led to improved immunisation by humans exposed to animal-borne diseases, and more technological advances. But can geography also explain differences in economic performance at the scale on which we usually think about them, say between different countries over decades or even a couple of centuries?

On some level, it is hard to think that the natural environment would not affect economic performance, on any time frame. Whether a country is in the middle of the Sahara desert, the Amazon rain forest, or some temperate climate zone must make some difference for the set of economic opportunities that it faces. This idea becomes more compelling when we look at the correlation between certain geographical variables and economic performance, as illustrated by the Figure (7.2), again taken from Acemoglu (2009). It is clear from that picture that countries that are closer to the equator are poorer on average. At the very least, any explanation for economic performance would have to be consistent with this stylised fact. The question, once again, is to assess to what extent these geographical differences underlie the ultimate performance, and this is not an easy empirical question.

Let us start by considering the possible conceptual arguments. The earliest version of the geography hypothesis has to do with the effect of the climate on the effort – the old idea that hot climates are not conducive to hard work. While this seems very naive (and not too politically correct) to our 21st century ears, the idea that climate (and geography more broadly) affects technological productivity, especially in agriculture, still sounds very plausible. If these initial differences in turn condition subsequent technological progress (as argued by Jared Diamond, as we have seen, and as we will see, in different forms, by Jeffrey Sachs), it just might be that geography is the ultimate determinant of the divergence between societies over the very long run.

A big issue with this modern version of the geography hypothesis is that it is much more appealing to think of geography affecting agricultural productivity, but modern growth seems to have a lot more to do with industrialisation. While productivity in agriculture might have conditioned the development of industry to begin with, once industrial technologies are developed we would have to explain why they are not adopted by some countries. Geography is probably not enough to account for that, at least in this version of the story.

Figure 7.2 Distance from the equator and income, from Acemoglu (2009)



Another version has to do with the effect of geography on the disease environment, and the effect of the latter on productivity. This is a version strongly associated with Jeffrey Sachs (2002), who argues that the disease burden in the tropics (malaria in particular) can explain a lot of why Africa is so poor. The basic idea is very simple: unhealthy people are less productive. However, many of these diseases have been (or could potentially be) tamed by technological progress, so the question becomes one of why some countries have failed to benefit from that progress. In other words, the disease environment that prevails in a given country is also a consequence of its economic performance. While this doesn't mean that there cannot be causality running in the other direction, at the very least it makes the empirical assessment substantially harder.

What does the evidence say, broadly speaking? Acemoglu et al. (2002) (henceforth AJR) make the argument of the *reversal of fortune* to suggest that geography cannot explain that much. Consider the set of countries that were colonised by the Europeans, starting in the 15th century. The point is that countries that were richer before colonisation eventually became poorer – think about Peru or Mexico versus Canada, Australia, or the U.S. (see Figures 7.3 and 7.4). But geography, if the concept is to mean anything, is largely constant over time! (At least over the time periods we talk about.)

But how about the version that operates through the disease environment? This might operate on a smaller scale than the one that is belied by the reversal of fortunes argument. To assess this argument, we want to have some exogenous variation in the disease environment, that enables us to disentangle the two avenues of causality. Acemoglu and Johnson (2007) use the worldwide technological shocks that greatly improved control over many of the world's worst diseases. They measure this exogenous impact, at the country level, by considering the date at which a technological breakthrough

Figure 7.3 Reversal of fortunes - urbanization, from Acemoglu (2009)





Figure 7.4 Reversal of fortunes - pop. density, from Acemoglu (2009)



was obtained against a given disease, such as tuberculosis or malaria, and the country's initial exposure to that disease. What they show is that, quite beyond not having a quantitatively important effect on output per capita, these health interventions actually seem not to have had any significant effect at all.⁹

Finally, another version of the geography argument relates to the role of natural resources and growth. Sachs and Warner (2001) tackle this issue and find a surprising result: countries endowed with natural resources seem to grow slower than countries that do not (think of Congo, Zambia or Iran, vs Japan and Hong Kong). How could this be so? Shouldn't having more resources be good? Sachs associates the poorer performance to the fact that societies that are rich in resources become societies of rent-seekers, societies where appropriating the wealth of natural resources is more important than creating new wealth. Another explanation has to do with relative prices. Commodity booms lead to a sustained real exchange rate appreciation that fosters the growth of non-tradable activities, where productivity growth seems a bit slower. Finally, commodity economies suffer large volatility in their real exchange rates, making economic activity more risky both in the tradable and non-tradable sectors. This decreases the incentives to invest, as we will see later in the book, and also hurts growth prospects. Obviously, this is not a foregone conclusion. Some countries like Norway or Chile have learnt to deal with the challenge of natural resources by setting sovereign wealth funds or investment strategies that try to diminish these negative effects. But then this, once again, pushes the question of this dimension of geography to that of institutions, to which we will shortly turn below.

7.2.3 | Culture

What do we mean by culture? The standard definition used by economists, as spelled out by Guiso et al. (2006), refers to “those customary beliefs and values that ethnic, religious, and social groups transmit fairly unchanged from generation to generation” (p. 23). In other words, culture is something that lives inside people’s heads – as opposed to being external to them – but it is not something idiosyncratic to individuals; it is built and, importantly, transmitted at the level of groups.

It is hard to argue against the assertion that people’s beliefs, values, and attitudes affect their economic decisions. It is just as clear that those beliefs, values and attitudes vary across countries (and over time). From this it is easy to conclude that culture matters for economic performance, an argument that goes back at least to Max Weber’s thesis that Protestant beliefs and values, emphasising hard work and thrift, and with a positive view of wealth accumulation as a signal of God’s grace, were an important factor behind the development of capitalism and the modern industrial development. In his words, the “Protestant ethic” lies behind the “spirit of capitalism”.

Other arguments in the same vein have suggested that certain cultural traits are more conducive to economic growth than others (David Landes is a particularly prominent proponent of this view, as in Landes (1998)), and the distribution of those traits across countries is the key variable to ultimately understand growth. “Anglo-Saxon” values are growth-promoting, compared to “Latin” or “Asian” values, and so on. More recently, Joel Mokyr (2018) has argued that Enlightenment culture was the key driving force behind the emergence of the Industrial Revolution in Europe, and hence of the so-called “Great Divergence” between that continent and the rest of the world.

A number of issues arise with such explanations. First, culture is hard to measure, and as such may lead us into the realm of tautology. A country is rich because of its favourable culture, and a favourable culture is defined as that which is held by rich countries. This doesn’t get us very far in understanding the causes of good economic performance. This circularity is particularly disturbing when the same set of values (say, Confucianism) is considered inimical to growth when Asian countries perform poorly, and suddenly becomes growth-enhancing when the same countries perform well. Second, even if culture is indeed an important causal determinant of growth, we still need to figure out where it comes from if we are to consider implications for policy and predictions for future outcomes.

These empirical and conceptual challenges have now been addressed more systematically, as better data on cultural attitudes have emerged. With such data, a vibrant literature has emerged, with economists developing theories and testing their predictions on the role that specific types of values (as opposed to a generic “culture” umbrella) play in determining economic performance. Many different types of cultural attitudes have been investigated: trust, collectivism, gender roles, beliefs about fairness, etc. This literature has often exploited historical episodes – the slave trade, the formation of medieval self-governing cities, colonisation, immigration, recessions – and specific cultural practices – religious rites, civic festivities, family arrangements – to shed light on the evolution of cultural attitudes and their impact on economic outcomes. Our assessment is that this avenue of research has already borne a lot of fruit, and remains very promising for the future. (For an overview of this literature, see the surveys by Guiso et al. (2006), Alesina and Giuliano (2015), and Nunn (2020).

As an example of this research, Campante and Yanagizawa-Drott (2015) address the question of whether one specific aspect of culture, namely religious practices, affects economic growth. They do so by focusing on the specific example of Ramadan fasting (one of the pillars of Islam). To identify a causal effect of the practice, they use variation induced by the (lunar) Islamic calendar: do (exogenously)

longer Ramadan fasting hours affect economic growth? The answer they find is yes, and negatively (in Muslim countries only, reassuringly). They find a substantial effect, beyond the month of Ramadan itself, which cannot be fully explained by toll exacted by the fasting, but that they attribute to changes in labour supply decisions. People also become happier, showing that there is more to life than GDP growth. These results are consistent with existing theory on the emergence of costly religious practices. They work as screening devices to prevent free riding, and the evidence shows that more religious people become more intensely engaged, while the less committed drop out. In addition, there is an effect on individual attitudes. There is a decline in levels of general trust, suggesting that religious groups may be particularly effective in generating trust. (Given that trust is associated with good economic outcomes, we may speculate about the possible long-term impact of these changes.) In short, this illustrates how we can try to find a causal effect of cultural practices on growth, as well as trying to elucidate some of the relevant mechanisms.

7.2.4 | Institutions

Last but not least, there is the view that institutions are a fundamental source of economic growth. This idea also has an old pedigree in economics, but in modern times it has been mostly associated, in its beginnings, with the work of Douglass North (who won the Nobel Prize for his work), and more recently with scholars such as Daron Acemoglu and James Robinson. From the very beginning, here is the million-dollar question: what do we mean by institutions?

North's famous characterisation is that institutions are "the rules of the game" in a society, "the humanly devised constraints that shape human interaction" (North (1990), p. 3). Here are the key elements of his argument:

- Humanly devised: Unlike geography, institutions are chosen by groups of human beings.
- Constraints: Institutions are about placing constraints on human behaviour. Once a rule is imposed, there is a cost to breaking it.
- Shape interactions: Institutions affect incentives.

OK, fair enough. But here is the *real* question: What *exactly* do we mean by institutions? A first stab at this question is to follow the Acemoglu et al. (2005) distinctions between economic and political institutions, and between *de facto* and *de jure* institutions.

The first distinction is as follows. Economic institutions are those that directly affect the economic incentives: property rights, the presence and shape of market interactions, and regulations. They are obviously important for economic growth, as they constitute the set of incentives for accumulation and technological progress. Political institutions are those that configure the process by which society makes choices: electoral systems, constitutions, the nature of political regimes, the allocation of political power etc. There is clearly an intimate connection between those two types, as political power affects the economic rules that will prevail.

The second distinction is just as important, having to do with formal vs informal rules. For instance, the law may state that all citizens have the right to vote, but in practice it might be that certain groups can have enough resources (military or otherwise) to intimidate or influence others, thereby constraining their right in practice. Formal rules, the *de jure* institutions, are never enough to fully characterise the rules of the game; the informal, *de facto* rules must be taken into consideration.

These distinctions help us structure the concepts, but we also hit the same issue that plagues the cultural explanations: since institutions are made by people, we need to understand where they come

from, and how they come about. Acemoglu et al. (2005) is a great starting point to survey this literature, and (Acemoglu and Robinson 2012) provides an extremely readable overview of the ideas.

How do we assess empirically the role of institutions as a fundamental determinant of growth? At a very basic level, we can start by focusing on one thing that generates discontinuous change in institutions, but not so much in culture, and arguably not at all in geography: borders. Consider the following two examples. Figure 7.5 shows a Google Earth image of the border between Bolivia (on the left) and Brazil. We can see how the Brazilian side is fully planted with (mostly) soybeans, unlike the Bolivian side. A better-known version showing the same idea, in even starker form, is the satellite image of the Korean Peninsula at night (Figure 7.6).

How can we do this more systematically? Here the main challenge is similar to the one facing the investigation on the effects of disease environment: is a country rich because it has good institutions, or does it have good institutions because it's rich? The seminal study here is Acemoglu et al. (2001), and it is worth going through that example in some detail – not so much for the specific answers they find, which have been vigorously debated for a couple of decades, at this point – but for how it illustrates the challenges involved, how to try and circumvent them, and the many questions that come from that process.

The paper explores the effects of a measure of institutional development given by an index of protection from expropriation. (What kind of institution is that? What are the problems with a measure like this?) The key challenge is to obtain credible exogenous variation in that measure – something that affects institutions, but not the outcome of interest (income per capita), other than through its effect on the former.

Their candidate solution for this identification problem comes again from the natural experiment of European colonisation. The argument is that current institutions are affected by the institutions that Europeans chose to implement in their colonies (persistence of institutions), and those in turn were affected by the geographical conditions they faced – in particular, the disease environment. In more inhospitable climates (from their perspective), Europeans chose not to settle, and instead set up extractive institutions. In more favourable climates they chose to settle and, as a result, ended up choosing institutions that protected the rights of the colonists. (Note that this brings in geography as a variable that affects outcomes, but *through its effect on institutions*. In particular, this helps explain the correlations with geographical variables that we observe in the data.) The key assumption is that the disease environment at the time of colonisation doesn't really affect economic outcomes today except

Figure 7.5 Border between Bolivia (left) and Brazil




Figure 7.6 The Korean Peninsula at night



through their effect on institutional development. If so, we can use variation in that environment to identify the causal effect of institutions.

Under these assumptions, they use historical measures of mortality of European settlers as an instrument for the measure of contemporaneous institutions (property rights protection), which allows them to estimate the impact of the former on contemporaneous income per capita. The resulting estimate of the impact of institutions on income per capita is 0.94. This implies that the 2.24 difference in expropriation risk between Nigeria and Chile should translate into a difference of 206 log points (approximately 8 times, since $e^{2.06} = 7.84$). So their result is that institutional factors can explain a lot of observed cross-country differences. Also, the results suggest that, once the institutional element is controlled for, there is no additional effect of the usual geographical variables, such as distance to the equator.

Their paper was extremely influential, and spawned a great deal of debate. What are some of the immediate problems with it? The most obvious is that the disease environment may have a direct impact on output (see geography hypothesis), and the disease environment back in the days

of colonisation is related to that of today. They tackle this objection, and argue that the mortality of European settlers is to a large extent unrelated to the disease environment for natives, who had developed immunity to a lot of the diseases that killed those settlers. An objection that is not as obvious is whether the impact of the European settlers was through institutions, or something else. Was it culture that they brought? They argue that accounting for institutions wipes out the effect of things such as the identity of the coloniser. Was it human capital? Glaeser et al. (2004) argue that what they brought was not good institutions, but themselves: the key was their higher levels of human capital, which in turn are what is behind good institutions. This is a much harder claim to settle empirically, so the question remains open.

Broadly speaking, there is wide acceptance in the profession, these days, that institutions play an important role in economic outcomes. However, there is a lot of room for debate as to which are the relevant institutions, and where they come from. How do societies choose different sets of institutions? Particularly if some institutions are better for economic performance than others, why do some countries choose bad institutions? Could it be because some groups benefit from inefficient institutions? If so, how do they manage to impose them on the rest of society? In other words, we need to understand the political economy of institutional development. This is a very open and exciting area of research, both theoretically and empirically.


As an example of the themes in the literature, Acemoglu and Robinson (2019) asks not only why certain countries develop a highly capable state and others don't, but also why, among those that do, some have that same state guarantee the protection of individual rights and liberties, while others have a state that tramples on those rights and liberties. Their argument is that institutional development proceeds as a race between the power of the state and the power of society, as people both demand the presence of the Leviathan enforcing rules and order, and resent its power. If the state gets too powerful relative to society, the result is a despotic situation; if the state is too weak, the result is a state incapable of providing the needed underpinnings for development. In the middle, there is the "narrow corridor" along which increasing state capacity pushes for more societal control, and the increased power of society pushes for a more capable (and inclusive) state. The dynamics are illustrated by Figure 7.7, and one crucial aspect is worth mentioning: small differences in initial conditions – say, two economies just on opposite sides of the border between two regions in the figure – can evolve into vastly different institutional and development paths.

7.3 | What have we learned?

When it comes to the proximate causes of growth, in spite of the limitations of each specific empirical approach – growth accounting, regression methods, and calibration – the message from the data is reasonably clear, yet nuanced: factor accumulation can arguably explain a substantial amount of income differences, and specific growth episodes, but ultimately differences in productivity are very important. This is a bit daunting, since the fact is that we don't really understand what productivity *is*, in a deeper sense. Still, it underscores the importance of the process of technological progress – and the policy issues raised in Chapter 6 – as a primary locus for growth policies.

How about the fundamental causes? There is certainly a role for geography and luck (multiple equilibria), but our reading of the literature is that culture and institutions play a key part. There remains a lot to be learned about how these things evolve, and how they affect outcomes, and these are bound to be active areas of research for the foreseeable future.

Figure 7.7 Weak, despotic and inclusive states, from Acemoglu and Robinson (2017)



7.4 | What next?

Once again, the growth textbook by Acemoglu (2009) is a superb resource, and it contains a more in-depth discussion of the empirical literature on the proximate causes of growth. It also has a very interesting discussion on the fundamental causes, but it's useful to keep in mind that, its author being one of the leading proponents of the view that institutions matter most, it certainly comes at that debate from that specific point of view.

Specifically on culture, the best places to go next are the survey articles we mentioned in our discussion. The survey by Guiso et al. (2006) is a bit outdated, of course, but still a great starting point. The more recent surveys by Alesina and Giuliano (2015), focusing particularly on the links between culture and institutions, and by Nunn (2020), focusing on the work using historical data, are very good guides to where the literature is and is going.

On institutions, there is no better place to go next than the books by Acemoglu and Robinson (2012) and Acemoglu and Robinson (2019). They are very ambitious intellectual exercises, encompassing theory, history, and empirical evidence, and meant for a broad audience – which makes them a fun and engaging read.

These being very active research fields, there are a lot of questions that remain open. Anyone interested in the social sciences, as the readers of this book most likely are, will find a lot of food for thought in these sources.

Notes

- ¹ We know, of course, that the NGM does not generate long-run growth, except through exogenous technical progress. However, keep in mind that we also live in the transition path!
- ² Measuring each of these variables is an art by itself, and hundreds of papers have tried to refine these measures. Capital stocks are usually computed from accumulating past net investment and human capital from accumulating population adjusted by their productivity, assessed through Mincer equations relating years of schooling and income.
- ³ This memorable phrase is attributed to Moses Abramovitz.
- ⁴ Check out the priceless first paragraph of his 1995 paper summarising his findings: “This is a fairly boring and tedious paper, and is intentionally so. This paper provides no new interpretations of the East Asian experience to interest the historian, derives no new theoretical implications of the forces behind the East Asian growth process to motivate the theorist, and draws no new policy implications from the subtleties of East Asian government intervention to excite the policy activist. Instead, this paper concentrates its energies on providing a careful analysis of the historical patterns of output growth, factor accumulation, and productivity growth in the newly industrializing countries (NICs) of East Asia, i.e., Hong Kong, Singapore, South Korea, and Taiwan” (p. 640).
- ⁵ As we mentioned, Kremer et al. (2021) have argued that the data has moved in the direction of absolute convergence across countries in the 21st century.
- ⁶ This allegiance is also behind their just as playful title, “A Contribution to the Empirics of Economic Growth”, which substitutes empirics for the theory from Solow’s original article.
- ⁷ For instance, the aforementioned work by Jones and Olken (2005) shows that the effect of leaders is present in non-democracies, but not in democracies, suggesting that luck of this sort may matter insofar as it interacts with (in this case) institutional features.
- ⁸ Diamond (2013).
- ⁹ How can that be? Think about what happens, in the context of the Solow model, when population increases.

References

- Acemoglu, D. (2009). *Introduction to modern economic growth*. Princeton University Press.
- Acemoglu, D. & Dell, M. (2010). Productivity differences between and within countries. *American Economic Journal: Macroeconomics*, 2(1), 169–88.
- Acemoglu, D. & Johnson, S. (2007). Disease and development: The effect of life expectancy on economic growth. *Journal of Political Economy*, 115(6), 925–985.
- Acemoglu, D., Johnson, S., & Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5), 1369–1401.
- Acemoglu, D., Johnson, S., & Robinson, J. A. (2002). Reversal of fortune: Geography and institutions in the making of the modern world income distribution. *The Quarterly Journal of Economics*, 117(4), 1231–1294.
- Acemoglu, D., Johnson, S., & Robinson, J. A. (2005). Institutions as a fundamental cause of long-run growth. *Handbook of Economic Growth*, 1, 385–472.
- Acemoglu, D. & Robinson, J. A. (2012). *Why nations fail: The origins of power, prosperity, and poverty*. Currency.

- Acemoglu, D. & Robinson, J. A. (2017). *The emergence of weak, despotic and inclusive states*. National Bureau of Economic Research.
- Acemoglu, D. & Robinson, J. A. (2019). *The narrow corridor: How nations struggle for liberty*. Penguin UK.
- Alesina, A. & Giuliano, P. (2015). Culture and institutions. *Journal of Economic Literature*, 53(4), 898–944.
- Barro, R. J. (1991). Economic growth in a cross section of countries. *The Quarterly Journal of Economics*, 106(2), 407–443.
- Barro, R. J. & Sala-i-Martin, X. (1992). Convergence. *Journal of Political Economy*, 100(2), 223–251.
- Baumol, W. J. (1986). Productivity growth, convergence, and welfare: What the long-run data show. *The American Economic Review*, 1072–1085.
- Campante, F. & Yanagizawa-Drott, D. (2015). Does religion affect economic growth and happiness? Evidence from Ramadan. *The Quarterly Journal of Economics*, 130(2), 615–658.
- De Long, J. B. (1988). Productivity growth, convergence, and welfare: Comment. *The American Economic Review*, 78(5), 1138–1154.
- Diamond, J. (2013). *Guns, germs and steel: A short history of everybody for the last 13,000 years*. Random House.
- Glaeser, E. L., La Porta, R., Lopez-de-Silanes, F., & Shleifer, A. (2004). Do institutions cause growth? *Journal of Economic Growth*, 9(3), 271–303.
- Guiso, L., Sapienza, P., & Zingales, L. (2006). Does culture affect economic outcomes? *Journal of Economic Perspectives*, 20(2), 23–48.
- Hall, R. E. & Jones, C. I. (1999). Why do some countries produce so much more output per worker than others? *The Quarterly Journal of Economics*, 114(1), 83–116.
- Jones, B. F. & Olken, B. A. (2005). Do leaders matter? National leadership and growth since World War II. *The Quarterly Journal of Economics*, 120(3), 835–864.
- Kremer, M., Willis, J., & You, Y. (2021). Converging to convergence. *NBER Macro Annual 2021*. <https://www.nber.org/system/files/chapters/c14560/c14560.pdf>.
- Krugman, P. (1994). The myth of Asia's miracle. *Foreign Affairs*, 73(6), 62–78.
- Landes, D. S. (1998). Culture counts. *Challenge*, 41(4), 14–30.
- Levine, R. & Renelt, D. (1991). *Cross-country studies of growth and policy: Methodological, conceptual, and statistical problems* (Vol. 608). World Bank Publications.
- Mankiw, N. G., Romer, D., & Weil, D. N. (1992). A contribution to the empirics of economic growth. *The Quarterly Journal of Economics*, 107(2), 407–437.
- Mokyr, J. (2018). *The British industrial revolution: An economic perspective*. Routledge.
- North, D. (1990). *Institutions, institutional change and economic performance*. Cambridge University Press: New York.
- Nunn, N. (2020). The historical roots of economic development. *Science*, 367(6485).
- Rodrik, D. (2012). Why we learn nothing from regressing economic growth on policies. *Seoul Journal of Economics*, 25.
- Sachs, J. & Malaney, P. (2002). The economic and social burden of malaria. *Nature*, 415(6872), 680–685.
- Sachs, J. D. & Warner, A. M. (2001). The curse of natural resources. *European Economic Review*, 45(4-6), 827–838.
- Sala-i-Martin, X. X. (1997). *I just ran four million regressions* (tech. rep.) National Bureau of Economic Research.

- Solow, R. M. (1957). Technical change and the aggregate production function. *The Review of Economics and Statistics*, 312–320.
- Summers, R. & Heston, A. (1988). A new set of international comparisons of real product and price levels estimates for 130 countries, 1950–1985. *Review of Income and Wealth*, 34(1), 1–25.
- Young, A. (1995). The tyranny of numbers: Confronting the statistical realities of the east asian growth experience. *The Quarterly Journal of Economics*, 110(3), 641–680.