

1 Endogenous growth: What matters are new ideas

So far, we have only considered models of exogenous technological progress. However, understanding the source of technological progress is paramount for at least three reasons. First, even in models of technological progress, such as the Solow model, technological progress is the ultimate source of economic flourishing over time. Understanding the source and, thus, possible policy interventions promises large benefits. Second, cross-country differences in technology are the main driver behind output per worker differences. That is, if we want Spain to raise its output per worker level to the level of the U.S., we need to understand how Spain could adopt more productive production technologies. Third, growth miracles result to a large part from fast technological growth. Hence, if we want Madagascar to follow the path of South Korea and escape poverty, we need to understand how Madagascar can adopt better production technologies

We will begin our analysis in understanding technological growth in developed economies, i.e., countries that operate at the technological frontier and adopt better technologies by inventing these. Afterward, we will study the case of a country that is below the technological frontier and has the possibility to use better technologies by copying them from more advanced economies. The model is based on [Romer \(1990\)](#), who won the [Nobel price](#) for his contribution to (endogenous) growth theory.

1.1 How ideas are discovered

If we think about what contributes to the discovery of new ideas, there are ample contributing factors. Some of the most important are the number of researchers and the research environment in a country, such as the quality of research institutes, universities, and the amount of knowledge exchange with other researchers. One may think that also the amount of capital investment matters. For example, the search for a nuclear fusion reactor has cost tens of billions of euros in physical capital investment over the years. Today, developing AI is costing billions in investment into microprocessors. However, research and development is a very labor

intensive sector overall and, to keep the model simple, we will assume it needs no capital investment. The research sector is very human capital intensive, thus, having a model of human capital accumulation may be desirable but, again, for simplicity we abstract from it.

A factor that is less obviously affecting the speed of discovering new ideas is the stock of existing ideas. There are two reasons why we may think that the stock of existing ideas is important for the discovery of new ideas. The first is the famous concept of *standing on the shoulders of giants*. The idea is that the stock of accumulated ideas makes it easier to develop new ideas. For example, it would have been impossible to calculate the trajectory of the rocket that eventually travel to the moon, if Newton and Leibniz had not invented calculus. The second is the concept of the *easiest fruits are already picked*. The idea is that idea discovery is not random. Instead, researchers, by chance or by choice, will discover first ideas that are relatively easy to discover. For example, in medicine, in 1847, the first institute implemented the policy that medical staff ought to disinfect their hands before assisting in child birth. The result was a reduction of maternal mortality from 10 to 2%, i.e, a factor of five. Today, in Spain, this rate is 3/100000, and it is hard to imagine that we find any time soon a technological innovation that would decrease this rate to 3/500000.

Given this discussion, we write the process of ideas development as

$$\dot{A}(t) = f\left(\theta, L_A(t), A(t)\right), \quad (1)$$

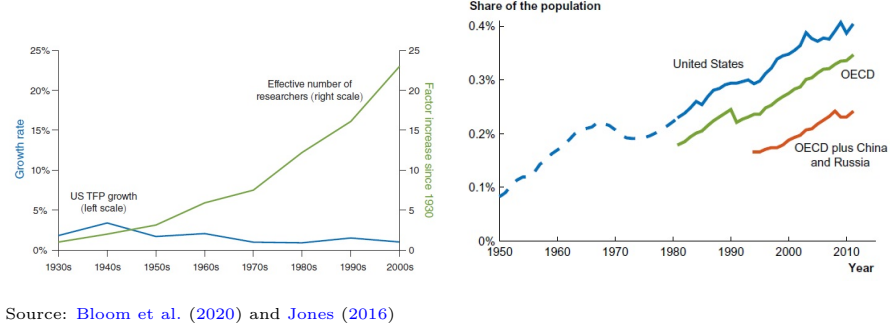
where θ is the research environment, and $L_A(t)$ are the number of researchers. We would like to learn from the data the properties of the function f . It is, however, simpler to look at the data in terms of growth rates, where I do not impose constant returns-to-scale of f :

$$\frac{\dot{A}(t)}{A(t)} = \frac{f\left(\theta, L_A(t), A(t)\right)}{A(t)}. \quad (2)$$

1.1.1 Macro-level evidence

We begin to study macroeconomic data to gain insights into the function f . The blue line in the left panel of Figure 1 shows a fact we have already seen before:

Figure 1: Technological progress in the data



In the U.S., over the last decades, the TFP growth rate is close to constant (with some slow-down since 1970). A constant growth rate may suggest that we can write

$$\dot{A}(t) = A(t)f(\theta, L_A(t)), \quad (3)$$

i.e., the growth rate does not depend on the current level of the stock of ideas. However, as the green line shows, this would be the wrong conclusion. We observe that $L_A(t)$ is growing over the same time period with an exponential process in the U.S. Hence, if the growth rate was independent of $A(t)$, the growth rate would have to grow over time given that the number of researchers is growing, instead of being constant. The only way these facts can be reconciled is if

$$\frac{\dot{A}(t)}{A(t)} = \frac{f(\theta, L_A(t), A(t))}{A(t)} \quad (4)$$

is a decreasing function of $A(t)$. Note, this does not mean that the stock of current ideas makes the discoveries of new ideas more difficult. To see this, consider a simple example with

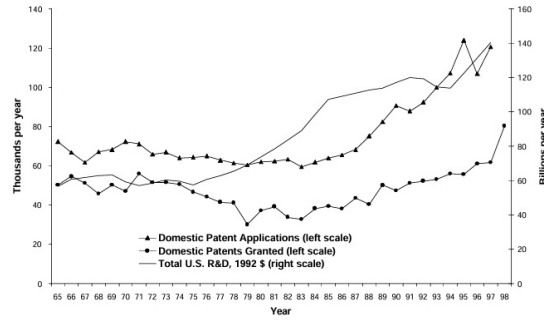
$$\dot{A}(t) = \sqrt{A(t)}. \quad (5)$$

Clearly, the stock of current ideas increases the discoveries of new ideas. However, the rate is too slow leading to a ever decreasing growth rate. The key is again to realize how extraordinary a constant growth rate is. For example, consider an economy where θ and $L_A(t)$ are constant such that $A(t)$ is the only variable in f that is time-varying, Assume our current stock of ideas is 10, and assume the

function f is such that we discover one additional idea, hence, the growth rate of ideas is 10%. Now imagine that our current stock reaches 100 ideas, and we discover 9 additional ideas. Clearly, $A(t)$ has a positive impact on the number of new ideas we are discovering. Nevertheless, the growth rate of new ideas is slowing down to 9% in this example.

So far, we measure the number of U.S. researchers. One may think the correct measure should include researchers that interact with U.S. researchers. For example, one of the key COVID-19 vaccines was developed in Germany for a U.S. company. The right panel shows that this makes the point only stronger: The number of researchers is growing across many countries of the world.

Figure 2: Patents and patent adoption




Source: [Jaffe \(1999\)](#)

Another objection is that TFP is a measure that is hard to interpret as it is effectively a residual. Technological innovations are often associated with patents, hence, the number of new patents may be a better measure of technological progress. The left panel of Figure 2 shows that patent data does not show a more positive picture. It shows that the number of patents granted to U.S. persons (green line) was basically flat between 1900 and 1980. It started to increase since then but much of the increase is associated with increasing R&D expenditures. Moreover, given that the increase does not translate to increasing TFP growth may suggest that the innovations have become more marginal.

Finally, one may worry that inventions are not what drives productivity but the adoption of those inventions into the economy. For example, AI will only increase productivity of our modern economies when people widely make use of it to improve their production processes. Figure 3 provides suggestive evidence that

Figure 3: Adoption of new ideas



Source: Jones (2016)

a slow-down in adoption of new technologies is, however, not behind the missing acceleration in productivity growth. Major technological innovations from the 18th century, such as the railroad, took about 100 years until they were widely adopted by the economy. Today, the adoption lag is closer to 10 years.

1.1.2 Micro-level evidence

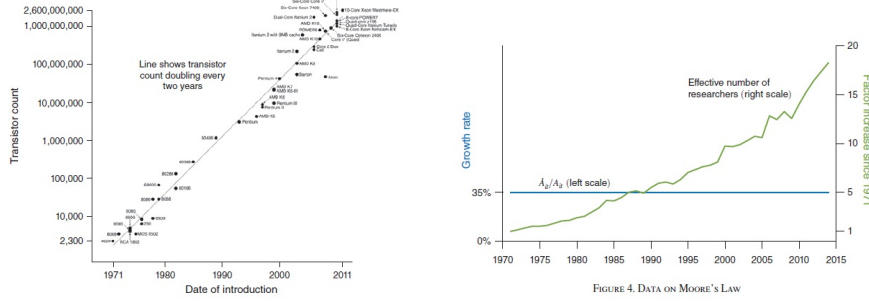
As with any evidence based on macroeconomic data, one may raise several questions: For example, is the quality of the researchers comparable over time? Are patents comparable over time? In a recent contribution, Bloom et al. (2020), instead of studying the macroeconomy as a whole, study three particular cases where they can observe the input and output in the ideas production function. hence, they can study changes in research productivity within a field over time. On average, they find that research productivity is falling by around 8% per year.

The first example comes from a well-known fact from computer science. Since the 1970s, the number of transistors on a CPU doubles every 2 years¹ implying an impressive example of a constant yearly technological growth rate of 35%. Importantly, the authors also have access to wage and *R&D* expenditure data from the major chip manufactures and find that today we need 18 times more input for the same growth than in the 1970s. Put differently, the growth rate per researcher, the research productivity, is falling drastically over time.

Their second example comes from major crops. Figure 5 shows that the increase

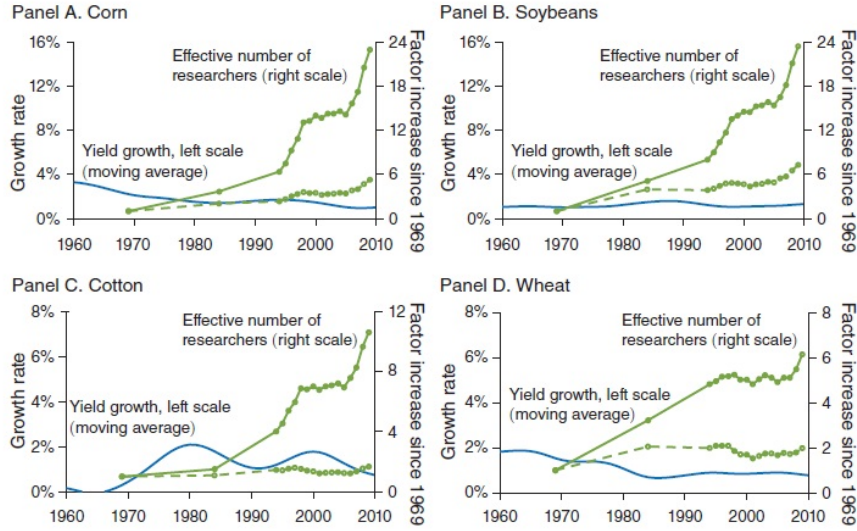
¹The number of transistors determine the number of operations a CPU can perform per cycle.

Figure 4: Moore's law



Source: Bloom et al. (2020)

Figure 5: Crop yields




Source: Bloom et al. (2020)

in crop yields was relatively constant since the 1960s. However, at the same time, the number of researchers trying to increase crop yields has risen sharply. Looking across the different crops, the data suggests a 6% yearly productivity decline.

Their third example comes from the field of health. In particular, they study the mortality rate from two common diseases in the U.S., breast cancer and heart disease. Both diseases show a close to linear decline in mortality rates since the 1970s. A linear decline in mortality rates implies that the number of years of life saved is *declining* over time, i.e., the health sector is generally not a sector

Figure 6: Disease mortality



Source: Bloom et al. (2020)

showing exponential growth. This is similar to our discussion about declining maternal mortality rates earlier. Even worse, the number of researchers active in the two fields, again, has increased drastically. This is irrespective of two different researcher measures: the number of clinical trials or the number of scientific publications.

Microeconomic evidence also has its problems when trying to understand a macroeconomic ideas production function. The most important one is substitution. Take the example of Moore's law. It may be true that fitting more and more transistors on a CPU chip becomes more and more difficult. However, we do not care by itself about the number of transistors but rather about the computing power of a chip. Once we think in this way, it is obvious that firms will find ways to keep increasing computing power without increasing the number of transistors on the CPU: decrease the time of a cycle that it takes a transistor to operate once (the Hertz), move away from CPU computing and towards GPU computing, or move away completely from classical computing towards quantum computing. However, given the joint evidence of macro measures and micro measures of ideas growth, we have to conclude that the data heavily suggests that a constant growth rate of ideas is only possible with ever increasing inputs.

1.2 Do we need a new framework

It may be tempting to think about technology as just another input we can accumulate (like capital) and just modify our Solow model by including a law-of-motion for productivity that depends on a savings decision. This, however, misses a key aspect of technology: Different from capital, technology is a [non-rivalrous](#) good. Non-rivalry means that by one person using a good, it does not diminish the usage for a different person. Applying this concept to ideas, if one person comes up with a better idea of producing, everyone could, in principle, use that idea. For example, when a drug company develops a new cancer drug, the usefulness of that drug does not decrease when all firms start producing that drug instead of inferior drugs.

The example, however, highlights an obvious problem with technological progress as a good in our current framework of the Solow model: Developing a new cancer drug is expensive, and firms will only be willing to engage in developing new drugs if they can sell them at a profit to recuperate their initial investment cost. However, in our Solow model, firms operate under perfect competition with constant returns to scale meaning they make zero profits on each unit they produce. With technological innovation it may be better to think of the production function looking something like


$$f(x) = 100(x - F), \tag{6}$$

where F are fixed-costs. Hence, $f(ax) > af(x)$, i.e., we have increasing returns to scale. Any market with fixed costs, where the fixed costs are not paid by a third party like the government, implies that firms need market power to make profits because average costs are always higher than marginal costs. For firms to make profits, we require that a good is excludable, i.e., the inventing firm can prevent other firms from using the new technology. If the innovation is not excludable, no private, profit-seeking firm will make the necessary investments and someone else has to provide the innovations. A classical example are innovations in national defense that are often provided by the government.

Innovations can be made excludable in various ways. One way are so called trade secrets. They simply reflect the fact that some new ideas are hard to copy. The most famous example is the recipe for Coca Cola that could not be copied as

long as chemistry was not developed far enough to chemically analyze its content. Also firm organization innovations, such as new management styles may be hard to observe and difficult to copy from one firm to another. In other cases, e.g., new drugs, reverse engineering and copying may be relatively simple. To give firms and individuals nevertheless an incentive to innovate, the government grants temporary monopolies to innovations. In case of tangible ideas these are patents. With a patent, the innovating firm publishes its idea openly, however, the government provides it with a temporary monopoly on the usage. For non-tangible ideas, such as music songs, the government similarly grants copyrights.

Figure 7: Growth take-off and patent law



Given the centrality of patents for providing firms incentives to innovate, economists have asked whether the introduction of patent law might be able to explain why output per worker started to grow in the 17th century after having be stalled for over a thousand years prior. The growth take-off, i.e., the industrial revolution, started in England and, indeed, England established a patent law in 1624. Critics observe that it took, however, at least another 50 years before growth in income per person started to occur, as Figure 7 shows. Those economists that see patent law as critical respond that the diffusion of the law was slow as many people did not learn about it for decades, and the implementation developed only slowly in the common-law system of England. In the following, we will not model patents explicitly, thus, we do not take a stance on their essence. Instead, we will simply assume that innovations as excludable.

1.3 Model set up

Our model follows the ideas outlined in [Romer \(1990\)](#). So far, our discussion centered around a firm innovating, producing, and selling a product, i.e., Apple inventing the Iphone and producing it. Conceptually, it will be easier to think of the three activities as different sectors, like a design company that sells its latest design to a coffeemaker which produces the coffee machine and sells it on Amazon. This makes our analysis a little more streamlined because it allows us to use prices. However, you could also think of Apple's sales and distribution departments paying its factories for the production of the phones, and the factories pay the R&D department for each innovation. Hence, in total, our economy will have four sectors: The *household sector* saves and accumulates the aggregate, homogeneous capital stock (like in the Solow model). The *research sector* develops new product designs taking the form of capital goods, e.g., the design of a new computer. The *capital goods sector* buys these designs and uses capital from the household sector to turn it into a productive capital good on which it holds an exclusive right. Finally, the *final goods sector* buys these capital goods and produces the final good under perfect competition.

1.3.1 Final good sector

To produce the final output good, the economy requires differentiated capital goods. For example, to produce a car, one needs assembly lines, robotic arms, stamping machines, welding tools, car chassis, and combustion engines. Another example is a window for which we require the frame, the handle, and the glass. One can make similar examples for other final consumption goods of which there are thousands. As before, for simplicity, we will assume, however, that there is only one final production good. You can think of this as the *representative* final good, or an artificial final output good that does indeed use *all* differentiated capital goods that an economy has.

The final goods producer uses labor and a measure of $A(t)$ available capital

goods to produce the final output good:

$$Y(t) = L_Y(t)^{1-\alpha} \int_0^{A(t)} x_j(t)^\alpha dj. \quad (7)$$

Note, the function has diminishing marginal returns to each individual capital good $x_j(t)$, i.e., increasing the number of welding tools while keeping all other capital goods and labor fixed will increase output, however, the additional output gains become smaller as we keep increasing the number of welding tools. Moreover, the production function has constant returns to scale with respect to capital and labor. That is, if we have one factory using a certain number of capital goods and labor to produce $Y(t)$, we can set another factory with the same numbers of capital goods and labor next to it and double output. Finally, note that the economy always employs all discovered capital goods $A(t)$, i.e., they do not become obsolete. This is true for some capital goods in practice, such as the wooden wheel. However, you can also find examples where this is not the case such as manual type-setting printing presses. We will see below a model where, instead of using all capital goods, the final good producer only uses the latest capital good.

As before, we assume the final good producer operates under perfect competition. Hence, the final goods producer choose the quantity of each capital good and labor to maximize profits while taking the wage, w , and the prices of each capital good, p_j , as given. We will normalize the price of the final output good to 1, and, for readability, omit the time index for the moment:

$$\max_{L_Y, x_j} \left\{ \Pi = L_Y^{1-\alpha} \int_0^A x_j^\alpha dj - wL_Y - \int_0^A p_j x_j dj \right\}. \quad (8)$$

The first-order-condition for labor reads

$$\frac{\partial \Pi}{\partial L_Y} = (1 - \alpha) L_Y^{-\alpha} \int_0^A x_j^\alpha dj - w = 0 \quad (9)$$

$$w = (1 - \alpha) \frac{Y}{L_Y}, \quad (10)$$

i.e., the firm hires labor until the wage bill is a constant fraction of output. Simi-

larly, the first-order-condition for the capital good j is:

$$\frac{\partial \Pi}{\partial x_j} = \alpha L_Y^{1-\alpha} x_j^{\alpha-1} - p_j = 0 \quad (11)$$

$$p_j = \alpha L_Y^{1-\alpha} x_j^{\alpha-1}. \quad (12)$$

This equation is an inverse demand function for the capital good j . The price the firm is willing to pay is decreasing in the quantity demanded, x_j , because of decreasing returns to the capital input, and, hence, a falling marginal product. Note, there are in total $A(t)$ first-order-conditions (and resulting demand functions) for the capital goods, one for each. However, each first-order-condition is exactly the same and, hence, it is sufficient to know the one for the j th capital good.

1.3.2 The capital goods sector

The downward-sloping demand function provides capital-goods producers with market power. As discussed above, this market power is necessary so that they generate profits that they can use to pay for research. In specific, the capital good producers buy patent designs from researchers at a fixed price P_A . Once they have the design, they produce the capital goods. The only factor of production is capital that they transform at a one-to-one rate into the capital good. Conceptually, it is easiest to assume that capital goods producers rent the required capital from the household sector at a price r instead of owning the capital themselves. This is similar to the Solow model, where we also assumed that firms rent the capital stock from households instead of owning it. hence, the problem of the capital good producer j is to choose the quantity of good j that maximizes profits:

$$\max_{x_j} \{ \pi_j = p_j(x_j)x_j - rx_j \}, \quad (13)$$

where $p_j(x_j)$ is given by the demand function (12). The first-order-condition is given by

$$\frac{\partial \pi_j}{\partial x_j} = \frac{\partial p_j(x_j)}{\partial x_j} x_j + p_j - r = 0 \quad (14)$$

$$0 = \frac{\frac{\partial p_j(x_j)}{\partial x_j} x_j}{p_j} + 1 - \frac{r}{p_j}. \quad (15)$$

From (12) we have that the elasticity is

$$\frac{\frac{\partial p_j}{\partial x_j} x_j}{p_j} = \frac{(\alpha - 1)\alpha L_Y^{1-\alpha} x_j^{\alpha-2} x_j}{\alpha L_Y^{1-\alpha} x_j^{\alpha-1}} \quad (16)$$

$$= \alpha - 1. \quad (17)$$

Putting the two equations together gives us the optimal price:

$$p_j = \frac{1}{\alpha} r > r. \quad (18)$$

Note, the optimal price is a mark-up over marginal costs r . hence, imperfect competition implies that the capital goods producers make a profit on each unit they sell:

$$\pi_j = \frac{1}{\alpha} r x_j - r x_j = x_j r \frac{1 - \alpha}{\alpha}. \quad (19)$$

It is this profit that allows them to buy the patent for their capital good which incentivizes the research in the first place. Given the solution for profits, we can now determine the value of a patent. To derive the value, we equalize the benefits and costs of holding the patent for one period. The implied costs for having a patent for one period is its price times the one-period return of an alternative investment, where the latter is given by the return on capital: rP_A . Its one-period benefit is the flow profit plus the change in the value of the patent over that period:

$$rP_A = \pi + \dot{P}_A. \quad (20)$$

One can show that this solves for

$$P_A = \frac{\pi}{r - n}. \quad (21)$$

Obviously, higher profits, $\pi(t)$, increase the value of the patent. Moreover, the opportunity costs of buying a patent, i.e., the return on the alternative investment capital, $r(t)$, decrease the value of a patent. What is less obvious is why the population growth rate, n , increases the value of a patent. The reason is that part of the value of the patent come from future sales, represented above by the change in the value of the patent over time. A larger economy, resulting from population growth, increases future sales and, thus, future profits from the patent.

1.3.3 The research sector

We now use the insights we developed above about the process of new ideas to model the research sector. The number of ideas that a single researcher discovers in a period, $\bar{\theta}$, depends on his productivity, θ , i.e., what we called above the research environment of an economy. Moreover, consistent with the idea of *Standing on the shoulders of giants* we model that the current stock of ideas has a positive effect on research output: $A(t)^\phi$ with $\phi > 0$. Finally, I allow research output to depend on the total number of researchers:

$$\bar{\theta} = \theta A(t)^\phi L_A(t)^{\lambda-1}. \quad (22)$$

λ determines the returns to scale to the number of researchers. If $\lambda < 1$, there is a stepping on your toes effect, otherwise, there are network effects. To get to the total change in the number of new ideas in an economy, $\dot{A}(t)$, we have to multiply with the number of researchers, L_A :

$$\dot{A}(t) = \theta L_A(t)^\lambda A(t)^\phi. \quad (23)$$

Given the ideas accumulation equation, we can determine the number of researchers in our economy. In our economy, households can freely choose between working as production workers, $L_Y(t)$, or being a researcher. The return from the

latter is $\bar{\theta}P_A$ while the return from the former is $w(t) = (1 - \alpha)\frac{Y(t)}{L_Y(t)}$. Hence, we must have that

$$\bar{\theta}P_A = (1 - \alpha)\frac{Y}{L_Y}. \quad (24)$$

In steady state, one can show that this solves for the fraction of households working as researchers being

$$s_R = \frac{1}{1 + \frac{r-n}{\alpha g^*}}, \quad (25)$$

where g^* is the steady state growth rate of technology that we will derive below. Little surprising a higher growth rate makes being a researcher more attractive as researches have more ideas. Note, a lower savings rate will increase the number of researchers by increasing the value of a patent.

1.3.4 Market clearing

Note, each capital good x_j has the same cost and the same benefit to the final good producer. Moreover, there are diminishing returns to using each individually. Hence, the equilibrium outcome is that all are demanded/produced in the same quantity:

$$x_j = x. \quad (26)$$

Let us denote the aggregate capital stock by $K(t)$ which, as eluded to above, will be provided by the households. Capital goods market clearing implies, that demand equals supply:

$$\int_0^A x_j dj = Ax = K \quad (27)$$

$$x = \frac{K}{A}. \quad (28)$$

Plugging the result into our final goods production function gives:

$$Y(t) = L_Y(t)^{1-\alpha} A(t) x(t)^\alpha \quad (29)$$

$$Y(t) = L_Y(t)^{1-\alpha} A(t) \left(\frac{K(t)}{A(t)} \right)^\alpha \quad (30)$$

$$Y(t) = (A(t) L_Y(t))^{1-\alpha} K(t)^\alpha \quad (31)$$

which is our familiar production function. Different from before, we now have a theory of productivity $A(t)$. An economy is more productive when it knows a larger variety of capital goods. The reason is that by having more capital varieties, the capital stock is more productive because the economy is better in avoiding diminishing marginal returns to each individual capital good. This idea captures partially our discussion at the beginning of the Solow model about differentiated capital goods in agricultural production: Adding more and more tractors to the production process while holding all other factors of production constant will run into diminishing marginal returns. However, adding a drone to agricultural production makes the tractor more and not less productive.² As a result, the production function has increasing returns to scale with regard to all *three* production factors.

Having derived the aggregate production function, we obtain the rental price of capital as

$$r(t) = \frac{\partial Y(t)}{\partial K(t)} = \alpha (A(t) L_Y(t))^{1-\alpha} K(t)^{\alpha-1} \quad (32)$$

$$= \alpha \frac{Y(t)}{K(t)}. \quad (33)$$

1.3.5 The household sector

The household provides labor, $L(t)$, either by working in the goods producing sector, L_Y , or the research sector, L_A . The amount of workers is growing at a

²The model captures this idea only partially, because in our model, the economy becomes more productive by reallocating capital away from tractors to drones, i.e., the capital goods are substitutes. Yet, in the example, the two capital goods are actually complements where one makes the other more productive.

constant rate:

$$L(t) = L_Y(t) + L_A(t) \quad (34)$$

$$L_A(t) = s_R L(t) \quad (35)$$

$$\dot{L}(t) = nL(t). \quad (36)$$

The household is the owner of the capital stock and obtains income, $\tilde{Y}(t)$, from renting out the capital stock, working in the research sector, or working in the final good production sector:

$$\tilde{Y}(t) = r(t)K(t) + \int \pi_j(t) dj + L_Y(t)w(t), \quad (37)$$

where I have used the fact that income from being a researcher must equal total profits from capital producers. Substituting in profits and the wage yields

$$\tilde{Y}(t) = r(t)K(t) + \int x_j(t)r(t)\frac{1-\alpha}{\alpha}dj + (1-\alpha)Y(t) \quad (38)$$

$$\tilde{Y}(t) = r(t)K(t) + K(t)r(t)\frac{1-\alpha}{\alpha} + (1-\alpha)Y(t) \quad (39)$$

$$\tilde{Y}(t) = \frac{r(t)}{\alpha}K(t) + (1-\alpha)Y(t) \quad (40)$$

$$\tilde{Y}(t) = Y(t), \quad (41)$$

where the last line follows from substituting in $r(t)$. Hence, as before, all production is redistributed to the households in form of income. As before, we assume households have a constant savings rate out of income. Hence, we have the already familiar law of motion for aggregate capital accumulation:

$$\dot{K}(t) = sY(t) - \delta K(t). \quad (42)$$

Thus, apart from $L_Y(t) < L(t)$, the aggregate economy looks the same as the Solow model. Note, however, the big difference: We have now a theory of the level of productivity, $A(t)$, and its law of motion $\dot{A}(t)$.

1.4 The steady state

Given our familiar production function and laws of motion, it is natural to think that we will again be able to find a steady state for the capital-to-output ratio. From the production function we obtain our familiar first steady state relationship:

$$z(t) = \frac{K(t)}{Y(t)} = \frac{K(t)}{K(t)^\alpha (A(t)L_Y(t))^{1-\alpha}} \quad (43)$$

$$= \left(\frac{K(t)}{A(t)L_Y(t)} \right)^{1-\alpha} \quad (44)$$

$$\frac{\dot{z}(t)}{z(t)} = (1-\alpha) \frac{\dot{K}(t)}{K(t)} - (1-\alpha)(n+g(t)) \quad (45)$$

$$\left(\frac{\dot{K}(t)}{K(t)} \right)^* = n + g(t). \quad (46)$$

Note, the key difference to the Solow model is that we do not know yet whether $g(t)$ is a constant in steady state. As always, our second steady state condition comes from the capital accumulation equation:

$$\dot{K}(t) = sK(t)^\alpha (A(t)L_Y(t))^{1-\alpha} - \delta K(t) \quad (47)$$

$$\frac{\dot{K}(t)}{K(t)} = \frac{s}{z(t)} - \delta \quad (48)$$

Combining the equations and yields

$$n + g(t) = \frac{s}{z^*} - \delta \quad (49)$$

$$z^* = \left(\frac{K(t)}{Y(t)} \right)^* = \frac{s}{n + g(t) + \delta}, \quad (50)$$

which is a constant if $g(t)$ has a steady state. To find this steady state, recall the ideas accumulation equation:

$$\dot{A}(t) = \theta L_A(t)^\lambda A(t)^\phi = \theta (s_R L(t))^\lambda A(t)^\phi. \quad (51)$$

It is easy to see that the growth rate of new ideas is constant, i.e., we have as steady state, when the current stock of ideas does not affect new idea creation,

$\phi = 1$ and the labor force is constant $L(t) = L$:

$$\frac{\dot{A}(t)}{A(t)} = \theta(s_R L)^\lambda, \quad (52)$$

i.e., we have constant exponential growth of ideas.

However, the data analyzed above heavily suggests that $\phi < 1$ and $s_R L(t)$ is increasing over time. In general, the ideas accumulation equation is given by

$$\frac{\dot{A}(t)}{A(t)} = \frac{\theta(s_R L(t))^\lambda}{A(t)^{1-\phi}}. \quad (53)$$

Note, the growth rate is still a constant if the numerator and denominator grow at the same rate, i.e.,

$$(1 - \phi) \frac{\dot{A}(t)}{A(t)} = \lambda \frac{\dot{L}(t)}{L(t)} = \lambda n \quad (54)$$

$$\frac{\dot{A}(t)}{A(t)} = \frac{\lambda n}{1 - \phi} = g^*, \quad (55)$$

which is indeed constant. Hence, our model has a steady state where the growth rate of technology is constant and, hence, we also have a steady state with a constant capital-to-output ratio. Note the implications of the steady-state technology growth, g^* . First, constant technological progress is only possible through population growth (or a growing share of people doing research). More people provide more ideas which increases output. This is obviously a much more positive view of population growth than the one one would obtain from the Solow model. Second, the share of researchers in the population does not affect the growth rate of technology. However, we will see below that it affects the level of technology. Third, network effects, λ , and stepping on the shoulders of giants, ϕ , create faster technological progress.

Next, we turn to output in steady state. Using the production function, and

substituting the steady state capital-to-output ratio, we have

$$Y(t) = K(t)^\alpha (A(t)L_Y(t))^{1-\alpha} \quad (56)$$

$$Y(t) = \left(\frac{K(t)}{Y(t)} \right)^{\frac{\alpha}{1-\alpha}} A(t)L_Y(t) \quad (57)$$

$$\frac{Y(t)}{L(t)} = \left(\frac{K(t)}{Y(t)} \right)^{\frac{\alpha}{1-\alpha}(1-s_R)} A(t) \quad (58)$$

$$\left(\frac{Y(t)}{L(t)} \right)^* = \left(\frac{s}{n + g^* + \delta} \right)^{\frac{\alpha}{1-\alpha}} (1 - s_R)A(t). \quad (59)$$

The final step is to derive $A(t)$ for which, again, we have now a theory. In particular, we need to solve the differential equation

$$\dot{A}(t) = \theta(s_R L(t))^\lambda A(t)^\phi. \quad (60)$$

To that end, define the auxiliary variable

$$v(t) = A(t)^{1-\phi} \quad (61)$$

with

$$\dot{v}(t) = (1 - \phi)A(t)^{-\phi} \dot{A}(t). \quad (62)$$

Substituting yields

$$\dot{v}(t) = (1 - \phi)\theta(s_R L(0))^\lambda \exp(\lambda n t). \quad (63)$$

Note, the antiderivative of $\exp(\lambda n t)$ is $\frac{1}{\lambda n} \exp(\lambda n t)$. Hence, the solution to the differential equation is

$$v(t) = \frac{(1 - \phi)}{n\lambda} \theta(s_R L(0))^\lambda \exp(\lambda n t) + l, \quad (64)$$

where l is an integration constant. Evaluating at $v(0)$ provides the integration constant:

$$l = v(0) - \frac{(1 - \phi)}{n\lambda} \theta(s_R L(0))^\lambda \quad (65)$$

Plugging in and substituting again $v(t) = A(t)^{1-\phi}$ yields:

$$A(t)^{1-\phi} = A(0)^{1-\phi} + \frac{(1-\phi)}{n\lambda} \theta (s_R L(0))^\lambda [\exp(\lambda n t) - 1] \quad (66)$$

$$A(t) = \left[A(0)^{1-\phi} + \frac{(1-\phi)}{n\lambda} \theta (s_R L(0))^\lambda [\exp(\lambda n t) - 1] \right]^{\frac{1}{1-\phi}}. \quad (67)$$

Little surprising, a higher initial level of productivity, $A(0)$, implies higher productivity today. One can also show that a higher population growth rate implies higher productivity today, i.e., $\frac{\partial A(t)}{\partial n} \geq 0$. Moreover, more productive researchers, θ , a higher share of researchers, s_R , or a larger initial workforce, $L(0)$, increase productivity.

Putting everything together, we have in steady state for output per worker:

$$\left(\frac{Y(t)}{L(t)} \right)^* = \left(\frac{s}{n + g^* + \delta} \right)^{\frac{\alpha}{1-\alpha}} (1 - s_R) \left[A(0)^{1-\phi} + \frac{(1-\phi)}{n\lambda} \theta (s_R L(0))^\lambda [\exp(\lambda n t) - 1] \right]^{\frac{1}{1-\phi}}. \quad (68)$$

Note, the population growth rate has an ambivalent effect on output per worker. On the one hand, it has a “Solow effect”: A higher n decreases the capital-to-output ratio which decreases output per worker. On the other hand, it has a “Romer effect”: A higher n increases the growth rate (and resulting stock) of ideas.

Also the share of researchers has an ambivalent effect on output per worker. On the one hand, more researchers reduce the pool of available workers to produce the final output good. On the other hand, more researchers lead to a higher stock of ideas and, hence, make those workers producing the final output good more productive.

1.5 Transition dynamics

Neither the share of researchers, s_R , nor their efficiency, θ , drive the long run growth rate in technology, g^* . Nevertheless, the law of motion for productivity implies that changes in these variables will temporarily change the growth rate

of technology. Before solving explicitly for the transition path, it is useful to understand qualitatively what is going to happen over time. With $\phi < 1$, after a one-time increase in either of the two, we have

$$(1 - \phi) \frac{\dot{A}(t)}{A(t)} > \lambda n, \quad (69)$$

i.e., the denominator in the growth rate of technology grows quicker than the numerator leading to temporarily higher growth. Over time, as the stock of ideas accumulates, $\frac{\theta(s_R L(t))^\lambda}{A(t)^{1-\phi}}$ falls, i.e., the rate of technological progress falls.

To solve for the transition path explicitly, we can use the solution for the level of productivity. Combining the two equations:

$$\frac{\dot{A}(t)}{A(t)} = \frac{\theta(s_R L(t))^\lambda}{A(t)^{1-\phi}} \quad (70)$$

$$A(t) = \left[A(0)^{1-\phi} + \frac{1-\phi}{n\lambda} \theta(s_R L(0))^\lambda [\exp(\lambda n t) - 1] \right]^{\frac{1}{1-\phi}} \quad (71)$$

yields

$$\frac{\dot{A}(t)}{A(t)} = \frac{\theta(s_R L(t))^\lambda}{A(0)^{1-\phi} + \frac{1-\phi}{n\lambda} \theta(s_R L(0))^\lambda [\exp(\lambda n t) - 1]} \quad (72)$$

$$= \frac{\lambda n}{1 - \phi} \frac{\theta(s_R L(t))^\lambda}{\frac{\lambda n}{1-\phi} A(0)^{1-\phi} + \theta(s_R L(t))^\lambda - \theta(s_R L(0))^\lambda} \quad (73)$$


Rearranging the equation slightly gives us

$$\frac{\dot{A}(t)}{A(t)} = \frac{\lambda n}{1 - \phi} \frac{1}{\frac{\lambda n}{1-\phi} \frac{A(0)^{1-\phi}}{\theta(s_R L(t))^\lambda} + 1 - \left(\frac{L(0)}{L(t)} \right)^\lambda}. \quad (74)$$

There are several points worth highlighting. First, when starting in steady state, i.e., $\frac{A(0)^{1-\phi}}{\theta(s_R L(0))^\lambda} = \frac{1}{g^*} = \frac{1-\phi}{\lambda n}$, and $L(0) = L(t)$, then $\frac{\dot{A}(t)}{A(t)} = \frac{\lambda n}{1-\phi}$. This point is obvious. All it says is that once we are in steady state, technology is growing at its steady state growth rate derived above. Second, when the labor force is growing, we will converge to steady state. To see this, note that for any $A(0), L(0)$, as time passes, $L(t)$ grows and $\frac{\dot{A}(t)}{A(t)} \mapsto \frac{\lambda n}{1-\phi}$. Third, the lower $A(0)$ is relative to steady state, the

faster is the initial growth rate. Finally, for any $A(0), L(0)$, increasing $s_R, \theta, \phi, \lambda$, or n increases $\frac{\dot{A}(t)}{A(t)}$.

Figure 8: Productivity dynamics research efficiency



The left panel of Figure 8 displays the growth rate of productivity after a permanent increase in research efficiency (in period 100). You may think of a reform such as improving the research centers in an economy. Initially, the growth rate of technology jumps but it converges back to its steady state over time. The right panel shows the implications for the (log) level of productivity. In the long-run, the trajectory runs parallel to the old trajectory, however, the economy has a permanently higher stock of ideas.

Our prime interest is in the transition dynamics of output per worker after changes in the environment such as a permanently higher research efficiency. It may be tempting to think that output per worker simply follows the transition dynamics as productivity. To see that this is not the case, let us write output per worker as a function of the capital to output ratio

$$y(t) = A(t)(1 - s_R) \left(\frac{K(t)}{Y(t)} \right)^{\frac{\alpha}{1-\alpha}} \quad (75)$$

$$\frac{\dot{y}(t)}{y(t)} = g(t) + \frac{\alpha}{1-\alpha} \frac{\dot{z}(t)}{z(t)} \quad (76)$$

with

$$\frac{\dot{z}(t)}{z(t)} = (1 - \alpha) \frac{s}{z(t)} - (1 - \alpha) (n + g(t) + \delta). \quad (77)$$

Note, that the growth rate of the capital-to-output ratio is negatively affected by

the growth rate in technology. Combining these equations yields

$$\frac{\dot{y}(t)}{y(t)} = g(t) + \alpha \left[\frac{s}{z(t)} - (n + g(t) + \delta) \right] \quad (78)$$

In steady state, $y(t)$ grows at rate g^* . A temporary increase in $g(t)$ increases the growth rate of $y(t)$ but it increases it initially by less than the increase in $g(t)$ because the capital to output ratio declines initially.

Figure 9: Output dynamics research efficiency





Figure 9 shows the effect graphically. Right after the reform, output per worker growth is lower than productivity growth as the capital-to-output ratio is falling. After an initial period of divergence, output per worker and productivity growth are approximately the same again and both converge back to their steady state growth rate.

Figure 10: Capital-to-output ratio and TFP in Europe



We end the discussion on the Romer model by looking at the data through the lens of the model. As discussed before, around 1970, TFP growth slowed down in several developed economies. The slowdown was particularly severe in several European countries. The left panel of Figure 10 shows this trend for three European countries. Note, the slowdown in productivity growth looks permanent and, hence, according to our model, cannot be caused by decreases in the share of researchers s_R or the research efficiency θ . Instead, a permanent slowdown can only be caused by a decline in the population growth rate n , decreasing network effects between researchers λ , or lower research returns on the current stock of ideas ϕ . We will revisit the cause of the productivity slowdown later. For now, we just ask whether the model-implied changes are consistent with data moments. In particular, the model predicts that the capital-to-output ratio should increase after a decline in the productivity growth rate. The right panel shows that this is indeed what occurs in the data, particularly in Spain and Italy where the ratio rose from around 4.2 in 1970 to about 6.5 in 2010. If the TFP growth rate permanently declined in one discrete step in 1970, the model would predict that the capital-to-output ratio should converge in a convex fashion to its new steady state which is not what we observe in the data. However, the left panel shows that this stylized thought experiment is not describing the data well. Instead, TFP growth fell much smoother throughout the 70s and the 80s.

1.6 Schumpeterian growth

In the Romer model, once discovered, designs of capital goods are used forever and, thus, increase the productivity of an economy forever. An alternative way to think of technological progress goes back to [Schumpeter](#): technological progress as creative destruction, i.e., [better designs replacing older designs](#). It is easy to come up with examples. In the case of smart phones, touchscreen phones have mostly replaced keyboard phones. In the case of aviation, the jet engine has mostly replaced propellers. Though the old designs may still be in use to some degree, the vast majority of a developed economy has fully adopted the newer, better design. The first economists that formalized these ideas in endogenous growth models were [Grossman and Helpman \(1991\)](#) and [Aghion and Howitt \(1992\)](#).

1.6.1 The different sectors

The model is set up analogously to the Romer model with four different sectors. Again the set up helps us conceptually because we obtain prices between sectors but one may think again about a single firm being fully vertically integrated.

Final good sector The final good producer use again production labor and a capital good for production. Different from the Romer model, they use only the latest version of the capital good, x_i , where previous innovation of the capital good are denoted by $x_{i-1}, x_{i-2}, \dots, x_0$. Using the example of aviation from above, x_i would be the jet engine, and x_{i-1} would be the propeller. A capital good has an associated productivity A_i with $A_i > A_{i-1}$. Again, using the aviation example, A_i is the speed, cost of production, and fuel efficiency of the jet engine, and A_{i-1} is the speed, cost of production, and fuel efficiency of the propeller. The final good is produced according to

$$Y(t) = (L_Y(t)A_i(t))^{1-\alpha} x_i(t)^\alpha. \quad (79)$$

Note, to produce the final output good, only one capital good is needed as input. One could imagine a more complicated setting where there are differentiated capital goods that are needed, but each “line” of capital goods, j , has a order of productivities. In the aviation example, $j = 1$ could be the engine with $x_i^1, x_{i-1}^1, \dots, x_0^1$ and $j = 2$ could be the airframe with $x_i^2, x_{i-1}^2, \dots, x_0^2$ and associated productivities. This generalization would, however, increase our notation significantly without adding major economic insights to the model.

As in the Romer model, the final good producers operate under perfect competition and take the final output price (normalized to 1) and wages as given. They choose the quantity of the latest capital good and labor to maximize profits:

$$\Pi = \max_{L_Y, x_i} \{ (L_Y A_i)^{1-\alpha} x_i^\alpha - w L_Y - p_i x_i \}. \quad (80)$$

The first order conditions yield

$$w = (1 - \alpha) \frac{Y}{L_Y} \quad (81)$$

$$p_i = \alpha (L_Y A_i)^{1-\alpha} x_i^{\alpha-1}. \quad (82)$$

Capital goods sector The capital good producing sector is basically identical to the Romer model. Producers buy patent designs from researchers at a fixed price P_A . Once they have the design, they produce capital goods by transforming the capital from the household sector into productive capital goods at a one-to-one rate. Hence, we have again our familiar profit equation:

$$\max_{x_i} \{ \pi_i = p_i(x_i)x_i - rx_i \}, \quad (83)$$

which solves again for

$$p_i = \frac{1}{\alpha} r. \quad (84)$$

The research sector We assume that each new innovation is a constant improvement over the productivity of the last innovation:

$$A_i(t) = (1 + \gamma) A_{i-1}(t). \quad (85)$$

Next, we need to specify again the law of motion for new ideas. We assume again that the probability that a single researcher discovers a new design depends on the research environment and the number of other active researchers through scale effects. To capture again the idea of decreasing returns to the current state of knowledge, we assume that it is a decreasing function of the quality of existing ideas:

$$\bar{\mu} = \frac{\theta L_A(t)^{\lambda-1}}{A_i(t)^{1-\phi}} \quad \phi < 1. \quad (86)$$

Hence, the probability that a new design is discovered by any of the active researchers is

$$\bar{\mu}L_A(t) = \frac{\theta L_A(t)^\lambda}{A_i(t)^{1-\phi}}. \quad (87)$$

Note, the probability that a new idea is discovered that will make the old idea obsolete with change the value of a patent relative to the Romer model. To find the value, we equate again the implied costs of holding the patent for one period to the implied benefits, where the change in the value of the patent now includes the probability that its entire value, P_A is lost:

$$rP_A = \pi + \underbrace{\dot{P}_A - \bar{\mu}L_AP_A}_{\text{change in value}}. \quad (88)$$

One can show that this solves for

$$P_A = \frac{\pi}{r - n + \bar{\mu}L_A(1 - \gamma)}. \quad (89)$$

The more innovative a new patent is, γ , the more value it has. The more likely that a new innovation comes along, $\bar{\mu}$, the lower is the value.

Given the value of a patent, we can solve again for the share of researchers by equalizing the return of being a researcher for one period, $P_A\bar{\mu}$, equals the return of being a production worker, w . One can show that this solved for

$$s_R = \frac{1}{1 + \frac{r-n+\bar{\mu}L_A(1-\gamma)}{\alpha\bar{\mu}L_A}}. \quad (90)$$

Note, research efficiency, $\bar{m}u$, has two effects on the number of researchers that work against each other. On the one hand, it makes becoming a researcher more attractive as it raises the probability to develop a new idea. On the other hand, it reduces the value of a patent because the patent is more likely to be replaced by

future research. To see which effect dominates, consider

$$\frac{\partial s_R}{\partial \bar{\mu}} = \frac{\frac{r-n}{\bar{\mu}^2}}{\left[1 + \frac{r-n+\bar{\mu}L_A(1-\gamma)}{\alpha\bar{\mu}L_A}\right]^2}. \quad (91)$$

Recall from the discussion of the Golden rule, that in the data, $r > n$ for developed economies. Hence, our model predicts that an increase in the research efficiency still increases the number of researchers even with Schumpeterian growth.

Market clearing and the aggregate production function As the final goods producer uses only the most productive capital good, we have that this capital good uses all the aggregate capital stock³

$$x_i = K. \quad (92)$$

Substituting this result into the final producer's production function and writing $A(t) = A_i(t)$, we have again the same aggregate production function:

$$Y(t) = (L_Y(t)A_i(t))^{1-\alpha} x_i(t)^\alpha \quad (93)$$

$$Y(t) = (A(t)L_Y(t))^{1-\alpha} K(t)^\alpha \quad (94)$$

As the aggregate production function is the same, we also have the same equation for the real rental price of capital:

$$r(t) = \alpha \frac{Y(t)}{K(t)}, \quad (95)$$

which implies all prices are the same as in the Romer model. The components of households' income are also unchanged:

$$\tilde{Y}(t) = r(t)K(t) + \int \pi_j(t)dj + L_Y(t)w(t). \quad (96)$$

Hence, we have again the result that all output is given to the households in form of

³Note, we assume that, once a new capital good is invented, we can seamlessly convert the capital embodied in the old capital good into the new capital good.

income $\tilde{Y}(t) = Y(t)$. Assuming again that the households have a constant savings rate yields the familiar law of motion for capital:

$$\dot{K}(t) = sY(t) - \delta K(t). \quad (97)$$

At this point, one may conjecture that the aggregate economy of the Schumpeter model behaves exactly as in the Romer model. This is not exactly correct because the law of motion for productivity is different. Instead of technological progress following a deterministic, smooth law of motion, we now have a stochastic process that is only smooth in expectations:

$$\mathbb{E}A(t) = \bar{\mu}L_A(t)(1 + \gamma)A(t-1) + (1 - \bar{\mu}L_A(t))A(t-1) \quad (98)$$

$$\mathbb{E}A(t) = A(t-1)(1 + \bar{\mu}L_A(t)\gamma). \quad (99)$$

Figure 11: Technological progress in the U.S.




Figure 11 displays TFP in the U.S. over time. Indeed, the data suggests a stochastic rather than a deterministic process. For example, technological progress was rapid in the early 2000s but slowed down substantially starting in the mid-2000s.

References

- AGHION, P. AND P. HOWITT (1992): “A model of growth through creative destruction,” *Econometrica*, 60.
- BLOOM, N., C. I. JONES, J. VAN REENEN, AND M. WEBB (2020): “Are ideas getting harder to find?” *American Economic Review*, 110, 1104–1144.
- GROSSMAN, G. M. AND E. HELPMAN (1991): “Innovation and growth in the global environment,” .
- JAFFE, A. B. (1999): “The US patent system in transition: policy innovation and the innovation process,” *NBER Working paper*, 7280.
- JONES, C. I. (2016): “The facts of economic growth,” in *Handbook of macroeconomics*, Elsevier, vol. 2, 3–69.
- ROMER, P. M. (1990): “Endogenous technological change,” *Journal of political Economy*, 98, S71–S102.