Chapter 4
The Schumpeterian Framework

November 16, 2006

1 Introduction

In this chapter we develop an alternative model of endogenous growth (see Aghion and Howitt 1988, 1992a) where growth is generated by a random sequence of quality improving (or "vertical") innovations that result from (uncertain) research activities.[1] This model grew out of modern industrial organization theory[2]. This model of growth with vertical innovations has the natural property that new inventions make old technologies or products *obsolete*. Hence it involves the force that Schumpeter (1942) called "creative destruction", with the resulting view that the growth process implies a permanent conflict between the old and the new, between previous innovators who have become the current incumbents and the new potential innovators. This social interaction has far reaching implications for the relationship between growth and institutions or organizations as we will see in following chapters of the book.

The Chapter is organized as follows. Section 2 provides a preview of some main implications of the Schumpeterian paradigm, which it contrasts with the predictions of the product variety model analyzed in the previous chapter. Section 3 analyzes in detail a simple discrete time version of the Schumpeterian growth model. Section 4 develops a variant of the Schumpeterian paradigm, where time is continuous and labor is the R&D input.

[1] The earliest attempts at providing a Schumpeterian approach to endogenous growth theory were by Segerstrom, Anant, and Dinopoulos (1990), who modeled sustained growth as arising from a succession of product improvements in a fixed number of sectors, but with no uncertainty in the innovation process, and Corriveau (1991), who produced a discrete-time model with uncertainty about cost-reducing process innovations.

[2] See Tirole (1988).

2 A road map to Schumpeterian growth economics

2.1 The paradigm in a nutshell

Three main assumptions underlie the Schumpeterian growth paradigm: 1) growth is primarily driven by technological innovations; 2) innovations are produced by entrepreneurs who seek monopoly rents from them; 3) new technologies drive out old technologies.

To put things a little more formally, the Schumpeterian paradigm begins with a production function specified at the industry level:

$$Y_{it} = A_{it}^{1-\alpha} x_{it}^{\alpha}, \qquad 0 < \alpha < 1 \tag{1}$$

where A_{it} is a productivity parameter attached to the most recent technology used in industry i at time t. In this equation, x_{it} represents the flow of a unique intermediate product used in this sector. Aggregate output is just the sum of the industry-specific outputs Y_{it}.

Each intermediate product is produced and sold exclusively by the most recent innovator. A successful innovator in sector i improves the technology parameter A_{it} and is thus able to displace the previous innovator as the incumbent intermediate monopolist in that sector, until displaced by the next innovator.

Where do innovations come from? As in the product variety model analyzed in the previous section, there are two main inputs to technological progress; namely the private expenditures made by the prospective innovator, and the stock of innovations that have already been made by past innovators. The latter input constitutes the publicly available stock of knowledge to which current innovators are hoping to add. However, in contrast with product variety model, the Schumpeterian growth paradigm is quite flexible in modeling the contribution of past innovations. It encompasses not only the case of frontier innovations but also the case in which technology advances by implementing some pre-existing technology, an activity which is often just as costly as frontier innovation[3]. One way to think about the innovation-implementation distinction[4], is to suppose that innovation either leapfrogs the best technology available before the innovation, resulting in a new technology parameter A_{it} in the innovating sector i, which is equal to some multiple γ of its pre-existing value, or that it catches up to a global technology frontier $\overline{A}_t$ which represents the stock of global technological knowledge available to innovators in all sectors of all countries.

For example, consider a country in which in any sector leading-edge innovations take

[3]See Nelson and Phelps (1966), Benhabib and Spiegel (1994), Evenson and Westphal (1995), and especially Howitt and Mayer (2002) who first introduced the dichotomy between innovation and implementation into the Schumpeterian growth paradigm.

[4]See Benhabib and Spiegel (1994) and Acemoglu, Aghion and Zilibotti (2002).

place at the frequency μ_n and implementation innovations take place at the frequency μ_m. Then the change in the economy's aggregate productivity parameter A_t will be:

$$A_{t+1} - A_t = \mu_n (\gamma - 1) A_t + \mu_m \left(\overline{A}_t - A_t\right)$$

and hence the growth rate will be:

$$g_t = \frac{A_{t+1} - A_t}{A_t} = \mu_n (\gamma - 1) + \mu_m \left(a_t^{-1} - 1\right) \tag{2}$$

where:

$$a_t = A_t / \overline{A}_t$$

is an inverse measure of "distance to the frontier."

Like with the product variety model, the Schumpeterian paradigm derives the innovation frequencies μ_m and μ_n from research arbitrage equations that equates the innovator's marginal cost of the corresponding research activities to the marginal benefit from these activities. The solution to the innovator's maximization problems (and thus to its research arbitrage equations) will turn depend upon institutional characteristics of the economy such as property rights protection and the financial system, and also upon government policy.

2.2 Three ideas

The following features distinguish the Schumpeterian paradigm from the product variety model in a fundamental way.

First idea: *new innovations displace previous technologies and therefore faster growth should imply a higher rate of firm turnover, with more entry of new innovators and more exit of former innovators.*

This conflict between the old and the new has implications for the relationship between growth and institutions, as already mentioned above. For example, if we take a normative point of view, how can one design institutions and policies that reward current incumbents for their past innovations but without deterring entry by new innovators. Or design policies that facilitate new entry but without penalizing old innovators and skills too much? Or, if we now take a positive approach, what guarantees that incumbent firms will not lobby the government for policies aimed at deterring entry by new innovators? As it turns out, a major source of slowdown in technical progress appears to lie in the existence of *vested interests* among individuals (firm owners, managers, workers,..) specialized in the old technologies and thus tempted to collude and exert political pressure to delay the arrival of new innovations that would destroy their rents. Economic historians like Mokyr (1990) have argued that the same forces that produced the first industrial revolution in Britain in the mid 1800s opposed further technological progress later on,thereby contributing to the industrial slowdown of the late 1800s.

A second idea, which follows immediately from the growth equation (2), is Gerschenkron's "advantage of backwardness":[5] that is, the further the country is behind the global technology frontier (i.e., the smaller is a_t) the faster it will grow, given the frequency of implementation innovations. As in Gerschenkron's analysis, the advantage arises from the fact that implementation innovations allow the country to make larger quality improvements the further it has fallen behind the frontier. Through this channel, the model can account for cross country convergence.

Second idea: *a country's growth performance should vary with its proximity to the technological frontier a_t and convergence results from knowledge spillovers generated by the frontier.*

Note that convergence here results from cross-country knowledge interaction, whereas in the neoclassical model convergence resulted from the decreasing returns from capital accumulation experienced by each country in isolation, whereby per capita GDP in each country would converge to its steady-state value.

A third idea, also embedded in growth equation (2), is Gerschenkron's notion of "appropriate institutions". Suppose indeed that the institutions or policies[6] that favor implementation innovations (inducing firms to emphasize μ_m at the expense of μ_n) are not the same as those that favor leading-edge innovations (by encouraging firms to focus on μ_n); then[7]:

Third idea: *far from the frontier a country or sector will maximize growth by choosing institutions or policies that facilitate implementation, however as it catches up with the technological frontier, to sustain a high growth rate the country will have to shift from implementation-enhancing institutions/policies to innovation-enhancing ones as the relative importance of leading-edge innovations for growth is also increasing.*

Failure to operate such a policy shift can prevent a country from catching up with the frontier level of per capita GDP, and Sapir et al (2003) argued that this failure largely explains why Europe stopped catching up with US per capita GDP since of the mid !970s. The paradigm can also explain why, since the mid 1990s, the EU is growing at a lower rate than the US. This idea will be explored in several chapters of the book.

[5] See Gerschenkron (1952).

[6] Here, we are not distinguishing between institutions and policies. However, one may think of policies as decisions (e.g education spending) that can be changed or reversed easily (e.g by the executive or by simple majority vote in parliament). Institutions (e.g the political constitution, the organization of the financial systems or that of labor markets) evolve more slowly; they either require constitutional changes, or sufficient consensus building among the constituency.

[7] See Acemoglu, Aghion and Zilibotti (2002).

3 A discrete-time model

3.1 The model

Consider the following simple model of endogenous growth with quality-improving innovations. Time is discrete, and the size of population is constant. There is one final good, produced each period by perfectly competitive firms using one intermediate product (x) as input[8], according to the production function:[9]

$$Y = A^{1-\alpha} x^{\alpha}, \quad 0 < \alpha < 1, \tag{3}$$

where A is the productivity of the intermediate input.

Growth now results from quality-improving innovations. Each innovation creates a new intermediate product, which is more productive than previous ones and renders the previous ones obsolete.[10] So if there is an innovation at date t then the productivity of the intermediate good in (3) goes from last period's value A_{t-1} up to $A_t = \gamma A_{t-1}$, where $\gamma > 1$ denotes the size of innovations. On the other hand, if there is no innovation at t then the intermediate product used in t will be the same one that was used in $t-1$, so A_t remains equal to A_{t-1}.

Innovations result from R&D, which is a costly activity that uses the final good as its only input. R&D is also an uncertain activity, which may fail to generate any innovation. But the more is spent on R&D the more likely it is that an innovation will occur. Specifically, the probability that an innovation occurs during any period is assumed to be:[11]

$$\lambda \left(R/A \right)^{\sigma}, \; 0 < \sigma < 1$$

where R is the amount of final good spent on R&D, A is the productivity of the new intermediate product that will result if the R&D succeeds, and λ is a parameter that measures the productivity of the R&D sector.

The reason why the probability of innovation depends inversely on A is that as technology advances it becomes more complex and thus harder to improve upon. So it is not the absolute amount of R&D expenditure R that matters for success but the productivity-adjusted expenditure R/A, which we denote by n.

It follows that the growth rate of A will equal $\frac{\gamma A - A}{A} = \gamma - 1$ with probability λn^{σ}, and it

[8] In the following chapters we shall use extended versions of this model with multiple intermediate sectors.

[9] We are implicitly assuming a constant-returns-to-scale production function in (x) and labor: $Y = A^{1-\alpha} L^{1-\alpha} x^{\alpha}$, where the aggregate supply of labor L is fixed and measured in units such that $L = 1$.

[10] It might not be obvious how just how the new intermediate product drives out the old one. We defer to subsection 3.2 below our analysis of this issue.

[11] We assume that the time period is too short for more than one innovation to occur.

will equal zero with probability $1-\lambda n^{\sigma}$. Therefore the expected growth rate of A depends on productivity-adjusted R&D according to:

$$g = \lambda n^{\sigma}(\gamma - 1). \tag{4}$$

So if we want to analyze the underlying determinants of growth we first have to analyze the underlying determinants of R&D.

People are motivated to undertake R&D by the prospect of monopoly profits. That is, someone who innovates in period t will be able to monopolize the intermediate sector for that period, producing his new intermediate product and selling it to final-good producers. No one else can do this, because the innovator keeps secret the technology for producing the new product or he gets a patent making it illegal for anyone else to sell the product. We assume that the monopoly only lasts for one period however, because the secret gets out or the patent expires.

The innovator's monopoly profit π is determined as follows. First, he needs to use the final good as input, with one unit needed for each unit of intermediate product he produces, so his total cost will be the quantity produced x. Next, since the final sector is perfectly competitive, the price at which he can sell his intermediate good is its marginal product in that sector, namely:

$$p(x) = \frac{\partial Y}{\partial x} = \alpha A^{1-\alpha} x^{\alpha-1}. \tag{5}$$

The monopolist's profit maximization problem is then:

$$\begin{aligned} \pi &= \max_x \{p(x)\, x - x\} \\ &= \max_x \left\{\alpha A^{\alpha-1} x^{\alpha} - x\right\} \end{aligned}$$

the first-order condition for this problem is:

$$\alpha^2 A^{1-\alpha} x^{\alpha-1} - 1 = 0$$

which solves for the equilibrium quantity:

$$x = A\alpha^{\frac{2}{1-\alpha}}. \tag{6}$$

Substituting for x in the price equation (5), we obtain the equilibrium price

$$p = \frac{1}{\alpha}.$$

This price equation shows that the parameter α is a measure of the degree of product-market competition. That is, since an incumbent's marginal cost is one, therefore the inverse $1/\alpha$ is the equilibrium markup of price over marginal cost, which is also known as the Lerner index of monopoly power. Using these last two equations to substitute for x and p in the expression for π yields:

$$\pi = (p-1)x = A\left(\frac{1-\alpha}{\alpha}\right)\alpha^{\frac{2}{1-\alpha}} = A\delta \tag{7}$$

The equilibrium level of R&D is determined by the condition that the marginal expected gain from one more unit of R&D expenditure should equal the marginal cost, which by definition is just unity. Note the total expected gain to R&D is just the probability of an innovation times the profit earned by a successful innovator:

$$\text{total expected gain } = \lambda\left(R/A\right)^{\sigma}\pi$$

So the marginal expected gain is the derivative of this with respect to R. Setting this derivative equal to unity produces the "research arbitrage" equation:

$$\sigma\lambda\left(1/A\right)\left(R/A\right)^{\sigma-1}\pi = \sigma\lambda n^{\sigma-1}\delta = 1,$$

which can be solved for the equilibrium level of productivity-adjusted R&D expenditure:

$$n = (\sigma\lambda\delta)^{\frac{1}{1-\sigma}} \tag{8}$$

According to the production function (3), what happens to output Y over time depends not just on the growth of productivity A but also on the quantity of intermediate input x. Substituting (6) into (3) yields:

$$Y_t = \alpha^{\frac{2}{1-\alpha}}A_t, \tag{9}$$

which is strictly proportional to productivity. So the long-run average growth rate of output will be the same as the long-run average growth rate of A_t, namely the expected growth rate g. Since the population is constant g is also the long-run average growth rate of per capita output.

3.2 A variant with non-drastic innovations

So far we implicitly assumed that the incumbent innovator in any intermediate good sector, could charge any price to the final good sector without fearing entry by a potential competitor in that sector. In theoretical IO (see Tirole (1988)) we refer to this case as the "drastic innovation" case.

Now, suppose that for each intermediate good sector there is a competitive fringe of firms that can produce a "knockoff" product which is perfectly substitutable for the incumbent monopolist's intermediate product but costs $\chi > 1$ units of final output to produce. Then the incumbent monopolist cannot charge more than the limit price

$$p(x) = \chi \tag{10}$$

in equilibrium, since otherwise the competitive fringe could profitably undercut his price. When $\chi > \frac{1}{\alpha}$ this limit price constraint is not binding since we saw in the previous subsection that the equilibrium monopoly price in the absence of a competitive fringe, was equal to $\frac{1}{\alpha}$. This case corresponds again to the drastic innovation case.[12]

However, when $\chi < \frac{1}{\alpha}$, the limit price constraint is binding. This we refer to as the "non-drastic innovation" case. Then, putting equation (10) together with equation (5), immediately yields

$$x = A(\frac{\chi}{\alpha})^{\frac{1}{\alpha-1}}.$$

Equilibrium profits are then equal to

$$\pi = p.x - x = (\chi - 1)A(\frac{\chi}{\alpha})^{\frac{1}{\alpha-1}} = \delta A,$$

where

$$\delta = (\chi - 1)(\frac{\chi}{\alpha})^{\frac{1}{\alpha-1}}$$

is an increasing function of χ.

Thus the equilibrium R&D investment

$$n = (\sigma\lambda\delta)^{\frac{1}{1-\sigma}}$$

is also increasing in χ, and so is the equilibrium growth rate

$$g = \lambda n^{\sigma}(\gamma - 1).$$

[12] There is also potential competition from the previous intermediate product in that sector, which can be produced at a marginal cost of unity unity but has a productivity parameter of only A/γ. We show in the appendix to this chapter that even if the incumbent monopolist charges the unconstrained profit-maximizing price $1/\alpha$, producers of the previous product cannot profitably compete if innovations are big enough, spcifically if $\gamma > 1/\alpha^{\frac{\alpha}{1-\alpha}}$. (For example, if $\alpha = 1/2$ then previous technologies can be safely ignored if $\gamma > 2$.) From here on we assume that this condition holds.

A higher χ may reflect stronger patent protection which increases the cost of imitating the current technology in any intermediate good sector. Or it may reflect a lower degree of competition in the intermediate good sectors. In either case, it should lead to more intense R&D as it raises the expected rents that accrue to a successful innovator. This in turn should result in higher growth.

3.3 Comparative statics

To determine the long-run average growth rate in terms of primitive variables of the system, just use equation (7) to substitute for δ in the equation (8) determining productivity-adjusted R&D, and then use this expression to substitute for n in the growth equation (4). Doing this yields the following comparative statics properties on this equilibrium growth rate:

1. Growth increases with the productivity of innovations λ. This result points to the importance of education, and particularly higher education, as a growth-enhancing device. Indeed countries that invest more in higher education will achieve a higher productivity of research activities and also reduce the (opportunity) cost of R&D by increasing the aggregate supply of skilled labor. An increase in the size of population should also bring about an increase in growth by raising L. This "scale effect" has been challenged in the literature and will be discussed in the next chapter.

2. Growth increases with the size of innovations, as measured by γ. This in turn points the existence of a wedge between private and social innovation incentives. If given the choice between increasing the frequency λ or increasing their size γ, the private individual will go for increasing frequency. This is because size increases the cost of innovation as well as the expected rents; the research arbitrage equation shows that these two effects cancel each other, leaving the equilibrium level of R&D independent of size.

3. Growth is increasing in the degree of property right protection as measured by higher value of χ.

4. Growth is decreasing with the degree of product market competition measured by a higher α or a lower χ. This prediction is at odds with most recent empirical studies, starting with the work of Nickell (1996) and Blundell et al (1999). In Chapter X we shall revisit the relationship between competition and growth.

4 The continuous time model with R&D labor

In this section we extending the previous model to the case of continuous time. This in turn will allow us to capture the creative destruction effect of innovations, that is, the negative externality that future innovations exert on current innovations.

4.1 The model

Time is now continuous, and individuals have linear intertemporal preferences: $u(y) = \int_0^\infty y_\tau e^{-r\tau} d\tau$, where r is the rate of time preference, also equal to the interest rate. There are L individuals in the economy, and each individual is endowed with one unit flow of labor, so L is also equal to the aggregate flow of labor supply.

Output of the consumption good is produced at any time using an intermediate good according to:

$$y = Ax^\alpha, \tag{11}$$

where $0 < \alpha < 1$ and x denotes the flow of intermediate good used in final good production.

Innovations consist of the invention of a new variety of intermediate good that replaces the old one, and whose use raises the technology parameter, A, by the constant factor, $\gamma > 1$.

The economy's fixed stock of labor L has two competing uses. It can produce the intermediate good one for one, and it can be used in research. Assuming labor market clearing at any period, we thus have:

$$L = x + n, \tag{L}$$

where x is the amount of labor used in manufacturing and n is the amount of labor used in research.

When the amount n is used in research, innovations arrive randomly with a Poisson arrival rate λn, where $\lambda > 0$ is a parameter indicating the productivity of the research technology. The firm that succeeds in innovating can monopolize the intermediate sector until replaced by the next innovator.

Box 2.1 here

As in the previous toy model and the models in the previous chapter, there are positive spillovers from current innovators to future innovators, in the sense that the current invention makes it possible for other researchers to begin working on the next innovation. However, now there is also a negative spillover in the form of a "business-stealing effect," whereby the successful monopolist destroys the surplus attributable to the previous generation of intermediate good by making it obsolete.

The amount of labor devoted to research is determined by the arbitrage condition

$$w_t = \lambda V_{t+1}, \tag{A}$$

where t here is not time but the number of innovations that have occurred so far, w_t the wage, and V_{t+1} the discounted expected payoff to the $(t+1)^{th}$ innovation. The left-hand side is the value of an hour in manufacturing, whereas the right-hand side is the expected value of an hour in research—the flow probability λ of an innovation times the value V_{t+1}.

This arbitrage equation governs the dynamics of the economy over its successive innovations. Together with the labor market equation (L), it constitutes the backbone of the model.

The value V_{t+1} is determined by the following asset equation:

$$rV_{t+1} = \pi_{t+1} - \lambda n_{t+1} V_{t+1},$$

which says that the expected income generated by a license on the $(t+1)^{th}$ innovation during a unit time interval, namely rV_{t+1}, is equal to the profit flow π_{t+1} attainable by the $(t+1)$ intermediate good monopolist minus the expected "capital loss" that will occur when the $(t+1)^{th}$ innovator is replaced by a new innovator and therefore loses V_{t+1}. The flow probability of this loss is the arrival rate λn_{t+1}. Put in slightly different terms, the value V_{t+1} of the $(t+1)^{th}$ innovation is the net present value of an asset that yields π_{t+1} until it disappears, which it does at the expected rate λn_{t+1}.[13]

Note that this equation presupposes that the incumbent innovator does not perform R&D, so that λn_{t+1} is indeed the probability of that innovator losing his or her monopoly rents. In fact, there is a simple reason why the incumbent innovator chooses to do no research: because all the other researchers have immediate access to the incumbent technology A_t as a benchmark for their own research, the value to the incumbent innovator of making the next innovation is $V_{t+1} - V_t$, which is strictly less than the value V_{t+1} to an outside researcher. This is an example of the "Arrow effect," or "replacement effect."

We thus have

$$V_{t+1} = \pi_{t+1}/(r + \lambda n_{t+1}). \tag{12}$$

The denominator of (12), which can be interpreted as the obsolescence-adjusted interest rate, shows the effects of creative destruction. The more research is expected to occur following the

[13] n_{t+1} is the amount of labor devoted to R&D after the $(t+1)^{th}$ innovation.

next innovation, the shorter the likely duration of the monopoly profits that will be enjoyed by the creator of the next innovation, and hence the smaller the payoff to innovating.

The model is now almost entirely specified, except for the profit flow π_t and also the flow demand for manufacturing labor x_t. Both are determined by the same profit-maximization problem solved by the intermediate producer that uses the t^{th} innovation. This producer could either be thought as being the innovator himself (who then sets up his own intermediate firm), or as an existing intermediate firm that purchases (at price V_t) the patent for the innovation from the t^{th} innovator. In either case, the t^{th} innovator is able to extract the whole expected net present value (NPV) of monopoly profits generated by that innovation during the lifetime of this innovation, namely V_t.

4.2 Steady-state equilibrium R&D

A steady-state (or balanced growth) equilibrium is simply defined as a stationary solution to the system of equations defined by (A) and (L), with $x_t \equiv x$ and $n_t \equiv n$. In other words, both the allocation of labor between research and manufacturing and the productivity-adjusted wage rate remain constant over time, so that wages, profit, and final output are all scaled up by the same $\gamma > 1$ each time a new innovation occurs.

To make use of the research arbitrage equation (A), we need to calculate the equilibrium profits π_t. Here, we proceed a long similar lines as in the previous section. First, given that the final good sector is competitive, the price of an intermediate good of quality A_t is equal to its marginal productivity in producing final output, namely:

$$p_t(x) = \frac{\partial y}{\partial x} = \alpha A_t x^{\alpha-1}.$$

This is the inverse demand curve faced by the intermediate good monopolist. Now, given that the intermediate good is produced one for one with labor, the intermediate good producer will choose x to solve:

$$\pi_t = \max_x \{p_t(x)x - w_t x\},$$

where w_t is the wage rate in the economy when the current technology is A_t.First-order conditions yield:

$$\alpha^2 A_t x^{\alpha-1} = w_t,$$

which in turn imply that the equilibrium profit π_t can be reexpressed as:

$$\pi_t = \frac{1-\alpha}{\alpha} w_t x,$$

in other words as being equal to the wage bill times $(\frac{1-\alpha}{\alpha})$.

The research equation (A) can thus be reexpressed as:

$$w_t = \lambda \frac{\frac{1-\alpha}{\alpha} w_{t+1} x}{r + \lambda n}.$$

Now, using the fact that output, profits and wages are all multiplied by γ each time a new innovation occurs, we have:

$$w_{t+1} = \gamma w_t.$$

Thus the research arbitrage equation can be written:

$$w_t = \lambda \frac{\frac{1-\alpha}{\alpha} \gamma w_t x}{r + \lambda n},$$

or, after dividing through by w_t :

$$1 = \lambda \frac{\frac{1-\alpha}{\alpha} \gamma x}{r + \lambda n}.$$

Now, combining the above research equation with the labor market clearing equation (L), we simply obtain:

$$1 = \lambda \frac{\gamma \frac{1-\alpha}{\alpha} (L - n)}{r + \lambda n}. \tag{13}$$

The λn term in the denominator captures the creative destruction effect, whereby more anticipated research expected after the next innovation, discourages current research.

This equation determines the steady-state equilibrium R&D, $\widehat{n}$, as a function of the parameters of the economy. In particular we see that n : (i) increases with λ or γ; in other words, a more productive R&D technology encourages R&D; (ii) decreases with α; that is, a more competitive market for the intermediate good discourages R&D by lowering the rents to a successful innovator; (iii) decreases with r; that is, a higher rate of time preference reduces the net present value of the rents accruing to a successful innovator which in turn discourages R&D.

4.3 Steady-state rate of growth

In a steady state the flow of consumption good (or final output) produced during the time-interval between the t^{th} and the $(t+1)^{th}$ innovation is

$$y_t = A_t \widehat{x}^{\alpha} = A_t (L - \widehat{n})^{\alpha},$$

where $\widehat{n}$ is determined by (13), which in turn implies that

$$y_{t+1} = \gamma y_t. \tag{14}$$

Now, the reader should remember that the variable "t" does not refer to real time, but rather to the sequence of innovations $t = 1, 2, 3$, and so on. What happens to the evolution of final output in real time, that is, as a function of τ?

From equation (14) we know that the log of final output $\ln y(\tau)$ increases by an amount equal to $\ln \gamma$ each time a new innovation occurs. However, the real time interval between two successive innovations is random.

figure 2.2 here

Therefore, the time path of the log of final output $\ln y(\tau)$ will itself be a random step function, with the size of each step being equal to $\ln \gamma > 0$ and with the time interval between each step being exponentially distributed with parameter $\lambda \widehat{n}$.

Taking a unit-time interval between τ and $\tau+1$, we have: $\ln y(\tau+1) = \ln y(\tau) + (\ln \gamma)\varepsilon(\tau)$, where $\varepsilon(\tau)$ is the number of innovations between τ and $\tau + 1$. Given that $\varepsilon(\tau)$ is distributed Poisson with parameter $\lambda \widehat{n}$, we have: $E(\ln y(\tau + 1) - \ln y(\tau)) = \lambda \widehat{n} \ln \gamma$, where the LHS is nothing but the average growth rate.

We thus end up with a very simple expression for the average growth rate in a steady state:

$$g = \lambda \widehat{n} \ln \gamma. \tag{G}$$

Combining this equation with the previous comparative-statics analysis on the steady-state level of research $\widehat{n}$, we are now able to sign the impact of parameter changes on the average growth rate. Increases in the size of the labor market L or a reduction of the interest rate r and in the degree of market competition α will increase $\widehat{n}$ and thereby also g. Increases in the size of innovation γ and/or in the productivity of R&D λ will also foster growth, directly (by increasing the factor $\lambda \ln \gamma$) and also indirectly through increasing $\widehat{n}$.

5 Conclusion

In this Chapter we have introduced the Schumpeterian paradigm where growth results from quality improving innovations. We started by providing a road map in which we laid out the main assumptions of the paradigm, sketched its main steps and equations, and gave a preview of some key ideas that distinguish the paradigm from alternative growth models. Then we developed in detail two variant of the model, one with discrete time and the final good as R&D input, the second with continuous time and labor as R&D input. In most subsequent chapters, we will use adaptations of the discrete time model, and it is only in the chapter on general purpose technologies that we come back to the continuous time model.

It may be useful to contrast again the Schumpeterian growth paradigm to the two alternative models of endogenous growth analyzed previously. The first version of endogenous growth theory, analyzed in Chapter 2, was the AK model whereby knowledge accumulation is a serendipitous by-product of capital accumulation by the various firms in the economy. Here thrift and the resulting capital accumulation were the keys to growth, not novelty and innovation. The second version of endogenous growth theory, analyzed in Chapter 3, was the product variety model. Innovation causes productivity growth in the product-variety paradigm by creating new, but not necessarily improved, varieties of products.

Compared to the AK model, both the Schumpeterian model and the product variety model have the advantage of presenting an explicit analysis of the innovation process underlying long-run growth. Compared to the product variety model, the Schumpeterian model has the advantage of allowing for entrepreneurs to make the choice between implementation and frontier innovation, and for this choice to vary with distance to the frontier; this allows the Schumpeterian model to generate context-specific policy implications and comparative-statics predictions, dependent particularly on a county's distance to the frontier. In addition, nothing in the product-variety model implies an important role for exit and turnover of firms and workers which is at the heart of the Schumpeterian approach and which, as we argued at the end of the previous chapter, is consistent with an increasing number of recent studies demonstrating that labor and product market mobility are key elements of a growth-enhancing policy near the technological frontier.

This is not to say that the Schumpeterian model is free of issues. Indeed much of the discussion in the chapters to follow will be concerned with addressing the issues raised by the model of this chapter. For example, there is considerable evidence that seems to refute the scale effect of increased population on growth predicted by this model (and by the product-variety model). We will show how this difficulty can be resolved in chapter 5. Another problem with the theory is the lack of financial market imperfections; in reality, R&D firms rely very much on capital markets, which seem to work much better in some countries than in other. The issue of financial constraints will be dealt with in Chapter 6 below. Another possible difficulty with the model is the absence of capital, which growth-accounting exercises (Jorgenson, 1995; Young, 1995) have shown to be quite important. Chapter 7 will show how to introduce capital into the analysis in such a way as to make it consistent with these exercises. Another possible issue is convergence, which we examine in Chapter 8 below, where we will show that the model implies a form of club-converge consistent with the evidence of Durlauf and Johnson (1995) and Quah (1993, 1997); the key is that in Schumpeterian theory convergence occurs through productivity, via a process of technology transfer, as well as through capital accumulation. Finally, as indicated above the implication of the first-generation of Schumpeterian models to the effect that more product-market competition is harmful to growth runs counter to much evidence. We address this issue in Chapter 9 below, where we show that various other effects can be found in more sophisticated versions of the theory, which imply a more complicated relationship between competition and growth, one that finds considerable support in the data.

Box 2.1 Poisson Processes

Throughout these chapters we will often assume that some random event X is governed by a "Poisson process," with a certain "arrival rate" μ. What this means mathematically is that the time T you will have to wait for X to occur is a random variable whose distribution is exponential with parameter μ:

$$F(T) \equiv \text{Prob}\{\text{Event occurs before } T\} = 1 - e^{-\mu T}.$$

So the probability density of T is

$$f(T) = F'(T) = \mu e^{-\mu T}.$$

That is, the probability that the event will occur sometime within the short interval between T and $T+dt$ is approximately $\mu e^{-\mu T}dt$. In particular, the probability that it will occur within dt from *now* (when $T=0$) is approximately μdt. In this sense μ is the probability per unit of time that the event will occur now, or the "flow probability" of the event.

For example, in the present chapter the event that an individual researcher discovers innovation number $t+1$ is governed by a Poisson process with the arrival rate λ. The expression λV_{t+1} on the right-hand side of the arbitrage equation (A) represents the expected income of an individual researcher, because over a short interval of length dt the researcher will make an innovation worth V_{t+1} with probability λdt.

If X_1 and X_2 are two distinct events governed by independent Poisson processes with respective arrival rates μ_1 and μ_2, then the flow probability that at least one of the events will occur is just the sum of the two independent flow probabilities $\mu_1 + \mu_2$, because the probability that both events will occur at once is negligible. In this sense, independent Poisson processes are "additive." This is why, in the present chapter, when n_t independent researchers each innovate with a Poisson arrival rate λ, the Poisson arrival rate of innovations to the economy as a whole is the sum λn_t of the individual arrival rates.

If a sequence of independent events takes place, each governed by the same independent process with the constant arrival rate μ, then the expected number of arrivals per unit of time is obviously the arrival rate μ. For example, in the present chapter the expected number of innovations per year in a balanced growth equilibrium is the arrival rate λn.

Moreover, the number of events x that will take place over any interval of length Δ

is distributed according to the "Poisson distribution" that you will find described in most statistics textbooks:

$$g\left(x\right)=\text{prob}\left\{x\text{ events occur}\right\}=\frac{\left(\mu\Delta\right)^{x}e^{\mu\Delta}}{x!},$$

whose expected value is the arrival rate times the length of the interval $\mu\Delta$. This distribution is used in section 2.3 to express the expected present value of future output in a balanced growth equilibrium.

Appendix

Here we show that if $\gamma > 1/\alpha^{\frac{\alpha}{1-\alpha}}$ then final good producers would prefer to buy the latest intermediate product at the unconstrained monopoly price $1/\alpha$ instead of buying the previous generation at a price equal to the per unit cost of unity. Because of this, the previous generation of products will never be produced in equilibrium and the incumbent monopolist need not take into account potential competition from the previous technology.

To see this, note that a final-good producer buying the most recent product at a price of $1/\alpha$ will earn a profit (gross of labor cost) equal to:

$$\Pi = \max_{x} A^{1-\alpha} x^{\alpha} - (1/\alpha)\, x$$

The first-order condition is:

$$\alpha A^{1-\alpha} x^{\alpha-1} - (1/\alpha) = 0$$

Solving this equation for

$$x = \alpha^{\frac{2}{1-\alpha}} A$$

and substituting back into the above expression for Π yields the profit:

$$\Pi = (1-\alpha)\, \alpha^{\frac{2\alpha}{1-\alpha}} A$$

If the same producer bought last period's product at a price of unity, his profit would be:

$$\Pi' = \max_{x'} \left(A/\gamma\right)^{1-\alpha} \left(x'\right)^{\alpha} - x'$$

The first-order condition is:

$$\alpha \left(A/\gamma\right)^{1-\alpha} \left(x'\right)^{\alpha-1} - 1 = 0$$

Solving this equation for

$$x' = \alpha^{\frac{1}{1-\alpha}} A/\gamma$$

and substituting back into the above expression for Π yields the profit:

$$\Pi' = (1-\alpha)\, \alpha^{\frac{\alpha}{1-\alpha}} A/\gamma$$

By inspection

$$\Pi \gtreqless \Pi' \text{ according to whether } \gamma \gtreqless \alpha^{\frac{-\alpha}{1-\alpha}}.$$