



What Do We Learn From Schumpeterian Growth Theory?

Philippe Aghion^{*}, Ufuk Akcigit[†], and Peter Howitt[‡]

^{*}Harvard University, NBER, and CIFAR, USA

[†]University of Pennsylvania and NBER, USA

[‡]Brown University and NBER, USA

Abstract

Schumpeterian growth theory has operationalized Schumpeter's notion of creative destruction by developing models based on this concept. These models shed light on several aspects of the growth process that could not be properly addressed by alternative theories. In this survey, we focus on four important aspects, namely: (i) the role of competition and market structure; (ii) firm dynamics; (iii) the relationship between growth and development with the notion of appropriate growth institutions; and (iv) the emergence and impact of long-term technological waves. In each case, Schumpeterian growth theory delivers predictions that distinguish it from other growth models and which can be tested using micro data.

Keywords

Creative destruction, Entry, Exit, Competition, Firm dynamics, Reallocation, R&D, Industrial policy, Technological frontier, Schumpeterian wave, General-purpose technology

JEL Classification Codes

O10, O11, O12, O30, O31, O33, O40, O43, O47



1.1. INTRODUCTION

Formal models allow us to make verbal notions operational and confront them with data. The Schumpeterian growth theory surveyed in this paper has “operationalized” Schumpeter's notion of creative destruction—the process by which new innovations replace older technologies—in two ways. First, it has developed models based on creative destruction that shed new light on several aspects of the growth process. Second, it has used data, including rich micro data, to confront the predictions that distinguish it from other growth theories. In the process, the theory has improved our understanding of the underlying sources of growth.

Over the past 25 years,¹ Schumpeterian growth theory has developed into an integrated framework for understanding not only the macroeconomic structure of growth but also the many microeconomic issues regarding incentives, policies, and organizations that interact with growth: who gains and who loses from innovations, and what the net rents from innovation are. These ultimately depend on characteristics such as property right protection; competition and openness; education; democracy; and so forth, and to a different extent in countries or sectors at different stages of development. Moreover, the recent years have witnessed a new generation of Schumpeterian growth models focusing on firm dynamics and reallocation of resources among incumbents and new entrants.² These models are easily estimable using micro firm-level datasets, which also bring the rich set of tools from other empirical fields into macroeconomics and endogenous growth.

In this paper, which aims to be accessible to readers with only basic knowledge in economics and is thus largely self-contained, we shall consider four aspects on which Schumpeterian growth theory delivers distinctive predictions.³ First, the relationship between growth and industrial organization: faster innovation-led growth is generally associated with higher turnover rates, i.e. higher rates of creation and destruction, of firms and jobs; moreover, competition appears to be positively correlated with growth, and competition policy tends to complement patent policy. Second, the relationship between growth and firm dynamics: small firms exit more frequently than large firms; conditional on survival, small firms grow faster; there is a very strong correlation between firm size and firm age; and finally, firm size distribution is highly skewed. Third, the relationship between growth and development with the notion of appropriate institutions: namely, the idea that different types of policies or institutions appear to be growth-enhancing at different stages of development. Our emphasis will be on the relationship between growth and democracy and on why this relationship appears to be stronger in more frontier economies. Four, the relationship between growth and long-term technological waves: why such waves are associated with an increase in the flow of firm entry and exit; why they may initially generate a productivity slowdown; and why they may increase wage inequality both between and within educational groups. In each case, we show that

¹ The approach was initiated in the fall of 1987 at MIT, where Philippe Aghion was a 1-year assistant professor and Peter Howitt a visiting professor on sabbatical from the University of Western Ontario. During that year they wrote their “model of growth through creative destruction” (see [Section 1.2](#) below), which was published as [Aghion and Howitt \(1992\)](#). Parallel attempts at developing Schumpeterian growth models include [Segerstrom et al \(1990\)](#) and [Corriveau \(1991\)](#).

² See [Klette and Kortum \(2004\)](#), [Lentz and Mortensen \(2008\)](#), [Akcigit and Kerr \(2010\)](#), and [Acemoglu et al. \(2013\)](#).

³ Thus, we are not looking at the aspects or issues that could be addressed by the Schumpeterian model and by other models, including [Romer’s \(1990\)](#) product variety model (see [Aghion and Howitt, 1998, 2009](#)). [Grossman and Helpman \(1991\)](#) were the first to point out the parallels between the two models, although using a special version of the Schumpeterian model.

Schumpeterian growth theory delivers predictions that distinguish it from other growth models and which can be tested using micro data.

The paper is organized as follows: [Section 1.2](#) lays out the basic Schumpeterian model; [Section 1.3](#) introduces competition and IO into the framework; [Section 1.4](#) analyzes firm dynamics; [Section 1.5](#) looks at the relationship between growth and development and in particular at the role of democracy in the growth process; [Section 1.6](#) discusses technological waves; and [Section 1.7](#) concludes.

A word of caution before we proceed: this paper focuses on the Schumpeterian growth paradigm and some of its applications. It is not a survey of the existing (endogenous) growth literature. There, we refer the reader to growth textbooks (e.g. [Acemoglu, 2009](#); [Aghion and Howitt, 1998, 2009](#); [Barro and Sala-i-Martin, 2003](#); [Galor, 2011](#); [Jones and Vollrath, 2013](#); [Weil, 2012](#)).



1.2. SCHUMPETERIAN GROWTH: BASIC MODEL

1.2.1 The Setup

The following model borrows directly from the theoretical IO and patent race literature (see [Tirole, 1988](#)). This model is Schumpeterian in that: (i) it is about growth generated by innovations; (ii) innovations result from entrepreneurial investments that are themselves motivated by the prospects of monopoly rents; and (iii) new innovations replace old technologies: in other words, growth involves creative destruction.

Time is continuous and the economy is populated by a continuous mass L of infinitely lived individuals with linear preferences, that discount the future at rate ρ .⁴ Each individual is endowed with one unit of labor per unit of time, which he or she can allocate between production and research: in equilibrium, individuals are indifferent between these two activities.

There is a final good, which is also the numeraire. The final good at time t is produced competitively using an intermediate input, namely:

$$Y_t = A_t \gamma_t^\alpha,$$

where α is between zero and one, γ_t is the amount of the intermediate good currently used in the production of the final good, and A_t is the productivity—or quality—of the currently used intermediate input.⁵

The intermediate good γ is in turn produced one for one with labor: that is, one unit flow of labor currently used in manufacturing the intermediate input produces one unit of intermediate input of frontier quality. Thus, γ_t denotes both the current production of the

⁴ The linear preferences (or risk neutrality) assumption implies that the equilibrium interest rate will always be equal to the rate of time preference: $r_t = \rho$ (see [Aghion and Howitt, 2009, Chapter 2](#)).

⁵ In what follows we will use the words “productivity” and “quality” interchangeably.

intermediate input and the flow amount of labor currently employed in manufacturing the intermediate good.

Growth in this model results from innovations that improve the quality of the intermediate input used in the production of the final good. More formally, if the previous state-of-the-art intermediate good was of quality A , then a new innovation will introduce a new intermediate input of quality γA , where $\gamma > 1$. This immediately implies that growth will involve creative destruction, in the sense that Bertrand competition will allow the new innovator to drive the firm producing the intermediate good of quality A out of the market, since at the same labor cost the innovator produces a better good than that of the incumbent firm.⁶

The innovation technology is directly drawn from the theoretical IO and patent race literatures: namely, if z_t units of labor are currently used in R&D, then a new innovation arrives during the current unit of time at the (memoryless) Poisson rate λz_t .⁷ Henceforth, we will drop the time index t , when it causes no confusion.

1.2.2 Solving the Model

1.2.2.1 The Research Arbitrage and Labor Market Clearing Equations

We shall concentrate our attention on balanced growth equilibria where the allocation of labor between production (y) and R&D (z) remains constant over time. The growth process is described by two basic equations.

The first is the labor market clearing equation:

$$L = y + z, \quad (\text{L})$$

reflecting the fact that the total flow of labor supply during any unit of time is fully absorbed between production and R&D activities (i.e. by the demand for manufacturing and R&D labor).

⁶ Thus, overall, growth in the Schumpeterian model involves both positive and negative externalities. The positive externality is referred to by [Aghion and Howitt \(1992\)](#) as a “knowledge spillover effect.” Namely, any new innovation raises productivity A forever, i.e. the benchmark technology for any subsequent innovation. However, the current (private) innovator captures the rents from his or her innovation only during the time interval until the next innovation occurs. This effect is also featured in [Romer \(1990\)](#), where it is referred to instead as “non-rivalry plus limited excludability.” But in addition, in the Schumpeterian model, any new innovation has a negative externality as it destroys the rents of the previous innovator. Following the theoretical IO literature, [Aghion and Howitt \(1992\)](#) refer to this as the “business-stealing effect” of innovation. The welfare analysis in that paper derives sufficient conditions under which the intertemporal spillover effect dominates or is dominated by the business-stealing effect. The equilibrium growth rate under laissez-faire is correspondingly suboptimal or excessive compared to the socially optimal growth rate.

⁷ More generally, if z_t units of labor are invested in R&D during the time interval $[t, t + dt]$, the probability of innovation during this time interval is $\lambda z_t dt$.

The second equation reflects individuals' indifference in equilibrium between engaging in R&D or working in the intermediate good sector. We call it the research-arbitrage equation. The remaining part of the analysis consists of spelling out this research-arbitrage equation.

More formally, let w_k denote the current wage rate conditional on there having already been $k \in \mathbb{Z}_{++}$ innovations from time 0 until current time t (since innovation is the only source of change in this model, all other economic variables remain constant during the time interval between two successive innovations). And let V_{k+1} denote the net present value of becoming the next $((k+1)\text{th})$ innovator.

During a small time interval dt , between the k th and $(k+1)$ th innovations, an individual faces the following choices: Either she employs her (flow) unit of labor for the current unit of time in manufacturing at the current wage, in which case she gets $w_t dt$. Or she devotes her flow unit of labor to R&D, in which case she will innovate during the current time period with probability λdt and then get V_{k+1} , whereas she gets nothing if she does not innovate.⁸ The research-arbitrage equation is then simply expressed as:

$$w_k = \lambda V_{k+1}. \quad (\text{R})$$

The value V_{k+1} is in turn determined by a Bellman equation. We will use Bellman equations repeatedly in this survey; thus, it is worth going slowly here. During a small time interval dt , a firm collects $\pi_{k+1} dt$ profits. At the end of this interval, it is replaced by a new entrant with probability $\lambda z dt$ through creative destruction; otherwise, it preserves the monopoly power and V_{k+1} . Hence, the value function is written as:

$$V_{k+1} = \pi_{k+1} dt + (1 - \lambda z dt) \left[\frac{\lambda z dt \times 0 + (1 - \lambda z dt) \times V_{k+1}}{(1 - \lambda z dt) \times V_{k+1}} \right].$$

Dividing both sides by dt , then taking the limit as $dt \rightarrow 0$ and using the fact that the equilibrium interest rate is equal to the time preference, the Bellman equation for V_{k+1} can be rewritten as:

$$\rho V_{k+1} = \pi_{k+1} - \lambda z V_{k+1}.$$

In other words, the annuity value of a new innovation (i.e. its flow value during a unit of time) is equal to the current profit flow π_{k+1} minus the expected capital loss $\lambda z V_{k+1}$ due to creative destruction, i.e. to the possible replacement by a subsequent innovator. If innovating gave the innovator access to a permanent profit flow π_{k+1} , then we know that

⁸ Note that we are implicitly assuming that previous innovators are not candidates for being new innovators. This in fact results from a replacement effect pointed out by [Arrow \(1962\)](#). Namely, an outsider goes from zero to V_{k+1} if he or she innovates, whereas the previous innovator would go from V_k to V_{k+1} . Given that the R&D technology is linear, if outsiders are indifferent between innovating and working in manufacturing, then incumbent innovators will strictly prefer to work in manufacturing. Thus, new innovations end up being made by outsiders in equilibrium in this model. This feature will be relaxed in the next section.

the value of the corresponding perpetuity would be π_{k+1}/r .⁹ However, there is creative destruction at rate λz . As a result, we have:

$$V_{k+1} = \frac{\pi_{k+1}}{\rho + \lambda z}, \quad (1.1)$$

that is, the value of innovation is equal to the profit flow divided by the risk-adjusted interest rate $\rho + \lambda z$ where the risk is that of being displaced by a new innovator.

1.2.2.2 Equilibrium Profits, Aggregate R&D, and Growth

We solve for equilibrium profits π_{k+1} and the equilibrium R&D rate z by backward induction. That is, first, for a given productivity of the current intermediate input, we solve for the equilibrium profit flow of the current innovator; then we move one step back and determine the equilibrium R&D using [Equations \(L\)](#) and [\(R\)](#).

Equilibrium profits Suppose that k_t innovations have already occurred until time t , so that the current productivity of the state-of-the-art intermediate input is $A_{k_t} = \gamma^{k_t}$. Given that the final good production is competitive, the intermediate good monopolist will sell his or her input at a price equal to its marginal product, namely:

$$p_k(\gamma) = \frac{\partial(A_k \gamma^\alpha)}{\partial \gamma} = A_k \alpha \gamma^{\alpha-1}. \quad (1.2)$$

This is the inverse demand curve faced by the intermediate good monopolist.

Given that inverse demand curve, the monopolist will choose γ to:

$$\pi_k = \max_{\gamma} \{p_k(\gamma)\gamma - w_k \gamma\}, \quad \text{subject to (1.2)} \quad (1.3)$$

since it costs $w_k \gamma$ units of the numeraire to produce γ units of the intermediate good. Given the Cobb-Douglas technology for the production of the final good, the equilibrium price is a constant markup over the marginal cost ($p_k = w_k/\alpha$) and the profit is simply equal to $\frac{1-\alpha}{\alpha}$ times the wage bill, namely:¹⁰

$$\pi_k = \frac{1-\alpha}{\alpha} w_k \gamma, \quad (1.4)$$

where γ solves [\(1.3\)](#).

Equilibrium aggregate R&D Combining [\(1.1\)](#), [\(1.4\)](#), and [\(R\)](#), we can rewrite the research-arbitrage equation as:

$$w_k = \lambda \frac{\frac{1-\alpha}{\alpha} w_{k+1} \gamma}{\rho + \lambda z}. \quad (1.5)$$

⁹ Indeed, the value of the perpetuity is:

$$\int_0^\infty \pi_{k+1} e^{-rt} dt = \frac{\pi_{k+1}}{r}.$$

¹⁰ To see that $p_k = w_k/\alpha$, simply combine the first-order condition of [\(1.3\)](#) with expression [\(1.2\)](#).

Using the labor market clearing condition (L) and the fact that on a balanced growth path all aggregate variables (the final output flow, profits, and wages) are multiplied by γ each time a new innovation occurs, we can solve (1.5) for the equilibrium aggregate R&D z as a function of the parameters of the economy:

$$z = \frac{\frac{1-\alpha}{\alpha} \gamma L - \frac{\rho}{\lambda}}{1 + \frac{1-\alpha}{\alpha} \gamma}. \quad (1.6)$$

Clearly, it is sufficient to assume that $\frac{1-\alpha}{\alpha} \gamma L > \frac{\rho}{\lambda}$ to ensure positive R&D in equilibrium. Inspection of (1.6) delivers a number of important comparative statics. In particular, a higher productivity of the R&D technology as measured by λ or a larger size of innovations γ or a larger size of the population L has a positive effect on aggregate R&D. On the other hand, a higher α (which corresponds to the intermediate producer facing a more elastic inverse demand curve and therefore getting lower monopoly rents) or a higher discount rate ρ tends to discourage R&D.

Equilibrium expected growth Once we have determined the equilibrium aggregate R&D, it is easy to compute the expected growth rate. First note that during a small time interval $[t, t + dt]$, there will be a successful innovation with probability $\lambda z dt$. Second, the final output is multiplied by γ each time a new innovation occurs. Therefore, the expected log-output is simply:

$$\mathbb{E}(\ln Y_{t+dt}) = \lambda z dt \ln \gamma Y_t + (1 - \lambda z dt) \ln Y_t.$$

Subtracting $\ln Y_t$ from both sides, dividing through dt , and finally taking the limit leads to the following expected growth:

$$\mathbb{E}(g_t) = \lim_{dt \rightarrow 0} \frac{\ln Y_{t+dt} - \ln Y_t}{dt} = \lambda z \ln \gamma,$$

which inherits the comparative static properties of z with respect to the parameters λ , γ , α , ρ , and L .

A distinct prediction of the model is:

Prediction 0: *The turnover rate λz is positively correlated with the growth rate g .*



1.3. GROWTH MEETS IO

Empirical studies (starting with [Nickell \(1996\)](#), [Blundell et al. \(1995, 1999\)](#)) point to a positive correlation between growth and product market competition. Also, the idea that competition—or free entry—should be growth-enhancing is also prevalent among policy advisers. Yet, non-Schumpeterian growth models cannot account for it: AK models assume perfect competition and therefore have nothing to say about the

relationship between competition and growth. And in Romer's product variety model, higher competition amounts to a higher degree of substitutability between the horizontally differentiated inputs, which in turn implies lower rents for innovators and therefore lower R&D incentives and thus lower growth.

In contrast, the Schumpeterian growth paradigm can rationalize the positive correlation between competition and growth found in linear regressions. In addition, it can account for several interesting facts about competition and growth that no other growth theory can explain.¹¹ We shall concentrate on three such facts. First, innovation and productivity growth by incumbent firms appear to be stimulated by competition and entry, particularly in firms near the technology frontier or in firms that compete neck-and-neck with their rivals, less so than in firms below the frontier. Second, competition and productivity growth display an inverted-U relationship: starting from an initially low level of competition, higher competition stimulates innovation and growth; starting from a high initial level of competition, higher competition has a less positive or even a negative effect on innovation and productivity growth. Third, patent protection complements product market competition in encouraging R&D investments and innovation.

Understanding the relationship between competition and growth also helps improve our understanding of the relationship between trade and growth. Indeed, there are several dimensions to that relationship. First is the scale effect, whereby liberalizing trade increases the market for successful innovations and therefore the incentives to innovate; this is naturally captured by any innovation-based model of growth, including the Schumpeterian growth model. But there is also a competition effect of trade openness, which only the Schumpeterian model can capture. This latter effect appears to have been at work in emerging countries that implemented trade liberalization reforms (for example, India in the early 1990s),¹² and it also explains why trade restrictions are more detrimental to growth in more frontier countries (see [Section 1.5](#) below).

1.3.1 From Leapfrogging to Step-By-Step Innovation¹³

1.3.1.1 The Argument

To reconcile theory with the evidence on productivity growth and product market competition, we replace the leapfrogging assumption of the model in the previous section (where incumbents are systematically overtaken by outside researchers) with a less radical *step-by-step* assumption. Namely, a firm that is currently m steps behind the technological leader in the same sector or industry must catch up with the leader before becoming a leader itself. This step-by-step assumption can be rationalized by supposing that an

¹¹ See [Aghion and Griffith \(2006\)](#) for a first attempt at synthesizing the theoretical and empirical debates on competition and growth.

¹² See, for instance, [De Loecker et al. \(2012\)](#), [Goldberg et al. \(2010\)](#), [Sivadasan \(2009\)](#), and [Topalova and Khandelwal \(2011\)](#).

¹³ The following model and analysis are based on [Aghion et al. \(1997\)](#), [Aghion et al. \(2001\)](#), [Aghion et al. \(2005\)](#), and [Acemoglu and Akgigit \(2012\)](#). See also [Peretto \(1998\)](#) for related work.

innovator acquires tacit knowledge that cannot be duplicated by a rival without engaging in its own R&D to catch up. This leads to a richer analysis of the interplay between product market competition, innovation, and growth by allowing firms in some sectors to be *neck-and-neck*. In such sectors, increased product market competition, by making life more difficult for neck-and-neck firms, will encourage them to innovate in order to acquire a lead over their rival in the sector. This we refer to as the escape–competition effect. On the other hand, in sectors that are not neck-and-neck, increased product market competition will have a more ambiguous effect on innovation. In particular, it will discourage innovation by laggard firms when these do not put much weight on the (more remote) prospect of becoming a leader and instead mainly look at the short run extra profit from catching up with the leader. This we call the Schumpeterian effect. Finally, the steady-state fraction of neck-and-neck sectors will itself depend upon the innovation intensities in neck-and-neck versus unleveled sectors. This we refer to as the composition effect.

1.3.1.2 Household

Time is again continuous and a continuous measure L of individuals work in one of two activities: as production workers and as R&D workers. We assume that the representative household consumes C_t , has logarithmic instantaneous utility $U(C_t) = \ln C_t$, and discounts the future at a rate $\rho > 0$. Moreover, the household holds a balanced portfolio of all the firms, \mathcal{A}_t . Hence, its budget constraint is simply $C_t + \dot{\mathcal{A}}_t = r_t \mathcal{A}_t + Lw_t$. These assumptions deliver the household's Euler equation as $g_t = r_t - \rho$. All costs in this economy are in terms of labor units. Therefore, the household's consumption is equal to the final good production $C_t = Y_t$, which is also the resource constraint of this economy.

1.3.1.3 A Multi-Sector Production Function

To formalize these various effects, in particular the composition effect, we obviously need a multiplicity of intermediate sectors instead of one, as in the previous section. One simple way to extend the Schumpeterian paradigm to a multiplicity of intermediate sectors is, as in [Grossman and Helpman \(1991\)](#), to assume that the final good is produced using a continuum of intermediate inputs, according to the logarithmic production function:

$$\ln Y_t = \int_0^1 \ln \gamma_{jt} dj. \quad (1.7)$$

Next, we introduce competition by assuming that each sector j is duopolistic with respect to production and research activities. We denote the two duopolists in sector j as A_j and B_j and assume, for simplicity, that γ_j is the sum of the intermediate goods produced by the two duopolists in sector j :

$$\gamma_j = \gamma_{A_j} + \gamma_{B_j}.$$

The above logarithmic technology implies that in equilibrium the same amount is spent at any time by final good producers on each basket y_j .¹⁴ We normalize the price of the final good to be 1. Thus, a final good producer chooses each γ_{Aj} and γ_{Bj} to maximize $\gamma_{Aj} + \gamma_{Bj}$ subject to the budget constraint: $p_{Aj}\gamma_{Aj} + p_{Bj}\gamma_{Bj} = Y$. That is, he or she will devote the entire unit expenditure to the least expensive of the two goods.

1.3.1.4 Technology and Innovation

Each firm takes the wage rate as given and produces using labor as the only input according to the following linear production function:

$$\gamma_{it} = A_{it}l_{it}, \quad i \in \{A, B\},$$

where l_{it} is the labor employed. Let k_i denote the technology level of duopoly firm i in some industry j ; that is, $A_i = \gamma^{k_i}$, $i = A, B$, and $\gamma > 1$ is a parameter that measures the size of a leading-edge innovation. Equivalently, it takes γ^{-k_i} units of labor for firm i to produce one unit of output. Thus, the unit cost of production is simply $c_i = w\gamma^{-k_i}$, which is independent of the quantity produced.

An industry j is thus fully characterized by a pair of integers (k_j, m_j) where k_j is the leader's technology and m_j is the technological gap between the leader and the follower.¹⁵

For expositional simplicity, we assume that knowledge spillovers between the two firms in any intermediate industry are such that neither firm can get more than one technological level ahead of the other, that is:

$$m \leq 1.$$

In other words, if a firm that is already one step ahead innovates, the lagging firm will automatically learn to copy the leader's previous technology and thereby remain only one step behind. Thus, at any point in time, there will be two kinds of intermediate sectors in the economy: (i) leveled or neck-and-neck sectors, where both firms are on a technological par with one another; and (ii) unleveled sectors, where one firm (the leader) lies one step ahead of its competitor (the laggard or follower) in the same industry.¹⁶

¹⁴ To see this, note that a final good producer will choose the y_j 's to maximize $u = \int \ln y_j dj$ subject to the budget constraint $\int p_j y_j dj = E$, where E denotes current expenditures. The first-order condition for this is:

$$\partial u / \partial y_j = 1/y_j = \lambda p_j \quad \text{for all } j,$$

where λ is a Lagrange multiplier. Together with the budget constraint this first-order condition implies:

$$p_j y_j = 1/\lambda = E \quad \text{for all } j.$$

¹⁵ The above logarithmic final good technology, together with the linear production cost structure for intermediate goods, implies that the equilibrium profit flows of the leader and the follower in an industry depend only on the technological gap, m , between the two firms. We will see this below for the case where $m \leq 1$.

¹⁶ Aghion et al. (2001) and Acemoglu and Acikgiz (2012) analyze the more general case where there is no limit to how far ahead the leader can get.

To complete the description of the model, we just need to specify the innovation technology. Here we simply assume that by spending the R&D cost $\psi(z) = z^2/2$ in units of labor, a leader (or frontier) firm moves one technological step ahead at the rate z . We call z the innovation rate or R&D intensity of the firm. We assume that a follower firm can move one step ahead with probability h , even if it spends nothing on R&D, by copying the leader's technology. Thus, $z^2/2$ is the R&D cost (in units of labor) of a follower firm moving ahead with probability $z + h$. Let z_0 denote the R&D intensity of each firm in a neck-and-neck industry, and let z_{-1} denote the R&D intensity of a follower firm in an unleveled industry; if z_1 denotes the R&D intensity of the leader in an unleveled industry, note that $z_1 = 0$, since our assumption of automatic catch-up means that a leader cannot gain any further advantage by innovating.

1.3.2 Equilibrium Profits and Competition in Leveled and Unleveled Sectors

We can now determine the equilibrium profits of firms in each type of sector and link them with product market competition. The final good producer in (1.7) generates a unit-elastic demand with respect to each variety:

$$y_j = \frac{Y}{p_j}. \quad (1.8)$$

Consider first an unleveled sector where the leader's unit cost is c . The leader's monopoly profit is:

$$\begin{aligned} p_1 y_1 - c y_1 &= \left(1 - \frac{c}{p_1}\right) Y \\ &= \pi_1 Y, \end{aligned}$$

where the first line uses (1.8) and the second line defines π_1 as the equilibrium profit normalized by the final output Y . Note that the monopoly profit is monotonically increasing in the unit price p_1 . However, the monopolist is constrained to setting a price $p_1 \leq \gamma c$, because γc is the rival's unit cost, so at any higher price the rival could profitably undercut his or her price and steal all their business. He or she will therefore choose the maximum possible price $p_1 = \gamma c$, such that the normalized profit in equilibrium is:

$$\pi_1 = 1 - \frac{1}{\gamma}.$$

The laggard in the unleveled sector will be priced out of the market and hence will earn a zero profit:

$$\pi_{-1} = 0.$$

Consider now a leveled (neck-and-neck) sector. If the two firms engaged in open price competition with no collusion, the equilibrium price would fall to the unit cost c

of each firm, resulting in zero profit. At the other extreme, if the two firms colluded so effectively as to maximize their joint profits and shared the proceeds, then they would together act like the leader in an unleveled sector, each setting $p = \gamma c$ (we assume that any third firm could compete using the previous best technology, just like the laggard in an unleveled sector), and each earning a normalized profit equal to $\pi_1/2$.

So in a leveled sector, both firms have an incentive to collude. Accordingly, we model the degree of product market competition inversely by the degree to which the two firms in a neck-and-neck industry are able to collude. (They do not collude when the industry is unleveled because the leader has no interest in sharing their profit.) Specifically, we assume that the normalized profit of a neck-and-neck firm is:

$$\pi_0 = (1 - \Delta) \pi_1, \quad 1/2 \leq \Delta \leq 1,$$

and we parameterize product market competition by Δ , that is, one minus the fraction of a leader's profits that the leveled firm can attain through collusion. Note that Δ is also the incremental profit of an innovator in a neck-and-neck industry, normalized by the leader's profit.

We next analyze how the equilibrium research intensities z_0 and z_{-1} of neck-and-neck and backward firms, respectively, and consequently the aggregate innovation rate, vary with our measure of competition Δ .

1.3.3 The Schumpeterian and Escape–Competition Effects

On a balanced growth path, all aggregate variables, including firm values, will grow at the rate g . For tractability, we will normalize all growing variables by the aggregate output Y . Let V_m (resp. V_{-m}) denote the normalized steady-state value of currently being a leader (resp. a follower) in an industry with technological gap m , and let $\omega = w/Y$ denote the normalized steady-state wage rate. We have the following Bellman equations:¹⁷

$$\rho V_0 = \max_{z_0} \left\{ \pi_0 + \bar{z}_0(V_{-1} - V_0) + z_0(V_1 - V_0) - \omega z_0^2/2 \right\}, \quad (1.9)$$

$$\rho V_{-1} = \max_{z_{-1}} \left\{ \pi_{-1} + (z_{-1} + h)(V_0 - V_{-1}) - \omega z_{-1}^2/2 \right\}, \quad (1.10)$$

$$\rho V_1 = \pi_1 + (z_{-1} + h)(V_0 - V_1), \quad (1.11)$$

where \bar{z}_0 denotes the R&D intensity of the other firm in a neck-and-neck industry (we focus on a symmetric equilibrium where $\bar{z}_0 = z_0$). Note that we already used $z_1 = 0$ in (1.11).

¹⁷ Note that originally the left-hand side is written as $rV_0 - \dot{V}_0$. Note that on a BGP, $\dot{V}_0 = gV_0$; therefore, we get $(r - g)V_0$. Finally, using the household's Euler equation, $r - g = \rho$, leads to the Bellman equations in the text.

In words, the growth-adjusted annuity value ρV_0 of currently being neck-and-neck is equal to the corresponding profit flow π_0 plus the expected capital gain $z_0(V_1 - V_0)$ of acquiring a lead over the rival plus the expected capital loss $\bar{z}_0(V_{-1} - V_0)$, if the rival innovates and thereby becomes the leader, minus the R&D cost $\omega z_0^2/2$. Similarly, the annuity value ρV_1 of being a technological leader in an unleveled industry is equal to the current profit flow π_1 plus the expected capital loss $z_{-1}(V_0 - V_1)$ if the leader is being caught up by the laggard (recall that a leader does not invest in R&D in equilibrium). Finally, the annuity value ρV_{-1} of currently being a laggard in an unleveled industry is equal to the corresponding profit flow π_{-1} plus the expected capital gain $(z_{-1} + h)(V_0 - V_{-1})$ of catching up with the leader, minus the R&D cost $\omega z_{-1}^2/2$.

Using the fact that z_0 maximizes (1.9) and z_{-1} maximizes (1.10), we have the first-order conditions:

$$\omega z_0 = V_1 - V_0, \quad (1.12)$$

$$\omega z_{-1} = V_0 - V_{-1}. \quad (1.13)$$

In Aghion et al. (1997) the model is closed by a labor market clearing equation that determines ω as a function of the aggregate demand for R&D plus the aggregate demand for manufacturing labor. Here, for simplicity we shall ignore that equation and take the wage rate ω as given, normalizing it at $\omega = 1$.

Then, using (1.12) and (1.13) to eliminate the V 's from the system of Equations (1.9)–(1.11), we end up with a system of two equations in the two unknowns z_0 and z_{-1} :

$$z_0^2/2 + (\rho + h)z_0 - (\pi_1 - \pi_0) = 0, \quad (1.14)$$

$$z_{-1}^2/2 + (\rho + z_0 + h)z_{-1} - (\pi_0 - \pi_{-1}) - z_0^2/2 = 0. \quad (1.15)$$

These equations solve recursively for unique positive values of z_0 and z_{-1} , and we are mainly interested in how equilibrium R&D intensities are affected by an increase in product market competition Δ . It is straightforward to see from Equation (1.14) and the fact that:

$$\pi_1 - \pi_0 = \Delta\pi_1,$$

that an increase in Δ will increase the innovation intensity $z_0(\Delta)$ of a neck-and-neck firm. This is the escape-competition effect.

Then, plugging $z_0(\Delta)$ into (1.15), we can look at the effect of an increase in competition Δ on the innovation intensity z_{-1} of a laggard. This effect is ambiguous in general: in particular, for very high ρ , the effect is negative, since then z_{-1} varies like:

$$\pi_0 - \pi_{-1} = (1 - \Delta)\pi_1.$$

In this case, the laggard is very impatient and thus looks at its short-term net profit flow if it catches up with the leader, which in turn decreases when competition increases. This

is the *Schumpeterian effect*. However, for low values of ρ , this effect is counteracted by an anticipated escape–competition effect.

Thus, the effect of competition on innovation depends on what situation a sector is in. In unleveled sectors, the Schumpeterian effect is at work even if it does not always dominate. But in leveled (neck-and-neck) sectors, the escape–competition effect is the only effect at work; that is, more competition induces neck-and-neck firms to innovate in order to escape from a situation in which competition constrains profits.

On average, an increase in product market competition will have an ambiguous effect on growth. It induces faster productivity growth in currently neck-an-neck sectors and slower growth in currently unleveled sectors. The overall effect on growth will thus depend on the (steady-state) fraction of leveled versus unleveled sectors. But this steady-state fraction is itself endogenous, since it depends on equilibrium R&D intensities in both types of sectors. We proceed to show under which condition this overall effect is an inverted U and, at the same time, derive additional predictions for further empirical testing.

1.3.3.1 Composition Effect and the Inverted U

In a steady state, the fraction of sectors μ_1 that are unleveled is constant, as is the fraction $\mu_0 = 1 - \mu_1$ of sectors that are leveled. The fraction of unleveled sectors that become leveled each period will be $z_{-1} + h$, so the sectors moving from unleveled to leveled represent the fraction $(z_{-1} + h)\mu_1$ of all sectors. Likewise, the fraction of all sectors moving in the opposite direction is $2z_0\mu_0$, since each of the two neck-and-neck firms innovates with probability z_0 . In the steady state, the fraction of firms moving in one direction must equal the fraction moving in the other direction:

$$(z_{-1} + h)\mu_1 = 2z_0(1 - \mu_1),$$

which can be solved for the steady-state fraction of unleveled sectors:

$$\mu_1 = \frac{2z_0}{z_{-1} + h + 2z_0}. \quad (1.16)$$

This implies that the aggregate flow of innovations in all sectors is¹⁸:

$$x = \frac{4(z_{-1} + h)z_0}{z_{-1} + h + 2z_0}.$$

One can show that for ρ large but h not too large, aggregate innovation x follows an inverted-U pattern: it increases with competition Δ for small enough values of Δ and

¹⁸ x is the sum of the two flows: $(z_{-1} + h)\mu_1 + 2z_0(1 - \mu_1)$. But since the two flows are equal, $x = 2(z_{-1} + h)\mu_1$. Substituting for μ_1 using (1.16) yields $x = \frac{4(z_{-1} + h)z_0}{z_{-1} + h + 2z_0}$.

decreases for large enough Δ . The inverted-U shape results from the composition effect whereby a change in competition changes the steady-state fraction of sectors that are in the leveled state, where the escape–competition effect dominates, versus the unleveled state, where the Schumpeterian effect dominates. At one extreme, when there is not much product market competition, there is not much incentive for neck-and-neck firms to innovate, and therefore, the overall innovation rate will be highest when the sector is unleveled. Thus, the industry will be quick to leave the unleveled state (which it does as soon as the laggard innovates) and slow to leave the leveled state (which will not happen until one of the neck-and-neck firms innovates). As a result, the industry will spend most of the time in the leveled state, where the escape–competition effect dominates (z_0 is increasing in Δ). In other words, if the degree of competition is very low to begin with, an increase in competition should result in a faster average innovation rate. At the other extreme, when competition is initially very high, there is little incentive for the laggard in an unleveled state to innovate. Thus, the industry will be slow to leave the unleveled state. Meanwhile, the large incremental profit $\pi_1 - \pi_0$ gives firms in the leveled state a relatively large incentive to innovate, so that the industry will be relatively quick to leave the leveled state. As a result, the industry will spend most of the time in the unleveled state where the Schumpeterian effect is the dominant effect. In other words, if the degree of competition is very high to begin with, an increase in competition should result in a slower average innovation rate.

Finally, using the fact that the log of an industry's output rises by the amount $\ln \gamma$ each time the industry completes two cycles from neck-and-neck ($m = 0$) to unleveled ($m = 1$) and then back to neck-and-neck, the average growth rate of final output g is simply equal to the frequency of completed cycles times $\ln \gamma$. But the frequency of completed cycles is itself equal to the fraction of time μ_1 spent in the unleveled state times the frequency ($z_{-1} + h$) of innovation when in that state. Hence, overall, we have:

$$g = \mu_1 (z_{-1} + h) \ln \gamma = \frac{x}{2} \ln \gamma.$$

Thus, productivity growth follows the same pattern as aggregate innovation with regard to product market competition.

1.3.4 Predictions

The main testable predictions are:

Prediction 1: *The relationship between competition and innovation follows an inverted-U pattern and the average technological gap within a sector (μ_1 in the above model) increases with competition.*

This prediction is tested by [Aghion et al. \(2005\)](#) (hereafter ABBGH) using panel data on UK firms spanning 17 two-digit SIC industries between 1973 and 1994. The chosen measure of product market competition is equal to 1 minus the Lerner index. The Lerner index, or price–cost margin, is itself defined by operating profits net of




Figure 1.1 Competition and innovation (regression lines).

depreciation, provisions and financial cost of capital, divided by sales, averaged across firms within an industry-year. [Figure 1.1](#) shows the inverted-U pattern, and it also shows that if we restrict attention to industries above the median degree of neck-and-neckness, the upward-sloping part of the inverted U is steeper than if we consider the whole sample of industries. ABBGH also show that the average technological gap across firms within an industry increases with the degree of competition the industry is subject to.

Prediction 2: *More intense competition enhances innovation in “frontier” firms but may discourage it in “non-frontier” firms.*

This prediction is tested by [Aghion et al. \(2009\)](#) (hereafter ABGHP). ABGHP use a panel of more than 5000 incumbent lines of businesses in UK firms in 180 four-digit SIC industries over the time period 1987–1993.

Taking the measure of technologically advanced entry of new foreign firms which ABGHP construct from administrative plant-level data as the proxy of competition, [Figure 1.2](#) (taken from ABGHP, 2009) illustrates the following two results. First, the upper line, depicting how productivity growth responds to entry in incumbents that are more-than-median close to the frontier, is upward sloping, and this reflects the escape-competition effect at work in neck-and-neck sectors. Second, the lower line, depicting how productivity growth responds to entry in incumbents that are less-than-median close to the frontier, is downward sloping, which reflects the Schumpeterian effect of competition on innovation in laggards. In the main empirical analysis, ABGHP also control for the influence of trade and average profitability-related competition measures, and address the issue that entry, as well as the other explanatory variables, can be endogenous to incumbent productivity growth, as well as incumbent innovation. To tackle entry endogeneity, in particular, instruments are derived from a broad set of UK and EU-level policy reforms.




Figure 1.2 Entry and growth (regression lines).

Prediction 3: *There is complementarity between patent protection and product market competition in fostering innovation.*

In the above model, competition reduces the profit flow π_0 of non-innovating neck-and-neck firms, whereas patent protection is likely to enhance the profit flow π_1 of an innovating neck-and-neck firm. Both contribute to raising the net profit gain ($\pi_1 - \pi_0$) of an innovating neck-and-neck firm; in other words, both types of policies tend to enhance the escape-competition effect. That competition and patent protection should be complementary in enhancing growth rather than mutually exclusive is at odds with [Romer's \(1990\)](#) product variety model, where competition is always detrimental to innovation and growth (as we discussed above) for exactly the same reason that intellectual property rights (IPRs) in the form of patent protection are good for innovation. Namely, competition reduces post-innovation rents, whereas patent protection increases these rents.¹⁹ Empirical evidence in line with [Prediction 3](#) has recently been provided. [Qian \(2007\)](#) uses the spreading of national pharmaceutical patent laws during the 1980s and 1990s to investigate the effects of patent protection on innovation. She reports that introducing national patent laws stimulates pharmaceutical innovation not on average across all countries, but, among others, in countries with high values of a country-level index of economic

¹⁹ Similarly, in [Boldrin and Levine \(2008\)](#), patenting is detrimental to competition and thereby to innovation for the same reason that competition is good for innovation. To provide support to their analysis the two authors build a growth model in which innovation and growth can occur under perfect competition. The model is then used to argue that monopoly rents and therefore patents are not needed for innovation and growth. On the contrary, patents are detrimental to innovation because they reduce competition.

freedom. The index is the Fraser Institute index, which aggregates proxies of freedom to trade, in addition to measures of access to money, legal structure, and property rights.

Aghion et al. (2013) (hereafter AHP) set out to study whether patent protection can foster innovation when being complemented by product market competition, using country-industry panel data for many industries in OECD countries since the 1980s. AHP find that the implementation of a competition-increasing product market reform, the large-scale European Single Market Program, has increased innovation in industries of countries with strong IPRs since the pre-sample period, but not so in those with weaker IPRs. Moreover, the positive response of innovation to the product market reform in strong IPR countries is more pronounced among firms in industries that rely more on patenting than in other industries. Overall, these empirical results are consistent with a complementarity between IPRs and competition.



1.4. SCHUMPETERIAN GROWTH AND FIRM DYNAMICS

One of the main applications of the Schumpeterian theory has been the study of firm dynamics. The empirical literature has documented various stylized facts using micro firm-level data. Some of these facts are: (i) the firm size distribution is highly skewed; (ii) firm size and firm age are highly correlated; (iii) small firms exit more frequently, but the ones that survive tend to grow faster than the average growth rate; (iv) a large fraction of R&D in the US is done by incumbents; and (v) reallocation of inputs between entrants and incumbents is an important source of productivity growth.

These are some of the well-known empirical facts that non-Schumpeterian growth models cannot account for. In particular, the first four facts listed require a new firm to enter, expand, then shrink over time, and eventually be replaced by new entrants. These and the last fact on the importance of reallocation are all embodied in the Schumpeterian idea of creative destruction.

We will now consider a setup that closely follows the highly influential work by Klette and Kortum (2004). This model will add two elements to the baseline model of Section 1.2: First, innovations will come from both entrants and incumbents. Second, firms will be defined as a collection of production units where successful innovations by incumbents will allow them to expand in product space. Creative destruction will be the central force that drives innovation, invariant firm size distribution, and aggregate productivity growth on a balanced growth path.

1.4.1 The Setup

Time is again continuous and a continuous measure L of individuals work in one of three activities: (i) as production workers, l ; (ii) as R&D scientists in incumbent firms, s_i ; and (iii) as R&D scientists in potential entrants, s_e . The utility function is logarithmic; therefore, the household's Euler equation is $g_t = r_t - \rho$. The final good is produced




Figure 1.3 Example of a firm.

competitively using a combination of intermediate goods according to the following production function:

$$\ln Y_t = \int_0^1 \ln y_{jt} dj, \quad (1.17)$$

where y_j is the quantity produced of intermediate j . Intermediates are produced monopolistically by the innovator who innovated last within that product line j , according to the following linear technology:

$$y_{jt} = A_{jt} l_{jt},$$

where A_{jt} is the product-line-specific labor productivity and l_{jt} is the labor employed for production. This implies that the marginal cost of production in j is simply w_t/A_{jt} where w_t is the wage rate in the economy at time t .

A firm in this model is defined as a collection of n production units (product lines) as illustrated in [Figure 1.3](#). Firms expand in product space through successful innovations. To innovate, firms combine their existing knowledge stock that they accumulated over time (n) with scientists (S_i) according to the following Cobb–Douglas production function:

$$Z_i = \left(\frac{S_i}{\zeta} \right)^{\frac{1}{\eta}} n^{1-\frac{1}{\eta}}, \quad (1.18)$$

where Z_i is the Poisson innovation flow rate, $\frac{1}{\eta}$ is the elasticity of innovation with respect to scientists, and ζ is a scale parameter. Note that this production function generates the

following R&D cost of innovation:

$$C(z_i, n) = \zeta w n z_i^\eta,$$

where $z_i \equiv Z_i/n$ is simply defined as the innovation intensity of the firm. When a firm is successful in its current R&D investment, it innovates over a random product line $j' \in [0, 1]$. Then, the productivity in line j' increases from $A_{j'}$ to $\gamma A_{j'}$. The firm becomes the new monopoly producer in line j' and thereby increases the number of its production lines to $n + 1$. At the same time, each of its n current production lines is subject to the creative destruction x by new entrants and other incumbents. Therefore, during a small time interval dt , the number of production units of a firm increases to $n + 1$ with probability $Z_i dt$ and decreases to $n - 1$ with probability $n x dt$. A firm that loses all of its product lines exits the economy.

1.4.2 Solving the Model

As before, our focus is on a balanced growth path, where all aggregate variables grow at the same rate g (to be determined). We will now proceed in two steps. First, we will solve for the static production decision and then turn to the dynamic innovation decision of firms, which will determine the equilibrium rate of productivity growth, as well as various firm moments along with the invariant firm size distribution.

1.4.2.1 Static Production Decision

As in [Section 1.3](#), the final good producer spends the same amount Y_t on each variety j . As a result, the final good production function in (1.17) generates a unit-elastic demand with respect to each variety: $y_{jt} = Y_t/p_{jt}$. Combined with the fact that firms in a single-product line compete à la Bertrand, this implies that a monopolist with marginal cost w_t/A_{jt} will follow limit pricing by setting its price equal to the marginal cost of the previous innovator $p_{jt} = \gamma w_t/A_{jt}$. The resulting equilibrium quantity and profit in product line j are:

$$y_{jt} = \frac{A_{jt} Y_t}{\gamma w_t} \quad \text{and} \quad \pi_{jt} = \pi Y_t, \quad (1.19)$$

where $\pi \equiv \frac{\gamma-1}{\gamma}$. Note that profits are constant across product lines, which will significantly simplify the aggregation up to the firm level. Note also that the demand for production workers in each line is simply $Y_t/(\gamma w_t)$.

1.4.2.2 Dynamic Innovation Decision

Next we turn to the innovation decision of the firms. The stock-market value of an n -product firm $V_t(n)$ at date t satisfies the Bellman equation:

$$rV_t(n) - \dot{V}_t(n) = \max_{z_i \geq 0} \left\{ \begin{array}{l} n\pi_t - w_t \zeta n z_i^\eta \\ + n z_i [V_t(n+1) - V_t(n)] \\ + n x [V_t(n-1) - V_t(n)] \end{array} \right\}. \quad (1.20)$$

The intuition behind this expression is as follows. The firm collects a total of $n\pi_t$ profits from n product lines and invests in total $w_t\zeta nz_i^\eta$ in R&D. As a result, it innovates at the flow rate $Z_i \equiv nz_i$, in which case it gains $V_t(n+1) - V_t(n)$. In addition, the firm loses each of its product lines through creative destruction at the rate x , which means that a production line will be lost overall at a rate nx , leading to a loss of $V_t(n) - V_t(n-1)$. It is a straightforward exercise to show that the value function in (1.20) is linear in the number of product lines n and proportional to aggregate output Y_t , with the form:

$$V_t(n) = nvY_t.$$

In this expression, $\nu = V_t(n)/nY_t$ is simply the average normalized value of a production unit that is endogenously determined as:

$$\nu = \frac{\pi - \zeta\omega z_i^\eta}{\rho + x - z_i}. \quad (1.21)$$

Note that this expression uses the Euler equation $\rho = r - g$ and that labor share is defined as $\omega \equiv w_t/Y_t$, which is constant on a balanced growth path. In the absence of incumbent innovation, i.e. $z_i = 0$, this value is equivalent to the baseline model (1.1). The fact that incumbents can innovate modifies the baseline value in two opposite directions: First, the cost of R&D investment is subtracted from the gross profit, which lowers the net instantaneous return $\pi - \zeta\omega z_i^\eta$. However, each product line comes with an R&D option value, that is, having one more production unit increases the firm's R&D capacity as in (1.18) and therefore the firm's value.

The equilibrium innovation decision of an incumbent is simply found through the first-order condition of (1.20):

$$z_i = \left(\frac{\nu}{\eta\zeta\omega} \right)^{\frac{1}{\eta-1}}. \quad (1.22)$$

As expected, innovation intensity is increasing in the value of innovation ν and decreasing in the labor cost ω .

1.4.2.3 Free Entry

We consider a mass of entrants that produce one unit of innovation by hiring ψ number of scientists. When a new entrant is successful, it innovates over a random product line by improving its productivity by $\gamma > 1$. It then starts out as a single-product firm. Let us denote the entry rate by z_e . The free-entry condition equates the value of a new entry $V_t(1)$ to the cost of innovation ψw_t , such that:

$$\nu = \omega\psi. \quad (1.23)$$

Recall that the rate of creative destruction is simply the entry rate plus an incumbent's innovation intensity, i.e. $x = z_i + z_e$. Using this fact, together with (1.21)–(1.23), delivers the equilibrium entry rate and incumbent innovation intensity:

$$z_e = \frac{\pi}{\omega\psi} - \frac{1}{\eta} \left(\frac{\psi}{\eta\zeta} \right)^{\frac{1}{\eta-1}} - \rho \quad \text{and} \quad z_i = \left(\frac{\psi}{\eta\zeta} \right)^{\frac{1}{\eta-1}}.$$

1.4.2.4 Labor Market Clearing

Now we are ready to close the model by imposing the labor market clearing condition. The equilibrium labor share ω equates the supply of labor L to the sum of aggregate labor demand coming from (i) production, $(\gamma\omega)^{-1}$, (ii) incumbent R&D, $\zeta \left(\frac{\psi}{\eta\zeta} \right)^{\frac{\eta}{\eta-1}}$, and (iii) outside entrants, $\frac{\pi}{\omega} - \zeta \left(\frac{\psi}{\eta\zeta} \right)^{\frac{\eta}{\eta-1}} - \psi\rho$. The resulting labor share is:

$$\omega = \frac{w_t}{Y_t} = \frac{1}{L + \rho\psi}.$$

1.4.3 Equilibrium Growth Rate

In this model, innovation takes place by both incumbents and entrants at the total rate of $x = z_i + z_e$. Hence, the equilibrium growth rate is:

$$g = x \ln \gamma = \left[\left(\frac{\gamma - 1}{\gamma} \right) \frac{L}{\psi} + \left(\frac{\eta - 1}{\eta} \right) \left(\frac{\psi}{\eta\zeta} \right)^{\frac{1}{\eta-1}} - \frac{\rho}{\gamma} \right] \ln \gamma.$$

In addition to the standard effects, such as the growth rate increasing in the size of innovation and decreasing in the discount rate, this model generates an interesting non-linear relationship between entry cost ψ and growth. An increase in the entry cost reduces the entry rate and therefore has a negative effect on equilibrium growth. However, this effect also frees up those scientists that used to be employed by outside entrants and reallocates them to incumbents, hence increasing innovation by incumbents and growth. This is an interesting trade-off for industrial policy. In a recent work, [Acemoglu et al. \(2013\)](#) analyze the effects of various industrial policies on equilibrium productivity growth, including entry subsidy and incumbent R&D subsidy, in an enriched version of the above framework.

1.4.4 Predictions

Now we go back to the initial list of predictions and discuss how they are captured by the above model.

Prediction 1: *The size distribution of firms is highly skewed.*

In this model, firm size is summarized by the number of product lines of a firm. Let us denote by μ_n the fraction of firms that have n products. The invariant distribution μ_n is found by equating the inflows into state n to the outflows from it:

$$\begin{aligned}\mu_1 x &= z_e, \\ (z_i + x) \mu_1 &= \mu_2 2x + z_e, \\ (z_i + x) n \mu_n &= \mu_{n+1} (n+1) x + \mu_{n-1} (n-1) z_i \quad \text{for } n \geq 2.\end{aligned}$$

The first line equates exits to entry. The left-hand side of the second line consists of outflows from being a one-product firm that happen when a one-product firm innovates itself and becomes a two-product firm or is replaced by another firm at the rate x . The right-hand side is the sum of the inflows coming from two-product firms or from outsiders. The third line generalizes the second line to n -product firms. The resulting firm size distribution is geometric as illustrated in Figure 1.4 and has the following exact form:

$$\mu_n (z_e/z_i) = \frac{z_e/z_i}{(1 + z_e/z_i)^n n},$$

and highly skewed as shown in a vast empirical literature (Simon and Bonini, 1958; Ijiri and Simon, 1977; Schmalensee, 1989; Stanley et al. 1995; Axtell, 2001; Rossi-Hansberg and Wright, 2007). Several alternative Schumpeterian models have been proposed after (Klette and Kortum, 2004) that feature invariant firm size distributions with a Pareto tail. (See Acemoglu and Cao (2011) for an example and a discussion of the literature.)

Prediction 2: *Firm size and firm age are positively correlated.*

In the current model, firms are born with a size of 1. Subsequent successes are required for firms to grow in size, which naturally produces a positive correlation between size and age. This regularity has been documented extensively in the literature. (For recent discussions and additional references, see Haltiwanger et al. (2010) and Akcigit and Kerr (2010)).

Prediction 3: *Small firms exit more frequently. The ones that survive tend to grow faster than average.*

In the above model, firm exit happens through the loss of product lines. Conditional on not producing a new innovation, a firm's probability of losing all of its product lines and exiting within a period is $(x\Delta t)^n$, which decreases in n . Clearly it becomes much more difficult for a firm to exit when it expands in product space.

The facts that small firms exit more frequently and grow faster conditional on survival have been widely documented in the literature (for early work, see Birch (1981, 1987) and Davis et al. (1996). For more recent work, see Haltiwanger et al. (2010), Akcigit and Kerr (2010), and Neumark et al. (2008)).

Prediction 4: *A large fraction of R&D is done by incumbents.*

There is an extensive literature that studies R&D investment and the patenting behavior of existing firms in the US (see, for instance, among many others, Acs and Audretsch




Figure 1.4 Firm size distribution.

(1988, 1991), Griliches (1990), Hall et al. (2001), Cohen (1995), and Cohen and Klepper (1996)). In particular, Freeman (1982), Pennings and Buitendam (1987), Tushman and Anderson (1986), Scherer (1984), and Akcigit and Kerr (2010) show that large incumbents focus on improving the existing technologies, whereas small new entrants focus on innovating with radical new products or technologies. Similarly, Akcigit et al. (2012) provide empirical evidence on French firms showing that large incumbents with a broad technological spectrum account for most of the private basic research investment.

On the theory side, Akcigit and Kerr (2010), Acemoglu and Cao (2011), and Acemoglu et al. (2012) have also provided alternative Schumpeterian models that capture this fact.

Prediction 5: *Both entrants and incumbents innovate. Moreover, the reallocation of resources among incumbents, as well as from incumbents to new entrants, is the major source of productivity growth.*

A central feature of this model is that both incumbents and entrants innovate and contribute to productivity growth. New entrants account for:

$$\frac{z_e}{z_e + z_i} = 1 - \left[\left(\left(\frac{\gamma - 1}{\gamma} \right) \frac{L}{\psi} - \frac{\rho}{\gamma} \right) \left(\frac{\eta \zeta}{\psi} \right)^{\frac{1}{\eta - 1}} + \frac{\eta - 1}{\eta} \right]^{-1},$$

percent of innovations in any given period. [Bartelsman and Doms \(2000\)](#) and [Foster et al. \(2001\)](#) have shown that 25% of productivity growth in the US is accounted for by new entry and the remaining 75% by continuing plants. Moreover, [Foster et al. \(2001, 2006\)](#) have shown that reallocation of resources through entry and exit accounts for around 50% of manufacturing and 90% of US retail productivity growth. In a recently growing cross-country literature, [Hsieh and Klenow \(2009, 2012\)](#), [Bartelsman et al. \(2009\)](#), and [Syverson \(2011\)](#) describe how variations in reallocation across countries explain differences in productivity levels. [Lentz and Mortensen \(2008\)](#) and [Acemoglu et al. \(2013\)](#) estimate variants of the baseline model in [Klette and Kortum \(2004\)](#) to quantify the importance of reallocation and study the impacts of industrial policy on reallocation and productivity growth.



1.5. GROWTH MEETS DEVELOPMENT

In this section, we argue that Schumpeterian growth theory helps bridge the gap between growth and development economics, by offering a simple framework to capture the idea that growth-enhancing policies or institutions may vary with a country's level of technological development. In particular, we will look at the role of democracy in the growth process, arguing that democracy matters for growth to a larger extent in more advanced economies.

1.5.1 Innovation Versus Imitation and the Notion of Appropriate Institutions

Innovations in one sector or one country often build on knowledge that was created by innovations in another sector or country. The process of diffusion, or technology spillover, is an important factor behind cross-country convergence. [Howitt \(2000\)](#) showed how this can lead to cross-country conditional convergence of growth rates in Schumpeterian growth models. Specifically, a country that starts far behind the world technology frontier can grow faster than one close to the frontier because the former country will make a larger technological advance every time one of its sectors catches up to the global frontier. In [Gerschenkron's \(1962\)](#) terms, countries far from the frontier enjoy an "advantage of backwardness." This advantage implies that, in the long run, a country with a low rate of innovation will fall behind the frontier but will grow at the same rate as the frontier; as they fall further behind, the advantage of backwardness eventually stabilizes the gap that separates them from the frontier.

These same considerations imply that policies and institutions that are appropriate for countries close to the global technology frontier are often different from those that are appropriate for non-frontier countries, because those policies and institutions that help a country to copy, adapt, and implement leading-edge technologies are not necessarily the same as those that help it to make leading-edge innovations. The idea of appropriate institutions was developed more systematically by [Acemoglu et al. \(2006\)](#), henceforth

AAZ, and it underlies more recent work, in particular, Acemoglu and Robinson's best-selling book *Why Nations Fail* (Acemoglu and Robinson (2012)), in which the authors rely on a rich set of country studies to argue that sustained growth requires creative destruction and therefore is not sustainable in countries with extractive institutions.

A particularly direct and simpler way to formalize the idea of appropriate growth policy is to move for a moment from continuous to discrete time. Following AAZ and more remotely (Nelson and Phelps, 1966), let A_t denote the current average productivity in the domestic country, and \bar{A}_t denote the current (world) frontier productivity. Then, think of innovation as multiplying productivity by factor γ , and of imitation as catching up with the frontier technology.

Then, if the fraction μ_n of sectors innovates and the fraction μ_m imitates, we have:

$$A_{t+1} - A_t = \mu_n (\gamma - 1) A_t + \mu_m (\bar{A}_t - A_t).$$

This in turn implies that productivity growth hinges upon the country's degree of "frontierness," i.e. its "proximity" $a_t = A_t/\bar{A}_t$ to the world frontier, namely:

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

In particular:

Prediction 1: *The closer to the frontier an economy is, that is, the closer to one the proximity variable a_t is, the more is growth driven by "innovation-enhancing" rather than "imitation-enhancing" policies or institutions.*

1.5.2 Further Evidence on Appropriate Growth Policies and Institutions

In Section 1.3 we mentioned some recent evidence for the prediction that competition and free-entry should be more growth-enhancing. Using a cross-country panel of more than 100 countries over the 1960–2000 period, AAZ regress the average growth rate on a country's distance to the US frontier (measured by the ratio of GDP per capita in that country to per capita GDP in the US) at the beginning of the period. Then, they split the sample of countries into two groups, corresponding respectively to countries that are more open than the median and to countries that are less open than the median. The prediction is:

Prediction 2: *Average growth should decrease more rapidly as a country approaches the world frontier when openness is low.*

To measure openness one can use imports plus exports divided by aggregate GDP. But this measure suffers from obvious endogeneity problems; in particular, exports and imports are likely to be influenced by domestic growth. To deal with the endogeneity problem, Frankel and Romer (1999) construct a more exogenous measure of openness that relies on exogenous characteristics such as land area, common borders, geographical

distance, population, etc. and it is this measure that AAZ use to measure openness in the following figures.

Figure 1.5A and B shows the cross-sectional regression. Here, average growth over the whole 1960–2000 period is regressed over the country's distance to the world technology frontier in 1965, respectively for less open and more open countries. A country's distance




Figure 1.5 Growth, openness and distance to frontier. A: less open countries (cross-section) B: more open countries (cross-section) C: less open countries (Panel) D: more open countries (panel).

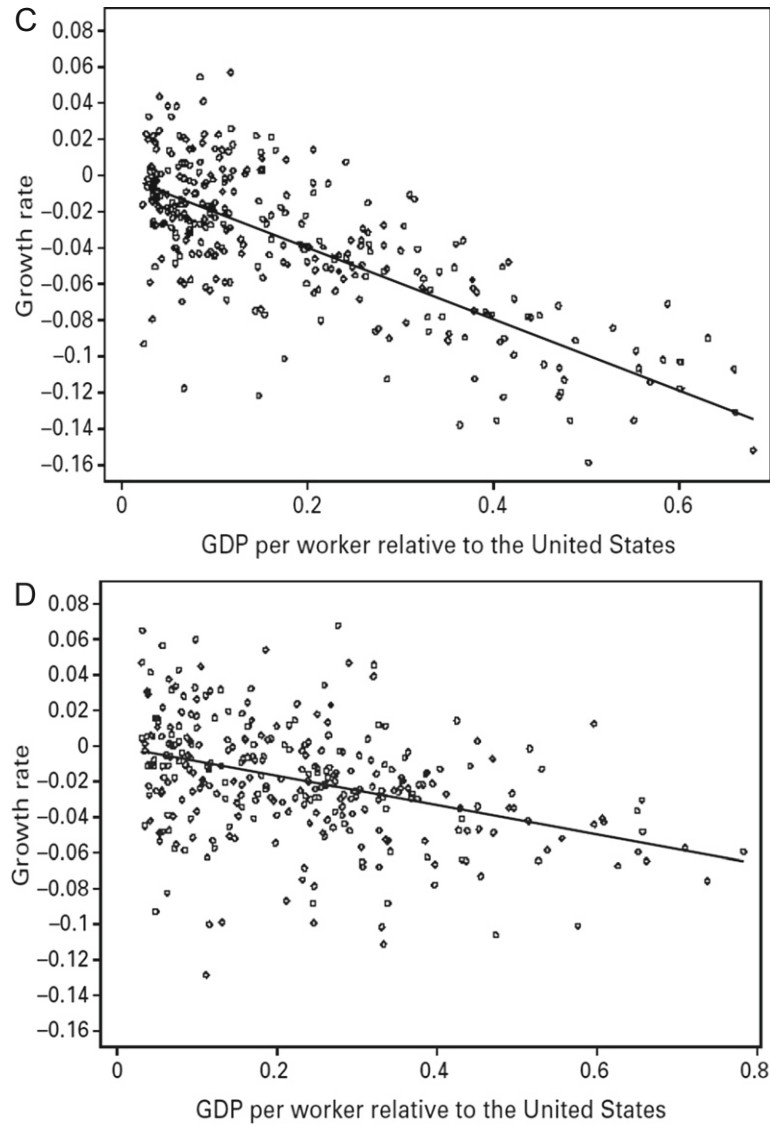


Figure 1.5 (Continued).

to the frontier is measured by the ratio between the log of this country's level of per capita GDP and the maximum of the logs of per capita GDP across all countries (which corresponds to the log of per capita GDP in the US).²⁰

²⁰ That the regression lines should all be downward sloping reflects the fact that countries farther below the world technology frontier achieve bigger technological leaps whenever they successfully catch up with

Figure 1.5C and D shows the results of panel regressions where AAZ decompose the period 1960–2000 in 5-year subperiods and then for each subperiod AAZ regress average growth over the period on distance to the frontier at the beginning of the subperiod, respectively for less open and more open countries. These latter regressions control for country fixed effects. In both cross-sectional and panel regressions we see that while a low degree of openness does not appear to be detrimental to growth in countries far below the world frontier, it becomes increasingly detrimental to growth as the country approaches the frontier.

AAZ repeat the same exercise using entry costs faced by new firms instead of openness. The prediction is:

Prediction 3: *High entry barriers become increasingly detrimental to growth as the country approaches the frontier.*

Entry costs in turn are measured by the number of days to create a new firm in the various countries (see Djankov et al. 2002). Here, the country sample is split between countries with high barriers relative to the median and countries with low barriers relative to the median. Figure 1.6A and B shows the cross-sectional regressions, respectively, for high and low barrier countries, whereas Figure 1.6C and D shows the panel regressions for the same two subgroups of countries. Both types of regressions show that while high entry barriers do not appear to be detrimental to growth in countries far below the world frontier, they indeed become increasingly detrimental to growth as the country approaches the frontier.

These two empirical exercises point to the importance of interacting institutions or policies with technological variables in growth regressions: openness is particularly growth-enhancing in countries that are closer to the technological frontier; entry is more growth-enhancing in countries or sectors that are closer to the technological frontier; below we will see that higher (in particular, graduate) education tends to be more growth-enhancing in countries or in US states that are closer to the technological frontier, whereas primary-secondary (possibly undergraduate) education tends to be more growth enhancing in countries or in US states that are farther below the frontier.

A third piece of evidence is provided by Aghion et al. (2009), who use cross-US-states panel data to look at how spending on various levels of education matter differently for growth across US states with different levels of frontierness as measured by their average productivity compared to frontier-state (Californian) productivity. The gray bars in Figure 1.7 do not factor in the mobility of workers across US states, whereas the solid black bars do. The more frontier a country or region is, the more its growth relies on frontier innovation and therefore our prediction is:

the frontier (this is the “advantage of backwardness” we mentioned above). More formally, for given μ_n and $\mu_{m,g,t} = \mu_n(\gamma - 1) + \mu_m(a_t^{-1} - 1)$ is decreasing in a_t .

Prediction 4: *The more frontier an economy is, the more growth in this economy relies on research education.*

As shown in the figure below, research-type education is always more growth-enhancing in states that are more frontier, whereas a bigger emphasis on 2-year colleges is more growth-enhancing in US states that are farther below the productivity frontier. This is not surprising: [Vandenbussche et al. \(2006\)](#) obtain similar conclusions using




Figure 1.6 Growth, entry and distance to frontier. A: high barrier countries (cross-section) B: low barrier countries (cross-section) C: high barrier countries (panel) D: low barrier countries (panel).

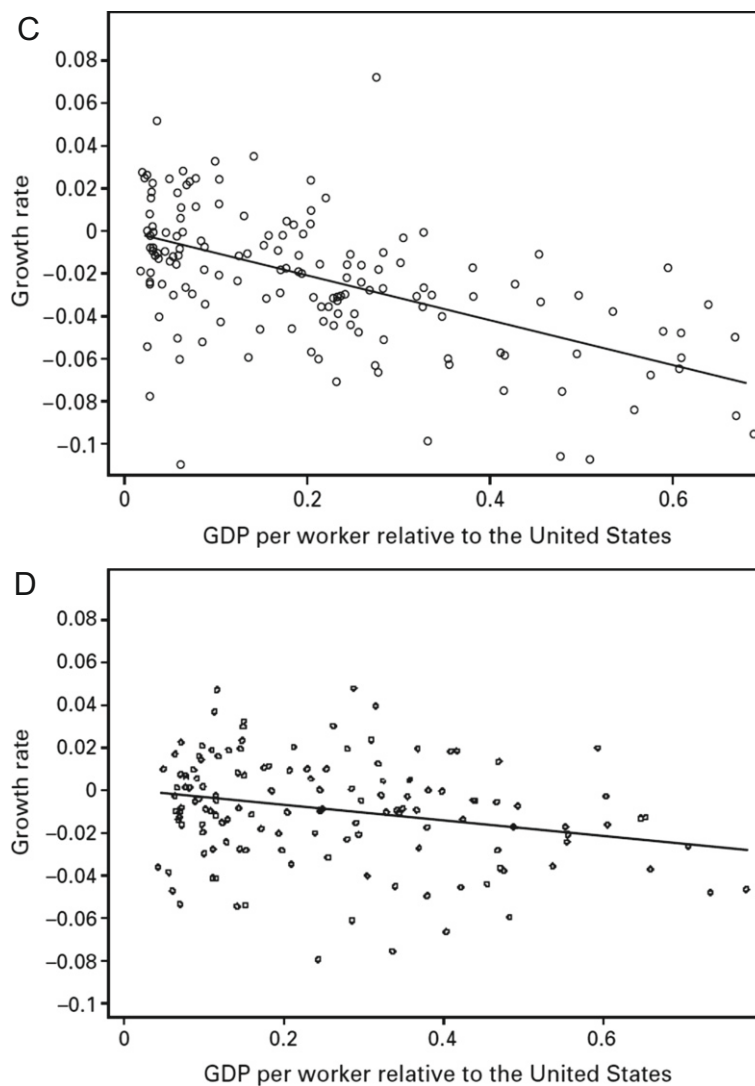


Figure 1.6 (Continued).

cross-country panel data, namely, that tertiary education is more positively correlated with productivity growth in countries that are closer to the world technology frontier.

1.5.3 Political Economy of Creative Destruction

Does democracy enhance or hamper economic growth? One may think of various channels whereby democracy should affect per capita GDP growth. A first channel is that democracy pushes for more redistribution from rich to poor, and that redistribution in

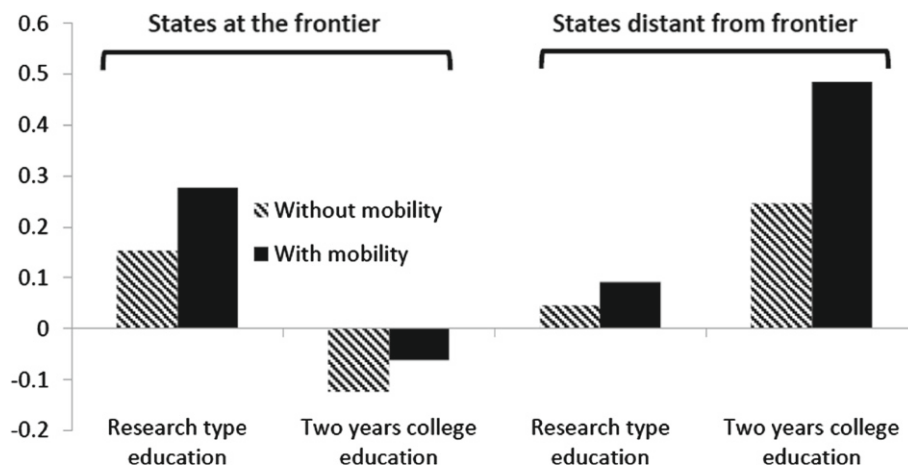


Figure 1.7 Growth, education, and distance to frontier.

turn affects growth. Thus, [Persson and Tabellini \(1994\)](#) and [Alesina and Rodrik \(1994\)](#) analyze the relationship between inequality, democratic voting, and growth. They develop models in which redistribution from rich to poor is detrimental to growth as it discourages capital accumulation. More inequality is then also detrimental to growth because it results in the median voter becoming poorer and therefore demanding more redistribution. A second channel, which we explore in this section, is Schumpeterian: namely, democracy reduces the scope for expropriating successful innovators or for incumbents to prevent new entry by using political pressure or bribes. In other words, democracy facilitates creative destruction and thereby encourages innovation.²¹ To the extent that innovation matters more for growth in more frontier economies, the prediction is:

Prediction 5: *The correlation between democracy and innovation/growth is more positive and significant in more frontier economies.*

The relationship between democracy, “frontierness” and growth, thus provides yet another illustration of our notion of appropriate institutions. In the next subsection we develop a simple Schumpeterian model that generates this prediction.

1.5.3.1 The Formal Argument

Consider the following Schumpeterian model in discrete time. All agents and firms live for one period. In each period t a final good (henceforth the numeraire) is produced in each state by a competitive sector using a continuum one of intermediate inputs,

²¹ [Acemoglu and Robinson \(2006\)](#) formalize another reason, also Schumpeterian, as to why democracy matters for innovation: namely, new innovations not only destroy the economic rents of incumbent producers, they also threaten the power of incumbent political leaders.

according to the technology:

$$\ln Y_t = \int_0^1 \ln y_{jt} dj, \quad (1.24)$$

where the intermediate products are produced again by labor according to:

$$y_{jt} = A_{jt} l_{jt}. \quad (1.25)$$

There is a competitive fringe of firms in each sector that are capable of producing a product with technology level A_{jt}/γ . So, as before, each incumbent's profit flow is:

$$\pi_{jt} = \pi Y_t,$$

where $\pi \equiv \frac{\gamma-1}{\gamma}$. Note that as in (1.19), each incumbent will produce using the same amount of labor:

$$l_{jt} = \frac{Y_t}{\gamma w_t} \equiv l, \quad (1.26)$$

where l is the economy's total use of manufacturing labor. We assume that there is measure one unit of labor that is used only for production. Therefore $l = 1$ implies:

$$w_t = \frac{Y_t}{\gamma}.$$

Finally, (1.24)–(1.26) deliver the final output as a function of the aggregate productivity A_t in this economy:

$$Y_t = A_t,$$

where $\ln A_t \equiv \int_0^1 \ln A_{jt} dj$ is the end-of-period- t aggregate productivity index.

Technology and entry Let \bar{A}_t denote the new world productivity frontier at date t and assume that:

$$\bar{A}_t = \gamma \bar{A}_{t-1},$$

with $\gamma > 1$ exogenously given. We shall again emphasize the distinction already made in the previous section between sectors in which the incumbent producer is neck-and-neck with the frontier and those in which the incumbent firm is below the frontier; at the beginning of date t , a sector j can either be at the current frontier, with productivity level $A_{jt}^b = \bar{A}_{t-1}$ (advanced sector) or one step below the frontier, with productivity level $A_{jt}^b = \bar{A}_{t-2}$ (backward sector). Thus, imitation—or knowledge spillovers—in this model means that whenever the frontier moves up one step from \bar{A}_{t-1} to \bar{A}_t , then backward sectors also automatically move up one step from \bar{A}_{t-2} to \bar{A}_{t-1} .

In each intermediate sector j , only one incumbent firm I_j , and one potential entrant E_j , are active in each period. In this model, innovation in a sector is made only by

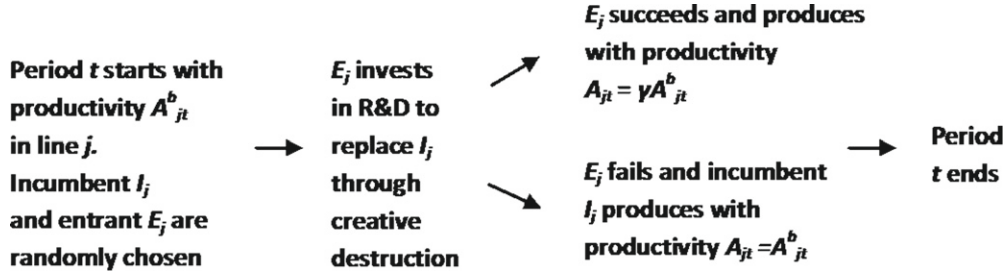


Figure 1.8 Timing of events.

a potential entrant E_j since innovation does not change the incumbent's profit rate. Before production takes place, potential entrant E_j invests in R&D in order to replace the incumbent I_j . If successful, it increases the current productivity of sector j to $A_{jt} = \gamma A_{jt}^b$ and becomes the new monopolist and produces. Otherwise, the current incumbent preserves its monopoly right and produces with the beginning-of-period productivity $A_{jt} = A_{jt}^b$ and the period ends. The timing of events is described in Figure 1.8.

Finally, the innovation technology is as follows: if a potential entrant E_j spends $A_t \lambda z_{jt}^2 / 2$ on R&D in terms of the final good, then she innovates with probability z_{jt} .

Democracy Entry into a sector is subject to the democratic environment in the domestic country. Similar to Acemoglu and Robinson (2006), we model democracy as freedom to enter. More specifically, in a country with democracy level $\beta \in [0, 1]$, a successful innovation leads to successful entry only with probability β , and it is blocked with probability $(1 - \beta)$. As a result, the probability of an unblocked entry is βz_j . An unblocked entrant raises productivity from A_{jt}^b to γA_{jt}^b and becomes the new monopoly producer.

Equilibrium innovation investments We can now analyze the innovation decision of the potential entrant E_j :

$$\max_{z_{jt}} \left\{ z_{jt} \beta \pi Y_t - A_t \lambda \frac{z_{jt}^2}{2} \right\}.$$

In equilibrium we get:

$$z_{jt} = \bar{z} = \frac{\beta \pi}{\lambda},$$

where we used the fact that $Y_t = A_t$. Thus, the aggregate equilibrium innovation effort is increasing in profit π and decreasing in R&D cost λ . Most important for us in this section, the innovation rate is increasing in the democracy level β :

$$\frac{\partial \bar{z}}{\partial \beta} > 0.$$

Growth Now we can turn to the equilibrium growth rate of average productivity. We will denote the fraction of advanced sectors by μ , which will also be the index for the frontierness of the domestic country. The average productivity of a country at the beginning of date t is:

$$A_{t-1} \equiv \int_0^1 A_{jt} dj = \mu \bar{A}_{t-1} + (1 - \mu) \bar{A}_{t-2}.$$

Average productivity at the end of the same period is:²²

$$A_t = \mu [\beta \bar{z} \gamma \bar{A}_{t-1} + (1 - \beta \bar{z}) \bar{A}_{t-1}] + (1 - \mu) \bar{A}_{t-1}.$$

Then the growth rate of average productivity is simply equal to:

$$g_t = \frac{A_t - A_{t-1}}{A_{t-1}} = \gamma \frac{\mu \beta \bar{z} (\gamma - 1) + 1}{\mu (\gamma - 1) + 1} - 1 > 0.$$

As is clear from the above expression, democracy is always growth-enhancing:

$$\frac{\partial g_t}{\partial \beta} = \left(\bar{z} + \frac{\partial \bar{z}}{\partial \beta} \beta \right) \frac{\gamma \mu (\gamma - 1)}{\mu (\gamma - 1) + 1} > 0.$$

Moreover, democracy is more growth enhancing the closer the domestic country is to the world technology frontier:

$$\frac{\partial^2 g_t}{\partial \beta \partial \mu} = \left(\bar{z} + \frac{\partial \bar{z}}{\partial \beta} \beta \right) \frac{(\gamma - 1) \gamma}{[\mu (\gamma - 1) + 1]^2} > 0.$$

This result is quite intuitive. Democratization allows for more turnover which in turn encourages outsiders to innovate and replace the incumbents. Since frontier countries rely more on innovation and benefit less from imitation or spillover, the result follows.

1.5.3.2 Evidence

A first piece of evidence supporting [Prediction 5](#) is provided by [Aghion et al. \(2007\)](#), henceforth AAT. The paper uses employment and productivity data at the industry level across countries and over time. Their sample includes 28 manufacturing sectors for 180 countries over the period 1964–2003. Democracy is measured using the Polity 4 indicator, which itself is constructed from combining constraints on the executive; the openness and competitiveness of executive recruitment; and the competitiveness of political participation. Frontierness is measured by the log of the value added of a sector divided by the maximum of the log of the same variable in the same sectors across all countries or

²² Here we make use of the assumption that backward sectors are automatically upgraded as the technology frontier moves up.

by ratio of the log of GDP per worker in the sector over the maximum of the log of per capita GDP in similar sectors across all countries. AAT take one minus these ratios as proxies for a sector's distance to the technological frontier. AAT focus on 5-year and 10-year growth rates. They compute rates over non-overlapping periods and in particular 5-year growth rates are computed over the periods 1975, 1980, 1985, 1990, 1995, and 2000. For the 10-year growth rates they use alternatively the years 1975, 1985, 1995, and the years 1980, 1990, and 2000.

AAT regress the growth of either value added or employment in an industrial sector on democracy (and other measures of civil rights), the country's or industry's frontierness, and the interaction term between the latter two. AAT also add time, country, and industry fixed effects.

The result is that the interaction coefficient between frontierness and democracy is positive and significant, meaning that the more frontier the industry is, the more growth-enhancing is democracy in the country for that sector. [Figure 1.9](#) below provides an illustration of the results. It plots the rate of value-added growth against a measure of the country's proximity to the technological frontier (namely, the ratio of the country's labor productivity to the frontier labor productivity). The dotted line shows the linear regression of industry growth on democracy for countries that are less democratic than the median country (on the democracy scale), whereas the solid line shows the corresponding relationship for countries that are more democratic than the median country. We see that growth is higher in more democratic countries when they are close to the technological frontier, but not when they are far below the frontier.




Figure 1.9 Growth, democracy, and distance to frontier (regression lines).



1.6. SCHUMPETERIAN WAVES

What causes long-term accelerations and slowdowns in economic growth and underlies the long swings sometimes referred to as Kondratieff cycles? In particular, what caused American growth in GDP and productivity to accelerate starting in the mid-1990s? The most popular explanation relies on the notion of general-purpose technologies (GPTs).

Bresnahan and Trajtenberg (1995) define a GPT as a technological innovation that affects production and/or innovation in many sectors of an economy. Well-known examples in economic history include the steam engine, electricity, the laser, turbo reactors, and more recently the information technology (IT) revolution. Three fundamental features characterize most GPTs. First, their pervasiveness: GPTs are used in most sectors of an economy and thereby generate palpable macroeconomic effects. Second, their scope for improvement: GPTs tend to underperform upon being introduced; only later do they fully deliver their potential productivity growth. Third, innovation spanning: GPTs make it easier to invent new products and processes—that is, to generate new secondary innovations—of higher quality.

Although each GPT raises output and productivity in the long run, it can also cause cyclical fluctuations while the economy adjusts to it. As David (1990) and Lipsey and Bekar (1995) have argued, GPTs like the steam engine, the electric dynamo, the laser, and the computer require costly restructuring and adjustment to take place, and there is no reason to expect this process to proceed smoothly over time. Thus, contrary to the predictions of real-business-cycle theory, the initial effect of a positive technology shock may not be to raise output, productivity, and employment, but to reduce them.²³

Note that GPTs are Schumpeterian in nature, as they typically lead to older technologies in all sectors of the economy being abandoned as they diffuse to these sectors. Thus, it is no surprise that Helpman and Trajtenberg (1998) used the Schumpeterian apparatus to develop their model of GPT and growth. The basic idea of this model is that GPTs do not come ready to use off the shelf. Instead, each GPT requires an entirely new set of intermediate goods before it can be implemented. The discovery and development of these intermediate goods is a costly activity, and the economy must wait until some critical mass of intermediate components has been accumulated before it is profitable for firms to switch from the previous GPT. During the period between the discovery of a new GPT and its ultimate implementation, national income will fall as resources are taken out of production and put into R&D activities aimed at the discovery of new intermediate input components.

²³ For instance, Greenwood and Yorukoglu (1974) and Hornstein and Krusell (1996) have studied the productivity slowdown during the late 1970s and early 1980s caused by the IT revolution.

1.6.1 Back to the Basic Schumpeterian Model

As a useful first step toward a growth model with GPT, let us go back to the basic Schumpeterian model laid out in [Section 1.2](#), but present it somewhat differently. Recall that the representative household has linear utility and the final good is produced with a single intermediate product according to:

$$Y_t = A_t \gamma^\alpha,$$

where γ is the flow of intermediate input and A is the productivity parameter measuring the quality of intermediate input γ .

Each innovation results in an intermediate good of higher quality. Specifically, a new innovation multiplies the productivity parameter A_k by $\gamma > 1$, so that:

$$A_{k+1} = \gamma A_k.$$

Innovations in turn arrive discretely with Poisson rate λz , where z is the current flow of research.

In the steady state the allocation of labor between research and manufacturing remains constant over time, and is determined by the research-arbitrage equation:

$$\omega_k = \lambda \gamma \nu_k, \quad (1.27)$$

where the LHS of (1.27) is the productivity-adjusted wage $\omega_k \equiv w_k/A_k$, which a worker earns by working in the manufacturing sector; $\nu_k \equiv V_k/A_k$ is the productivity-adjusted value and $\lambda \gamma \nu_k$ is the expected reward from investing one unit flow of labor in research.²⁴ The productivity-adjusted value ν_k of an innovation is in turn determined by the Bellman equation:

$$\rho \nu_k = \tilde{\pi}(\omega_k) - \lambda z \nu_k, \quad (1.28)$$

where $\pi(\omega_k) = A_k [1 - \alpha] \alpha^{\frac{1+\alpha}{1-\alpha}} \omega_k^{\frac{\alpha}{\alpha-1}}$ is the equilibrium profit and $\tilde{\pi}(\omega_k) \equiv \pi(\omega_k)/A_k$ denotes the productivity-adjusted flow of monopoly profits accruing to a successful innovator and we used the fact that $r_t = \rho$. In (1.28) the term $(-\lambda z \nu)$ corresponds to the capital loss involved in being replaced by a subsequent innovator. In the steady state, the productivity-adjusted variables ω_k and ν_k remain constant; therefore, the subscript k will henceforth be dropped.

The above arbitrage equation, which can now be re-expressed as:

$$\omega = \lambda \gamma \frac{\tilde{\pi}(\omega)}{\rho + \lambda z},$$

²⁴ Equation (1.27) is just a rewrite of Equation (R) in [Section 1.2](#). Recall that the latter is expressed as:

$$w_k = \lambda V_{k+1};$$

using the fact that $V_{k+1} = \gamma V_k$, this immediately leads to [Equation \(1.27\)](#).

the labor-market clearing condition:

$$\gamma(\omega) + z = L,$$

where $\gamma(\omega)$ is the manufacturing demand for labor, jointly determine the steady-state amount of research z as a function of the parameters $\lambda, \gamma, L, \rho, \alpha$.

In a steady state the flow of the final good produced between the k th and $(k + 1)$ th innovation is:

$$Y_k = A_k [L - z]^\alpha.$$

Thus, the log of final output increases by $\ln \gamma$ each time a new innovation occurs. Then the average growth rate of the economy is equal to the size of each step $\ln \gamma$ times the average number of innovations per unit of time, λz : i.e.:

$$\mathbb{E}(g) = \lambda z \ln \gamma.$$

Note that this is a one-sector economy where each innovation corresponds by definition to a major technological change (i.e. to the arrival of a new GPT), and thus where growth is uneven with the time path of output being a random step function. But although it is uneven, the time path of aggregate output does not involve any slump. Accounting for the existence of slumps requires an extension of the basic Schumpeterian model, which brings us to the GPT growth model.

1.6.2 A Model of Growth with GPTs

As before, there are L workers who can engage either in the production of existing intermediate goods or in research aimed at discovering new intermediate goods. Again, each intermediate good is linked to a particular GPT. We follow [Helpman and Trajtenberg \(1998\)](#) in supposing that before any of the intermediate goods associated with a GPT can be used profitably in the final-goods sector, some minimal number of them must be available. We lose nothing essential by supposing that this minimal number is one. Once the good has been invented, its discoverer profits from a patent on its exclusive use in production, exactly as in the basic Schumpeterian model reviewed earlier.

Thus, the difference between this model and our basic model is that now the discovery of a new generation of intermediate goods comes in two stages. First, a new GPT must come, and then the intermediate good must be invented that implements that GPT. Neither can come before the other. You need to see the GPT before knowing what sort of good will implement it, and people need to see the previous GPT in action before anyone can think of a new one. For simplicity we assume that no one directs R&D toward the discovery of a new GPT. Instead, the discovery arrives as a serendipitous by-product of learning-by-doing with the previous one.

The economy will pass through a sequence of cycles, each having two phases, as indicated in [Figure 1.10](#). GPT_{*i*} arrives at time t_i . At that time, the economy enters phase




Figure 1.10 Phases of GPT cycles.

1 of the i th cycle. During phase 1, the amount z of labor is devoted to research. Phase 2 begins at time $t_i + \Delta_i$ when this research discovers an intermediate good to implement GPT_i . During phase 2, all labor is allocated to manufacturing until GPT_{i+1} arrives, at which time the next cycle begins. Over the cycle, output is equal to $A_{i-1}F(L - z)$ during phase 1 and to $A_i F(L)$ during phase 2. Thus, the drawing of labor out of manufacturing and into research causes output to fall each time a GPT is discovered, by an amount equal to $A_{i-1}[F(L) - F(L - z)]$.

A steady-state equilibrium is one in which people choose to do the same amount of research each time the economy is in phase 1; that is, z is constant from one GPT to the next. As before, we can solve for the equilibrium value of z using a research-arbitrage equation and a labor-market-equilibrium condition. Let ω_j be the (productivity-adjusted) wage, and v_j the expected (productivity-adjusted) present value of the incumbent (intermediate good) monopolist when the economy is in phase $j \in \{1, 2\}$. In a steady state these productivity-adjusted variables will all be independent of which GPT is currently in use.

Because research is conducted in phase 1 but pays off when the economy enters into phase 2 with a productivity parameter raised by the factor γ , the following research-arbitrage condition must hold in order for there to be a positive level of research in the economy:

$$\omega_1 = \lambda \gamma v_2.$$

Suppose that once we are in phase 2, the new GPT is delivered by a Poisson process with constant arrival rate μ . Then the value v_2 is determined by the Bellman equation:

$$\rho v_2 = \tilde{\pi}(\omega_2) + \mu [v_1 - v_2].$$

By analogous reasoning, we have:

$$\rho v_1 = \tilde{\pi}(\omega_1) - \lambda z v_1.$$

Combining the above three equations yields the research-arbitrage equation:

$$\omega_1 = \frac{\lambda \gamma}{\rho + \mu} \left[\tilde{\pi}(\omega_2) + \frac{\mu \tilde{\pi}(\omega_1)}{\rho + \lambda z} \right]. \quad (1.29)$$

Because no one does research in phase 2, we know that the value of ω_2 is determined independently of research, by the market clearing condition:

$$L = \gamma(\omega_2).$$

Thus, we can take this value as given and regard the preceding research-arbitrage condition (1.29) as determining ω_1 as a function of z . The value of z is then determined, as in the previous subsection, by the labor-market equation:

$$L - z = \gamma(\omega_1).$$

The average growth rate will be the frequency of innovation times the size $\ln \gamma$, for exactly the same reason as in the basic model. The frequency, however, is determined a little differently than before because the economy must pass through two phases. An innovation is implemented each time a full cycle is completed. The frequency with which this implementation occurs is the inverse of the expected length of a complete cycle. This in turn is just the expected length of phase 1 plus the expected length of phase 2: $1/\lambda z + 1/\mu = [\mu + \lambda z] / \mu \lambda z$. Thus, the growth rate will be:

$$g = \ln \gamma \frac{\mu \lambda z}{\mu + \lambda z}$$

which is positively affected by anything that raises research. Note also that growth tapers off in the absence of the arrival of new GPTs, i.e. if $\mu = 0$. This leads [Gordon \(2012\)](#) to predict a durable slowdown of growth in the US and other developed economies as the ITC revolution is running out of steam.

The size of the slump $\ln(F(L)) - \ln(F(L - z))$ that occurs when each GPT arrives is also an increasing function of z , and hence it will tend to be positively correlated with the average growth rate.

One further property of this cycle worth mentioning is that the wage rate will rise when the economy goes into a slump. That is, because there is no research in phase 2, the normalized wage must be low enough to provide employment for all L workers in the manufacturing sector; whereas, with the arrival of the new GPT, the wage must rise to induce manufacturers to release workers into research. This brings us directly to the next subsection on wage inequality.

1.6.3 GPT and Wage Inequality

In this subsection we show how the model of the previous section can account for the rise in the skill premium during the IT revolution. We modify that model by assuming that there are two types of labor. Educated labor can work in both research and manufacturing, whereas uneducated labor can only work in manufacturing. Let L^s and L^u denote the supply of educated (skilled) and uneducated (unskilled) labor, let ω_1^s and ω_1^u denote their

respective productivity-adjusted wages in phase 1 of the cycle (when research activities on complementary inputs actually take place), and let ω_2 denote the productivity-adjusted wage of labor in phase 2 (when new GPTs have not yet appeared and therefore labor is entirely allocated to manufacturing).

If in equilibrium the labor market is segmented in phase 1, with all skilled labor being employed in research while unskilled workers are employed in manufacturing, we have the labor-market-clearing conditions:

$$L^s = z, \quad L^u = \gamma(\omega_1^u), \quad \text{and} \quad L^s + L^u = \gamma(\omega_2),$$

and the research-arbitrage condition:

$$\omega_1^s = \lambda \gamma v_2, \tag{1.30}$$

where v_2 is the productivity-adjusted value of an intermediate producer in stage 2. This value is itself determined as before by the two Bellman equations:

$$\rho v_2 = \tilde{\pi}(\omega_2) + \mu [v_1 - v_2],$$

and:

$$\rho v_1 = \tilde{\pi}(\omega_1^u) - \lambda z v_1.$$

Thus, the above research-arbitrage Equation (1.30) expresses the wage of skilled labor as being equal to the expected value of investing (skilled) labor in R&D for discovering complementary inputs to the new GPT.

The labor market will be truly segmented in phase 1, if and only if, ω_1^s defined by research-arbitrage condition (1.30) satisfies:

$$\omega_1^s > \omega_1^u,$$

which in turn requires that L^s not be too large. Otherwise the labor market remains unsegmented, with $z < L^s$ and:

$$\omega_1^s = \omega_1^u,$$

in equilibrium. In the former case, the arrival of a new GPT raises the skill premium (from 0 to $\omega_1^s/\omega_1^u - 1$) at the same time as it produces a productivity slowdown because labor is driven out of production.

1.6.4 Predictions

The above GPT model delivers the following predictions.²⁵

²⁵ While Jovanovic and Rousseau (2005) provide evidence for the first three predictions, we refer the reader to Acemoglu (2002, 2009), Aghion et al. (1999), and Aghion and Howitt (2009) for evidence on growth and wage inequality. In particular, Aghion and Howitt contrast the GPT explanation with alternative explanations based on trade, deunionization, or directed technical change considerations.

Prediction 1: *The diffusion of a new GPT is associated with an increase in the flow of firm entry and exit.*

This results from the fact that the GPT is Schumpeterian in nature; thus it generates quality-improving innovations, and therefore creative destruction, in any sector of the economy where it diffuses.

Prediction 2: *The arrival of a new GPT generates a slowdown in productivity growth; this slowdown is mirrored by a decline in stock-market prices.*

The diffusion of a new GPT requires complementary inputs and learning, which may draw resources from normal production activities and may contribute to future productivity in a way that cannot be captured easily by current statistical indicators. Another reason why the diffusion of a new GPT should reduce growth in the shortrun is by inducing the obsolescence of existing capital in the sectors it diffuses to (see [Aghion and Howitt, 1998, 2009](#)).

Prediction 3: *The diffusion of a new GPT generates an increase in wage inequality both between and within educational groups.*

An increase in the skill premium occurs as more skilled labor is required to diffuse a new GPT to the economy, as we saw above. The other and perhaps most intriguing feature of the upsurge in wage inequality is that it took place to a large extent within control groups, no matter how narrowly those groups are identified (e.g. in terms of experience, education, gender, industry, occupation). One explanation is that skill-biased technical change enhanced not only the demand for observed skills as described earlier but also the demand for unobserved skills or abilities. Although theoretically appealing, this explanation is at odds with econometric work ([Blundell and Preston, 1999](#)) showing that the within-group component of wage inequality in the United States and the United Kingdom is mainly transitory, whereas the between-group component accounts for most of the observed increase in the variance of permanent income. The explanation based on unobserved innate abilities also fails to explain why the rise in within-group inequality has been accompanied by a corresponding rise in individual wage instability (see [Gottschalk and Moffitt, 1994](#)). Using a GPT approach, [Aghion et al. \(2002\)](#) argue that the diffusion of a new technological paradigm can affect the evolution of within-group wage inequality in a way that is consistent with these facts. The diffusion of a new GPT raises within-group wage inequality primarily because the rise in the speed of embodied technical progress associated with the diffusion of the new GPT increases the market premium to those workers who adapt quickly to the leading-edge technology and are therefore able to survive the process of creative destruction at work as the GPT diffuses to the various sectors of the economy.²⁶

²⁶ In terms of the preceding model, let us again assume that all workers have the same level of education but that once a new GPT has been discovered, only a fraction α of the total labor force can adapt quickly enough to the new technology so that they can work on looking for a new component that complements the GPT. The other workers, who did not successfully adapt have no alternative but to work in



1.7. CONCLUSION

In this paper, we argued that Schumpeterian growth theory—where current innovators exert positive knowledge spillovers on subsequent innovators as in other innovation-based models, but where current innovators also drive out previous technologies—generates predictions and explains facts about the growth process that could not be accounted for by other theories.

In particular, we saw how Schumpeterian growth theory manages to put IO into growth and to link growth with firm dynamics, thereby generating predictions on the dynamic patterns of markets and firms (entry, exit, reallocation, etc.) and on how these patterns shape the overall growth process. These predictions and the underlying models can be confronted with micro data and this confrontation in turn helps refine the models. This back-and-forth communication between theory and data has been key to the development of the Schumpeterian growth theory over the past 25 years.²⁷

Also, we argued that Schumpeterian growth theory helps us reconcile growth with development, in particular, by bringing out the notion of appropriate growth institutions and policies, i.e. the idea that what drives growth in a sector (or country) far below the world technology frontier is not necessarily what drives growth in a sector or country at the technological frontier where creative destruction plays a more important role. In particular, we pointed to democracy being more growth enhancing in more frontier economies. The combination of the creative destruction and appropriate growth institutions ideas also underlies the view²⁸ that “extractive economies,” where creative destruction is deterred by political elites, are more likely to fall into low-growth traps.

manufacturing. Let ω_1^{adapt} denote the productivity-adjusted wage rate of adaptable workers in phase 1 of the cycle, and let ω_1 denote the wage of non-adaptable workers. Labor market clearing implies: $\alpha L = z$; $[1 - \alpha] L = \gamma(\omega_1)$; $L = \gamma(\omega_2)$, whereas research arbitrage for adaptable workers in phase 1 implies $\omega_1^{adapt} = \lambda \gamma \nu_2$. When α is sufficiently small, the model generates a positive adaptability premium: $\omega_1^{adapt} > \omega_1$.

²⁷ For example, when analyzing the relationship between growth and firm dynamics, this back-and-forth process amounts to what one might call a layered approach. Here, we refer the reader to Daron Acemoglu’s panel discussion at the Nobel Symposium on Growth and Development (September 2012). The idea here is that of a step-by-step estimation method, where at each step a small subset of parameters are being identified in their neighborhood. Thanks to the rich set of available micro data, one can first identify a parameter and its partial equilibrium effect as well as some of its industry equilibrium effects. Next, one can test the predictions of the model using moments in the data that were not directly targeted in the original estimation. Then one can check that the model also satisfies various out-of-sample properties and reach a macro-aggregation by building on detailed micro moments. Schumpeterian models are well suited for this type of approach as they are able to generate realistic firm dynamics with tractable aggregations.

²⁸ See [Acemoglu and Robinson \(2012\)](#).

Beyond enhancing our understanding of the growth process, Schumpeterian growth theory is useful in at least two respects. First, as a tool for the design of growth policy: departing from the Washington consensus view whereby the same policies should be recommended everywhere, the theory points to appropriate growth policies, i.e. policies that match the particular context of a country or region. Thus, we saw that more intense competition (lower entry barriers), a higher degree of trade openness, and more emphasis on research education are all more growth-enhancing in more frontier countries.²⁹

The Schumpeterian growth paradigm also helps us assess the relative magnitude of the counteracting partial equilibrium effects pointed out by the theoretical IO literature. For example, there is a whole literature on competition, investments, and incentives³⁰ that points to counteracting partial equilibrium effects without saying much about when one particular effect should be expected to prevail. In contrast, [Section 1.3](#) illustrated how aggregation and the resulting composition effect could help determine under which circumstances the escape–competition effect would dominate the counteracting Schumpeterian effect. Similarly, [Section 1.4](#) showed the importance of reallocation for growth; thus, policies supporting entry or incumbent R&D could contribute positively to economic growth in partial equilibrium, yet in general equilibrium [Section 1.4](#) showed that this is done at the expense of reduced innovation by the rest of the economy.

Where do we see the Schumpeterian growth agenda being pushed over the next years? A first direction is to look more closely at how growth and innovation are affected by the organization of firms and research. Thus, over the past 5 years Nick Bloom and John Van Reenen have popularized fascinating new datasets that allow us to look at how various types of organizations (e.g. more or less decentralized firms) are more or less conducive to innovation. But firms' size and organization are in turn endogenous, and in particular, they depend on factors such as the relative supply of skilled labor or the nature of domestic institutions. Future studies should try to model and then test the relationship from skill endowment and the institutional environment to firm organization and then from firm organization to innovation and growth.

A second and related avenue for future research is to look at growth, firm dynamics, and reallocation in developing economies. Recent empirical evidence (see [Hsieh and Klenow, 2009, 2012](#)) has shown that the misallocation of resources is a major source of the productivity gap across countries. What are the causes of misallocation, and why do these countries lack creative destruction that would eliminate the inefficient firms? Schumpeterian theory with firm dynamics could be an invaluable source to shed light on these important issues that lie at the core of the development puzzle.

A third avenue is to look at the role of finance in the growth process. In [Section 1.5](#) we pointed to equity finance being more growth-enhancing in more frontier economies.

²⁹ Parallel studies point to labor market liberalization and stock-market finance being more growth-enhancing in more advanced countries or regions.

³⁰ See the recent analytical surveys by [Gilbert \(2006\)](#), [Vives \(2008\)](#), and [Schmutzler \(2010\)](#).

More generally, we still need to better understand how different types of financial instruments map with different sources of growth and different types of innovation activities. Also, we need to better understand why we observe a surge of finance during the acceleration phase in the diffusion of new technological waves, as mentioned in [Section 1.6](#), and how financial sectors evolve when the waves taper off. These and many other microeconomic aspects of innovation and growth await further research.

ACKNOWLEDGMENT

This survey builds on a presentation at the Nobel Symposium on Growth and Development (September 2012) and was subsequently presented at the Schumpeter Lecture at the Swedish Entrepreneurship Forum (January 2013). We thank Pontus Braunerhjelm, Mathias Dewatripont, Sam Kortum, Michael Spence, John Van Reenen, David Warsh, and Fabrizio Zilibotti for helpful comments and encouragements, and Sina Ates, Salome Baslandze, Jonathan Roth, and Felipe Saffie for outstanding editing work. This paper is produced as part of the project Growth and Sustainability Policies for Europe (GRASP), a collaborative project funded by the European Commission's Seventh Research Framework Programme, Contract No. 244725.

REFERENCES

- Acemoglu, D., 2002. Technical change, inequality, and the labor market. *Journal of Economic Literature* 40, 7–72.
- Acemoglu, D., 2009. *Introduction to Modern Economic Growth*. Princeton University Press.
- Acemoglu, D., Aghion, P., Zilibotti, F., 2006. Distance to Frontier, Selection, and economic growth. *Journal of the European Economic Association*, 37–74.
- Acemoglu, D., Akcigit, U., 2012. Intellectual property rights policy, competition and innovation. *Journal of the European Economic Association* 10, 1–42.
- Acemoglu, D., Akcigit, U., Bloom, N., Kerr, W., 2013. Innovation, Reallocation and Growth. NBER Working Paper 18993.
- Acemoglu, D., Akcigit, U., Hanley, D., Kerr, W., 2012. Transition to Clean Technology. University of Pennsylvania, mimeo.
- Acemoglu, D., Cao, D., 2011. Innovation by Entrants and Incumbents. MIT, mimeo.
- Acemoglu, D., Robinson, J., 2006. Economic backwardness in political perspective. *American Political Science Review* 100, 115–131.
- Acemoglu, D., Robinson, J., 2012. *Why Nations Fail*, Crown Business.
- Acs, Z., Audretsch, D., 1988. Innovation in large and small firms: an empirical analysis. *American Economic Review* 78, 678–690.
- Acs, Z., Audretsch, D., 1991. Innovation and size at the firm level. *Southern Economic Journal* 57, 739–744.
- Aghion, P., Alesina, A., Trebbi, F., 2007. Democracy, technology and growth. In: Helpman, E. (Eds.), *Institutions and Economic Performance*, Cambridge University Press.
- Aghion, P., Bloom, N., Blundell, R., Griffith, R., Howitt, P., 2005. Competition and innovation: an inverted-U relationship. *Quarterly Journal of Economics* 120, 701–728.
- Aghion, P., Blundell, R., Griffith, R., Howitt, P., Prantl, S., 2009. The effects of entry on incumbent innovation and productivity. *Review of Economics and Statistics* 91, 20–32.
- Aghion, P., Boustan, L., Hoxby, C., Vandenbussche, J., 2009. Exploiting States' Mistakes to Identify the Causal Effects of Higher Education on Growth. Harvard, mimeo.
- Aghion, P., Caroli, E., Garcia-Penalosa, C., 1999. Inequality and economic growth: the perspective of the new growth theories. *Journal of Economic Literature* 37, 1615–1660.
- Aghion, P., Griffith, R., 2006. *Competition and Growth: Reconciling Theory and Evidence*, MIT Press.
- Aghion, P., Harris, C., Howitt, P., Vickers, J., 2001. Competition, imitation and growth with step-by-step innovation. *Review of Economic Studies* 68, 467–492.

- Aghion, P., Harris, C., Vickers, J., 1997. Competition and growth with step-by-step innovation: an example. *European Economic Review, Papers and Proceedings*, 771–782.
- Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. *Econometrica* 60, 323–351.
- Aghion, P., Howitt, P., 1998. *Endogenous Growth Theory*, MIT Press.
- Aghion, P., Howitt, P., 2009. *The Economics of Growth*, MIT Press.
- Aghion, P., Howitt, P., Prantl, S., 2013. Patent Rights, Product Market Reforms and Innovation. NBER Working Paper 18854.
- Aghion, P., Howitt, P., Violante, G., 2002. General purpose technology and wage inequality. *Journal of Economic Growth* 7, 315–345.
- Akcigit, U., Hanley, D., Serrano-Velarde, N., 2012. Back to Basics: Basic Research Spillovers, Innovation Policy and Growth. University of Pennsylvania, mimeo.
- Akcigit, U., Kerr, W., 2010. Growth Through Heterogeneous Innovations. NBER Working Paper 16443.
- Alesina, A., Rodrik, D., 1994. Distributive politics and economic growth. *Quarterly Journal of Economics* 109, 465–490.
- Arrow, K., 1962. The economic implications of learning by doing. *Review of Economic Studies* 29, 155–173.
- Axtell, R., 2001. Zipf distribution of US firm sizes. *Science* 293, 1818–1820.
- Barro, R.J., Sala-i-Martin, X., 2003. *Economic Growth*. McGraw-Hill.
- Bartelsman, E., Doms, M., 2000. Understanding productivity: lessons from longitudinal microdata. *Journal of Economic Literature* 38, 569–594.
- Bartelsman, E., Haltiwanger, J., Scarpetta, S., 2009. Cross-country Differences in Productivity: The Role of Allocation and Selection NBER. Working Paper 15490.
- Birch, D., 1981. Who Creates Jobs? The Public Interest 65, 3–14.
- Birch, D., 1987. *Job Creation in America: How Our Smallest Companies Put the Most People to Work*. Free Press.
- Blundell, R., Griffith, R., Van Reenen, J., 1995. Dynamic Count Data Models of Technological Innovation. *Economic Journal* 105, 333–344.
- Blundell, R., Griffith, R., Van Reenen, J., 1999. Market share, market value and innovation in a panel of British manufacturing firms. *Review of Economic Studies* 66, 529–554.
- Blundell, R., Preston, I., 1999. Inequality and Uncertainty: Short-Run Uncertainty and Permanent Inequality in the US and Britain. University College London, mimeo.
- Boldrin, M., Levine, D., 2008. *Against Intellectual Monopoly*. Cambridge University Press.
- Bresnahan, T., Trajtenberg, M., 1995. General purpose technologies: engines of growth? *Journal of Econometrics*, 65, 83–108.
- Cohen, W., 1995. Empirical studies of innovative activity. In: Stoneman, Paul (Ed.), *Handbook of the Economics of Innovations and Technological Change*, Blackwell.
- Cohen, W., Klepper, S., 1996. Firm size and the nature of innovation within industries: the case of process and product R&D. *Review of Economics and Statistics*, 232–243.
- Corriveau, L., 1991. *Entrepreneurs, Growth, and Cycles*. PhD Dissertation, University of Western, Ontario.
- David, P., 1990. The dynamo and the computer: an historical perspective on the modern productivity paradox. *American Economic Review* 80, 355–361.
- Davis, S., Haltiwanger, J., Schuh, S., 1996. *Job Creation and Destruction*, MIT Press.
- De Loecker, J., Goldberg P., Khandelwal, A., Pavcnik N., 2012. Prices, Markups and Trade Reform. Yale, Mimeo.
- Djankov, S., La Porta, R., Lopez-de-Silanes, F., Shleifer, A., 2002. The regulation of entry. *Quarterly Journal of Economics* 117, 1–37.
- Foster, L., Haltiwanger, J., Krizan, C., 2001. Aggregate productivity growth: lessons from microeconomic evidence. In: Dean, E., Harper, M., Hulten, C. (Eds.), *New Directions in Productivity Analysis*. University of Chicago Press.
- Foster, L., Haltiwanger, J., Krizan, C., 2006. Market selection, reallocation, and restructuring in the U.S. retail trade sector in the 1990s. *Review of Economics and Statistics* 88, 748–758.
- Frankel, J., Romer, D., 1999. Does trade cause growth? *American Economic Review* 89, 379–399.

- Freeman, C., 1982. *The Economics of Industrial Innovation*. MIT Press.
- Galor, O., 2011. *Unified Growth Theory*. Princeton University Press.
- Gerschenkron, A., 1962. *Economic Backwardness in Historical Perspective: A Book of Essays*. Belknap Press of Harvard University Press.
- Gilbert, R., 2006. Looking for Mr Schumpeter: Where Are We in the Competition-Innovation Debate? In: Lerner, J., Stern, S. (Eds.), *Innovation Policy and Economy*. NBER, MIT Press.
- Goldberg, P., Khandelwal, A., Pavcnik, N., Topalova, P., 2010. Imported intermediate inputs and domestic product growth: evidence from India. *Quarterly Journal of Economics*, 125, 1727–1767.
- Gordon, R., 2012. Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds. NBER Working Paper 18315.
- Gottschalk, P., Moffitt, R., 1994. The growth of earnings instability in the US labor market. *Brookings Papers on Economic Activity* 2, 217–272.
- Greenwood, J., Yorukoglu, M., 1974. Carnegie-Rochester Series on Public Policy 46, 49–95.
- Griliches, Z., 1990. Patent statistics as economic indicators: a survey. *Journal of Economic Literature* 28, 1661–1707.
- Grossman, G., Helpman, E., 1991. Quality ladders in the theory of growth. *Review of Economic Studies* 58, 43–61.
- Hall, B., Jaffe, A., Trajtenberg, M., 2001. The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. NBER Working Paper 8498.
- Haltiwanger, J., Jarmin, R., Miranda, J., 2010. Who Creates Jobs? Small vs. Large vs. Young. NBER Working Paper 16300.
- Helpman, E., Trajtenberg, M., 1998. A time to sow and a time to reap: growth based on general purpose technologies. In: Helpman, E. (Eds.), *General Purpose Technologies and Economic Growth*, MIT Press.
- Hornstein, A., Krusell, P., 1996. Can Technology Improvements Cause Productivity Slowdowns? *NBER Macroeconomics Annual*, 11, 209–259.
- Howitt, P., 2000. Endogenous growth and cross-country income differences. *American Economic Review* 90, 829–46.
- Hsieh, C., Klenow, P., 2009. Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics* 124, 1403–1448.
- Hsieh, C., Klenow, P., 2012. The Life Cycle of Plants in India and Mexico. NBER Working Paper 18133.
- Ijiri, Y., Simon, H., 1977. *Skew Distributions and the Sizes of Business Firms*. North-Holland.
- Jones, C., Vollrath, D., 2013. *Introduction to Economic Growth*. W. W. Norton & Company.
- Jovanovic, B., Rousseau, P., 2005. General purpose technologies. In: Aghion, P., Durlauf, S. (Eds.), *Handbook of Economic Growth*. Elsevier, North-Holland.
- Klette, T., Kortum, S., 2004. Innovating firms and aggregate innovation. *Journal of Political Economy* 112, 986–1018.
- Lentz, R., Mortensen, D., 2008. An empirical model of growth through product innovation. *Econometrica* 76, 1317–1373.
- Lipsey, R., Bekar, C., 1995. A structuralist view of technical change and economic growth. *Bell Canada Papers on Economic and Public Policy* 3, 9–75.
- Nelson, R., Phelps, E., 1966. Investment in humans, technological diffusion, and economic growth. *American Economic Review* 61, 69–75.
- Neumark, D., Wall, B., Zhang, J., 2008. Do Small Businesses Create More Jobs? New Evidence for the United States from the National Establishment Time Series. NBER Working Paper 13818.
- Nickell, S., 1996. Competition and corporate performance. *Journal of Political Economy*, 104, 724–746.
- Pennings, J.M., Buitendam, A., 1987. *New Technology As Organizational Innovation: The Development and Diffusion of Microelectronics*. Bollinger.
- Peretto, P., 1998. Technological change, market rivalry, and the evolution of the capitalist engine of growth. *Journal of Economic Growth* 3, 53–80.
- Persson, T., Tabellini, G., 1994. Is inequality harmful for growth? *American Economic Review* 84, 600–621.
- Qian, Y., 2007. Do national patent laws stimulate domestic innovation in a global patenting environment? *Review of Economics and Statistics* 89, 436–453.

- Romer, P., 1990. Endogenous technical change. *Journal of Political Economy* 98, 71–102.
- Rossi-Hansberg, E., Wright, M., 2007. Establishment size dynamics in the aggregate economy. *American Economic Review* 97, 1639–1666.
- Scherer, F.M., 1984. *Innovation and Growth: Schumpeterian Perspectives*. MIT Press.
- Schmalensee, R., 1989. Inter-industry studies of structure and performance. In: Schmalensee, R., Willig, R.D. (Eds.), *Handbook of Industrial Organization*, vol. 2. North-Holland.
- Schmutzler, A., 2010. Is competition good for innovation? a simple approach to an unresolved question. *Foundations and Trends in Microeconomic Analysis* 5, 355–428.
- Segerstrom, P., Anant, T., Dinopoulos, E., 1990. A schumpeterian model of the product cycle. *American Economic Review* 88, 1077–1092.
- Simon, H., Bonini, C., 1958. The size distribution of business firms. *American Economic Review* 48, 607–17.
- Sivadasan, J., 2009. Barriers to competition and productivity: evidence from India. *B.E. Journal of Economic Analysis and Policy* 9, 42.
- Stanley, M., Buldyrev, S., Havlin, S., Mantegna, R., Salinger, M., Stanley, E., 1995. Zipf plots and the size distribution of firms. *Economic Letters* 49, 453–57.
- Syversen, C., 2011. What determines productivity. *Journal of Economic Literature* 49, 326–365.
- Tirole, J., 1988. *The Theory of Industrial Organization*. MIT Press.
- Topalova, P., Khandelwal, A.K., 2011. Trade liberalization and firm productivity: the case of India. *Review of Economics and Statistics* 93, 995–1009.
- Tushman, M.L., Anderson, P., 1986. Technological discontinuities and organizational environments. *Administrative Science Quarterly* 31, 439–465.
- Vandenbussche, J., Aghion, P., Meghir, C., 2006. Growth, distance to frontier, and composition of human capital. *Journal of Economic Growth* 11, 97–127.
- Vives, X., 2008. Innovation and competitive pressure. *Journal of Industrial Economics* 56, 419–469.
- Weil, D., 2012. *Economic Growth*. Prentice Hall.