Innovating Firms: Evidence and Theory

Tor Jakob Klette* and Samuel Kortum†

January 11, 2001

Preliminary and incomplete draft.

Abstract: We develop a model of innovating firms that is rich enough to confront the detailed evidence yet parsimonious enough to be analytically tractable. The model describes the dynamic behavior of individual heterogenous firms, captures industry behavior through distributions of outcomes, and aggregates into a general equilibrium model of endogenous growth. Along the way it gives new interpretations to the empirical evidence. According to the model the widely studied empirical relationship between R&D and productivity contains little information about the impact of R&D on innovation. The model explains why R&D effort is strongly related to productivity and patenting yet essentially unrelated to firm size or growth. In the end, the model demonstrates that the firm-level evidence, while seemingly puzzling, supports the view that R&D is crucial for aggregate growth.

JEL classification: xx
Keywords: xx

1 Introduction

Recent theoretical work has emphasized technological innovations generated by R&D investment as the primary engine of growth. Aggregate comparisons of growth, innovation, and R&D lend support to these theories, but the evidence is more suggestive than conclusive.[1] To obtain a sharper understanding of R&D's role may require exploiting the detailed observations available at the firm level.

Firm level evidence, however, does not provide as clean a picture as one might have hoped. First, a number of firms seem to survive and prosper in high-tech industries while reporting little or no R&D.

*Department of Economics, University of Oslo, and CEPR. *E-mail:* t.j.klette@econ.uio.no, http://www.uio.no/~torjk/, *Tel:* +47 22 85 51 34. *Fax:* +47 22 85 50 35.
†Department of Economics, Boston University, and NBER. *E-mail:* kortum@bu.edu .

[1] On the theoretical side we have in mind Aghion and Howitt (1992), Grossman and Helpman (1991), and Romer (1990) while on the empirical side Cabbalero and Jaffe (1993), Coe and Helpman (1995), Kortum (1997), Eaton and Kortum (1999), and Howitt (2000).

Second, there is little correlation between R&D intensity and firm size or growth. This second fact is particularly puzzling in light of the strong correlation across firms between R&D effort and both firm productivity levels and patenting activity. Adding to the problem is that the evidence on R&D, industry characteristics, and firm level performance is scattered across separate lines of research. There has been no coherent summary of the causal relationships between R&D, firm size, patenting and productivity. The available theories on R&D, innovation, and growth are of little help in making sense of these firm-level observations since they typically do not contain a theory of the firm. In most such models, there is a one-to-one correspondence between products and firms, with each new good introduced by a new firms. This tight relationship between a firm and a product is hard to reconcile with firm-level data.

To address these problems, this paper presents a theoretical framework for firm-level analysis. We use evidence from a variety of sources to motivate our model of innovating firms. The evidence includes statistics on research effort, innovation and productivity, and firm growth. These three sets of empirical observations, that previously have been considered separately, are given a unified interpretation with our model. The model is rich enough to confront this evidence yet parsimonious enough to be analytically tractable. It describes the dynamic behavior of individual heterogenous firms, captures industry behavior through distributions of outcomes, and aggregates into a general equilibrium model of endogenous growth.

Our model builds on three theoretical literatures that have developed along surprisingly independent paths. The first is models of firm dynamics of the sort put forward by Penrose (1959) and Ijiri and Simon (1977) and recently revived by Sutton (1998) and Amaral et al. (1998). Although these models are able to match stylized facts about firm growth and size distributions, they are criticized for ascribing overly mechanical behavior to firms. The second is reduced form models used to motivate firm-level econometric studies of R&D, patenting, and productivity, as summarized in Griliches (1990) and Griliches (1995). While much has been learned from this work about the relationship between innovative outputs and R&D effort, a shortcoming is that the analytical framework does not yield a satisfactory explanation of what drives the variation across firms in R&D effort. The third is the models of endogenous technological change developed recently by Aghion and Howitt (1992), Grossman and Helpman (1991), and Romer (1990). This work gets at the determinants of aggregate innovative effort, but it lacks a compelling theory of firms as noticed above. As a result it is difficult to relate the parameters of these models to the vast amount of empirical work at the firm level.

While each of these three approaches has its problems, we believe that by borrowing what is best from each we can address the shortcomings. From the first literature we borrow the idea that firms grow by taking over new markets. From the second we borrow the idea of the firm-level stock of knowledge and the innovation production function. From the third we borrow the idea of an economy characterized by differentiated products with monopolistically competitive firms. The result is a model of research and exploration in the spirit of Ericson and Pakes (1995), but with no strategic interaction between firms[2]. In fact, the only interaction between firms in our model is that while each one takes the hazard of losing a market as given, in equilibrium this hazard is the result of their combined innovative effort, as in Grossman and Helpman (1991). While the lack of strategic interaction limits the applicability of our model in industries with a small number of large firms, it does enable us to go quite far in explicitly characterizing firm dynamics[3].

The next section surveys empirical studies which have shown considerable variation across firms within narrowly defined industries almost in every dimension including sales, productivity, profit margins, patenting, and R&D investment (see e.g. Cohen (1995) and Geroski (1998)). Much of this heterogeneity is quite persistent over a number of years. The section summarizes the available evidence in a series of stylized facts.

These stylized facts form the target for our theoretical model, which is developed over Sections 3-7. Section 3 introduces our model of firm-level innovation while Section 4 derives the implications for firm dynamics. Section 5 moves to the industry level, describing the entry condition and deriving the size distribution of firms. Section 6 extends the model to incorporate firm-level heterogeneity in R&D intensity. Section 7 returns to the set of stylized facts and reviews them in light of the theory. Section 8 moves to the aggregate level and solves the model in general equilibrium. Section 9 concludes.

2 Evidence on Innovating Firms

This section presents a comprehensive list of empirical regularities or facts which have emerged from a large number of studies of firm-level data. The theoretical framework presented in the subsequent sections is aimed at providing a coherent interpretation of these facts.

We have listed the facts that are robust and economically significant. Following Cohen and Levin (1989) and Schmalensee (1989), we have ignored relationships (even when statistically significant) if they seem obscure.

2.1 R&D, Productivity, and Patents

Fact # 1 *There is a positive relationship between measured productivity and various measures of R&D activity across firms. Firm-level productivity growth is not strongly related to measures of R&D activity.*

There is a vast literature verifying a positive and statistically significant relationship between measured productivity and R&D activity at the firm level; see e.g. Griliches (1995, 1998 ch. 12) and Hall

[2] Two other related studies are Hopenhayn (1992) and Mitchell (2000). The first study treats innovations as exogenous processes, while the second study emphasizes costless learning rather than R&D-investment and innovation as the dynamic force in firm growth.

[3] We are not the first to try and bring realistic features of firms into models of endogenous technological change. Previous efforts, such as Thompson (1996), Peretto (1998), and Klette and Griliches (2000) have all assumed a firm produces a single product. Our approach here is more like earlier models of firm dynamics, however, in which firms grow by acquiring new products. A shortcoming of our approach is that we cannot capture the fact that firms often improve on their own products.

(1996). This positive relationship has been consistently verified in a number of studies focusing on cross-sectional differences across firms. The longitudinal relationship between firm-level differences in R&D and productivity *growth,* which control for permanent differences across firms, has turned out to be fragile and typically insignificant statistically.

Fact # 2 *There is approximately a constant returns relationship between patents and R&D in the cross-sectional dimension, while the longitudinal dimension suggests diminishing returns.*

The relationship between innovation, patents, and R&D has been surveyed by Griliches (1990). He emphasizes that there is quite a strong relationship between R&D and the number of patents received across firms, with a median R-square around 0.9 . For larger firms the patents-R&D relationship is close to linear while there is a reasonably large number of smaller firms that exhibit significant patenting while reporting very little R&D. That is to say, small firms appear to be more efficient, receiving a larger number of patents per R&D dollar. Cohen and Klepper (1996) emphasize this high patent-R&D ratio among the small firms, and interpret it as evidence for smaller firms being more innovative.

Griliches, on the other hand, argues: "the appearance of diminishing returns at the cross-sectional level is due, I think, primarily to two effects: selectivity and the differential role of formal R&D and patents for small and large firms (pg. 1675)." There is a selectivity bias since small firms in available samples are not representative but are typically more innovative than the average small firm. Furthermore, small firms are likely to be doing relatively more informal R&D while reporting less of it and hence providing the *appearance* of more patents per R&D dollar.[4] Hence, Griliches suggests that in terms of patents per R&D dollar, there is little evidence for diminishing returns in the cross-sectional dimension.

There is also a robust patents-R&D relationship in the longitudinal (within-firm across-time) dimension: "the evidence is quite strong that when a firm changes its R&D expenditures, parallel changes occur also in its patent numbers [Griliches (1990), pg. 1674]." A *diminishing returns* relationship between patents and R&D is more pronounced in the longitudinal dimension than in the cross section. Referring to Hall et al. (1986) and other studies, Griliches suggests that the patent elasticity of R&D is between 0.3 and 0.6. Using recent econometric techniques on the same sample as Hall et al., Blundell et al. (1999) report somewhat higher estimates, in the range 0.6 to 0.9.

2.2 R&D Investment

Fact # 3 *R&D intensity is independent of firm size.*

The large literature relating R&D expenditures to firm size is surveyed by Cohen (1995) and Cohen and Klepper (1996). Cohen and Klepper state that among firms doing R&D: "in most industries it has not been possible to reject the null hypothesis that R&D varies proportionately with size across the

[4] This effect has been shown to be operative by Kleinknecht (1987) and others.

entire firm size distribution (pg. 929)." On the other hand, they also point out: "The likelihood of a firm reporting positive R&D effort rises with firm size, and approaches one for firms in the largest size ranges (pg. 928)." While the first statement supports Fact # 3, the second seems to contradict it.

As pointed out above, Griliches (1990) interprets the appearance of less R&D among small firms as an artifact of the available data rather than a reflection of differences in real innovative activity between large and small firms. That is to say, the higher fraction of small firms reporting no *formal* R&D is offset by small firms doing more *informal* R&D. Furthermore, smaller firms tend to have a lower absolute level of R&D, and R&D surveys often have a reporting threshold related to the absolute level of R&D. Similarly, the innovative activity being singled out in a firm's accounts as formal R&D is related to the absolute level of R&D.

Fact # 4 *The R&D intensity distribution is highly skewed, and a considerable fraction of firms report zero R&D even in high-tech industries.*

A number of studies have reported substantial variation in R&D intensities across firms within the same industry; see Cohen (1995). Cohen and Klepper (1992) show that the R&D intensity distribution exhibit a regular pattern across industries, in accordance with Fact # 4. The R&D intensity distributions they present are all unimodal, positively skewed with a long tail to the right, and with a large number of R&D non-performers. Klette and Johansen (1998) report the same pattern of a unimodal and skewed R&D intensity distribution based on a sample of Norwegian firms.

Fact # 5 *Differences in R&D intensity across firms are highly persistent.*

Scott (1984) shows that in a large longitudinal sample of US firms about 50 percent of the variance in business unit R&D intensity is accounted for by firm fixed effects. Klette and Johansen (1998), considering a panel of Norwegian firms in high-tech industries, confirm that differences in R&D intensity are highly persistent over a number of years, and that R&D investment is far more persistent than investment in physical capital.

Fact # 6 *Firms' R&D-investments are close to geometric random walks.*

In a study of US manufacturing firms, Hall et al. (1986) conclude by describing "R&D investment [in logs] within a firm as essentially a random walk with an error variance which is small (about 1.5 percent) relative to the total variance of R&D expenditures between firms (pg. 281)." Similarly, Klette and Griliches (2000) report zero correlation between changes in log R&D and the level of R&D for Norwegian firms.

2.3 Entry, Exit, Growth, and the Size Distribution of Firms

Fact # 7 *The size distribution of firms is highly skewed.*

This fact has been recognized for a long time, see Ijiri and Simon (1977) and Schmalensee (1989)[5]. According to Audretsch (1995), "virtually no other economic phenomenon has persisted as consistently as the skewed asymmetric firm-size distribution. Not only is it almost identical across every manufacturing industry, but it has remained strikingly constant over time (at least since the Second World War) and even across developed industrialized nations (pg. 65)."

Fact # 8 *Growth rates are unrelated to past growth or to size, at least for larger firms. Smaller firms have a lower probability of survival, but those that survive grow proportionally faster than larger firms.*

Fact # 8 has emerged from a number of empirical studies from the last 10-15 years as a refinement of Gibrat's law, which states that firm sizes and growth rates are uncorrelated[6]. Our statement corresponds to the summaries of the literature on Gibrat's law by Geroski (1998) and Sutton (1998).

Fact # 9 *The variance of growth rates is higher for smaller firms.*

This pattern has been recognized in a large number of studies discussed in Sutton (1997) and Caves (1998).

Fact # 10 *New firms are more likely to exit, but the hazard rate declines with age. An entering cohort's market share tends to decline slowly as it ages.*

Caves (1998) reviews the empirical literature on patterns among new entrant firms.

After we present our model in the next sections, we will return to these empirical facts and interpret them in light of the theoretical framework.

3 The Model of an Innovating Firm

This section and the four that follow it present our theoretical model. We start here by describing the innovation process for an individual firm. In the following sections we derive firm dynamics, introduce entry and derive aggregate dynamics, introduce heterogeneity in firms' innovative capabilities, and connect the model to endogenous growth theory.

We assume that a firm operates in a single industry or aggregate economy described by a continuum of distinct markets $j \in [0, 1]$. We interpret these markets as differentiated goods as in theories of monopolistic competition. In what follows we refer to a market as a good.

[5] A recent contribution to the empirical literature is Stanley et al. (1995).

[6] Theory and evidence related to Gibrat's law have recently been surveyed by Sutton (1997) and Caves (1998).

3.1 The Firm

We assume that each good is produced by a single firm, generates a unit flow of revenues and yields a profit flow $0 < \bar{\pi} < 1$. A firm is defined by its collection of goods. Since goods are anonymous, we only need to keep track of the number of distinct goods $n = 1, 2, 3, \ldots$ that a firm produces. The size of a firm, n, is a suitable state variable since the firm's revenues are simply n and its profits are $n\bar{\pi}$.

To add new goods a firm invests in innovative activity such as R&D. If successful, such investment leads to a product or process innovation for a particular good. With this innovation, the firm can successfully compete against the incumbent producer of the good so that the incumbent loses the product to the innovating firm. Expenditures on R&D could yield an innovation relevant to any product with equal probability, i.e. the actual product is drawn from the Uniform density on $[0, 1]$. Since each firm is infinitesimal relative to the continuum of products, it will never improve on a good it is currently producing.

The incumbent firm loses a product when some other firm comes up with an innovation related to one of the incumbent's goods. The Poisson hazard rate per product is μ. This parameter, which we refer to as the intensity of creative destruction, is taken as given by each firm, and is assumed to be constant over time. In later sections we will endogenize μ.

To summarize, a firm can be described in terms of the number of goods n that it produces. A size n firm receives a flow of profits $\bar{\pi}n$. It loses goods at a Poisson rate μn. It acquires new goods at a Poisson rate I which depends on its R&D investment and its innovation capability, which we now consider.

3.2 The Innovation Technology

We assume that a firm's Poisson innovation rate I depends on both its investment in R&D, denoted by R, and its knowledge capital. Differences between firms in their available knowledge capital is assumed to be summarized by n. Notice that n counts the firm's past innovations that have not yet been superseded by new innovations[7]. The knowledge capital is useful for the firm as it tries to come up with new innovations[8].

The innovation production function is,

$$I = F(R, n). \tag{1}$$

We place the following restrictions on F: (i) strictly increasing in R, (ii) strictly concave in R, (iii) strictly increasing in n, (iv) homogeneous of degree one in I and n. The second restriction captures the idea of decreasing returns to expanding research effort, and allows us to tie down the research investment of an

[7] n may be a useful measure of the firm's available knowledge capital not only because the knowledge embodied in the innovations themselves are useful in making new innovations. n may also reflect differences in the quality of firms' laboratories, which may evolve over time as indicated by changes in n.

[8] That R&D activities introduce scope economies in the development of related products has also been emphasized by Jovanovic (1993) in a static model of firm formation.

individual firm. The third embodies the idea that firms' knowledge capital facilitates innovation. The last restriction neutralizes the effect of firm size on the innovation process.[9]

For working out the firm's optimal R&D policy, it is useful to rewrite the innovation production function in the form of a cost function. Given the assumption made above, we get

$$R = C(I,n) = nc(I/n), \tag{2}$$

where C inherits the homogeneity property of F. Exploiting homogeneity, the intensive form of the cost function $c(x)$ is increasing and convex in x. To avoid a number of uninteresting qualifications in the analysis below we assume that $c(0) = 0$ and that $c(x)$ is differentiable on $x \geq 0$.

Our model of firm growth is consistent with arguments by Penrose (1959), that there are no constraints on firm size, but there are constraints on firm growth.[10] Interestingly, as in the model here, Penrose argues that firms grow through diversification into new markets.

3.3 The Innovation Decision

A firm with $n \geq 1$ products receives a flow of profits $\bar{\pi}n$ and faces a Poisson hazard μn of becoming a firm of size $n-1$. By spending on research it influences the Poisson hazard I of becoming a firm of size $n+1$. The Bellman equation for the value $V(n)$ of a firm with $n \geq 1$ products is

$$rV(n) = \max_{I} \left\{ \bar{\pi}n - C(I,n) + I\left[V(n+1) - V(n)\right] - \mu n \left[V(n) - V(n-1)\right] \right\}, \tag{3}$$

where r is the firm's discount rate. As shown in the appendix, the solution to the Bellman equation is,

$$\begin{aligned} V(n) &= vn \\ I(n) &= \lambda n, \end{aligned}$$

where, for the case of $\lambda > 0$, v and λ solve,

$$\begin{aligned} c'(\lambda) &= v \\ (r + \mu - \lambda)v &= \bar{\pi} - c(\lambda). \end{aligned} \tag{4}$$

It follows that innovation intensity λ is increasing in $\bar{\pi}$ and decreasing in $r + \mu$ or in a proportional increase in the cost function c. If $c'(0) > \frac{\bar{\pi}}{r+\mu}$ then the solution is $\lambda = 0$. We assume that $c'(\mu) > \frac{\bar{\pi} - c(\mu)}{r}$ so that the solution for λ is strictly less than μ. The reason for this restriction will be clearer below when we endogenize μ.

Our assumptions about knowledge capital are reflected in the firm's value function. Consider the value a firm obtains from a single product. If the firm were to simply market this product without

[9] A similar innovation production function has been used by Hall and Hayashi (1989) and Klette (1996).

[10] Penrose's arguments have partly been formalized by Uzawa (1969). Our formulation of how a firm's knowledge capital expands is also closely related to the formulation of physical capital accumulation in Lucas and Prescott (1971)

exploiting the knowledge capital to engage in new research, its expected present discounted value would be $v_1 = \bar{\pi}/(r+\mu)$. The actual value v exceeds v_1 by the amount $v_2 = [\lambda\bar{\pi}/(r+\mu) - c(\lambda)]/(r+\mu-\lambda)$ which cannot be negative since the firm is free to choose $\lambda = 0$ (in which case $v_2 = 0$). The term v_2 is value of the firm's option to exploit its knowledge capital n by innovating in the future.

In this model R&D activity and innovation are determinant at the firm level yet there is a certain decoupling of innovation and the size of firms. Since innovation intensity λ is independent of firm size, two firms of size n together will innovate at the same rate as one firm of size $2n$. Similarly research intensity (the fraction of firm revenues spent on R&D) is independent of firm size, $R/n = nc(\lambda)/n = c(\lambda)$. Thus, total R&D performed by a group of firms (and their total amount of innovation) depends on the overall size of the group but not on how firm size is distributed within the group. This property of the model will keep the analysis very simple in what follows.

4 Firm Evolution

Given the intensity of creative destruction μ, and having solved for innovation intensity λ, we can characterize the growth process for an individual firm. Consider a firm of size n. At any instant of time it will either remain in its current state, acquire a product and grow to size $n+1$ or lose a product and shrink to size $n-1$. A firm of size one that loses a product has exited (state zero is absorbing).

Let $p_n(t;n_0)$ denote the probability that a firm is size n at date t given that it is size n_0 at date 0. The rate at which this probability changes over time $\dot{p}_n(t;n_0)$ depends on the probabilities that the firm makes an innovation or loses one of its products. As formally derived in Appendix B, firm evolution is described by the following system of equations:

$$\dot{p}_n(t;n_0) = (n-1)\,\lambda p_{n-1}(t;n_0) + (n+1)\mu p_{n+1}(t;n_0) - n\,(\lambda+\mu)\,p_n(t;n_0) \qquad n \geq 1. \tag{5}$$

The reasoning is as follows: (i) if the firm had $n-1$ products then with a hazard $I(n-1) = (n-1)\lambda$ it innovates and becomes a size n firm, (ii) if the firm had $n+1$ products it faces a hazard $(n+1)\mu$ of losing one and becoming a size n firm, but (iii) if the firm already had n products it might either innovate or lose a product in which case it moves to one of the adjoining states. The equation for state $n=0$ is,

$$\dot{p}_0(t;n_0) = \mu p_1(t;n_0). \tag{6}$$

Exit is an absorbing state.

The solution to the set of coupled difference-differential equations (5) and (6) can be summarized by the probability generating function (pgf). In the appendix we show that the pgf is

$$H(z,t;n_0) = \sum_{n=0}^{\infty} p_n(t;n_0) z^n = \left[\frac{\mu(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)}{\lambda(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)}\right]^{n_0}.$$

By repeated differentiation of the pgf we can recover the entire density $p_n(t;n_0)$ for each date t and conditional on any initial size n_0, as shown in the appendix. To get some intuition about what is going on, we consider a few particular properties of this distribution.

4.1 Small Firms

We assume below that firms begin with a single product. To track the distribution at date t of a firm entering at date 0 we set $n_0 = 1$. The pgf then yields,

$$\begin{aligned} p_0(t;1) &= \frac{\mu}{\lambda}\gamma(t), \\ p_1(t;1) &= [1-p_0(t;1)][1-\gamma(t)], \\ p_n(t;1) &= p_{n-1}(t;1)\gamma(t), \qquad n = 2,3,\ldots, \end{aligned} \tag{7}$$

where

$$\gamma(t) = \frac{\lambda(1-e^{-(\mu-\lambda)t})}{\mu - \lambda e^{-(\mu-\lambda)t}}.$$

This term satisfies $\gamma(0) = 0$, $\gamma'(t) > 0$, and $\lim_{t\to\infty}\gamma(t) = \lambda/\mu < 1$. Figure 1 illustrates how a firm starting with only one product evolves stochastically over time (fixing $\lambda = 0.33$ and $\mu = 0.35$). The most steeply increasing curve corresponds to the probability of exit, while the six other curves corresponds to the firm having become the producer of $n = 1, 2, 3, 5, 10$ and 20 products, conditional on survival. Notice that $\lim_{t\to\infty} p_0(t,1) = 1$, i.e. all new firms will eventually exit.

Conditioning on survival, we get the simple Geometric distribution (shifted 1 to the right) at any fixed date,

$$\frac{p_n(t;1)}{1-p_0(t;1)} = [1-\gamma(t)]\gamma(t)^{n-1}, \qquad n = 1,2,\ldots.$$

The parameter of this distribution $\gamma(t)$ increases in t making the distribution grow stochastically larger over time.

4.2 Firm Age

A firm enters at a size of 1. Let A denote the random age of the firm when it eventually exits. The probability that the firm exits before age a is simply $p_0(a;1)$. That is, the cumulative distribution function of firm age is $\Pr[A \le a] = p_0(a;1)$. The expected length of life of a firm is thus,

$$E[A] = \int_0^\infty [1-p_0(a;1)]da = \frac{\ln\frac{\mu}{\mu-\lambda}}{\lambda},$$

which is decreasing in the intensity of creative destruction, μ, holding λ fixed. Holding μ fixed, as λ gets close enough to μ the expected life of a firm rises.

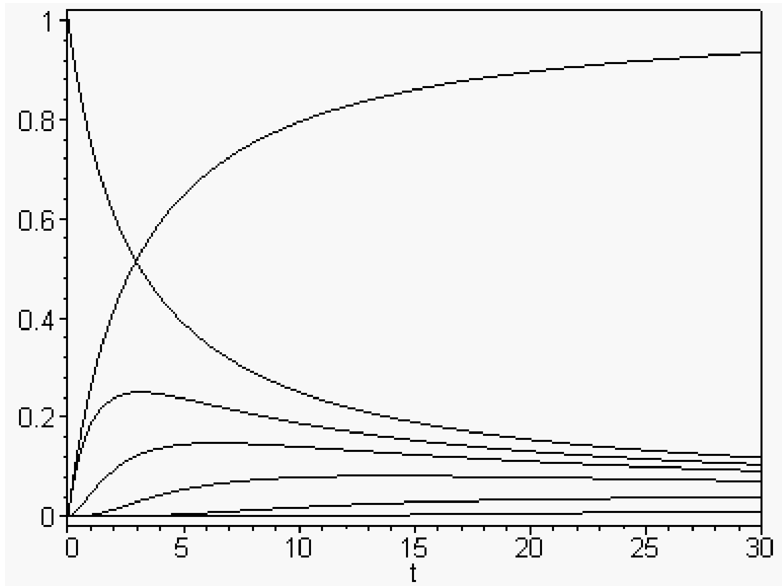

Figure 1: The evolution of a firm starting with one product when $\lambda = 0.33$ and $\mu = 0.35$. The most steeply increasing curve corresponds to the probability of exit, while the six other curves corresponds to the firm having become the producer of $n = 1, 2, 3, 5, 10$ and 20 products, conditional on survival.

The hazard rate of exit is

$$h(a) = \frac{\dot{p}_0(a; 1)}{1 - p_0(a; 1)} = \frac{\mu(\mu - \lambda)}{\mu - \lambda e^{-(\mu-\lambda)a}} = \mu[1 - \gamma(a)]. \tag{8}$$

The last equality shows that hazard rate is simply the product of the rate of creative destruction and the probability of being in state $n = 1$ for a firm that has survived to age a. It follows that $h(0) = \mu$, $h'(a) < 0$, and $\lim_{a\to\infty} h(a) = \mu - \lambda$. As a firm ages, implying that it has survived, it tends to get bigger and is therefore less likely to exit. The expected firm size, conditional on survival, is

$$\sum_{n=1}^{\infty} n \frac{p_n(a; 1)}{1 - p_0(a; 1)} = \frac{1}{1 - \gamma(a)}, \tag{9}$$

which increases with age, a.

4.3 Large Firms

It is convenient to think of a firm of size n_0 as if it consists of n_0 divisions each starting out at size 1. The form of the pgf implies that the evolution of the entire firm is obtained by summing the evolution of these independent divisions, each behaving as would a size 1 firm. Conditional on a single division surviving to date t, we already derived that its size is drawn from the Geometric distribution (shifted one to the right) with parameter $\gamma(t)$. If exactly $0 < k \leq n_0$ divisions survive to date t, the total size of the firm is the sum of k independent random variables drawn from the Geometric distribution, which is simply the Negative Binomial distribution (shifted k to the right) $\binom{n-1}{k-1}\gamma(t)^{n-k}[1-\gamma(t)]^k$. The probability that exactly $0 \le k \le n_0$ divisions of the firm survive to date t is given by the Binomial density $\binom{n_0}{k}[1-p_0(t;1)]^k p_0(t;1)^{n_0-k}$. For the case of $k=0$ we get the probability that the firm exits by date t, $p_0(t;1)^{n_0}$. Hence, since $\lim_{t\to\infty} p_0(t,1)=1$, all firms will eventually exit. To get the probability that the firm is size $n>0$ at date t we take a Binomial mixture of Negative Binomial distributions. Although the resulting distribution is complicated, the moments of the distribution are easy to obtain. We use them below to analyze firm growth.

4.4 Firm Growth

Without solving for $p_n(t;n_0)$ we can derive the moments of firm growth directly from the pgf , as shown in the appendix. Letting the random variable N_t denote the size of a firm at date t, its growth over the period from date 0 to t is $G_t = (N_t - N_0)/N_0$. It turns out that our model is consistent with Gibrat's law, i.e. expected firm growth given initial size is

$$E\left[G_t \middle| N_0 = n_0\right] = e^{-(\mu-\lambda)t} - 1, \tag{10}$$

which is independent of initial size.

The variance of firm growth given initial size is

$$\text{Var}\left[G_t | N_0 = n_0\right] = \frac{\lambda+\mu}{n_0(\mu-\lambda)} e^{-(\mu-\lambda)t}\left(1-e^{-(\mu-\lambda)t}\right), \tag{11}$$

which declines in initial firm size. The growth of a larger firm is an average of the growth of its independent components, hence the variance of growth is inversely proportional to the firm's initial size.

In the derivations above we have included firms which exit during the period. It is also possible to condition on survival. As shown above, the probability that a firm of size n_0 at date 0 survives to date t is $1-p_0(t;n_0) = 1-p_0(t;1)^{n_0} = 1-(\mu/\lambda)^{n_0}\gamma(t)^{n_0}$, which is clearly increasing in initial size. Expected growth *conditional on survival* is

$$E\left[G_t | N_t > 0, N_0 = n_0\right] = \frac{e^{-(\mu-\lambda)t}}{1-[(\mu/\lambda)\gamma(t)]^{n_0}} - 1,$$

which is a declining function of initial size. Knowing that an initially small firm has survived suggests that it has grown relatively fast. For firms which are initially very large the probability of survival over the period is close to 1 anyway, so Gibrat's law will be a very good approximation.

5 Entry and Industry Evolution

In this section we will examine the dynamic process characterizing an industry with many competing firms. Each firm within the industry innovates according to the model presented in the previous section.

In considering an industry, we assume that every product is being produced by some firm. The number of goods produced by any given firm is countable, hence there must be a mass of firms to account for the unit continuum of products. We can describe the population of firms in terms of the measure of firms of each size. There is no randomness at the industry level.

We denote, for date t, the measure of firms in the industry with n products by $M_n(t)$. Because there is a unit mass of products, and each product is produced by exactly one firm, $\sum_{n=1}^{\infty} nM_n(t) = 1$.

Taken as a whole, industry incumbents innovate at rate

$$\sum_{n=1}^{\infty} M_n(t)I(n) = \sum_{n=1}^{\infty} M_n(t)n\lambda = \lambda,$$

where λ is the innovation intensity of incumbent firms. Innovative activity is related to firm size and yet the size distribution of firms has no implications for the total amount of innovation carried out by incumbents.

Although each firm takes the intensity of creative destruction μ as given, this magnitude is determined endogenously for the industry. One component of it is the rate of innovation by incumbents λ. The other component is the rate of innovation by entrants η, so that

$$\mu = \eta + \lambda. \tag{12}$$

We now turn to the determination of η.

5.1 Firm Entry

As mentioned above, we assume that entrants begin with a single product. The value of entering is therefore $V(1) = v$. We assume that there is a fixed cost of entry F and a mass of potential entrants. If there is any entry, equilibrium requires $F = v$.

The entry condition, combined with the optimal research policy of incumbents, pins down the innovation intensity of incumbents:

$$c'(\lambda) = F, \tag{13}$$

or $\lambda = 0$ if $c'(0) > F$. If the cost of entry is higher incumbents are shielded more effectively from competition and therefore invest more in innovation. On the other hand, with active entry, (13) shows that the innovation intensity of incumbents is unrelated to demand side incentives such as the profit flow from an innovation, $\bar{\pi}$. The reason is that entry responds to these demand side incentives, exactly neutralizing any potential influence on incumbents.

To pin down the rate of innovation by entrants, we return to the expression for the value function from (27). Rearranging it under the assumption that there is active entry, i.e. $F = v$, and using (12),

$$\eta = \frac{\bar{\pi} - c(\lambda)}{F} - r.$$

To guarantee $\eta > 0$, we assume $c'(\lambda) > F$, where λ solves $\frac{\bar{\pi}-c(\lambda)}{r} = c'(\lambda)$. An implication is that $\lambda < \mu$ as we assumed above.

Having pinned down the entry rate, we can examine the behavior of a cohort of firms defined by when the firms enter. Of the cohort entering at data $t-a$, the fraction surviving until t is $1-p_0(a;1)$. The market share for this age a cohort is

$$\eta \sum_{n=0}^{\infty} n p_n(a;1) = \eta \frac{1-\frac{\mu}{\lambda}\gamma(a)}{1-\gamma(a)}. \tag{14}$$

Since, $\mu > \lambda$, the market share of a cohort declines steadily with a. Even though the surviving firms in the cohort tend to get larger, this effect is more than offset by exit from the cohort.

5.2 The Size Distribution

We now have expressions for η, λ, and hence $\mu = \eta + \lambda$. These parameters are all that matter for analyzing the size distribution of firms. The configuration of the industry is summarized by the measure of firms with $1, 2, 3, \ldots$ products. We define $M(t) = \sum_{n=1}^{\infty} M_n(t)$ to be the measure of firms in the industry, and solve for it below.

Since there is a mass of firms of any given size, we can describe its rate of change over time by a non-stochastic differential equation. Flowing into the mass of firms with n products are firms with $n-1$ products that just acquired a new product and firms with $n+1$ products that just lost one. Flowing out of the mass of firms with n products are firms that were of that size and that either just acquired or just lost a product. Thus, for $n \geq 2$,

$$\dot{M}_n(t) = (n-1)\lambda M_{n-1}(t) + (n+1)\mu M_{n+1}(t) - n(\lambda+\mu)M_n(t). \tag{15}$$

For $n = 1$ we have

$$\dot{M}_1(t) = \eta + 2\mu M_2(t) - (\lambda+\mu)M_1(t). \tag{16}$$

To obtain the steady state distribution we can simply set all the time derivatives to zero. In the appendix, we show that the solution is

$$M_n = \frac{\lambda^{n-1}\eta}{n\mu^n} = \frac{\phi}{n}\left(\frac{1}{1+\phi}\right)^n, \qquad n \geq 1, \tag{17}$$

where $\phi = \eta/\lambda$. The mass of large firms is greatest as ϕ approaches zero, i.e. when there is little entry. The total mass of firms is

$$M = \phi \sum_{n=1}^{\infty} \frac{\left(\frac{1}{1+\phi}\right)^n}{n} = \phi \ln\left(\frac{1+\phi}{\phi}\right). \tag{18}$$

In a steady state, the mass of firms in the industry is an increasing function of the entry rate. The mass approaches one as the entry rate gets arbitrarily large (or λ gets arbitrarily small). In this case almost all firms are new entrants with just one product. As the rate of entry approaches zero (or λ gets arbitrarily large), the mass of firms in the industry gets very small, indicating that the average firm is large, i.e. has many products.[11]

More generally, we can analyze distribution outside of a steady state by integrating over the history of cohorts of entrants while taking account of how these cohorts evolve. Starting with the size distribution, $M_{n_0}(0)$, at date 0,

$$M_n(t) = \sum_{n_0=1}^{\infty} p_n(t;n_0)M_{n_0}(0) + \eta \int_0^t p_n(s;1)ds = \sum_{n_0=1}^{\infty} p_n(t;n_0)M_{n_0}(0) + \frac{\eta}{n\lambda}\gamma(t)^n. \tag{19}$$

All firms in existence at date 0 will eventually exit, hence

$$\lim_{t\to\infty} p_n(t,n_0) = 0, \qquad \forall\, [n,n_0] \geq 1.$$

It can be shown that the first summation in (19) disappears as $t \to \infty$, while the second term converges to $\frac{\eta}{n\lambda}(\lambda/\mu)^n = M_n$. Thus, the system converges to the steady state distribution (17), and in fact we can trace out its evolution during this process of convergence using (19).

6 Innovators and Imitators

We showed above that a firm's research intensity, its research expenditure as a fraction of revenues, is $R(n)/n = c(\lambda)$. This result is attractive in the sense that research intensity is pinned down at the firm level, persistent over time, and unrelated to the size of the firm. However, measures of research intensity display considerable cross-sectional variability. In fact, the variation of research intensity across firms within industries is at least as great as the variation of average research intensity across different industries.

Incorporating heterogeneity in research intensity into the model, we need to be careful not to build in a result that research intensive firms become big firms. If this were the case then size would be a good predictor of research intensity, which it is not (see, Fact # 3). To avoid this implication we seek to unhinge the research process from the process of revenue growth.

We assume that firms differ in how innovative they are. Empirical research in economics, sociology and management science has emphasized the high degree of persistency in differences in innovative strategies

[11] We refer to M_n for $n = 1, 2, \ldots$ as the size distribution, although for it to be a proper distribution we should normalize by the total mass of firms. Doing so we get

$$P_n = \frac{\left(\frac{1}{1+\phi}\right)^n}{n \ln\left(\frac{1+\phi}{\phi}\right)} \qquad n = 1, 2, \ldots .$$

This is the well known logarithmic distribution, as discussed in Johnson et al. (1993). The mean of the distribution (the average number of products per firm) is $\phi^{-1}/\ln(1+\phi^{-1})$, which is decreasing in ϕ. The logarithmic distribution is discussed in the context of firm sizes by Ijiri and Simon (1977).

across firms[12]. To capture such differences, we relax the assumption that all firms receive the same flow of profits $\bar{\pi}$ from marketing a good. Instead we allow the flow of profits π to vary across firms. An firm making drastic innovations obtains a flow of profits close to one while a less innovative firm obtains π closer to zero. A firm with π close to zero may be considered more of an imitator than an innovator[13].

We index firms by $0 < \pi < 1$. The index not only affects the flow of profits but also the cost of doing research. To obtain independence between inventiveness and firm growth, we assume that the research cost function for a firm of type π is $C_\pi(I,n) = \frac{\pi}{\bar{\pi}} C(I,n)$. It follows that the intensive form of the cost function is $c_\pi(I/n) = \frac{\pi}{\bar{\pi}} c(I/n)$. The distribution of π in the pool of entrants is denoted $h(\pi)$, with the mean of the distribution being $\bar{\pi}$. Our original assumption of a fixed flow of profits is obtained as the special case when $h(\pi)$ has zero variance.

Returning to a firm's optimal innovation decision, we modify the conditions only slightly. In particular, the solution is $v_\pi = \frac{\pi}{\bar{\pi}} v$ and $\lambda_\pi = \lambda$, where v and λ were the solutions from before. Innovative intensity and hence firm growth is unrelated to π. On the other hand, the value of the firm rises in proportion to π.

We have to modify the entry condition only slightly. We assume that entrants do not know their type when they pay the entry cost, but they know the distribution of types in the pool of entrants. Under these conditions, if there is active entry we still have the condition $F = E[v_\pi] = v$. The entry rate η remains unchanged.

We have allowed for firm heterogeneity in a way that has no effect on firm dynamics. It follows that the distribution of π in the population of incumbents is h, whether or not we condition on firm size. Research intensity rises in proportion to π since $c_\pi(\lambda) = \frac{\pi}{\bar{\pi}} c(\lambda)$. Thus, the distribution h determines the distribution of research intensity among firms of any size.

The assumption that differences across firms in their research capability do not give rise to different innovation rates may seem restrictive, and we examine a more general specification in Appendix D.

7 Revisiting the Evidence

We conclude by returning to the evidence on innovating firms. To what extent have we succeeded in capturing the various facts? To what extent does the theoretical model lead to a different interpretation of these facts?

7.1 R&D Investment

Consistent with Fact # 3, our model predicts R&D investment in proportion to firm size. Equivalently, R&D intensity is independent of size. The reason is that larger firms take advantage of their greater

[12] See e.g. Henderson (1993), Cohen (1995), Langlois and Robertson (1995), Carroll and Hannan (2000), Cockburn et al. (2000) and Jovanovic (2000). Geroski et al. (1998) presents a contradictory view (see their stylized fact 5).

[13] This story is consistent with arguments in Nelson (1988).

knowledge capital to offset the diminishing returns to R&D. They therefore scale up their research investments in proportion.

As generalized in Section 6, the model allows for heterogeneity in research intensity, in line with Fact # 4. This heterogeneity can be quite general, simply reflecting the distribution $h(\pi)$. Firms which make greater innovative steps (as captured by π) are also more research intensive. One might expect such firms to grow faster as well, yet if they did we would predict a strong positive correlation between firm size and research intensity. To avoid contradicting Fact # 3 requires the strong assumption, imposed in Section 6, that research costs rise in proportion to the profitability of innovations.

Although the model predicts that all firms will do at least some R&D, firms with a very small value of π may have little innovative activity that is reported as formal R&D. Hence, our interpretation of the second half of Fact # 4 is that there are a considerable number of firms with low π's, our imitators, even in high-tech industries.

The factor determining R&D intensity is modeled as being a characteristic of the firm itself. Thus R&D intensity is highly persistent over time, in line with Fact # 5. In line with Fact # 6, the level of R&D investment by a firm follows close to a geometric random walk as it stays in step with the changing number of goods produced by the firm.

7.2 R&D and Patenting

The firm-level innovation production function, (1), is essential to the workings of the model. Yet, it turns out to be quite plausible on its own. It embodies decreasing returns to expanding research at the firm level, which is consistent with the second half of Fact # 2. On the other hand, there are also benefits to size. Considering firms of different sizes (but with no heterogeneity in π), their optimal R&D policy is $R(n) = nc(\lambda)$. Plugging this R&D policy into the innovation production function (1) implies that a Poisson regression of the number of innovations by a firm on its R&D expenditure should yield a Poisson parameter proportional to R&D. When the number of innovations by the firm is proxied by its number of patents, this prediction is borne out as noted in Fact # 2.

In the case of heterogeneous research intensity, the cross-sectional relationship between innovations (as captured by patenting) and R&D becomes less clear cut. Now R&D will vary across firms not only due to heterogeneity in firm size n but also due to heterogeneity in π. The expected number of innovations rises in proportion to the former but is unrelated to the latter. Nonetheless larger innovative steps may be associated with more patentable inventions. If so, the model may still be roughly consistent with the cross-sectional relationship between patents and R&D.

7.3 R&D and Productivity

Let us start by presenting a somewhat surprising argument which we think has some interest beyond the framework presented above. Productivity in a firm i, A_i, is defined as output Y_i relative to an index of input X_i, or $A_i = Y_i/X_i$. We assume constant returns to scale for simplicity. Physical output is rarely observed, and researchers typically use nominal output, P_iY_i, instead of real output when they construct their measure of productivity. Hence, measured productivity is $\tilde{A}_i = P_iY_i/X_i$. With imperfect competition, the price is $P_i = m_iC_i'$, where m_i is the markup and C_i' is marginal costs. It follows that $P_i = m_i W_i/A_i$, where W_i is the price per unit of factor input. Using this expression for P_i, measured productivity is

$$\begin{aligned}\tilde{A}_i &= P_i\frac{Y_i}{X_i} = m_i\frac{w_i}{A_i}A_i \\ &= m_iW_i.\end{aligned}$$

That is to say, differences in *measured* productivity do not directly reflect differences in *real* productivity, but only indirectly through differences in markups and factor prices if they are correlated with differences in real productivity! In other words, to the extent that firms capture their innovative rents through large markets shares rather than high markups, differences in measured productivity may contain little information about differences in firms' innovative outputs.[14]

Returning to our framework, more innovative firms with high π will tend to have higher markups, and they will also have higher research costs. Hence, in a cross section of firms, there will be positive correlation between measured productivity and any measure of R&D expenditures (i.e. either an R&D stock or R&D intensity), consistent with Fact # 1.

Notice, however, that the relationship between measured productivity and R&D contains little information about the rate of return to R&D expenditures. For instance, consider the case where longitudinal variations in the π's for any firm are just random noise, uncorrelated with R&D costs. In this case, there will be zero correlation between measured productivity and R&D expenditures, yet R&D will, on average, earn the required rate of return.

Finally, with our assumption that, for any firm, π does not vary over time, our model suggests that there will be no variation in measured productivity in the longitudinal dimension (apart from noise). The empirical evidence supports the model's implication of persistent differences across firms in profitability and productivity, see Geroski (1998). Hence, longitudinal movements in measured productivity are uninformative about innovations, as pointed out in Fact # 1.

[14] Bernard et al. (2000) develop this argument in the context of the productivity advantage of exporting plants.

7.4 Entry, Exit, and Growth

As shown in equation (8), our model implies that the hazard rate is steadily declining with firm age, consistent with Fact # 10. Our model is also consistent with the second half of Fact # 10, stating that firm size, conditional on survival, is increasing, as shown in equation (9), and that a cohort's total market share declines over time, as shown in equation (14).

The declining hazard rate has been widely interpreted as evidence in favor of Jovanovic's (1982) model of selection of firms based on given differences in fitness. According to his model, new firms receive noisy signals about their fitness, and these signals adds more information when the firms are young. A large number of ill-fated firms recognize their unfitness upon receiving their first signals. In our model, the declining hazard arises without any *ex ante* differences in fitness. Instead the declining hazard is due to selection based on *ex post* innovation outcomes. The *surviving* firms are on average growing, giving them a less vulnerable market position and hence a steadily declining hazard rate.

Our model predicts that firm growth, conditional on survival, is a declining functions of initial size, as shown in (10). For larger firms, the model predicts that Gibrat's law is a very good approximation. The model is consequently consistent with Fact # 8. We have also shown in (11) that the model is consistent with Fact # 9[15], i.e. the variance of firm growth rates is declining in initial size.

7.5 Size Distribution

The size distribution derived in (17) is highly skewed, and consequently consistent with Fact # 7 in Section 2.3. The distribution has a long tail of large firms, corresponding to higher market concentration, if there is a low entry rate of new firms relative to the rate of innovation by incumbents (i.e. ϕ is small).

According to (17), the size distribution of firms and, in particular, market concentration is determined by the parameter $\phi = \eta/\lambda$. Recall that η and λ are the rate of innovation by entrants and incumbents, and both parameters have been derived as functions of more basic parameters in Section 5.1. There is a large empirical literature searching across industries for the relationship between market concentration, the rate of entry of new firms, and the rate of innovation. Most, if not all, studies have searched for bi-variate relationships between either the rate of innovation and market concentration, or innovation and entry. The results have not been very consistent or robust, according to Cohen (1995).

The evidence, if anything, suggests (i) a positive relationship between R&D intensities and market concentration (Cohen, pg. 192), and (ii) a modest positive association between innovation and entry (Cohen, pg. 194). In view of our framework, it is not surprising that the empirical findings have been rather weak, as our result in (17) suggests that there is a more complex *tri-variate* relationship. Still, our relationship $\lambda = \eta/\phi$ together with (17), is consistent with the empirical findings suggesting that the

[15] Although Amaral et al. (1998) show that the variance decline more moderately with size than predicted by (11). See also Sutton (2000).

rate of innovation by incumbents is positively related to both the rate of entry and market concentration. Notice that, according to our framework, the relationship between innovation, market concentration, and entry is not causal, but derived as simultaneous outcomes of more basic determinants of firm behavior and industry structure, i.e. demand, entry costs and innovative opportunities.

8 General Equilibrium and Growth

We now show how our model of a firm and an industry can be imbedded in a general equilibrium model of endogenous technological change. In particular, with a few additional assumptions, the general equilibrium appears much like the quality ladders model in chapter 4 of Grossman and Helpman (1991).

For simplicity we consider the special case in which research intensity does not vary, i.e $h(\pi)$ has zero variance. (*We believe the results will generalize, with a few added complications.*) In the context of a quality ladders model our innovations take the form of product improvements. The inventive step $q > 1$ represents the factor improvement in quality so that if quality was x, after an innovation it rises to qx.

8.1 The Market Equilibrium

The representative consumer has preferences over the continuum of goods:

$$U_t = \int_t^{\infty} e^{-\rho(\tau - t)} \left\{ \int_0^1 \ln \left[x_\tau(j) q^{J_\tau(j)} \right] dj \right\} d\tau,$$

where ρ is the discount rate, $x_\tau(j)$ is consumption of good j at date τ, and $J_\tau(j)$ is the number of innovations for good j that have taken place by date τ. Following Grossman and Helpman, we choose aggregate expenditures to be the numeraire, setting them to one in each period. This normalization delivers our assumption above that the flow of revenues per good is always one. Another implication of this normalization is that $r = \rho$.

An innovator will engage in Bertrand competition with the incumbent he displaces. The result of this competition is that the innovator takes over the market and charges a markup of price over marginal cost equal to the inventive step, q. The profit flow to an innovator is therefore $\bar{\pi} = 1 - q^{-1}$.

There is a fixed endowment of L workers. Workers can either produce output (as do L^X), create start-up firms (as do L^S), or perform research within incumbent firms (as do L^R). Thus $L = L^X + L^S + L^R$. We will examine an equilibrium in which all activities take place so that the wage w for each activity is the same.

We derive labor demand for each of the three activities starting with production work. Given unit expenditure per product, the quantity produced of each good is $x(j) = 1/p(j) = 1/(qw)$. Choosing units of output so that production workers have productivity one, we get $L^X = 1/(qw)$.

We now turn to demand for researchers at incumbent firms. We assume that all research costs are in the form of researcher wages. Thus, the number of researchers employed at any firm of size n is

$$L^R(n) = R(n)/w = nc(\lambda)/w = n\bar{c}(\lambda),$$

which implicitly defines $\bar{c}(\lambda) = c(\lambda)/w$. (A researcher at an incumbent firm gets an innovation at a Poisson arrival rate of $\lambda/\bar{c}(\lambda)$.) Aggregating over firms of different sizes,

$$L^R = \sum_{n=1}^{\infty} M_n L^R(n) = \sum_{n=1}^{\infty} M_n n\bar{c}(\lambda) = \bar{c}(\lambda).$$

Finally, a worker trying to launch a firm gets an idea for an innovation at a Poisson arrival rate ψ. Thus the cost of entry is the opportunity cost of the expected time until the first innovation, $F = w/\psi$. The value of a start-up firm (with one product) is still v. The equilibrium condition (assuming active entry) becomes $w = \psi v$ or $v = w/\psi$. Since $c'(\lambda) = v$ it follows that $\bar{c}'(\lambda) = \psi^{-1}$. Thus, the innovation rate of incumbents is tied down by the productivity ψ of entrants. The equation for the entry rate becomes,

$$\eta = \frac{\bar{\pi} - c(\lambda)}{w/\psi} - r = \frac{(1-q^{-1})\psi}{w} - \psi\bar{c}(\lambda) - \rho.$$

Given η, we can read off the demand for start-up labor from $\eta = \psi L^S$.

We can now solve for the wage that equates the labor endowment L and the three sources of demand for labor, assuming a steady state with active entry. Simplifying and rearranging we obtain $w = \frac{\psi}{L\psi+\rho}$. Substituting this value of the wage into our expression for the rate of entry,

$$\eta = (1-q^{-1})L\psi - \frac{\rho}{q} - \psi\bar{c}(\lambda).$$

The equilibrium growth rate of the economy is $\mu \ln q$, i.e. the rate of innovation by both entrants and incumbents multiplied by the percentage size of inventive steps. The overall innovation rate is

$$\mu = \eta + \lambda = (1-q^{-1})L\psi - \frac{\rho}{q} + \lambda - \psi\bar{c}(\lambda),$$

where, as noted above, λ solves $\bar{c}'(\lambda) = \psi^{-1}$. The first two terms of the expression for μ are identical to Grossman and Helpman (1991). The last two terms are new. They arise because of the knowledge capital of incumbent firms.

To illustrate, we can eliminate the knowledge capital of incumbents by assuming $\bar{c}(\lambda) = \lambda\bar{c}$. The productivity of researchers at incumbent firms is thus $1/\bar{c}$. If we assume the productivity of researchers at start-up firms exceeds the productivity of researchers at incumbent firms then incumbent firms will choose $\lambda = 0$ and all research will be performed by entrants. Furthermore, the last two terms in the expression for μ drop out, yielding the Grossman and Helpman result exactly.[16]

[16] Since our model yields aggregate predictions nearly identical to the model of Grossman and Helpman (1991), it therefore suffers from the empirical problems pointed out by Jones (1995). We believe that these empirical shortcomings at the aggregate level could be eliminated by modifying the model along the lines of Kortum (1997), Eaton and Kortum (1999), or Howitt (2000).

8.2 Welfare

Is the market equilibrium efficient? In the appendix we solve the social planner's problem. It turns out that the market equilibrium produces the social optimal innovation intensity for incumbents. The deviation between the social planners solution and the market equilibrium arises from differences in the entry rate. Denoting the social planner's solution with a *, we have,

$$\mu^* - \mu = \eta^* - \eta = \frac{\rho}{q}\left(\frac{\psi L}{\rho} + 1 - \frac{q}{\ln q}\right).$$

The market economy grows too slowly when the inventive step q is in an intermediate range. On the other hand, if q is either very small or very large then the market economy grows too quickly. This deviation between the innovation rates in the market economy and the planner's problem is identical to that obtained by Grossman and Helpman (1991). Although their model does not have research by incumbents, it turns out that the incumbents always do the optimal amount of research in the market economy.

9 Conclusions

This paper develops a theoretical model of innovating firms that is rich enough to confront the detailed evidence yet parsimonious enough to be analytically tractable. The model provides a simple and coherent framework to account for the many empirical findings on firm growth and size, market concentration, and R&D inputs and outputs. The model describes the dynamic behavior of individual heterogenous firms, captures industry behavior through distributions of outcomes, and aggregates into a general equilibrium model of endogenous growth.

Our approach gives new interpretations to some of the empirical evidence. According to the model the widely studied empirical relationship between R&D and productivity contains little information about the impact of R&D on innovation. The model explains why R&D effort is strongly related to productivity and patenting yet essentially unrelated to firm size or growth. In the end, the model demonstrates that the firm-level evidence, while seemingly puzzling, supports the view that R&D is crucial for aggregate growth.

The model has been developed in order to confront, at least qualitatively, the main empirical evidence on innovating firms. We intend to extend the empirical confrontation in two directions. First, the model contains further empirical implications for correlations across industries and countries in entry rates, the evolution pattern of new cohorts of firms, various statistics of the firm-size distributions, and innovative inputs and outputs. This line of research would more carefully consider the quantitative and not just the qualitative predictions from the model. The second extension of the empirical confrontation is to consider in detail firm level data from an individual industry to see whether a calibrated version of the model can match facts about the industry's evolution.

Appendix A: The Firm's Optimization Problem

According to dynamic programming, the intertemporal optimization problem for a firm with $n \geq 1$ products, facing the constraints described in Sections 3.1-3.2, can be formulated as

$$V(n) = \max_{I} \left\{ [\pi n - nc(I/n)] \Delta t + \frac{1}{1 + r\Delta t} E[V(n')|I] \right\}, \tag{20}$$

where n' is the state of the firm after some small length of time Δt has elapsed. We treat a firm in state $n = 0$ as having permanently exited the industry, so that $V(0) = 0$. The entry decision for a new firm is modeled in Section 5.1, and will not be considered here.

The firm increases its number of products with propensity I, while losing products at a rate μn. The firm considers μ to be exogenous, but at the industry level there is a tight link between μ and the total amount of innovation across all firms as shown in Section 5.

Letting Δt be arbitrarily small, we need to worry about the firm either acquiring or shedding at most one variety. Thus,

$$E[V(n')|I] = [1 - I\Delta t - \mu n \Delta t] V(n) + I\Delta t V(n+1) + \mu n \Delta t V(n-1). \tag{21}$$

Inserting (21) into (20), multiplying both sides by $(1 + r\Delta t)$, and neglecting terms of second or higher order in Δt, it follows that

$$rV(n) = \max_{I} \{\bar{\pi} n - nc(I/n) + I[V(n+1) - V(n)] - \mu n [V(n) - V(n-1)]\}, \tag{22}$$

The first order condition for the firm's optimal level of research is

$$c'\left(\frac{I^*}{n}\right) = V(n+1) - V(n), \tag{23}$$

where I^* is the optimal level of R&D. Since we assume that $c'' < 0$, it follows the second order condition is satisfied. Consider a solution of the form

$$V(n) = vn \tag{24}$$

$$I^* = \lambda n. \tag{25}$$

Then (23) can be rewritten

$$c'(\lambda) = v. \tag{26}$$

Substituting (24) into (22) and dividing both sides by n,

$$v = \frac{\pi - c(\lambda)}{r + \mu - \lambda}, \tag{27}$$

which confirms our conjecture that v does not depend on n. Equations (24) and (25), together with (26) and (27), identify the value function and the optimal policy solving the firm's dynamic optimization problem.

Appendix B: Solving the System of Difference-Differential Equations

The derivation of equation (6) can be seen as follows. The probability for the firm having $n \geq 1$ products at time $t + \Delta t$ satisfies the relationship

$$\begin{aligned} P\left[N(t+\Delta t)=n\right] &= (n-1)\,\lambda\Delta t\, P\left[N(t)=n-1\right] \\ &\quad + (n+1)\,\mu\Delta t\, P\left[N(t)=n+1\right] \\ &\quad + \left[1-n\left(\lambda+\mu\right)\Delta t\right]\; P\left[N(t)=n\right] + \mathcal{O}\left(\Delta t\right) \end{aligned}$$

Following standard techniques described e.g. in ch. 6 in Taylor and Karlin (1998) and Section 2 in Goel and Richter-Dyn (1974), we find that

$$\begin{aligned} \frac{\partial P\left[N(t)=n\right]}{\partial t} &= \lim_{\Delta t\to 0}\frac{P\left[N(t+\Delta t)=n\right]-P\left[N(t)=n\right]}{\Delta t} \\ &= (n-1)\,\lambda\; P\left[N(t)=n-1\right]+(n+1)\,\mu\; P\left[N(t)=n+1\right] \\ &\quad -n\left(\lambda+\mu\right)\; P\left[N(t)=n\right] \end{aligned}$$

or in more compact notation

$$\dot{p}_n(t;n_0) = (n-1)\,\lambda p_{n-1}(t;n_0) + (n+1)\mu p_{n+1}(t;n_0) - n\left(\lambda+\mu\right)p_n(t;n_0) \qquad n \geq 1. \tag{28}$$

The probability of exit, i.e. hitting the absorbing state $n = 0$, is described by:

$$\dot{p}_0(t;n_0) = \mu p_1(t;n_0). \tag{29}$$

The solution to the set of coupled difference-differential equations (5) and (6) involves the probability generating function[17], defined as

$$H(z,t) = \sum_{n=0}^{\infty} p_n(t)\; z^n. \tag{30}$$

Hence

$$\begin{aligned} \frac{\partial H(z,t)}{\partial z} &= \sum_{n=0}^{\infty} n p_n(t)\; z^{n-1} \\ &= \sum_{n=1}^{\infty} n p_n(t)\; z^{n-1}. \end{aligned} \tag{31}$$

Consider

$$\frac{\partial H(z,t)}{\partial t} = \sum_{n=0}^{\infty} \dot{p}_n(t)\; z^n = \dot{p}_0(t) + \sum_{n=1}^{\infty} \dot{p}_n(t)\; z^n. \tag{32}$$

[17] The solution procedure is described e.g. in Goel and Richter-Dyn (1974) (in particular chapter 2 and appendix B)

The first term on the right hand side can be restated by (6), while the second term can be restated by multiplying (5) by z^n and summing over n from 1 to ∞, which, after some rearrangements of terms gives

$$\begin{aligned}\frac{\partial H(z,t)}{\partial t} &= -(\lambda+\mu)\sum_{n=1}^{\infty} n p_n(t)\, z^n \\ &\quad +\lambda\sum_{n=1}^{\infty}(n-1)\, p_{n-1}(t)\, z^n \\ &\quad +\mu\left[p_1+\sum_{n=1}^{\infty}(n+1)\, p_{n+1}(t)\, z^n\right]. \end{aligned} \tag{33}$$

Using (31) on each of the three sums, we find that (33) can be restated as

$$\frac{\partial H(z,t)}{\partial t} = \left[\lambda z^2-(\lambda+\mu)z+\mu\right]\frac{\partial H(z,t)}{\partial z}. \tag{34}$$

This is a partial differential equation of the Lagrangian type, and its solution is discussed in Goel and Richter-Dyn (1974).

The solution $H(z,t)$ has to fulfil an initial condition. If we wanted to consider new entrants, we would set $p_1(0)=1$. More generally, to analyze a firm starting in state n_0, we set $p_{n_0}(0)=1$ and $p_n(0)=0$ for $n\neq n_0$. From (30) it follows that

$$H(z,0;n_0)=\sum_{n=0}^{\infty} p_n(0)\, z^n = z^{n_0}. \tag{35}$$

With the condition (35), the solution to (34) can be written

$$H(z,t;n_0)=\left[\frac{\mu(z-1)e^{-(\mu-\lambda)t}-(\lambda z-\mu)}{\lambda(z-1)e^{-(\mu-\lambda)t}-(\lambda z-\mu)}\right]^{n_0}. \tag{36}$$

By Taylor expanding this expression for $H(z,t;n_0)$ in z at $z=0$, we derive the probability distribution $p_n(t)$ (as it evolves for a firm in state n_0 at date 0) as the coefficients in the Taylor series as can be seen from (30).

We notice that the condition $\sum_{n=0}^{\infty} p_n(t)=1$, $\forall t$, requires that

$$H(1,t;n_0)=\sum_{n=0}^{\infty} p_n(t)=1, \tag{37}$$

which is satisfied by the expression on the right hand side in (36).

The analytical expression (36) contains a full characterization of the probability distribution for the firm size at any date ($t\geq 0$), conditional on the firm's size at $t=0$. In particular, we obtain the probability distribution as

$$\begin{aligned} p_n(t;n_0) &= \frac{1}{n!}\frac{\partial^n}{\partial z^n}H(z,t;n_0)\Big|_{z=0}, \qquad n=1,2,\ldots \\ p_0(t;n_0) &= H(0,t;n_0). \end{aligned}$$

Furthermore, from (36) one can easily derive an analytical expression for the k-th moment for probability distribution describing firm size:

$$\sum_{n=0}^{\infty} n^k \, p_n(t; n_0) = \left[\frac{\partial^k H(e^s, t; n_0)}{\partial s^k} \right]_{s=0} . \tag{38}$$

Appendix C: The Size Distribution

>From (15) and (16), we have that

$$\dot{M} = \eta - \mu M_1. \tag{39}$$

In steady state all the time-derivatives in (15), (16), and (39) are zero. Starting with (39), the steady state requires

$$M_1 = \frac{\eta}{\mu}. \tag{40}$$

Substituting (40) into (16), we see that the steady state also requires,

$$M_2 = \frac{\lambda \eta}{2\mu^2}. \tag{41}$$

Applying (15), it is straightforward to prove by induction that the general condition for the steady state is,

$$M_n = \frac{\lambda^{n-1} \eta}{n \mu^n}, \qquad n \geq 1. \tag{42}$$

Appendix D: Differences in Research Efficiency

Change notation - just started in the first equation!

The assumption that differences across firms in their research capability do not give rise to different innovation rates seems restrictive, and we will now explore a more general specification. Assume that there are different types that differ in their research efficiency, i.e. with different research costs $c(\lambda)$ where the subscript θ indexes the type of firm, with the set of all types denoted by Θ. More specifically, we define

$$\theta = \frac{\pi}{c_\pi(\lambda)}.$$

Hence, there will be differences in the innovation rates

$$\lambda_\theta = \left[\frac{v(\theta, 1, r + \mu)}{1 + \alpha} \right]^{1/\alpha} .$$

Among entrants, we denote the frequency of type θ firms by $g(\theta)$, hence $\sum_{\theta \in \Theta} g(\theta) = 1$.

Allowing for firms with different λ's complicates the analysis of the firm size distribution only slightly. We can think about the different types of firms just as superimposed populations and they only affect each other through the rate of creative destruction. That is to say, the whole analysis in Sections 4 and 5 goes through if we consider the results to be conditional on a given type of firm. The only changes that must be made in the previous analysis is the derivation of μ and the free entry condition.

Given a specific type of firm, θ, the steady state analysis in Section 5 goes through. That is, in steady state, we have from Section 5 that the mass of firms of size n and type θ is given as

$$M_n(\theta) = \eta g(\theta) \frac{\lambda_\theta^{n-1}}{n\mu^n}, \tag{43}$$

where $\eta g(\theta)$ is the entry rate of type θ firms. It follows that, given any type of firms θ, the fraction of firms with n products is

$$\begin{aligned} P_n(\theta) &\equiv \frac{M_n(\theta)}{\sum_{n=1}^{\infty} M_n(\theta)} \\ &= \frac{1}{n}\left(\frac{\lambda_\theta}{\mu}\right)^n \left[\ln\left(\frac{\mu}{\mu-\lambda_\theta}\right)\right]^{-1}. \end{aligned}$$

We can calculate the average firm size for the different types of firms

$$\bar{n}(\theta) = \sum_{n=1}^{\infty} n P_n(\theta) = \frac{\lambda_\theta}{\mu-\lambda_\theta}\left[\ln\left(\frac{\mu}{\mu-\lambda_\theta}\right)\right]^{-1} \tag{44}$$

i.e. the average size is larger for high θ-firms, not surprisingly. Equation (44) implies that high-θ firms, with high λ_θ, will tend to be larger. Formally, this is natural as firms with higher research efficiency will grow faster and consequently be larger. But, this implication goes against the second half of Fact # 8. That is, the empirical evidence suggests that smaller firms grow faster (or at least not more slowly) than larger firms. Hence, there seems to be a case for allowing for differences in innovative capabilities along the lines outlined in Section 6 rather allowing for differences in research efficiency.

However, studies of the relationship between growth and firm size are typically confined to individual countries. Hence, different λs might be consistent with Gibrat's law if the differences in λ are across countries. Typically Gibrat's law is examined separately across countries. Then our model could be used to explain differences in firm sizes across countries; see, Kumar et al. (1999), and the role of R&D subsidies with international competition.

With differences in research efficiency as modelled in this section, research intensities will be higher for larger firms (*verify this??*), which also contradicts the evidence summarized in Fact # 3.

To complete the analysis with different types of firms, we recalculate the rate of creative destruction. Using the definition of μ with different types of firms, we have that

$$\mu = \sum_{\theta\in\Theta} \eta g(\theta) + \sum_{n=1}^{\infty}\sum_{\theta\in\Theta} n\lambda_\theta M_n(\theta). \tag{45}$$

Inserting (43) into (45), and using that

$$\sum_{n=1}^{\infty}\left(\frac{\lambda_\theta}{\mu}\right)^n = \frac{\mu}{\lambda_\theta+\mu}, \tag{46}$$

we find that

$$\begin{aligned}
\mu &= \sum_{\theta\in\Theta}\left[\eta g(\theta)+\sum_{n=1}^{\infty} n\lambda_\theta \frac{\lambda_\theta^{n-1}\eta g(\theta)}{n\mu^n}\right] \\
&= \sum_{\theta\in\Theta}\eta g(\theta)\frac{\mu}{\mu-\lambda_\theta} \\
&\Rightarrow\ 1=\sum_{\theta\in\Theta}\eta g(\theta)\frac{1}{\mu-\lambda_\theta}. \qquad (47)
\end{aligned}$$

Denote the right hand side of the last equation as $x(\mu)$. Notice that $\lim_{\mu\to\infty} x(\mu)=0$ and $\lim_{\mu\downarrow\lambda^{\max}} x(\mu)=\infty$, where $\lambda^{\max}$ is the largest innovation rate. Hence, since $x(\mu)$ is a continuously declining function in μ, (45) has a solution with $\mu>\lambda^{\max}$ [18]. The inequality $\mu>\lambda^{\max}$ was implicitly assumed to hold in (46).

Finally, we should consider the free entry condition when there are different types of firms. Assuming that the entrants do not know their types prior to sunk the entry costs[19], the free entry condition is given by

$$\begin{aligned}
F &= \sum_{\theta\in\Theta}\int_1^{\infty} v\left(\pi(q),\tilde{c}_\theta(q),r+\mu\right)\ g(\theta)h_\theta(q)\ dq \\
&= \sum_{\theta\in\Theta} v\left(\theta,1,r+\mu\right)g(\theta)\int_1^{\infty}\pi(q)h_\theta(q)dq.
\end{aligned}$$

Appendix E: The Social Planner's Problem

The planner's objective function is the same as the representative consumer's, and can be simplified as

$$U_t=\int_t^{\infty} e^{-\rho(\tau-t)}\left[\ln x_\tau + J_\tau\ \ln q\right]d\tau.$$

Here, the stock of innovations is given by $J_\tau=\int_0^\tau \mu(s)ds$ so that $\dot{J}=\mu=\lambda+\eta$. The labor constraint can be rewritten to yield,

$$\eta=\psi\left(L-x-\bar{c}(\lambda)\right).$$

The current value Hamiltonian is

$$\mathcal{H}=\ln x + J\ln q + \theta\left[\lambda+\psi\left(L-x-\bar{c}(\lambda)\right)\right].$$

The first order conditions are

$$\frac{\partial\mathcal{H}}{\partial x} = 0 \qquad (48)$$

$$\Rightarrow\ \theta=\frac{1}{x\psi} \qquad (49)$$

$$\frac{\partial\mathcal{H}}{\partial\lambda} = 0 \qquad (50)$$

$$\Rightarrow\ \bar{c}(\lambda^*)=\psi^{-1} \qquad (51)$$

[18] We need to assume that $\eta(\lambda_\theta)>0$ for all λ_θ in the population of firms to get this result.

[19] The alternative assumption, that entrants know their type prior to sinking their costs, involves more complicated algebra, and is not considered here.

The co-state variable satisfies

$$\dot{\theta} = \rho\theta - \frac{\partial \mathcal{H}}{\partial J} = \rho\theta - \ln q.$$

In steady state $\dot{\theta} = 0$ which implies $x = \frac{\rho}{\psi \ln q}$. Hence,

$$\eta^* = \psi L - \frac{\rho}{\ln q} - \psi \bar{c}(\lambda^*).$$

References

Aghion, P. and P. Howitt (1992): A model of growth through creative destruction. *Econometrica*, 60, 323–51.

Aghion, P., C. Harris, P. Howitt, and J. Vickers (2000): Competition, imitation and growth with step-by-step innovation. mimeo, Brown University.

Amaral, L., S. Buldyrev, S. Havlin, M. Salinger and H. Stanley (1998): Power law scaling for a system of interacting units with complex internal structure. *Physical Review Letters*, 80, 1385–1388.

Audretsch, D. (1995): *Innovation and Industry Evolution*. Cambridge (U.S.): MIT Press.

Bernard, A., J. Eaton, J. Jensen and S. Kortum (2000): Plants and productivity in international trade. NBER Working paper no. 7688.

Blundell, R., R. Griffith and F. Windmeijer (1999): Individual effects and dynamics in count data models. IFS Working Paper W99/3 (London).

Cabbalero, R. and A. Jaffe (1993): How high are the giants' shoulders: An empirical assessment of knowledge spillovers and creative destruction in a model of economic growth. In O. Blanchard and S. Fischer (eds.), *NBER Macroeconomic Annual 1993*, 15–73. Cambridge, Mass.: MIT Press.

Carroll, G. and T. Hannan (2000): *The Demography of Corporations and Industries*. New Jersey: Princeton University Press.

Caves, R. (1998): Industrial organization and new findings on the turnover and mobility of firms. *Journal of Economic Literature*, 36, 1947–82.

Cockburn, I., R. Henderson and S. Stern (2000): Disentagnling the origins of competitive advantage. Mimeo. (Forthcoming in *Stratetgic Management Journal*, Special Issue on the Evolution of Firm Capabilities).

Coe, D. and E. Helpman (1995): International R&D spillovers. *European Economic Review*, 39, 859–897.

Cohen, W. (1995): Empirical studies of innovative activity. In P. Stoneman (ed.), *Handbook of the Economics of Innovation and Technological Change*, 182–264. Oxford: Blackwell.

Cohen, W. and S. Klepper (1992): The anatomy of industry R&D intensity distributions. *American Economic Review*, 82, 773–99.

— (1996): A reprise of size and R&D. *Economic Journal*, 106, 925–51.

Cohen, W. and R. Levin (1989): Empirical studies of innovation and market structure. *ch. 18 in R. Schmalensee and R. Willig (eds): Handbook of Industrial Organization, II, North Holland Publ. Co..*

Dunne, T., M. Roberts and L. Samuelson (1989): The growth and failure of US manufacturing plants. *Quarterly Journal of Economics*, 104, 671–698.

Eaton, J. and S. Kortum (1999): International technology diffusion: Theory and measurement. *International Economic Review*, 40, 537–70.

Ericson, R. and A. Pakes (1995): Market perfect industry dynamics: A framework for empirical analysis. *Review of Economic Studies*, 62, 53–82.

Evans, D. S. (1987): Tests of alternative theories of firm growth. *Journal of Political Economy*, 95, 657–674.

Geroski, P. (1998): An applied econometrician's view of large company performance. *Review of Industrial Organization*, 13, 271–93.

Geroski, P. A., I. Small and C. F. Walters (1998): Agglomeration economies, technology spillovers and company productivity growth. CEPR Discussion Paper No. 1867, London.

Goel, N. and N. Richter-Dyn (1974): *Stochastic Models in Biology*. New York: Academic Press.

Griliches, Z. (1990): Patent statitics as economic indicators: A survey. *Journal of Economic Litterature 28, 1661-1797..*

— (1995): R&D and productivity: Econometric results and measurement issues. In P. Stoneman (ed.), *Handbook of the Economics of Innovation and Technical Change*. Oxford: Blackwell.

— (1998): *R&D and Productivity: The Econometric Evidence*. Chicago: University of Chicago Press and NBER.

Grossman, G. and E. Helpman (1991): *Innovation and Growth in the Global Economy*. Cambridge, (U.S.): MIT Press.

Hall, B. (1996): The private and social returns to research and development. In B. Smith and C. Barfield (eds.), *Technology, R&D, and the Economy*, 140–162. Washington DC: Brookings Institution and AEI.

Hall, B. (1987): The relationship between firm size and firm growth in the US manufacturing sector. *Journal of Industrial Economics*, 3, 583–606.

Hall, B., Z. Griliches and J. Hausman (1986): Patents and R&D: Is there a lag? *International Economic Review*, 27, 265–83.

Hall, B. and F. Hayashi (1989): Research and development as an investment? NBER Working paper no. 2973.

Henderson, R. (1993): Underinvestment and incompetence as responses to radical innovations: Evidence from the photolithographic industry. *RAND Journal of Economics*, 24(2).

Hopenhayn, H. (1992): Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60, 1127–50.

Howitt, P. (2000): Endogenous growth and cross-country income differences. *American Economic Review*, 90, 829–46.

Ijiri, Y. and H. Simon (1977): *Skew distributions and the sizes of business firms*. Amsterdam: North Holland Publ.Co.

Johnson, N., S. Kotz and A. Kemp (1993): *Univariate Discrete Distributions*. New York: Wiley.

Jones, C. (1995): Models of R&D and endogenous growth. *Journal of Political Economy*, 103, 759–84.

Jovanovic, B. (1982): Selection and evolution of industries. *Econometrica*, 50, 649–70.

— (1993): The diversification of production. *Brooking Papers on Economic Activity*, (Microeconomics), 197–235.

— (2000): Fitness and age. *Journal of Economic Literature*, (forthcoming).

Kleinknecht, A. (1987): Measuring R&D in small firms: How much are we missing? *Journal of Industrial Economics*, 36, 253–6.

Klette, T. (1996): R&D, scope economies and plant performance. *RAND Journal of Economics*, 27, 502–22.

Klette, T. and Z. Griliches (2000): Empirical patterns of firm growth and R&D-investment: A quality ladder model interpretation. *Economic Journal*, 110, 363–87.

Klette, T. and F. Johansen (1998): Accumulation of R&D-capital and dynamic firm performance: a not-so-fixed effect model. *Annales d'Economie et de Statistique*, (49/50), 389–419.

Kortum, S. (1997): Research, patenting, and technological change. *Econometrica*, 65, 1389–1419.

Kumar, K., R. Rajan and L. Zingales (1999): What determines firm size? NBER Working Paper No. 7208.

Langlois, R. and P. Robertson (1995): *Firms, Markets and Economic Change*. New York: Routledge.

Lucas, R. (1978): On the size distribution of business firms. *Bell Journal of Economics*, 9, 508–523.

Lucas, R. and E. Prescott (1971): Investment under uncertainty. *Econometrica*, 39, 659–681.

Mitchell, M. (2000): The scope and organization of production: Firm dynamics over the learning curve. *RAND Journal of Economics*, 31, 180–205.

Nelson, R. (1988): Modelling the connections in the cross section between technical progress and R&D intensity. *RAND Journal of Economics*, 19, 478–85.

Pakes, A. and R. Ericson (1998): Empirical implications of alternative models of firm dynamics. *Journal of Economic Theory*, 79, 1–45.

Penrose, E. (1959): *The Theory of the Growth of the Firm*. Oxford: Basil Blackwell.

Peretto, P. (1998): Cost reduction, entry, and the interdependence of market structure and economic growth. *Journal of Monetary Economics*, (forthcoming).

Romer, P. (1990): Are nonconvexities important for understanding growth? *AER Papers and Proceedings 80, 97-103*.

Schmalensee, R. (1989): Inter-industry studies of structure and performance. In R. Schmalensee and R. Willig (eds.), *Handbook of Industrial Organization, Vol. II*. Amsterdam: North Holland Publ. Co.

Scott, J. (1984): Firm versus industry variability in R&D intensity. In Z. Griliches (ed.), *R&D, Patents and Productivity*. Chicago: University of Chicago Press and NBER.

Stanley, M. H., S. V. Buldyre, S. Havlin, R. N. Mantegna, M. A. Salinger and H. E. Stanley (1995): Zipf plots and the size distribution of firms. *Economics Letters*, 49, 453–7.

Sutton, J. (1997): Gibrat's legacy. *Journal of Economic Literature*, 35, 40–59.

— (1998): *Market Structure and Technology*. Cambridge, Mass.: MIT Press.

— (2000): The variance of firm growth rates: The 'scaling' puzzle. Mimeo, London School of Economics.

Taylor, H. and S. Karlin (1998): *An Introduction to Stochastic Modelling, Third Edition*. San Diego: Academic Press.

Thompson, P. (1996): Technological opportunity and the growth of knowledge: A Schumpeterian approach to measurement. *Journal of Evolutionary Economics*, 6, 77–97.

Uzawa, H. (1969): Time preference and the Penrose effect in a two-class model of economic growth. *Journal of Political Economy*, 77, 628–52.