

NBER WORKING PAPER SERIES

INNOVATING FIRMS AND AGGREGATE INNOVATION

Tor Jakob Klette
Samuel Kortum

Working Paper 8819
<http://www.nber.org/papers/w8819>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2002

We thank John Dagsvik, Jonathan Eaton, Charles Jones, Boyan Jovanovic, Kalle Moene, Jarle Møen and Atle Seierstad for helpful comments. This paper has also benefitted from comments by participants at seminars at Boston University, Brandeis University, Georgetown University, University of Minnesota, University of Oslo, Pennsylvania State University, the Norwegian School of Economics and Business Administration, Royal Institute of Technology Stockholm, the NBER Summer Institute 2000 and 2001, and the International Conference on Technological Policy and Innovation, Paris, 20-22 November 2000. Klette gratefully acknowledges support from the Research Council of Norway and Kortum from the National Science Foundation. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2002 by Tor Jakob Klette and Samuel Kortum. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Innovating Firms and Aggregate Innovation
Tor Jakob Klette and Samuel Kortum
NBER Working Paper No. 8819
March 2002
JEL No. L11, O31

ABSTRACT

We develop a parsimonious model of innovating firms rich enough to confront firm-level evidence. It captures the dynamic behavior of individual heterogeneous firms, describes the evolution of an industry with simultaneous entry and exit, and delivers a general equilibrium model of technological change. While unifying the theoretical analysis of firms, industries, and the aggregate economy, the model yields insights into empirical work on innovating firms. It accounts for the persistence over time of firms' R&D investment, the concentration of R&D among incumbent firms, and the link between R&D and patenting. Furthermore, it explains why R&D as a fraction of revenues is strongly related to firm productivity yet largely unrelated to firm size or growth.

Tor Jakob Klette
Department of Economics
University of Oslo
and CEPR
t.j.klette@econ.uio.no
<http://folk.uio.no/torjk>

Samuel Kortum
Department of Economics
University of Minnesota,
Federal Reserve Bank of Minneapolis,
and NBER
kortum@econ.umn.edu
<http://www.econ.umn.edu/~kortum>

1 Introduction

Endogenous growth theory has sketched the bare bones of an aggregate model of technological change.¹ Firm-level studies of R&D, productivity, patenting, and firm growth could add flesh to these bones.² So far they have not.

Exploiting firm-level findings for this purpose raises difficult questions. For instance, studies of how R&D affects productivity and patenting do not address why firms conduct R&D on very different scales in the first place. What are the sources of this heterogeneity in innovative effort across firms, and why is it so persistent? Why do some firms prosper while doing little or no R&D?³ Is the estimated relationship between R&D and productivity causal, or no more than a correlation? Does a firm's productivity adequately measure innovative performance given that many innovations appear in the form of new products? Is patenting a superior indicator of performance?

What about firm growth? Empirical studies of growth, entry, exit, and the firm size distribution should complement the literature on R&D, productivity and patenting.⁴ In fact, these two lines of research have developed independently. They deserve an integrated treatment.

We construct a model of innovating firms rich enough to match stylized facts from firm-level studies yet simple enough to yield analytical results both for the dynamics of individual firms and for the behavior of the economy as a whole. Underlying the model is a Poisson process for a firm's innovations with an arrival rate a function of its current R&D and knowledge generated by its past R&D. This specification of the innovation process is consistent with the empirical relationship between patents and R&D at the firm level. We derive the optimal R&D-investment rule that, together with the innovation function, delivers firm growth rates independent of size, i.e., Gibrat's law. The stochastic process

¹The seminal papers are Aghion and Howitt (1992), Grossman and Helpman (1991), and Romer (1990).

²Much of this literature stems from the work of Zvi Griliches and his co-workers. See Griliches (1990, 1995).

³Cohen and Klepper (1996) have highlighted these puzzles.

⁴The classic reference is Ijiri and Simon (1977), but this literature on firm growth has recently been revived by Sutton (1998) and Amaral et al. (1998) among others.

for innovation leads to heterogeneity in the size of firms. The R&D rule, together with size heterogeneity, captures the observed persistent differences in R&D across firms.⁵

The foundation of our approach is related Penrose's (1959) theory of the growth of the "innovating, multiproduct, 'flesh-and-blood' organizations that businessmen call firms" (p. 13). A firm of any size can expand into new markets, but in any period such growth depends on the firm's resources. While Penrose stresses managerial and entrepreneurial services, our model emphasizes knowledge resources. In our formulation, a firm of any size adds new products by innovating, but in any period its likelihood of success depends on its knowledge capital accumulated through past product innovations. Although firms grow by innovating in products new to them, the economy grows as innovations raise the quality of a given set of products.

The Schumpeterian force of creative destruction is pervasive in our analysis. Firms' innovative successes always come at the expense of competitors. A firm may be driven out of business when hit by a series of destructive shocks. The model predicts that exit is, in fact, the fate of any firm. While each firm follows a stochastic life cycle, industry equilibrium generally involves simultaneous entry and exit, with a stable, skewed firm size distribution. In this sense our model captures some of the richness of the framework developed by Ericson and Pakes (1995). While their model is suitable for the analysis of industries with a few competitors, we get much further analytically by taking each firm to be small relative to the industry. Solving our model in general equilibrium extends the work of Aghion and Howitt (1992) and Grossman and Helpman (1991) by incorporating the contribution of incumbent research firms to aggregate innovation.

The paper begins by summarizing the firm-level findings. The list of stylized facts in section 2 is 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 implications for the dynamics of a firm. Section 5 discusses heterogeneity in firms' innovative capability.

⁵The work of Thompson (1996, 2001), Peretto (1999), and Klette and Griliches (2000) has a similar motivation. Each attempts to bring realistic features of firms into an aggregate model of technological change. Our approach differs from these earlier contributions by building in the multiproduct nature of firms. The resulting model is particularly tractable.

Section 6 moves to the industry level, describing entry and deriving the size distribution of firms. Section 7 solves the model in general equilibrium. In the light of our theory, we return to the stylized facts that motivated our analysis in section 8 and offer some suggestions for future work.

2 Evidence on Innovating Firms

This section presents a comprehensive list of empirical regularities or stylized facts which have emerged from a large number of studies using 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 empirical regularities that are robust and economically significant. Following Cohen and Levin (1989) and Schmalensee (1989), we have ignored findings (even when sometimes statistically significant) if they appear fragile and not robust. The stylized facts on which we focus are largely summarized in surveys by others, including Caves (1998), Cohen (1995), Griliches (1990, 1998, 2000) and Sutton (1997). Appendix A contains a discussion of the stylized facts with more detailed references to our sources.

Our first two stylized facts summarize the relationship between firm R&D (measured as expenditure, as a stock, or as a fraction of firm revenues) and innovations, measured in terms of patents or productivity. A large number of studies have documented a significant positive relationship between productivity and R&D, but the relationship is only robust in the cross-firm dimension.

Stylized Fact 1 *Productivity and R&D across firms are positively related, while productivity growth is not strongly related to firm R&D.*

The evidence has shown that R&D is highly correlated with patenting in cross sections of firms, and firms that increase their R&D produce more patents. The empirical studies also point to the shape of these relationships:

Stylized Fact 2 *Patents vary proportionally with R&D across firms, while there are diminishing returns to R&D in the longitudinal dimension.*

The empirical evidence on patterns of R&D investment is presented in stylized facts 3 through 6. There is a large literature studying whether large firms are more R&D intensive (i.e., devote a higher fraction of revenues to R&D) than small firms. At least among R&D-reporting firms, the evidence suggests that R&D increases in proportion to sales:

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

Many firms report no formal R&D activity even in high-tech industries, and R&D intensity varies substantially across firms within narrowly defined industries:

Stylized Fact 4 *The distribution of R&D intensity is highly skewed, and a considerable fraction of firms report zero R&D.*

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

Stylized Fact 6 *Firm R&D investment follows essentially a geometric random walk.*

Our last set of stylized facts considers entry, exit, growth, and the size distribution of firms:

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

Stylized Fact 8 *Smaller firms have a lower probability of survival, but those that survive tend to grow faster than larger firms. Among larger firms, growth rates are unrelated to past growth or to firm size.*

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

Stylized Fact 10 *Younger firms have a higher probability of exiting, but those that survive tend to grow faster than older firms. The market share of an entering cohort of firms generally declines as it ages.*

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

3 The Model of an Innovating Firm

We begin our presentation of the theoretical model by describing the innovation process for an individual firm. A firm operates in an economy with 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 sometimes refer to a market as a good.

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 the portfolio of goods that it produces. Since goods all generate the same flow of revenue and profit to the producer, we need only keep track of the number of distinct goods $n = 1, 2, 3, \dots$ that a firm produces. A firm with n goods has revenues equal to n and profits of $\bar{\pi}n$.

To add new goods to its portfolio a firm invests in innovative effort, which we term 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 so that the innovating firm takes over production of the good. Expenditures on R&D could yield an innovation relevant to any good with equal probability; i.e., the product to which it applies is drawn from the uniform distribution on $[0, 1]$. Since a firm is infinitesimal relative to the continuum of goods, we can ignore the possibility that it improves on a good it is currently producing.

An incumbent producer loses a good from its portfolio when some other firm comes up with an innovation applying to that good. The Poisson hazard rate per good is $\mu > 0$. The parameter μ , which we call the intensity of creative destruction, is taken as given by each firm and is taken as constant over time. In section 6 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, as explained below. The simplest interpretation of these assumptions about demand and competition is the quality ladder model of Grossman and Helpman (1991). They also assume a unit measure of differentiated goods. Aggregate

expenditure is the numeraire and preferences over the differentiated products are Cobb-Douglas so that expenditure per good is unity. Innovations come in the form of quality improvements, allowing the most recent innovator, through Bertrand competition, to capture the market for a particular good. Innovations raise quality by a fixed factor $q > 1$ so that $\bar{\pi} = 1 - q^{-1}$. The aggregate rate of innovation is also the rate of creative destruction, μ . We now turn to our specification of the innovation process, which is where our model deviates substantially from the existing literature.

3.1 The Innovation Technology

We assume that a firm's innovation rate depends on both its investment in R&D, denoted by R , and its knowledge capital. The firm's knowledge capital stands for all the skills, techniques, and know-how that it draws upon as it attempts to innovate. We view knowledge capital as a crucial element of what Penrose (1959) refers to as the internal resources of the firm that can be devoted to expansion. In her analysis, these resources evolve as the firm grows so that, while they constrain growth in any period, they have no implication for an optimal size of firm. To capture the rather abstract concept of knowledge capital in the simplest way, we assume that it is summarized by n , the firm's past innovations that have not yet been superseded.⁶

With knowledge capital measured by n , the innovation production function is

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

We assume G is (i) strictly increasing in R , (ii) strictly concave in R , (iii) strictly increasing in n , and (iv) homogeneous of degree one in R and n . The second restriction captures decreasing returns to expanding research effort, allowing us to tie down the research investment of an individual firm and limiting firm growth in any period. The third captures the idea that a firm's knowledge capital facilitates innovation. The last restriction

⁶Although a crude measure of a complex entity like knowledge capital, n does capture the idea that past innovations provide fodder for new innovations. Scope economies in the development of related products has been emphasized by Jovanovic (1993) in a static model of firm formation. A related interpretation is that n reflects differences in the quality of firms' laboratories. The set of currently commercially viable innovations having come out of a lab would then be our proxy for the lab's quality.

neutralizes the effect of firm size on the innovation process: A firm that is twice as large will expect to innovate twice as fast by investing twice as much in R&D.⁷

In anticipation of working out the firm's optimal R&D policy, it is convenient to rewrite the innovation production function in the form of a cost function. Given the assumption made above, the firm's R&D costs are a homogenous function of the Poisson arrival rate of innovations I and its stock of knowledge n ,

$$R = C(I, n) = nc(I/n), \quad (2)$$

where $c(x) = C(x, 1)$. The intensive form of the cost function $c(x)$ is increasing and strictly convex in x . To avoid a number of uninteresting qualifications in the analysis below we assume that $c(0) = 0$, $c(x)$ is twice differentiable for $x \geq 0$, and $[\bar{\pi} - c(\mu)]/r \leq c'(x) < \bar{\pi}/(r + \mu)$ for $x \in [0, \mu]$.

3.2 The Firm's R&D Policy

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 R&D it influences the Poisson hazard I of becoming a firm of size $n + 1$. We assume that a firm of size n chooses an innovation policy $I(n)$ [or equivalently an R&D policy $R(n) = C(I(n), n)$] to maximize its expected present value $V(n)$, given a fixed interest rate r . We treat a firm in state $n = 0$ as having permanently exited, so that $V(0) = 0$.

The firm's Bellman equation is

$$rV(n) = \max_I \{ \bar{\pi}n - C(I, n) + I[V(n+1) - V(n)] - \mu n[V(n) - V(n-1)] \}. \quad (3)$$

It is easy to verify that the solution is

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

⁷A similar innovation production function has been used by Hall and Hayashi (1989) and Klette (1996). Our justification was based on knowledge capital as an input to the innovation process. Another justification comes by analogy with accumulation of physical capital with costs of adjustment. See Lucas (1967) and Uzawa (1969), who both use a formulation like (1). According to this interpretation, we have simply imposed convex costs on the firm as it adjusts its stock of knowledge capital.

where v and λ solve

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

We refer to $\lambda = I(n)/n$ as the firm's innovation intensity. Note that the innovation intensity is independent of firm size. In Appendix B we show that innovation intensity satisfies $0 < \lambda \leq \mu$. Furthermore, innovation intensity is increasing in $\bar{\pi}$, is decreasing in r and μ , and is also decreasing in an upward shift of c' .

The firm's R&D policy is $R(n) = C(I(n), n) = nc(\lambda)$. A firm scales up its R&D expenditure in proportion to its knowledge capital.⁸ Research intensity, i.e., the fraction of firm revenue spent on R&D [$R/n = c(\lambda)$], is independent of firm size, in line with stylized fact 3. Hence, the total amount of R&D performed by a group of firms depends on the overall size of the group but not on how firm size is distributed within the group. Similarly, the expected number of innovations made by the group will also be unrelated to the size distribution.

The value of a firm of any size can be decomposed into the value of each of its products, v . We can further decompose the value of a single product. If the firm simply marketed the product, without exploiting the associated knowledge capital to engage in new research, the 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 we interpret as the value of a unit of knowledge capital. It is therefore possible to think of a firm of size n as consisting of a production division whose value is v_1n and a research division with value v_2n .

It is a property of the model that a firm of size n can be analyzed in terms of its individual products or in terms of its marketing and research divisions. As stated, there

⁸Using the firm's R&D policy, we can link our concept of the firm's stock of knowledge n to the measure proposed by Griliches (1979). For Griliches, the stock of knowledge was the discounted sum of past R&D by the firm, which he denoted K . According to our theory, the expected value of our proposed measure of the stock of knowledge conditional on past R&D expenditures is (up to a constant) equal to K as well: $E[n_t | \{R_s\}_{t_0}^t] = E \int_{t_0}^t e^{-\mu(t-s)} I_s ds = a \int_{t_0}^t e^{-\mu(t-s)} R_s ds = aK_t$, where t_0 is the date when the firm was born, $a = G(1, 1/c(\lambda))$. Note that the appropriate depreciation rate on past R&D is the rate of creative destruction.

are neither gains nor losses to merging or breaking up firms. Since we wish to focus on the internal growth of firms, we will assume that there is some cost either to merging two firms or to spinning off an individual product from an existing firm. Any such cost, even if tiny, will eliminate all mergers and divestitures. Such a cost is quite plausible due to the difficulty of making transactions involving technology. Exploring the exact nature of this cost is beyond the scope of our paper, but we recognize that until we do we cannot address interesting questions related to spin-offs, firm disintegration, mergers, or acquisitions.

4 The Life Cycle of a Firm

We can now characterize the growth process for an individual firm, having solved for its innovation intensity λ and taking as given the intensity of creative destruction μ . Consider a firm of size n . At any instant of time it will 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 exits if it loses its product; i.e., $n = 0$ is an absorbing state.

Let $p_n(t; n_0)$ denote the probability that a firm is size n at date t given that it was 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 C, firm evolution is described by the following system of equations:

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

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

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

which reflects that 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 Appendix C 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 distribution $p_n(t; n_0)$ for each date t and conditional on any size n_0 at date 0 (see Appendix C). To get some intuition about this solution, we consider a few properties of the implied probability distribution.

4.1 New Firms

We assume that firms begin with a single product. To track the size distribution at date t of a firm entering at date 0 we set $n_0 = 1$. (This analysis also applies to the subsequent evolution of any firm that at some date reaches a size of $n = 1$, whether or not it just entered.) The pgf then yields

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

where

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

This last term satisfies $\gamma(0) = 0$, $\gamma'(t) > 0$, and $\lim_{t \rightarrow \infty} \gamma(t) = \lambda/\mu$. For the case of $\lambda = \mu$ we can use l'Hopital's rule to get $\gamma(t) = \mu t / (1 + \mu t)$. Notice that in any case $\lim_{t \rightarrow \infty} p_0(t; 1) = 1$; i.e., with the passage of time, the probability of exit approaches one.

Conditioning on survival, we get the simple geometric distribution (shifted 1 to the right) for the size of the firm at date t ,

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

The parameter of this distribution is $\gamma(t)$; hence, the distribution grows stochastically larger over time. As time passes, and conditional on the firm surviving, there is an increasingly high probability that the firm has become very large.

4.2 Firm Age

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

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

Expected age is decreasing in the intensity of creative destruction, μ , holding λ fixed. Holding μ fixed, the expected age of a firm increases in λ , becoming infinite for $\lambda = \mu$.

The hazard rate of exit is

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

The last equality shows that the 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 the initial hazard rate of exit is μ . The hazard rate declines steadily with age and approaches $\mu - \lambda$ for a very old firm. As a firm ages, implying that it has survived, it tends to get bigger and is therefore less likely to exit, in line with stylized fact 10.⁹

The expected size of a firm, conditional on survival, is

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

⁹The probability that a firm of size n exits within t periods is $p_0(t; 1)^n$. Hence, larger firms have a lower hazard of exiting. As in Hopenhayn (1992), our model implies that the reason firm age matters for exit is because it is a proxy for firm size. Conditional on size, younger firms are no more likely to exit, according to these models. In the data, there appears to be an independent role for age in firm exit rates. See Dunne et al. (1988).

which increases with age, a . Consider a cohort of firms all entering at the same date. As the cohort ages, the number of firms in the cohort diminishes. On the other hand, the survivors are on average bigger. Let m denote the initial number of firms in the cohort. Then the expected number of products produced by it (i.e., the total revenues it generates) at age a is

$$m \sum_{n=0}^{\infty} n p_n(a; 1) = m \frac{1 - \frac{\mu}{\lambda} \gamma(a)}{1 - \gamma(a)}. \quad (10)$$

If $\mu > \lambda$, the share of the cohort in the overall market declines steadily with a , as in stylized fact 10.

4.3 Large Firms

A firm of size $n_0 > 1$ at date 0 will evolve as if it consists of n_0 divisions of 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 firm starting with a single product. In this sense, our formal analysis of size 1 firms carries over to the behavior of each division.

4.4 Firm Growth

We can derive the moments of firm growth directly from the pgf, as shown in Appendix C. Let the random variable N_t denote the size of a firm at date t . Then 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

$$E[G_t | N_0 = n_0] = e^{-(\mu - \lambda)t} - 1, \quad (11)$$

which is independent of initial size. Taking the limit of $E[G_t | N_0 = n_0]/t$ as t approaches 0, we get the common expected instantaneous rate of growth $-(\mu - \lambda)$. If $\mu > \lambda$, any given firm will tend to decline relative to the economy. As we show in section 7, the economy may grow fast enough so that firms on average experience positive growth even

as they decline relative to the entire economy (our normalization of unit revenue comes from choosing aggregate expenditure as the numeraire).

The variance of firm growth given initial size is

$$\text{Var} [G_t | N_0 = n_0] = \frac{\lambda + \mu}{n_0 (\mu - \lambda)} e^{-(\mu - \lambda)t} (1 - e^{-(\mu - \lambda)t}), \quad (12)$$

which declines in initial firm size, in line with stylized fact 9.¹⁰ 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.¹¹

In the derivations above we have included firms that exit during the period (and whose growth is therefore -1). 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

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

which is a decreasing 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 to date t is close to 1 anyway, so Gibrat's law will be a very good approximation, in line with stylized fact 8.

5 Innovators and Imitators

We showed above that a firm's research intensity, its research expenditure as a fraction of revenue, 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, none of which is captured by the model.

¹⁰For the case of $\lambda = \mu$, the variance expression reduces to $2\mu t/n_0$. For any $\lambda \leq \mu$, the limit as t approaches 0 of $\text{Var} [G_t | N_0 = n_0] / t$ is simply $(\lambda + \mu) / n_0$.

¹¹It is well known since Hymer and Pashigian (1962) that the variance of firm growth does not fall as steeply as the inverse of firm size. More recently Amaral et al. (1998) and Sutton (2001) have developed models that more closely match the actual rate of decline.

Accounting for heterogeneity in research intensity is challenging because we want to avoid 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 stylized fact 3). To avoid this implication we seek to unhinge the research process from the process of revenue growth while introducing another dimension of firm heterogeneity.

Empirical research in economics, sociology and management science has emphasized the large and highly persistent differences in innovative strategies across firms¹². 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 endow each firm with a profit level $0 < \pi < 1$. A firm making major 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, as in Nelson (1988). In the quality ladders model, these differences in profits would reflect differences in the inventive steps q of different firms, $\pi = 1 - q^{-1}$. We pursue this interpretation below.

A firm's type π remains fixed throughout the life of the firm. The type not only affects the firm's flow of profits from an innovation but also its cost of doing research. To obtain independence between innovative ability and firm growth, we assume that the cost of making larger innovations rises in proportion to the greater profitability of such innovations. That is, a firm of type π has a research cost function $C_\pi(I, n) = \frac{\pi}{\bar{\pi}}C(I, n)$. For a firm of type $\pi = \bar{\pi}$, this specification reduces to the research cost function $C(I, n)$ used above.

Returning to a firm's R&D policy, we need to modify the analysis only slightly. The intensive form of the cost function becomes $c_\pi(I/n) = \frac{\pi}{\bar{\pi}}c(I/n)$. The solution to the Bellman's equation (4) remains unchanged for a firm of type $\bar{\pi}$. More generally, the value per product of a firm of type π is simply $v_\pi = \frac{\pi}{\bar{\pi}}v$. Innovation intensity λ is the same for firms of any type; i.e., firm growth is unrelated to inventiveness, π . On the other hand, R&D intensity for a firm of type π is $C_\pi(I, n)/n = \frac{\pi}{\bar{\pi}}c(I/n) = \frac{\pi}{\bar{\pi}}c(\lambda)$. Thus, heterogeneity

¹²See e.g. Henderson (1993), Cohen (1995), Langlois and Robertson (1995), Klepper (1996), Carroll and Hannan (2000), Cockburn et al. (2000) and Jovanovic (2001). Geroski et al. (1997) presents a contradictory view (see their stylized fact 5).

in π carries over to R&D intensity, with the more innovative firms being more research intensive.

6 Entry and Industry Evolution

In this section we examine the dynamic process characterizing an industry with many competing firms. 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 state of the industry in terms of the measure of firms of each size. There is no randomness at the industry level.

We denote the measure of firms in the industry with n products at date t by $M_n(t)$. The total measure of firms in the industry is $M(t) = \sum_{n=1}^{\infty} M_n(t)$. 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

$$\mu = \eta + \lambda. \tag{13}$$

We now turn to the determination of η .

6.1 Firm Entry

As mentioned above, we assume that entrants begin with a single product. Their type π is drawn from a distribution $\Phi(\pi)$, whose mean value is $\bar{\pi}$. There is a fixed cost of entry F and a mass of potential entrants. Entrants do not know their type π when they pay the entry cost, but they know the distribution of types. Under these conditions, if there is active entry we have the condition $F = E[v_\pi] = v$.

When there is active entry, i.e., $\eta > 0$, the optimal R&D policy of incumbents pins down their innovation intensity:

$$c'(\lambda) = F, \tag{14}$$

or else, if $c'(0) > F$ then $\lambda = 0$. 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, (14) shows that the innovation intensity of incumbents is unrelated to demand side incentives, such as the mean 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 (4). Rearranging it under the assumption that there is active entry, i.e., $F = v$, and using (13), we get

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

Of course, for this solution to be valid we require F to be small enough so that the equation above yields $\eta > 0$, in which case it follows that $\lambda < \mu$.

Alternatively, for F large enough, there will be no entry. In this case we can set $\mu = \lambda$ in the solution to the Bellman's equation (4), which yields λ as the solution to $c'(\lambda) = v = [\bar{\pi} - c(\lambda)]/r$ and $\eta = 0$.

6.2 The Size Distribution

We now have expressions for η , λ ; hence, $\mu = \eta + \lambda$. These parameters are all that matter for analyzing the size distribution of firms. The state of the industry is summarized by

the measure of firms with $1, 2, 3, \dots$ products.¹³

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 $n \geq 2$,

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

For $n = 1$ we have

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

In the case of no entry, $\eta = 0$, the mass of firms of any particular size will not settle down. We can still study the evolution of the industry, however, using the analysis of section 4. Suppose at date 0 the industry consists of a mass $M(0)$ of firms, all of size 1. By date t , there will be a mass $M(t) = [1 - p_0(t; 1)]M(0)$, among whom the size distribution will be geometric with a parameter $\mu t / (1 + \mu t)$. Thus, without entry, the size distribution becomes ever more skewed and the industry becomes ever more concentrated.

In the case of positive entry, $\eta > 0$, we show below that the industry will converge to a steady state with a constant mass of firms and fixed size distribution. To solve for this steady state, set all the time derivatives to zero in (15) and (16). In Appendix D, we show that the solution is

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

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

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

¹³Since firms of any type π choose the same innovation intensity λ (and thus follow the same stochastic growth process) the distribution of types Φ among entrants carries over to the distribution of types among firms of any particular size n .

In the 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.

The steady state size distribution, $P_n = M_n/M$, can be written as

$$P_n = \frac{\left(\frac{1}{1+\theta}\right)^n}{n \ln\left(\frac{1+\theta}{\theta}\right)}, \quad n = 1, 2, \dots$$

This is the well-known logarithmic distribution, as discussed in Johnson et al. (1993). The distribution is highly skewed, in line with stylized fact 7. The mean of the distribution, i.e., the average number of products per firm, is $\theta^{-1}/\ln(1 + \theta^{-1})$, which is decreasing in θ . The logarithmic distribution is discussed in the context of firm sizes by Ijiri and Simon (1977).

We can analyze industry dynamics outside of the steady state, and establish convergence to the 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,

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

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

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

This result suggests that the first summation in (19) disappears as $t \rightarrow \infty$, which is argued more formally in Appendix D. If $\eta > 0$, the second term converges to $\frac{\eta}{n\lambda}(\lambda/\mu)^n = M_n$. Thus, with entry, 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).¹⁴

¹⁴Notice that λ and μ are constants even out of the steady state, as shown in section 6.1.

7 General Equilibrium and Growth

In the introduction, we asserted that studies of innovating firms could add flesh to the bones of endogenous growth theory. To evaluate that assertion, in this section we conduct our analysis in a general equilibrium setting. The resulting model of growth differs from its predecessors in describing the dynamics, the heterogeneity, and the innovative output of incumbent research firms. Capturing the role of incumbent research firms is not only important for describing reality, as only a small fraction of R&D is actually performed by entrants, but is also critical for evaluating policies that will often affect incumbents differently than entrants.

By closing the model in general equilibrium, we can also clarify our seemingly arbitrary assumptions made above about the size of the market for a good, the constancy of the interest rate, and the inputs entering the innovation technology. We therefore make explicit the close link between the structure of our model and the quality ladder model in Grossman and Helpman (1991, chapter 4).¹⁵

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 y , after an innovation it rises to qy . The profit flow from making such an innovation is $\pi = 1 - q^{-1}$. A firm of type π is thus a firm taking innovative steps of size $q = (1 - \pi)^{-1}$. By chance, firms of various types will make innovations related to any particular good j . We let $q(j, k)$ denote the step size of innovation k for good j .

The representative consumer has preferences over the continuum of goods:

$$U_t = \int_t^\infty e^{-\rho(\tau-t)} \ln C_\tau d\tau,$$

$$\ln C_\tau = \left\{ \int_0^1 \ln [x_\tau(j) z_\tau(j)] dj \right\},$$

¹⁵Like the model in Grossman and Helpman (1991), our model suffers from the empirical shortcomings pointed out by Jones (1995). In particular, it cannot account for the observed upward trend in aggregate R&D while at the same time accounting for the lack of any observed upward trend in productivity growth. We believe that the model could be modified to address these issues, along the lines of either Kortum (1997), Eaton and Kortum (1999), or Howitt (2000). We have not introduced such a modification here because we do not think the firm-level data have anything to say about which one, if any, is most appropriate.

where ρ is the discount rate while $x_\tau(j)$ is the quantity and $z_\tau(j)$ is the quality of consumption good j at date τ . The quality level of good j , $z_\tau(j)$, develops through a series of $J_\tau(j)$ product improving innovations

$$z_\tau(j) = \prod_{k=1}^{J_\tau(j)} q(j, k).$$

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 $r = \rho$, delivering our assumption of a constant discount rate faced by the firms.

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. Let $q_\tau(j) = q(j, J_\tau(j))$ denote the inventive step for the *last* innovation made prior to date τ . It follows that the price per (physical) unit of good j is

$$p_\tau(j) = w q_\tau(j).$$

Given unit expenditure per product, the quantity produced of good j is

$$x_\tau(j) = \frac{1}{p_\tau(j)} = \frac{1}{w q_\tau(j)}.$$

The profit flow from product j at time τ is $\pi_\tau(j) = [p_\tau(j) - w] x_\tau(j) = 1 - [q_\tau(j)]^{-1}$, as asserted above.

We can aggregate across goods by exploiting the law of large numbers while noting that $J_\tau(j)$ is distributed Poisson with parameter $\mu\tau$ (if the innovation process begins at date 0):

$$\ln C_\tau = (\mu\tau - 1) E[\ln q] - \ln w,$$

where

$$E[\ln q] = \int \ln q d\Psi(q)$$

and $\Psi(q)$ is the cumulative distribution function of the step sizes.¹⁶ The growth rate of the market economy is $d \ln C_\tau / d\tau = \mu E [\ln q]$. To complete the analysis of growth for this market economy, we derive μ in terms of parameters of the model.

There is a fixed endowment of L workers. Workers can 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. An implication is that the wage w for all activities is the same. The case of no entry is a simple extension.

We derive labor demand for each of the three activities starting with production work. Choosing units of output so that production workers have productivity one, labor demand for producing good j is $l_X(j) = x(j) = [q(j) w]^{-1}$. Total demand for production workers is consequently

$$L_X = \int_0^1 l_X(j) dj = E [q^{-1}] / w.$$

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

$$L_R(n) = R(n)/w = nc(\lambda)/w = n l_R(\lambda),$$

which implicitly defines the number of researchers per good as $l_R(\lambda) = c(\lambda)/w$. Aggregating over firms of different sizes, we get

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

Finally, a worker trying to launch a firm gets an idea for an innovation at a Poisson arrival rate h . Thus the cost of entry is the opportunity cost of the expected time until

¹⁶Replacing average values across the infinite set of goods by the expected value, we get $\int_0^1 \ln x_\tau(j) dj = E [\ln x_\tau]$ and $\int_0^1 \left[\sum_{k=1}^{J\tau(j)} \ln q_k(j) \right] dj = (\mu\tau) E [\ln q_\tau]$. The expectations are over the distribution of q , which is closely linked to the distribution of π introduced in section 5: $\Phi(\pi) = \Psi([1 - \pi]^{-1})$. As noted earlier, since the stochastic process of survival and growth is the same across firms of any type, the distribution of types among incumbents is the same as the distribution among entrants.

the first innovation, $F = w/h$. The value of a start-up firm (with one product) is still v . The equilibrium condition (assuming active entry) becomes $w = hv$ or $v = w/h$. Since $c'(\lambda) = v$ it follows that $l'_R(\lambda) = h^{-1}$. Thus, the innovation rate of incumbents is tied down by the productivity h of entrants. The equation for the entry rate becomes

$$\eta = \frac{\bar{\pi} - c(\lambda)}{w/h} - r = \frac{(1 - E[q^{-1}])h}{w} - hl_R(\lambda) - \rho.$$

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

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 $w = \frac{h}{Lh+\rho}$. Substituting this value of the wage into our expression for the rate of entry yields

$$\eta = (1 - E[q^{-1}])Lh - \frac{\rho}{q} - hl_R(\lambda).$$

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

$$\mu = \eta + \lambda = (1 - E[q^{-1}])Lh - \frac{\rho}{q} - hl_R(\lambda) + \lambda,$$

where, as noted above, λ solves $l'_R(\lambda) = h^{-1}$. The first two terms of the expression for μ are identical to Grossman and Helpman (1991).¹⁷

The term $\lambda - hl_R(\lambda)$ is new. It represents the contribution of incumbents to aggregate innovation less the innovation that would have occurred if the researchers employed by incumbent firms were instead employed in start-ups. The term is positive if the knowledge capital accumulated by incumbent firms is a productive resource in generating new innovations. In equilibrium, the innovative productivity of the last researcher at an incumbent

¹⁷As a special case of the model, the remaining two terms are zero and our model collapses to Grossman and Helpman (1991), with all research performed by entrants. In particular, assume that knowledge capital does not enter the innovation production function: $I = G(R, n) = h_R R/w$, with h_R representing the productivity of researchers at incumbent firms. If $h_R < h$ then incumbents will choose $\lambda = 0$ (only entrants will do research) and $\mu = (1 - E[q^{-1}])Lh - \frac{\rho}{q}$.

firm is equated to her productivity at a start-up. But the inframarginal researchers are more productive at incumbent firms than at start-up firms.

In Appendix E we examine whether the market equilibrium is efficient. In the special case with no variation in q across firms, we replicate the welfare conclusion by Grossman and Helpman (1991): 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 fast. However, this welfare conclusion is sensitive to deviations from the assumed Cobb-Douglas preferences, as pointed out by Li (2001), who examines a quality ladder model with CES preferences.

8 Discussion

Having laid out the theory, we ask, what are its implications for firm-level studies of R&D, productivity, and patenting? What has our analysis contributed to the existing theories of firm and industry dynamics? Where do we go from here? We conclude by addressing these questions in turn.

8.1 Interpreting Firm-Level Indicators of Innovation

We have commented along the way when our model fit one of the stylized facts listed in section 2. Here we return to the questions that motivated our work: Why does R&D vary so much across firms, and how do these differences in research input show up in measures of innovative output?

8.1.1 R&D Investment

A central prediction of our model is that R&D intensity (R&D as a fraction of sales) is independent of firm size. While a firm faces diminishing returns to expanding R&D at a point in time, a larger firm has more knowledge capital to devote to the innovation process. When R&D investment scales with firm size, as it does when firms choose R&D optimally, these two effects exactly offset each other, leaving both the average and the marginal productivity of research the same across different sizes of firms. Since the

model generates substantial heterogeneity in firm size, it predicts large differences in R&D investment across firms. Since the model generates substantial persistence in firm size, it also predicts persistence in R&D investment, in line with stylized fact 6.

A second source of differences in R&D across firms arises from heterogeneity in research intensity (as required by stylized fact 4). We posit exogenous permanent differences across firms in the size of the inventive steps embodied in their innovations, with research costs increasing in the size of the step. Although larger inventive steps are more profitable, their increased cost is just enough so that all firms optimally choose to innovate at the same rate. Nonetheless, firms that take big steps are more research intensive than more imitative firms taking tiny steps. Since they all innovate at the same rate, the more R&D intensive firms do not grow faster, and hence end up being no larger on average, than the more imitative firms. This last result maintains the observed independence of firm size and R&D intensity (stylized fact 3). Since a firm always takes inventive steps of a particular size, we capture the persistence of differences in R&D intensity (stylized fact 5). Of course, our explanation prompts the question of why a firm is endowed with the ability to take inventive steps of a particular size. We defer this question for future work.

8.1.2 R&D and Patenting

We can relate the firm-level innovation production function (1) to data under the assumption that each innovation is patented. Our specification of decreasing returns to R&D in the innovation production function then appears to contradict the observation of constant returns in firm-level regressions of patents on R&D (stylized fact 2). This contradiction is only apparent since knowledge capital is also an input in our innovation production function and R&D is endogenous in the model. In the case of equal sized inventive steps, the optimal R&D policy implies that firms doing more R&D have more knowledge capital. Looking across firms that have optimally chosen their R&D investment, the model implies that patents should follow a Poisson distribution with a parameter proportional to firm R&D, much like what has been observed.

Allowing for heterogeneous research intensity, but continuing to assume one patent

per innovation, the model predicts that patents should rise less than proportionately with R&D. The reason is that variation in R&D then reflects not only differences in knowledge capital but also heterogeneity in the size of firms' inventive steps. The expected number of patents rises in proportion to the former but is unrelated to the latter. If, however, innovations embodying larger inventive steps are more likely to be patented or to receive more than one patent, the model could still match the observed constant returns relationship.

8.1.3 R&D and Productivity

How can we relate a firm's productivity to its innovative performance if, as in our model, the firm innovates by extending its product line? It turns out that while patents are an indicator of the number of innovations, higher productivity reflects larger inventive steps. Since a firm's R&D intensity is also increasing in the size of its inventive steps, we predict the observed positive correlation between R&D intensity and productivity (stylized fact 1).

To make this argument precise, consider a firm taking inventive steps of size q . The step size gives the firm market power which it exploits by setting the prices of its products equal to a markup q over its constant unit labor cost. Summing across the firm's products, the ratio of its total revenue to its total labor cost is also equal to q . We would typically measure the firm's productivity a as the value of its output divided by employment. Thus $a = qw$, where w is the market wage rate. Variation in the size of inventive steps across firms produces heterogeneity in this measure of productivity. Since we assume that the size of inventive steps is a characteristic of a firm, we predict persistent differences in productivity.¹⁸

8.2 Interpreting Firm and Industry Dynamics

We argued in the introduction that firm innovation and firm growth deserve an integrated treatment. It turns out that our model of innovating firms has the key elements found in

¹⁸Our argument about how to interpret measures of firm-level productivity borrows from earlier work by Klette and Griliches (1996) and Bernard et al. (2000).

existing models of firm and industry dynamics: heterogeneous firms, simultaneous exit and entry, optimal investments in expansion, explicit individual firm dynamics, and a steady-state firm size distribution. In contrast to the existing literature, including Simon and Bonini (1958), Jovanovic (1982), Hopenhayn (1992), Ericson and Pakes (1995), and Sutton (1998), our model captures all these elements while remaining analytically tractable.

The fundamental source of firm heterogeneity in the model is the luck of the draw in R&D. A firm grows if it innovates and shrinks if a competitor innovates by improving on one of the firm's products. The firm's optimal R&D strategy has it innovate at a rate proportional to its size. A firm enters if the expected value of a new product covers the entry cost, and it exits when it loses its last product to a competitor. Together, these elements of the model capture (i) exit probabilities that are decreasing in firm size (and age), (ii) firm growth rates that are decreasing in size among surviving small firms, and (iii) Gibrat's law holding as a good approximation for large firms. The dispersion in firm sizes converges to a stable skewed distribution.

Of course, our model has set aside some important aspects of reality. We assume, as in Hopenhayn (1992), a continuum of firms and no aggregate shocks. Hence, we have ruled out aggregate uncertainty as well as strategic investment behavior, features likely to be important in an industry with just a few competitors. The Ericson and Pakes (1995) framework is much richer in these respects, but at the cost of substantial complexity.

8.3 Directions for Future Work

Our goal has been to establish a connection between theories of aggregate technological change and findings from firm-level studies of innovation. The potential payoff is twofold. We have attempted to demonstrate above that our fully articulated equilibrium model can clarify the interpretation of firm-level empirical findings. Furthermore, building on the firm-level stylized facts, the resulting aggregate model is likely to be more credible both as a description of reality and as a tool for policy analysis.

One direction for future research is to analyze a specific industry in which innovation plays a major role. With firm-level panel data on R&D, patenting, employment, and

revenue from such an industry, we could subject the model to a more detailed quantitative assessment. If it survives such an assessment, the model could help explore difficult questions about the interactions between industry evolution and technological change.

Another direction is to pursue the model's implications for policy. Unlike in its predecessors, in our model incumbent research firms play an important role in driving aggregate technological change. This feature is essential for evaluating the impact of actual R&D subsidies, which, as emphasized by Mansfield (1986), are often explicitly designed to act on the marginal expenditures of R&D-doing firms. We see a potential for extending the analysis here to address questions that frequently arise concerning policies to promote innovation.

Appendix A: Discussion and references for the evidence of innovating firms

R&D, Productivity, and Patents

Stylized Fact 1 *Productivity and R&D across firms are positively related, while productivity growth is not strongly related to firm R&D.*

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 (1998, ch.12; 2000, ch.4) and Hall (1996). This positive relationship has been consistently verified in a number of studies focusing on cross-sectional differences across firms. The longitudinal (within-firm, across-time) relationship between firm-level differences in R&D and productivity *growth*, which controls for permanent differences across firms, has turned out to be fragile and typically not statistically significant.

Stylized Fact 2 *Patents vary proportionally with R&D across firms, while there are diminishing returns to R&D in the longitudinal dimension.*

The relationship between innovation, patents, and R&D has been surveyed by Griliches (1990). He emphasizes that there is quite a strong relationship across firms between R&D and the number of patents received. For larger firms the patents-R&D relationship is close to linear, whereas 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 (1990), on the other hand, argues that “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” (p. 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” (p. 1676).¹⁹ 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 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), p. 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.

R&D Investment

Stylized 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 entire firm size distribution” (p. 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” (p. 928). While the first statement supports stylized 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,

¹⁹See Kleinknecht (1987).

the innovative activity being singled out in a firm's accounts as formal R&D is related to the absolute level of R&D.

Stylized Fact 4 *The distribution of R&D intensity is highly skewed, and a considerable fraction of firms report zero R&D.*

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 exhibits a regular pattern across industries, in accordance with stylized fact 4. The R&D intensity distributions they present are all unimodal, are positively skewed with a long tail to the right, and have 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.

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

Scott (1984) shows that in a large longitudinal sample of U.S. 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.

Stylized Fact 6 *Firm R&D investment follows essentially a geometric random walk.*

In a study of U.S. manufacturing firms, Hall et al. (1986) concludes 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" (p. 281). Similarly, Klette and Griliches (2000) report zero correlation between changes in log R&D and the level of R&D for Norwegian firms.

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

Stylized 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)²⁰. 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” (p. 65).

Stylized Fact 8 *Smaller firms have a lower probability of survival, but those that survive tend to grow faster than larger firms. Among larger firms, growth rates are unrelated to past growth or to firm size.*

Stylized 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. Our statement corresponds to the summaries of the literature on Gibrat’s law by Sutton (1997), Caves (1998) and Geroski (1998).

Stylized 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).

Stylized Fact 10 *Younger firms have a higher probability of exiting, but those that survive tend to grow faster than older firms. The market share of an entering cohort of firms generally declines as it ages.*

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

²⁰A recent contribution to the empirical literature is Stanley et al. (1995).

Appendix B: The Firm's Innovation Policy

To derive properties of the firm's innovation policy, it is useful to construct the function $f(x) = [\bar{\pi} - c(x)]/[r + \mu - x]$. The solution to the Bellman's equation (3) implies $f(\lambda) = v = c'(\lambda)$, where λ is the firm's optimal innovation intensity.

Differentiating, we get $f'(x) = [f(x) - c'(x)]/[r + \mu - x]$. Our assumptions on the cost function imply $f(0) > c'(0)$ and $f(\mu) \leq c'(\mu)$. Combining these results, $f'(0) > 0$ and $f'(\mu) \leq 0$. Thus, there is some intermediate $x^* \in (0, \mu]$ such that $f'(x^*) = 0$. Since $f'(x) > 0$ for any $x \in [0, x^*)$ and $f'(x) < 0$ for any $x \in (x^*, \mu + r)$, it follows that x^* is unique. Of course, the optimal innovation intensity of the firm is simply $\lambda = x^*$.

To see how λ depends on parameters, plot $f(x)$ and $c'(x)$ for $x \in [0, \mu]$. The unique intersection of these curves determines λ . An increase in $\bar{\pi}$ shifts up $f(x)$, leading to an increase in λ . Similarly, an increase in r or μ shifts $f(x)$ down, leading to a decrease in λ . A shift up in c' also leads to an increase in c and a resulting shift down in $f(x)$. The shift up in c' and shift down in f lead to an unambiguous decline in λ .

Our assumptions on the cost function $[\bar{\pi} - c(\mu)]/r \leq c'(x) < \bar{\pi}/(r + \mu)$ for $x \in [0, \mu]$ appear arbitrary. If we dropped the upper bound on c' , then we get the solution $\lambda = 0$ in the case of $c'(0) \geq \frac{\bar{\pi}}{r + \mu}$. The lower bound on c' is more critical, as we want to analyze a firm's dynamics under the restriction $\lambda \leq \mu$, a restriction which will arise naturally when we consider the industry equilibrium.

Appendix C: Solving the System of Difference-Differential Equations

Let N_t be the random variable giving the size of a firm at date t . The probability of the firm having n products ($n \geq 1$) at time $t + \Delta t$ satisfies the relationship

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

where $\lim_{\Delta t \rightarrow 0} \mathcal{O}(\Delta t) / \Delta t = 0$. Following standard techniques described, e.g., in Karlin and Taylor (1975, ch. 4) and in Goel and Richter-Dyn (1974, section 2), we find that

$$\begin{aligned} \frac{\partial P[N_t = n]}{\partial t} &= \lim_{\Delta t \rightarrow 0} \frac{P[N_{t+\Delta t} = n] - P[N_t = n]}{\Delta t} \\ &= (n-1)\lambda P[N_t = n-1] + (n+1)\mu P[N_t = n+1] \\ &\quad - n(\lambda + \mu) P[N_t = n]. \end{aligned}$$

Letting $p_n(t) = P[N_t = n]$, we obtain a more compact expression

$$\dot{p}_n(t) = (n-1)\lambda p_{n-1}(t) + (n+1)\mu p_{n+1}(t) - n(\lambda + \mu) p_n(t), \quad n \geq 1. \quad (20)$$

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

$$\dot{p}_0(t) = \mu p_1(t). \quad (21)$$

The set of coupled difference-differential equations (20) and (21) can be solved with the aid of the probability generating function²¹, defined as

$$H(z, t) = \sum_{n=0}^{\infty} p_n(t) z^n. \quad (22)$$

Hence,

$$\begin{aligned} \frac{\partial H(z, t)}{\partial z} &= \sum_{n=0}^{\infty} n p_n(t) z^{n-1} \\ &= \sum_{n=1}^{\infty} n p_n(t) z^{n-1}. \end{aligned} \quad (23)$$

Consider

$$\frac{\partial H(z, t)}{\partial t} = \sum_{n=0}^{\infty} \dot{p}_n(t) z^n = \dot{p}_0(t) + \sum_{n=1}^{\infty} \dot{p}_n(t) z^n. \quad (24)$$

The first term on the right-hand side can be restated by (21), while the second term can be restated by multiplying (20) by z^n and summing over n from 1 to ∞ , which, after

²¹The solution procedure is described in detail in Kendall (1948) and in Goel and Richter-Dyn (1974) (in particular chapter 2 and appendix B).

some rearrangements of terms, gives

$$\begin{aligned} \frac{\partial H(z, t)}{\partial t} = & -(\lambda + \mu) \sum_{n=1}^{\infty} n p_n(t) z^n \\ & + \lambda \sum_{n=1}^{\infty} (n-1) p_{n-1}(t) z^n \\ & + \mu \left[p_1 + \sum_{n=1}^{\infty} (n+1) p_{n+1}(t) z^n \right]. \end{aligned} \quad (25)$$

Using (23) on each of the three sums, we find that (25) can be restated as

$$\frac{\partial H(z, t)}{\partial t} = [\lambda z^2 - (\lambda + \mu)z + \mu] \frac{\partial H(z, t)}{\partial z}. \quad (26)$$

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

In order to solve for $H(z, t)$, we require some initial condition. To analyze a firm that was in state n_0 at date 0, we set $p_{n_0}(0) = 1$ and $p_n(0) = 0$ for $n \neq n_0$. From (22) it follows that

$$H(z, 0; n_0) = \sum_{n=0}^{\infty} p_n(0) z^n = z^{n_0}. \quad (27)$$

With the initial condition (27), the solution to (26) can be written

$$H(z, t; n_0) = \left[\frac{\mu(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)}{\lambda(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)} \right]^{n_0}. \quad (28)$$

A Taylor series expansion of $H(z, t; n_0)$ around $z = 0$ yields the probability distribution $p_n(t; n_0) = P[N_t = n | N_0 = n_0]$ as the coefficients in the series, i.e.,

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

Note that

$$H(1, t; n_0) = \sum_{n=0}^{\infty} p_n(t; n_0) = 1, \quad (29)$$

so that for any t and n_0 , we get a proper distribution on the nonnegative integers. Furthermore, (28) also provides expressions for the moments of the distribution:

$$\sum_{n=0}^{\infty} n^k p_n(t; n_0) = \left[\frac{\partial^k H(e^s, t; n_0)}{\partial s^k} \right]_{s=0}. \quad (30)$$

In the case when $\mu \rightarrow \lambda$ (the case with no entry) both the numerator and the denominator inside the square brackets on the right-hand side of (28) approach zero. Using l'Hopital's rule, we can also obtain the pgf in this special case:

$$\lim_{\mu \rightarrow \lambda} H(z, t; n_0) = \left[\lim_{\mu \rightarrow \lambda} \frac{\mu(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)}{\lambda(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)} \right]^{n_0} \quad (31)$$

$$= \left[\frac{\mu t(z-1) - 1}{\mu t(z-1) - z} \right]^{n_0}. \quad (32)$$

Appendix D: The Size Distribution

Derivation of (17): From (15) and (16), we have that

$$\dot{M} = \eta - \mu M_1. \quad (33)$$

In steady state all the time-derivatives in (15), (16), and (33) are zero. Starting with (33), the steady state requires

$$M_1 = \frac{\eta}{\mu}. \quad (34)$$

Substituting (34) into (16), we see that the steady state also requires

$$M_2 = \frac{\lambda \eta}{2\mu^2}. \quad (35)$$

Applying (15) and setting $\dot{M}_n(t) = 0$, it is straightforward to prove by induction that the general condition for the steady state is

$$M_n = \frac{\lambda^{n-1} \eta}{n \mu^n}, \quad n \geq 1. \quad (36)$$

Derivation of (19): We want to prove that

$$\lim_{t \rightarrow \infty} \sum_{m=1}^{\infty} p_n(t; m) M_m(0) = 0, \quad \forall n \geq 1.$$

Since $p_n(t; m)$ is a sequence of terms bounded between 0 and 1 and $\sum_{m=1}^{\infty} M_m(0)$ is finite (equal to the total mass of firms at date 0), it follows that for any $\varepsilon > 0$, $\sum_{m=k}^{\infty} p_n(t; m) M_m(0) \leq \sum_{m=k}^{\infty} M_m(0) \leq \varepsilon$ if k is sufficiently large. Hence, we can interchange the summation and the limit-operation:

$$\begin{aligned} & \lim_{t \rightarrow \infty} \sum_{m=1}^{\infty} p_n(t; m) M_m(0) \\ &= \sum_{m=1}^{\infty} \lim_{t \rightarrow \infty} p_n(t; m) M_m(0) \\ &= 0. \end{aligned}$$

The third equality reflects that $\lim_{t \rightarrow \infty} p_n(t; m) = 0$, $\forall (m, n) \geq 1$, since all firms eventually exit as shown in section 4.3. Q.E.D.

Appendix E: The Social Planner's Problem

We consider the social planner's problem as limited to choosing the allocation of workers between production and the two research activities, start-ups and incumbent research. The planner is constrained to accept the market allocation of production workers across goods. That is, the planner can subsidize production or research work as such, but cannot target subsidies to the production of individual goods.

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

$$U_t = \int_t^{\infty} e^{-\rho(\tau-t)} \ln C_{\tau} d\tau.$$

Now, $x(j) = [p(j)]^{-1} = [q(j)w]^{-1}$ and $L_X = \int_0^1 l_X(j) dj = \int_0^1 x(j) dj = E[q^{-1}]/w$; i.e., $w = E[q^{-1}]/L_X$. Furthermore,

$$\begin{aligned} \ln C_{\tau} &= J_{\tau} E[\ln q] + \int_0^1 \ln x(j) dj \\ &= J_{\tau} E[\ln q] - \ln(E[q^{-1}]) - E[\ln q] + \ln L_X. \end{aligned}$$

The stock of innovations for any good is given by $J_\tau = \int_0^\tau \mu(s)ds$ so that $\dot{J} = \mu = \lambda + \eta$. From the labor constraint

$$\eta = h(L - L_X - L_R).$$

Hence, the current value Hamiltonian is

$$\mathcal{H} = J_\tau E[\ln q] - \ln(E[q^{-1}]) - E[\ln q] + \ln L_X + \vartheta \{\lambda + h[L - L_X - l_R(\lambda)]\},$$

where we have used that $L_R = l_R(\lambda)$. The first-order conditions, where we denote the social planner's solution with a *, are

$$\frac{\partial \mathcal{H}}{\partial L_X} = 0 \Rightarrow L_X^* = \frac{1}{\vartheta h} \quad (37)$$

$$\frac{\partial \mathcal{H}}{\partial \lambda} = 0 \Rightarrow l'_R(\lambda^*) = h^{-1}. \quad (38)$$

Notice that the last relationship is the same as for the market solution, which implies that incumbent research is the same for the planner and the market economy. The co-state variable satisfies

$$\dot{\vartheta} = \rho\vartheta - \frac{\partial \mathcal{H}}{\partial J} = \rho\vartheta - E[\ln q].$$

The only non-exploding solution for ϑ requires $\dot{\vartheta} = 0$, which implies $\vartheta = E[\ln q] / \rho$. Hence,

$$L_X^* = \frac{\rho}{hE[\ln q]}.$$

The planner will choose a rate of innovation from new firms given by

$$\eta^* = h \left\{ L - \frac{\rho}{hE[\ln q]} - l_R(\lambda^*) \right\}.$$

As derived in section 7.1, the market will choose $L_X = E[q^{-1}][L + \rho/h]$. Hence,

$$\begin{aligned} \mu^* - \mu &= \eta^* - \eta = h(L_X - L_X^*) \\ &= \rho E[q^{-1}] \left(\frac{hL}{\rho} + 1 - \frac{1}{E[q^{-1}]E[\ln q]} \right). \end{aligned}$$

In the special case of only a single type of firm $\pi = \bar{\pi}$, and hence no variation in q , we get

$$\mu^* - \mu = \frac{\rho}{q} \left(\frac{hL}{\rho} + 1 - \frac{q}{\ln q} \right),$$

which is exactly the same expression as Grossman and Helpman (1991) derive. 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.

References

- Aghion, P. and P. Howitt (1992): A model of growth through creative destruction. *Econometrica*, 60, 323–51.
- 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, Mass.: 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).
- Carroll, G. and T. Hannan (2000): *The Demography of Corporations and Industries*. Princeton, 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): Untangling the origins of competitive advantage. *Strategic Management Journal*, 21, 1123–45. Special Issue on the Evolution of Firm Capabilities.
- 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.*
- Dunne, T., M.J. Roberts, and L. Samuelson (1988): Patterns of firm entry and exit in U.S. manufacturing industries. *Rand Journal of Economics* 19, 495-515.
- 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.
- Geroski, P. (1998): An applied econometrician's view of large company performance. *Review of Industrial Organization*, 13, 271-93.
- Geroski, P., J. Van Reenen and C. Walters (1997): How persistently do firms innovate? *Research Policy*, 26, 33-48.
- Goel, N. and N. Richter-Dyn (1974): *Stochastic Models in Biology*. New York: Academic Press.
- Griliches, Z. (1979): Issues in assessing the contribution of research and development to productivity growth. *Bell Journal of Economics*, 10, 92-116.
- (1990): Patent statistics as economic indicators: A survey. *Journal of Economic Literature* 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.
- (2000): *R&D, Education, and Productivity: A Retrospective*. Cambridge, Mass.: Harvard University Press.
- Grossman, G. and E. Helpman (1991): *Innovation and Growth in the Global Economy*. Cambridge, Mass.: 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., 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.
- Hymer, S. and P. Pashigian (1962): Firm size and rate of growth. *Journal of Political Economy*, 70, 556–569.
- Ijiri, Y. and H. Simon (1977): *Skew distributions and the sizes of business firms*. Amsterdam: North-Holland.
- 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. *Brookings Papers on Economic Activity*, (Microeconomics), 197–235.
- (2001): Fitness and age: Review of Carroll and Hannan’s “Demography of corporations and industries.” *Journal of Economic Literature*, 39, 105–9.
- Karlin, S. and H. Taylor (1975): *A First Course in Stochastic Processes. Second Edition*. San Diego: Academic Press.
- Kendall, D. (1948): On the generalized “birth-and-death” process. *Annals of Mathematical Statistics*, 19, 1–15.
- Kleinknecht, A. (1987): Measuring R&D in small firms: How much are we missing? *Journal of Industrial Economics*, 36, 253–6.
- Klepper, S. (1996): Entry, exit, growth and innovation over the product life cycle. *American Economic Review*, 86, 562–83.

- Klette, T. (1996): R&D, scope economies and plant performance. *RAND Journal of Economics*, 27, 502–22.
- Klette, T. and Z. Griliches (1996): The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics*, 11, 343–61.
- (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.
- Langlois, R. and P. Robertson (1995): *Firms, Markets and Economic Change*. New York: Routledge.
- Li, C. (2001): On the policy implications of endogenous technological progress. *Economic Journal*, 111, C164–79.
- Lucas, R. (1967): Adjustment costs and the theory of supply. *Journal of Political Economy*, 74, 321–34.
- Mansfield, E. (1986): The R&D tax credit and other technology policy issues. *American Economic Review, Papers and Proceedings*, 76, 190–94.
- Nelson, R. (1988): Modelling the connections in the cross section between technical progress and R&D intensity. *RAND Journal of Economics*, 19, 478–85.
- Penrose, E. (1959): *The Theory of the Growth of the Firm*. Oxford: Basil Blackwell.
- Peretto, P. (1999): Cost reduction, entry, and the interdependence of market structure and economic growth. *Journal of Monetary Economics*, 43, 173–95.
- Romer, P. (1990): Endogenous technological change. *Journal of Political Economy* 98, S71-S102.
- 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.
- 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.

- Simon, H. A. and C. P. Bonini (1958): The size distribution of business firms. *American Economic Review*, 48, 607–17.
- 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.
- (2001): The variance of firm growth rates: The ‘scaling’ puzzle. Discussion Paper no. EI/24, STICERD, The London School of Economics and Political Science.
- Thompson, P. (1996): Technological opportunity and the growth of knowledge: A Schumpeterian approach to measurement. *Journal of Evolutionary Economics*, 6, 77–97.
- (2001): The microeconomics of an R&D-based model of endogenous growth. *Journal of Economic Growth*, (forthcoming).
- Uzawa, H. (1969): Time preference and the Penrose effect in a two-class model of economic growth. *Journal of Political Economy*, 77, 628–52.