

Innovating Firms and Aggregate Innovation

Tor Jakob Klette and Samuel Kortum

*University of Oslo and Centre for Economic Policy Research; University of Minnesota,
Federal Reserve Bank of Minneapolis, and National Bureau of Economic Research*

ABSTRACT

We develop a parsimonious model of innovation to confront firm-level evidence. It captures the dynamics of individual heterogeneous firms, describes the behavior of an industry with firm 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 of firms' R&D investment, the concentration of R&D among incumbents, the link between R&D and patenting, and why R&D as a fraction of revenues is positively correlated with firm productivity but not with firm size or growth. (JEL classification: L11, O31)

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 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 innovative output?

Firm growth is yet another measure of innovative performance. Empirical studies of firm growth, entry, exit, and the size distribution could 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 to address these issues. The model is 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. At the root of 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 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 closely related to Penrose’s (1959) theory 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 internal resources. While Penrose stresses managerial and entrepreneurial resources of the firm, 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.

The Schumpeterian force of creative destruction is pervasive in our analysis. Firms grow by making innovations in products new to them, but the economy grows as innovations raise the quality of a given set of products. Thus, 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. In fact, the model typically predicts that exit is the fate of any firm.

While innovating firms follow a stochastic life cycle, industry equilibrium typically involves simultaneous entry and exit, with a stable, skewed firm size distribution. In this sense our model captures some of the features 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 following the strategy of Hopenhayn (1992) in which each firm is small relative to the industry. Solving our model in general equilibrium extends the work of Grossman and Helpman (1991) and Aghion and Howitt (1992) by incorporating the contribution of incumbent research firms to aggregate innovation.

The paper begins by summarizing the firm-level findings, or stylized facts, that are the target for our theoretical model. The model itself is developed in Section 3. In the light of the

model, we return to the stylized facts that motivated our analysis in Section 4 and offer some suggestions for future work.

2 Evidence on Innovating Firms

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

We have listed only the empirical regularities that are robust and economically significant. Following Cohen and Levin (1989) and Schmalensee (1989), we have ignored findings (despite their statistical significance) if they appear fragile or not economically significant. The stylized facts on which we focus are largely summarized in surveys by others, including Griliches (1990, 1998, 2000), Cohen (1995), Sutton (1997), and Caves (1998). 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 innovative output, 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. As for firm-level patenting and R&D, the positive relationship is very robust in either dimension.

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

Stylized Fact 2 *Patents and R&D are positively related both across firms at a point in time and across time for given firms.*

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

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

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 section, we will return to these stylized facts and interpret them in light of the theoretical framework.

3 The Model

We present the model in steps starting with the innovation process for an individual firm. We then, in turn, analyze firm dynamics, introduce exogenous heterogeneity in firms' research intensity, describe entry and the size distribution of firms, and solve for aggregate innovation in general equilibrium.

The economy consists of a unit continuum of differentiated goods. Consumers have symmetric Cobb-Douglas preferences across these goods so that the same amount is spent on each one. We let total expenditure be the numeraire and set it to one. Time is continuous; hence, there is a unit flow of expenditure on each good.

3.1 The Innovating Firm

A firm is defined by the portfolio of goods that it produces. Competition between firms, described in detail below, results in each good being produced by a single firm and yielding a profit flow $0 < \pi < 1$. Note that the profit flow is strictly less than the revenue flow of one. In Section 3.3 we will introduce heterogeneity in π across firms, but for now it suffices to consider a fixed value $\pi = \bar{\pi}$. Since each good generates the same flow of revenue and profit to the producer, in describing the state of a firm we need only keep track of a scalar, the number of goods $n = 1, 2, 3, \dots$ that it produces. We do not, for example, need to keep track of the n distinct locations on the unit interval indicating the particular set of goods produced by the firm. 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*.

In particular, the firm's R&D investment determines the Poisson rate I at which its next product innovation arrives. With an innovation for a particular good, the firm can successfully compete against the incumbent producer. The most recent innovator takes over the market for that good. Expenditures on R&D could yield an innovation relevant to any good with equal probability; i.e., the good to which it applies is drawn from the uniform distribution on $[0, 1]$.

Because each firm produces only an integer number of goods, when we consider the economy as a whole in Sections 3.4–3.5, we will be dealing with a continuum of firms. For now we focus on the behavior of one such firm. Since any firm is infinitesimal relative to the continuum of goods, we can ignore the possibility that it innovates on a good it is currently producing. A firm does, however, face the possibility that some other firm will innovate on a good it is currently producing. In this case, the incumbent producer loses the good from its portfolio. 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.

The natural interpretation of our assumptions about innovation, demand, and competition is the quality ladder model of Grossman and Helpman (1991). The unit cost of production is a constant across all goods and firms. Innovations come in the form of quality improvements, raising the quality of some good by a factor $q > 1$. With Bertrand competition between the latest innovator and the incumbent producer, the innovator captures the entire market for a particular good, charges a price q times unit cost, and obtains a flow of profit $\pi = 1 - q^{-1}$ for that good. The value of μ is determined by the aggregate rate of innovation. This interpretation will be formalized in Section 3.5, where we lay out the general equilibrium. First, we turn to our specification of the firm's innovation possibilities, which is where our model deviates substantially from the existing literature.

3.1.1 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 abstract concept of knowledge capital in the simplest way, we assume that it is summarized by n , the number of goods produced by the firm, which is also equal to all 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 first condition implies that R&D is a productive investment. The second 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 condition 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 assumptions made above,

the firm's R&D costs are a homogeneous 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 . We assume that (i) $c(0) = 0$, (ii) $c(x)$ is twice differentiable for $x \geq 0$, (iii) $c'(0) < \bar{\pi}/r$, and (iv) $[\bar{\pi} - c(\mu)]/r \leq c'(\mu)$. Restriction (iii) is needed in Section 3.4 (with μ treated as endogenous) to guarantee $\mu > 0$, while (iv) is needed in Section 3.2 (with μ taken as given) to guarantee $\lambda \leq \mu$.

3.1.2 The R&D 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 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

$$V(n) = vn$$

$$I(n) = \lambda n,$$

where v and λ solve

$$\begin{aligned} c'(\lambda) &= v \quad [\text{or } c'(0) > v \text{ and } \lambda = 0] \\ (r + \mu - \lambda)v &= \bar{\pi} - c(\lambda). \end{aligned} \tag{4}$$

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

The firm's associated 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.

3.2 The Firm's Life Cycle

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) derived in Appendix C. We now turn to the economic implications of that solution.

3.2.1 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.) In this case, the pgf

from Appendix C yields

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

where

$$\gamma(t) = \frac{\lambda(1 - e^{-(\mu-\lambda)t})}{\mu - \lambda e^{-(\mu-\lambda)t}} = \frac{\lambda}{\mu} p_0(t; 1).$$

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 one 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, conditional on the firm surviving, there is an increasingly high probability that the firm has become very large. (Of course, if $\lambda = 0$ a surviving firm is always size $n = 1$.)

3.2.2 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 analysis of size 1 firms can be applied to how a larger firm evolves. For example, the probability that a firm of size n exits within t periods is $p_0(t; 1)^n$. Larger firms have a lower hazard of exiting, in line with stylized fact 8.

3.2.3 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

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

(For the case of $\lambda = 0$, apply l'Hopital's rule to get $E[A] = \mu^{-1}$.) 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 intensity 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, in line with stylized fact 10, and approaches $\mu - \lambda$ for a very old firm. The intuition is that as a firm ages, implying that it has survived, it tends to get bigger (for $\lambda > 0$) and is therefore less likely to exit.⁹

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

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

which increases with age. Consider a cohort of m firms all entering at the same date. Over time the number of firms in the cohort diminishes while the survivors grow bigger on average. The expected number of products produced by the cohort (i.e., the total revenues it generates) at age a is

$$m \sum_{n=0}^{\infty} n p_n(a; 1) = m \frac{1 - p_0(a; 1)}{1 - \frac{\lambda}{\mu} p_0(a; 1)}. \quad (10)$$

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

3.2.4 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 shrink as measured by the number of products it produces or, equivalently, its revenue. (Due to our choice of numeraire, firm revenue is measured relative to aggregate expenditure in the economy, a

point which is taken up again at the end of Section 3.5.)

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 $1 - p_0(t; n_0) = 1 - p_0(t; 1)^{n_0}$, which is clearly increasing in initial size. Expected growth *conditional on survival* is

$$E [G_t | N_t > 0, N_0 = n_0] = \frac{e^{-(\mu - \lambda)t}}{1 - [p_0(t; 1)]^{n_0}} - 1,$$

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 that are initially very large the probability of survival to date t is close to one anyway, so Gibrat's law will be a very good approximation, in line with stylized fact 8.

3.3 Heterogeneous Research Intensity

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 because 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, which is not captured by the model.

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 as we introduce another dimension of firm heterogeneity.

To do so 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 value $0 < \pi < 1$ representing the profit flow that it obtains from each of its innovations (i.e., each good that it sells). It is notationally convenient to set the mean value of π equal to $\bar{\pi}$. (In Section 3.4.1 we are more explicit about the distribution of π .) A firm's profit type π remains fixed throughout its life. Below we will model π in terms of the size of the innovative step made by the firm so that a larger π corresponds to a more innovative firm.¹²

The firm's type not only affects its flow of profits from an innovation but also its cost of doing research. 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 $C(I, n)$ from (2).

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 function for a size n firm of type π is $V_\pi(n) = v_\pi n$, where the value per product of such a firm is simply $v_\pi = \frac{\pi}{\bar{\pi}}v$. Firms of any type π choose the same innovation intensity λ so that firm growth is unrelated to the firm's

type. The analysis of a firm's life cycle in Section 3.2 thus applies without modification to firms of any type π . However, 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 in π carries over to heterogeneity in R&D intensity, with the more innovative firms (higher type π firms) being more R&D-intensive.

3.4 Industry Behavior

In this section we examine the dynamic process characterizing an industry with many competing firms. In considering an industry, we assume that every good is being produced by some firm. The number of goods produced by any given firm is countable; hence, there must be a continuum of firms if, taken as a whole, they cover the production of the entire unit continuum of goods. 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, $\sum_{n=1}^{\infty} nM_n(t) = 1$. In accounting for the measure of firms of different sizes, we can totally ignore differences in firms of type π since firms of any type will choose the same λ .

Taken as a whole, industry incumbents innovate at rate

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

Thus, the innovative intensity of each incumbent also has an aggregate interpretation. 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 entry process, which yields a simple expression for both λ and η .

3.4.1 Entry

There is a mass of potential entrants. A potential entrant must invest at rate F in return for a Poisson hazard 1 of entering with a single product. The type π of the entrant is drawn from a distribution $\Phi(\pi)$ with mean $\bar{\pi}$. A potential entrant knows the distribution of types but only learns his type after entering.

If there is active entry, i.e. $\eta > 0$, we have the condition $F = E[v_\pi] = E[\frac{\pi}{\bar{\pi}}v] = v$. In this case the innovation intensity of incumbents is determined by

$$c'(\lambda) = F \quad [\text{or } c'(0) > F \text{ and } \lambda = 0]. \tag{14}$$

Given (14), a higher entry cost shields incumbents from competition, leading them to invest more in innovation.

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

$$\eta = \frac{\bar{\pi} - c(\lambda)}{F} - r \quad [\text{or } \frac{\bar{\pi} - c(\lambda)}{F} - r \leq 0 \text{ and } \eta = 0],$$

where λ in the expression above is from (14). For large enough $\bar{\pi}$ and small enough F and r , there will be positive entry.

If there is not entry ($\eta = 0$ in the expression above), then (14) need not hold. Instead, we can set $\mu = \lambda$ in the solution to the Bellman's equation (4), which yields λ as the solution to $c'(\lambda) = v = [\bar{\pi} - c(\lambda)]/r$. The requirement that $\mu > 0$ is assured by our assumption in Section 3.1.1 that $c'(0) < \bar{\pi}/r$.

3.4.2 The Size Distribution

We now have expressions for η and λ , and hence $\mu = \eta + \lambda$. These magnitudes 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, ... 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

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

Case of $\eta = 0$: With no entry, 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 individual firm dynamics from Section 3.2.1. Suppose at date 0 the industry consists of a unit mass of size 1 firms. By date t , there will be a mass $M(t) = 1 - p_0(t; 1) = 1/(1 + \mu t)$, among which the size distribution will be geometric with a parameter $\mu t/(1 + \mu t)$. Thus, without entry, the mass of firms continually declines, the average size of surviving firms becomes ever larger, and the size distribution of survivors becomes ever more skewed.

Case of $\lambda = 0$: If incumbents do no research, then all innovation is carried out by entrants. Entrants start out as size $n = 1$ and never grow. Hence, the size distribution is $M_1 = 1$ and $M_n = 0$ for all $n > 1$.

Case of $\lambda > 0, \eta > 0$: In the case of innovation by both entrants and incumbents, we can show that the industry will converge to a steady state with a constant mass of firms and a 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 increases as θ gets small since when there is relatively less entry (and relatively more innovation by incumbents) large incumbent firms less often lose goods to small entrants. The total mass of firms is

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

The mass of firms in the industry is an increasing function of θ , as with relatively more entry (and relatively less innovation by incumbents) there tend to be many size 1 firms.

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).

To analyze industry dynamics and to establish convergence to the steady state, we integrate over past 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 $t \rightarrow \infty$, as we establish formally in Appendix D. Since $\eta > 0$, 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).¹⁴

3.5 Aggregate Innovation in General Equilibrium

We can now close the model in general equilibrium. Doing so will clarify our seemingly arbitrary assumptions about the firm's revenue, the firm's profit per good, the cost function for innovation, the constant rate of creative destruction, and the constant interest rate. The cost of this general equilibrium setup will be some additional notation allowing us to be more explicit about the nature of innovations.

3.5.1 The Aggregate Setting

The economy consists of a unit measure of goods indexed by $j \in [0, 1]$ and a mass L of labor. A measure L_X of labor produces goods, a measure L_S innovates to launch new firms, and a measure L_R innovates at incumbent firms. Labor receives a wage w in any of these three activities. One worker produces a unit flow of any good j , hence the unit cost of production is w . A team of h researchers get an idea for a new firm at a Poisson rate 1; hence, the fixed cost of entry (introduced in Section 3.4.1) is $F = wh$. At an incumbent firm of size n and type π , it takes $\frac{\pi}{n}nl_R(x)$ researchers to generate innovations of type π at a Poisson rate nx . (We will be more explicit about these types in what follows.) We assume $l_R(0) = 0$, $l'_R(x) > 0$, and $l''_R(x) > 0$, consistent with the fact that the intensive form of the research cost function (introduced in equation (2)) is $c(x) = wl_R(x)$.

Each innovation (whether it spawns a new firm or a new good for an incumbent) is a quality improvement applying to a good drawn at random from the unit interval. The innovator can produce a higher quality substitute at the same unit cost w . For any particular good j , quality improvements arrive at a Poisson rate μ . The number arriving between date 0 and date t , denoted $J_t(j)$, is thus distributed Poisson with parameter μt .

For simplicity, we establish conditions at date 0 that lead immediately to a stationary

equilibrium. The highest quality of good j available at date 0 is denoted $z(j, 0)$, which is itself an advance over a default quality $z(j, -1) = 1$. Thus, by date t good j will be available in $J_t(j) + 2$ versions with quality levels $z(j, -1) < z(j, 0) < z(j, 1) < \dots < z(j, J_t(j))$. The innovative step of the k th innovation is $q(j, k) = z(j, k)/z(j, k - 1) > 1$, for $0 \leq k \leq J_t(j)$. The innovative steps are drawn independently from the distribution $\Psi(q)$.

The representative household has preferences over all versions $k \in \{-1, 0, 1, \dots, J_t(j)\}$ of each good $j \in [0, 1]$ at each date $t \geq 0$:

$$U = \int_0^\infty e^{-\rho t} \ln C_t dt,$$

$$\ln C_t = \int_0^1 \left[\ln \sum_{k=-1}^{J_t(j)} x_t(j, k) z(j, k) \right] dj,$$

where ρ is the discount rate and $x_t(j, k)$ is consumption of version k of good j at date t . Note that different versions of the same good are perfect substitutes once quality is taken into account.

The household owns all the firms and finances all the potential entrants. The average value of a size n firm at date t is $E[V_\pi(n)] = nE[v_\pi] = nv$. Thus, the value of all the firms in the economy at date t is $\sum_{n=1}^\infty M_n(t)nv = v$. Given an interest rate r , the household gets income rv from these assets. The cost of financing potential entrants is wL_S , yielding a flow of assets in new firms worth $(L_S/h)v$. The household's total income is therefore $Y = wL + rv + (v/h - w)L_S$.

Firms producing different versions of a good engage in Bertrand competition. We apply the standard tie-breaking rule that the household buys the highest quality version of a good if it costs no more than any other in quality-adjusted terms. Log preferences across goods imply that the household spends the same amount on each good j . In equilibrium only the highest quality version of good j is sold. Its price is $p_t(j) = wq_t(j)$, where $q_t(j) \equiv q(j, J_t(j))$ is the size of the last

innovative step.

Taking aggregate expenditure as the numeraire, we have expenditure per good $p_t(j)x_t(j) = 1$, where $x_t(j) \equiv x_t(j, J_t(j))$. Thus, the quantity of good j produced is $x_t(j) = [wq_t(j)]^{-1}$, and the profit flow to the last firm innovating in good j is $\pi_t(j) = [p_t(j) - w]x_t(j) = 1 - [q_t(j)]^{-1}$. The heterogeneity in profit flow π is linked to heterogeneity in the innovative step q according to $\pi = 1 - q^{-1}$. A firm of type π is thus a firm taking innovative steps of size $q = (1 - \pi)^{-1}$. The distribution over firm types $\Phi(\pi)$ (introduced in Section 3.4.1) is therefore linked to the distribution over innovative steps $\Psi(q)$ by the equality $\Phi(\pi) = \Psi([1 - \pi]^{-1})$. Averaging across goods and exploiting the law of large numbers, $\bar{\pi} = \int_0^1 [1 - q_t(j)^{-1}] dj = 1 - \int_1^\infty q^{-1} d\Psi(q)$.

Note that $\bar{\pi}$ is the profit share of income while $1 - \bar{\pi}$ is the share of income going to production labor. Throughout the analysis we have posited $\mu > 0$. To guarantee that this restriction is satisfied, with μ endogenous, we assume $\bar{\pi}L > (1 - \bar{\pi})\rho l'_R(0)$. The need for this restriction becomes apparent in our analysis of the stationary equilibrium.

3.5.2 Stationary Equilibrium

A stationary equilibrium of the economy involves constant values for the interest rate r , wage w , firm value v (for a size 1 firm of type $\bar{\pi}$), innovation by incumbents λ , and innovation by entrants η such that (i) potential entrants break even in expectation, (ii) incumbent firms maximize their value, (iii) the representative consumer maximizes utility subject to his intertemporal budget constraint, and (iv) the labor market clears.

A stationary equilibrium is simple to characterize. From the entry condition

$$v = wh \quad [\text{or } v < wh \quad \text{and} \quad L_S = 0].$$

From the incumbent's problem

$$v = wl'_R(\lambda) \quad [\text{or } v < wl'_R(0) \quad \text{and } \lambda = 0].$$

The demand for start-up researchers is therefore $L_S = \eta h$. The expected demand for incumbent researchers at a firm of size n (integrating over the distribution of types π) is $L_R(n) = nl_R(\lambda)$, and hence the aggregate demand for incumbent researchers is $L_R = \sum_{n=1}^{\infty} M_n n l_R(\lambda) = l_R(\lambda)$. The demand for production workers is $L_X = (1 - \bar{\pi})/w$. The household is willing to spend the same amount each period only if $r = \rho$. Since aggregate expenditure is the numeraire and start-up research breaks even, the consumer's budget constraint reduces to $wL + \rho v = 1$. The labor market clears if $L_S + L_R + L_X = L$. We can now solve for the equilibrium allocation of labor.

Case of Active Entry: Suppose $L_S > 0$. Combining the entry condition and the consumer's budget constraint, the wage is $w = 1/(L + \rho h)$. Combining the wage, demand for production workers, and labor market clearing, the measure of researchers at incumbents and potential entrants is $L_R + L_S = \bar{\pi}L - (1 - \bar{\pi})\rho h$.

From the incumbent's problem, if $h \leq l'_R(0)$, then $\lambda = 0$. That is, incumbents do not innovate. In that case $L_R = 0$ and $L_S = \bar{\pi}L - (1 - \bar{\pi})\rho h$ [hence, $L_S \geq \bar{\pi}L - (1 - \bar{\pi})\rho l'_R(0) > 0$].

Alternatively, if $h > l'_R(0)$, then $\lambda > 0$ solves $h = l'_R(\lambda)$. The measure of researchers at incumbents is $L_R = l_R(\lambda)$ and at potential entrants is $L_S = \bar{\pi}L - (1 - \bar{\pi})\rho h - l_R(\lambda)$. If this expression does not yield a positive L_S , then we must consider the case of no entry.

Case of No Entry: Suppose $L_S = 0$. We will conjecture and then confirm that in this case $L_R > 0$. Combining the incumbent's condition and the consumer's budget constraint, the wage is $w = 1/(L + \rho l'_R(\lambda))$. Combining the expressions for the wage, the demand for production workers,

and labor market clearing, the measure of researchers at incumbents is determined by the value of λ that solves $l_R(\lambda) = \bar{\pi}L - (1 - \bar{\pi})\rho l'_R(\lambda)$ [since $l_R(\lambda) \geq \bar{\pi}L - (1 - \bar{\pi})\rho l'_R(0) > 0$, the solution yields $L_R > 0$].

3.5.3 Aggregate Innovation and Growth

Having determined the equilibrium allocation of labor, the aggregate rate of innovation is simply $\mu = \lambda + \eta = l_R^{-1}(L_R) + L_S/h$. In all cases the rate of innovation is increasing in the size of the labor force, L , and in the profit share, $\bar{\pi}$, while it is decreasing in the discount rate, ρ , and in the size of the research team required for entry, h . (Of course, h becomes irrelevant if there is no entry.)

The consequence of aggregate innovation for growth requires returning to the representative consumer's preferences. Only the highest quality version of good j is produced in equilibrium. Its quality is $z_t(j) \equiv z(j, J_t(j)) = \prod_{k=0}^{J_t(j)} q(j, k)$. Taking account of how much is consumed, its contribution to utility is

$$\ln \sum_{k=-1}^{J_t(j)} x_t(j, k) z(j, k) = \ln [x_t(j) z_t(j)] = \left[\sum_{k=0}^{J_t(j)-1} \ln q(j, k) \right] - \ln w.$$

Looking across goods and exploiting the law of large numbers, the average value of $J_t(j)$ is μt and the geometric mean of $q(j, k)$ is $\bar{q} \equiv \exp \int_1^\infty \ln q \, d\Psi(q)$. Thus,

$$\ln C_t = \mu t \ln \bar{q} - \ln w.$$

Since the wage is constant, the equilibrium growth rate of the economy is $g = \mu \ln \bar{q}$, i.e. the rate of innovation times the expected percentage size of an innovative step.

In the case of no research by incumbents, the rate of growth is $g = [\bar{\pi}L/h - (1 - \bar{\pi})\rho] \ln \bar{q}$.

This expression is identical to that in the quality ladders model of Grossman and Helpman (1991, ch. 4) if innovative steps are a constant size $q = \bar{q}$, so that $\bar{\pi} = 1 - q^{-1}$. However, if $L_R > 0$ and $L_S > 0$, then there is an additional term $\lambda - l'_R(\lambda)/h$ in the expression for μ , representing the contribution of incumbents to aggregate innovation less the innovation that would have occurred if the researchers employed by incumbents were instead employed by potential entrants. This term is positive since the knowledge capital accumulated by incumbent firms is a productive resource in generating new innovations. The innovative productivity of the last researcher at an incumbent firm is equated to her productivity at a start-up, $l'_R(\lambda) = h$, but inframarginal researchers are more productive at incumbent firms than at start-up firms.

Our goal in this general equilibrium analysis is not to evaluate all the properties of yet another model of aggregate growth. This one has a number of obvious shortcomings that would make welfare analysis rather dubious.¹⁵ Instead, our aim is to demonstrate that our model of innovating firms hangs together in general equilibrium and, in the process, to clarify our earlier partial equilibrium assumptions. One additional point of clarification concerns our choice of aggregate expenditure as numeraire (so that a firm's revenue per good is always one). In terms of this numeraire we determined that the expected instantaneous growth rate of a firm's revenue is $-(\mu - \lambda) = -\eta$; i.e., with entry incumbent firms are expected to shrink. In a growing economy, however, we would not typically measure firm revenue as a share of GDP. A more natural measure of firm revenue would deflate by the appropriate price index for the economy, $P_t = 1/C_t = w\bar{q}^{-\mu t}$. Since this price index falls at rate $g = \mu \ln \bar{q}$, expected firm revenue growth in real terms is $-\eta + \mu \ln \bar{q}$. If the economy grows rapidly it is quite possible for incumbent firms to be expanding on average in real terms even as the expected number of goods each one produces is shrinking.¹⁶

4 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.

4.1 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?

4.1.1 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 (measured by the number of past innovations of the firm that are still in use or, equivalently, the number of goods produced by the firm) 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). Here the model is somewhat tentative and ad hoc. We

posit exogenous permanent differences across firms in the size of the innovative steps embodied in their innovations, with research costs increasing in the size of the step. Although larger innovative 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 innovative 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 innovative steps of a particular size. We leave this difficult question for future work.

4.1.2 R&D and Patenting

If we equate patents with innovations, we can use the firm-level innovation production function (1) to interpret the patent-R&D relationship. To do so we use the firm's R&D policy $R(n)$ to substitute knowledge capital n out of the innovation production function. In the case of equal sized innovative steps, $I = RG(1, 1/c(\lambda))$. Since innovation intensity λ is a constant, the model implies that patents should have a Poisson distribution with a mean proportional to firm R&D.

In the case of heterogeneous research intensity, but continuing to assume one patent per innovation, the model predicts that patents should rise less than proportionally with R&D across firms. The reason is that variation in R&D then reflects not only differences in knowledge capital but also heterogeneity in the size of firms' innovative steps q ; i.e., $R_q(n) = n(1-q^{-1})c(\lambda)/\bar{\pi}$. Holding q fixed, an increase in R leads to a proportional increase in I . But, holding n fixed, an increase in

R due to a higher innovative step acts only through the first argument of the innovation production function. In either case the model predicts a strong positive relationship between patents and R&D both across firms and over time for a given firm, consistent with stylized fact 2.

4.1.3 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? While a firm's patents may indicate how many innovations it has made, higher productivity reflects larger innovative steps. Since a firm's R&D intensity is also increasing in the size of its innovative steps, we predict the observed positive correlation between R&D intensity and productivity across firms (stylized fact 1).

To make this argument precise, consider a firm taking innovative steps of size q . The step size gives the firm market power, which it exploits by setting the price of its products equal to a markup q over its constant unit labor cost w . 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, so that $a = qw$. Heterogeneity in the size of innovative steps across firms produces variation in this measure of productivity, and this variation is positively correlated with variation in R&D intensity. Since we assume that the size of innovative steps is a characteristic of a firm, we predict persistent differences in productivity.¹⁷

4.2 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 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 outcomes. 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.

4.3 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 set of industries in which innovation plays a major role. With firm-level panel data on R&D, patenting, employment, and revenue from such industries, we could subject the model to a more detailed quantitative assessment. If it survives such an assessment, the model could be used to explore difficult questions about the interactions between industry evolution and technological change.

Another direction is to pursue the model's implications for policy. Unlike many of its predecessors, in our model incumbent research firms play an important role in driving aggregate technological change.¹⁸ 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 firms that do R&D. We see a potential for extending the analysis here to address questions that frequently arise concerning policies to promote innovation.

APPENDIX A: Discussion of Evidence on 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., Hall (1996) and Griliches (1998, ch. 12; 2000, ch. 4). 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 and R&D are positively related both across firms at a point in time and across time for given firms.*

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 proportional, whereas many smaller firms exhibit significant patenting while reporting very little R&D. 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), however, argues that small firms in available samples are not representative but are typically more innovative than the average small firm. Furthermore, he notes that “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

There is also a robust patents-R&D relationship in the within-firm 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]). Summarizing the results in Hall et al. (1986) and other studies, Griliches reports that the elasticity of patents with respect to R&D is between 0.3 and 0.6. Revisiting the evidence with new econometric techniques, Blundell et al. (2002) report a preferred estimate of 0.5.

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). However, 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, 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 (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 nonperformers. 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) 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” (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.*

The skewed size distribution of firms has been recognized for a long time and is discussed in Ijiri and Simon (1977), Schmalensee (1989), and Stanley et al. (1995). As noted by 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 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). It is the focus of recent research by Amaral et al. (1998) and Sutton (2001).

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 these patterns among new entrant firms.

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 $f(\lambda) = v$, where λ is the firm's optimal innovation intensity. Furthermore, either $\lambda = 0$ and $f(0) \leq c'(0)$ or else $\lambda > 0$ and $f(\lambda) = c'(\lambda)$.

The uniqueness of λ can be shown as follows. Differentiating $f(x)$, we get $f'(x) = [f(x) - c'(x)]/[r + \mu - x]$. Thus, for $x \in [0, \mu]$, if $f(x) > c'(x)$ then $f'(x) > 0$, and if $f(x) < c'(x)$ then $f'(x) < 0$. Furthermore, by assumption $c''(x) > 0$ and $f(\mu) \leq c'(\mu)$. We must consider three cases: (i) If $f(0) \leq c'(0)$ then $f(x) < c'(x)$ for all $x \in (0, \mu]$, and hence $\lambda = 0$. (ii) If $f(\mu) = c'(\mu)$ then $f(x) > c'(x)$ for all $x \in [0, \mu)$, and hence $\lambda = \mu$. (iii) The remaining case is when both $f(0) - c'(0) > 0$ and $f(\mu) - c'(\mu) < 0$. By the intermediate value theorem there will be a number $\lambda \in (0, \mu)$ such that $f(\lambda) - c'(\lambda) = 0$. This λ is unique since for all $x \in [0, \lambda)$ we have $f(x) - c'(x) > 0$ and for all $x \in (\lambda, \mu]$ we have $f(x) - c'(x) < 0$.

Suppose $0 < \lambda < \mu$. 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 a fall in λ .

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 date $t + \Delta t$ satisfies the relationship

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

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

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

The pgf is analog to the moment-generating function for a distribution of the continuous type. It contains all the information in the underlying distribution, but is much more convenient to work with in carrying out certain derivations.²⁰

Differentiating the pgf with respect to z yields

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

while differentiating with respect to t gives

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

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 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

$$H(z, t; n_0) = \sum_{n=0}^{\infty} p_n(t; n_0) z^n = \left[\frac{\mu(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)}{\lambda(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)} \right]^{n_0}, \quad (28)$$

where $p_n(t; n_0) = P[N_t = n | N_0 = n_0]$.

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

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

$$p_0(t; n_0) = H(0, t; n_0).$$

Equation (28) also provides expressions for the moments of the distribution:

$$\sum_{n=0}^{\infty} n^k p_n(t; n_0) = \left[\frac{\partial^k H(e^s, t; n_0)}{\partial s^k} \right]_{s=0}. \quad (29)$$

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:

$$\lim_{\mu \rightarrow \lambda} H(z, t; n_0) = \left[\lim_{\mu \rightarrow \lambda} \frac{\mu(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)}{\lambda(z-1)e^{-(\mu-\lambda)t} - (\lambda z - \mu)} \right]^{n_0} = \left[\frac{\mu t(z-1) - 1}{\mu t(z-1) - z} \right]^{n_0}. \quad (30)$$

APPENDIX D: The Size Distribution

Derivation of (17):

In this derivation we assume both $\eta > 0$ and $\lambda > 0$. From (15) and (16), we have that

$$\dot{M} = \eta - \mu M_1. \quad (31)$$

In the steady state all the time derivatives in (15), (16), and (31) are zero. Starting with (31), the steady state requires

$$M_1 = \frac{\eta}{\mu}. \quad (32)$$

Substituting (32) into (16), we see that the steady state also requires

$$M_2 = \frac{\lambda \eta}{2\mu^2}. \quad (33)$$

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

$$M_n = \frac{\lambda^{n-1}\eta}{n\mu^n}, \quad n \geq 1. \quad (34)$$

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 earlier.

REFERENCES

Aghion, Phillipe, and Howitt, Peter. "A Model of Growth Through Creative Destruction." *Econometrica* 60 (March 1992): 323–51.

- Amaral, Luís A. Nunes; Buldyrev, Sergey V.; Havlin, Shlomo; Salinger, Michael A.; and Stanley, H. Eugene. “Power Law Scaling for a System of Interacting Units with Complex Internal Structure.” *Physical Review Letters* 80 (1998): 1385–88.
- Audretsch, David B. *Innovation and Industry Evolution*. Cambridge: MIT Press, 1995.
- Barro, Robert J. and Sala-i-Martin, Xavier. *Economic Growth*. New York: McGraw-Hill, 1995.
- Bernard, Andrew B., Eaton, Jonathan; Jensen, J. Bradford; and Kortum, Samuel. “Plants and Productivity in International Trade.” *American Economic Review* 93 (September 2003): 1268–90.
- Blundell, Richard; Griffith, Rachel; and Windmeijer, Frank. “Individual Effects and Dynamics in Count Data Models.” *Journal of Econometrics* 108 (May 2002): 113–31.
- Carroll, Glenn R., and Hannan, Michael T. *The Demography of Corporations and Industries*. Princeton: Princeton University Press, 2000.
- Caves, Richard E. “Industrial Organization and New Findings on the Turnover and Mobility of Firms.” *Journal of Economic Literature* 36 (December 1998): 1947–82.
- Cockburn, Iain; Henderson, Rebecca; and Stern, Scott. “Untangling the Origins of Competitive Advantage.” *Strategic Management Journal* 21 (2000): 1123–45.
- Cohen, Wesley. “Empirical Studies of Innovative Activity.” In *Handbook of the Economics of Innovation and Technological Change*, edited by Paul Stoneman. Oxford: Blackwell, 1995.
- Cohen, Wesley, and Klepper, Steven. “The Anatomy of Industry R&D Intensity Distributions.” *American Economic Review* 82 (September 1992): 773–99.
- . “A Reprise of Size and R&D.” *Economic Journal* 106 (July 1996): 925–51.
- Cohen, Wesley and Levin, Richard. “Empirical Studies of Innovation and Market Structure.” In *Handbook of Industrial Organization, Vol. II*, edited by Richard Schmalensee and Robert Willig. Amsterdam: North-Holland, 1989.
- Dunne, Timothy; Roberts, Mark J.; and Samuelson, Larry. “Patterns of Firm Entry and Exit in U.S. Manufacturing Industries.” *RAND Journal of Economics* 19 (Winter 1988): 495–515.

- Eaton, Jonathan, and Kortum, Samuel. "International Technology Diffusion: Theory and Measurement." *International Economic Review* 40 (August 1999): 537–70.
- Ericson, Richard, and Pakes, Ariel. "Markov Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62 (January 1995): 53–82.
- Geroski, Paul A. "An Applied Econometrician's View of Large Company Performance." *Review of Industrial Organization* 13 (June 1998): 271–93.
- Goel, Narendra S., and Richter-Dyn, Nira. *Stochastic Models in Biology*. New York: Academic Press, 1974.
- Griliches, Zvi. "Issues in Assessing the Contribution of Research and Development to Productivity Growth." *Bell Journal of Economics* 10 (Spring 1979): 92–116.
- . "Patent Statistics as Economic Indicators: A Survey." *Journal of Economic Literature* 28 (December 1990): 1661–1707.
- . "R&D and Productivity: Econometric Results and Measurement Issues." In *Handbook of the Economics of Innovation and Technical Change*, edited by Paul Stoneman. Oxford: Blackwell, 1995.
- . *R&D and Productivity: The Econometric Evidence*. Chicago: University of Chicago Press, 1998.
- . *R&D, Education, and Productivity: A Retrospective*. Cambridge: Harvard University Press, 2000.
- Grossman, Gene M., and Helpman, Elhanan. *Innovation and Growth in the Global Economy*. Cambridge: MIT Press, 1991.
- Hall, Bronwyn. "The Private and Social Returns to Research and Development." In *Technology, R&D, and the Economy*, edited by B. Smith and C. Barfield. Washington, D.C.: Brookings Institution and AEI, 1996.
- Hall, Bronwyn H.; Griliches, Zvi; and Hausman, Jerry A. "Patents and R and D: Is There a Lag?" *International Economic Review* 27 (June 1986): 265–83.

- Hall, Bronwyn, and Hayashi, Fumio. "Research and Development as an Investment." NBER *Working Paper no. 2973*. Cambridge, Mass.: NBER, May 1989.
- Henderson, Rebecca. "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Equipment Industry." *RAND Journal of Economics* 24 (Summer 1993): 248–70.
- Hopenhayn, Hugo A. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (September 1992): 1127–50.
- Howitt, Peter. "Endogenous Growth and Cross-Country Income Differences." *American Economic Review* 90 (September 2000): 829–46.
- Hymer, Stephen, and Pashigian, Peter. "Firm Size and Rate of Growth." *Journal of Political Economy* 70 (December 1962): 556–69.
- Ijiri, Yuri, and Simon, Herbert. *Skew Distributions and the Sizes of Business Firms*. Amsterdam: North-Holland, 1977.
- Johnson, Norman L.; Kotz, Samuel; and Kemp, Adrienne W. *Univariate Discrete Distributions*. New York: Wiley, 1993.
- Jones, Charles I. "R&D-Based Models of Economic Growth." *Journal of Political Economy* 103 (August 1995): 759–84.
- Jovanovic, Boyan. "Selection and the Evolution of Industry." *Econometrica* 50 (May 1982): 649–70.
- . "The Diversification of Production." *Brookings Papers on Economic Activity* (Microeconomics), (1993): 197–235.
- . "Fitness and Age: Review of Carroll and Hannan's 'Demography of Corporations and Industries.'" *Journal of Economic Literature* 39 (March 2001): 105–19.
- Karlin, Samuel, and Taylor, Howard. *A First Course in Stochastic Processes. Second Edition*. San Diego: Academic Press, 1975.
- Kendall, David G. "On the Generalized 'Birth-and-Death' Process." *Annals of Mathematical Statistics* 19 (1948): 1–15.

- Kleinknecht, Alfred. "Measuring R&D in Small Firms: How Much Are We Missing?" *Journal of Industrial Economics* 36 (December 1987): 253–56.
- Klepper, Steven. "Entry, Exit, Growth, and Innovation Over the Product Life Cycle." *American Economic Review* 86 (June 1996): 562–83.
- Klepper, Steven, and Thompson, Peter. "Submarkets and the Evolution of Market Structure." Manuscript. Pittsburgh: Carnegie Mellon University, 2002.
- Klette, Tor Jakob. "R&D, Scope Economies and Plant Performance." *RAND Journal of Economics* 27 (Autumn 1996): 502–22.
- Klette, Tor Jakob, and Griliches, Zvi. "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous." *Journal of Applied Econometrics* 11 (July–August 1996): 343–61.
- . "Empirical Patterns of Firm Growth and R&D-Investment: A Quality Ladder Model Interpretation." *Economic Journal* 110 (April 2000): 363–87.
- Klette, Tor Jakob, and Johansen, Frode. "Accumulation of R&D-Capital and Dynamic Firm Performance: A Not-So-Fixed Effect Model." *Annales d'Economie et de Statistique* 49/50 (January/June 1998): 389–419.
- Kortum, Samuel S. "Research, Patenting, and Technological Change." *Econometrica* 65 (November 1997): 1389–1419.
- Langlois, Richard N., and Robertson, Paul L. *Firms, Markets and Economic Change: A Dynamic Theory of Business Institutions*. New York: Routledge, 1995.
- Li, Chol-Won. "On the Policy Implications of Endogenous Technological Progress." *Economic Journal* 111 (May 2001): C164–79.
- Lucas, Robert E., Jr. "Adjustment Costs and the Theory of Supply." *Journal of Political Economy* 74 (1967): 321–34.
- Mansfield, Edwin. "The R&D Tax Credit and Other Technology Policy Issues." *American Economic Review, Papers and Proceedings* 76 (May 1986): 190–94.

- Nelson, Richard R. “Modelling the Connections in the Cross Section Between Technical Progress and R&D Intensity.” *RAND Journal of Economics* 19 (Autumn 1988): 478–85.
- Penrose, Edith. *The Theory of the Growth of the Firm*. Oxford: Basil Blackwell, 1959.
- Peretto, Pietro F. “Cost Reduction, Entry, and the Interdependence of Market Structure and Economic Growth.” *Journal of Monetary Economics* 43 (February 1999): 173–95.
- Romer, Paul M. “Endogenous Technological Change.” *Journal of Political Economy* 98 (October 1990): S71-S102.
- Schmalensee, Richard. “Inter-Industry Studies of Structure and Performance.” In *Handbook of Industrial Organization, Vol. II*, edited by Richard Schmalensee and Robert Willig. Amsterdam: North-Holland, 1989.
- Scott, John T. “Firm Versus Industry Variability in R&D Intensity.” In *R&D, Patents and Productivity*, edited by Zvi Griliches. Chicago: University of Chicago Press, 1984.
- Segerstrom, Paul S. “Endogenous Growth Without Scale Effects.” *American Economic Review* 88 (December 1998): 1290–1310.
- Simon, Herbert A., and Bonini, Charles P. “The Size Distribution of Business Firms.” *American Economic Review* 48 (September 1958): 607–17.
- Stanley, Michael H. R.; Buldyre, S. V.; Havlin, S.; Mantegna, R. N.; Salinger, M. A.; and Stanley, H. E. “Zipf Plots and the Size Distribution of Firms.” *Economics Letters* 49 (October 1995): 453–7.
- Sutton, John. “Gibrat’s Legacy.” *Journal of Economic Literature* 35 (March 1997): 40–59.
- . *Technology and Market Structure: Theory and History*. Cambridge: MIT Press, 1998.
- . “The Variance of Firm Growth Rates: The ‘Scaling’ Puzzle.” Discussion Paper no. EI/24, STICERD. London: London School of Economics and Political Science, 2001.
- Thompson, Peter, and Waldo, Doug. Growth and Trustified Capitalism. *Journal of Monetary Economics* 34 (December 1994): 445–462.
- Thompson, Peter. “Technological Opportunity and the Growth of Knowledge: A Schumpeterian Approach to Measurement.” *Journal of Evolutionary Economics* 6 (February 1996): 77–97.

———. “The Microeconomics of an R&D-Based Model of Endogenous Growth.” *Journal of Economic Growth* (December 2001): 263–83.

Uzawa, Hirofumi. “Time Preference and the Penrose Effect in a Two-Class Model of Economic Growth.” *Journal of Political Economy* 77 (July/August 1969): 628–52.

NOTES

Tor Jakob Klette died in August 2003. Tor Jakob was a wonderful friend and partner in research. It has been very difficult to make the final revisions to this paper without him. We thank John Dagsvik, Jonathan Eaton, Charles Jones, Boyan Jovanovic, Kalle Moene, Jarle Møen, Atle Seierstad, Peter Thompson, Galina Vereshchagina, and numerous seminar participants for helpful comments. 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 Federal Reserve Bank of Minneapolis or the Federal Reserve System.

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

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 Amaral et al. (1998) and Sutton (1998), among others.

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.

⁶Although it is 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 have 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.

⁷A similar innovation production function has been used by Hall and Hayashi (1989) and Klette (1996). Our justification is 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. Lucas (1967) and Uzawa (1969) 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.

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 is the discounted sum of past R&D by the firm, which he denotes K . 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 and $a = G(1, 1/c(\lambda))$. Note that the appropriate depreciation rate on past R&D is the rate of creative destruction.

⁹Our model implies that firm age matters for exit because it is a proxy for firm size. Conditional on size, younger firms are no more likely to exit. In the data, however, there appears to be an independent negative effect of age on firm exit rates (see Dunne et al. [1988]). Klepper and Thompson (2002) provide a promising resolution to this empirical shortcoming. They develop a model of firm growth which, like ours, is driven by firms expanding into new goods (new submarkets in their terminology). They allow these goods to come in random sizes so that the firm's

total revenue (size) is no longer equal to the number of goods it produces. Firm age and firm size therefore play independent roles as proxies for the underlying number of goods, which, as in our model, is what matters for the firm's survival.

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.

¹²In the general equilibrium version of the model developed below, differences in π arise from exogenous differences in the innovative step q according to $\pi = 1 - q^{-1}$. A firm taking large steps obtains a flow of profits close to one while a firm taking tiny innovative steps is essentially an imitator, and its π is close to zero. Empirical research in economics, sociology, and management science has emphasized the large and highly persistent differences in innovative strategies across firms. See, e.g., Henderson (1993), Cohen (1995), Langlois and Robertson (1995), Klepper (1996), Carroll and Hannan (2000), Cockburn et al. (2000), and Jovanovic (2001). Here we accommodate such differences but shed no light on why they arise.

¹³Firms of any type π choose the same innovation intensity λ and thus follow the same stochastic growth process. Therefore, the distribution of firm types π among entrants, given by Φ , carries over to the distribution of types among firms of any size.

¹⁴Note that λ and μ can be treated as constants, even out of the steady state, as shown in Section 3.4.1.

¹⁵As pointed out by Li (2001), welfare conclusions in this class of growth models are very sensitive to deviations from the restrictive logarithmic preferences that we have assumed. Further-

more, as with many growth models, ours suffers from the empirical shortcomings pointed out by Jones (1995). In particular, it cannot account for the observed upward trend in R&D while at the same time accounting for the lack of any upward trend in productivity growth. We believe that the model could be modified to address these issues, along the lines of Kortum (1997), Segerstrom (1998), Eaton and Kortum (1999), or Howitt (2000). Our reason for suspecting such generalizations of the model are possible is that they introduce decreasing returns to R&D in a manner external to the firm. We have not introduced such modifications here because we do not think the firm-level data have anything to say about which one is most appropriate.

¹⁶Similarly, although the wage w is constant in the steady state, the real wage w/P_t grows at rate $\mu \ln \bar{q}$.

¹⁷Our argument about how to interpret measures of firm-level productivity borrows from earlier work by Klette and Griliches (1996) and Bernard et al. (2003).

¹⁸Thompson and Waldo (1994), Barro and Sala-i-Martin (1995), Peretto (1999) have also introduced incumbent researchers into models of endogenous technological change.

¹⁹See Kleinknecht (1987) for more on this issue.

²⁰The solution procedure is described in detail in Kendall (1948) and in Goel and Richter-Dyn (1974) (in particular, ch. 2 and Appendix B).