Entry and Schumpeterian Profits: How Technological Complexity Affects Industry Evolution*

Philip E. Auerswald†
School of Public Policy, George Mason University
Kennedy School of Government, Harvard University
auerswald@gmu.edu
December 2007

Abstract

A large empirical literature has documented differences in Schumpeterian profits both among firms in single industries and between firms in different industries. Theorists have proposed various institutional and strategic factors to account for such differences, but have had relatively little to say about the manner in which technology affects entry and profits. In this paper I present a model in which persistent intra-industry differences in firm profitability arise as the outcomes of learning and imitation, and inter-industry differences in the persistence of above normal profits arise solely from production being more technologically complex in some industries than in others.

JEL Classification Numbers: D20, D21, L10, L11, L23, 014, 030. *Keywords*: complexity, intrafirm externalities, industry evolution, organizational learning, persistence of profits, production recipes, technology.

---

* I thank Karl Shell for helpful discussions and correspondence regarding this research at various stages. I also thank seminar audiences at the University of Washington, Los Alamos National Laboratory, the Western Economic Association annual meetings, and the New England Complex Systems Institute (NECSI) International Conference on Complex Systems for comments and suggestions. All errors are mine.

† Assistant Professor, School of Public Policy, George Mason University and Research Associate, Belfer Center for Science and International Affairs, Kennedy School of Government, Harvard University. Postal : 4400 University Drive, MS 3C6, Fairfax, VA, 22030-4444 USA. tel : 703 993-3787. fax : 703 993-2284.

1 Introduction

A new firm in an industry where the technology of production is simple may learn quickly to be highly efficient. However, the very characteristics of the technology that make learning rapid in such a "low-tech" industry are likely also to make imitation easy. At the other extreme, if the technology of production is too complex, learning will be so slow as to confer to incumbents no substantive edge over new entrants; any advantage a firm acheives over competitors will last only if the firm is able to institute rigid internal protocols preventing any deviation from its successful practices.

Employing reasoning with antecedents in Shell (1973) and Lippman and Rumelt (1982), in this paper I present a model in which persistent quasi-rents accrue to firms in industries where technology is of intermediate complexity—that is, where learning is rapid enough to confer competitive advantage, but imitation is sufficiently uncertain to deter entry.[1] I define the complexity of production as the extent to which a technical decision by one unit within the firm affects the productive efficiency of other units. Complexity affects the persistence of profits above the norm in two ways: (1) by making the firm learning process both slower and less regular; and (2) by decreasing the probability that an industry leader's method of production will be successfully duplicated by a rival (or rivals). Taken together, these effects create purely technological barriers to entry in industries in which production is of intermediate complexity, allowing incumbent firms to earn above normal profits that persist over time.

Learning by doing and imitation drive the dynamics. An incumbent firm experiments at each time step with a single small change to its method of production, or "production recipe." If the change improves the firm's overall efficiency, it is incorporated into the firm's production recipe; if it reduces overall efficiency the change is dropped and forgotten. The process of learning by doing is subject to strongly diminishing returns over time, as improvements to existing production methods become increasing difficult to find. A new entrant, in turn, can begin production by imperfectly imitating current best practice.

The specific stochastic structure I employ to model the firm learning process is a reduced form of that introduced in Auerswald, Kauffman, Lobo and Shell (2000) and Auerswald (2007), based in turn upon Kauffman and Levin's (1987) NK model of "fitness landscapes.[2] In Auerswald et al. (2000), the stochastic learning model is calibrated to empirical data on firm learning curves. It is in part with reference to this previous work that I argue that the model of learning in this paper is better grounded, both theoretically and empirically, than is the case in competing models.

An important point of reference for the current paper is Jovanovic (1982), which presented the first (and still among the most persuasive) theoretical accounts of the evolution of a competitive industry driven by the resolution of uncertainty at the firm level. In that paper, firm learning is a process of signal extraction. Each entrant into the industry is endowed with a fixed and unchanging θ, the level of efficiency. The true value of θ is obscured by noise. In the course of production the firm is able to reduce the variance in its beliefs about its own parameter θ. Since the firms do not effect price, they are unconcerned about the efficiencies of their competitors. With this simple model Jovanovic is able broadly to match evidence on firm entry, growth, and exit within a generic industry.[3] However, the Bayesian model that Jovanovic employs is arguably parsimonious to a fault: its sparse form leaves little room for any theory or discussion of the organizational structure underlying the learning process. Moreover, the application of Bayes' rule to incoming production signals suggests not the work of a team of production workers engaged in shop-floor experimentation—as envisioned by Arrow (1962)[4] and suggested by empirical evidence[5]—but rather that of an adept back-office administrator. In this sense the Bayesian model is best characterized as "learning *and* doing" rather than "learning *by* doing". In contrast the model of firm learning in this paper can, at least in theory, be mapped directly to a detailed graphical representation of organizational structure, which in turn has the potential to be compared directly with plant-level data on actual production processes.[6]

Jovanovic and MacDonald (1994) present an even more closely related model in which firms may improve their technology through both innovation and through imitation.[7] Motivating the paper is the intuitively appealing and empirically supported hypothesis that barriers to the acquisition of information are a determining factor in the persistence of technological diversity within industries.[8] A key result is that, even when imitation is possible (though costly), the long term equilibrium may involve a diversity of levels of technology within an industry. The main shortcoming of the paper is that it rules out entry, exit, and learning by doing.

This paper seeks to advance on these prior contributions by focusing on the manner in which entry and technology may combine to determine industry structure. Specifically, I suggest a link between technological complexity of an industry and the persistence of profits. Section 2 presents a motivating example. Section 3 develops the learning model with reference to fundamentals of microeconomic theory, and establishes the stochastic structure for learning by a single-plant firm. Section 4 describes the industry model, including rules for entry and exit; labor market equilibrium at each time period; and conditions for an inter-termporal equilibrium. Section 5 characterizes the dynamics for industries typified by simple and complex production processes and discusses the persistence of profits in the case of intermediate complexity. Section 6 concludes.

2 A Motivating Example

Gordon Moore, co-founder of Intel in 1968 and universally known as the originator of "Moore's Law," has offered the following recollection of the early years of Fairchild Semiconductor (the immediate predecessor to Intel) that motivates a number of the core concepts in the paper:

> Eight of us left [Shockley Semiconductor] and ended up starting Fairchild Semiconductor... We advertised for a general manager, and buried among the many salesmen who believed they could manage was an application from Ed Baldwin, the engineering manager for the Hughes Semiconductor operation, then one of

> the largest semiconductor companies in the world... Baldwin taught us most of the simple things every M.B.A. knows. He taught us that different parts of the organization should be established with different responsibilities... And everything was working fine: the development and the pre-production engineering for our process and first products was complete; we had a thick process-spec book that recorded all the detailed recipes; and we had interested customers.
>
> But then Fairchild had the first of its Silicon Valley spin-offs. One day we came to work and discovered that Baldwin, along with a group of people he had suggested we hire, were leaving to set up a competing semiconductor company (Rheem) just down the road. He and his group took with them the 'recipes' for manufacturing that we had developed. But they left behind something more valuable. He had graduated a class of engineer-managers who now had the ability to figure out how to do it alone. And we did. Our response to their departure was to compete technologically by improving upon (and manufacturing) the products that Baldwin had guided us to develop. Although a court eventually ruled that they had to return the copy of the spec book to us, it no longer mattered. (Moore and Davis 2001: pp. 4-5)

With his reference to Fairchild's "thick process-spec book that recorded all the detailed recipes" for production of semiconductors, Moore offers a tangible example of the concept of the "production recipe" as employed in this paper. The existence of such a recipe book points to a missing element in the textbook representation of technological change as simple outward movements of a production possibilities frontier (PPF).[9] Moore's comments suggest that improvements in the recipes used to combine inputs into a reliable product that meets market-determined specifications underlie shifts in the PPF. Production recipes can be thought of literally as the sort of "process spec-book" alluded to by Moore; they may also be thought of more generally as complete vectors of decisions made at any time period by the core productive units of the firm. As Moore states, the principle that "different parts of the organization should be established with different responsibilities" is among the "simple things every M.B.A. knows."[10]

Moore also provides two examples of the central role of (imperfect) imitation of best practice in the process of new firm creation: first, the hiring by the "traitorous eight"[11] of Ed Baldwin, an experienced manager from the leading incumbent firm, to guide the creation of Fairchild; subsequently, Baldwin's departure from Fairchild, along with a cadre of the firm's engineers, to form a competing "spin-off" firm. Linking both examples to the concept of a (broadly conceived, firm-wide) production recipe, we can describe one core strategy for spin-off firms as the attempt to copy the production recipe of a successful incumbent firm. The extent to which such copying can be successful is constrained by a number factors including, but not limited to, the underlying complexity of the production process, the strength and enforceability of trade secret laws, the relative importance of tacit knowledge in the production process.[12] In this paper, I focus in particular on the complexity of production as a factor limiting a new firm's ability to copy the practices of an incumbent firm.

Moore also emphasizes the overriding importance to firm survival of the ability to improve upon existing products and processes, stating that the response of Fairchild "loyalists" to Baldwin's departure was to "compete technologically by improving upon (and manufacturing) the products that Baldwin had guided us to develop." From the standpoint of growth and competitiveness in a dynamic environment, the mix of inputs that enter the production plant is less significant than the development of firm specific technology—that is, the firm's complete production recipe(s). In a later section of the same narrative, Moore refers to a fundamental decision early in the history of Intel "to avoid [the] split between R&D and manufacturing. We'd be willing to accept less efficient manufacturing for a more efficient transfer process. We made the R&D people actually do their development work right in the production facility and we have continued that (with some variation) ever since." (Moore and Davis 2000: p. 14) Taken in combination these statements suggest that the conventional distinction between discontinuous technical change that occurs as the result of explicit investment (R&D) and incremental "shop floor" learning by doing may, at least in this particular case, be exaggerated.[13]

3 Firm Learning

3.1 The Relationship Between Production Recipes and Neoclassical Technology

Let us start with a very simple production process involving just three distinct operations. To simplify further, let us say that each of these operations can be carried out in one of only two ways. In this case, the eight possible "production recipes" may be enumerated as the set $\{000, 001, 010, 011, 100, 101, 110, 111\}$. Each of these recipes will be associated with its own scalar measure of effectiveness. We will call the level of effectiveness the "organizational capital" associated with the recipe. For example, recipe 010 might be associated with organizational capital θ, and recipe 101 might be associated with organizational capital θ'. Let us arbitrarily say that recipe 010 is the best of the bunch, so its associated level of organizational capital is greater than the organizational capital associated with any of the other recipes.

How then do such production recipes relate to a standard, neoclassical, production function? As has been famously described by Nelson and Phelps (1966), Winter (1967), and Nelson and Winter (1982), and others, the neoclassical economic theory of production focuses on the relationship between inputs and outputs, and not the manner in which inputs are combined to create outputs—the recipe. A neoclassical production plan is a particular input-output relationship; in its simplest rendition, it is a point (x, y) where $x \geq 0$ is the quantity of the input and where $y \geq 0$ is the quantity of the output. Figure 1 shows the production possibilities of the firm, the shaded area $\mathbf{T}$, and three specific possible production plans labeled A, B, and C. The production function in this figure

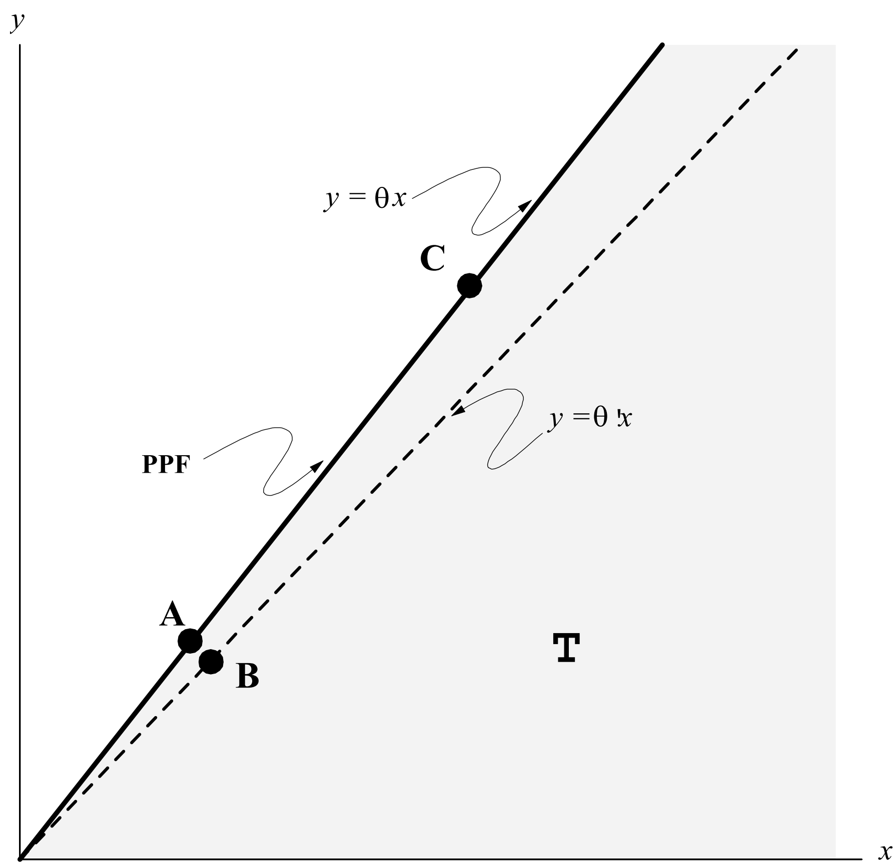

Figure 1: Neoclassical Technology

exhibits constant returns to scale, such that the best a firm can do is

$$y = \theta x, \tag{1}$$

where θ is a positive scalar that can be thought of as the organizational capital of the most productively efficient firm.[14] The production function is comprised of the set of input-output pairs that lie on the boundary of the production possibilities set. In this case, eq. 1 is the unique production function associated with the production possibilities given by the entire shaded area **T**.

In our simple, three-operation example, eq. 1 presumes that all firms (incumbents and potential entrants) are using production recipe 010. For example, the input-output pair A that lies on the boundary line as defined by eq. 1 clearly "dominates" input-output pair B: the firm using recipe 010 produces more output with less input than the firm using recipe 101. In general terms, the neoclassical production function requires that all firms have knowledge of the elements of the set of potentially usable recipes; (2) that all firms are aware of the effectiveness of each recipe in actual production; and, as a consequence, (3) that, in practice, all firms employ the same, most effective, recipe.

Consider figure 1 again, now with the objective of distinguishing technological from economic "distance." From an economic standpoint, input-output pair A is close to input-output pair B, but distant from input-output pair C. However, from the standpoint of technology, pairs A and C are the same, as they are produced with the same recipe (010); input-output pair B is maximally different from both A and C, in that the recipe use to produce B differs in every operation from that used to produce pairs A and C. Taking one operation at a time: $0 \neq 1; 1 \neq 0; 0 \neq 1$. Since there are three operations in all, and the two recipes differ in every operation, the technological distance between the two recipes is 3. We will use this notion of technological distance to describe the process of firm learning.

3.2 Production Experiments

Let us now consider more generally a firm in which production is comprised of N distinct operations, each of of which may be carried out in one of s ways. The total number of potential recipes, then, is given by s^N; in the simple case above with $N = 3$ and $s = 2$, the total number of possible recipes was $2^3 = 8$.

The industry begins at time 0 comprised of a finite number of firms employing recipes selected at random. In each time period thereafter, the firms systematically conduct small scale experiments to find improvements. As a consequence of organizational inertia, the experiments are incremental. Specifically, with each batch of production, exactly one of the N operating units within each incumbent firm undertakes a production experiment. (For example, a firm employing the recipe 010 might now try the recipe 011.) The trial production recipe is "sampled" without loss or gain in terms of current production.[15] Thus, if a firm starts production at time t with efficiency θ_t, it will maintain efficiency θ_t throughout the time period, regardless of the outcome of the production experiment that occurs during time t.[16] If the organizational capital associated with the trial recipe is given as θ_t^{lbd}, then the learning rule can be specified as

$$\theta_{t+1} = \begin{cases} \max\left\{\theta_t, \theta_t^{\text{lbd}}\right\} & \text{with probability } \eta; \\ \theta_t^{\text{lbd}} & \text{with probability } 1-\eta. \end{cases} \tag{2}$$

The probability η represents the likelihood that a change is made to the production recipe, but not noticed by management. In this case, the trial production recipe is adopted in the next time period even if it does not improve overall efficiency. If management's monitoring of the firm's production is perfect, then $\eta = 1$; if management is absent, then $\eta = 0$.

Note, importantly, that in the spirit of Arrow (1962) and related literature, the learning process in this paper is happenstance, analogous to that in Levinthal (1997). It stands in contrast to optimal learning approaches, such as that in Kauffman et. al (2000).

3.3 Intrafirm Externalities and Technological Complexity

When entrepreneurs or managers act to create or expand a firm they "internalize externalities," incorporating into the firm precisely those activities for which contracts are difficult to negotiate, for example due to multiple contingencies or high degrees of intrinsic uncertainty. This internalization of externalities implies some degree of inter-dependency between distinct production units brought together in any given firm. As a consequence, finding the optimal configuration of a firm's activities is much like finding the solution to a Rubik's cube puzzle: the creation, expansion, and management of the firm is made difficult by the fact that modification to the practices of one unit will affect the effectiveness of other units. Indeed, if one particular unit of a firm is not linked to any other via such "intrafirm externalities," then we can reasonably wonder why that unit is part of the firm to begin with (rather than, for example, acting as an outside contractor). Entrepreneurs and firm managers are thus typically charged with solving complex coordination problems.[17]

The complexity parameter, e, describes the magnitude of such internalized, intrafirm externalities, that are typical of a given industry. Higher values of e mean greater technological complexity.[18] This is the key parameter in the paper.[19]

To focus attention on the manner in which different levels of technological complexity drive inter-industry differences in the persistence of profits, in this paper I assume that e is exogenous—determined by the engineering and other technical principles underlying production in a given industry. With this assumption the parameter e can serve to distinguish one industry from another. Three types of industries are possible:

- $e = 1$ **(zero intrafirm externalities; simple production).** One limiting case is that in which there are no intrafirm externalities: A change in the production method employed by one of the N production units within the firm affects the efficiency only of that single unit. Each unit is "linked" to exactly one unit: itself. The average level of interconnection of the firm's production units, e, therefore is equal to 1.

- $1 < e \ll N$ **(partial intrafirm externalities; intermediate complexity).** Values of e such that $1 < e \ll N$ characterize production over a range of industries where a change in the production method by one of the N production units in the firm effects the efficiency of that unit, as well as some, but not all, of the other $N-1$ production units. In this range of industries the level of complexity of production (the average linkage of the firm's production units) is increasing in e. The argument above suggests that most industries fall in this category. If monitoring is perfect ($\eta = 1$), the firm's organizational capital grows monotonically over time as the result of learning by doing, following the process decribed above; alternately, so long as $\eta > 0$, organizational capital is stochastically increasing as a function of time.

- $e = N$ **(total intrafirm intrafirm externalities; maximal complexity).** The limiting value $e = N$ represents the case of maximal complexity: a change in the production method by one of the N production units within the firm affects not only that unit, but all of the other production units as well?

The results that follow are presented in terms of these three categories of industries.

The correlation between θ_t (current efficiency) and θ_t^{lbd} (the efficiency associated with the production experiment) depends on e in the following manner:

$$\theta_t^{\text{lbd}} = \underbrace{(N-e)\,\theta_t}_{\text{unaffected by time } t \text{ lbd}} + \underbrace{\sum_{i=1}^{e} \phi^i}_{\text{affected by time } t \text{ lbd}} \tag{3}$$

where the ϕ^i's are i.i.d. random variables drawn from a distribution $g(\phi)$ with mean μ which is common knowledge to all firms.

The first term on the RHS of equation 3 represents the component of firm's efficiency (or stock of organizational capital) that is unaffected by the experiment. This unaffected component represents roughly a fraction $\frac{N-e}{N}$ of the firm's total efficiency level. Note, importantly for the results, that the fraction of the firm's efficiency that is unaffected by the experiment is decreasing in the complexity of production. The second term on the RHS of equation 3 represents the contribution to the firm's efficiency (or stock of organizational capital) of the e units that are affected by the experiment. I model the contribution to the firm's efficiency of the affected components simply as the summation of e independent and identically distributed random variables, drawn from a distribution that is unchanging and known to all firms. Implicit in this construction is the view that firms experience as random events shocks to efficiency at the level of the production unit (that is, at the organizational scale that lies below that of the firm as a whole) resulting from incremental changes in production.[20]

4 Industry Evolution

The neoclassical baseline in the study of entry, profits, and industry dynamics is the partial equilibrium model of Viner (1932). In an industry with many firms, all sharing the same costs curves, free entry (or exit) in response to industry-wide profits (or losses) leads to a stable long-run equilibrium in which production is divided evenly among surviving firms, each operating at a scale corresponding to the minimum of the average cost curve. In a such a model scope for entry is created entirely on the demand side, with the optimal scale of any single firm limited by diminishing returns to scale.

In the current paper, firms have different cost curves. Entrants create new firms that almost perfectly reflect existing best practice. More precisely, "imitation" is minimally imperfect in the sense that the new entrant deviates from the best practice of the leading firm in the industry in modifying practices relating to only 1 of the N production units.

4.1 Static Profit Maximization

As my primary goal is to show that differences between industries in technological complexity may alone generate significant differences in industry dynamics, I construct the model so as to keep to a minimum sources of competitive friction other than those attributable to firm-specific innovation by doing. I follow a broad precedent in the literature on learning and industry evolution[21] by assuming that, even as learning occurs, firms remain price takers at all times. This, of course, does not directly imply that firms earn zero profits, even under conditions of free entry. (Price taking, of course, is a necessary, but not suffcient condition for perfect competition.) Consequently, there is no uncertainty in the model regarding the time-paths of Θ_t, the aggregate (industry-level) productivity level, or w_t, the wage rate in the industry; the only uncertainty pertains to θ_t productive efficiency at the firm level.

Recall that the firm's level of efficiency is fixed in any time period. With prices and efficiency fixed, employment is the only margin on which firms can operate. For this reason, the action in the model takes place in the labor market. With the level of efficiency fixed, production exhibits decreasing returns to scale with respect to labor (the only variable input).[22] More efficient firms are thus able to hire more workers before the productivity of the marginal worker falls to the market wage rate. Labor is perfectly mobile, so the market wage is the same for all firms when the labor market is in equilibrium.

Specifically, a firm with productive efficiency θ_t employs labor input l_t to produce x_t units of output during time period t, as given by the following production function:[23]

$$x_t = l_t^{\alpha} \left(\theta_t\right)^{1-\alpha},$$

Solving the first order conditions of firm i's static profit maximization problem with the output price (normalized to 1) given yields the following equation for a single firm's labor demand:

$$l_t = \theta_{i,t} \left(\frac{\alpha}{\mathrm{w}_t \left(\Theta_t ; \alpha\right)} \right)^{\frac{1}{1-\alpha}}.$$

Equilibrium in the labor market at time t occurs when the sum of all labor demands l_t equals aggregate labor supply. Substituting the equilibrium wage and equilibrium labor demand function into the profit function yields the indirect profit function:

$$\pi_t^* = \frac{\theta_t}{\left(\Theta_t\right)^{\alpha}} \left(1-\alpha\right) - c^{\text{fixed}} \tag{4}$$

Equation 4 indicates that the firm's time t profits are increasing in θ_t and decreasing in Θ_t: an incumbent firm's profitability increases when it improves its own efficiency, and decreases when rivals' improve their efficiency and/or when successful entry occurs.

4.2 Imitation

The difference in practice between the entrant's recipe and the leading incumbent may be deliberate or due to inherent imperfections in the process of imitation (e.g. due to tacit knowledge). The point of this assumption is to highlight scenarios under which profits may persist even when processes are legally protected by trade secret or patents.

Despite the fundamental similarity between the productions recipes employed by the leading incumbent and the potential entrant, the potential entrant's organizational capital may be very distant from incumbent's. This is due to the presence of intrafirm externalities. The entrant's modification of a single operating unit will affect the performance of exactly other units within the firm. The organizational capital level associated with the entrant's production recipe thus takes on the following stochastic form:.

$$\theta_t^{\text{entrant}} = \underbrace{\frac{(N-e)}{N}\theta_t^{\text{leading incumbent}}}_{\text{unaffected by imitation}} + \frac{1}{N}\underbrace{\sum_{i=1}^{e}}_{\text{affected by imitation}}\phi^i \tag{5}$$

where the ϕ^i's are i.i.d. random variables drawn from a distribution $g(\phi)$ with mean μ which is common knowledge to all firms. The higher the value of e, the lower the correlation between the incumbent's stock of organizational capital and that of the nearly-perfectly imitating entrant. Consequently the higher the value of e, the greater the difficultly (ease) of finding an improvement to a high (low) efficiency production method. The distribution of outcomes from imitation is (i) stochastically nondecreasing in θ_t for $e < N$; (ii) not a function of e for $e = N$; (iii) stochastically nondecreasing in e for $\theta_t > \mu$, and stochastically nonincreasing in e for $\theta_t < \mu$.[24] This learning rule represents a limiting case, allowing imitating firms to be as effective in copying the approach of the leading firm as is possible short of perfection.

The reader may note that, on the surface, new entrants are at an advantage over incumbents in the model as specified in that entrants can imitate best practice, while incumbents can not. This apparent asymmetry can be addressed without changing the model by simply categorizing as effective new entrants those incumbents who make the radical changes to their production method that are required to imitate current best practice. To approach radical organizational change as new entry is likely closer to corporate reality than would be the opposite assumption that corporate inertia is zero and dramatic changes to internal practices are possible at no direct cost to the firm.

4.3 Equilibrium

An industry consists of incumbents and potential entrants. In a given time period, an incumbent firm makes the following decisions simultaneously: first, considering the current time period, how much output to produce; second, considering future time periods, whether to remain in the industry after the end of the current period or to accept a privately known scrap value κ. Temporary shut-down is not an option. A potential entrant will enter the industry if the expected value of a newly created firm is greater than the cost of entry, κ_{entry}.

A competitive equilibrium in the model consists of a sequence of wages (w_t),[25] entry decisions, and exit decisions that jointly satisfy the static labor market equilibrium condition[26] and the following entry condition:

$$E_t\left[\left(\pi^*\left(\theta_t^{\text{entry}}\right)+(1+r)^{-1}V\left(\theta_{t+1}^{\text{lbd}},\Theta_t|1\right)\right)\right]-\kappa=0 \tag{6}$$

The intuition behind the process of equilibrium determination is as follows. Let the industry start out at time t in some arbitrary initial state (i.e. distribution of firms and efficiencies). Incumbent firms will calculate their expected profits for time $t+1$, taking into account the actions of potential entrants, and the conditional probability distributions governing learning and imitation. Potential entrants will, in turn, calculate their expected profits in the next period, taking into account the actions of incumbent firms. As there is no aggregate uncertainty in the model (due to the large number of firms), beliefs on which entry and exit decisions are conditioned will be self-fulfilling. The stochastic learning system specified above determines the realizations of θ_t^{ibd} (and therefore θ_{t+1}, efficiency levels in the coming period) for all incumbent firms. Meanwhile, current firm efficiency levels (θ_t for all firms) and the requirement of period-by-period labor market equilibrium determine time $t+1$ wage levels and profits. At this point the composition of the industry going into period $t+1$ is defined. The process then iterates forward one time period.

It is critical to note again that none of these conditions require profits to be zero. Rather, I assume only that the labor market clears and that incumbents and entrants are not systematically wrong in their perceptions of the future. The fact that current organizational capital is fixed builds a rigidity into the model that allows for persistence of profits, even as firms are price takers.

5 Persistence of Profits

Above I established that the complexity of production—the magnitude of intrafirm externalities—affects the accumulation of efficiency improvements within an industry in two related ways: first, by making learning more difficult, and second by making imitation more difficult. In this section I discuss how differences between industries in the accumulation of efficiency improvements translate into differences in the persistence of above normal profits.

5.1 Industries with simple production processes ($e = 1$)

In industries where production is simple ($e = 1$), the technological "landscape" faced by incumbents and new entrants alike has a single global maximum. This is the same as saying that any firm that follows the innovation by doing process specified above will soon arrive at the single best production recipe.[27] This well-known result differentiates the $e = 1$ case from all cases for which $e > 1$.[28]

Understanding the relative brevity of the transient during which innovation by doing takes place in the $e = 1$ case is important to understanding the results in the model. In the $e = 1$ case, learning for an early entrant is rapid, but terminates quickly once the optimal recipe is identified. With nearly perfect imitation of current best practice by later entrants, once a single incumbent firm has reached the global maximum characterized by organizational capital $\theta^{\max}$, a new entrant will require fewer than s time steps to arrive at the production recipe that maps to $\theta^{\max}$, where s is the number of "states" that a given operation can assume. (For example, any given production operation can be performed either "fast" or "slow," then $s = 2$). As a time step in the model corresponds with the production of a single "batch," the transient time is negligible. Disregarding the time spent learning, the equilibrium number of firms F^* solves the the equation:

$$E_t \left[\left(\pi^* \left(\theta_t^{\text{entry}} \right) + (1+r)^{-1} V \left(\theta_{t+1}^{\text{lbd}}, F_t^* \theta^{\max} | 1 \right) \right) \right] - \kappa = 0$$

The industry is comprised of a large number of firms, each using the same method of production and operating at a scale determined by the level of organizational capital $\theta^{\max}$. All firms will earn economic profits equal in expectation to the mean of the distribution of fixed costs of entry. As this mean approaches zero, the equilibrium becomes exactly that of Viner (1932). Zero profits here follow not by construction, but as a consequence of the simplicity of the underlying production problem "solved" by firms in the industry.

5.2 Industries with production processes of maximal complexity ($e = N$)

Now consider the other limiting case of an industry in which production is maximal complex. Like the $e = 1$ industry, the $e = N$ industry can be fully characterized analytically. In such an industry, any change carried out within a single production unit affects the productive efficiency of all other units in the firm. In terms of the model, this means that every time a production experiment takes place, a firm draws a new trial level of organization capital that is completely uncorrelated with its current level of organizational capital. As a consequence, "innovation by doing" takes the form of repeated, uncorrelated draws from the distribution of values of organizational capital.

Without correlation between successive draws, the model of technological search reduces to that of Muth (1986). However, as described above, simulation of the learning model over a wide variety of parameter values establishes that in the $e = N$ case, learning is very slow and irregular. In the additive specification in this model, the standard deviation of the sampled distribution decreases as e grows. The proof is a simple application of the law of large numbers. As a consequence, repeated draws in this case are less likely to draw values of θ far from the mean of ϕ that is the case for industries with lower values of e.

Of course, in such a maximally complex environment, imitation of existing best practice is effectively impossible. Any firm that draws a particularly high value of organizational capital (either upon entry or in the process of learning) will be able to earn persistent profits under so long as the firm is able to retain the the high value recipe over time without error ($\eta = 1$).

However, if monitoring by management is even slighly imperfect (that is, for any value of η less than 1), the situation is quite different. Experience suddenly counts for little. If at any point the firm "forgets" part of the recipe, or otherwise makes a modfication to practice that is not recoverable, then the firm will revert to the mean of the population in expectation. Profits thus do not persist.

5.3 Industries with production processes of intermediate complexity $(1 < e \ll N)$

Now consider the most interesting case of an industry in which production is of intermediate complexity ($1 < e \ll N$). In this case, engineering and other exogenously given fundamentals in the industry are such that technical decisions made by one production unit within a firm will necessarily affect the efficiency of some, but not all, of the other production units within the same firm. As in the case of an industry characterized by simple and maximally complex technology, firms will innovate by doing in the process of production. However, due to the presence of intrafirm externalities, the technological landscape searched by firms in industries characterized by $e > 1$ will have many local optima. Learning by firms in such industries will be slower and more irregular that in industries for which the underlying technology is simple ($e = 1$), but more rapid than industries for which the underlying technology is maximally complex ($e = N$). Furthermore, a new entrant that imperfectly imitates a successful firm will find its starting efficiency to be poorly correlated with that of the incumbent leader. Thus, even when (imperfect) imitation is possible, incentives to enter will diminish over time as the industry is increasingly populated by high efficiency firms. Where above normal profits rapidly dissipate in the low complexity case, in the intermediate complexity case above normal profits persist.

The formal result depends on the difference between the mean organizational capital of incumbents, which I denote by $\left\langle \theta^{\text{lop}} \right\rangle$ (where "lop" stands for "local optimum"), and the organizational capital expected by a new firm upon entry.[29] New entrants compete in the labor market with highly efficient incumbents, and will expect first period losses approximated as

$$E_t\left(\pi^0\right) \simeq \left(\theta^{\text{entry}} - \left\langle \theta^{\text{lop}} \right\rangle\right) \frac{p\left(1-\alpha\right)}{\left(F^* \left\langle \theta^{\text{lop}} \right\rangle\right)^{\alpha}}$$

The degree to which prospective new entrants will be deterred here is a function of the magnitude of the difference between θ^{entry} and the mean productive efficiency of "established" incumbents, $\left\langle \theta^{\text{lop}} \right\rangle$. The difference

$$\theta^{\text{entry}} - \left\langle \theta^{\text{lop}} \right\rangle \tag{7}$$

in turn is determined by the outcomes of learning and imitation.[30] The structure of the imitation process in particular directly implies that the difference given in equation 7 will be greater (in expectation) when e, the magnitude of intrafirm externalities (the complexity of production), is an intermediate value between 1 and N. The presence of intrafirm externalities thus creates purely technological barriers to entry which allow firms with high levels of organizational capital to earn above normal profits that persist over time.

6 Conclusion

Industrial economists have analyzed and documented numerous determinants of firm profitability. Some, when present, are likely to be common to all incumbent firms in a given industry: economies of scale;[31] sunk costs such as R&D and/or advertising intensity;[32] informational impediments to imitation;[33] and the extent of first-mover advantages, such as those due to consumer switching costs, network externalities and/or learning by doing.[34] Others are firm-specific, such as managerial ability, direct costs of production, and organizational capital.[35] In this paper I introduce technological complexity as an alternate way of explaining persistent interindustry differences in the persistence of profits. In industries where production processes are simple, I find that profits rapidly converge upon the norm, particular when imitation is possible. In industries where production processes are more complex, persistent profits accrue to surviving firms. Such profits are greatest in industries in the early stages of industries where technology is of intermediate complexity—that is, where learning is rapid enough to confer competitive advantage, but imitation is sufficiently uncertain to deter later entry.

The paper suggests at least two interesting paths of further research into the complexity of production and firm profitability. One path would be to endogenize the boundedness of learning by requiring that firms commit some fraction of production to the experimentation process (if experimentation is to occur). Imposing such a cost of experimentation brings the model closer to the realities of process R&D.

Another path would be to endogenize the complexity of production as defined here in terms of intrafirm externalities. This would involve explicitly modeling the process by which the boundaries of a "typical" firm are defined over time by the profit maximizing decisions of managers seeking to internalize (beneficial) external effects. Rather than assuming, as I have here, that the magnitude of intrafirm externalities is constant and the same for all firms in the industry, one could consider a distribution of levels of intrafirm externalities that evolves over time as the result of mergers, acquisitions and the formation of alliances.

Pursued along these lines, the results in this paper may provide a starting point for formal analysis of the process by which a “high-tech” product becomes a commodity, as famously occurred in the disk drive industry.[36] The suggestion is that the evolution of technology and practices through learning and standardization may transform an industry in which imitation is difficult into one in which imitation is easy, with profits dissipating accordingly.

This paper also opens the door to study of contexts in which production may be decentralized among different firms, with organizational learning by any single firm in the network involving a subset of all of the routines in the production process. In such settings, industry-wide learning may occur as a combination of incremental learning by existing firms, and discontinuous technological change may drive a wave of entry and exit within the network.

7 Appendix

Normalizing the aggregate supply of labor to 1, we have by substitution

$$L \equiv 1 = \sum_i l_{i,t} \tag{8}$$

$$= \sum_i \theta_{i,t} \left(\frac{\alpha p}{\mathrm{w}_t \left(\theta_{i,t}, \Theta_t \right)} \right)^{\frac{1}{1-\alpha}} \tag{9}$$

$$= \Theta_t \left(\frac{\alpha p}{\mathrm{w}_t \left(\theta_{i,t}, \Theta_t \right)} \right)^{\frac{1}{1-\alpha}} \tag{10}$$

where

$$\Theta_t \equiv \sum \theta_{i,t}. \tag{11}$$

Note that relatively unproductive firms enter equation 11 with the same weight as productive firms. At first glance, this could seem strange, as we might expect productive firms to bid up the wage, while unproductive firm will not be able to bid it down. While it is true that a more productive firm is in a position to offer a higher wage than an unproductive firm, it can offer that wage only to a limited number of workers. The productive firm *prefers* to hire *additional* workers so long as the productivity of the marginal worker exceeds the market clearing wage rate. Given the process of wage determination implicit in equation 8 (i.e. the action of a standard Walrasian auctioneer) the identity expressed in equation 11 is straight algebraic substitution, and requires no additional or particular assumptions.

Rearranging the terms in equation 10, we can express the equilibrium wage for a single period as a function of the price and aggregate productivity levels:

$$\mathrm{w}_t^* \left(\Theta_t; p, \alpha\right) = \left(\Theta_t\right)^{1-\alpha} p\alpha, \tag{12}$$

Similarly, the equilibrium labor demand is given by

$$l_{i,t}^* = \frac{\theta_{i,t}}{\Theta_t} \tag{13}$$

At the wage rate given in equation 12, the labor demands of all firms which commit to producing in period t are satisfied and there is no unemployment.

Notes

[1]See also Auerswald (1999). Rivkin (2000) developed in parallel a very similar model conveying the insight that complex strategies deter imitation. Klepper (2002) describes the manner in which firm-level experience with technology drove the evolution of oligoploly in the cases of autos, tires, televisions, and pencillin.

[2]The NK-model has been applied previously to production theory, organizational theory, and industrial economics. In addition to Auerswald (1999) and Rivkin (2000), other applications of the NK-model to industrial economics and organizational theory include Kauffman (1988), Levinthal (1997), Kauffman, Lobo, and Macready (2000). In the model that follows, my parameter N is directly analogous to N in the NK model, and my parameter e is directly analogous to $K+1$ in the NK model.

[3]In particular, smaller firms in the Jovanovic model have higher and more variable growth rates than larger firms, in accordance with the evidence from Mansfield (1962) and others.

[4]See also Solow (1997). He contrasts activities "that one would clearly define as research" with the "'continuous improvement' of products and processes," which consists of "an ongoing series of minor improvements in the design and manufacture of standard products." He elaborates (p. 20):

> These improvements usually arise somewhere close to the factory floor. They may have nothing to do with the sort of people who are engaged in research and they are not the product of a research activity. If one had to give a shorthand description of the process of continuous improvement, 'learning by doing' would serve pretty well.

[5]See for example the survey article on learning curves by Dutton and Thomas (1984) which presents a typology of sources of firm learning.

[6]The analysis by Ulrich and Pearson (1998) of the cost components involved in the manufacture of drip coffee makers suggests one such mapping.

[7]Jovanovic and MacDonald (1994) solve for a stationary equilibrium with the following characteristics: the distribution of know-how converges (asymptotically) to a long-run distribution; product price decreases over time; average output stochastically increases with time; the likelihood of learning a better technique is decreasing in θ (smaller firms will innovate more frequently, but they will not necessarily grow faster); and a central planner prefers more learning effort than competition yields.

[8]Adjustment cost/vintage capital models are the leading alternative explanation.

[9]In an earlier expression of a similar concept, Hirsch (1969) observes that "[l]ike a cake recipe, technology specifies ingredients as well as sequence of operations. In the neoclassical representation of technological change, improvements in knowledge and techniques result in an alteration of the productivity of inputs; a new production function results." Romer (1996: 204) employs the term "recipe" in reference to the fundamental encoding of economically relevant technical ideas that transform things from lower to higher valued configurations: "Ideas are nonrival goods that could be stored in a bit string. Things are rival goods with mass (or energy). With ideas and things, one can explain how economic growth works. Nonrival ideas can used to rearrange things, for example, when one follows a recipe and transforms noxious olive into tasty and healthful olive oil. Economic growth arises from the discovery of new recipes and the transformation of things from low to high value configurations."

[10]In analogous formulations Nelson and Winter (1982) have described the practices of the different parts of an organization acting jointly as the firm's production "routine."

[11]"Traitorous eight" is the term used by Moore to refer to the eight engineers, including himself, who left Shockley Semiconductor to found Fairchild.

[12]Klepper and Sleeper (2005) explore in detail the types of firms that spawn spinoffs, the market conditions conducive to spinoff, and the relationship of spinoffs to their "parents".

[13]As the first mover and dominant firm in an industry where dynamic returns to scale through learning are well documented, Intel may have had a particularly strong strategic incentive to focus upon speed to market of incremental improvements on its existing products.

[14]Prescott and Visscher's (1980) concept of "organization capital", which may be thought of as a cardinal measure of the level of the firm's knowledge about production.

[15]The model presented here can be considered a limiting case in which only one production experiment can be conducted per "batch" of output, and producing a "batch" of requires a fixed time regardless of the batch size. The batch is small as compared with the firm's output, so the firm does not commit to the change in its production recipe.

[16]Under these assumptions, there is little room for strategic price setting, quantity setting or other entry-deterring behaviors broadly characteristic of imperfectly competitive models–particularly monopolistic or oligopolistic models of learning by doing. See e.g. Spence (1981), Fudenberg and Tirole (1983), Smiley and Ravid (1983), and Cabral and Riordan (1994). Firms in this model do, in theory, have the latitude to set wages in a strategic manner. In the absence of intertemporal economies of scale, however, they have little incentive to do so.

[17]A classic paper by Reiter and Sherman (1962) entitled "Allocating Indivisible Resources Affording External Economies or Diseconomies," anticipates recent work on the firm as a solver of hard combinatorial optimization problems.

[18]Clearly, other definitions are possible. For example, Jovanovic and Nyarko (1997) associate "complexity" with the number of tasks carried out by the firm, and Radner (1993) emphasizes the information-processing demands on managers.

[19]The e here is directly analogous to the K in the NK-model. The parameter e can take on integral values between 1 and N.

[20]Here the model bears some similarity to the less purposive approaches of Simon and Bonini (1958), Hopenhayn (1992), and Atkeson and Kehoe (1997). However, where in these papers purely stochastic shocks occur at the scale of the firm as a whole, here they occur of the scale of the production unit.

Auerswald et al. (2000) analyze the comparative dynamics of such a model, emphasizing the effects of varying the following basic parameters of the model: the number of operations (N); the number of instructions per operation (s); the "intra-firm externality" parameter (e); the maximum number of changes in a recipe per trial; the number of trials per measured batch; and the length in time of the production run. Calibration of the model is based primarily upon two quantitative points of reference: the sample mean of the estimated progress ratios, and the sample standard deviation of the estimated progress ratios. The results pertaining to e are the most pertinent for this paper. If $e = 1$ (simple production), it is possible to demonstrate analytically that, regardless of the values of the other parameters, any firm searching the technology landscape will, relatively rapidly, find the global low-cost recipe. When the number of elements in the recipe space is small (e.g., when $s = 2$). Increasing technological complexity actually increases rates of firm learning for small e. In all other cases, increasing technological complexity slows learning and makes the learning process less regular. For a range of values of the other parameters, the best fit of a calibrated model to observed learning curves, both in terms of rates of learning and of other observed characteristics (plateauing, intermittent progress, and initial curvature in log-log plot) comes for production processes in which each production unit is linked via "intrafirm externalities" to approximately 5% of other production units in the firm (for example, if $N = 100$, $e \simeq 5$).

[21]This includes Jovanovic (1982), Hopenhayn (1992), Jovanovic and MacDonald (1994), Atkeson and Kehoe (1997), Petrakis, Rasmusen and Roy (1997), and Mitchell (2000).

[22]The assumption of price-taking in the presence of learning by doing is plausible only if production exhibits decreasing returns to scale. Accordingly, production in this model exhibits decreasing returns to scale when holding fixed the level of organizational capital, θ_t. As observed by Petrakis, Rasmusen and Roy (1997, p. 249): "[I]f the average cost curve at any point of time is constant in current output, then learning introduces an intertemporal economy of scale that creates a natural monopoly. This need not be the case, however, if the technology displays su.ciently decreasing returns. In that case, learning does not lead to a natural monopoly and is, in fact, compatible with competition."

[23]Rosen (1972), Lucas (1978) and Atkeson and Kehoe (1997) employ analogous functional forms: standard Cobb-Douglas production functions (homogeneous of degree 1) with some measure of productive efficiency included as one of the inputs. For Lucas θ is the managerial "span of control;" for Atkeson and Kehoe it is "organization capital".

[24]See Auerswald (2007).

[25]The cost of supervision as determined in the industry equilibrium can be thought of as the wage paid to managers, whose abilities are distinct from production workers. In this interpretation, there are two separate labor markets.

[26]Equation 13 in the Appendix.

[27]The relative time required to reach a local optimum as a function of the level of intrafirm externalities, e, has been formally derived in the context of the NK model, for example by Kauffman and Levin (1987) and Macken and Perelson (1989).

[28]See Durrett and Limic (2001) for a thorough characterization of optima in NK models (of which the technology model presented in this paper is an instance).

[29]Again, Durrett and Limic (2001) offer the formal characterization.

[30]The probability that a potential entrant will draw a starting level of organizational capital that is below the mean the organizational capital of incumbents is given by

$$\Pr\left(\theta^{\text{entry}} \leq \left\langle \theta^{\text{lop}} \right\rangle ; e\right) = H\left(z \mid \left\langle \theta^{\text{lop}} \right\rangle ; e\right) = \int_0^z h\left(z \mid \left\langle \theta^{\text{lop}} \right\rangle ; e\right) dz$$

It follows directly that the $H\left(z \mid \left\langle \theta^{\text{lop}} \right\rangle ; e\right)$ is (i) stochastically non-decreasing in $\left\langle \theta^{\text{lop}} \right\rangle$ for $e < N$; (ii) non a function of $\left\langle \theta^{\text{lop}} \right\rangle$ for $e = N$; (iii) stochastically non-decreasing in e for $\left\langle \theta^{\text{lop}} \right\rangle > \mu$, and stochastically non-increasing in e for $\left\langle \theta^{\text{lop}} \right\rangle < \mu$. With probability approaching 1 as t grows, the probability that a new entrant will draw a value below that of an incumbent is stochastically non-decreasing in e; however, for $e = N$, the expected difference between θ^{entry} and θ^{lop} is not a function of e.

[31]See e.g. Bain (1956) and Schmalensee (1981).

[32]See e.g. Dixit (1981) and Sutton (1991).

[33]See e.g. Lippman and Rumelt (1982), and Mansfield (1985).

[34]See e.g. Schmalensee (1982), Arthur (1989), and Smiley and Ravid (1983). Waring (1996) tests for the dependence of the persistence of above normal firm returns on a set of variables representing these and other theoretically relevant factors, and finds that following industry-level variables yield positive and significant estimates: skill, specialization, unionization, consumer purchases as a percentage of output, R&D intensity, and economies of scale.

[35]See, respectively, Jovanovic (1982), Lucas (1978), and Atkeson and Kehoe (1997).

[36]Christensen (1997).

*

References

Arrow, K. J.: 1962, The economic implications of learning by doing, *Review of Economic Studies* **29**, 155–73.

Arthur, W. B.: 1989, Competing technologies, increasing returns and lock-in by historical events, *Economic Journal* **99**, 116–131.

Atkeson, A. and Kehoe, P.: 1997, Industry evolution and transition: A neoclassical benchmark, *Working Paper 6005*, NBER.

Auerswald, P.: 1999, *Organizational Learning, Intrafirm Externalities and Industry Evolution.*, University of Washington Ph.D. thesis.

Auerswald, P.: 2007, Entrepreneurship in the theory of the firm, *Small Business Economics* p. forthcoming.

Auerswald, P., Kauffman, S., Lobo, J. and Shell, K.: 2000, The production recipes approach to modeling technological innovation: An application to learning by doing, *Journal of Economic Dynamics and Control* **24**, 389–450.

Bain, J.: 1956, *Barriers to New Competition*, Harvard University Press.

Cabral, L. and Riordan, M.: 1994, The learning curve, market dominance and predatory pricing, *Econometrica* **62(5)**, 1115–1140.

Christensen, C. M.: 1997, *The Innovator's Dilema: When New Technologies Cause Great Firms to Fail*, Boston: Harvard Business School Press.

Dixit, A. K.: 1981, The role of investment in entry deterrence, *Economic Journal* **90(357)**, 95–106.

Durrett, R. and Limic, V.: 2001, Population genetic consequences of evolution on random landscapes, *mimeo*, Department of Mathematics Cornell University, Ithaca NY.

Dutton, J. M. and Thomas, A.: 1984, Treating progress functions as a managerial opportunity, *Academy of Management Review* **9(2)**, 235–247.

Fudenberg, D. and Tirole, J.: 1983, Learning by doing and market performance, *Bell Journal of Economics* **14**, 522–530.

Hirsch, W. Z.: 1969, Technological progress and microeconomic theory, *American Economic Review* **59(2)**, 36–42.

Hopenhayn, H.: 1992, Exit, entry, and firm dynamics in long run equilibrium, *Econometrica* **60(5)**, 1127–1150.

Jovanovic, B.: 1982, Selection and the evolution of an industry, *Econometrica* **50**, 659–670.

Jovanovic, B. and MacDonald, G.: 1994, Competitive diffusion, *Journal of Political Economy* **102(1)**, 24–52.

Jovanovic, B. and Nyarko, Y.: 1997, Learning by doing and the choice of technology, *Econometrica* **64(6)**, 1299–1310.

Kauffman, S.: 1988, The evolution of economic webs, *in* P. W. Anderson, K. J. Arrow and D. Pines (eds), *The Economy as an Evolving Complex System*, Addison Wesley, pp. 125–146.

Kauffman, S. and Levin, S.: 1987, Toward a general theory of adaptive walks on rugged landscapes, *Journal of Theoretical Biology* **128**, 11–45.

Kauffman, S., Lobo, J. and Macready, W.: 2000, Optimal search on a technology landscape, *Journal of Economic Behavior and Organization* **43**, 141–166.

Klepper, S.: 2002, Firm survival and the evolution of oligopoly, *RAND Journal of Economics* **33(1)**, 37–61.

Klepper, S. and Sleeper, S.: 2005, Entry by spinoffs.

Levinthal, D. A.: 1997, Adaptation on rugged landscapes, *Management Science* **43**, 934–950.

Lippman, S. and Rumelt, R.: 1982, Uncertain imitability: An analysis of interfirm differences in efficiency under competition, *The Bell Journal of Economics* **13**, 418–438.

Lucas Jr., R. E.: 1978, On the size and distribution of business firms, *The Bell Journal of Economics* **9**, 508–523.

Macken, C. A. and Perelson, A. S.: 1989, Protein evolution on rugged landscapes, *Proceedings of the National Academy of Sciences* **86**, 6191–6195.

Mansfield, E.: 1962, Entry, Gibrat's law, innovation, and the growth of firms, *American Economic Review* pp. 1022–1051.

Mansfield, E.: 1985, How rapidly does new industrial technology leak out?, *Journal of Industrial Economics* **34(2)**, 217–223.

Mitchell, M.: 2000, The scope and organization of production: Firm dynamics over the learning curve, *RAND Journal of Economics* **31(1)**, 180–205.

Moore, G. and Davis, K.: 2001, Learning the silicon valley way, *Working Paper 00-45*, Stanford Institute for Economic Policy Research.

Muth, J.: 1986, Search theory and the manufacturing progress function, *Management Science* **32**, 948–962.

Nelson, R. R. and Phelps, E. S.: 1966, Investment in humans, technological diffusion, and economic growth, *American Economic Review* **56**, 69–82.

Nelson, R. R. and Winter, S.: 1982, *An Evolutionary Theory of Economic Change*, Belknap, Harvard U. Press.

Petrakis, E., Rasmusen, E. and Roy, S.: 1997, The learning curve in a competitive industry, *RAND Journal of Economics* **28(2)**, 248–268.

Prescott, E. C. and Visscher, M.: 1990, Organization capital, *Journal of Political Economy* **88(3)**, 446–61.

Radner, R.: 1993, The organization of decentralized information processing, *Econometrica* **61(5)**, 1109 – 1146.

Reiter, S. and Sherman, G. R.: 1962, Allocating indivisible resources affording external economies or diseconomies, *International Economic Review* **3(1)**, 108–135.

Rivkin, J.: 2000, Imitation of complex strategies, *Management Science* **46**, 824–844.

Romer, P. M.: 1996, Why, indeed, in America? Theory, history and the origins of modern economic growth, *American Economic Review* **86(2)**, 202–206.

Rosen, S.: 1972, Learning by experience as joint production, *Quarterly Journal of Economics* **86**.

Schmalensee, R.: 1981, Economies of scale and barriers to entry, *Journal of Political Economy* **89(6)**, 1228–1238.

Schmalensee, R.: 1982, Product differentiation advantages of pioneering brands, *American Economic Review* **72(3)**, 349–365.

Shell, K.: 1973, Inventive activity, industrial organisation and economic growth, *in* J. A. Mirrlees and N. H. Stern (eds), *Models of Economic Growth*, John Wiley and Sons, pp. 77–100.

Simon, H. A. and Bonini, C.: 1958, The size distribution of business firms, *American Economic Review* **48(4)**, 607–617.

Smiley, R. H. and Ravid, S. A.: 1983, The importance of being first: Learning price and strategy, *Quarterly Journal of Economics* pp. 353–362.

Solow, R. M.: 1997, *Learning from 'Learning by Doing': Lessons for Economic Growth*, Stanford University Press.

Spence, A.: 1981, The learning curve and competition, *The Bell Journal of Economics* **12**, 49–70.

Sutton, J.: 1991, *Sunk Costs and Market Structure: Price Competition, Advertising and the Evolution of Concentration*, MIT Press.

Ulrich, K. T. and Pearson, S.: 1998, Assessing the importance of design through product archaeology, *Management Science* **44(3)**, 352–369.

Viner, J.: 1932, Cost curves and supply curves, *Zeitschrift für Nationalökonomie* **3**, 23–46.

Waring, G. F.: 1996, Industry differences in the persistence of firm-specific rents, *American Economic Review* **86(5)**, 1353–1364.

Winter, S. G.: 1968, Toward a neo-Schumpeterian theory of the firm, *Working Paper P-3802*, RAND.