Learning, Contingencies and Firm Differentiation

Hugo A. Hopenhayn
UCLA

William Lovejoy
University of Michigan

June 9, 2009

1 Introduction

This paper studies organizational learning in a stable task environment, a term that will be defined below. The medium is a formal mathematical representation of a learning model introduced by March and Simon (1958), in which innovative problem solving yields to routinized problem solving over time. In our model, firms decide how to respond to profit opportunities faced, whether investing to acquire new specific knowledge or resorting to accumulated expertise. The decision of when and what to learn is done with perfect foresight, taking into account the dynamic consequences of these investments. Thus, our model provides a link between organizational learning and capital theory. The model is consistent with some empirical observations in a variety of literatures, as described below. Our model also offers some interesting additional insights.

First, the investment decisions made by a firm at any point in time are affected by the specific profit opportunities faced by the firm in the period and by the type of investments made in the past. Accordingly, specific opportunities and investments faced and adopted by a firm in the past have lasting effects on its characteristics and performance. As time evolves firms become differentiated according to the idiosyncratic opportunities confronted. This differentiation is persistent in that it may not close over time. In this way, the model provides a *contingency theory of organizational differentiation*.

Second, in our model the average cost of firms falls over time at a decreasing rate, as suggested by the empirical evidence on learning curves.[1]The rate of learning is determined by the specific challenges faced by firms and it is thus linked to the process of firm differentiation. One implication of this model is that industry learning curves are considerably smoother than plant learning curves, a feature which is also consistent with empirical evidence.[2]

The extent of learning is also affected by the degree to which the specific contingencies faced by a firm can be anticipated and the speed of implementation of new solutions. These factors affect the timing of investment

[1]This results from the stable task environment assumed, where the returns to learning decrease over time. This feature is shared by models of Bayesian learning in stable task environments (e.g. Jovanovic and Nyarko 1994, 1995), where the returns to experimentation also decrease with the precision of knowledge.

[2]See review articles by Conley (1970), Yelle (1979), Muth (1986) and Epple et al (1991). Auerswald, Kauffman, Lobo and Shell (1996) provide a general model with these properties and give a thorough account of the empirical literature.

decisions and also the heterogeneity of firm values. Thus, the model suggests a new set of characteristics of the learning environment that may contribute to explaining the large dispersion in learning rates across different industries, as evidenced by the data.

Third, our model has the feature that older firms, when confronted with new opportunities, tend to invest less in new specialized capital and resort more to their accumulated investments in order to profit from these opportunities. This is consistent with the March and Simon model, which predicts that formalized processes will, over time, replace innovation as the dominant firm activity. Older firms will invest less frequently in new assets (which can be knowledge assets, i.e. learning, or capital assets) than younger firms, on average, because they are already better adapted to their environment. This is the result of their longer history of accumulated investments in environment-specific assets.

Our model provides a mechanism for the observed heterogeneity of firm values, and the persistence of those values over time. Recent empirical studies indicate that most US industries are characterized by a large degree of firm heterogeneity which is mostly the result of firm level uncertainty (Davis et al 1996, Dunne et. al 1990, Evans 1987, Leonard 1987). According to these studies, establishment employment growth follows a mean reverting process with persistence increasing with the age of firms. If we hypothesize a correlation between firm value and employment level, our model provides a mechanism for explaining these observations. It is interesting to note that, in contrast to other models in the literature, this persistence is obtained even when the opportunities faced by firms are assumed to be independent across time.[3]

Our model also generates a mean reverting profitablity distribution in an industrial cohort. That is, the distribution of profitability in a cohort of firms begins trivially as each firm has an identical cost structure. But, due to specialized investments and the particular sample path of opportunities each firm faces, the variance of firm profits increases over time, to a point.

[3] The dominant economic model used to explain these facts has been Jovanovic's (1982) selection model. Firms' optimal employment level is affected by a firm-specific underlying productivity parameter, which is not known at entry. As the firm learns about this parameter over time, its estimate becomes more precise and its targeted employment level more persistent. The firm's evolving perceptions of its underlying parameter are autocorrelated via Bayesian updating. In contrast, persistence in our model comes from the interaction of specialized and irreversible investment.

Eventually, as firms reach a point of asset saturation, the profitability distribution across firms in an industry narrows once again. This variability does not disappear, however, as the history-dependent nature of each firm's investments exerts a persistent influence on its profitability and firm value.

We also investigate the consequences of time lags in implementing solutions to encountered problems. If there is no autocorrelation in the time series of problems faced by a firm, and if there is a time lag for a firm to respond to problem types, then firms do not (indeed, cannot) react to the current problem and must instead anticipate future problems. Since all firms face the same distribution of future problem types, firms will be more similar in their chosen investments and firm values will exhibit no heterogeneity. Hence, there is an interdependence between firm heterogeneity and the response time of firms in an industry.

Organizations researchers, starting with the surveys of Scottish and English firms by Woodward (1965) and Burns and Stalker (1961), associate "mechanistic" firms with stable task environments. A mechanistic firm is one that operates via formal rules and established procedures, and a stable task environment is one in which the distribution of problem types faced by the firm is stable over time. The learning model proposed by March and Simon (1958) suggests a dynamic by which firms evolve to mechanistic behaviors. Their model can be summarized as follows.

Consider a firm that is completely new to its business. This firm, like all firms in the industry, will face a series of problems that it must solve. Faced with a new problem, the firm will solve the problem by working through the following sequence of activities: evoking possible solutions, evaluating these alternatives, and implementing one of them based on the evaluations. Once a problem is solved in this fashion, and a solution implemented, the problem and its adopted solution become part of the memory of the firm. As time goes on, the firm will face more and more problems, and expand its list of problem/solution pairs in its repertory (that is, the firm becomes more experienced in its task).

It is less costly to implement an old solution than search for a new one. Hence, faced with a new problem, the firm will first peruse its repertory of problem/solution pairs to see if any are applicable to the new challenge. If a solution has already been found, then the firm can simply implement that solution program. March and Simon call this "reproductive" problem solving, because it relies on the reproduction of past programs of action rather than the innovation of new ones. If the task environment is stable, then over time the repertory of solutions maintained by the firm will expand to cover most of the problems the firm is likely to face, and reproductive problem solving is invoked more frequently than innovative problem solving. Over time, decision rules will evolve with the following character: "If a task of type A is encountered, then implement program C." Firm members simply follow these rules. This is an efficient way to operate in a stable task environment, because prior time-consuming searches for solution alternatives are leveraged to identify good responses to the problems of the day. Our model of firm differentiation derives from the above evolutionary dynamic, and elaborates on it in the context of a formal mathematical model.

The paper is organized as follows. Section 2 describes the model mathematically. Section 3 presents some attributes of the optimal solution. Section 4 presents some computational results and section 5 contains some concluding remarks.

2 The model

At discrete points in time a firm is confronted with a profit opportunity, which we represent by a point s on the unit circle S. The firm can respond by investing in some specific assets, which can be people, equipment, procedures, or knowhow. The space of possible investments is also assumed to be the unit circle. There may be a time lag $l \geq 0$ for a new investment to be implemented and operational. Investments do not depreciate with time, so that the capital stock of a firm at any time t will be the union of all investments initiated up through time $t - l$. This union can be represented by a set $k = \{k^1, k^2, ...\}$ where $k^j \in S$ for all j.

We assume that the profit opportunities follow a first order Markov process. At the beginning of time period t a firm will be cognizant of its current capital stock set k_{t-1} and will be presented with a new profit opportunity $s_t \in S$. The firm then has the following decisions to make: whether to invest in new assets or not, and what type, and if and how to apply its available assets to opportunity s_t. Let $i_t \subset S$ denote the new investments initiated in period t. It costs a fixed amount κ to install any single new asset or capability, so that the cost of investment i_t will be $\kappa n(i_t)$ where $n(i_t)$ is the cardinality of i_t. If $l > 0$ then the new investments i_t will not be available to attend to s_t. If $l = 0$ then new investments are immediately operational and can be used to address s_t.

Let $k_t = k_{t-1} \cup i_{t-l}$ be the set of available assets after the investment decision is taken. Suppose the firm decides to take the given profit oppor-tuninty s_t, and chooses a specific asset $k^j \in k_t$ to apply to it. Then the firm receives profits that depend on the distance between s_t and k^j given by $\phi(|s_t - k^j|)$ where ϕ is a bounded, continuous decreasing function. Naturally, given this framework, if the firm takes the profit opportunity, it will always choose to apply the asset in k_t closest to s_t. For any capital stock k_t and profit opportunity s_t this defines a profit function $\pi(s_t, k_t)$. The firm receives this profit, carries capital stock k_t into the next time period, receives another profit opportunity, and the process begins anew.

The decision problem of the firm has the following interpretation. The set k_t represents the accumulated capabilities (in the form of capital assets, procedues, experience and expertise, etc.) embodied in the firm. We will use the terms "capabilities" and "assets" interchangeably below. s_t is the specific profit opportunity encountered by the firm in time t. The firm can respond to opportunities by investing in new assets and/or specifying an available asset to apply to the current opportunity. We assume that a firm will behave optimally (defined more completely below). Hence, as noted above, if profitable, the firm will always activate the capability closest to the current opportunity. Since the investment behavior of the firm is affected by the opportunities or contingencies it confronts, the history of the firm will have a permanent effect on its final asset base. As time evolves, firms become differentiated according to the idiosyncratic opportunities they confront. In this way, the model provides a theory of firm differentiation. The manner in which firms evolve and become differentiated from each other will depend on the specific opportunity stream they encounter, as well as the industry parameters such as the implementation lag (l) and autocorrelation structure of $\{s_t\}$.

3 Optimal investment in assets and procedures

We now turn to the optimal investment problem of the firm. We assume that firms discount cash flows at the rate $\delta = \frac{1}{1+r}$ for some opportunity rate r, and behave so as to maximize their expected discounted profits over an infinite time horizon:

$$supE \sum_{t=1}^{\infty} \delta^{t-1} \left[\pi(s_t, k_t) - \kappa n(i_t) \right] \tag{1}$$

where $k_t = k_{t-1} \cup i_{t-l}$. All firms are assumed to begin with no capital stock, $k_0 = \emptyset$, and we define $\pi(s, \emptyset) = 0$ for any $s \in S$.

Since current investments will not become available for l periods, at any decision epoch the profits for the next $l-1$ time periods depend on historical actions and an exogenous stochastic process only, and are independent of current policy decisions. A reformulation of this problem, standard in the inventory literature with delivery lag times (cf Heyman and Sobel 1984) shows that an equivalent objective to (1) is the following:

$$supE \sum_{t=1}^{\infty} \delta^{t-1} \left[\delta^l \pi(s_{t+l}, k_{t+l-1} \cup i_t) - \kappa n(i_t) \right]. \tag{2}$$

For notational convenience we will redefine k to be the capital stock "position," that is, procedures already on hand or in the process of being implemented. In our case, what was referred to as k_{t+l} above will henceforth be referred to as k_t. This is analogous to the notion of an "inventory position" in the inventory literature, which refers to inventory on hand and on order. Note that the capital stock position is the stock that will be fully implemented and usable l periods hence with no additional investment activity. It is well known that in problems of this type, the stock position is sufficient to prescribe optimal investments, and is the natural way to define the state space for the problem (cf Heyman and Sobel). With this change of notation, the problem is to find an investment policy to maximize

$$E \sum_{t=1}^{\infty} \delta^{t-1} \left[\delta^l \pi(s_{t+l}, k_{t-1} \cup i_t) - \kappa n(i_t) \right] \tag{3}$$

where $k_0 = \emptyset$ and $k_t = k_{t-1} \cup i_t$.

Letting

$$R(s, k) = \delta^l E \left(\pi(s_{t+l}, k) | s_t = s \right),$$

the dynamic program for this investment problem is given by

$$V(s, k) = sup_{i \subset S} \left[R(s, k \cup i) - \kappa n(i) + \delta E \left(V(s_{t+1}, k \cup i) | s_t = s \right) \right]$$

To proceed with the analysis, we first introduce a symmetry condition on the Markov process s_t. For any arbitrary $x \in S$ and $a \in [0, 2\pi)$ let $x + a$ denote the element in S that is arc length a from x in the clockwise direction (since S is a circle of radius 1, $x + 2\pi = x$). For $y = x + a$ let $F(y, x)$ be the probability that $s_{t+1} \in [x, x + a)$ given $s_t = x$. We assume the following:

Assumption 1: For all x, y in S and $a \in [0, 2\pi)$, $F(x+a, x) = F(y+a, y)$.

Assumption 2: F is a continuous function and the process defined by F is recurrent.

Assumption 1 implies that, in a scale relative to the current point s, the transition probabilities are independent of s. It follows from Assumption 1 that the firm's value will be unaffected by a simple rotation on S of the current (s, k) configuration. Assumption 2 implies that the process has a unique stationary distribution, which by Assumption 1 is the uniform distribution. In the appendix we establish that, subject to these assumptions, the optimal value of (2) will satisfy the above Bellman equation and the dynamic program based on this equation will generate that optimal value.

$V(s, k)$ is the maximum possible expected discounted profit available to a firm from a starting position (s, k). This formulation naturally yields the result that shorter implementation lags are always better, since any investment policy available to a firm with lag l can be reproduced by a firm with lag $l' \leq l$ simply be delaying the investments $l - l'$ time periods, thereby achieving the same capital stock trajectory but at discounted cost. So we have the following intuitive result:

Remark 1: The value of a firm is nonincreasing in the implementation lag l.

Intuitively, this models the increased value of being able to respond flexibly to market opportunities. Also, although procedures incur fixed costs to implement, once in place a firm cannot be worse off for having them there.

Remark 2: If $k' \subset k$ then for any $s \in S, V(s, k) \geq V(s, k')$.

In particular, it follows from Remark 2 that a firm's value will be nondecreasing over time, since capital is added but never removed. It also follows that older firms in an industry, with more environment specific assets already in place, have a profitability advantage over younger firms.

A firm with capital stock k can be identified with a partition $\mathcal{L}$ of S by $\mathcal{L}(k) = \{[k^1, k^2), [k^2, k^3), ...\}$. The optimal sequential investment problem is equivalent to that of determining an optimal sequential partition of S.

Each set in the partition is an interval determined by two endpoints which correspond to neighboring capital types that the firm has invested in. We will use ℓ to denote an interval in some partition $\mathcal{L}$ of S. The optimal investment problem is one of optimally refining this partition.

This dynamic programming problem appears to be a complex one, being that the state space is isomorphic to the set of all possible partitions of the unit circle. However, as we now show, this problem can be decomposed into simple subproblems.

Because the probability that a profit opportunity falls within any specific interval is independent of the firm's investment policy, and a firm will always use the asset closest to the current opportunity, the value of a claim to operate an interval, $[k_1, k_2)$ say, will be independent of investments outside that interval (that is, further away than the endpoints k^1 or k^2). One may conjecture, then, that intervals in any partition of S can be valued independently, and that the total value of the firm can be attained by summing up such claims. This decomposition is established in Proposition 1 below.

Consider an interval ℓ. Define $\pi_\ell(s,k)$ to be $\pi(s,k)$ if $s \in \ell$ and zero otherwise. That is, π_ℓ is a restriction of the function π to the interval ℓ. Let k_ℓ denote the set of (two) endpoints ℓ. Finally, let

$$v(s,\ell) = sup_{i \subset \ell} \left\{ \delta^l E_{s_{t+l}|s_t=s} \pi_\ell(s_{t+l}, k_\ell \cup i) - \kappa n(i) + \delta \sum_{\ell' \in \mathcal{L}(k_\ell \cup i)} E_{s_{t+1}|s_t=s} v(s_{t+1}, \ell') \right\}.$$

This gives the value of a claim to operate the subset of profit opportunities in the interval ℓ.

Proposition 1 $V(s,k) = \sum_{\ell \in \mathcal{L}(k)} v(s,\ell)$.

Proof. See appendix.

Assumption 1 will imply that starting with $k_0 = \emptyset$ firms are indifferent to the location on S of the first opportunity s_1. In particular, if a firm chooses not to invest during the first time period, it will never do so, no other firm will do so, and all firms will have value zero (essentially, nobody chooses to participate in the industry). So for the remainder of this paper, assume that firms find it profitable to invest in period 1. Since the first problem s_1 will differ among firms only in a rotation on S, all firms will make investments i_1 that differ only by such a rotation. However, starting with the profit opportunity in period 2 firm values will differ based on the idiosyncratic position of the opportunities they face relative to their installed asset base. Essentially, the first investments on S remove one's indifference to rotations of s on S, since such rotations can move the opportunity s closer to or farther from an installed procedure in k. So, firms may differentiate themselves after period 1.

After the first period, not all opportunities will lead to investments. For example, if an opportunity arises very close to an existing asset, that asset can be used instead of incurring the cost of installing a new procedure. Investments will continue to occur at times, provided the length of the largest interval in the partition exceeds a critical number $\epsilon > 0$. For example, starting with any partition $\mathcal{L}$, an upper bound on the value gain for investment anywhere would be $\frac{1}{1-\delta}(\phi(0) - \phi(|\ell|_{max})$ where $|\ell|_{max}$ denotes the interval of maximal length in the partition $\mathcal{L}$. Consequently, if $\frac{1}{1-\delta}(\phi(0) - \phi(|\ell|_{max}) < \kappa$, no further investment will take place. In this case $\epsilon \geq \phi^{-1}(\phi(0) - \kappa(1-\delta))$. If all intervals have length below the critical value, no further investment occurs and the final configuration of the firm is attained. If an interval in the partition has a length exceeding the critical value, then there will exist a region of positive (probabilty) measure within the interval such that an opportunity falling within this region will result in an asset being acquired, and the interval subdivided. It is apparent that all intervals will eventually, and with probability one, be subdivided into subintervals of length less than or equal to ϵ. That is, the firm or organization will, inevitably, cease innovative learning and rely solely on installed assets.

We now further characterize an optimal investment policy in some specific cases of this general framework.

A. No implementation lags for new investments

With no lags in implementing new procedures ($l = 0$), we will never invest in advance of usage. This is because with discounted costs any investment not used immediately can be installed in the next time period at a discounted cost, and with no disbenefit for the delay. So, it is always better to delay any investment that will not be used immediately.

Remark 3. No investment will be made unless the newly installed procedure (asset) is used to address the current opportunity.

Since only the closest asset to an opportunity is used, the next result follows immediately.

Remark 4. At most one asset will be acquired each time period.

Now we can track the evolution of the firm as the sequential partitioning of intervals. Starting with no assets, suppose the first period results in an investment at k^1. Suppose that an opportunity is eventually received that results in a new investment, and assume that the optimal policy prescribes an investment at the point k^2. After this investment takes place, the state of the firm is isomorphic to the partition $\mathcal{L}(k) = \{[k^1, k^2), [k^2, k^1)\}$ of S. The problem can now be decomposed into the separate operation of these two segments. Suppose a new problem eventually arises at $s \in [k^1, k^2)$ that leads to a new investment at point k^3. Then by Remark 3, we must have $|s - k^3| < \min(|s - k^1|, |s - k^2|)$. In particular, this implies that $k^3 \in [k^1, k^2)$, resulting in a subdivision of this interval. The new partition for the firm is then $\mathcal{L}(k) = \{[k^1, k^2), [k^2, k^3), [k^3, k^1)\}$. We can summarize this in the following remark.

Remark 5. If an opportunity leads to investment, it will subdivide the interval in the partition currently containing the opportunity.

When all investments have ceased, a firm will resort to its repertory of installed procedures to solve all new problems. Formalized processes have replaced innovation as the dominant firm activity and the firm has evolved from an innovating organization into a mechanistic bureaucracy. This is efficient, indeed optimal, in the sort of stable task environment that we have assumed here. The correlation between mechanistic behaviors and firm age is driven by the fact that older firms have a larger set of sunk investments that they can draw from, and consequently derive lower returns from new investments. As we have noted, the value of a firm is monotonically increasing over time, so that older firms will be on average the most profitable ones.

The model implies a stochastic process for the value of a firm which converges to an invariant distribution. Letting k_∞ denote the limiting capital configuration of the firm, we can define the long run value of this firm by

$$\overline{V}(k_\infty) = \frac{\delta^l E\pi(s, k_\infty)}{1 - \delta}$$

where the expectation is taken with respect to a uniform distribution since by assumption this is the limiting distribution for s.

B. Positive implementation lags $(l > 0)$.

The presence of implementation lags can change the nature of the investment policy considerably. Consider again the situation where s_t is an i.i.d. process, that is, F is uniform with support S. Since investment in period t cannot be used for the problem s_t, investment decisions will not be affected by any particular realization. In fact, investment can only be justified as setting the firm up gracefully for the anticipated uniform distribution of opportunities, the current problem being irrelevant to this decision. This situation is unaffected by the passage of time, with the following result.

Proposition 2 *If $l > 0$ and s_t is i.i.d., then each firm will make all of its investments in the first period, and the optimal policy is to partition S symmetrically into n pieces, where n depends on the parameters of the problem. Consequently, the value of a firm will be constant (after this initial investment) and all firms will have identical values.*

Proof. See appendix.

When s_t is not i.i.d., but implementation lags are greater than zero, one may conjecture the following mixed policy will result. Investments will take place in chunks to cover a neighborhood of the highest probability regions for s_{t+l} given s_t. The number and distribution of these investments will be affected by the parameters of the problem. As an example, consider the following scenario. Suppose the profit opportunities fall in either of two sets, S_1 and S_2, so that $S = S_1 \cup S_2$. The transition probability between the two sets is $0 < \varepsilon < 1$ and the conditional distribution for the states in each set is uniform. If the initial state $s_1 \in S_1$, and ε is close to zero, the first period investment i_1 will consist of a set of points in S_1, as discussed above. There will be no more investments, until an opportunity arrives in S_2. At that time, a new set of investments will be made, this time in S_2. From there on, there will be no more investetmtents. In the more general case, the interaction between the statistics of the opportunity process and implementation lags in determining firm value is likely to be complex, and we have no results on this. The exception is the case described above and also the extreme case where the first opportunity s_1 recurs indefinitely. In the latter case, the only investment a firm will make will be at that point, and firms may have different assets but all will have identical values of $\frac{\delta^l}{1-\delta}\phi(0)$.

Tractable computations using this model are made possible by the decomposition results above. The following section presents some numerical results that illustrate the dynamic process of firm learning and differentiation.

4 Computational results

We consider a version of the above model with $\phi(|s_t - k^j|) = A - \alpha\, |s_t - k^j|^2$, where α and A are constants independent of s_t and k^j. There are no implementation lags. The interpretation of this function is as follows. The firm has activities which give a constant payoff A but faces problems (s_t) which must be solved using some of its accumulated procedures (k^j) or investing and using a new one. In one way or another, all problems must be solved. We assume that s_t is a sequence of *i.i.d.* random variables which are uniformly distributed in S. The value maximization problem is obviously equivalent to the minimization of the expected discounted cost of solving these problems, so the program we consider is given by

$$C(s,k) = min\{\min_{k^j \in k} \alpha(s-k^j)^2 + \delta \int C(s',k)\, ds,$$
$$\kappa + \min_{i \in S} \alpha(s-i)^2 + \delta \int C(s', k \cup i)\ \}$$

where integration is with respect to the uniform distribution on S. The unconditional expected cost $\bar{C}(k) = \int C(s,k)\, ds$ and the average expected cost per period is $(1-\delta)\bar{C}(k)$. Note that

$$\bar{C}(k) = \int \min\{\min_{k^j \in k} \alpha(s-k^j)^2 + \delta \bar{C}(k),$$
$$\kappa + \min_{i \in S} \alpha(s-i)^2 + \delta \bar{C}(k \cup i)\}\, ds$$

By an argument similar to Proposition 1 it can also be established that

$$\bar{C}(k) = \sum_{\ell \in \mathcal{L}(k)} \bar{c}(\ell)$$

where $\bar{c}(\ell)$ is the expected cost of operating the interval ℓ.

To solve this problem, we discretized the state space by subdividing the unit circle into 400 elementary points. We assumed that the distribution for problems was uniform on these points, with probabilities $\frac{1}{400}$ of occurring. The value function and the optimal policy rule were calculated using the following iterative scheme.

An interval consisting of n points is interpreted as one where zero cost solutions exist for the endpoints only. Consequently, the firm starts with a single interval consisting of 400 points. On the other extreme, a subinterval consisting of only two points is one where all possible problems that may arrive have zero cost solutions, having thus a zero cost of operation. The value of operating an interval with n points is given by

$$\begin{aligned}\bar{c}^n = \quad & \frac{1}{400}\sum_{i=1}^{n} \min\{\alpha \min\left[\left(\frac{n-i}{400}\right)^2, \left(\frac{i-1}{400}\right)^2\right] + \delta\bar{c}^n, \\ & \qquad \kappa + \min_{1<m<n} \alpha\left(\frac{m-i}{400}\right)^2 + \delta\,\left(\bar{c}^m + \bar{c}^{n-m}\right)\} \\ & + \left(1 - \frac{n}{400}\right)\delta\,\bar{c}^n,\end{aligned}$$

where $\bar{c}^2 = 0$.

Several features of the above formula deserve clarification. The value is multiplied by $\frac{1}{400}$ because this is the probability that a problem arises in each point of the interval. When confronted with a problem $i = 1, ..., n$ the firm must first decide whether it will subdivide the interval. The case of no subdivision is given first, with a discounted continuation cost of $\delta\bar{c}^n$. The case of investment is given by the last term in brackets. m represents the point chosen for the new investment, which without loss of generality is the one used to solve the current problem. After this subdivision takes place the discounted continuation cost is the sum of the cost of operating the subintervals of size m and $n - m$. Finally, if no problem occurs in this interval of size n the continuation value is $\delta\bar{c}^n$. The probability of this event is $\left(1 - \frac{n}{400}\right)$.

Starting from $\bar{c}^2 = 0$, the $\bar{c}^n$ are successively calculated. The optimal policies for each case are also obtained as indicated above. Table 1 provides some results on the characteristics of the process under the optimal policy. The results given were obtained from 200 independent sample paths for $\{s_t\}$ of 100 periods each. As a first approximation, abstracting from competitive interactions, it is convenient to think of this as representing the evolution of a cohort of 200 firms that enter the industry at a given date. The statistics computed refer to the expected cost per period, i.e. $(1 - \delta)\bar{C}(k)$, described above.

According to column 2, the average expected cost decreases at a rate which is itself decreasing. This is consistent with standard observations on learning curves. Statistics of the variability of these average costs are given in columns 3 to 6. The industry starts with no variability (e.g. all firms are identical). After a few periods, some significant variability (and thus differentiation) develops. Both the standard deviation and the coefficient of variation of expected costs per period are monotonically increasing up to a point and then decrease. This is consistent with observations on the evolution of the performance of firms in new industries as measured by size, where concentration increases at early stages and then decreases. There is a substantial degree of asymptotic variability as measured by the coefficient of variation (16.8%) and also reflected in the fact that the least efficient firm has more than twice the average cost of the most efficient one.

Column 7 gives the autocorrelation between expected costs in periods t and $t-1$, which provides a measure of persistence of interfirm differences. Persistence increases with age, as has been observed for the case of differences in employment in the data. Consequently, the mobility of firms' level of expected costs decreases monotonically through time. This results from the interaction between irreversible and specialized investments.

5 Final Remarks

This paper developed a model of firm differentiation which captures the idea that contingencies faced by firms can have persistent effects on their performance. The key element in our analysis is the combination of specialized and irreversible investments. The model is broadly consistent with a variety of observations on firm evolution.

Firms in our model are portfolios of known procedures, each acquired at a cost to the firm and intended for use to solve problems posed by the task environment. In this way our model formalizes the notion of firms as repositories of knowledge (cf Teece 1982). In our model older firms have accumulated more specific capital and thus face lower average returns to new investments. Consequently, these firms turn more often to established expertise in order to confront new opportunities, rather than investing in new specific knowledge.

We now consider several possible extensions. Our paper follows a decision theoretic approach, modeling the optimal investment behavior of a firm in isolation. The particular decision problem confronted by the firm could be incorporated into an equilibrium model. Then, the profits of an individual firm would also depend on the particular configuration of investments made by its counterparts. It would then be possible to embed such an environment in an entry/exit model.

In our model we consider the case where shocks faced by firms are independently distributed over time. As we establish, in the presence of implementation lags firms do not respond to shocks and make identical upfront investments. An interesting extension is to consider the case of correlated shocks with implementation lags. We have conjectured that in this case investment may occur in “chunks": in the presence of a new distinct set of opportunities, firms may decide to invest in anticipation of related challenges that may subsequently emerge. It is worth emphasizing that recent empirical studies of plant investment also suggest that individual firms' investments tend to be concentrated over short time periods (Cooper et al 1995, Doms and Dunne 1994)

This paper assumes a stable environment, and this is why older firms may, on average, have an advantage over their younger counterparts. An interesting extension is to consider the situation that arises when environments change. For instance, we may relate an environment to a certain distribution for the profit opportunities faced by firms. A change in environment implies a change in this probability distribution, thus depreciating investments made by older firms.

A Justifying the Bellman equation

The dynamic programming problem is defined by

$$V(s,k) = \sup_{i \in S_f} \{R(s, k \cup i) - \kappa n(i) + \delta \int V(s', k \cup i) F(ds'|s)\}$$

where S_f is the set of finite subsets of S and $k \in S_f$. Note that $S_f = \cup_n S^n$ where S^n is the set of subsets of S with cardinality less or equal to n. Let $n(\emptyset) = 0$.

The return function for this DP problem $H : S \times S_f \times S_f \to \Re$ is defined by

$$H(s,k,i) = R(s, k \cup i) - \kappa n(i).$$

Using Assumption 2 it is straightforward to show that $R(s,k)$ is continuous in s, k. Endowing S_f with the Hausdorff metric, it can be shown that $n(i)$ is lower semi-continuous on S_f so that H is continuous in s and k and upper semi-continuous in i.

It can also be shown that each S^n is compact with the Hausdorff metric, so that S_f is the union of compact sets. It then follows then from Proposition 9.8 and Corollary 9.17.3 of Bertsekas and Shreve (1978) that the dynamic programming algorithm will generate a sequence of value functions that will converge to the optimal value function, which will be upper semi-continuous; the optimal value function will be a fixed point of the Bellman equation; and we can restrict attention to nonrandomized, stationary Markov policies.

B Proof of Proposition 1: additivity of DP program.

We establish the additivity of the value function by induction. Let V_t denote the tth iterate value function in the dynamic programming recursion, which we know converges to the optimal value function. We show that the additive property is preserved by the recursion, and hence in the limit the optimal value function has the additive form. Let $\mathcal{L}(k)$ denote the set of intervals forming the partition of S induced by assets k, so that $S = \cup_{\ell \in \mathcal{L}(k)} \ell$.

Assume inductively that the value function $V_{t+1} : S \times S_f \to \Re$ has the form

$$V_{t+1}(s,k) = \sum_{\ell \in \mathcal{L}(k)} \nu_{t+1}(s,\ell)$$

where

$$\nu_t(s,\ell) = \sup_{i \subset \ell} R(s, k_\ell \cup i) - \kappa n(i) + \delta \sum_{\ell' \in \mathcal{L}(k_\ell \cup i)} E_{s_{t+1}|s} \nu_{t+1}(s_{t+1}, \ell')$$

and i denotes the investments in interval ℓ. Consider now the functional equation defining V, that is

$$V_t(s,k) = \sup_{i \subset S} \left[R(s, k \cup i) - \kappa n(i) + \delta E_{s_{t+1}|s} V_{t+1}\left(s_{t+1}, k \cup i\right) \right]$$

By the definition of R and π_ℓ, the induction hypothesis and noting that $k = \sum_{\ell \in \mathcal{L}(k)} k_\ell$, it follows that

$$\begin{aligned} V_t(s,k) &= \sup_{i \subset S} \sum_{\ell \in \mathcal{L}(k)} \delta^l E \pi_\ell(s_{t+l}, k_\ell \cup i) - \kappa \sum_{\ell \in \mathcal{L}(k)} n(i) \\ &\quad + \delta E \sum_{\ell' \in \mathcal{L}(k \cup i)} \nu_{t+1}(s_{t+1}, \ell')] \\ &= sup_{\{i_\ell, \ell \in \mathcal{L}(k)\}} \sum_{\ell \in \mathcal{L}(k)} \Big[\delta^l E \pi_\ell(s_{t+l}, k_\ell \cup i_\ell) - \kappa n(i_\ell) \\ &\quad + \delta \sum_{\ell' \in \mathcal{L}(k_\ell \cup i \ell)} E \nu_{t+1}(s_{t+1}, \ell') \Bigg] \end{aligned}$$

where all expectations are conditional on $s_t = s$. Because each investment set i_ℓ can be chosen independently for each $\ell \in \mathcal{L}(k)$, this equals

$$\begin{aligned} &\sum_{\ell \in \mathcal{L}(k)} sup_{i_\ell \in \ell} \Big\{ \delta^\ell E \pi_\ell(s_{t+l}, k_\ell \cup i_\ell) - \kappa n(i_\ell) \\ &\quad + \delta \sum_{\ell' \in \mathcal{L}(k_\ell \cup i \ell)} E \nu_{t+1}(s_{t+1}, \ell') \Bigg\} \\ &= \sum_{\ell \in \mathcal{L}(k)} \nu_t(s, k_\ell) \end{aligned}$$

completing the induction.

C Proof of Proposition 2: Lags with no persistence

We now establish that if the lag $l > 0$ and the problem instances are *i.i.d* (and hence uniformly distributed on S by our assumptions) then all firms will invest in identical assets in period 1 and will refrain from any further investments. In this case, firm values will all be identical. Note that if there is no autocorrelation among problem instances then the distribution of the problem to be faced l periods from now is independent of the problem currently faced. Hence, $R(s,k) = \delta^l E\left(\pi(s_{t+l},k)|s_t = s\right)$ is independent of s. Likewise, $\delta \int V(s', i' \cup i)F(ds'|s) = \delta \int V(s', i' \cup i)F(ds')$ is independent of s. Consider now the very first time period for any firm. The firm will invest in the battery i^* of assets that maximizes

$$R(s, i^*) - \kappa n(i^*) + \delta \int V(s', i^*)F(ds'|s)$$

$$= R(i^*) - \kappa n(i^*) + \delta \int V(s', i^*)F(ds')$$

because all costs are independent of s. Now, suppose at some future period the firm finds it beneficial to invest in additional assets i. Then at that time the firm must believe that

$$R(i \cup i^*) - \kappa n(i) + \delta \int V(s', i^* \cup i)F(ds')$$

$$> R(i^*) + \delta \int V(s', i)F(ds')$$

but this implies that

$$R(i \cup i^*) - \kappa n(i) - \kappa n(i^*) + \delta \int V(s', i^* \cup i)F(ds')$$

$$> R(i^*) - \kappa n(i^*) + \delta \int V(s', i)F(ds')$$

which contradicts the optimality of i^* in period 1. That is, the firm will invest in assets in period 1 and will not invest at any future time. Since all firms face the same situation in period 1, all firms will invest identically (up to a translation) and all firms will have the same value.

References

[1] Anthony, R., J. Dearden and V. Govindarajan. *Management Control Systems.* Irwin, Homewood, IL, 1992.

[2] Auerswald, Phil, Stuart Kauffmand, José Lobo and Karl Shell, "A Microeconomic Theory of Learning-by-Doing: An Application of the Nascent Technology Approach," mimeo, Cornell University, 1996.

[3] Barnett, W. "The Dynamics of Competitive Intensity." Research paper 1227, Stanford University Graduate School of Business, Stanford, CA. 1992.

[4] Bertsekas, D. and S. Shreve. *Stochastic Optimal Control: The Discrete Time Case.* Academic Press, N.Y. 1978.

[5] Burns, T. and G. Stalker. *The Management of Innovation.* Tavistock Publications, London, 1961.

[6] Cooper, R., J. Haltiwanger and L. Power. "Machine Replacement and the Business Cycle: Lumps and Bumps," mimeo July 1995.

[7] Conley, P. "Experience Curves as a Planning Tool," *IEEE Spectrum* 7, 1970, 63-68.

[8] Davis, S., J. Haltiwanger and S. Schuh. *Job Creation and Destruction,* MIT Press, Cambridge, MA 1996. Doms, M. an T. Dunne, "Capital Adjustment Patterns in Manufacturing Plants," mimeo, Center of Economic Studies, U.S. Census Bureau, 1994.

[9] Dunne, T., M. Roberts and L. Samuelson. "The Growth and Failure of US Manufacturing Plants," *Quarterly Journal of Economics* 104, 1990, 671-698.

[10] Epple, D., L. Argote and R. Devadas. "Organizational Learning Curves: A Method for Investigating Intra-plant Transfer of Knowledge Acquired Through Learning by Doing," *Organizational Science* 2, 1991, 58-70.

[11] Ericson, R., and A. Pakes. "An Alternative Theory of Firm and Industry Dynamics." Working paper, Department of Economics, Columbia University, N.Y., N.Y., 1989.

[12] Evans, D.S. The Relationship Between Firm Growth, Size and Age: Estimates for 100 Manufacturing Industries, *Journal of Industrial Economics*, 1987 vol. 35, 567-581.

[13] Hannan, M. and J. Freeman. "The Population Ecology of Organizations." *American Journal of Sociology* 82, 1977, 929-964.

[14] Hayes, R. and S. Wheelwright, *Restoring our Competitive Edge*, John Wiley and Sons, New York, 1984.

[15] Heyman, D. and M. Sobel. *Stochastic Models in Operations Research: volume II.* McGraw Hill, N.Y. 1984.

[16] Jovanovic, B. Selection and the Evolution of Industry, Econometrica, 1982, vol. 50, 649-670.

[17] Jovanovic, B. and Y. Nyarko. "The Bayesian Foundations of Learning by Doing," working paper 4739, NBER, 1994.

[18] Jovanovic, B. and Y. Nyarko. "A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves," Brookings Papers on Economic Activity: Micro issue no. 1, Brookings Institution 1995.

[19] Leonard, J. "In the Wrong Place at the Wrong Time: The Extent of Frictional and Structural Unemployment," in *Unemployment and the Structure of Labor Markets*, edited by K. Lang and J. Leonard. New York, Basil Blackwell, 1987.

[20] March, J. and H. Simon. *Organizations.* John Wiley and Sons, N.Y. 1958.

[21] Muth, J. "Search Theory and the Manufacturing Progress Function," *Management Science* 32, 1986, 948-962.

[22] Nelson, R. and S. Winter. *An Evolutionary Theory of Economic Change.* Harvard University Press, Cambridge, MA, 1982.

[23] Stalk, G. and T. Hout. *Competing Against Time: How Time-based Competition is Reshaping Global Markets.* The Free Press, N.Y. 1990.

[24] Teece, D. "Towards an Economic Theory of the Multiproduct Firm." *Journal of Economic Behavior and Organization* 3, 1982, 39-63.

[25] Utterback, J. *Mastering the Dynamics of Innovation.* Harvard Business School Press, Boston, 1994.

Woodward, J. *Industrial Organization: Theory and Practice.* Oxford University Press, London, 1965. Yelle, L.E. "The Learning Curve: Historical Review and Comprehensive Survey," *Decision Sciences* 10, 1979, 302-328.

Table 1. Evolution of Firms' Expected Cost

Period	Mean	Std. Dev.	coef. var.	min	max	Corr. (t,t-1)
1	215.6	0.0	0.0	215.6	215.6	.....
2	199.3	7.0	3.5	195.4	215.6	0.000
3	185.0	10.0	5.4	175.5	215.6	0.560
4	171.6	12.2	7.1	155.7	198.3	0.700
5	160.1	14.1	8.8	136.3	196.2	0.753
10	120.0	18.3	15.3	65.0	159.8	0.877
20	82.4	17.2	20.9	43.2	124.8	0.925
30	69.2	14.9	21.5	39.1	120.1	0.967
40	61.8	12.5	20.1	33.0	106.1	0.983
50	57.9	11.1	19.2	33.0	99.3	0.985
100	51.7	8.7	16.8	32.0	81.4	1.000

Note: Parameter values: $\alpha = 1600$, $\kappa = 40$ and $\delta = 0.90$.