A Bayesian Approach to Managing Learning-curve Uncertainty

Joseph B. Mazzola • Kevin F. McCardle
Fuqua School of Business, Duke University, Durham, North Carolina 27708-0120

This paper introduces a Bayesian decision theoretic model of optimal production in the presence of learning-curve uncertainty. The well-known learning-curve model is extended to allow for random variation in the learning process with uncertainty regarding some parameter of the variation. A production run generates excess value (above its current revenue) for a Bayesian manager in two ways: it pushes the firm further along the learning curve, increasing the likelihood of lower costs for future runs; and it provides information, through the observed costs, that reduces the uncertainty regarding the rate at which costs are decreasing. We provide conditions under which one of the classical deterministic learning-curve results—namely, that optimal production exceeds the myopic level—carries over to the extended framework. We demonstrate that another classical deterministic learning-curve result—namely, that optimal production increases with cumulative production—does not hold in the Bayesian setting.
(*Learning Curve; Stochastic Learning; Bayesian Analysis; Dynamic Programming Model*)

1. Introduction

The value of process innovations, such as material requirements planning (MRP), just-in-time manufacturing (JIT), total quality management (TQM), and business process reengineering (BPR), is captured only in part by the immediate reduction in unit production costs brought about by the innovation's adoption. Indeed, the immediate effect of the adoption may be an increase in costs. The bulk of the value comes through production learning—the rate at which production costs decline with cumulative output. Historically, the existence of production-cost decline phenomena (also known as learning-curve or experience-curve phenomena) has been established for a wide range of industries (Wright 1936, Yelle 1979). The typical learning-curve model posits that costs decrease deterministically at a constant exponential rate for every doubling of cumulative production. It is well-known empirically, however, that the rate of the observed cost declines tends to vary widely, both across and within industries, and even across departments within a firm.

Suppose that a Bayesian manager has embarked upon the production of a new product line or has adopted a process innovation and knows that the associated costs will follow a learning curve. How should she plan production to optimize the learning-curve effect, i.e., to optimize the value of the innovation? If costs decrease deterministically and she knows the rate at which costs will decline, the answer is well established: produce in excess of the short run maximizing level (Hax and Majluf 1982). If the cost decrease is random, but the distribution of the variation is known, the answer is similar (Mazzola and McCardle 1995). This paper answers the manager's question when the cost decrease is random and she has some uncertainty regarding the distribution of the variation, i.e., we characterize optimal monopolistic production planning when costs follow a learning curve but some parameters of the learning-curve cost function are random variables. In particular, we develop a discrete-time stochastic dynamic programming model of a learning-curve production problem. We consider general cost functions, but focus attention on the case in which costs are (on average) exponentially de-

0025-1909/96/4205/0680$01.25

creasing in cumulative production with some random variation in the observed costs. The firm has some uncertainty about the parameters of the learning-curve cost function, represented by a prior probability distribution, and its uncertainty is partially resolved through time as costs are observed and the prior beliefs are updated in a Bayesian manner.

The literature based on a deterministic model of the learning curve is extensive. Several general results arise from this varied literature. The core results are that in the presence of learning-curve cost dynamics the optimal production path is increasing in cumulative production, and in each period optimal production is above the myopic level specified by short-run marginal costs (see, e.g., Hax and Majluf 1982). Applications of the learning curve arise in strategy where its use was championed by the Boston Consulting Group (Henderson and Zakon 1980), marketing (Robinson and Lakhani 1975, Dolan and Jeuland 1981, Kalish 1983), operations (Ebert 1976, Hiller and Shapiro 1986, Karwan et al. 1988, Fine and Porteus 1989, Dada and Srikanth 1990, Gaimon 1992), and economics (Clarke et al. 1982).

While empirical studies that rely on regression analysis implicitly assume a stochastic, though not Bayesian, learning model, e.g., Lieberman (1984), Adler and Clark (1991), the analysis tends to be descriptive, not prescriptive. There are a few papers that formally incorporate uncertainty into the learning-curve production problem. For example, Majd and Pindyck (1989) develop a learning model in which costs decline deterministically with cumulative production but prices vary stochastically. They show that optimal production follows a "bang-bang" policy: produce at capacity when price exceeds a reservation level, and shut down when price falls below that level. Spence (1981) and Fudenberg and Tirole (1983) develop competitive models in which the uncertainty is strategic. Optimal production in these models need not increase monotonically. Mazzola and McCardle (1995) introduce a stochastic model of the learning curve that allows random variation in production costs that follow a learning curve. Unlike the setting considered here, the model in their paper is static in the sense that the distribution of the random variable affecting costs is known and does not change over time.

Bayesian models of optimal firm behavior in the face of uncertainty abound. This literature includes application of Bayesian methodology to two of the areas addressed in the current paper: production planning and innovation adoption. For example, Harpaz et al. (1982) and Thompson and Horowitz (1993) develop a model of optimal multiperiod production when demand is random and some parameter of the demand distribution is unknown. Harpaz et al. show that production by a competitive Bayesian firm exceeds that of a nonexperimenting firm, while Thompson and Horowitz include a comparison with cooperative experimenting firms and stress that the benefit to experimentation is small. Lippman and McCardle (1991) analyze a model of technology choice when there are a set of available technological innovations, and the values of the technologies are unknown but can be discovered through a process of experimentation. Optimal experimentation and innovation adoption policies are characterized.

In §2 we introduce an infinite-horizon Bayesian framework for optimal production in the presence of learning-curve uncertainty. We distinguish two types of learning: the first is production learning epitomized by the deterministic learning curve; the second is Bayesian learning regarding the parameters of the cost function. To help distinguish between the two, henceforth we will use the term *experience curve* to refer to production learning, and *learning* to refer to Bayesian updating.

If uncertainty exists regarding some parameter of the distribution of the random variable defining costs, then observed costs will carry information about that parameter. Current-period production pushes the firm further along the experience curve and offers an opportunity to learn more about the experience curve. When embarking on the production of a new product, it seems reasonable that there will be considerable uncertainty regarding the rate at which costs will decline with cumulative production. Empirical evidence suggests that the decline rates vary widely across industries, across firms within an industry, and even across functions within a firm. In addition to the uncertainty regarding the parameters of the experience curve, we assume that there is a natural variation in costs around that curve. This variation in costs arises in a number of ways. For example, short-term changes in process technology such as the introduction of new equipment, revised job design, changes in process layout, as well as engineering design changes in subprocesses or in component

parts, among other factors, each contributes to this variation. In addition, random variation arises from the imprecision of the accounting procedures used to measure costs.

To model this uncertainty and variation, the state space must include not only cumulative production to date but also the distribution over the unknown parameters of the cost function. The first dimension of the state space (cumulative production) increases deterministically according to each period's production, while the second dimension (uncertainty regarding the parameters of the experience curve) is updated via Bayes Theorem. We specify sufficient conditions under which uniqueness of the value function and existence of an optimal policy hold for the Bayesian learning model. Two variants of an exponentially decreasing cost function are considered in further detail. In both cases, we provide conditions under which optimal production exceeds the myopic level, thus extending one of the key results from the deterministic experience-curve literature.

To further our understanding of optimal production behavior in the Bayesian setting, we shift attention to the finite-horizon case in §3. In addition to providing valuable insights into optimal behavior, our analysis demonstrates that some well-known properties of deterministic experience-curve models do not apply in the Bayesian environment. Explicit solution of a two-period model provides answers to questions such as: does optimal production increase with cumulative production? does increased variance decrease expected value? does optimal production exceed the level set by a non-Bayesian firm, i.e., a firm which does not update the distribution for the random variation? Furthermore, a second example shows that the difference in returns between the Bayesian and myopic firm can be substantial.

2. The Infinite-horizon Model

In this section we develop a discrete-time, infinite-horizon, stochastic dynamic programming model of the experience curve. The model allows not only random variation in the amount by which costs decrease each period but also uncertainty regarding the distribution of that variation. Production yields two types of learning: the production learning typically modelled in the experience-curve literature and Bayesian learning regarding the parameters of the model. By observing costs over time, the firm learns about the way those costs are changing—the firm learns how quickly it gains experience.

The firm has a constant marginal cost function $c(Q_{j-1}, Z_j)$ in period j, where Q_{j-1} is cumulative production at the beginning of period j and Z_j is a random variable with observed value z_j. We assume that $c(Q, Z)$ is decreasing in cumulative production Q and that the random variable Z is independent of Q. We first establish an existence result for general cost functions before analyzing two specific stochastic versions of the standard exponentially decreasing cost function. The firm's decision variable q_j in period j represents current period production. Let $\Gamma = [0, \bar{q}]$ be the feasible decision space, where $\bar{q}$ is an upper bound on production capacity; hence, $q_j \in \Gamma$ for all j. We assume that the random variable Z_j is independent of q_j. The inverse demand curve $p(q_j)$ is bounded and slopes downward. The period j profit function $\pi(Q_{j-1}, z_j, q_j)$ is given by

$$\pi(Q_{j-1}, z_j, q_j) = [p(q_j) - c(Q_{j-1}, z_j)]q_j. \tag{2.1}$$

The per-period discount factor is denoted by δ, and we assume that $0 < \delta < 1$.

The firm's problem is to choose a production policy $(q_1(Q_0, z_1), q_2(Q_1, z_2), \cdots)$ to maximize expected discounted profits:

$$E \sum_{j=1}^{\infty} \delta^{j-1}\pi(Q_{j-1}, Z_j, q_j(Q_{j-1}, Z_j)), \tag{2.2}$$

subject to the constraint that $Q_j = Q_{j-1} + q_j$. The expectation in (2.2) is with respect to the distributions of the Z_j. The recursive process works as follows: the firm starts period j with Q_{j-1} units of cumulative production and known cost of production $c(Q_{j-1}, z_j)$. If it produces q_j units, at the end of period j cumulative production has increased to $Q_j = Q_{j-1} + q_j$. With the period's production complete, the firm is able to observe a new random variable (the cost $c(Q_j, z_{j+1})$ for next period's production).

Mazzola and McCardle (1995) consider the stochastic optimization problem that results from (2.2) when the Z_j are independent and identically distributed random variables with known distribution. In this paper we assume that some parameter θ of the distribution of Z is

unknown. The firm has a subjective prior distribution ν over the values of θ. The support of θ is Θ. A random draw Z gives the firm information about the unknown value θ, and the firm uses this information to update its prior distribution via Bayes' rule.

For example, let c_j be the costs in period j corresponding to cumulative production Q_{j-1}. A commonly assumed cost function in the experience-curve literature is one that exponentially decreases with cumulative production, i.e.,

$$c_j = \underline{c} + c_0(Q_{j-1}/Q_0)^{-b},$$

or on taking logarithms, $\ln(c_j - \underline{c}) = \ln c_0 - b \ln(Q_{j-1}/Q_0)$. This cost function has initial cost $\underline{c} + c_0$ that decreases exponentially at rate $b > 0$ to the limiting cost $\underline{c}$. Random variation can be incorporated into the exponentially decreasing cost model as follows:

$$\ln(c_j - \underline{c}) = \ln c_0 - b \ln(Q_{j-1}/Q_0) + Z_j \quad \text{or}$$
$$c_j = \underline{c} + c_0(Q_{j-1}/Q_0)^{-b}e^{Z_j}. \tag{2.3}$$

If the values of Q_0, c_0, $\underline{c}$, and b are known, for example, uncertainty regarding some parameter of the distribution of Z is equivalent to uncertainty regarding the intercept of the log-linear cost function.

Specific distributional assumptions about Z_j are possible. For example, suppose costs c_j in period j satisfy (2.3) where Z_j is normally distributed with unknown mean θ and known precision s (where precision = 1/variance), and the prior on θ is normal with mean a and precision t. The posterior distribution for θ after costs c_j are observed at production level Q_{j-1} is normal with mean

$$\{ta + s[\ln((c_j - \underline{c})/c_0) + b \ln(Q_{j-1}/Q_0)]\}/(t + s)$$

and precision $t + s$. Extending to the case where costs $c_1, c_2, \ldots, c_j$ are observed at associated cumulative production levels $Q_0, Q_1, \ldots, Q_{j-1}$ yields a normal posterior distribution for θ with mean

$$\left\{ta + s \sum_{i=1}^{j} [\ln((c_i - \underline{c})/c_0) + b \ln(Q_{i-1}/Q_0)]\right\}/(t + js)$$

and precision $t + js$. Note that the variance of the posterior distribution converges to 0 as the number of periods increases. Of course, other distributional assumptions and other cost functions are possible.

Before proceeding to more detailed analysis of specific cases, we concern ourselves with the general issue of verifying the optimality equation and proving the existence of an optimal policy in a Bayesian experience-curve model. Let ν represent the prior distribution of the unknown parameter θ, and let $B(\nu, z)$ generate the posterior distribution of θ given the observation Z. That is, if $f(z|\theta)$ is the density for Z, then $B(\nu, z)$ generates a posterior distribution ν^+ for θ via

$$\nu^+(A) = \frac{\int_A f(z|\theta)d\nu(\theta)}{\int_\Theta f(z|\theta)d\nu(\theta)},$$

for all Borel sets A in Θ.

The firm's state $s = (Q, \nu, z) \in S$ is captured by its cumulative production to date, Q; the distribution ν over the unknown parameter; and the observed value z of the current period's random variable. The distribution ν has not been updated to account for this period's observed z. The return function $v(Q, \nu, z)$ is given by

$$v(Q, \nu, z) = \max_{q \in \Gamma} [\pi(Q, z, q) + \delta E v(Q + q, B(\nu, z), Z')]. \tag{2.4}$$

The state next period is determined by cumulative production to date, $Q + q$; next period's beliefs, $B(\nu, z)$, updated by this period's observed value of z; and next period's random shock Z', where the belief regarding the distribution of Z' is represented by $B(\nu, z)$. In the Bayesian model of (2.4), next period's Z' is only conditionally independent of this period's z, where the conditioning is with respect to the parameter θ. The appropriate conditioning information is carried in $B(\nu, z)$.

Let

$$T(Q, \nu, z, q) = (Q + q, B(\nu, z), Z')$$

be the transition mapping from $S \times \Gamma$ into the state space S. To prove the existence of an optimal policy for (2.4), we will need (among other requirements) to establish the continuity of T. Clearly, T is continuous in Q and q. Hence, the continuity of T is determined by $B(\nu, z)$ and Z'. Let $\phi_1, \phi_2, \cdots$ be a sequence of probability measures weakly converging to a probability measure ϕ; let $x_1, x_2, \cdots$ be a sequence of points converging to x; and let X_i' and X' be the random variables associated with $B(\phi_i, x_i)$ and $B(\phi, x)$, respectively. We say

that T is *continuous* if $B(\phi_i, x_i)$ converges weakly to $B(\phi, x)$, and X_i' converges in distribution to X'. Because the distribution of X_i' is determined by $B(\phi_i, x_i)$, it suffices to establish that $B(\phi_i, x_i)$ converges weakly to $B(\phi, x)$ to prove the continuity of T. That is, if Bayesian updating is continuous, then the transition function is continuous.

Unfortunately, Bayesian updating is not always continuous; discontinuities may occur, for example, if $f(z|\theta)$ is unbounded or discontinuous and the support of Z and θ is not compact. Easley and Kiefer (1988) prove that if $f(z|\theta)$ is jointly continuous, the space Θ is compact, the support of Z is compact, and if $f(z|\theta)$ is strictly positive on the support of Z for all θ, then $B(\nu, z)$ is continuous. The compactness assumptions of Easley and Keifer rule out the case described above in which Z has a normal distribution. The existence of a bound on $f(z|\theta)$, however, can substitute for compactness of the state space, as Lemma 1 demonstrates.

LEMMA 1. *If $f(z|\theta)$ is jointly continuous, strictly positive on the support of Z for all θ, and bounded, then $B(\nu, z)$ is continuous.*

PROOF. See appendix.

In addition to the example of a normal prior distribution with normally distributed observations, there are a number of natural conjugate-pair distributions that satisfy the assumptions of Lemma 1 but violate the compactness assumptions in Easley and Kiefer. For example, a Gamma prior distribution with exponentially distributed observations, and a Gamma prior distribution with Poisson distributed observations. (See Degroot 1970, Chapter 9 for a description of these and other natural conjugate pairs.) Taken together, Lemma 1 above and Theorem 1 of Easley and Kiefer prove that Bayesian updating is continuous for most distributions of interest.

The main dynamic programming theorem on which we rely, due to Maitra (1967), is that if the current period reward function π is bounded and upper semicontinuous,[1] if the decision space Γ is compact, and if the transition function $T(Q, \nu, z, q)$ is continuous, then there is a unique bounded upper semicontinuous function v^* satisfying (2.4), v^* equals the optimal value of the problem in (2.2), and there is an optimal stationary policy q^* generated from v^*. (See also Bertsekas and Shreve 1978, Chapter 9.) We state this as Theorem 1.

THEOREM 1. *If current period profits are bounded and upper semicontinuous, and if Bayesian updating is continuous, then there is a unique function v^* satisfying (2.4), v^* is bounded and upper semicontinuous, and it yields an optimal policy.*

There are several ways in which a firm can act myopically. Of these, two are of particular interest. The firm can ignore all future effects of its current decision and optimize current period profits only. We call this behavior *fully myopic* and denote the resulting policy by $\hat{q}(Q, z)$; i.e., $\hat{q}(Q, z)$ maximizes $\pi(Q, z, q)$. This corresponds to the familiar notion of myopic behavior in the dynamic programming literature. Another way that the firm can be shortsighted is to ignore the information content of this period's random shock. We call this behavior *learning myopic*, and denote the resulting policy by $\tilde{q}(Q, z)$. That is, the learning-myopic firm correctly anticipates that next period's cumulative production increases to $Q + q$ and considers the effects of this period's decision on the discounted expectation of all future payoffs, but does not anticipate updating the distribution over the uncertainty. The learning-myopic policy $\tilde{q}$ optimizes the dynamic recursion

$$\tilde{v}(Q, z) = \max_{q \in \Gamma} [\pi(Q, z, q) + \delta E\tilde{v}(Q + q, Z')], \quad (2.5)$$

where the distribution of Z' is assumed to be the same as that which generated z. Equation (2.5) is a stationary version of (2.4). The fully myopic policy $\hat{q}(Q, z)$ and the learning-myopic policy $\tilde{q}(Q, \nu, z)$ provide natural benchmarks against which to compare the optimal policy $q^*(Q, \nu, z)$.

Two standard results of the deterministic experience-curve literature are that optimal production exceeds the myopic level (the meaning of myopic is unambiguous in the deterministic case) and that optimal production increases with cumulative production. The value of production learning is positive, causing the firm to produce more than the myopic amount that equates short run marginal costs and short run marginal revenues. Al-

[1] Recall that a function $f(x)$ is *upper semicontinuous* if for any sequence x_n converging to x, $\limsup f(x_n) \leq f(x)$. Continuous functions are upper semicontinuous.

though the positive value of production learning is decreasing with cumulative production, this effect is outweighed by the direct cost effect, with the result that optimal production increases with cumulative production. To determine whether similar results hold in our Bayesian model, further characterization of the profit function is necessary. We now analyze two variants of the exponentially decreasing cost experience curve in more detail.

2.1. Exponentially Decreasing Costs I

In the model with costs described by (2.3), profits are bounded above because costs are bounded below by $\underline{c}$ and production is bounded above by $\bar{q}$. Call these maximum available single period profits $\bar{\pi} \equiv \max_{q\in[0,\bar{q}]}[(p(q) - \underline{c})q]$. If the support of Z is unbounded above, then costs are unbounded above and profits are not necessarily bounded below, violating one of the assumptions of Theorem 1. We begin by showing that with a simple restriction of the policy space, a restriction that is without loss of optimality, the boundedness of the profit function can be recovered and Theorem 1 applied. Because profits are bounded above by $\bar{\pi}$, the most that the firm could hope to gain from next period forward under any policy is $\delta\bar{\pi}/(1 - \delta)$. If the firm chooses current period production that generates current period profits less than $-\delta\bar{\pi}/(1 - \delta)$, then the value function is negative. Such a production policy is dominated by one which produces 0 from this period forward, yielding a value of 0. Thus, without loss of optimality, we can restrict attention to policies that do not let current period profits fall below $-\delta\bar{\pi}/(1 - \delta)$. Henceforth, we assume that the set Γ of feasible policies includes this restriction. With this restriction, profits are bounded, and the existence results of Theorem 1 hold.

We are able to use the first-order conditions, which are necessary conditions for an optimal policy, generated from (2.4) and (2.3) to evaluate the impact of Bayesian learning on the optimal policy. Differentiating and rearranging terms yields the first-order condition at the jth production decision:

$$p'(q_j)q_j + p(q_j) - \underline{c} - c_0\left(\frac{Q_{j-1}}{Q_0}\right)^{-b} e^{z_j}$$
$$= -\delta\frac{bc_0}{Q_0}\left(\frac{Q_{j-1} + q_j}{Q_0}\right)^{-(b+1)} E(e^{Z_{j+1}}q_{j+1}). \quad (2.6)$$

As noted earlier, Z_{j+1} is independent of z_j only when conditioned on θ. Because z_j carries information regarding θ, the optimal Bayesian policy q_{j+1} in period $j + 1$ depends on z_j.

When evaluated at the fully myopic policy $\hat{q}$, the left-hand side of (2.6) is equal to zero; i.e., the fully myopic policy ignores the future effects of current period decisions as captured on the right-hand side of (2.6). Note that the right-hand side of (2.6) is negative. If $p''(q)q + 2p'(q) < 0$, then marginal revenues are decreasing, i.e., $p'(q)q + p(q)$ is decreasing in q. With this assumption, it follows that optimal production exceeds the fully myopic level, extending to this particular Bayesian framework the relation between optimal and myopic policies that holds in the deterministic and stochastic models (see Mazzola and McCardle 1995). We summarize in the following proposition.

PROPOSITION 1. *If costs are exponentially decreasing as in* (2.3), *there exists a unique, bounded, upper semicontinuous value function* v^* *solving* (2.4); *and* v^* *generates an optimal policy* q^*. *In addition, if marginal revenues are decreasing, then the optimal policy exceeds the myopic policy; i.e.,* $q^*(Q, \nu, z) \geq \hat{q}(Q, z)$.

Comparison with the learning-myopic policy is more difficult. Recall that the while the fully myopic policy ignores all future effects of current period decisions, the learning-myopic policy ignores only the information effects. Thus, the learning-myopic policy would not solve for the zero of the left-hand side of (2.6), but instead, it would compute the right-hand side differently from the optimal policy. For example, if in (2.3) the prior on θ at stage j is normal with mean a and precision t, and the conditional distribution of Z_j is normal with mean θ and precision s, then the unconditional distribution of Z_{j+1} after z_j is observed is normal with mean $(ta + sz_j)/(t + s)$ and precision $s(t + s)/(t + 2s)$. The optimal policy uses this distribution to compute $E(e^{Z_{j+1}}q_{j+1})$ on the right-hand side of (2.6). The learning-myopic policy, on the other hand, assumes that Z_{j+1} has the same unconditional distribution as Z_j, i.e., normal with mean a and precision $st/(s + t)$. Thus, it is easy to see why an unambiguous comparison is not possible, even with an assumption of normality: the distributions against which the expectation of $(e^{Z_{j+1}}q_{j+1})$ is taken have different means and different variances.

Computational results on a two-period model presented in §3 reinforce the ambiguity of the relationship between optimal production and learning-myopic production. The question of whether the optimal policy $q^*(Q, \nu, z)$ is monotonically increasing in cumulative production Q is delayed until §3.

2.2. Exponentially Decreasing Costs II

Rather than having uncertainty regarding the intercept of the log-linear cost function as in §2.1, we now consider a model with costs

$$\ln(c_j - \underline{c}) = \ln c_0 - \beta \ln(Q_{j-1}/Q_0) + U_j, \quad (2.7)$$

where there is uncertainty regarding the value of β. In the model described by (2.7) the slope of the log-linear cost function, i.e., the cost decline rate, is unknown. The U_j term captures the random variation in the costs, while β captures the firm's uncertainty about the rate at which costs will decline on average. If there were no variation in costs, i.e., if $U_j \equiv 0$, then the firm would learn the true value of β after one period. Because there is variation in costs, it takes time for the firm to resolve its uncertainty about β.

The boundedness assumption of Theorem 1 for current period profits π is circumvented in the same manner as in §2.1. Hence, there is a unique bounded upper semicontinuous value function v^* satisfying (2.4) with costs as in (2.7), and v^* yields an optimal policy.

Having established the existence of an optimal policy satisfying (2.4), consider the corresponding first-order conditions. The observed random shock in period j is $[(Q_{j-1}/Q_0)^{-\beta_j}e^{u_j}]$. The firm cannot decouple β_j from u_j; i.e., the firm does not observe two random components β_j and u_j and combine them, but rather the firm observes costs. Based on its knowledge of the distributions of the random variables, it is able to update the distribution of β. (The example considered shortly will clarify this issue.) The first-order conditions are given by

$$p'(q_j)q_j + p(q_j) - \underline{c} - c_0\left(\frac{Q_{j-1}}{Q_0}\right)^{-\beta_j} e^{u_j}$$
$$= -\delta E\left[\frac{c_0\beta_{j+1}}{Q_0}\left(\frac{Q_{j-1}+q_j}{Q_0}\right)^{-(\beta_{j+1}+1)} e^{U_{j+1}}q_{j+1}\right]. \quad (2.8)$$

As in equation (2.6), with the assumption of decreasing marginal revenues, the right-hand side of (2.8) determines whether optimal production exceeds fully myopic production. Unfortunately, the right-hand side can be positive, in which case the optimal policy falls below the fully myopic level. This term could be positive, for instance, if the prior next period has a negative mean for β. Thus, if the cost decline rate is unknown and the firm uses observations of costs to learn about the cost decline rate, it does *not* necessarily follow that optimal production exceeds the myopic amount. Note, however, that an assumption that the domain of β is nonnegative is sufficient to guarantee that optimal production exceeds fully myopic levels. Also observe that a negative expectation for β implies that the firm believes the experience curve is upward sloping, i.e., costs are increasing with cumulative production. We summarize in Proposition 2.

PROPOSITION 2. *If costs are exponentially decreasing as in* (2.7), *there exists a unique, bounded, upper semicontinuous value function* v^* *solving* (2.4), *and* v^* *generates an optimal policy* q^*. *In addition, if marginal revenues are decreasing and if the domain of* β *is nonnegative, then the optimal policy exceeds the fully myopic policy; i.e.,* $q^*(Q, \nu, z) \geq \hat{q}(Q, z)$.

Again, further insight is gained by considering a specific example. As is common in regression models fitting data to experience curves, suppose U_j is normally distributed with mean 0 and precision s. Assume that the prior on β_j is normal with mean b_j and precision t. It is notationally convenient to consider the random variable $Z_j = [\beta_j \ln(Q_{j-1}/Q_0) - U_j]$. As noted above, Z_j is the random variable observed by the firm, and the firm uses the functional relationship of (2.7) to update its prior distribution. From the assumptions on U_j and β_j, it follows that Z_j is normally distributed with unknown mean θ and precision s. The prior on θ is normal with mean $b_j \ln(Q_{j-1}/Q_0)$ and precision $t/\ln^2(Q_{j-1}/Q_0)$. It is possible to substitute (b, t) for the distribution ν in the functional equation

$$v^*(Q, (b, t), z) = \max_{q\in\Gamma} [[p(q) - \underline{c} - c_0e^{-z}]q + \delta Ev^*(Q + q, B(b, t, z), Z')], \quad \text{where}$$

$$B(b, t, z) = \left(\frac{tb + sz\ln(Q/Q_0)}{t + s\ln^2(Q/Q_0)}, t + s\ln^2(Q/Q_0)\right).$$

That is, if the firm has a prior mean of $-b$ for the slope of the log-linear equation and observes a value of z, the updated mean for the slope is

$$-[tb + sz \ln(Q/Q_0)]/[t + s \ln^2(Q/Q_0)].$$

Because the domain of Z includes negative values in the normally distributed case, we cannot guarantee that the updated mean for the slope remains negative even if the prior mean $-b < 0$.

3. The Finite-horizon Model

To provide an explicit solution to the first-order conditions (2.6) or (2.8), it is natural to restrict the model to a finite number of periods and calculate the optimal policy for the finite-horizon version of the firm's problem. Indeed, much of the research cited in the introduction assumes a finite horizon (see, e.g., Spence 1981, Harpaz et al. 1982, Fudenberg and Tirole 1983, Dada and Srikanth 1990). Characterizing the resulting finite-horizon optimal policy and value function, in addition to being worthwhile in its own right, gives insight into the characteristics of the optimal policy and value function for the infinite-horizon problem. In this section we develop a finite-horizon version of the problem in (2.2) and (2.4) and provide explicit numerical solutions of two-period versions of the model developed in §2.1 with the associated normality assumptions.

With a finite number N of periods, the firm's problem is to find a policy $q_1(Q_0, z_1), q_2(Q_1, z_2), \ldots, q_N(Q_{N-1}, z_N)$ to maximize discounted expected profits

$$E \sum_{j=1}^{N} \delta^{j-1}\pi(Q_{j-1}, Z_j, q_j(Q_{j-1}, Z_j)), \tag{3.1}$$

subject to the constraint that $Q_j = Q_{j-1} + q_j$. As in (2.2), the expectation in (3.1) is with respect to the distributions of the Z_j. The jth period value function generated from (3.1) is given by

$$v_j(Q_{j-1}, \nu_{j-1}, z_j) = \max_{q \in \Gamma} [\pi(Q_{j-1}, z_j, q) + \delta E v_{j+1}(Q_{j-1} + q, B(\nu_{j-1}, z_j), Z_{j+1})]. \tag{3.2}$$

To complete the recursion, we assume that the value function in period $N + 1$ is identically zero: $v_{N+1} \equiv 0$.

With the same assumptions as made in §2, we immediately have the next Theorem (see, e.g., Bertsekas and Shreve 1978, Chapter 8).

THEOREM 2. *If current period profits are bounded and upper semicontinuous, and if Bayesian updating is continuous, then the jth stage optimal value function v_j^* satisfies (3.2), v_j^* is bounded and upper semicontinuous, and an optimal policy exists.*

Because the first-order conditions are the same in both the finite and infinite horizon, with the exception of the last period, the following proposition follows from Propositions 1 and 2.

PROPOSITION 3. *If marginal revenues are decreasing, and costs are given either by (2.3) or by (2.7) with the additional assumption that the domain of β is positive, then the finite-horizon optimal policy exceeds the fully myopic policy.*

Armed with these results, we now compute the optimal policy and value function for an $N = 2$ period model. Note that the characterization of the optimal policy and value function we obtain in the two-period model immediately carries over to a finite-horizon model of any length.

Let demand be given by $p(q) = 10 - q$; costs be given by (2.3) with $\underline{c} = 0$, $c_0 = 1$, and $Q_0 = 1$, i.e.,

$$c(Q, Z) = Q^{-b}e^{Z}; \tag{3.3}$$

and the discount factor $\delta = 0.9$. There are two production periods: 1 and 2. Assume that the random shocks Z_1 and Z_2 each has a normal distribution with mean θ and precision s, and the prior on θ before Z_1 is observed is normally distributed with mean 0 and precision t. Then the marginal distribution of Z_1 is normal with mean 0 and precision $st/(s + t)$, and the marginal distribution of Z_2 given z_1 is normal with mean $sz_1/(t + s)$ and precision $s(t + s)/(t + 2s)$. Because marginal revenues are decreasing, optimal production exceeds fully myopic production as per Proposition 3.

In period 2, the firm chooses q_2^* to maximize

$$\begin{aligned}\pi(Q_1, z_2, q_2) &= [p(q_2) - c(Q_1, z_2)]q_2 \\ &= [10 - q_2 - Q_1^{-b}e^{z_2}]q_2.\end{aligned} \tag{3.4}$$

Solving yields[2]

$$q_2^*(Q_1, z_2) = [5 - Q_1^{-b}e^{z_2}/2]^+. \tag{3.5}$$

[2] $x^+ = \max\{0, x\}$.

Table 1 **Bayesian Learning Model: Optimal Production Quantity ($b = 0.1$, $s = 1.0$)**

$z_1/\sigma \backslash t$	0.001	0.01	0.1	0.5	1	1.5	2.5	5
−2	5	5	5	4.99	4.98	4.98	4.98	4.97
−1.5	5	5	5	4.97	4.96	4.95	4.94	4.93
−1	5	5	4.99	4.93	4.9	4.89	4.88	4.86
−0.5	5	5	4.92	4.82	4.78	4.77	4.76	4.75
0	4.54	4.54	4.54	4.54	4.54	4.54	4.54	4.54
0.5	0	0	2.45	3.86	4.03	4.09	4.14	4.18
1	0	0	0	2.26	3.01	3.24	3.42	3.56
1.5	0	0	0	0	0.97	1.64	2.14	2.49
2	0	0	0	0	0	0	0	0.69

Table 2 **Bayesian Learning Model: Optimal Production Quantity ($b = 0.5$, $s = 1.0$)**

$z_1/\sigma \backslash t$	0.001	0.01	0.1	0.5	1	1.5	2.5	5
−2	5	5	5	5	5.01	5.01	5.02	5.02
−1.5	5	5	5	4.99	4.99	4.99	4.99	4.99
−1	5	5	4.99	4.96	4.95	4.94	4.93	4.93
−0.5	5	5	4.94	4.87	4.84	4.83	4.82	4.82
0	4.63	4.63	4.63	4.63	4.62	4.62	4.62	4.62
0.5	0	0	2.73	4.01	4.16	4.21	4.25	4.27
1	0	0	0	2.54	3.21	3.41	3.57	3.68
1.5	0	0	0	0	1.4	1.96	2.37	2.67
2	0	0	0	0	0	0	0.55	1.08

Plug (3.5) into (3.4) and call the resulting function $\pi_2^*(Q_1, z_2)$. Recalling that $Q_0 = 1$ and $Q_1 = Q_0 + q_1$ generates the first-period problem of choosing q_1^* to maximize

$$\pi(1, z_1, q_1) + \delta E_{Z_2}\pi_2^*(1 + q_1, Z_2), \qquad (3.6)$$

where the distribution of Z_2 is normal with mean $sz_1/(t + s)$ and precision $s(t + s)/(t + 2s)$ as specified in the previous paragraph. The fully myopic policy $\hat{q}_1$ maximizes $\pi(1, z_1, q_1)$, ignoring the second period payoff. The learning-myopic policy $\tilde{q}_1$ maximizes (3.6) under the assumption that Z_2 has a normal distribution with mean 0 and precision $st/(s + t)$.

It is also useful to introduce the limiting fully myopic policy as a benchmark for comparison. As cumulative production Q increases to infinity, costs $c(Q, z)$ as given in (3.3) converge to 0 for all values of z. This is the lowest cost that the firm could ever hope to achieve. Let $\hat{q}_\infty$ maximize the limiting profit function with zero costs. Solving for our example yields $\hat{q}_\infty = 5$. Mazzola and McCardle (1995) show that if costs are decreasing deterministically, then $\hat{q}_\infty$ exceeds the optimal production in every period, and optimal production is monotonically increasing in Q. We now consider whether similar results hold for the Bayesian policy.

Summary information from a computational analysis of (3.6) is contained in Tables 1 and 2 and Figure 1. Tables 1 and 2 report the optimal production quantity $q_1^*(z_1)$ as a function of the prior precision t and the standardized value of the observed shock z_1 (standardized by dividing z_1 by its standard deviation) for two different cost decline rates, $b = 0.1$ and $b = 0.5$. A cost decline rate of $b = 0.1$ implies that for every doubling of cumulative production, costs go down by 7% (i.e., a 93% experience curve), while a cost decline rate of $b = 0.5$ implies that costs go down by 29% for every doubling of cumulative production (a 71% experience curve). In both of these figures the variance of Z conditioned on θ is 1, i.e., $s = 1$. Figure 1 plots the difference between optimal production and learning-myopic production.

Observe from Tables 1 and 2 that the precision t of the prior affects the optimal first-period policy in different ways depending on the observed value z_1 of the first-period's random shock and on the cost decline rate b. In both Tables 1 and 2, for large positive values of z_1, optimal first-period production q_1^* increases with the

Figure 1 **Bayesian Learning ($b = 0.1$, $s = 1.0$)**

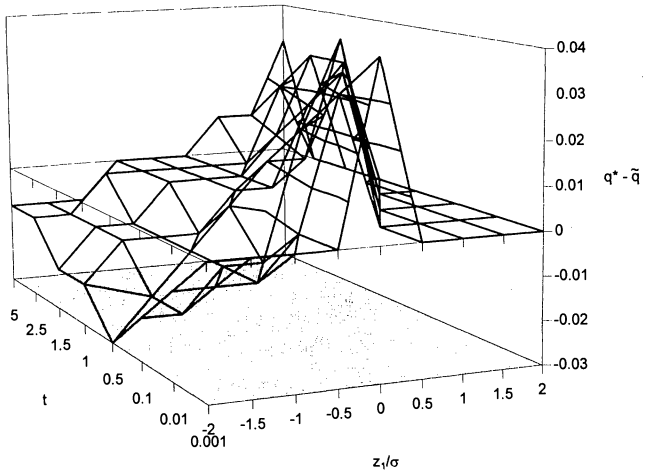

precision t; while for large negative values of z_1, optimal first-period production is decreasing in Table 1 and increasing in Table 2. Thus, there is no unambiguous monotonic relationship between the informativeness of the prior (as measured by the precision) and optimal production in the finite-horizon model of (3.2) with costs as given in (2.3).

Of more import is the fact that in Table 2, the optimal first-period production q_1^* exceeds the limiting fully myopic policy $\hat{q}_\infty = 5$ for some values of z_1. For instance, $q_1^* = 5.02$ for $t = 5$ and a standardized z_1 of -2. Because the limiting fully myopic value is also the limiting optimal policy (the limit as Q converges to infinity), we see that q^* is not increasing in Q. This provides a counterexample in the finite-horizon case to the conjecture left open at the end of §2.1.

Figure 1 compares the optimal policy with the learning-myopic policy; it, too, shows that an unambiguous ordering is unlikely. For large positive values of z_1, the optimal policy tends to exceed the updating-myopic policy; for large negative values of z_1, the optimal policy tends to fall below the learning-myopic policy. Figure 1 also shows that there are values of z where both policies choose the same production. Overall, however, the data show no clear unambiguous comparative relationship exists similar to the relation between the optimal policy and the fully myopic policy. Thus, while Proposition 3 provides conditions under which $q^*(Q, \nu, z) \geq \hat{q}(Q, z)$ for all Q, ν, and z, such an ordering does not exist between $q^*(Q, \nu, z)$ and $\tilde{q}(Q, \nu, z)$ in the finite-horizon case.

To provide some intuition regarding the ambiguity of the relation depicted in Figure 1, consider that a learning-myopic firm, when it observes a large positive value of z_j, concludes that it has just had a hit of bad luck, and decreases current period production accordingly. No revision of future expectations is undertaken, and the learning-myopic firm assumes that tomorrow will likely be a better day. On the other hand, if an optimizing firm observes a large positive value of z_j, the posterior mean of θ is shifted upward, with the result that the firm expects costs overall to be higher. Higher costs in the future lead to lower production levels in the future. To balance against the negative future expectations, the firm compensates by decreasing current period production less than the learning-myopic firm. Similarly, the learning-myopic firm views a large negative value of z_j as good luck, and immediately takes full advantage of it. The next period is not likely to be so lucky. For the optimizing firm, however, for large negative values of z_j, the posterior mean of θ is shifted downward, with the result that the firm expects costs overall to be lower. Lower costs lead to greater production levels in the future, decreasing the need to take advantage of this period's low costs.[3]

Thus, we have shown in a finite-horizon stochastic model with uncertainty regarding a parameter of the experience curve that optimal production is not bounded above by the limiting fully myopic production, that optimal production is not increasing in cumulative production, and that optimal production does not exceed learning-myopic production. Recall that two of the core results of the deterministic experience-curve literature are that optimal production is bounded above by the limiting myopic production and optimal production is increasing in cumulative production. Thus, these results are not robust to the determinism assumption.

Because information has value, a simple intuition might lead one to expect that optimal (Bayesian) production would exceed the learning-myopic levels. This intuition is confirmed, for example, in the model analyzed by Harpaz et al. (1982) and Thompson and Horowitz (1993). In that model, however, there is no link between the current period's decision and the next period's state, corresponding to the transition in our model from Q to $Q + q$. Thus, there is no difference in their model between a fully myopic and a learning-myopic manager, and the comparison they present between the Bayesian and nonexperimenting manager is equivalent to the comparison we present between the Bayesian and fully myopic manager.

The differences reported in Figure 1 appear relatively small, and Thompson and Horowitz (1993) show that the difference between Bayesian and myopic production is relatively small in their model. This leads to a natural question regarding the relative value of

[3] Of course, neither of the situations just described has taken into account the additional effects of updating on the variance of the posterior distribution as well as the effects of decreased variance on production.

performing the more complex Bayesian analysis. We complete our discussion of the finite-horizon model by exploring the potential benefits that can accrue from the application of the Bayesian approach to managing experience-curve uncertainty. Again, we consider a two-period example.

Suppose that two firms face identical production-planning decisions and that one firm employs Bayesian learning and the other uses the learning-myopic model. We assume the same normally distributed uncertainty structure as in the example detailed in Tables 1 and 2. Using (3.2), the Bayesian Firm A determines its optimal production policy q^*, as specified by $q_j^*(Q, z)$, $j = 1, 2$. Similarly, the learning-myopic Firm B determines its optimal policy $\tilde{q}$, as specified by $\tilde{q}_j(Q, z)$, $j = 1, 2$.

Observe that for given values of Q_0 and z, it is not meaningful to compare directly the corresponding value functions of these two policies, since the value functions depict the state of the world differently. To assess the potential benefits of the Bayesian approach, it is necessary to establish a means for comparing these two policies. Because a key value of a Bayesian approach is that it can correct an incorrect prior belief through updating, in order to compare the approaches, we check what happens when both firms are wrong in their initial forecasts of the means of Z_1 and Z_2. We examine the policies q^* and $\tilde{q}$ in the light of their respective performance with respect to some true underlying distribution Z^P, known only to an omniscient observer.

In the absence of uncertainty concerning the true distribution, the omniscient observer seeks a policy q that maximizes

$$v_P(q) = E \sum_{j=1}^{N} \delta^{j-1} \pi(Q_{j-1}, Z^P, q_j(Q_{j-1}, Z_j^P)), \quad (3.7)$$

subject to $Q_j = Q_{j-1} + q_j$. Denote such an optimal policy by q^P, and let $v_P^{**} = v_P(q^P)$ be the expected return to the omniscient observer from following an optimal policy. We can then use (3.7) to assess the policies q^* and $\tilde{q}$. In particular, let $v_P^* = v_P(q^*)$, and let $\tilde{v}_P = v_P(\tilde{q})$. The ratio,

$$\frac{v_P^* - \tilde{v}_P}{v_P^{**}}, \quad (3.8)$$

then measures the increase in returns attributable to the Bayesian approach relative to the omniscient observer.

Considering now a specific example, let demand be given by $p(q) = 400 - 0.1q$, and costs be given by (2.3) with $\underline{c} = 0$, $c_0 = 400$, $Q_0 = 1$, and $b = 0.5$ (representing a 71% experience curve, as in Table 2). The maximum production quantity $\bar{q} = 2500$; because of the size of the problem, the production quantity q is assumed to be an integer. The distributions of Z_1 and Z_2 are as described above with $s = 1$ and $t = 0.1$. The discount factor $\delta = 0.9$. Assume that both firms are off by one standard deviation in their initial estimates of the means of Z_1 and Z_2. Specifically Z^P is normally distributed with mean 1 and precision s, while both firms assess the mean to be 0.

Solving the problem and determining the optimal policies for the Bayesian firm, the learning-myopic firm, and the omniscient observer, we have that $v_P^* - \tilde{v}_P = 9596.3$ and $v_P^{**} = 146536.1$. Using (3.8), this represents a 6.5% increase in profits resulting from the use of the Bayesian model. This increase is representative of the results we have obtained using a variety of parameter values.

This simple two-period example shows that the benefits from employing the Bayesian approach can be substantial. In competitive markets, increased profit margins of this magnitude can often determine the firm's viability in the marketplace.

4. Conclusion

Throughout this paper we have focused on solving for and characterizing an optimal production policy when costs follow an experience curve, there is random variation in the costs, and some parameter of the distribution of those costs is unknown. Our first result provides conditions under which Bayesian updating is continuous (in the sense of weak convergence of measures). Armed with this result, we prove the existence of a unique bounded value function satisfying a dynamic stochastic recursion and the existence of an optimal policy. To obtain further characterizations of the optimal policy, we consider both finite and infinite-horizon examples and calculate the optimal policy explicitly for a two-period model. While we are able to provide conditions guaranteeing that the optimal policy exceeds the myopic

policy, the question of whether optimal production increases with cumulative production in the infinite-horizon case remains open. In §3, however, we demonstrate the property does not hold for the finite-horizon case.

Returning to the manager's problem posed in the introduction, suppose there is a fixed cost K for adopting the new process technology. By having solved the optimal production policy, and knowing her prior distribution ν over the uncertainty, the manager can compare the expected value $v^*(Q_0, \nu, z_1)$ of adopting the new technology to the fixed costs of adoption and decide whether or not adoption is worthwhile. Because they directly affect the current period reward, changes in several of the parameters directly affect the value function and have an impact on the adoption decision. For example, in the model with costs as given in (2.3), an increase in the cost decline rate b leads to an increase in profits, hence an increase in the value function (though not necessarily in the optimal policy). Reductions in the initial cost or in the limiting cost also increase the optimal value.[4]

[4] The authors gratefully acknowledge the support of the Business Associates Fund of the Fuqua School of Business. We thank Professor James E. Smith, the referees, and the Area Editor for detailed comments on an earlier draft of this paper. Seminar participants at INSEAD provided helpful comments.

Appendix

PROOF OF LEMMA 1. Let $\nu_1, \nu_2, \cdots$ be a sequence of probability measures on Θ converging weakly to the probability measure $\hat{\nu}$, and let $x_1, x_2, \cdots$ be a sequence of points converging to $\hat{x}$.[5] The assumption that $f(x|\theta) > 0$ on the support of x for all θ ensures that $B(\nu, x)$ is well defined. The goal is to show that $B(\nu_j, x_j)$ converges weakly to $B(\hat{\nu}, \hat{x})$. Equivalently, if ν_j^+ is the measure generated from $B(\nu_j, x_j)$, and $\hat{\nu}^+$ is the measure generated by $B(\hat{\nu}, \hat{x})$, then $\nu_j^+(A)$ converges to $\hat{\nu}^+(A)$ for all $\hat{\nu}^+$ continuity sets A; i.e., all Borel sets A with $\hat{\nu}^+$ boundary measure zero. (See Billingsley 1979, Section 25.)

Let A be a $\hat{\nu}$ continuity set. Recall that

$$\nu_j^+(A) = \frac{\int_A f(x_j|\theta)d\nu_j(\theta)}{\int_\Theta f(x_j|\theta)d\nu_j(\theta)} \quad \text{and} \quad \hat{\nu}^+(A) = \frac{\int_A f(\hat{x}|\theta)d\hat{\nu}(\theta)}{\int_\Theta f(\hat{x}|\theta)d\hat{\nu}(\theta)}.$$

[5] The measures ν_j are not necessarily the ones generated by Bayesian updating. We make the notational substitution of x for z to avoid confusion.

Concentrating on the numerators first,

$$\begin{aligned}
&\left|\int_A f(x_j|\theta)d\nu_j(\theta) - \int_A f(\hat{x}|\theta)d\hat{\nu}(\theta)\right| \\
&\quad \le \left|\int_A f(x_j|\theta)d\nu_j(\theta) - \int_A f(x_j|\theta)d\hat{\nu}(\theta)\right| \\
&\quad + \left|\int_A f(x_j|\theta)d\hat{\nu}(\theta) - \int_A f(\hat{x}|\theta)d\hat{\nu}(\theta)\right| \\
&\quad \le \int_A f(x_j|\theta)|d\nu_j(\theta) - d\hat{\nu}(\theta)| + \int_A |f(x_j|\theta) - f(\hat{x}|\theta)|d\hat{\nu}(\theta).
\end{aligned}$$

There exists an M such that $f(x|\theta) \le M$ for all x and θ. Hence, the first term in the last expression of the above inequality is less than or equal to $M|\nu_j(A) - \hat{\nu}(A)|$, which converges to 0 by the weak convergence of ν_j to $\hat{\nu}$. The second term converges to 0 by the pointwise convergence of $f(x_j|\theta)$ to $f(\hat{x}|\theta)$ (recall that $f(x|\theta)$ is jointly continuous), the boundedness of $f(x|\theta)$, and the fact that $\hat{\nu}$ can be viewed as a measure on A.

Similar arguments hold for the denominator. Hence, $\nu_j^+(A)$ converges to $\hat{\nu}^+(A)$ for all $\hat{\nu}$ continuity sets A. Observing that the measures $\hat{\nu}$ and $\hat{\nu}^+$ have the same continuity sets completes the proof. □

References

Adler, P. S. and K. B. Clark, "Behind the Learning Curve: A Sketch of the Learning Process," *Management Sci.*, 37, 3 (1991), 267–281.

Bertsekas, D. P. and S. E. Shreve, *Stochastic Optimal Control: The Discrete Time Case*, Academic Press, New York, 1978.

Billingsley, P., *Probability and Measure*, John Wiley & Sons, New York, 1979.

Clarke, F. H., M. N. Darrough, and J. M. Heineke, "Optimal Pricing Policy in the Presence of Experience Effects," *J. Business*, 55, 4 (1982), 517–530.

Dada, M. and K. N. Srikanth, "Monopolistic Pricing and the Learning Curve: An Algorithmic Approach," *Oper. Res.*, 38, 4 (1990), 656–666.

DeGroot, M. H., *Optimal Statistical Decisions*, McGraw-Hill, New York, 1970.

Dolan, R. J. and A. P. Jeuland, "Experience Curves and Dynamic Demand Models: Implications for Optimal Pricing Strategies," *J. Marketing*, 45 (1981), 52–62.

Easley, D. and N. M. Kiefer, "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56, 5 (1988), 1045–1064.

Ebert, R. J., "Aggregate Planning with Learning Curve Productivity," *Management Sci.*, 23, 2 (1976), 171–182.

Fine, C. H. and E. L. Porteus, "Dynamic Process Improvement," *Oper. Res.*, 37, 4 (1989), 580–591.

Fudenberg, D. and J. Tirole, "Learning-by-doing and Market Performance," *Bell J. Economics*, 14 (1983), 522–530.

Gaimon, C., "Incremental Capacity Expansion: The Effect of Lead Time and Learning," Georgia Institute of Technology Working Paper, 1992.

Harpaz, G., W. Y. Lee, and R. L. Winkler, "Learning, Experimentation, and the Optimal Output Decisions of a Competitive Firm," *Management Sci.*, 28, 6 (1982), 589–603.

Hax, A. C. and N. S. Majluf, "Competitive Cost Dynamics: The Experience Curve," *Interfaces*, 12, 5 (1982), 50–61.

Henderson, B. D. and A. J. Zakon, "Pricing Strategy: How to Improve It (the Experience Curve)," in Kenneth Albert (Ed.), *Handbook of Business Problem Solving*, McGraw Hill, New York, 1980.

Hiller, R. S. and J. F. Shapiro, "Optimal Capacity Expansion Planning When There Are Learning Effects," *Management Sci.*, 32, 9 (1986), 1153–1163.

Kalish, S., "Monopolist Pricing with Dynamic Demand and Production Cost," *Marketing Sci.*, 2, 2 (1983), 135–159.

Karwan, K. R., J. B. Mazzola, and R. C. Morey, "Production Lot Sizing Under Setup and Worker Learning," *Naval Res. Logistics*, 35 (1988), 159–175.

Lieberman, M. B., "The Learning Curve and Pricing in the Chemical Processing Industry," *Rand J. Economics*, 15, 2 (1984), 213–228.

Lippman, S. A. and K. F. McCardle, "Uncertain Search: A Model of Search among Technologies of Uncertain Value," *Management Sci.*, 37, 11 (1991), 1474–1490.

Maitra, A., "Discounted Dynamic Programming on Compact Metric Spaces," *Sankhya: Indian J. Statistics*, 30A (1968), 211–216.

Majd, S. and R. S. Pindyck, "The Learning Curve and Optimal Production under Uncertainty," *Rand J. Economics*, 20, 3 (1989), 331–343.

Mazzola, J. B. and K. F. McCardle, "The Stochastic Learning Curve: Optimal Production in the Presence of Learning Curve Uncertainty," *Oper. Res.*, (1995, forthcoming).

Robinson, B. and C. Lakhani, "Dynamic Price Models for New-product Planning," *Management Sci.*, 21, 10 (1975), 1113–1122.

Spence, A. M., "The Learning Curve and Competition," *Bell J. Economics*, 12 (1981), 49–70.

Thompson, P. and I. Horowitz, "Experimentation and Optimal Output Decisions: The Cooperative versus the Entrepreneurial Firm," *Management Sci.*, 39 (1993), 46–53.

Wright, T. P., "Factors Affecting the Cost of Airplanes," *J. Aeronautical Sci.*, 3 (1936), 122–128.

Yelle, L. E., "The Learning Curve: Historical Review and Comprehensive Survey," *Decision Sci.*, 10, 2 (1979), 302–328.

Accepted by Robert T. Clemen; received June 27, 1994. This paper has been with the authors 7 months for 1 revision.