Dynamic Process Improvement

by

Charles H. Fine*
and
Evan L. Porteus**

WP #1952-87 October 1987

* Sloan School of Management
Massachusetts Institute of Technology

**Graduate School of Business
Stanford University

Dynamic Process Improvement

Charles H. Fine and Evan L. Porteus

ABSTRACT

This paper explores the economics of investing in gradual process improvement. We seek to understand better the economics of the well known Just-in-Time and Total Quality Control philosophies, both of which promote gradual technological improvement as a means of achieving increased productivity and manufacturing competitiveness. In addition, there is empirical evidence that the cumulative effect of minor technical changes on productivity is often at least as important as the effect of major technological innovations.

We formulate a general model of dynamic process improvement as a Markov decision process, analyze it, and apply it to the problem of setup reduction and process quality improvement. Instead of a deterministic model with a one-time investment opportunity for a large predictable technological advance, this paper builds a stochastic model with opportunities for many smaller investments over time, and in which potential process improvements may be unpredictable.

We use a somewhat nonstandard formulation in which the immediate return from an investment is the infinite horizon expected present value of the reduction in operating costs of the system, less the amount of that investment. The policy that simply maximizes the immediate return, called the last chance policy, reveals a great deal about the optimal policy. The optimal policy invests no more than the last chance policy. The optimal policy invests in process improvement if and only if the last chance policy does. There is a target state such that both the optimal and last chance policy invest in process improvement if that state has not yet been reached and invest nothing if it has been. We derive fairly restrictive conditions that must be met for optimality of the policy of investing forever in process improvements, to zero defects or zero setup costs. Decreasing the uncertainty of the process (making the potential improvements more predictable) has a desirable effect: the total return is increased and the target state increases, so the ultimate system is more productive. Numerical examples are presented and analyzed.

This paper seeks to understand better the economics of policies that promote continuous gradual technological evolution, as opposed to discontinuous technological change through infrequent adoption of radical innovations. According to Rosenberg (1982, p. 62), "... a large portion of the total growth in productivity takes the form of a slow and often almost invisible accretion of individually small improvements in innovations. The difficulty of perception [of this phenomenon] seems to be due to a variety of causes ... the small size of individual improvements ... a frequent preoccupation with what is *technologically* spectacular rather than *economically* significant ... ." Rosenberg goes on to cite several empirical studies, including that of Hollander (1965) who sought to explain the sources of cost reduction in du Pont's rayon plants. In Rosenberg's (1982, p. 68) words, Hollander found that "the cumulative effect of minor technical changes on cost reduction was actually greater than the effect of major technical change."

Emphasis on gradual process improvement is a central tenet of the Japanese Just-in-Time (JIT) and Total Quality Control (TQC) philosophies, as related for example by Schonberger (1982, 1986) and Hall (1983). These philosophies embody the policy of continual investment in quality improvement and setup cost reduction towards the goals of zero defects and zero cycle inventories. Hayes (1981) and Schonberger (1982), among others, have reported that an important factor in the manufacturing ascendancy of many large Japanese corporations is their application of these philosophies. Schonberger (1987) advocates especially strongly the merits of emphasizing continuous gradual improvement over major technology acquisition. This paper seeks to understand better the economics of gradual process improvement.

To study gradual process improvement, we formulate a stochastic, dynamic programming model. We present numerous general results for our model, including conditions under which it is optimal to pursue process improvements forever, toward zero defects or zero setup costs. These conditions are quite restrictive, and, in general, there is a finite target state for a firm's setup costs and process quality, such that it is optimal to stop investing once the target state has been reached.

In our model formulation, the immediate return from investment in process improvement is the infinite horizon expected net (of investment) present value of the reduction in operating costs, that will be achieved if no further investments are made. The policy that maximizes this immediate return is called the last chance policy, because it gives the optimal amount to invest if there is only one last chance to do so. It is optimal to invest a strictly positive amount in process improvement if and only if the last chance policy invests a strictly positive amount. Thus, the target state for the optimal policy can be found simply by finding the target state for the last chance policy. Furthermore, it is never optimal to invest more than the last chance policy.

We focus primarily on the problem of investing in the improvement of one attribute, such as setup costs. We provide several analytical results for this problem. We formulate and provide numerical results for an example of the problem of investing simultaneously in the improvement of two attributes: setup cost and process quality.

This work is related to previous work by us and several other authors. Porteus (1986) characterizes optimal simultaneous investment in setup cost reduction and quality improvement in a deterministic model with a one-time investment opportunity and logarithmic investment cost functions. Rosenblatt and Lee (1986) formulate and analyze a model similar to that in Porteus (1986), although they do not consider investment in process improvements. Fine (1986) presents a deterministic continuous-time model of investment in quality improvement, where such investments can help firm learning and cost reductions. That work provides conditions under which it is optimal to continue investing toward the attainment of perfect quality (zero defects). Fine (1987) uses a stochastic dynamic programming model to characterize optimal inspection policies, where inspection activities function as investments in future process quality improvements. Tapiero (1987) links together optimal quality inspection policies and the resulting improvements in manufacturing costs. See also Pierskalla and Voelker (1976) and Monahan (1982) for surveys of earlier work on quality models, and Fine and Bridge (1986) for a survey of some managerial literature on quality improvement.

In the next section we formulate the single attribute process improvement investment model. Section 2 describes how our model fits several applications, including setup cost reduction and process quality improvement. Section 3 presents our analytical results for the single attribute model. Section 4 formulates the problem of simultaneously optimizing set up costs and process quality, and describes the results for a numerical example. Section 5 contains concluding remarks.

1. Formulation of the Single Attribute Model

We first consider a discrete time dynamic process improvement problem with a single process attribute. (We formulate and analyze multiattribute problems in section 4.) Let x denote the state of the process being managed. It will be convenient to think of x as the number of "nominal" improvements that have taken place since time zero. (Thus, at time zero, the state equals zero.) In our examples and applications, each nominal improvement represents a fixed percentage reduction in a critical parameter, such as the setup cost in an EOQ model, but other representations are possible. These improvements need not be integer-valued. Indeed, we assume that x can be any nonnegative real number, even if all such numbers cannot be realized with positive probability. Let $g(x)$ denote the discounted present value of the operating costs over an infinite horizon if x (nominal) improvements have been achieved and no further improvements are made.

If an amount a is invested in process improvement in a period, then, after the period has elapsed, a *random* number of nominal improvements is achieved. We assume that only nonnegative amounts can be invested. The state of the process at the beginning of a period is given by the cumulative sum of the improvements achieved since time zero. Given a history of prior investments and improvements, we assume that the number of (additional) improvements achieved depends solely on the amount newly invested in process improvement. Let $\Phi(z,a)$ denote the probability that no more than z improvements are achieved in a period if the amount a is invested on process improvement in that period. It is convenient to work with the inverse distribution function $z(a,\omega)$, which gives, for each investment amount a and fractile ω of the distribution function, the number of improvements achieved. For example, if $\Phi(z,a)$ is continuous in z and a is invested, then the probability of achieving no more than $z(a,\omega)$ improvements is $\omega: \quad \Phi(z(a,\omega),a) = \omega$. In general, $\Phi(z(a,\omega),a) \geq \omega$ and $\Phi(z,a) < \omega$ for every $z < z(a,\omega)$. For each possible investment amount a, there is a random variable, W, uniformly distributed on $[0,1]$, such that if $W = \omega$, then the number of improvements achieved is $z(a,\omega)$.

Let i denote the fractional opportunity cost of capital per period, so that $\alpha = 1/(1+i)$ is the one period discount factor. The natural objective is to minimize the net present value of the operating and investment costs over an infinite horizon. However a formulation of the problem as a return maximization problem enhances tractability and interpretation. This formulation is in the spirit of the approach taken by Bell (1971), and discussed by Stidham (1972) and Harrison (1975), in the context of queueing control problems. Rather than accumulating the operating costs of the system over time, we compute the expected present value of operating the system in its current configuration over an infinite horizon. Whenever a change is made to the system, we compute the infinite horizon costs of operating under the new configuration. The reduction in these infinite horizon costs (the old value less the new one) is the savings that are associated with making the change. If that savings starts at the beginning of the next period, we discount it into present value terms. In particular, suppose that x improvements have been achieved to date, the amount a is invested immediately, and $z(a,\omega)$ improvements are achieved (by the end of the period), then the conditional expected present value of the future operating costs, over an infinite horizon starting at the beginning of the next period, will be $g(x + z(a,\omega))$. If no improvements were achieved, then the comparable cost would be $g(x)$. Thus, a savings of $g(x) - g(x + z(a,\omega))$ has been achieved, which must be discounted by α to put it in terms that are comparable to the investment a, which is being made now. After subtracting the cost of that current investment, we arrive at the following expression for the immediate return from taking action a:

$$r(x,a) := -a + \alpha \int_{\omega=0}^{1} [g(x) - g(x + z(a,\omega))]d\omega.$$

Thus, $r(x,a)$ is the infinite horizon net (of investment) expected present value of the *reduction* in operating costs, resulting from the random improvements achieved from the current investment. We interpret $r(x,a)$ as the immediate net return from making an investment of a when the state is x in a Markov decision process. Thus, savings achieved due to further investments in the future will be identified then and discounted back to the present.

We assume that investing nothing yields no improvement: $z(0,1) = 0$. Therefore, investing nothing yields no immediate return: $r(x, 0) = 0$ for all x. Our problem is equivalent to a stopping problem in that if it is optimal to invest zero in a given state, then the process will be in that state again at the beginning of the next period, and no change will ever take place thereafter. Thus, investing zero at some point in time can be thought of as deciding to stop investing (for all time) in process improvement.

2. Examples for Applications of the Model

2.1. Setup Reduction

Consider the inventory control problem of managing a single item over an infinite horizon. We use an EOQ model as extended by Porteus (1986) and Rosenblatt and Lee (1986) to include the effects of process quality. Thus, m denotes the (uniform) demand rate per period (a standard unit of time such as a year), K denotes the setup cost, h the sum of physical and financial holding costs per unit per period, q the probability that the production process will begin producing defectives (go out of control) with the current unit being produced (given the process was in control before the unit was started), and $c_{\mathbf{R}}$ the unit rework cost for defective units. Each production run is assumed to begin with the process in control. Each time an additional unit is processed, there is an independent chance that the process will go out of control and begin to produce defectives. There is no opportunity to monitor the process during the production run, so once the process goes out of control, the remaining units processed in the batch are defective. Then, by Porteus (1986), the operating cost per period (after optimizing over the lot size) is, assuming q is small, approximately

$$\sqrt{2mK(h + mc_{\mathbf{R}}q)}.$$

Consider the option of investing to reduce the setup cost. This option was also considered in Porteus (1985). The difference here is that the extent of the improvement that will be obtained with each investment is uncertain. We assume that at the beginning of each period, the current setup cost is known. The operating cost for that period is assumed, for convenience, to equal the optimal operating cost per period given above, corresponding to the current setup cost. (This assumption corresponds to assuming that a relatively large number of setups occur per period, so that the long run average cost can be used to represent the per period costs.) The decision is whether to invest in setup reduction, and, if so, how much. We assume that the probability distribution of possible improvements for each investment level is known and that at the end of each period, the actual improvements achieved are known.

Let K_0 denote the initial setup cost. Let β denote a fraction between zero and one that represents a nominal improvement in the setup cost. That is, if one (nominal) improvement is made, then the setup cost becomes βK_0. If x improvements are accumulated over time, then the setup cost becomes $\beta^x K_0$. This rescaling loses no generality because we can allow for fractional improvements to be made. Let $\Phi(z,a)$ denote the probability that no more than z nominal improvements will be achieved by the end of the period if the amount a (≥ 0) is invested in setup reduction. To fit this example into the framework presented in the previous section, let

$$g(x) := \sqrt{2m\beta^x K_0(h + mc_{\mathbf{R}}q)}/i, \tag{1}$$

which is the present value of the infinite horizon operating costs if x improvements have been achieved and no further investments are made. (To facilitate comparison in section 4 with the deterministic model of Porteus (1986), costs are assumed to be incurred at the end of each period.)

2.2. Process Quality Improvement

This example is similar to that of the previous section, except that we allow investment in process quality improvement, rather than in setup cost reduction. Let $\gamma \in (0,1)$ represent the nominal process quality improvement factor and q_0 the initial quality level. Thus, if x improvements have been accumulated, quality has improved to $\gamma^x q_0$ from q_0. For this problem,

$$g(x) := \sqrt{2mK(h + mc_{\mathbf{R}}\gamma^x q_0)}/i. \tag{2}$$

We consider the problem of simultaneously investing in setup reduction and quality improvement in section 4.

2.3. Other Examples

Many other examples, such as investing to reduce the unit cost c, or to increase the demand rate m, fit our model. Variations of the EOQ model used in the examples above can be used for these examples as well.

Other examples that require the use of different cost functions can be treated as well. For example, other process attributes that might be changed are: average yield, mean component leadtime, variance in component leadtime, and variance in setup time.

3. Analysis of the General Model

3.1. Assumptions

We assume that $g(x)$, the discounted present value of the operating costs over an infinite horizon if x improvements have been achieved and no further ones are, is positive $(g(x) \geq 0)$ and decreasing $(g(x) \leq g(z)$ if $x \geq z)$. (We say *strictly positive* and *strictly decreasing* when the inequalities are strict.) We also assume that g is convex. These and subsequent assumptions are satisfied by our examples and facilitate our obtaining interesting results.

We assume that investing more yields more improvement, in a first order stochastic dominance sense: $z(a, \omega)$ is increasing in a for every fixed ω. We assume that either all nonnegative investment amounts are admissible, in which case we also assume that $z(a, \omega)$ is continuous in a for each ω, or that the admissible investment amounts form a discrete set. To avoid more notation, we denote the admissible investment amounts by $a \geq 0$ in either case.

We assume, for convenience, that there exists a unique maximizer, $a(x)$, of $r(x, a)$, for each x. This assumption, which holds in the examples described above, allows the use of simpler expository language and stronger results. The policy that chooses to invest $a(x)$, the investment amount that maximizes solely the one period net return is called the *last chance* policy. Using conventional terminology, as in Ignall and Veinott (1969) and Heyman and Sobel (1984), for example, such a policy would be called *myopic*. However, we feel that the usual connotation of myopic being short sighted is misleading in our context. We would say that a policy is shortsighted if it invests too little in process improvement, because only the immediate effects of that policy are considered when making it. In our model, a last chance policy does indeed ignore the consequences of any future decisions that might be made. However, it completely accounts for the infinite horizon consequences of the current decision. Indeed, the last chance policy is just the opposite of a shortsighted policy: it invests *too much* in process improvements, even though only the immediately identifiable effects of that policy are taken into account. The optimal policy in the dynamic setting may be to make a large number of small investments in process improvements over time, each of which yields a high ratio of expected number of improvements to dollars invested. If only one last chance at investing is available, it will be optimal to invest at least as much as would be optimal in any single period. Details and further explanation come later.

We use primes to denote right hand derivatives. We denote right hand partial differentiation by the use of subscripts, which indicate the argument(s). For example, $r_1(x,a)$ denotes the right hand partial derivative of r with respect to its first argument, evaluated at the point (x,a).

3.2. General Results

As formulated, the problem of maximizing the net present value of the net cost reduction is a discounted infinite horizon, stationary Markov decision process with continuous state and action spaces. A policy specifies a decision for each state for a given period (stage). A strategy specifies a policy for each period. A stationary strategy specifies the same policy for every period. An optimal strategy achieves the maximum expected present value starting at any state and any period. The proof of the following result can be found the Appendix.

Lemma 1.
(a) $r(\cdot,a)$ *is decreasing for each fixed* a.

(b) The last chance policy $a(x)$ *is decreasing in* x *and approaches zero as* x *gets arbitrarily large.*

(c) The optimal return function f *exists, is positive, and is decreasing.*

(d) An optimal stationary strategy exists.

(e) f *is the unique solution to the Bellman optimality equations*

$$f(x) = \max_{a \geq 0} \left(r(x,a) + \alpha \int_{\omega=0}^{1} f(x + z(a,\omega))d\omega \right).$$

(f) The optimality condition holds: a strategy is optimal if and only if it achieves the maximization in the optimality equations at every state.

(g) If g *has a negative third derivative, then* f *is convex.*

(h) If $r(x,a) < 0$, *then no optimal policy invests* a *when the state is* x.

Part (a) says that as the number of cumulative improvements achieved over time increases, the immediate return from investing a fixed amount in process improvement decreases. Part (b) provides what will be an upper bound on the optimal investment amount. Part (c) confirms our intuition that the optimal savings from having the opportunity to invest in process improvements cannot be negative, and as the number of improvements accumulated increases, the value of having that opportunity decreases. Part (d) allows us to restrict consideration to stationary strategies without loss of generality. The optimal such policy is what then matters.

Theorem 1.

(a) Every optimal policy invests less than the last chance policy, at every state.

(b) Every optimal policy invests a strictly positive amount at a state if and only if the last chance policy does.

(c) There exists a target state x_T, *possibly infinite, such that both the optimal and last chance policies invest a strictly positive amount if the state is strictly less than* x_T, *and invest nothing if the state is strictly greater than* x_T.

(d) It is optimal to never stop investing in process improvement if and only if, for every $x \geq 0$, *there exists* $a > 0$ *such that* $r(x,a) > 0$.

(e) If there is only a finite number of admissible investment amounts, then the target state is finite.

(f) If there is only one admissible investment amount, $\bar{a}$, *other than no investment at all, then the last chance policy is optimal. (Invest if and only if* $r(x,\bar{a}) > 0$ *.)*

Proof.

(a) Fix x and consider investing $a > a(x)$. Then

$$r(x,a) + \alpha \int_{\omega=0}^{1} f(x + z(a,\omega))d\omega \le r(x,a) + \alpha \int_{\omega=0}^{1} f(x + z(a(x),\omega))d\omega$$
$$< r(x,a(x)) + \alpha \int_{\omega=0}^{1} f(x + z(a(x),\omega))d\omega.$$

The first inequality is a standard stochastic dominance result. It follows directly since f is decreasing and we assume that $z(\cdot,\omega)$ is increasing. The second follows from the definition of $a(x)$. Thus, by Lemma 1(f), investing such an amount cannot be optimal.

(b) Part (a) says that if the last chance policy invests nothing, then any optimal policy will also invest nothing. Suppose that the last chance policy invests a positive amount at a particular state x and that an optimal policy invests nothing there. Then, under the optimal policy, no further investments will ever be made, so the optimal return function is zero there. But if the last chance policy invests a strictly positive amount, then the immediate return from that amount is strictly positive, so investing nothing cannot be optimal, by Lemma 1(f).

(c) By Lemma 1(b), the last chance policy $a(x)$ decreases in x. If $a(x) > 0$ for every x, then let $x_T = \infty$. Otherwise, let x_T be the "smallest" point such that $a(x)$ equals zero: $x_T := \inf\{\, x \mid a(x) = 0 \,\}$. The result then follows from (b).

(d) Clearly, $a(x) > 0$ if and only if there exists $a > 0$ such that $r(x,a) > 0$. The result then follows from (c).

(e) Let $\bar{a}$ denote the smallest of the strictly positive admissible investment amounts. Let $\bar{x}$ be a state that satisfies $g(\bar{x}) < g(\infty) + \bar{a}$. Clearly $\bar{a} > 0$ and $\bar{x}$ exists and is finite. Suppose $x \ge \bar{x}$ and $a \ge \bar{a}$. Then $r(x,a) \le -a + \alpha[g(x) - g(\infty)] < 0$, so by Lemma 1(h), it is not optimal to invest any of the strictly positive admissible amounts when the state is greater than $\bar{x}$.

(f) The result follows directly from (b).

□

The last chance policy, which is relatively easy to determine, yields a great deal of information about the optimal policy, which may not be very easy to determine. If the last chance policy invests a strictly positive amount at a state, every optimal policy will do so also, and the optimal amount will be no more than the last chance amount. We have not shown that the optimal investment amount decreases as a function of the number of accumulated improvements. However, the optimal investment amount is bounded above by the last chance amount, which, by Lemma 1(b), is decreasing, and decreases to zero as the number of improvements gets arbitrarily large. Thus, the optimal investment amount also decreases to zero as the number of improvements gets arbitrarily large. However, it is possible that it is always optimal to invest something, albeit a very small amount. That occurs when the target state is infinite. In that case, it is optimal to pursue forever the goal of process improvement (zero cycle stocks in the setup reduction case and zero defects in the quality improvement case). Such an outcome is not possible if there is only a finite number of admissible investment amounts.

3.4. One Improvement at a Time

If, for each investment, the only possible outcomes are (1) a single nominal improvement with probability $\theta(a)$, and (2) no improvement otherwise, then the formula for the immediate return simplifies to:

$$r(x,a) := -a + \alpha\theta(a)[g(x) - g(x+1)].$$

Theorem 2. *Suppose that $\theta'(0)$ exists. An optimal policy will never stop investing in process improvement if and only if $\theta'(0) = \infty$.*

Proof. If $\theta'(0) = \infty$, then

$$r_2(x,a) = -1 + \alpha\theta'(a)[g(x) - g(x+1)].$$

Since g is strictly decreasing, $g(x) - g(x+1)$ is always strictly positive, so $r_2(x,0) = \infty$. Therefore, for every x, there must exist an $a > 0$ such that $r(x,a) > 0$. Thus, by Theorem 1(d), an optimal policy nevers stops investing in process improvement.

If it is optimal to never stop investing, then, by Theorem 1(d), for every x, there exists $a > 0$ such that $r(x,a) > 0$. Suppose, for the purpose of deriving a contradiction, that $\theta'(0)$ is finite, equal to, say, M. Given $\epsilon > 0$, there exists $\delta > 0$ such that $\theta(a)/a \le M + \epsilon$ for $0 < a < \delta$. Let $N := \max(M + \epsilon, 1/\delta)$. Since $\theta(a) \le 1$ for all a, it follows that $\theta(a)/a \le N$ for all $a > 0$. By picking a cutoff point $\bar{x}$ large enough, it follows that $r(x,a) < 0$ for $x \ge \bar{x}$ and all $a > 0$, which is a contradiction. Thus, $\theta'(0)$ must be infinite.

□

3.5. Increasing Risk

We now examine how the amount of uncertainty (risk in the stochastic improvement process) affects the optimal investment policy. Consider two environments, one of which has higher risk. Let $Z^1(a)$ and $Z^2(a)$ denote the respective (random) numbers of improvements achieved in the different environments if the amount a is invested. We use Rothschild and Stiglitz's (1970) *mean preserving spread* definition of increasing risk: The second environment has higher risk than the first if, for each a, there exists a nondegenerate random variable X such that $Z^2(a)$ and $Z^1(a) + X$ have the same probability distribution and $E(X \mid Z^1(a) = z) = 0$. for every z. That is, the second environment yields the same expected number of improvements as the first for the same level of investment, but suffers more variability.

Theorem 3. *The target state is a decreasing function of the risk.*

Proof. Using obvious notation, we have

$$\begin{aligned} r^2(x,a) &= -a + \alpha \int_{\omega=0}^{1} [g(x) - g(x + z^2(a,\omega))]d\omega \\ &\le -a + \alpha \int_{\omega=0}^{1} [g(x) - g(x + z^1(a,\omega))] = r^1(x,a)d\omega. \end{aligned}$$

Since $-g$ is concave, the inequality follows from Rothschild and Stiglitz's (1970) result that $EU(Z^2(a)) \le EU(Z^1(a))$ for every concave U is an *equivalent* formulation of the statement that the second environment is riskier than the first. Thus, if it is optimal for the last chance policy to invest a strictly positive amount in the second environment ($r^2(x,a) > 0$ for some a), then the same is true in the first environment: $r^1(x,a) \ge r^2(x,a) > 0$. Thus, by Theorem 1(b), the target state is smaller for the riskier environment. □

Theorem 3 says we are better off if we can make our process improvement process less risky. With less risk, the optimal investment policy yields a greater total payoff and invests in process improvements that would not be worth attempting in a riskier environment. The result also provides an easy way to get an upper bound on the target state: If it is not worth investing in process improvements in the risk-free environment, then it is not worth investing in them in a risky environment.

3.6. Discussion

The optimal policy will invest less than the last chance policy if it is economical to carry out a long series of small investments. This, in turn, depends on the relationship between investment amounts and the random number of improvements that will be achieved. Consider the return on investment ratio of expected number of improvements to investment amount. If that ratio decreases as the investment amount increases, then it may be economical to invest a small amount in each of many consecutive periods. For example, if there are many process engineers of differing abilities assigned to process improvement, then the better engineers are likely to achieve more improvements than the less talented ones. One interpretation of our model is that it determines how many process engineers (and other related indirect laborers) should be utilized.

When the return on investment ratio gets infinitely large for arbitrarily small investment amounts, it is optimal to invest a positive amount in every period, regardless of the number of improvements accumulated over time. This situation seems fairly unrealistic. Thus, our formulation suggests that an optimal investment policy will cease investment after some finite number of improvements has been accumulated. Despite this observation, we cannot claim that the JIT and TQC philosophies of continual pursuit of zero inventories and zero defects are nonoptimal in practice. Our model is not comprehensive and does not incorporate some realistic possibilities, such as the appearance of steadily more attractive investment opportunities over time. Experience in carrying out many process improvements might lead to such opportunities. Even if it is optimal to stop investing after a certain state has been reached, the optimal implementation plan may consist of stating that "we will never stop," to encourage the identification of as many good opportunities as is economical, and to announce the cessation of such investments only when the target state has been reached. Indeed, a manager of a Toyota plant in Japan responded to one of the authors that the current setup (changeover) time of one of their die presses was at such a low level that no active effort was being made to reduce it further, other than watching for new opportunities that might arise.

4. Simultaneous Setup Reduction and Process Quality Improvement

In this section we focus on the investment in simultaneous improvement of setup cost and process quality. When either attribute is analyzed individually, g is defined as in section 2, by (1) and (2), respectively. In either case, g is not only decreasing, positive, and convex, but also has a negative third derivative. Thus, all of the general results (Lemma 1 and Theorems 1, 2, and 3) apply to these examples when we restrict consideration to investing in one attribute at a time.

As explained in Schonberger (1982) and demonstrated in Porteus (1986) and Rosenblatt and Lee (1986), there can be a significant interaction between setup reduction and quality improvement. Let x and y denote the number of nominal improvements achieved to date in setup reduction and quality improvement, respectively, and let a and b denote the respective amounts invested in improving each attribute. Also let

$$g(x,y) := \sqrt{2m\beta^x K_0(h + mc_{\mathbf{R}} q_0 \gamma^y)}/i,$$

the discounted present value of the infinite horizon operating costs, given that x and y improvements have been achieved, respectively, in setup reduction and quality improvement, and no further investments are made.

We explore the two attribute problem through a numerical example. We begin our analysis with the restriction that only one improvement can be achieved in each attribute in each period. We assume that the investment effects and improvement returns for each attribute are independent of each other. The probability of achieving an improvement in setup reduction [process quality] given an investment of a [b] is $\theta(a)$ [$\phi(b)$]. With this assumption, the immediate net return is

$$\begin{aligned} r(x,y,a,b) = & -a - b + [g(x,y) - \theta(a)\phi(b)g(x+1,y+1) - \theta(a)[1-\phi(b)]g(x+1,y) \\ & - [1-\theta(a)]\phi(b)g(x,y+1) - [1-\theta(a)][1-\phi(b)]g(x,y)]. \end{aligned}$$

In addition the optimal value function satisfies

$$v(x,y) = \max_{a,b}\Big\{r(x,y,a,b) + \alpha[\theta(a)\phi(b)v(x+1,y+1) + \theta(a)[1-\phi(b)]v(x+1,y) + [1-\theta(a)]\phi(b)v(x,y+1) + [1-\theta(a)][1-\phi(b)]v(x,y)]\Big\}.$$

For our numerical examples, we choose parameters to facilitate comparison with the numerical example in Porteus (1986). Thus, for the quality-extended EOQ model, the demand rate per period is $m = 1000$, the initial setup cost is $K_0 = 100$, the rework cost per defective unit is $c_R = 25$, the holding cost per unit per period is $h = 10.5$, and the initial probability that the process goes out of control with manufacture of each additional unit is $q_0 = 0.0004$. For the investment return functions, we set $\theta(a)$ such that the expected number of improvements from investing amount a in setup cost reduction in our model is equal to the deterministic number of improvements that would be achieved in Porteus' (1986) model. In the latter, improving the setup cost to βK_0 from K_0, requires an investment of the fixed amount $-B\ln(\beta)$. Thus, to achieve an equivalent expected number of improvements, we set $\theta(a) = -a/B\ln(\beta)$. To assure that $\theta(a) \leq 1$, we require that $a \leq -B\ln(\beta)$. Similarly, for returns to investment in quality improvement, we have $\phi(b) = -b/\bar{B}\ln(\gamma)$, and we require that $b \leq -\bar{B}/\ln(\gamma)$. As in Porteus (1986), we assume that $B = 1898$, so that each ten percent reduction in the setup costs requires an investment of \$200, and that $\bar{B} = 190$, so that each ten percent improvement in process quality costs \$20.

With these data, Porteus (1986) found that the optimal investment policy for joint setup reduction and quality improvement reduces the setup cost from \$100 to \$18.20 and reduces the process failure probability from 400 ppm (parts per million) $(q = 0.0004)$ to 36 ppm $(q = 0.000036.)$ Because Porteus' model is static, the entire investment is made at one time and the benefits accrue immediately. For the example presented here, at most one nominal improvement is achieved per period, so that the amount of time to reach one of the target states depends on the values of β and γ, the improvement factors, and on the outcome of the stochastic process governing improvement returns. For the examples to follow, we let $\beta = .8$ and $\gamma = .75$.

The optimal policy for this formulation would follow a bang-bang solution because $\theta(a)$ and $\phi(b)$ (and therefore $r(x,y,a,b)$) are linear in the investment levels. In practice one might expect that diminishing returns to investment would be observed. Therefore, for our numerical example, we make the improvement probabilities concave by discretizing the act space and assuming there are maximum attainable one-step improvement probabilities $\bar{\theta}$ and $\bar{\phi}$. Thus, in the example that follows, we have

$$\theta(a) = \min\{-a/B\ln(\beta), \bar{\theta}\} \text{ and}$$
$$\phi(b) = \min\{-b/\bar{B}\ln(\gamma), \bar{\phi}\}$$

Let $\underline{a}$ and $\bar{a}$ be the adjacent feasible setup investment levels on either side of the break point: $\theta(\underline{a}) = -\underline{a}/B\ln(\beta)$ and $\theta(\bar{a}) = \bar{\theta}$. Then one of three acts will be optimal for each state $x : a(x) = 0, a(x) = \underline{a},$ or $a(x) = \bar{a}$. Similarly, we will get $b(y) = 0, \underline{b},$ or $\bar{b}$, where $\underline{b}$ is the largest feasible act such that $\phi(\underline{b}) < \bar{\phi}$, and $\bar{b}$ is the next larger feasible action. For our examples, we set $\bar{\theta} = \bar{\phi} = .8$.

We use a policy improvement algorithm to calculate optimal policies. Using the logic of Theorem 1, it is easy to generate upper bounds on the state and action spaces. They are bounded below by zero. For the action space, \$450 is an upper bound on per period investment in setup reduction and \$60 is a corresponding bound for quality improvement investment. Each of these action spaces is divided into eleven equally-spaced investment levels. So, for example, the feasible investment levels for setup reduction are $0, 45, 90, \ldots, 450$. For both setup cost and process quality the state space for the number of improvements achieved is the set of nonnegative integers in $[0, 10]$.

Exhibit 1 here

Exhibit 1 displays a stationary optimal policy. For this example, $\underline{a} = 315$, $\bar{a} = 360$, $\underline{b} = 42$, and $\bar{b} = 48$. The corresponding improvement probabilities are $\theta(\bar{a}) = .74$, $\theta(\bar{a}) = \bar{\theta} = .8$, $\phi(\underline{b}) = .77$ and $\phi(\bar{b}) = \bar{\phi} = .8$ Let the optimal investment levels be denoted by $a^*(x,y)$ and $b^*(x,y)$. An entry in the exhibit in cell (x,y) gives the values of $a^*(x,y)$ and $b^*(x,y)$. An entry of $\underline{a}\bar{b}$ in the $(0,1)$ cell, for example, means $a^*(0,1) = \underline{a}$ and $b^*(0,1) = \bar{b}$. In addition, an entry $\underline{b}$ means $a^*(x,y) = 0, b^*(x,y) = \underline{b},$ and an entry of "-" means $a^*(x,y) = b^*(x,y) = 0$. Whereas the diminishing returns cause this optimal policy never to use $\bar{a}$, the optimal quality investment level is $\bar{b}$ when both the quality level and the setup cost level are in low states, with much improvement to go.

States in the target region (where it is optimal to stop investing) can easily be observed from Exhibit 1. There are three such target states that can be reached starting from the point $(0,0)$: $(7,9)$, $(8,7)$, and $(8,8)$. The point $(7,9)$ is unlikely to occur, as it requires that nine quality improvements be achieved before getting the second setup reduction, even though investment is being made to pursue both goals. The point $(8,8)$ is next least likely to occur. It requires getting eight quality improvements before getting the sixth setup reduction. The remaining point $(8,7)$ is therefore most likely to occur. Since $\beta = 0.8$, and $\gamma = 0.75$, this point corresponds to a setup cost of 16.7 $(= 100 \times 0.8^8)$ and a process quality of 53 ppm. In Porteus (1986), in the simultaneous joint investment problem, the optimized setup cost is 18.2 and the optimized process quality is 36 ppm. Because of the discreteness of our example, we can only achieve setup costs of 21.0 and 16.8. When the setup cost is 21.0, it is still optimal to invest in further setup reduction. We must shoot past 18.2 to 16.8 before we stop investing. This result suggests that we are obtaining very close to the same target state information on setup costs in both the deterministic and stochastic models. However, on the quality side, we are likely to stop at 53 ppm. One more quality improvement would put us at 40 ppm, which is still higher than the optimal quality level in the deterministic model.

We now explore the sensitivity of the above solution to further increased risk in the stochastic improvement process. To do this, we must enrich the probability model somewhat. We want to find optimal policies for the improvement process where, for each investment amount, the expected number of improvements is constant, but the risk increases. With the two point probability distribution, we cannot change the risk without also changing the mean.

We therefore extend the model in the following way. Consider an investment of a in setup reduction. In our example above, we assumed that $\theta(a)$ was the probability that exactly one improvement would take place. We now interpret $\theta(a)$ as the probability that *some* improvement is possible. In particular, we assume that if some improvement is possible (but not guaranteed), then two improvements will occur with probability p_a, one will occur with probability $1-2p_a$, and none will occur with probability p_a. Thus, the unconditional expected number of improvements is still $\theta(a)$. If $p_a = 0$, then we have our earlier model. And if p_a is increased, in the range $(0, .5]$, the variance of the number of improvements increases. Similarly, for quality improvement, we assume that there is a parameter p_b such that, if some improvement is possible, which occurs with probability $\phi(b)$, there will be two improvements with probability p_b, one improvement with probability $1-2p_b$, and no improvement with probability p_b.

Exhibit 2 here

Exhibit 2 shows an optimal policy for the case of maximal variance: $p_a = p_b = 0.5$. Although this policy does not differ dramatically from that of Exhibit 1, the direction of change is clear. With more variable returns, it is optimal to stop investing earlier and to achieve less total improvement in setup costs and process quality. In light of Theorem 3, which predicts such a result for the single attribute case, this is no surprise.

We made two additional interesting comparisons with these examples. Despite the lower target states, the expected amount of time to reach a stopping state, starting from any (x, y), is larger in the riskier environment. For example, starting from $(0, 0)$, the expected time to stop is 9.7 periods for the low variance $(p_a = p_b = 0)$ case and 11.7 periods for the high variance case. In addition, the expected discounted investment required over the lifetime of the process is also larger in the riskier environment. For example, again starting from $(0, 0)$, the expected discounted investment required is \$2005 and \$2029, respectively, for the low and high variance cases.

These results illustrate how variability in the stochastic improvement process discourages firms from pursuing setup reduction and quality improvement. Reducing the variability of these processes would therefore seem to be a worthwhile pursuit for some firms. One might reduce the variability in the improvement process through, for example, better training of production workers and manufacturing engineers who work on process improvement, through thoughtful design of experiments in the improvement process, or through policies that assure well maintained and reliable manufacturing equipment.

5. Summary and Concluding Discussion

The model and analysis of this paper seek to understand better the economics of investing in gradual process improvement, which seems to be an important source of total productivity growth and technological progress. Toward this end, we formulated a dynamic, stochastic model of process improvement investment and obtained qualitative results on the optimal policies for our model. We obtained analytical results for the model of single attribute improvement and solved a two attribute example, involving setup reduction and process quality improvement. In our model, the policy of perpetual investment toward zero defects or zero cycle stocks is optimal only for a fairly narrow set of conditions. In general, there is a stopping state beyond which it is uneconomical to invest and that stopping state is likely to be finite in practical examples. The last chance policy for our model reveals this stopping state and bounds the optimal policy.

We presented a numerical example in which the likely optimal stopping state for investment in gradual improvement is close to but not as ambitious as the optimal outcome in Porteus' (1986) model of major technical leaps. The example also illustrated how increasing the uncertainty in the improvement process (in the sense of Rothschild and Stiglitz (1970)) impairs total improvement.

In addition to adding to the theoretical literature on investment in process improvement, our model is potentially useful to management practitioners. Managers of U.S. businesses sometimes direct attention to corporate requirements of "cost justifying" all investments in process improvement.

In order to use our model as a tool for justifying investment projects that yield sequences of small, stochastic improvements, one needs to estimate the parameters of the model. Estimating EOQ-like costs, although not trivial by any means, is discussed in most operations management textbooks. For our model, estimating the relationship between investment and number of improvements may be more difficult. One must estimate the probability distribution of the number of nominal improvements, for each possible level of investment. In practice, such estimates are, of necessity, highly subjective. However, in many cases, using a rough estimate of the probability distributions will be adequate.

In most companies, investment decisions for small gradual technological improvements are not made through the same capital budgeting procedures used for major technology acquisitions. Rather, investments in gradual process improvement consists of expenditures for process engineers, quality assurance personnel, and other indirect labor. Expenditures on such white collar labor is quite substantial in many firms, and we know of no other models that attempt to model optimal expenditures of this type. We hope that this work will stimulate others to study the optimal use of indirect labor for process improvement.

Acknowledgements

The authors wish to thank Arnold Barnett, Gabriel Bitran, Stephen Graves, and William Lovejoy for their helpful comments and suggestions.

Appendix. Proof of Lemma 1

Proof. Part (a) follows from the convexity of g.

(b) We give details for the differentiable case. By direct differentiation,

$$r_{12}(x,a) = -\int_{\omega=0}^{1} g''(x+z(a,\omega))z_1(a,\omega)d\omega,$$

which is negative: $g'' \geq 0$ because g is convex, and $z_1 \geq 0$ by assumption (stochastic dominance). Therefore $a(x)$ is decreasing. Let $g(\infty) := \lim_{x\to\infty} g(x)$, which exists. Thus, $\lim_{x\to\infty} r(x,a) = -a$. A simple argument by contradiction then shows that $\lim_{x\to\infty} a(x) = 0$.

(c), (d), and (e) We apply the results of Maitra (1968). It follows from the definition of $r(x,a)$ and since g is decreasing that $r(x,a) < 0$ for $a > g(0)$. We temporarily limit a to be no larger than $g(0)$. The remaining results of this lemma are contingent on this temporary limit. We first solve the problem conceptually with the limit imposed. We then see, by part (h) of this lemma, that if the limit is allowed to be violated, no greater return can be achieved. Thus, the limit can be made without loss of generality.

Thus, the action space is (effectively) compact. Since g is convex, it is continuous on its interior and upper semi-continuous (u.s.c.) on its domain $[0,\infty)$. We have assumed that the transition probability function is continuous on the action space. Thus, by Maitra (1968), the optimal return function exists and is u.s.c., there exists an optimal stationary strategy, and (e) holds. Here, $f(x) \geq r(x,0) = 0$, so f is positive.

Suppose that v is u.s.c., positive, and decreasing. Let A denote the optimal return operator, so that

$$Av(x) = \max_{a\geq 0}\left(r(x,a) + \alpha\int_{\omega=0}^{1} v(x+z(a,\omega))d\omega\right).$$

Suppose $x_1 \leq x_2$ and that action a_2 achieves the maximization above when the state is x_2. Then, using Lemma 1(a), $Av(x_1) \geq r(x_1,a_2) + \alpha\int_{\omega=0}^{1} v(x_1+z(a_2,\omega))d\omega \geq r(x_2,a_2) + \alpha\int_{\omega=0}^{1} v(x_2+z(a_2,\omega))d\omega = Av(x_2)$. Thus, Av is decreasing: decreasingness is preserved by the optimal return operator. Thus, by Porteus (1982), f is decreasing.

(f) This result follows from Porteus (1975, 1982). The sufficiency part, called the *optimality criterion* in Porteus (1975, 1982), can also be pieced together from the results in Blackwell (1965).

(g) Suppose that v is convex. Direct differentiation reveals that $r(\cdot,a)$ is convex for each fixed a. Then

$$Av(x) = \max_{a\geq 0}\left(r(x,a) + \alpha\int_{\omega=0}^{1} v(x+z(a,\omega))d\omega\right),$$

which is the pointwise maximum of a sum of convex functions, which is therefore also convex. Hence, convexity is preserved by the optimal return operator. Thus, by Porteus (1982), f is convex.

(h) Suppose $r(x,a) < 0$. Then, using the facts that $r(x,0) = 0$ and that f is decreasing, we have $Af(x) \geq r(x,0) + \alpha f(x) > r(x,a) + \alpha f(x) \geq r(x,a) + \alpha \int_{\omega=0}^{1} f(x + z(a,\omega))d\omega$. That is, action a does not achieve the maximization in the optimality equations when the state is x. Thus, by (f), action a cannot be optimal when the state is x.

□

	$y=0$	$y=1$	$y=2$	$y=3$	$y=4$	$y=5$	$y=6$	$y=7$	$y=8$	$y=9$	$y=10$
$x=0$	$\underline{a}$ $\bar{b}$	$\underline{a}$ $\bar{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$
$x=1$	$\underline{a}$ $\bar{b}$	$\underline{a}$ $\bar{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$
$x=2$	$\underline{a}$ $\bar{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=3$	$\underline{a}$ $\bar{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=4$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=5$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=6$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=7$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$ $\underline{b}$	$\underline{a}$	$\underline{a}$	-	-
$x=8$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	-	-	-	-
$x=9$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	-	-	-	-
$x=10$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	-	-	-	-	-

Exhibit 1 A Stationary Optimal Policy

	$y=0$	$y=1$	$y=2$	$y=3$	$y=4$	$y=5$	$y=6$	$y=7$	$y=8$	$y=9$	$y=10$
$x=0$	$\underline{a}\ \bar{b}$	$\underline{a}\ \bar{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}$	$\underline{a}$
$x=1$	$\underline{a}\ \bar{b}$	$\underline{a}\ \bar{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=2$	$\underline{a}\ \bar{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=3$	$\underline{a}\ \bar{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=4$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=5$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=6$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}$	$\underline{a}$	$\underline{a}$	$\underline{a}$
$x=7$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{a}\ \underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	-	-	-	-
$x=8$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	-	-	-	-
$x=9$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	-	-	-	-
$x=10$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	$\underline{b}$	-	-	-	-	-

Exhibit 2 A Stationary Optimal Policy for the High Variance Improvement Case

References

Bell, C. 1971. Characterization and Computation of Optimal Policies for Operating an M/G/1 Queuing System with Removable Server. *Operations Research* **19**, 208-218.

Blackwell, D. 1965. Discounted Dynamic Programming. *Annals of Mathematical Statistics* **36**, 226-235.

Fine, C. 1986. Quality Improvement and Learning in Productive Systems. *Management Science* **32**, 1301-15.

———1987. A Quality Control Model with Learning Effects. To appear in *Operations Research*.

Fine, C., and D. Bridge 1986. Managing Quality Improvement. in Sepehri, M. (ed), *Quest for Quality*, Industrial Engineering and Management Press, Atlanta.

Hall, R. 1983. *Zero Inventories*. Dow Jones-Irwin, Homewood, Ill..

Harrison, J. M. 1975. A Priority Queue with Discounted Linear Costs. *Operations Research* **23**, 260-269.

Hayes, R. 1981. Why Japanese Factories Work. *Harvard Business Review* **81**, No. 4, 56-66.

Heyman, D., and M. Sobel 1984. *Stochastic Models in Operations Research, Vol. II.* McGraw-Hill, New York.

Hollander, S. 1965. *The Source of Increased Efficiency: The Study of Dupont Rayon Plants.* Cambridge, Massachusetts.

Ignall, E., and A. Veinott, Jr. 1969. Optimality of Myopic Inventory Policies for Several Substitute Products. *Management Science* **15**, 284-304.

Maitra, A. 1968. Discounted Dynamic Programming on Compact Metric Spaces. *Sankhya* **30A**, 211-216.

Monahan, G. 1982. A Survey of Partially Observable Markov Decision Processes. *Management Science* **28**, 1-16.

Pierskalla, W., and J. Voelker 1976. A Survey of Maintenance Models. *Naval Research Logistics Quarterly* **23**, 353-388.

Porteus, E. 1975. On the Optimality of Structured Policies in Countable Stage Decision Processes. *Management Science* **22**, 148-157.

———1982. Conditions for Characterizing the Structure of Optimal Strategies in Infinite-Horizon Dynamic Programs. *Journal of Optimization Theory and Applications* **36**, 419-432.

———1985. Investing in Reduced Setups in the EOQ Model. *Management Science* **31**, 998-1010.

———1986. Optimal Lot Sizing, Process Quality Improvement and Setup Cost Reduction. *Operations Research* **34**, 137-44.

Rosenberg, N. 1982. *Inside the Black Box: Technology and Economics.* Cambridge University Press, New York.

Rosenblatt, M. and H. Lee 1986. Economic Production Cycles with Imperfect Production Processes. *IIE Transactions* **86**, 48-55.

Rothschild, M. and J. Stiglitz 1970. Increasing Risk: I. A Definition. *Journal of Economic Theory* **2**, 225-243.

Schonberger, R. 1982. *Japanese Manufacturing Techniques.* The Free Press, New York.

———1986. *World Class Manufacturing.* The Free Press, New York.

———1987. Frugal Manufacturing. *Harvard Business Review* **65**, No. 5, 95-100.

Stidham, S., Jr. 1972. $L = \lambda W$: A Discounted Analogue and a New Proof. *Operations Research* **20**, 1115-1126.

Tapiero, C. 1987. Production Learning and Quality Control. To appear in *IIE Transactions.*