A Budget-Sensitive Approach to Scheduling Maintenance in a Total Productive Maintenance (TPM) Program

Abhijit Gosavi, Missouri University of Science and Technology
Susan L. Murray, Missouri University of Science and Technology
V. Manojramam Tirumalasetty, Ungerboeck Systems International
Shreerang Shewade, Analytical Mechanics Associates

Abstract: Scheduling planned maintenance activities is key to the success of Total Productive Maintenance (TPM) in reducing the mean and variability of production lead time. Most existing maintenance-scheduling models are risk-neutral, striving to control long-run costs. Some are variance-penalizing, addressing both average cost and cost variance. Neither addresses budget constraints. We present two models that use semi-variance combined with the mean to simultaneously optimize maintenance with respect to long-run costs and short-term budgets. The first model, geared for individual pieces of equipment (e.g., a pump or dryer), uses renewal theory. The second model presented is based on Markov decision processes and is appropriate for manufacturing systems composed of several units, any one of which can fail. An application of each model is presented. Beneficial operational costs variance is not penalized in this approach, which is more appropriate.

Keywords: Preventative Maintenance Scheduling, Semi-variance, Budget-based Scheduling

EMJ Focus Areas: Quantitative Methods and Models

Total Productive Maintenance (TPM) is a program that began in the 1970s in Japan. Seiichi Nakajima popularized TPM throughout Japan and is widely recognized as a pioneering practitioner of this field. One of the main goals of TPM is to maximize equipment effectiveness and availability. In the 1980s, TPM started gaining popularity in the United States. The initial interest may have been due to competition from Japan and the need for cutting costs. It quickly became clear that a well-designed TPM program would lead to less lead time variability, which in turn, significantly lessened the need to carry finished goods inventory. Since the late 1990s, many manufacturers have transitioned from make-to-stock (push) to either make-to-order (pull) or delayed differentiation strategies. To be competitive, make-to-order requires significantly reduced lead time and is less tolerant of unexpected machine failures disrupting production. Consequently, the need for TPM has increased in recent years. TPM is no longer viewed just as a continuous improvement tool to cut costs but is also considered a critical tool in keeping lead time in check, which is absolutely essential for survival in a world where customers demand low prices and rapid delivery.

TPM is oftentimes implemented in multiple phases, the first phase being linked to the philosophy of being pro-active in order to improve equipment performance. A modification of a popular slogan in industry captures the TPM philosophy: "If it ain't broke, fix it anyway!" TPM's overall goal is to prevent the unexpected failure that can disrupt a production schedule. Unexpected failures can reduce throughput—especially if the failure affects the process bottleneck. Furthermore, when a failure occurs, it usually takes longer to correct than a scheduled maintenance activity would, resulting in significantly higher costs. It has been empirically shown that preventive maintenance can reduce the frequency of unexpected failures and, if done at appropriate time intervals, can reduce the overall costs (Askin and Goldberg, 2002). TPM is now viewed as an integral part of regular operations in production firms. With the advent of computers and the cheap availability of personal computers in the last couple of decades, computerized maintenance systems have become increasingly popular in industry (Westerkamp, 2006). Such systems make it easy to collect and maintain historical data of machine failures, their frequencies, and down-times due to repairs and maintenance. These databases can be used to determine parameters (such as distribution type, mean, etc.) of system failures, providing an excellent basis to model and improve the maintenance process. It is no exaggeration to state that production systems cannot remain healthy and productive without a good TPM program; however, developing effective operational strategies for TPM can be quite challenging because of numerous complicating factors, such as random failures of the different machines and pieces of equipment in a system, randomness in repair/maintenance times due to variability in the availability of spare parts and repairpersons, and the complex stochastic dynamics of production systems. The manager has to analyze the underlying stochastic processes, costs, and revenues, and a host of other factors in order to develop an effective TPM program.

Problem Statement

Production managers must balance the need to reduce lost production costs due to equipment failures with preventive maintenance costs. Managers do this by developing preventative maintenance schedules. Most statistical models in reliability and preventive maintenance (PM) attempt to optimize the system over the long-run (i.e., an infinite time horizon) to minimize long-run and average costs. While this leads to good long-run performance, in a finite time horizon such a scheduling policy can exceed short-run maintenance department budgets. Our goal in this research is to develop maintenance schedules that minimize the long-run average costs over the infinite time horizon, but at the same time minimize the chances of costs exceeding a predetermined daily or weekly budget. To accomplish this we employ a relatively less known measure of risk called semi-variance that measures variability of costs rising above a predetermined threshold (target or ceiling). The advantage of this metric is that when suitably combined with the long-run average cost metric, it can produce solutions that seek to keep the long-run average cost under

Refereed Research Manuscript. Accepted by Editor Doolen.

control and at the same time minimize the chances of costs rising above the threshold.

We present two statistical models, both of which are based on well-known stochastic models. The first is based on renewal theory (see Askin and Goldberg, 2002, for applications in a lean manufacturing setting); the second is based on semi-Markov decision processes (SMDPs) theory (Howard, 1971). The SMDP model is a generalized version of the more well-known MDP (Markov decision process) model. The theory of renewal processes captures the dynamics of cyclical phenomena that occur in stochastic systems. The renewal reward theorem (Ross, 1997) is a key tool that supplies the much needed cost-and-reward structure to the renewal process and thereby provides a mechanism to measure the performance of systems that can be modeled with renewal processes. The SMDP model is based on construction of a Markov chain, which assumes that the system transitions randomly from one state to another. Embedded within the Markov chain is a reward and cost structure that can be used to optimize the system.

The renewal-theoretic model that we develop can be applied with less effort than the SMDP model. The SMDP model is more applicable to more complicated systems consisting of several units that can fail. Renewal-theoretic models in preventive maintenance have been widely used in industry (Askin and Goldberg, 2002) because they require less work for the analyst and are very transparent. In the context of this article, the renewal-theoretic model is likely to be more useful for components of manufacturing systems, such as conveyor systems or heat-treatment units. The MDP model is applicable for a larger system such as a production line consisting of numerous machines or a fleet of automated guided vehicles (AGVs). The ground work in the SMDP model consists of identifying system states and calculating transition probabilities. While the SMDP model is more sophisticated, it has seen fewer applications in the real world—probably because of the additional modeling effort it demands.

After developing the models, we tested them numerically. For the renewal-theoretic model, we present numerical results gathered from a New York manufacturer. The data have been modified for proprietary reasons. For the SMDP model, the data used are from the literature. Our numerical results show that our models can produce maintenance schedules that with a slight increase in the average long-run cost minimize the variability of costs rising above a predetermined budget. The computational formulas are simple enough to be implemented using spreadsheet software, such as Microsoft Excel.

Literature Review

The literature on TPM is vast, and it has been adequately surveyed in review articles (McCall, 1965; Valdez-Flores and Feldman, 1989; McKone and Weiss, 1998; Wang, 2002; Ahuja and Khamba, 2006), which have appeared with an increasing frequency. This regularity of appearance is indicative of TPM's widespread use and importance in industry. The focus of this article is on those aspects of TPM that are related to determination of the interval for preventive maintenance. We now describe some of the related literature, linking the different streams of research that have been performed.

Early Days of PM

In the early days of the industrial revolution, maintenance was synonymous with repair. In other words, resources were allocated to repairing machines *after* they failed rather than doing anything to minimize the chances of failure. The science of statistics offered a path for better understanding failure mechanisms, which led to the birth of a field now called reliability-centered maintenance (RCM). RCM was invented formally at Boeing in the early 1970s. Barlow and Proschan (1965) published the first textbook to have addressed this issue as a statistical science. It gradually became clear that one could reduce the total costs of repair and maintenance if one were to proactively maintain machines, e.g., lubricate bearings that can cause failure or do a complete overhaul by shutting down a machine that is functioning properly for maintenance after an appropriate time interval. This led to the modern era of TPM. The term TPM was popularized first at Toyota Motor Company.

Stochastic Processes

A great majority of the papers that study the statistical aspects of the failure mechanism exploit stochastic processes, either renewal processes or some variants of the Markov decision process (MDP). Some of the earlier models for RCM were based on the assumption of "shocks" in the form of failures that lead to Markov chains and were aimed at minimizing the total discounted costs over an infinite time horizon (Chitke and Deshmukh, 1981; Anderson, 1981). Dada and Marcellus (1994) modeled the problem as an MDP and used notions of out-of-control and in-control processes. Such decision-theoretic models have since been used widely in modeling the RCM problem—see Das and Sarkar (1999) and references therein. For a textbook description of the use of renewal theory for maintenance, which is also called *age replacement*, see Gertsbakh (2000), which provides a comprehensive account of the related literature.

Industrial Implementation

Repair and maintenance costs can account for about 15-40% of production costs in industries (McKone and Weiss, 1998). As a result, one finds a number of case studies of TPM applications in the literature (Shimburn, 1995; Wilmeth and Usrey, 2000). A detailed analysis of how widespread TPM is within industry can be found in McKone et al., (1999); also see Nayak and Shayan (1998).

Systems Viewpoint

A branch of the related literature seeks to integrate TPM with other functions in production management. Some of this work includes integrating quality (Panagiotidou and Tagaras, 2007), lead time constraints (Sheu and Chien, 2004), and job scheduling constraints within PM scheduling (Cassady and Kutanoglu, 2005). This has lead to some new advances in PM scheduling that adopt a systems viewpoint that seeks to minimize PM costs and at the same time minimize other costs incurred in the production system.

Recent Work

Much of the recent work has attempted to overcome some of the classical obstacles to implementing TPM in industry. Some of these obstacles have come in the form of the assumptions made in many TPM models: (i) the machine is as good as new after repair or maintenance, (ii) the machine operates as a solitary unit in the production system (whereas in most systems machines are a part of a production line and their operations are inter-dependent), and (iii) operators are available all the time for maintenance. We now cite some papers that seek to address these issues. El-Ferik and Ben-Daya (2006) study a scenario where the machine ages after every repair or maintenance. Quan et al. (2007) sought to model operator (workforce) availability constraints in an aircraft maintenance facility within a TPM schedule, while Dellagi et al. (2007) studied TPM of two interconnected machines, whose operations are dependent on each other. TPM is a topic that continues to attract research interest.

Risk-Averse Preventive Maintenance

All of the models described above attempt to minimize the long-run average costs with no consideration of the variability or "risks" involved in these policies. Such models will be referred to as risk-neutral models. The first paper that considers risk in the literature is Chen and Jin (2003), where they use variance as a measure of risk. They measure risk in a renewal process but set the time of each transition to one. Subsequently, Gosavi (2006) pursued a model that employed *cyclical* variance using renewal processes and one-step variance in the MDP model. Both papers seek to combine, within the performance measure, the expected costs and the variability of the costs. Here, our goal is to use a different measure of risk, namely *target semi-variance* (or semi-variance for short), within the RCM framework both for renewal processes and SMDPs. Target semi-variance has a notable advantage over variance in that it accounts for variability in costs incurred ***above*** the budget set by senior management. Using variance in a model, on the other hand, considers situations where the costs ***below and above*** the mean are of the same importance. In application, a manager would be pleased with costs below the department's budget and concerned with those exceeding the budget. Penalizing variance leads to penalizing variability that is actually welcome. Using semi-variance to penalize variance, on the other hand, leads to budget-sensitive behavior that more closely represents the concerns of engineering managers.

Porter (1974) provides an early definition of semi-variance. Target semi-variance, or simply semi-variance, is now used in financial portfolio analysis as an alternative to variance in measuring risks. The mathematical development of MDP theory using semi-variance is the subject of a companion paper (Gosavi, 2010) that is likely to be of interest to theoreticians. In this article, we focus on the application and the implementation aspects that are of more interest to the practitioner.

Models for Budget-Sensitive PM Scheduling

As we present the mathematical details of the two models, some definitions and assumptions are necessary. *Time for failure* is the amount of time it takes a new machine or a machine that is newly repaired or maintained to fail. We assume that the machine is as good as new after every maintenance or repair. This assumption implies that after every maintenance or repair, the time for failure of the machine has the same distribution. This is a standard assumption in the literature (Gertsbakh, 2000) and representative of many industrial situations. It is assumed that the modeler has sufficient knowledge concerning the equipment involved to develop appropriate failure distribution(s); often these distributions can be determined from historical machine maintenance data. It is also assumed that the costs of repair and maintenance are known.

Renewal-Theoretic Model

We first describe the physics of the renewal process and then derive a formula that measures the semi-variance of the cost per unit time. The "age" of a machine is assumed to be the time elapsed since it was repaired or maintained. As stated above, we assume that the unit is as good as new after repair or maintenance. That is represented by the machine's age being set to zero after repair or maintenance. The unit ages only when it is in operation. The following notation is used in the model:

X: the time for failure of the system, which is a random variable,
T: the unit's age when PM is performed
$F(\cdot)$: the cumulative distribution function (*cdf*) of X
$f(\cdot)$: the probability density function (*pdf*) of X
C_r: the expected cost of one repair
C_m: the expected cost of one PM activity
t_r: the expected time required to perform one repair
t_m: the expected time required to perform one PM activity
τ: the budget *rate* for TPM
$E(R)$: the expected cost in one renewal cycle
$E(L)$: the expected cycle time in the renewal cycle

In renewal theory, one uses the notion of cycles to describe stochastic events. Here we will assume that a new cycle starts when production begins on a new machine, or a machine that has just been repaired or maintained. One of the following two events end the cycle: (i) the machine fails and a repair is performed, and (ii) the machine is preventively maintained before it fails. When a cycle ends and production starts, a new cycle begins. In renewal theory, our attempt is to measure the mean net cost incurred in the cycle, denoted by $E(R)$, and the mean time spent in one cycle, $E(L)$. According to the renewal reward theorem (Ross, 1997), the mean (expected) cost rate (in \$/hr), denoted by ρ, is given by

$$E(Cost\ Rate) = \rho = \frac{E(R)}{E(L)} \ . \qquad (1)$$

Minimizing the mean cost rate leads to a PM policy that minimizes the average long-run cost, but may lead to a situation where the cost in unit time exceeds a budget (ceiling or target) τ. Usually, the budget is available for the year or month. From this one can compute the hourly budget. For instance, if B denotes the annual budget for TPM, and if H denotes the number of hours for which machines are used each year, then

$$\tau = \frac{B}{H} \ . \qquad (2)$$

The goal is to derive a performance measure that attempts to keep the cost rate under control but at the same time also minimizes the chances of the hourly (or weekly or monthly) costs exceeding τ. Semi-variance is used to achieve this goal as opposed to variance, which does not account for performance with respect to budgets. Variance measures the variability above and below the mean. Semi-variance measures variability above a given budget (threshold) rate, τ, for costs. The multi-objective metric in risk-sensitive formulations that seek to simultaneously minimize the average cost and the variability can be defined as follows:

$$\text{Minimize } E\ (Cost\ Rate) + \theta\ Risk\ Rate \qquad (3)$$

where the term *Risk Rate* could be a function of variance or semi-variance depending on how the manager wants to measure risk (depending on the risk metric preferred) and θ denotes the risk-sensitivity factor, which is indicative of how risk-averse the manager is. Typically, θ in Equation 3 assumes small values in the range of (0,0.5). The greater the value of this factor, the higher the risk-averseness; however, very high values of θ (i.e., above 0.5) will lead to a solution in which the risk becomes more important than the expected cost, and can produce solutions which are very expensive but have low variability. A value of 0 for θ implies that the manager is risk-neutral. This kind of a formulation was invented by Markowitz (1952) and is now widely used in portfolio optimization. The *Risk Rate* in Equation 3 is defined as follows:

$$Risk\ Rate = \frac{Risk}{E(L)} \ . \tag{4}$$

The precise definition of risk in Equation 4 will depend on the metric chosen, and details will be provided in the model development. Using the standard mechanism for building the model for machine maintenance (e.g., Askin and Goldberg, 2002), results in the following expression for the expected cost:

$$E(R)= C_r\, P(X<T) + C_m\, P(X\geq T) = C_r\, F(T) + C_m(1\text{-}F(T)). \tag{5}$$

Equation 5 can be explained as follows: X (a random variable) denotes the time at which failure occurs and T (a deterministic variable) denotes the age at which the machine is maintained. Hence the cycle ends either with preventative maintenance being performed (when $X>T$), which costs C_m, or failure (when $X<T$), which costs C_r. Hence the expected cost of any cycle is the cost of a failure (C_r) times the probability of failure plus the cost of preventative maintenance times the probability of preventative maintenance being performed in that period. Similarly, the mean cycle time in the renewal process (see Ross, 1997) can be estimated as follows:

$$E(L)= \int_0^T (x+t_r)f(x)dx+(T+t_m)(1-F(T)). \tag{6}$$

To define the terms for semi-variance, we first introduce the following notation: $[a]_+ = \max(0,a)$. Then, in a manner analogous to the formula in Equation 5, the semi-variance of R can be defined as:

$$Svar(R) = \int_0^T [C_r-\tau(x+t_r)]_+^2 f(x)dx + [C_m- \tau\,(x + t_m)]_+^2 (1 - F(T)). \tag{7}$$

Following Equations (3), (1), and (4); the budget-sensitive performance metric will be defined as:

$$g\,(T, \theta) = \frac{E(R)}{E(L)} + \theta\ \frac{Svar(R,L)}{E(L)} \tag{8}$$

where the quantities *E(R)*, *E(L)*, and *Svar(R,L)* are defined in Equations 5, 6, and 7, respectively. The notation $g(T,\theta)$ will be used to denote that this is a function of *T,* the age at which the machine is maintained, and θ, the risk-averseness factor; note that T and θ are inputs for Equation 8.

The intent here is to compare the performance of the semi-variance model with the variance models in the literature. When cyclical variance in a renewal process is used, one has that

$$Var_1(R) = \int_0^T [C_r- E(R)]^2 f(x)dx + [C_m- E(R)]^2 (1 - F(T)). \tag{9}$$

Then, using Equations 3, 1, and 4, the variance-penalized performance metric will be defined as:

$$g\,(T, \theta) = \frac{E(R)}{E(L)} + \theta\ \frac{Var_1(R)}{E(L)} \tag{10}$$

The cyclical variance $Var_1(R)$ can be replaced by the more general *asymptotic* variance. To use asymptotic variance (Brown and Solomon, 1975; Gosavi, 2008), instead of $Var_1(R)$ in Equation 10, we use $Var_2(R,L)$, which is defined as follows:

$$Var_2(R,L) = E(R^2) - 2\rho^2(E(L))^2 + (\rho^2E(L^2). \tag{11}$$

in which ρ is as defined as in Equation 1 and *E(L)* is defined as in Equation 6,

$$E(L^2) = \int_0^T (x + t_r)^2 f(x)dx + (T + t_m)^2\,(1 - F(T)), \tag{12}$$

and

$$E(R^2) = C_r^{\,2}F(T) + C_m^{\,2}\,(1 - F(T)). \tag{13}$$

Equation 10 can be used again as the variance-penalized objective function after replacing $Var_1(R)$ by $Var_2(R,L)$.

The formulas presented above will be used to optimize the system with respect to Equation 8 in order to determine the budget-sensitive optimal time, T^*, for preventative maintenance. This time should be compared to that obtained from a risk-neutral optimization—again via Equation 8 but with θ = 0. The system will also be optimized with respect to Equation (10). The variance-penalized metric in Equation 10 is already available in the literature (Chen and Jin, 2003; Gosavi, 2006) for deriving variability-penalized policies demonstrating its usefulness as a budget-sensitive metric.

SMDP Model

The SMDP model seeks to capture the behavior of a more complex system with several components that can fail. A Markov chain is constructed to model the system dynamics. A cost structure is then added to the Markov chain. Like a discrete-event simulator of the system, the Markov chain can be used to predict the values of system performance measures. One can easily combine the Markov chain to an optimizer to determine the optimal values of the decision variable(s). The process of optimization can be performed either with linear programming or dynamic programming. For additional details on Markov chains and MDPs, see Bertsekas (2007).

In order to define a problem as an MDP or SMDP, one must define the system states and determine the optimal action to be chosen in each state. For the TPM problem, the states are defined by the age of the system. For this model, the age is assumed to be a discrete variable rather than a continuous variable as was used in the Renewal-Theoretic Model. When the machine is maintained or repaired, the age will be automatically set to be zero. Further it is assumed that age will increase in discrete steps and the unit of days is used. The action in each state will either be *produce* (do no maintenance) for one more cycle or *maintain* the system. The optimal solution for this problem is then defined by the optimal action in each state. Thus for instance, if one possible solution is to produce until the age is five days, the implication is that one must produce at all ages of the system until the system's age becomes five or more at which time the machine should be maintained. Performance metrics (in terms of dollars and hours) are associated with each possible solution. The goal of this approach is to determine the optimal solution with respect to the performance metrics.

Decisions are made at the end of a production cycle, after repairs, and after preventative maintenance. When the decision to produce is made, the system either completes the production successfully and its age is increased by the duration of the production cycle, or it may fail during the production. If the system fails, it is repaired, which takes a random amount of time, at the end of which the system is assumed to be as good as new and the age is reset to zero. The randomness of the repair time stems from the fact that every failure is unique, and significant fluctuations in the time required to determine the problem and fix it are common. When the decision to perform preventive maintenance is made, the maintenance work also takes a random amount of time; however, the mean and variability for preventative maintenance times tends to be less than that of repair time. At the end of the PM, the system's age returns to zero. Generally, the probability of failure starts increasing with the age. This is called increasing failure rate in reliability. The notation for this model follows:

- S: finite set of states
- $A(i)$: set of action allowed in state i
- $p(i, a, j)$: transition probability of going from state i to j under action a
- $t(i, a, j)$: time of transition in going from i to j under action a
- $c(i, a, j)$: cost of transition from i to j under action a
- $\bar{c}(i, a) = \sum_{j\in s} p(i, a, j)c(i, a, j)$: the expected cost of transitioning from i under action a
- $v(i, a, j) = c(i, a, j) - \tau \times t(i, a, j)]^2_+$: the budget-sensitive semi-variance in cost incurred in transitioning from i to j under action a
- $\bar{v}(i, a) = \sum_{j\in s} p(i, a, j)v(i, a, j)$: the expected budget-sensitive semi-variance in transitioning from i under action a
- d: policy or solution of the decision process, where $d(i)$ will denote the action prescribed by d in state i
- The purpose of the SMDP model is to

$$\text{Minimize:}\quad \frac{\sum_{j\in S}\pi_d(i)\bar{c}(i, d(i))}{\sum_{j\in S}\pi_d(i)\bar{t}(i, d(i))} + \theta\,\frac{\sum_{j\in S}\pi_d(i)\bar{v}(i, d(i))}{\sum_{j\in S}\pi_d(i)\bar{t}(i, d(i))} \qquad (14)$$

where $\pi_d(i)$ denotes the steady-state probability of being in state i when policy d is used. The term in both denominators represents the expected time per transition of the system when policy d is pursued. The numerator $\sum_{j\in S}\pi_d(i)\bar{c}(i, d(i))$ in the first term denotes the expected cost per transition when policy d is pursued while the numerator $\sum_{j\in S}\pi_d(i)\bar{v}(i, d(i))$ in the second term denotes the expected semi-variance per transition when policy d is pursued. The goal is to determine the policy that minimizes the function in Equation 14. In order to use this to solve the problem, one would have to exhaustively determine the steady-state probabilities of each policy. Fortunately, it is not necessary to search over all solutions to solve this problem; the optimal solution can be obtained by solving a linear program (LP) as shown in Tijms (2003) for the SMDP (the original formulation for the MDP is from D'Epeneoux [1960] and Manne [1960]). The LP is formulated as follows:

$$\text{Minimize}\ \sum_{i\in S}\sum_{a\in A(i)}\bar{c}(i, a)x(i, a) + \theta\sum_{i\in S}\sum_{a\in A(i)}\bar{v}(i, a)x(i, a) \qquad (15)$$

such that

$$\sum_{a\in A(j)}x(j, a) - \sum_{i\in S}\sum_{a\in A(i)}p(i, a, j)x(i, a) = 0 \quad \text{for all } j\in S \qquad (16)$$

$$\sum_{i\in S}\sum_{a\in A(i)}x(i, a)\bar{t}(i, a) = 1 \quad \text{and } x(i, a) \geq 0 \quad \text{for all } i\in S \text{ and } a \in A(i) \qquad (17)$$

The optimal policy can be obtained as follows from the solution of the above LP as follows:

$$d^*(i, a) = \frac{x^*(i, a)}{\sum_{a\in A(i)}x^*(i, a)} \quad \text{for all } i\in S \text{ and } a\in A(i) \qquad (18)$$

where $x^*(i,a)$ denotes the optimal solution of the LP above. It can be shown (Tijms, 2003) that $d^*(i,a)$ will return a 0 or a 1, which can be interpreted as: if $d^*(i,a) = 1$, action a is optimal for state i, and if $d^*(i,a) = 0$, it is not optimal for i.

In order to benchmark the performance of the budget-sensitive algorithm, two other approaches are used for comparison: (i) a risk-neutral approach, which follows the above with $\theta = 0$ and (ii) the variance-penalized approach. For the variance-penalized approach, one needs to solve the following quadratic program (Filar et al., 1989). The quadratic program is as follows:

$$\text{Minimize}\ \sum_{i\in S}\sum_{a\in A(i)}[\bar{c}(i, a) + \theta\,\bar{c}^2(i, a)]x(i, a) - \theta\sum_{i\in S}\sum_{a\in A(i)}\bar{c}(i, a)x(i, a) \qquad (19)$$

subject to constraints (16) and (17) above.

Application of the Budget-Sensitive Models

In order to test the renewal theory model, data was gathered from a New York manufacturer of alternators and starter motors. The data came from an alternator series that had the greatest contribution to the company in terms of sales and warranty returns. It was determined by the company that a significant reduction in the warranty return cost could be realized with PM of the alternator production line. The data presented in this article were modified to protect the manufacturer's privacy. The nature of original data was retained. The data were collected over five years. The data for the second Markov-chain based model are a modification of the industrial data used in Gosavi (2006). The original data were collected from a different manufacturer in New York and have been significantly modified to protect the interests of the firm.

Renewal-Theoretic Model

Exhibit 1 shows that the budget-sensitive objective function (see Equation 8) is a convex function with a unique minimum for one particular case. This indicates that this problem is well posed and that by selecting the optimal time of maintenance, the manager can expect to achieve the best behavior in terms of the budget-sensitive performance metric. We would also like to note that both forms of variance, cyclical (defined in Equation 9) and asymptotic (defined in Equation 11), produced the same solution numerically in our experiments. Hence, in what follows, we present results from only one variance model.

The input parameters for the systems tested are enumerated in Exhibit 2, which defines the nine cases that we have studied. Cases 1 and 2 are based closely on real-world data obtained from a New York manufacturer. In order to generate additional data to test our models, seven other cases were developed by modification of one or more parameters in Cases 1 and 2. C_r is the cost to repair the equipment once it has broken, and it is set to \$33 or \$83 for the various cases. By comparison the cost of preventative maintenance, C_m, is significantly less at either \$2 or \$5. If the costs were comparable, then there would be no reason to implement TPM. Typically a repair is going to be unscheduled, is more difficult, and requires more time than a preventative maintenance activity. The time difference is reflected in the differences between t_r and t_m. The budget, τ, was varied over a range of values to test

Exhibit 1. Objective Function vs. Time Since Last Repair or Maintenance

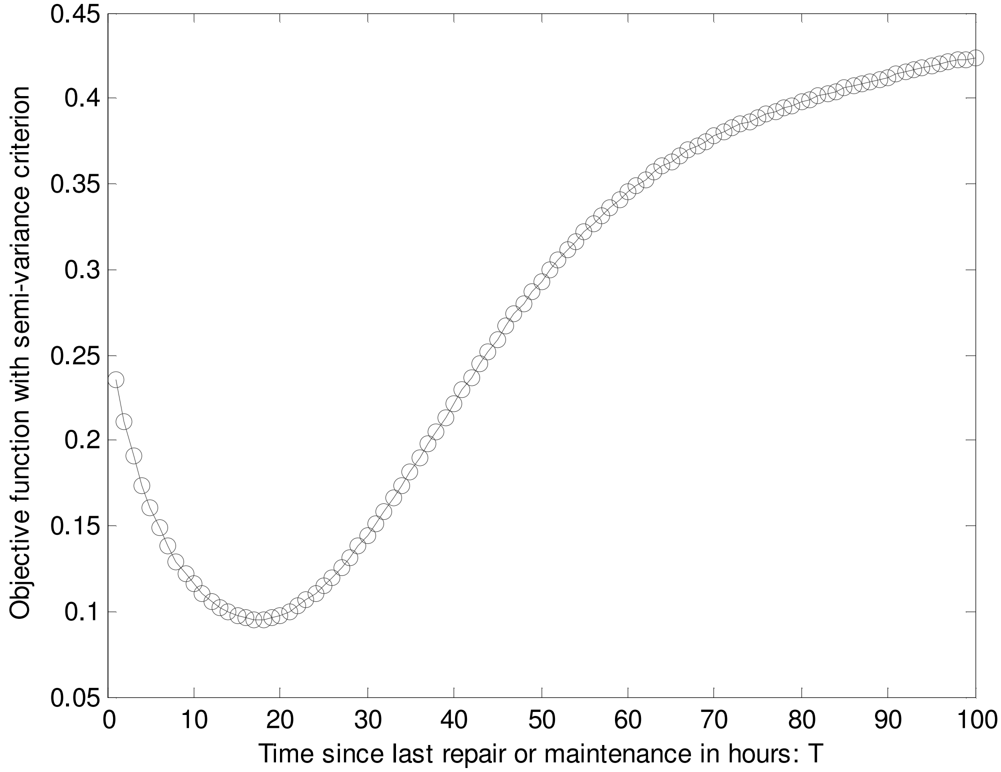

Exhibit 2. Input Parameters for the Renewal-Theoretic Model

Case	Gamma (r, λ)	C_r ($)	C_m ($)	τ ($/hrs)	t_r (hrs)	t_m (hrs)	θ
1	(6,12.5)	33	2	0.3	25	7.5	0.2
2	(8,12.5)	83	2	0.45	50	15	0.2
3	(4,12.5)	83	5	1.8	25	7.5	0.3
4	(12,8.33)	83	5	0.7	50	15	0.3
5	(6,12.5)	33	2	0.5	25	7.5	0.3
6	(9,10)	33	2	0.16	50	15	0.3
7	(10,11.11)	83	5	0.7	25	7.5	0.2
8	(11,6.67)	83	5	0.65	50	15	0.2
9	(10,10)	33	2	0.3	25	7.5	0.3

the performance of this model. The risk sensitivity factor was varied between the values of 0.2 and 0.3.

The time between successive failures is assumed to have a gamma distribution, whose parameters are (r,λ). The means the random variable is $r\lambda$. The gamma distribution ensures that the system has an increasing failure rate. Recall that Equations 8 and 10 are functions of T, the age of the machine since repair or maintenance. The goal is to determine the optimal value of T, the time when the manager should shut down the machine for preventive maintenance. The calculations in Equations 8 and 10 require only simple integration, which can be done with Microsoft Excel or MATLAB. The equations have to be evaluated over a range of T values from 0 to a number K, where K is taken to be five times the mean length of the machine's life; the reason for choosing K in this style is that for the gamma distribution, at five times the mean machine life, failure is almost certain. A finite set of T values is selected in the range from 0 to K; the values being $(0, 0.1, 0.2, 0.3, \ldots, K)$. In other words, the function in Equation (8) is evaluated over possible values of T. The optimal value of T is determined for the variance-penalized case (Equation 10) and the risk-neutral case (when $\theta = 0$ in Equation 8). The optimal times for maintenance obtained from these three models for each system are shown in Exhibit 3.

A pattern emerges from the solutions presented in Exhibit 3 and the results provide insights for managers. The risk-neutral policy seems to wait too long to perform PM, while the variance-penalized policy appears to perform premature maintenance. It needs to be kept in mind that the optimal value, which is provided by the semi-variance metric, seeks to minimize the average costs and at the same time limit the semi-variance. Hence it is necessary to determine the deviation of the risk-neutral and variance-penalized policies from the optimal

Exhibit 3. Optimal Time, T*, for PM in the Renewal-Theoretic Model Using the Three Criteria

Case	T* with Semi-Variance Model (hrs)	T* with Risk Neutral Model (hrs)	T* with Variance Model (hrs)
1	17	24	15
2	21	30	18
3	8	12	4
4	37	43	28
5	20	24	14
6	25	34	23
7	35	46	30
8	24	30	20
9	34	41	28

objective function (i.e., Equation 8). Exhibit 4 shows the values of $g(T,\theta)$ where T is obtained from optimizing for either semi-variance, risk-neutrality, or variance-penalties, along the expected costs obtained when θ =0. The values of the objective function of the cyclical variance and asymptotic variance are the same. This result shows that optimizing with respect to variance or using a risk-neutral objective can significantly affect the semi-variance-penalized score, which should be ideally minimized in order to obtain a budget-sensitive policy. The improvement in percent in the variance-penalized score over the risk-neutral and variance case is defined as follows:

$$Improvement\ (Risk-neutral) = \frac{g(T^*(Risk-neutral),\theta) - g(T^*(Risk-variance),\theta)}{g(T^*(Risk-neutral),\theta)} \times 100; \quad (20)$$

$$Improvement\ (Variance) = \frac{g(T^*(Variance),\theta) - g(T^*(Semi-variance),\theta)}{g(T^*(Variance),\theta)} \times 100. \quad (21)$$

When θ = 0, the negative of the value of the formula in Equation 20 represents the increase (in %) in the expected cost due to the use of the semi-variance policy. This value reflects the downside of using a risk-penalized policy from the perspective of expected costs. When a manager looks at the gains obtained from using a risk-penalized policy, he or she may also be interested in examining the increase in the expected cost that will result from using the risk-averse policy. Exhibit 5 shows the improvements and the increases in costs. The improvements in the budget-sensitive score range from 2.9% (Case 5) to 27.9% (Case 2) over the risk-neutral policy. The improvements in the same score over the variance-penalized policy range from 1.26% (Case 6) to 10.13 % (Case 3). The increase in expected cost can range from 3.38 % (Case 5) to 11.92 % (Case 2). It should be noted that a given percent of increase in the expected cost cannot be equated or compared to the same increase in percent terms of the semi-variance score, since the two do not have the same units. In fact, the semi-variance score contains two terms with different units ($/hour and $\$^2$/hour). And yet, the improvement in the score and increase in the expected cost (both in percent) should for very useful metrics for guiding the manager in decision-making. It is also important to understand that the actual improvement will depend on two other factors: (i) the risk-averseness of the manager (i.e., value of θ) and the magnitude of the inherent randomness (variability) in the system. The latter is hard to quantify; however, clearly a system with very little randomness should perform similarly under the risk-neutral and the risk-averse criteria. The results in Exhibit 5 show that the improvement in the semi-variance score of the semi-variance policy over the risk-neutral policy can vary significantly: from 2.93% (Case 5) to 27.97% (Case 2); thus, a manager in a system such as the one in Case 5 may be less willing to embrace a risk-averse policy especially if the expected cost increase is significant. Overall, the results show that managers can expect significant gains from using budget-sensitive policies, but should also consider the corresponding increases in expected costs. The improvements over risk-neutral policies, which tend to be the most popular policies in TPM, are more significant than those over the variance-penalized policies. While this is expected from our explanation of variance and semi-

Exhibit 4. The Scores and Expected Costs Associated with the Three Criteria

Case	$g(T^*,\theta)$ (semi-variance model)	Cost (semi-variance model)	$g(T^*,\theta)$ (risk-neural model)	Cost (risk-neutral model)	$g(T^*,\theta)$ (variance model)	Cost (variance model)
1	0.0953	0.0850	0.1103	0.0767	0.0972	0.0909
2	0.0618	0.0563	0.0858	0.0503	0.0631	0.0609
3	0.4009	0.3423	0.4375	0.3189	0.4461	0.4369
4	0.107	0.0993	0.1116	0.0952	0.1179	0.1166
5	0.0892	0.0793	0.0919	0.0767	0.0977	0.0945
6	0.0548	0.0508	0.0667	0.0458	0.0555	0.0531
7	0.1348	0.1205	0.1616	0.1080	0.1402	0.1344
8	0.1394	0.1306	0.155	0.1221	0.146	0.1435
9	0.0555	0.0502	0.0602	0.0472	0.0587	0.0569

Exhibit 5. Improvement in the Score of the Semi-Variance Policy over that of the Risk-Neutral and Variance Policies and Increase in Cost of the Semi-Variance-Optimal Policy over that of Risk-Neutral Policy

Case	Improvement of score over risk-neutral policy	Improvement of score over variance policy	Increase of expected cost over risk-neutral policy
1	13.59%	1.95%	10.82%
2	27.97%	2.06%	11.92%
3	8.36%	10.13%	7.34%
4	4.12%	9.24%	4.3%
5	2.93%	8.70%	3.38%
6	17.84%	1.26%	10.91%
7	16.58%	3.85%	11.57%
8	10.06%	4.52%	6.96%
9	7.80%	5.45%	9.63%

variance, the results also show that the manager should be careful about how risk is measured for TPM purposes, and whether risk should be a consideration in the decision-making.

SMDP Model

For the SMDP model, some assumptions are made that are justifiable for a large transfer line in an automotive plant. It is assumed that the state is modeled by the number of days since equipment was last repaired or maintained. Under action "produce," the transition probability is defined as follows:

$$p(i, produce, i+1) = 1 - p(i, produce, 0) \tag{22}$$

where $p(i,produce,0)$ denotes the probability of failure after the decision to produce is made in state *i*. For a complex system with numerous components and random variables, it is easier to determine the transition probability of failure from historical data. To do this let $K(i)$ denote the number times the system transitions from day *i* to *i+1* without failure in the historical data set. Let $K'(i)$ denote the number of times the system fails during the day after a decision to produce is made on day *i*. Then for every *i*,

$$p(i, produce, 0) = \frac{K'(i)}{K(i) + K'(i)}\ . \tag{23}$$

Clearly, as the denominator starts becoming large, the estimates of this probability improve. The mean time taken by the production cycle is denoted by t_p. It is assumed that when a machine fails it is repaired after a time interval whose mean duration since the start of the production cycle is $M1*t_p$ time units. Also when the machine is maintained, the maintenance is complete after a time interval whose mean duration since the start of the production cycle is approximately $M2*t_p$ time units. The values of M1 and M2 depend on the system modeled. This assumption about the duration of the repair and maintenance time is rather general and allows our model to be applicable over a large range of systems.

Typically preventive maintenance is useful *only* in systems with increasing failure rates (Lewis, 1995), i.e., systems in which the probability of failure increases as the system ages. The value of $p(i,produce,0)$, which is the probability of failure when its age is *i*, depends on the numerous random variables, such as the time between failures of all the interacting system components, repair times, and maintenance times. In order to capture the general behavior of *any* complex system that has an increasing failure rate, a *generic* model that provides the probability, $p(i,produce,0)$, will be used in the numerical experiments. The advantage of the generic model is that it can be applied to any complex production line regardless of the number of components in it that can fail. The generic model is as follows:

$$p(i,produce,0) = 1 - \psi^i \tag{24}$$

where ψ is a small positive number whose value will depend on the system being studied, and the precise value of ψ must be estimated from historical data of past failures of the system studied. In the generic model above, as i increases, $p(i,produce,0)$ increases and tends to 1 as i tends to infinity; this essentially captures the phenomenon of the probability of failure increasing with an aging system and the fact that ultimately an unattended system fails with probability 1.

It is not hard to see that $c(i,produce,0)=C_r$ and $c(i,maintain,0) =C_m$. The example used to illustrate this model is based on a large automobile assembly line in New York. The advantage of the generic model in Equation 24 is that it is widely applicable to a range of complex production lines and is highly compatible with the SMDP solution methodology.

For the computational study, the LP and the quadratic program described above are used. The expression in Equation 15 or 19, which is the objective function for the SMDP model, can also be expressed as $S(criterion)$. Note that $S(.)$ is a function of the criterion used for determining the optimal maintenance policy. For the semi-variance criterion, Equation 15 is used with a strictly positive value for θ. For the risk-neutral criterion, Equation 15 is used with $\theta = 0$. For the variance criterion, Equation 19 is used with a strictly positive θ. The expected cost, γ, for risk-neutral case equals the value in Equation 15 with $\theta = 0$; for the variance criterion, it is the value of the sum of terms without θ in Equation 19, and for the semi-variance criterion, it is the first term in Equation 15. Improvements of the semi-variance criterion over the risk-neutral and variance criteria and the increase in expected cost over the risk-neutral criterion are defined in a manner analogous to that for the renewal theory model.

Exhibit 6 describes the input parameters for the different cases evaluated by the SMDP model. Note that τ = 0.15 in dollars per day was used for all the cases. The data here is based on data in Gosavi (2006), which in turn was based on historical failures in an automotive plant in New York. Exhibit 7 lists the results of optimizing with respect to the semi-variance criterion (Svar), the risk-neutral (RN) criterion, and the variance (Var) criterion. The exhibit also lists the value of *i*, the age in number of days, at which the system should be maintained, depending on the criterion chosen. It can be seen that like the renewal-theoretic model, the optimization models recommend premature maintenance with the variance criterion and late maintenance with the risk-neutral criterion. Again, this is due to the fact that the risk-neutral criterion ignores the variability and occasional failures that cause the costs to exceed the budget. The variance criterion is over-sensitive to variability, since it measures variability both above and below the mean. This causes it to recommend maintenance early in order to minimize the chances of any variation, including perhaps the variation of the costs below the targets, which are actually welcome. Exhibit 8 shows the resulting semi-variance scores and expected costs of the policies that optimize with respect to the three criteria. Exhibit 9 contains the improvements in the scores range from 1.57% (Case 3) to 50.21% (Case 5) over the risk-neutral criterion and from 0.08% (Case 3) to 4.27% (Case 5) over the variance criterion, and the increases in the expected costs that range from 0.22% (Case 3) to 12.00% (Case 4). As stated in the context of the previous model, the manager must carefully weigh the downside of increased cost against the reduced variability before implementing a policy.

Both models, the renewal reward model and the SMDP model, can be solved using either a spreadsheet package (such as Microsoft Excel) or a computational software package (such as MATLAB). Models were also coded in MATLAB and are available upon request from the first author.

Conclusions and Managerial Insights

This article presented improvements in planning maintenance scheduling for Total Productive Maintenance (TPM). It has been recognized over the last five decades that TPM saves thousands of dollars in lost production, reduces production lead time, makes systems less variable, and can extend equipment lifespan. Much of the existing maintenance scheduling literature seeks to minimize the long-run average costs with no regard to ongoing maintenance budgets. These risk-neutral policies are based on variance. This article presented two models that exploit a less-known measure of risk, called semi-variance. This combined with long-run cost criterion results in preventive maintenance policies that minimize long-run average costs and at the same time reduce the frequency of costs exceeding the budget. The performance of

Exhibit 6. Input Parameters for the SMDP Model

Case	ψ	C_r ($)	C_m ($)	θ
1	0.94	5	2	0.2
2	0.92	6	4	0.2
3	0.91	7	5	0.1
4	0.88	8	5	0.3
5	0.93	6	2	0.2
6	0.92	7	5	0.2
7	0.89	6	4	0.3
8	0.96	6	2	0.2
9	0.90	5	2	0.2
10	0.95	10	7	0.1

(Note: Number of States is 100 and *M1=2*, *M2=1.25*, and $t_p=15$ hours. A day is assumed to be composed of 16 hours.)

Exhibit 7. Optimal Time for Maintenance under Various Policies

Case	Policy (Variance)	Policy (semi-variance)	Policy (Risk-neutral)
1	3 days	4 days	5 days
2	5 days	6 days	9 days
3	7 days	8 days	10 days
4	1 day	3 days	6 days
5	2 days	3 days	4 days
6	6 days	7 days	10 days
7	4 days	5 days	7 days
8	3 days	4 days	5 days
9	2 days	3 days	4 days
10	8 days	9 days	12 days

Exhibit 8. The SMDP Model Performance with Varying Policies after Optimizing

Case	S(Semi-variance)	γ(Semi-variance)	S(Variance)	γ(Variance)	S(Risk-neutral)	γ(Risk-neutral)
1	0.0559	0.0495	0.0775	0.0531	0.0568	0.0491
2	0.0970	0.0752	0.1077	0.0775	0.0976	0.0741
3	0.1123	0.0911	0.1141	0.0915	0.1124	0.0909
4	0.2109	0.1302	0.7472	0.2667	0.2199	0.1161
5	0.0694	0.0592	0.1394	0.0681	0.0725	0.0584
6	0.1277	0.0875	0.1363	0.0891	0.1280	0.0866
7	0.1211	0.0860	0.1410	0.0897	0.1221	0.0845
8	0.0564	0.0473	0.1030	0.0516	0.0579	0.0465
9	0.0676	0.0608	0.1018	0.0689	0.0691	0.0604
10	0.1529	0.1021	0.1613	0.1031	0.1538	0.1014

Exhibit 9. Improvement in the Score of the Semi-Variance Policy over that of the Risk-Neutral and Variance Policies and increase in Cost of the Semi-Variance-Optimal Policy over that of Risk-Neutral policy

Case	Improvement of score over risk-neutral policy	Improvement of score over variance policy	Increase of expected cost over risk-neutral policy
1	1.58%	27.87%	0.81%
2	0.61%	9.93%	1.48%
3	0.09%	1.57%	0.22%
4	4.09%	71.77%	12.00%
5	4.28%	50.21%	1.36%
6	0.23%	6.31%	1.04%
7	0.82%	14.11%	1.77%
8	2.59%	45.24%	1.72%
9	2.17%	33.60%	0.66%
10	0.59%	5.21%	0.69%

the semi-variance-penalized criterion, also called the budget-sensitive criterion, was compared to that of the risk-neutral criterion and to the variance-penalized criterion. Numerical results indicate that the budget-sensitive criterion outperforms the risk-neutral and the variance-penalized criterion.

TPM is an important tool in the engineering manager's arsenal. It is useful in reducing lead time and keeping system variability in check. Our results show that the commonly used policies based on the industry standard, which happens to be risk-neutral, has the potential to exceed budgets. Although variance is a popular measure of risk and variability, using variance-penalized policies may actually turn out to be worse than the risk-neutral standard from the perspective of meeting budgets, as was shown in some of the cases presented in this article.

Many production operations collect a wealth of historical data. The models presented can be easily implemented with data related to equipment failures, repairs, and maintenance activities. For simpler systems composed of one or two pieces of equipment, the Markov chain model with its transition probabilities may be deemed too time consuming. The simpler renewal reward model will provide improved scheduling with a manageable amount of effort. The Markov chain model is best suited for a complex system composed of numerous pieces of equipment. An interim approach could be to model a key piece of equipment (the process bottleneck) using the renewal reward model and verify the resulting improvements before modeling the entire process using the semi-Markov decision process model.

The budget-sensitive aspect of the semi-variance metric lends itself more naturally to a finite horizon analysis of the problem. Hence, although PM problems for production machines are typically studied over the infinite horizon when modeled as MDPs, one potential direction for future research could be the development of a finite horizon model for the problem studied here. The target in this model could be defined as the total maintenance budget for the time horizon, which could have the duration of a few months, as is typically the case with budgets set by production firms.

Acknowledgements

The authors would like to thank the three anonymous reviewers whose comments significantly improved the quality of this work and its presentation. The first author would also like to acknowledge support from the NSF via grant ECS: 0841055.

References

Ahuja, Inder, and Jana. S. Khamba, "Total Productive Maintenance: Literature Review and Directions," *The International Journal of Quality & Reliability Management*, *25*:7 (2008), pp. 709-756.

Anderson, Michael Q., "Monotone Optimal Preventive Maintenance Policies for Stochastically Failing Equipment," *Naval Research Logistics Quarterly*, *28* (1981), pp. 347-358.

Askin, Ronald G., and Jeffrey B. Goldberg, *Design and Analysis of Lean Manufacturing Systems*, John Wiley (2002).

Gosavi, Abhijit, "A Risk-Sensitive Approach to Total Productive Maintenance," *Automatica*, *42* (2006), pp. 1321-1330.

Gosavi, Abhijit, "On Step Sizes, Stochastic Shortest Paths, and Survival Probabilities in Reinforcement Learning," *Proceedings of the Winter Simulation Conference*, (2008).

Gosavi, Abhijit, "Semi-Variance Sensitive Control of Markov and Semi-Markov Processes," Working Paper at the Missouri University of Science and Technology (2010).

Barlow, Richard E., and Frank Proschan, *Mathematical Theory of Reliability*, John Wiley (1965).

Bertsekas, Dimitri, *Dynamic Programming and Optimal Control*, Volume II, 3rd Ed., Athena Scientific (2007).

Brown, Mark, and Herbert Solomon, "A Second-Order Approximation for the Variance of a Renewal Reward Process," *Stochastic Processes and their Applications*, *3* (1975), pp. 301-314.

Cassady, Richard, and Erhan Kutanoglu, "Integrating Preventive Maintenance and Production Scheduling for a Single Machine," *IEEE Transactions on Reliability*, *54*:2 (2005).

Chen, Yong, and Jonghua Jin, "Cost-Variability-Sensitive Preventive Maintenance Considering Management Risk," *IIE Transactions*, *35*:12 (2003), pp. 1091-1102.

Chitke, Satya, and Satish Deshmukh, "Preventive Maintenance and Replacement under Additive Damage," *Naval Research Logistics Quarterly*, *28* (1981), pp. 33-46.

Dada, Maqbool, and Richard Marcellus, "Process Control with Learning," *Operations Research*, *42*:2 (1994), pp. 323-336.

Das, Tapas K., and Sudeep Sarkar, "Optimal Preventive Maintenance in a Production Inventory System," *IIE Transactions*, *31* (1999), pp. 537–551.
Dellagi, Sofiane, Nidhal Rezg, and Xiao-Lan Xie, "Preventive Maintenance of Manufacturing Systems under Environmental Constraints," *International Journal of Production Research*, *45*:5 (2007), pp. 1233-1254.
D'epenoux, F., "Sur un Probleme de Production et de Stockage dans l'Aleatoire," Revue Francaise de Recherche Operationnelle, *14* (1960).
El-Ferik, Sami, and Mohamed Ben-Daya, "Age-Based Hybrid Model for Imperfect Preventive Maintenance," *IIE Transactions*, *38* (2006), pp. 365-375.
Filar, Jerzy, Lodewijk Kallenberg, and H. Lee. "Variance-Penalized Markov Decision Processes," *Mathematics of Operations Research*, *14*:1 (1989), pp. 147-161.
Gertsbakh Ilya, *Reliability Theory with Applications to Preventive Maintenance*, Springer (2000).
Howard, Ronald A., *Dynamic Probabilistic Systems, Vol. II: Semi-Markov and Decision Processes*, John Wiley & Sons (1971).
Lewis, Elmer, *Introduction to Reliability Engineering 2nd Ed.*, John Wiley and Sons (1995).
Manne, A., "Linear Programming and Sequential Decisions," *Management Science*, *6* (1960), pp. 259-267.
Markowitz, Harry, "Portfolio Selection," *Journal of Finance*, *7*:1 (1952), pp. 77-91.
McCall, John J., "Maintenance Policies for Stochastically Failing Equipment: A Survey," *Management Science, 11*:5 (1965), pp. 493-524.
McKone, Kathleen E., Roger G. Schroeder, and Kristy O. Cua, "Total Productive Maintenance: A Contextual View," *Journal of Operations Management*, *17* (1999), pp. 123-144.
McKone, Kathleen E., and Eliott N. Weiss, "TPM: Planned and Autonomous Maintenance: Bridging the Gap Between Practice and Research," *Production and Operations Management*, *7*:4 (1998), pp. 335-351.
Nayak, Bijay, and Ebrahim Shayan, "Maintenance Improvement Opportunities and Benchmarketing Practices in Australian Manufacturing Industry," *Engineering Management Journal*, *10*:2 (1998), pp. 33-37.
Panagiotidou, Sofia, and George Tagaras, "Optimal Preventive Maintenance for Equipment with Two Quality States and General Failure Time Distributions," *European Journal of Operational Research*, *180* (2007), pp. 329-353.
Porter, Burr, "Semi-Variance and Stochastic Dominance," *American Economic Review*, *64* (1974), pp. 200-204.
Quan, Gang, Garrison W. Greenwood, Donglin Liu, and Sharon Hu, "Searching for Multi-Objective Preventive Maintenance Schedules: Combining Preferences with Evolutionary Algorithms," *European Journal of Operational Research*, *177* (2007), pp. 1969-1984.
Ross, Sheldon, *Introduction to Probability Models 6th Ed.*, Academic Press (1997).
Sheu, S.H., and Y.H. Chien, "Optimal Age-Replacement Policy of a System Subject to Shocks with Random Lead-Time," *European Journal of Operational Research, 159*:1 (2004), pp. 132-144.
Shimbun, Nikkan (Ed.), *TPM Case Studies*, Productivity Press (1995).
Tijms, Henk, *A First Course in Stochastic Models*, Wiley (2003).
Valdez-Flores, Ciriaco, and Richard M. Feldman, "A Survey of Preventive Maintenance Models for Stochastically Deteriorating Single-Unit Systems," *Naval Research Logistics Quarterly*, *36* (1989), pp. 419-446.
Wang, Hengzhou, "A Survey of Maintenance Policies of Deteriorating Systems," *European Journal of Operational Research*, *139* (2002), pp. 469-489.
Westerkamp, Thomas, "Maintaining Maintenance," *IIE Solutions, 38*:7 (July 2006), pp. 37-42.
Wilmeth, Randall G., and Michael W. Usrey, "Reliability-Centered Maintenance: A Case Study," *Engineering Management Journal, 12*:4 (2000), pp. 25-31.

About the Authors

Abhijit Gosavi has a BE and an MTech in mechanical engineering and a PhD in industrial engineering from the University of South Florida. He joined the Department of Engineering Management and Systems Engineering at the Missouri University of Science and Technology (Missouri S&T) in 2008 as an assistant professor. His research interests are in simulation, reinforcement learning, and manufacturing.

Susan L. Murray, PhD, PE is a professor of Engineering Management and Systems Engineering at Missouri University of Science and Technology. She received her PhD (1994) and BS in industrial engineering from Texas A&M University and an MS degree in IE from the University of Texas-Arlington. Her research interests include safety engineering, human systems integration, industrial engineering, and human factors.

V. Manojramam Tirumalasetty has a BE in electrical and electronics from Anna University, Chennai, India. He obtained an M.S. in engineering management from Missouri S&T in 2010. He was the recipient of the best graduate research assistant award from the Department of Engineering Management and Systems Engineering at Missouri S&T in 2009. He is currently employed as a systems analyst in Ungerboeck Systems International in St. Louis, MO.

Shreerang Shewade has a BE in mechanical engineering from the University of Pune, India. He obtained an MS in operations research from the University at Buffalo (SUNY) in 2006, where his research was focused on semi-variance modeling and data collection for machine failures. He is currently employed as an engineering associate in Analytical Mechanics Associates in Plano, TX.

Contact: Dr. Susan Murray, Missouri S&T, EMSE, 600 West 14th Street, Engineering Management Bldg., Rolla, MO 65409; phone: 573-341-4038; murray@mst.edu