

# Optimization of Inspection and Maintenance Decisions for Infrastructure Facilities under Performance Model Uncertainty: A Quasi-Bayes Approach\*

Pablo L. Durango-Cohen<sup>†</sup>      Samer M. Madanat<sup>‡</sup>

## Abstract

We present an optimization model to find joint inspection and maintenance policies for infrastructure facilities under performance model uncertainty. The objective in the formulation is to minimize the total expected social cost of managing facilities over a finite planning horizon. As in recent optimization models, performance model uncertainty is accounted for by representing facility deterioration as a mixture of known models taken from a finite set. The mixture proportions are assumed to be continuous random variables, with probability densities that are updated over time. In this paper, we relax the assumptions of fixed and error-free inspections. We present a parametric study to analyze the effect of initial performance model uncertainty and bias on the expected total cost of managing a facility. The main observation is that reducing the initial variance in model uncertainty may be more important than reducing the initial bias. Our study also shows that cost savings can result from relaxing the constraint of a fixed inspection schedule.

**Keywords:** Infrastructure, Systems management, Stochastic models, Adaptive control, Finite Mixtures, Quasi-Bayes

---

\*This paper was accepted for publication in *Transportation Research Part A: Policy and Practice* in September of 2006.

<sup>†</sup> Assistant Professor, Department of Civil and Environmental Engineering & Transportation Center, Northwestern University, Evanston, IL 60208-3109, Tel: (847) 491-4008, Fax: (847) 491-4011, Email: [pdc@northwestern.edu](mailto:pdc@northwestern.edu).

<sup>‡</sup> Corresponding Author, Professor, Department of Civil and Environmental Engineering, University of California, Berkeley, CA 94720, Tel: (510) 643-1084, Fax: (510) 642-1046, Email: [madanat@ce.berkeley.edu](mailto:madanat@ce.berkeley.edu).

# 1 Introduction

Infrastructure systems play a vital role in achieving a productive and competitive national economy and improving citizens' overall quality of life. The United States has historically made an extraordinary investment in its infrastructure. For instance, the federal government has spent an average of about \$59 billion annually since the 1980s on the nation's civilian infrastructure. This spending showed a slightly upward trend through the 1990s. Similarly, spending by state and local governments continued an upward trend that began in the 1980s and exceeded federal spending in certain areas (GAO, 2001).

However, due to extensive use and insufficient maintenance, the public infrastructure, including highways, bridges, pipelines, mass transit, water supply, energy and wastewater treatment facilities are rapidly deteriorating. The recent emphasis in infrastructure investment has shifted toward maintenance rather than new construction. Of the total expenditure on public works improvements, an increasingly larger proportion is being spent on maintenance. Government statistics show that the proportion of public non-capital spending for infrastructure increased from 39% in 1960 to 57% in 1994 (CBO, 1999), a clear indication of the rising importance of maintenance and rehabilitation (M&R) vs. new facility construction. However, the magnitude of M&R investment has been far from sufficient to meet the growing demands for M&R. Therefore, the critical issue facing public works agencies today is how to allocate the limited resources that are available for M&R so as to obtain the best return for their expenditure.

Infrastructure management is the process through which agencies collect and analyze data about infrastructure systems and make decisions concerning M&R of facilities over a planning horizon. In developed countries, where much of the transportation infrastructure is mature and portions are nearing the end of their service lives and need to be replaced, M&R decisions are increasingly important. This is due to both the negative social and economic impact of deficient infrastructure, as well as the scale of management expenditures, which in the United States, for example, are on the order of tens of billions of dollars per year. In turn, the importance of transportation infrastructure management has, over the last 40 years, spawned a great deal of research in developing models that can aid agencies in evaluating and selecting M&R policies (see for example Gendreau and Soriano (1998) and Van Noortwijk and Frangopol (2004) and the references therein).

Models to support infrastructure management are used to evaluate both the short and long-term economic consequences associated with M&R decisions. This evaluation involves processing data related to current infrastructure condition and forecasting the effect of M&R decisions on future condition. Data related to infrastructure condition consist of distress evaluations and measurements. Condition forecasts are generated with deterioration models, which in this paper correspond to statistical expressions that relate condition to explanatory variables such as design characteristics, traffic loading, environmental factors, and history of M&R activities. The economic consequences of M&R decisions are then estimated with a function that maps condition forecasts to costs. The motivation for the work presented herein is that, typically, several competing models can be used to represent deterioration. The multiplicity of models results from different underlying assumptions about the physical deterioration process, or they may represent different expert opinions about the future performance of a facility. This leads to a situation where agencies face uncertainty in the choice of an appropriate model. This type of uncertainty is referred to as epistemic or performance model uncertainty. Note that this uncertainty is not a result of the underlying randomness in the deterioration process, i.e., aleatory uncertainty. Indeed, each of the competing models may be probabilistic.

The developers of modern infrastructure management systems have recognized the presence of uncertainty in the choice of deterioration models used in practice. They have included a model-updating step in these management systems, where data collected as part of condition surveys are used to update deterioration model parameters or the relative weights placed on the competing models. For example, Harper and Majidzadeh (1991) use Bayesian methods to update the parameters of their deterioration models. Likewise, in the popular bridge management system Pontis, transition probability matrices are updated over time (Golabi and Shepard, 1997).

In the proposed framework, we represent a facility’s deterioration process as a mixture of known models taken from a finite set. We assume that the mixture proportions are continuous random variables. Performance model uncertainty is captured by the probability density function associated with the mixture proportions. From a “Bayesian” perspective, the probability density function represents an agency’s beliefs about deterioration. The density function can be updated in response to condition data gathered during the management process. As is shown in Durango and Madanat

(2002), over time, such an adaptive scheme can lead to estimates that provide an adequate characterization of the physical deterioration process. It is clear that, in practice, the success of such an adaptive scheme, and consequently the choice of M&R actions, depends critically on two closely-related managerial decisions: the inspection frequency and the choice of technology used to inspect a facility.

In this paper we present an extension that simultaneously optimizes M&R and inspection decisions. In this context, facility inspections serve two purposes: First, they provide an assessment of the current facility condition. Second, they provide information that is used to update an agency's beliefs, which in turn determine predictions about future condition. Earlier adaptive M&R optimization models, such as Durango and Madanat (2002) or Madanat et al. (2006), assume that facilities are inspected in every period and that the inspection process is error-free, i.e., that it reveals the true facility condition. This motivates the development of a methodology that captures both the optimal timing of inspections, and the capabilities of the technologies that are available to inspect a facility.

## 2 Literature Review

Optimization models to support M&R decisions have been developed from different perspectives to address numerous applications. Extensive surveys of M&R models and applications appear in McCall (1965), Pierskalla and Voelker (1976) and Dekker (1996). With few exceptions, discrete-time M&R optimization models, consistent with the periodic review nature of infrastructure management, are formulated as finite (state and action) Markov Decision Processes (MDPs). In the remainder of this section, we provide a review of MDP models, on which we build. However, we emphasize that our contribution is general because uncertainties in selecting performance models or in measuring facility condition are not unique to the MDP framework (or to the problem of obtaining a joint inspection and M&R policy). Examples of closely-related M&R models building on other frameworks are presented in Tsunokawa and Schofer (1994), Li and Madanat (2002), Ouyang and Madanat (2004) and Suzuki and Pautsch (2005).

## 2.1 Markov Decision Process Formulations

Derman (1962) and Klein (1962) were first to propose the MDP formulations for M&R problems. The first adaptation of the methodology to support the management of transportation infrastructure is discussed in Golabi et al. (1982), where a mixed-criteria, constrained MDP is proposed for pavement management in the state of Arizona (a network of 12,000 km. of highways). The same optimization model drives Pontis (Golabi and Shepard, 1997), a bridge management system licensed to more than 45 state DOTs and other agencies nationally and internationally, and perhaps the most widely used implementation of a M&R optimization model. MDP formulations are appealing to manage infrastructure because they provide a rigorous framework to account for uncertainty, because optimal policies can be obtained by solving a linear program<sup>1</sup>, and perhaps most importantly, because they have been successfully implemented. For example, savings of \$14 million were reported in the first year of implementation of the Arizona Pavement Management System, and \$101 million was forecast for the subsequent four years. Not surprisingly, the methodology has been adopted and implemented to support transportation infrastructure in various countries, e.g., Finland (Thompson et al., 1987), Denmark (Ullidtz, 1987) and Portugal (Golabi and Pereira, 2003), and has also been used to support the management of other infrastructure systems, e.g., buildings (Van Winden and Dekker, 1998), dikes (Van Noortwijk and Van Gelder, 1996), and sewers (Wirahadikusumah, 1999).

MDP models to support M&R decision-making consider the management of facilities over a planning horizon of length  $T$ . At the start of each decision-making stage,  $t = 1, 2, \dots, T$ , an agency reviews the facility's state, represented with a random variable  $X_t$  (defined over a finite set  $\mathcal{S}$ ), and selects a M&R activity  $A_t$  (from a finite set  $\mathcal{A}$ ). Thus, the history of the process up to  $t$  can be represented with the vector  $H_t \equiv \{X_1, A_1, X_2, A_2, \dots, X_{t-1}, A_{t-1}, X_t\}$  (defined over the set  $\mathcal{H}_t \equiv \mathcal{S} \times [\mathcal{A} \times \mathcal{S}]^{t-1}$ ).<sup>2</sup> If activity  $A_t = a$  is chosen at the start of  $t$  when a facility's condition is  $X_t = i$ , then:

1. A cost  $g(X_t = i, A_t = a) \in \mathbb{R}$  is incurred. This cost structure can be used to capture both agency and user costs. At the end of the planning horizon, (i.e.: for  $t = T + 1$ ) a facility is assumed to have a salvage/residual value,  $s(X_{T+1}) \in \mathbb{R}$ , that depends on its terminal state.

---

<sup>1</sup>Finite horizon MDP models are usually solved by backwards induction.

<sup>2</sup>The symbol “ $\times$ ” denotes a Cartesian product. We use the power function applied to a set to represent multiple Cartesian products of the set with itself, e.g.,  $\mathcal{S}^2 \equiv \mathcal{S} \times \mathcal{S}$ .

2. The facility’s subsequent state is given by a *deterioration model*,  $\mathcal{D}(X_t = i, A_t = a)$ , as follows:

$$X_{t+1} = \mathcal{D}(X_t = i, A_t = a) = \{j; \text{with probability } \pi_{ij}(a)\} \quad (1)$$

where  $\pi_{ij}(a)$  is the conditional probability that the facility deteriorates from state  $i$  to state  $j$  during  $t$ , given that it receives activity  $a$ . References such as Carnahan et al. (1987), Jiang and Sinha (1989), Madanat and Wan Ibrahim (1995); Madanat et al. (1997) and Mishalani and Madanat (2002) present different approaches to estimate transition probabilities.

The objective in M&R optimization is to choose a set of actions, a *policy*, that minimizes the expected discounted sum of agency and user costs incurred over the planning horizon. Because both the costs and the deterioration are *Markovian*, i.e., depend only on the last state and decision, there is an optimal policy that can be represented as a mapping,  $\mu$ , that prescribes a decision in  $\mathcal{A}$  for every stage-state pair in the set  $\{1, 2, \dots, T\} \times \mathcal{S}$ , i.e.,  $\mu : \{1, 2, \dots, T\} \times \mathcal{S} \mapsto \mathcal{A}$ .

Two important limitations of MDP models, as described above, are related to the implicit assumptions that (i) the true facility condition is revealed at the start of every period, and that (ii) agencies can select (without uncertainty) a deterioration model that provides a perfect, albeit stochastic, representation of the physical deterioration process. In the following subsections, we discuss how these limitations have been addressed individually.

### 2.1.1 Inspection and Maintenance Optimization with Imperfect Information

The possibility of a flexible inspection schedule and the presence of measurement errors constitute a violation of the basic premise of MDPs that the true facility condition is revealed at the end of each period. Incorporating the decision to inspect in M&R optimization models is relatively straightforward (cf. Klein (1962)). However, these models do not allow for the presence of measurement errors. Madanat (1993) explains that in the context of M&R optimization, measurement errors need to be accounted for because they can lead to the selection of inappropriate actions when a policy specifies different actions for the true condition and the measured condition. Due to the presence of measurement errors, the measured facility condition is now only probabilistically related to its true condition. We build on Madanat and Ben-Akiva (1994) and assume that the distribution of the measurement relative to the true state is known and depends on the technology used. The

relationship between measured and true states is given by:

$$\epsilon_{jk}^c \equiv P\left(\hat{X}_t = k | X_t = j, C_t = c\right) \quad (2)$$

where  $\hat{X}_t$  (defined over  $\mathcal{S}$ ) is the measured state at the start of stage  $t$ ,  $C_t$  is the technology used to measure and  $c$  is an element of the set of available technologies  $\mathcal{C}$ . The necessary probabilities can be derived from empirical measurement error models such as those presented in Humplick (1992). In particular, the results presented in that study show that, indeed, the capabilities associated with inspection technologies are a primary source of measurement errors, even though, in general, they also depend on factors such as the facility being inspected, equipment operators, or the nature and magnitude of the underlying distress.

Having to rely on error-prone measurements (imperfect information) has significant implications on the management process (and on the problem’s computational complexity) because control policies can only rely on information available to a decision-maker. In this case, the set of available information, from the start of the planning horizon can be written as follows:

$$I_t \equiv \left\{ I_1, A_1, \dots, \hat{X}_{t-1}, A_{t-1}, \hat{X}_t \right\} \quad (3)$$

$I_1$  represents the initial information available at the start of the planning horizon (including measurement  $\hat{X}_1$ ).

### 2.1.2 Maintenance Optimization under Performance Model Uncertainty

As in other problems of decision under uncertainty, M&R policies are evaluated and selected under the assumption that a single, “best choice” performance model provides a perfect representation of deterioration. This approach ignores the uncertainty that decision-makers face when choosing an appropriate model. This uncertainty stems from unobserved heterogeneities, as well as epistemic and parametric uncertainties. Unobserved heterogeneities refer to the presence of persistent, facility-specific, but unobserved factors such as construction quality. Epistemic and parametric uncertainties, respectively, refer to uncertainty about the structural assumptions that are made to specify and estimate performance models, and to the confidence levels that apply to the parameter estimates. The problem with ignoring performance model uncertainty in M&R optimization is that incorrect representations of deterioration can lead to the implementation of inefficient policies.

Indeed, Carnahan et al. (1987) provides examples of situations where M&R policies are sensitive to the transition probabilities.

MDP problems with unknown/uncertain transition probabilities are a special class of adaptive control problems that have been studied extensively. Kumar and Varaiya (1986) presents an excellent overview and provides extensive references. The fundamental attributes of these models include the characterization of uncertainty, and the scheme used to learn based on observed data. The model presented herein can be seen as a generalization of Durango and Madanat (2002); Guillaumot et al. (2003). In our earlier work, facility deterioration is represented as an unknown element from a finite set of models. In the current paper, we allow for deterioration to be represented as an unknown mixture of known performance models taken from a finite set. The estimation problem involves finding a mixture proportion that adequately represents the physical deterioration process. In the proposed framework, maximum likelihood *a posteriori* estimates of the mixture proportions are obtained using the Quasi-Bayes approach of Smith and Makov (1980). Finite mixture models provide a rigorous and flexible statistical approach to model unknown distributional shapes. Moreover, from a practical perspective, the approach is appealing because it mimics an agency's choice from a (finite) set of performance models, each estimated under different structural assumptions, or using different data, or relying on the opinion of different experts. We do note, however, that other characterizations of performance model uncertainty have been used in models to support M&R decisions for transportation infrastructure. In particular, Madanat et al. (2006) and Kuhn and Madanat (2005) present models where every row in the transition probability matrices is assumed to be a random vector. The former model relies on maximum likelihood estimation to update estimates of the probabilities. The latter proposes robust optimization to obtain policies that are appropriate for ranges of the transition probabilities.

### 3 Proposed Model

We propose a modeling framework where facility performance is represented by a finite mixture of MDPs. We begin this section by introducing the model parameters.



### 3.1 Performance Models

A set of transition probabilities is specified for each of the (stationary, Markovian) deterioration models in the finite set  $\mathcal{R}$ . Their transition probabilities are denoted:

$$\pi_{ij}^r(a) \equiv P(X_{t+1}|Y = r, X_t = i, A_t = a) \quad (4)$$

where  $Y = r$  indicates that deterioration is governed by model  $r$ .

We assume that a facility's true/physical deterioration is given by a mixture of models in the finite set  $\mathcal{R}$ . That means that the transition probabilities that govern the deterioration process are as follows:

$$\pi_{ij}^*(a) = \sum_{r \in \mathcal{R}} \lambda_r^* \pi_{ij}^r(a) \quad (5)$$

where  $\lambda_r^* \geq 0$  and  $\sum_{r \in \mathcal{R}} \lambda_r^* = 1$ , i.e.,  $\vec{\lambda}^*$  is an element of the  $(|\mathcal{R}| - 1)$ -dimensional unit simplex, which we denote  $\mathcal{S}$ . Another interpretation is that  $Y$  is a latent random variable defined over  $\mathcal{R}$ . An agency's estimates of  $\vec{\lambda}^*$  at  $t$ ,  $\hat{\lambda}^t$ , correspond to the probability mass function of  $Y$ .

### 3.2 Measurement Errors

The notation for representing measurement errors is the same as the one presented earlier. The inspection decision is represented by a choice between different types of technology: one with the measurement precision associated with each inspection technology, and the other with a measurement error of infinite variance. The model can accommodate a set of different technologies. However, in the computational study we reduce the choice to a binary decision: ( $C_t = 1$  for inspection and  $C_t = 0$  for no inspection).  $C_t = 0$  refers to a technology where for each state the probability of measuring any state is uniformly distributed. That is,

$$\epsilon_{jk}^0 = P(\hat{X}_t = k | X_t = j, C_t = 0) = \frac{1}{|\mathcal{S}|} \quad (6)$$

where  $|\mathcal{S}|$  denotes the cardinality of the set  $\mathcal{S}$ . This case, where every condition state is equally likely to be observed regardless of the true state, is shown to be equivalent to not inspecting in Madanat and Ben-Akiva (1994). The associated cost is set to zero.

### 3.3 Costs

We generalize the definition of  $g(X_t, A_t, C_{t+1})$  as the generic cost incurred during period  $t$  associated with activity  $A_t$  on a facility in state  $X_t$  and choosing to use inspection technology  $C_{t+1}$  at the beginning of next period. We consider a cost structure that includes components for user costs, inspection costs, costs associated with M&R activities, and a salvage cost at the end of the horizon. The projected salvage (residual) value of the facility at the end of the planning horizon is  $s(i)$ .

### 3.4 Dynamic Programming Formulation

Decision-making involves the choice of action to perform during period  $t$ , as well as whether to inspect or not at the beginning of  $t + 1$ . Not having perfect information about a facility's state means that the decisions taking place at the start of stage  $t$ ,  $A_t$  and  $C_{t+1}$ , depend on the information available to a decision-maker, which we represent with the vector

$I_t \equiv \{I_1, A_1, C_2, \hat{X}_2, A_2, C_2, \dots, \hat{X}_{t-1}, A_{t-1}, C_t, \hat{X}_t\}$  (defined over  $\mathcal{I}_t \equiv \mathcal{I}_1 \times [\mathcal{C} \times \mathcal{A} \times \mathcal{S}]^{t-1}$ ). A

control policy for the problem,  $\mu_t : \mathcal{I}_t \mapsto \mathcal{A} \times \mathcal{C}, t = 1, 2, \dots, T$ , is a function that maps the set of possible information sets to the possible combinations of M&R actions for the current period and technology choices for the following period. An agency's objective is to minimize the expected sum of discounted costs incurred over a planning horizon. The problem of obtaining such a policy can be formulated as a dynamic program as follows:

$$v_t(I_t) = \min_{A_t \in \mathcal{A}; C_{t+1} \in \mathcal{C}} \left\{ \mathbb{E}_{X_t|I_t} [g(X_t, A_t, C_{t+1})] + \delta \mathbb{E}_{I_{t+1}|I_t} [v_{t+1}(I_{t+1})] \right\}^3 \quad (7)$$

$$v_{T+1}(I_{T+1}) = \mathbb{E}_{X_{T+1}|I_{T+1}} [-s(X_{T+1})] \quad (8)$$

where:

$v_t(I_t)$ : is referred to as the optimal objective value function and corresponds to the minimum expected sum of discounted costs from the start of stage  $t$  until the end of the planning horizon (start of stage  $T + 1$ ), given that the set of information available at the start of stage  $t$  is  $I_t$ .

$\delta$ : Discount amount factor associated with the time value of money ( $\delta < 1$ ).

---

<sup>3</sup> $\mathbb{E}_Y$  is the expectation operator. The expectation is taken with respect to  $Y$ . Similarly,  $\mathbb{E}_{Y|W}[\cdot]$  represents the conditional expectation with respect to  $Y$  given  $W$ .

The first term in the recurrence relation, Equation (7), corresponds to the expected costs incurred in the current period. The second term corresponds to the expected minimum costs incurred from the start of the subsequent stage,  $t + 1$ , until the end of the planning horizon. Equation (8) is the boundary condition. It assigns the salvage value to a facility based on its expected condition at the end of the planning horizon. An optimal policy can be obtained by using the recurrence relation to evaluate the optimal objective value function for every stage-state combination in reverse chronological order ( $T$  to 1). The minimum expected costs for the problem are given by evaluating the optimal objective value function for a given, initial set of information,  $I_1$ .

There are two significant practical obstacles with solving the above problem (as posed). The first one is that the state-space of the problem,  $\bigcup_{t=1}^T \mathcal{I}_t$ , with the unattractive feature that the number of states per stage grows exponentially with the index  $t$ . The second problem has to do with specifying and updating probability distribution functions to evaluate the expectations in the above expressions. We consider a Bayesian framework where the distributions are taken to represent a decision-maker's beliefs regarding deterioration and regarding the precision and accuracy of the available technologies. Inspections, in this context, reveal information about the current condition of a facility. This, in turn, provides information about the measurement errors associated with a given technology, as well as information about the deterioration process. Our approach to address the aforementioned computational difficulties is described in the following section.

### **3.5 A Quasi-Bayes Approach for the Estimation of Finite Mixtures of MDPs with Imperfect Information**

The problem of estimating a mixture of a finite set of possible models has recently received a great deal of attention in the statistics and econometrics literature. An extensive and recent review of this literature appears in McLachlan and Peel (1997). The control strategy proposed in this paper builds on the Quasi-Bayes estimation procedure proposed by Smith and Makov (1980). A critical difficulty in estimating finite mixtures is that no reproducing (natural conjugate) densities exist. As a result, the formal Bayes estimation procedure has sometimes been regarded as of little practical use in this context, due to the complex form of the basic likelihood. Below, we explain this point further as we consider an approximation solution approach for the problem.

The Quasi-Bayes approach to estimate finite mixtures introduced by Smith and Makov (1980)

considers the formal Bayesian solution in the particular situation where the prior density,  $f(\vec{\lambda})$ , for the unknown parameter  $\vec{\lambda}$  is taken to have the form of a Dirichlet density<sup>4</sup>:

$$f(\vec{\lambda}) = \frac{\Gamma(\alpha_1^{(1)} + \alpha_2^{(1)} + \dots + \alpha_R^{(1)})}{\Gamma(\alpha_1^{(1)})\Gamma(\alpha_2^{(1)})\dots\Gamma(\alpha_R^{(1)})} \lambda_1^{\alpha_1^{(1)}-1} \lambda_2^{\alpha_2^{(1)}-1} \dots \lambda_R^{\alpha_R^{(1)}-1} \quad (9)$$

which we denote  $D(\vec{\lambda}; \alpha_1^{(1)}, \alpha_2^{(1)}, \dots, \alpha_R^{(1)})$ , with  $\alpha_r^{(1)} > 0, \forall r$ . Recalling that the Dirichlet distribution is the conjugate prior for the parameters of the multinomial distribution allows us to update the probability density as follows when condition data are collected at the start of stage 2:

$$f(\vec{\lambda}|I_2) = \mathbb{E}_{Z_1} [f(\vec{\lambda}|I_2, Z_1)] \quad (10)$$

$$f(\vec{\lambda}|I_2) = w_1^2 D(\vec{\lambda}; \alpha_1^{(1)} + 1, \alpha_2^{(1)}, \dots, \alpha_R^{(1)}) + w_2^2 D(\vec{\lambda}; \alpha_1^{(1)}, \alpha_2^{(1)} + 1, \alpha_3^{(1)}, \dots, \alpha_R^{(1)}) + \dots + w_R^2 D(\vec{\lambda}; \alpha_1^{(1)}, \alpha_2^{(1)}, \dots, \alpha_R^{(1)} + 1) \quad (11)$$

where  $w_r^2 \equiv P(Z_1 = r|I_2)$ .  $Z_1$  is a random variable defined over  $\mathcal{R}$ , where  $Z_1 = r$  means that the first transition was generated by model  $r$ . Note that  $Z_1$  is a latent variable that follows a multinomial distribution. We derive an expression to compute  $w_r^2, \forall r \in \mathcal{R}$  in Appendix A. It is easily seen that, in general, the density function  $f(\vec{\lambda}|I_t)$  builds up as weighted averages of Dirichlet densities. To avoid the computational problems, it is natural to consider approximating the subsequent posterior densities by suitable Dirichlet densities. Having just observed  $\hat{X}_2$  consider the first step. If we were informed of the true model generating the measurement, i.e., if we were able to observe  $Z_1$ , the distribution of  $\vec{\lambda}$  would be independent of  $\hat{X}_2$  and would be given by:

$$D(\vec{\lambda}; \alpha_1^{(1)} + \Delta_{11}, \alpha_2^{(1)} + \Delta_{12}, \dots, \alpha_R^{(1)} + \Delta_{1R}) \quad (12)$$

where  $\Delta_{1r} = 1$  if  $Z_1 = r$  and 0 otherwise. Since we are not informed of the true source, we need to estimate  $\Delta_{1r}$ . In the Quasi-Bayes Learning procedure,  $\Delta_{1r}$  in the above expression are replaced by their expectations,  $w_r^1, \forall r$ .  $P(\vec{\lambda}|I_2)$  is thus approximated by  $D(\vec{\lambda}; \alpha_1^{(1)} + w_1^2, \alpha_2^{(1)} + w_2^2, \dots, \alpha_R^{(1)} + w_R^2)$ . Because the posterior distribution has the same form as the prior, the procedure can easily be repeated as additional data are collected. The complete iterative procedure is summarized below:

---

<sup>4</sup>We assume that the prior density is specified at the start of the first stage, after  $\hat{X}_1$  is collected.

$$\hat{\lambda}_r^{t-1} = \frac{\alpha_r^{(t-1)}}{\sum_{s \in \mathcal{R}} \alpha_s^{(t-1)}} \quad (13)$$

$$w_r^t = \frac{\sum_{j \in \mathcal{S}} \epsilon_{jk}^{C_t} \sum_{i \in \mathcal{S}} \pi_{ij}^r(A_{t-1}) P(X_{t-1} = i | I_{t-1}) \hat{\lambda}_r^{t-1}}{\sum_{s \in \mathcal{R}} \sum_{j \in \mathcal{S}} \epsilon_{jk}^{C_t} \sum_{i \in \mathcal{S}} \pi_{ij}^s(A_{t-1}) P(X_{t-1} = i | I_{t-1}) \hat{\lambda}_s^{t-1}} \quad (14)$$

$$\alpha_r^{(t)} = \alpha_r^{(t-1)} + w_r^t \quad (15)$$

Notice that the Equation (14) requires the conditional state distribution. Each component of the associated probability mass function is obtained as follows:

$$\begin{aligned} P(X_t = j | I_t) &= \sum_{r \in \mathcal{R}} P(X_t = j | Y = r, I_t) P(Y = r | I_t) \\ &= \sum_{r \in \mathcal{R}} \hat{\lambda}_r^t \frac{P(X_t = j | Y = r, I_{t-1}, A_{t-1}, C_t) \epsilon_{jk}^{C_t}}{\sum_{l \in \mathcal{S}} P(X_t = l | Y = r, I_{t-1}, A_{t-1}, C_t) \epsilon_{lk}^{C_t}} \\ &= \dots \\ &= \sum_{r \in \mathcal{R}} \hat{\lambda}_r^t \frac{\sum_{i \in \mathcal{S}} \pi_{ij}^r(A_{t-1}) P(X_{t-1} = i | I_{t-1}) \epsilon_{jk}^{C_t}}{\sum_{l \in \mathcal{S}} \sum_{i \in \mathcal{S}} \pi_{il}^r(A_{t-1}) P(X_{t-1} = i | I_{t-1}) \epsilon_{lk}^{C_t}} \end{aligned} \quad (16)$$

Smith and Makov (1980) show that any sequence of estimates,  $\hat{\lambda}^t, t = 1, 2, \dots$ , generated by the above procedure converges to the maximum likelihood estimates of the mixture proportions. This constitutes a significant asymptotic performance guarantee from the point of view of classical statistics.

Equations (13) through (16) describe how the information available to a decision-maker evolves over time, i.e., they describe what are referred to as the state transitions in the dynamic programming formulation presented in the previous section. The above development also reveals that the parameters of the Dirichlet distribution,  $\alpha_r^{(t)}, \forall r \in \mathcal{R}$ , and the conditional state distribution,  $P(X_t = i | I_t), \forall i \in \mathcal{S}$  are needed in order to execute the above algorithm at time  $t+1$ . We refer to the above parameters as a sufficient statistic for the information set  $I_t$ , i.e., they contain the essential information contained in  $I_t$ . For technical reasons it is desirable to replace the set of  $\alpha$ 's with the  $\lambda$ 's in the sufficient statistic. Appendix A shows how the state transitions can be carried out without relying on the the  $\alpha$ 's (except for the initial set). Thus, we have that, for the purpose of solving the dynamic program presented earlier,  $I_t$  is equivalent to  $(\hat{\lambda}_r^t, \forall r \in \mathcal{R}; P(X_t = i | I_t), \forall i \in \mathcal{S})$ , which we denote  $(\hat{\lambda}^t; P_t | I_t)$ . This, in turn, allows us to rewrite the recurrence relation (Equation (7)) as

follows:

$$v_t(\hat{\lambda}^t; P_t | I_t) = \min_{A_t \in \mathcal{A}, C_{t+1} \in \mathcal{C}} \left\{ \sum_{i \in \mathcal{S}} P(X_t = i | I_t) \left[ g(i, A_t, C_{t+1}) + \delta \sum_{r \in \mathcal{R}} \hat{\lambda}_r^t \sum_{j \in \mathcal{S}} \pi_{ij}^r(A_t) \sum_{k \in \mathcal{S}} \epsilon_{jk}^{C_{t+1}} v_{t+1}(\hat{\lambda}^{t+1}; P_{t+1} | I_{t+1}) \right] \right\} \quad (17)$$

## 4 Case Study: An Application of the Methodology to Pavement Management

We present a set of numerical examples to illustrate the implementation of the methodology presented herein. The objective is to develop insight about the effect of performance model uncertainty (and bias) on the optimal life-cycle costs of managing transportation facilities. We also evaluate the potential for savings from allowing agencies to adopt flexible inspection schedules. In particular, we consider instances of managing pavement over a planning horizon of 15 years and a discount rate  $\rho = 5\%$ , where  $\delta = 1/(1 + \rho)$ . As in Carnahan et al. (1987), we assume that pavement condition is represented by eight states, each representing 12.5 points on the PCI scale of 100. In our examples we set the initial pavement condition to state 5. The agency can choose from the following M&R actions: (1) do-nothing, (2) routine maintenance, (3) 1-in overlay, (4) 2-in overlay, (5) 4-in overlay, (6) 6-in overlay, and (7) reconstruction. Three possible deterioration models are considered: (1) slow, (2) medium, and (3) fast. Each model is characterized by a set of seven transition probability matrices (one for each action). The models are taken from Durango and Madanat (2002) and are such that:

- The effect of M&R actions on transitions is assumed to follow a truncated normal distribution with the mean depending on the action and the model and the variance depending on the model;
- Actions are less effective in improving pavement condition under faster deterioration models; and
- Faster deterioration models have higher variance in forecasting.

The means and standard deviations of the effects of actions are presented in Table 1. The transition probabilities are presented in Durango (2002).

Table 1: Means and standard deviations of action effects on change in pavement condition

	<b>Deterioration Model:</b>		
	<b>Slow</b>	<b>Medium</b>	<b>Fast</b>
<b>Std. Dev.</b>	0.30	0.50	0.70
<b>Action</b>	<b>Mean Effects</b>		
1	-0.25	-0.75	-1.75
2	0.50	0.00	-0.50
3	1.75	1.00	0.25
4	3.00	2.00	1.00
5	4.25	3.00	1.75
6	5.50	4.00	2.50
7	8.00	6.00	4.00

If an inspection is performed, the agency is said to have “perfect state information”. That is, we ignore measurement error in this study, i.e.,  $\epsilon_{kj}^1 = \{1 \text{ if } k = j; 0 \text{ otherwise}\}$ . This assumption was made to reduce the number of parameters and simplify the interpretation of the results. The general case, with inspection technologies of different precisions, is considered in the computational studies presented in Madanat (1993) and Madanat and Ben-Akiva (1994). The close relationship between our model and the Latent MDP formulations used in those studies leads us to believe that the insights extend to our model.

The costs include inspection costs, user costs, and costs of applying M&R actions. We restrict the facility condition to the “non-failed” states by setting the user cost of a failed pavement (state 1) to infinity. In order to prevent the facility from deteriorating too far at the end of the planning horizon we set the salvage costs of pavements in states 1 – 4 to infinity. Table 2 summarizes the costs considered which are taken from Carnahan et al. (1987). The inspection cost is \$0.065/lane-yard and is taken from Madanat and Ben-Akiva (1994).

Prior to presenting the results of the study, we note that numerical solutions for the dynamic program require a finite grid approximation to the continuous state-space. In our implementation, we discretize the  $[0, 1]$  interval into 101 points.

Table 2: Costs (\$/lane-yard)

Pav. State	Maintenance & Repair Actions							User Costs
	1	2	3	4	5	6	7	
1	0.00	6.90	19.90	21.81	25.61	29.42	25.97	$\infty$
2	0.00	2.00	10.40	12.31	16.11	19.92	25.97	25.00
3	0.00	1.40	8.78	10.69	14.49	18.30	25.97	22.00
4	0.00	0.83	7.15	9.06	12.86	16.67	25.97	14.00
5	0.00	0.65	4.73	6.64	10.43	14.25	25.97	8.00
6	0.00	0.31	2.20	4.11	7.91	11.72	25.97	4.00
7	0.00	0.15	2.00	3.91	7.71	11.52	25.97	2.00
8	0.00	0.04	1.90	3.81	7.61	11.42	25.97	0.00

### 4.1 Computational Study and Results

Figure 1 (2) compares the expected costs when the physical process corresponds to the slow (fast) model. The initial information set is consistent with the initial pavement condition, i.e.,  $P_1|I_1 = (0, 0, 0, 0, 1, 0, 0, 0)$ . For the case of “slow” initial beliefs, we set the belief vector to  $\hat{\lambda}^1 = (0.8, 0.1, 0.1)$ .<sup>5</sup> That is, a probability of 0.8 is assigned to the event that the physical process is governed by the slow model, 0.1 to the medium, and 0.1 to the fast model. Similarly, for “fast” initial beliefs we set the vector such that:  $\hat{\lambda}^1 = (0.1, 0.1, 0.8)$ . We also consider an initial belief vector that corresponds to a case of high model uncertainty. This case is labeled “no” which stands for the non-informative initial beliefs  $\hat{\lambda}^1 = (0.33, 0.33, 0.33)$ .

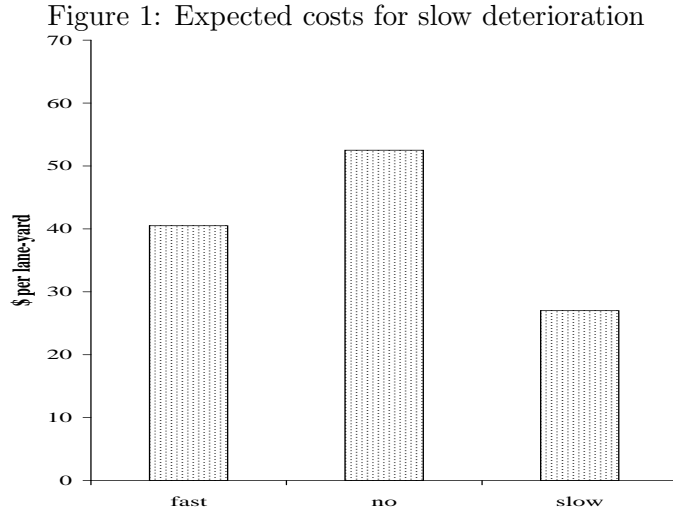
As expected, in both instances, whether the deterioration is slow or fast, when the initial beliefs are close to the physical process, the expected cost is the smallest. The expected costs are higher in Figure 2 than in Figure 1 because it is costlier to maintain a pavement that deteriorates faster. An interesting result is that the non-informative initial beliefs are the worst in both instances. This result indicates that inaccurate beliefs with low variance are preferred to less biased beliefs of higher variance.

To understand this surprising result, we conducted a simulation study whose results are presented in Figure 3. Instead of computing the expected costs analytically we generate instances for the case where the physical deterioration process corresponds to the fast model. The beliefs assigned to the fast model in each period,  $\hat{\lambda}_3^t = P(Y = 3|I_t)$ , are averaged over one thousand sim-

---

<sup>5</sup>In all cases we set  $\alpha^{(1)} = \hat{\lambda}^1$ .






ulation runs. We plot the trajectory of the average over time.

The average  $\hat{\lambda}_3^t$  converges much faster when the initial beliefs are biased (slow), than when they are non-informative, and therefore the actions taken in the non-informative situation are not as efficient as those taken when the initial beliefs are wrong. Hence the higher expected cost when the initial beliefs about deterioration have a higher variance attached to them. The faster convergence of the beliefs in the biased case compared to the non-informative case can be explained qualitatively by the contrast between the observations and the expectations. This contrast is augmented by the action taken in both cases: when the initial beliefs are biased, the M&R actions taken will be mild compared to the non-informative case. Therefore, worse states are more likely to be observed. Such unexpected outcomes provide feedback that leads to drastic and prompt revision of the beliefs in the biased case. This result may not generalize to all situations as it depends on the parameters used in the study (which were chosen from the pavement management literature). Characterizing (analytically or by simulation) the conditions (on the parameters) that are sufficient for this result to hold seems like an interesting direction for future study.

Finally, we compare the formulation we introduce to the *closed-loop* control formulation with fixed (yearly) inspections presented in Durango and Madanat (2002). The results are presented in Figures 4 and 5 below:

Figure 2: Expected costs for fast deterioration




We observe that relaxing the constraint of annual inspections leads to a reduction in expected costs. This is due to the fact that an inspection is performed only when it provides information that will improve future decisions. As a result, the expected number of inspections is reduced. A noteworthy feature of both figures is that the larger reduction in expected cost is observed when the initial beliefs are adequate. This indicates that the benefit of a flexible inspection schedule is greater when inspections provide less information, which is an intuitive result.

## 5 Conclusions and Discussion

In this paper we have presented an adaptive optimization model for the problem of finding joint inspection and maintenance policies for infrastructure facilities. The model simultaneously relaxes the assumption of a fixed inspection schedule and accounts for uncertainties both in the choice or specification of a performance model (to represent deterioration), and in the process of measuring facility condition. The methodology we present is referred to as adaptive because the information from measurements of condition is used to obtain an adequate representation of a facility's physical deterioration process over time, i.e., to learn about deterioration. In addition, the formulation captures the value in inspecting facilities to assess current condition or to reduce the measurement error associated with the technology used for inspections. The contribution of the paper is that it integrates two frameworks that address decision-making in infrastructure management: latent Markov decision process formulations and adaptive control formulations.

Figure 3: Probability assigned to fast model




A case study in pavement management leads to several insights about the problem. The results show that reducing the initial variance in model uncertainty may be more important than reducing the initial bias. This means that providing the wrong information may be less costly than providing no information about deterioration. The reason for this result is that the beliefs about deterioration can be adjusted drastically and quickly in response to unexpected events. Another result is that substantial benefits can be achieved by implementing a flexible inspection schedule when the initial beliefs are adequate. The reason is that inspections are providing very little information and so it is not necessary to perform them as frequently.

The scope of this research was purposely limited to the facility-level of the M&R problem. An immediate extension is to adapt the formulation to the network-level problem with administrative restrictions. A possible approach to incorporate network-level constraints is to formulate the model developed herein using randomized policies and to solve it using linear programming.

A significant practical advantage of an adaptive approach such as that presented in this paper is that it allows public agencies to start managing their infrastructure systems without having to develop detailed deterioration models of their facilities. Agencies can start managing their facilities with a set of existing models, possibly those developed by other state or national highway agencies, and let the adaptive optimization method described in this paper select the best mixture

Figure 4: Fixed vs. flexible inspection schedule: slow deterioration



of these models. This model refinement is performed over time, simultaneously with the selection of maintenance and repair decisions. This practical advantage is the reason why the most widely used infrastructure management systems in use today such as Pontis (Golabi and Shepard, 1997) are based on adaptive optimization.


## A Derivations

We derive expressions for used throughout the paper.

### A.1 Derivation of $w_r^2$

$$\begin{aligned}
 w_r^2 &= P(Z_1 = r|I_2) = P\left(Z_1 = r|\hat{X}_1, A_1, C_2, \hat{X}_2\right) \\
 &= \int_{\mathcal{S}} P\left(Z_1 = r|\vec{\lambda}, \hat{X}_1, A_1, C_2, \hat{X}_2\right) f\left(\vec{\lambda}|I_2\right) d\vec{\lambda} \\
 &= \int_{\mathcal{S}} \frac{P\left(Z_1 = r, \hat{X}_2|\vec{\lambda}, \hat{X}_1, A_1, C_2\right)}{P\left(\hat{X}_2|\vec{\lambda}, \hat{X}_1, A_1, C_2\right)} f\left(\vec{\lambda}|I_2\right) d\vec{\lambda}
 \end{aligned} \tag{18}$$

Figure 5: Fixed vs. flexible inspection schedule: fast deterioration



From Bayes' Law,  $f(\vec{\lambda}|I_2)$  can be obtained from the prior,  $f(\vec{\lambda})$ , as follows:

$$f(\vec{\lambda}|I_2) = \frac{f(\vec{\lambda}) P(\hat{X}_2|\vec{\lambda}, \hat{X}_1, A_1, C_2)}{\int_{\mathcal{S}} f(\vec{\lambda}) P(\hat{X}_2|\vec{\lambda}, \hat{X}_1, A_1, C_2) d\vec{\lambda}} \quad (19)$$

Letting  $\hat{\lambda}_r^1 = \int_{\mathcal{S}} \lambda_r f(\vec{\lambda}) d\vec{\lambda}$ , notice that  $\hat{\lambda}_r^1 = \frac{\alpha_r^{(1)}}{\sum_{s \in \mathcal{R}} \alpha_s^{(1)}}$ . Now, replacing (19) into (18) we have the following:

$$\begin{aligned} w_r^2 &= \frac{\int_{\mathcal{S}} P(Z_1 = r, \hat{X}_2|\vec{\lambda}, \hat{X}_1, A_1, C_2) f(\vec{\lambda}) d\vec{\lambda}}{\int_{\mathcal{S}} P(\hat{X}_2|\vec{\lambda}, \hat{X}_1, A_1, C_2) f(\vec{\lambda}) d\vec{\lambda}} \\ &= \frac{\int_{\mathcal{S}} P(\hat{X}_2|Z_1 = r, \vec{\lambda}, \hat{X}_1, A_1, C_2) P(Z_1 = r|\vec{\lambda}, \hat{X}_1, A_1, C_2) f(\vec{\lambda}) d\vec{\lambda}}{\int_{\mathcal{S}} \sum_{s \in \mathcal{R}} P(\hat{X}_2|Z_1 = s, \vec{\lambda}, \hat{X}_1, A_1, C_2) P(Z_1 = s|\vec{\lambda}, \hat{X}_1, A_1, C_2) f(\vec{\lambda}) d\vec{\lambda}} \\ &= \frac{P(\hat{X}_2|Z_1 = r, \hat{\lambda}^1, \hat{X}_1, A_1, C_2) \hat{\lambda}_r^1}{\sum_{s \in \mathcal{R}} P(\hat{X}_2|Z_1 = s, \hat{\lambda}^1, \hat{X}_1, A_1, C_2) \hat{\lambda}_s^1} \end{aligned} \quad (20)$$

where  $P(\hat{X}_2 = k|Z_1 = r, \hat{\lambda}^1, \hat{X}_1, A_1, C_2)$

$$\begin{aligned} &= \sum_{j \in \mathcal{S}} P(\hat{X}_2 = k|X_2 = j, Z_1 = r, \hat{\lambda}^1, \hat{X}_1, A_1, C_2) \cdot P(X_2 = j|Z_1 = r, \hat{\lambda}^1, \hat{X}_1, A_1, C_2) \\ &= \sum_{j \in \mathcal{S}} \epsilon_{jk}^{C_2} \sum_{i \in \mathcal{S}} P(X_2 = j|Z_1 = r, X_1 = i, A_1) \cdot P(X_1 = i|I_1) \\ &= \sum_{j \in \mathcal{S}} \epsilon_{jk}^{C_2} \sum_{i \in \mathcal{S}} \pi_{ij}^r(A_1) P(X_1 = i|I_1) \end{aligned}$$

Thus, given  $\hat{X}_2 = k$ ,  $w_r^2$  can be written as:

$$\frac{\sum_{j \in \mathcal{S}} \epsilon_{jk}^{C_2} \sum_{i \in \mathcal{S}} \pi_{ij}^r(A_1) P(X_1 = i | I_1) \hat{\lambda}_r^1}{\sum_{s \in \mathcal{R}} \sum_{j \in \mathcal{S}} \epsilon_{jk}^{C_2} \sum_{i \in \mathcal{S}} \pi_{ij}^s(A_1) P(X_1 = i | I_1) \hat{\lambda}_s^1} \quad (21)$$

## A.2 On the Equivalence of Sufficient Statistics

We show that, in the Quasi-Bayes approach, the sequence of  $\hat{\lambda}$ 's can be generated without computing the  $\alpha$ 's.

$$\begin{aligned} \hat{\lambda}_r^t &= \frac{\alpha_r^{(t)}}{\sum_{s \in \mathcal{R}} \alpha_s^{(t)}} \\ &= \frac{\alpha_r^{(t-1)} + w_r^t}{\sum_{s \in \mathcal{R}} \alpha_s^{(t-1)} + w_s^t} \\ &= \frac{\alpha_r^{(t-1)}}{1 + \sum_{s \in \mathcal{R}} \alpha_s^{(t-1)}} + \frac{w_r^t}{1 + \sum_{s \in \mathcal{R}} \alpha_s^{(t-1)}} \\ &= \hat{\lambda}_r^{t-1} - \frac{\alpha_r^{(t-1)}}{\left(1 + \sum_{s \in \mathcal{R}} \alpha_s^{(t-1)}\right) \sum_{s \in \mathcal{R}} \alpha_s^{(t-1)}} + \frac{w_r^t}{1 + \sum_{s \in \mathcal{R}} \alpha_s^{(t-1)}} \\ &= \hat{\lambda}_r^{t-1} + \frac{1}{1 + \sum_{s \in \mathcal{R}} \alpha_s^{(t-1)}} \left[ w_r^t - \hat{\lambda}_r^{t-1} \right] \\ &= \hat{\lambda}_r^{t-1} + \frac{1}{t + \sum_{s \in \mathcal{R}} \alpha_s^{(1)}} \left[ w_r^t - \hat{\lambda}_r^{t-1} \right] \end{aligned}$$

We see from the last expression that the sequence of  $\lambda$ 's can be generated without having to compute the  $\alpha$ 's, except for the initial set, which can be included in the prior set of information,  $I_1$ .

## References

- Carnahan, J., Davis, W., Shahin, M., Keane, P., Wu, M., 1987. Optimal maintenance decisions for pavement management. *ASCE Journal of Transportation Engineering* 113 (5), 554–573.
- Congressional Budget Office (CBO), 1999. Trends in public infrastructure spending. Washington, DC.
- Dekker, R., 1996. Applications of maintenance optimization models: A review and analysis. *Reliability Engineering and System Safety* 51, 229–240.

- Derman, C., 1962. On sequential decisions and Markov chains. *Management Science* 9 (1), 16–24.
- Durango, P., 2002. Adaptive optimization models for infrastructure management. Ph.D. thesis, University of California, Berkeley.
- Durango, P., Madanat, S., 2002. Optimal maintenance and repair policies in infrastructure management under uncertain facility deterioration rates: An adaptive control approach. *Transportation Research Part A* 36 (9), 763–778.
- United States General Accounting Office (GAO), 2001. U.S. infrastructure: Funding trends and federal agencies investment estimates. GAO-01-986T, Washington, DC.
- Gendreau, M., Soriano, P., 1998. Airport pavement management systems: An appraisal of existing methodologies. *Transportation Research Part A* 32 (3), 197–214.
- Golabi, K., Kulkarni, R., Way, G., 1982. A statewide pavement management system. *Interfaces* 12 (6), 5–21.
- Golabi, K., Pereira, P., 2003. An innovative pavement management and planning system for the road network of Portugal. *ASCE Journal of Infrastructure Systems* 9 (2), 75–80.
- Golabi, K., Shepard, R., 1997. Pontis: A system for maintenance optimization and improvement of u.s. bridge networks. *Interfaces* 27 (1).
- Guillaumot, V., Durango-Cohen, P., Madanat, S., 2003. Adaptive optimization of infrastructure maintenance and inspection decisions under performance model uncertainty. *ASCE Journal of Infrastructure Systems* 9 (4), 133–139.
- Harper, W., Majidzadeh, K., 1991. Use of expert opinion in two pavement management systems. *Transportation Research Record* 1311, 242–247.
- Humplick, F., 1992. Highway pavement distress evaluation: Modeling measurement error. *Transportation Research Part B* 26 (2), 135–154.
- Jiang, Y., Sinha, K., 1989. A dynamic optimization model for bridge management systems. *Transportation Research Record* (1211), 92–100.
- Klein, M., 1962. Inspection-maintenance-replacement schedules under Markovian deterioration. *Management Science* 9 (1), 25–32.

- Kuhn, K., Madanat, S., 2005. Model uncertainty and the management of a system of infrastructure facilities. *Transportation Research Part C* 13 (4), 391–404.
- Kumar, P., Varaiya, P., 1986. *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Prentice Hall, Inc., Englewood Cliffs, NJ.
- Li, Y., Madanat, S., 2002. A steady-state solution for the optimal pavement resurfacing problem. *Transportation Research Part A* 36 (5), 347–365.
- Madanat, S., 1993. Optimal infrastructure management decisions under uncertainty. *Transportation Research Part C* 1 (1), 77–88.
- Madanat, S., Ben-Akiva, M., 1994. Optimal inspection and repair policies for infrastructure facilities. *Transportation Science* 28 (1), 55–61.
- Madanat, S., Karlaftis, M., McCarthy, P., 1997. Probabilistic infrastructure deterioration models with panel data. *ASCE Journal of Infrastructure Systems* 3 (1), 4–9.
- Madanat, S., Park, S., Kuhn, K., 2006. Adaptive optimization and systematic probing of infrastructure system maintenance policies under model uncertainty. *ASCE Journal of Infrastructure Systems* 12 (3), 192–198.
- Madanat, S., Wan Ibrahim, W., 1995. Poisson regression models of infrastructure transition probabilities. *ASCE Journal of Transportation Engineering* 121 (3), 267–272.
- McCall, J., 1965. Maintenance policies for stochastically failing equipment: A survey. *Management Science* 11 (1), 493–524.
- McLachlan, G., Peel, D., 1997. *Finite Mixture Models*. Wiley, New York, NY.
- Mishalani, R., Madanat, S., 2002. Computation of infrastructure transition probabilities using stochastic duration models. *ASCE Journal of Infrastructure Systems* 8 (4), 139–148.
- Ouyang, Y., Madanat, S., 2004. Optimal scheduling of rehabilitation activities for multiple pavement facilities: Exact and approximate solutions. *Transportation Research Part A* 38 (5), 347–365.
- Pierskalla, W., Voelker, J., 1976. Survey of maintenance models: The control and surveillance of deteriorating systems. *Naval Research Logistics Quarterly* 23 (3), 353–388.



- Smith, A., Makov, U., 1980. A quasi-Bayes sequential procedure for mixtures. *Journal of the Royal Statistical Society B* 40 (1), 106–111.
- Suzuki, Y., Pautsch, G., 2005. A vehicle replacement policy for motor carriers in an unsteady economy. *Transportation Research Part A* 39 (5), 463–480.
- Thompson, P., Neumann, L., Miettinen, M., Talvitie, A., 1987. A micro-computer Markov dynamic programming system for pavement management in Finland. *Proceedings of the Second North American Conference on Managing Pavements Toronto, Canada.*
- Tsunokawa, K., Schofer, J., 1994. Trend curve optimal control model for highway pavement maintenance: Case study and evaluation. *Transportation Research Part A* 28 (2), 151–166.
- Ullidtz, P., 1987. A danish pavement management system. *Proceedings of the Second North American Conference on Managing Pavements Toronto, Canada.*
- Van Noortwijk, J., Frangopol, D., 2004. Deterioration and maintenance models for insuring safety of civil infrastructures at lowest life-cycle cost. In: Frangopol, D. M., Bruhwiler, E., Faber, M. H., Adey, B. (Eds.), *Life-Cycle Performance of Deteriorating Structures: Assessment, Design and Management*. ASCE, Reston, Virginia, pp. 384–391.
- Van Noortwijk, J., Van Gelder, P., 1996. Optimal maintenance decisions for berm breakwaters. *Structural Safety* 18 (4), 419–446.
- Van Winden, C., Dekker, R., 1998. Rationalisation of building maintenance by Markov decision models: A pilot case study. *Journal of the Operational Research Society* 49 (3), 293–309.
- Wirahadikusumah, R., 1999. Optimization modeling for management of large combined sewer networks. Ph.D. thesis, Purdue University.