

UC Berkeley

Earlier Faculty Research

Title

Uncertainty in Infrastructure Deterioration Modeling and Robust Maintenance Policies for Markovian Management Systems

Permalink

<https://escholarship.org/uc/item/9pq1w5d3>

Author

Kuhn, Kenneth David

Publication Date

2006

**Uncertainty in Infrastructure Deterioration Modeling and Robust
Maintenance Policies for Markovian Management Systems**

by

Kenneth David Kuhn

B.A. (Johns Hopkins University) 2001
M.S. (University of California, Berkeley) 2002

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering-Civil and Environmental Engineering

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Samer Madanat, Chair
Professor Laurent El Ghaoui
Professor Raja Sengupta

Fall 2006

UMI Number: 3253936

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform 3253936

Copyright 2007 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

**Uncertainty in Infrastructure Deterioration Modeling and Robust
Maintenance Policies for Markovian Management Systems**

Copyright 2006

by

Kenneth David Kuhn

Abstract

Uncertainty in Infrastructure Deterioration Modeling and Robust Maintenance
Policies for Markovian Management Systems

by

Kenneth David Kuhn

Doctor of Philosophy in Engineering-Civil and Environmental Engineering

University of California, Berkeley

Professor Samer Madanat, Chair

Infrastructure management systems help public works agencies decide when and how to maintain, repair, and rehabilitate infrastructure facilities in a cost effective manner. An integral part of an infrastructure management system is a model describing how the different infrastructure facilities to be managed deteriorate over time and with use. Many sources of error limit the ability of management systems to accurately predict how built systems will deteriorate. This dissertation introduces and examines different techniques for considering error and uncertainty in deterioration modeling within an infrastructure management system.

Techniques used include robust optimization and adaptive control. In the context of robust optimization, both MAXIMIN and Hurwicz decision criteria are studied. Computational studies involving the simulated management of pavement systems illustrate the

strengths and weaknesses of the proposed approaches. These studies involve short-term, limited horizon planning, as well as indefinite-term, infinite horizon planning. Single facility infrastructure management problems are presented alongside more complex problems involving the management of a network of an arbitrarily large number of related facilities.

It is found that both robust optimization and adaptive control formulations have certain comparative advantages. Some discussion is included of the possibility of combining the robust and adaptive frameworks to create a new hybrid approach. Regardless of what approach is used, this work makes clear that consideration of uncertainty in deterioration modeling during decision-making can alter 'optimal' maintenance strategies selected and change the potential user and agency costs of infrastructure management.

Professor Samer Madanat
Dissertation Committee Chair

To my parents,
Steven and Betsey Kuhn,
for their love and support.

Contents

List of Figures	iv
List of Tables	vi
1 Introduction	1
1.1 Outline of the Dissertation	2
2 Infrastructure Management	5
2.1 Maintenance, Repair, and Rehabilitation (MR&R) Decision-Making	7
2.1.1 Single Facility Markov Decision Problem (MDP) Formulation	10
2.1.2 System Level MDP Formulation	13
2.1.3 Conclusion	17
3 Infrastructure Deterioration Modeling	18
3.1 Survival Analysis	21
3.2 The Weibull Distribution	24
3.3 The Curse of Uncertainty	30
3.3.1 Parametric vs. Epistemic Uncertainty	30
3.3.2 Quantifying Uncertainty	31
3.4 Conclusion	38
4 Robust Infrastructure Management: At the Single Facility Level	39
4.1 Introduction to Robust Optimization	40
4.2 The Maximin Criterion	41
4.3 The Hurwicz Criterion	45
4.4 Computational Studies	49
4.4.1 Problem Specifications	50
4.4.2 Results	53
4.5 Conclusion	61

5	Robust Infrastructure Management: At the System Level	62
5.1	The Maximin Criterion	63
5.1.1	Solution Strategy	64
5.2	The Hurwicz Criterion	65
5.3	Computational Studies	68
5.3.1	Results	69
5.4	Conclusion	77
6	Comparing Adaptive and Robust Infrastructure Management	79
6.1	Adaptive Control in Infrastructure Management	79
6.1.1	An Example Adaptive Control Approach	81
6.1.2	Systematic Probing	83
6.2	Computational Studies	84
6.2.1	Results	86
6.3	Qualitative Comparisons	92
6.4	Towards A Hybrid Management System	94
6.4.1	An Example Hybrid Approach	94
7	Conclusion	98
7.1	Future Directions	100
7.2	Final Thoughts	102

List of Figures

3.1	Deterioration model fit: Pdf of the time a bridge deck will spend in state 8	28
3.2	Deterioration model fit: Probability of a bridge deck remaining in state 8 for 1 year as a function of the time already spent in state 8	28
3.3	Deterioration model uncertainty: Histograms of estimated parameter values	34
3.4	Deterioration model uncertainty: 95% confidence interval pdf of time in state 8	36
3.5	Deterioration model uncertainty: 95% confidence interval of probability of remaining in state 8 for 1 year	37
4.1	Single facility robust optimization: Relative and absolute benefit of using maximin robust optimization in worst-case conditions	55
4.2	Single facility robust optimization: Cost ranges of non-robust asset management	56
4.3	Single facility robust optimization: Cost ranges of Maximin robust asset management	57
4.4	Single facility robust optimization: Cost ranges of Hurwicz (optimism = 0.5) robust asset management.	60
5.1	System level robust optimization: The complexity of Hurwicz criterion based robust optimization	67
5.2	System level robust optimization: Worst case costs of maximin robust and non-robust asset management	69
5.3	System level robust optimization: Expected case costs of maximin robust and non-robust asset management	73
5.4	System level robust optimization: Cost ranges of maximin robust and non-robust asset management	74
5.5	System level robust optimization: Cost ranges of hurwicz robust (for $\beta = 0.5$) and non-robust asset management	76
6.1	OLFC Adaptive control: Agency responsibilities over time.	82
6.2	Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = fast, assumed deterioration = slow	87

6.3	Adaptive vs. Robust: Cost ranges of robust management, deterioration = fast, assumed deterioration = slow	87
6.4	Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = slow, assumed deterioration = fast	88
6.5	Adaptive vs. Robust: Cost ranges of robust management, deterioration = slow, assumed deterioration = fast	88
6.6	Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = fast, assumed deterioration = medium	89
6.7	Adaptive vs. Robust: Cost ranges of robust management, deterioration = fast, assumed deterioration = medium	89
6.8	Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = slow, assumed deterioration = medium	90
6.9	Adaptive vs. Robust: Cost ranges of robust management, deterioration = slow, assumed deterioration = medium	90
6.10	Adaptive and robust infrastructure management: Flowchart of agency responsibilities over time.	96

List of Tables

3.1	Deterioration model fit: Bridge deck characteristics related to deterioration	25
3.2	Deterioration model fit: Estimated values of parameters related to bridge deck deterioration	27
3.3	Deterioration model uncertainty: Counts of bridge deck sample population by hw class and wear surface	33
3.4	Deterioration model uncertainty: Confidence intervals of model parameters	35
4.1	Agency costs (dollars per lane-yard) of performing different MR&R actions on pavement with different condition ratings.	50
4.2	User costs (dollars per lane-yard) associated with pavement characterized by different condition ratings.	51
4.3	Mean and standard deviation of effects of actions on the condition rating of a pavement section.	52
4.4	Single facility robust optimization: Maximin robust optimal actions when different uncertainty levels are considered	53
4.5	Single facility robust optimization: Worst-case future discounted costs of managing a like new pavement	54
4.6	Single facility robust optimization: Hurwicz optimal actions at different uncertainty and optimism levels	58
5.1	System level robust optimization: Fractional Variables of Condition State and Actions Taken	71

Acknowledgments

I would like to thank Professor Samer Madanat, for guiding me through this research and being a considerate advisor. I have looked forward to our regular afternoon coffee breaks and discussions about research, soccer, politics, and life. More than just my advisor, I consider Samer my mentor, friend, and (at times) role model. His door has always been open, despite his large and ever expanding list of responsibilities on campus.

I would also like to thank the other members of my committee, Professors Laurent El Ghaoui and Raja Sengupta, for their time and advice, which also helped shape this research effort. Additionally, I must acknowledge Professors Mike Cassidy and Elizabeth Deakin, who spoke with me at critical moments in my graduate school career and helped me find my place at Berkeley. If it wasn't for Professors Cassidy and Deakin, I'm not sure I would be graduating with a Ph.D. today. The same can be said for Professors Benjamin Hobbs and Chuck ReVelle of Johns Hopkins University. The late Professor ReVelle in particular was an invaluable asset to me at a crucial point in my life.

Outside of academia, I have to acknowledge the support of my family. My mother and father's unconditional love and support has given me considerable strength. They have always been there for me and have even given me detailed advice on how to navigate graduate school.

I'd finally like to acknowledge the friends that have supported me through my academic career. My best friends from high school, Erik Taylor and Wayland Lee, continue to harass me even though I have moved 3,000 miles away. Similarly I rely on my college buddy Minn Yang, and one day I'll get out to the other side of the world to see him. At

Berkeley, I have to acknowledge a host of people who have made my time at Berkeley among the happiest times in my life. Charles-Antoine Robelin, Kristine Bolt, Marco Zennaro, and many other classmates of mine have made me look forward to coming to campus. Ni Liu, Naseem Ehsan, David Lee, Cherry Chaicharn and many others have been good friends outside of school. Finally, I have to acknowledge my best friend Hiren Nisar. Hiren came to my apartment just after I arrived in Berkeley to watch tv and didn't leave for almost 4 years. He has been a source of color and happiness in my life.

Chapter 1

Introduction

The United States has historically made an extraordinary investment in its infrastructure. For instance, according to the GAO (2001) the federal government has spent an average of about \$59 billion annually since the 1980s on the nation's civilian infrastructure. The emphasis of infrastructure investment has shifted in the past 30 years toward maintenance rather than new construction. According to the CBO (1999), a larger and larger proportion is being spent on the maintenance of the total expenditure on public works improvements, with the proportion of public non-capital spending for infrastructure increasing from 39% in 1960 to 57% in 1994. However, the magnitude of Maintenance, Repair, and Rehabilitation (MR&R) investment has been far from sufficient. Therefore, the critical issue facing public works agencies today is how to allocate limited resources that are available for MR&R so as to obtain the best return for their expenditure.

The MR&R resource allocation problem has come to be known as the infrastructure management problem, and infrastructure management systems have been developed

to help public agencies make optimal decisions regarding MR&R expenditures. The recommendations of infrastructure management systems are useful only if the assumptions that underlie such tools are reasonably correct. In particular, such systems make assumptions relating to how facilities deteriorate that are key to the maintenance policies they recommend.

The research presented here explains how and why assumptions made by management systems relating to deterioration may be prone to uncertainty and error. A number of techniques are presented as alternatives that are well suited to minimizing the extra costs uncertainty may add to infrastructure maintenance. Several of these methodologies that fall under the heading of robust optimization are applied to infrastructure management for the first time here. The overall aims of the work presented here are to illustrate that uncertainty in deterioration modeling can be a problem in infrastructure management, and that there are different techniques for dealing with this uncertainty. No one technique is chosen as the ‘optimal’ technique for dealing with decision-making in the presence of uncertainty at the end of this text.

1.1 Outline of the Dissertation

Reading through a lengthy scientific discourse like this can be a challenge. The reader must focus on the technical details of the current section while recalling the underlying methodology of the current chapter and the overall aims of the entire text. It is easy to lose the plot, to lose track of the story being told or the characters doing the telling. An outline of this text is presented here. The hope is that this outline will allow the reader

to quickly and easily grasp the aims of the text as a whole and the motivations behind the various parts of the research.

This dissertation begins with two introductory chapter entitled *Infrastructure Management* and *Infrastructure Deterioration Modeling*. The research presented in this dissertation is in the field of infrastructure management, and relates to how built infrastructure facilities are maintained. This field will be introduced to the the reader in some detail. Infrastructure management is bifurcated into one area relating to statistical descriptions of infrastructure deterioration and another relating to maintenance decision-making given deterioration models. The final part of the introductory chapter describes uncertainty in deterioration modeling. This is the central ‘problem’ that motivates the research presented in this text.

The following two chapters describe a novel approach to dealing with the problem of uncertainty in deterioration modeling, bringing the consideration of this uncertainty into maintenance decision-making and making use of nascent robust optimization techniques. One chapter relates to the management of single infrastructure facilities, while the other considers management of a system of facilities. There are important methodological and computational differences between single facility and system level infrastructure management. The intuitive meaning of various robust decision-making criteria and the results of computational case studies involving simulated management of infrastructure facilities are emphasized.

The next chapter compares robust optimization based methodologies to adaptive control approaches. Adaptive control is the framework currently used in situations where

there is significant uncertainty in deterioration modeling. It is neither clearly inferior nor superior to robust optimization based approaches. Computational studies and qualitative comparisons reveal the relative strengths and weaknesses of the two approaches. Some discussion is presented describing how adaptive control and robust optimization could be combined to create a hybrid robust and adaptive methodology for infrastructure management.

The final chapter of this text summarizes the key findings of the research presented here. The results of the computational studies performed in various chapters are revisited in the context of the overall findings of the research effort. Directions for future research, of both a theoretical and a computational nature, are presented.

References to this outline have been added to the various chapters of this dissertation. It is hoped that this outline and the references to it make the research easier to follow.

Chapter 2

Infrastructure Management

Built infrastructure facilities, like power lines and paved roads, provide the foundations needed for modern society to function. These facilities deteriorate over time and, if nothing is done, may degrade to the point where they become unusable. Public agencies monitor infrastructure facilities and take action to slow or reverse deterioration processes through a process known as infrastructure management. The decisions made in infrastructure management include the selection and scheduling of Maintenance, Repair, and Rehabilitation (MR&R) actions. There are two, often competing, objectives of infrastructure management. The primary objective is to provide the best possible service to the users of the infrastructure in question. In the case of managing a single bridge, this might involve minimizing the chance that the bridge in question collapses. The secondary goal of infrastructure management is to minimize expenditures. In order to minimize the chance a bridge falls down, it may be optimal to rebuild the bridge every year. However, this policy forces government agencies to spend a tremendous amount of money on bridge reconstruction that

might be better spent elsewhere.

Infrastructure management systems are computer programs that help public agencies collect information and make decisions related to infrastructure deterioration (Ferreira et al, 2002). The dual responsibilities of information storage and decision-making go together well; in order to develop efficient MR&R policies for infrastructure facilities, it is crucial to have large amounts of data regarding the condition of these facilities. Data regarding past and current conditions, provided by facility inspection, are input into deterioration models to predict future conditions. Infrastructure management systems use deterioration models to estimate how effective different MR&R actions will be at upgrading the condition levels of facilities. The costs and benefits of different MR&R actions can then be compared and 'optimal' policies found to maximize future levels of service provided and/or minimize future management expenses.

Unfortunately, predicting the effects of performing MR&R activities cannot be done with complete accuracy. Available data is limited and some statistical uncertainty surrounds the parameters of deterioration models. Furthermore, there exists additional uncertainty related to the choice of deterioration models and the assumptions these models make to simplify complex real-world phenomena.

This chapter introduces the reader to the infrastructure management decision-making process, while the following chapter focuses on deterioration modeling to support decision-making, and the problem of uncertainty in deterioration modeling. Wherever possible detailed examples are provided. Two explicit mathematical formulations of infrastructure management problems are presented. In the next chapter, a specific, state-of-the-art

deterioration model is presented with some discussion of how the outputs of the deterioration model can be used in conjunction with the mathematical formulations of the management problem. Finally, an analysis of uncertainty in the example deterioration model is performed. It is hoped that this level of detail will inform discussions, and prepare the reader for the following chapters of this text.

2.1 Maintenance, Repair, and Rehabilitation (MR&R)

Decision-Making

Intuition will tell you that maintenance activities should be undertaken whenever deterioration reaches a certain point, i.e. repave the road when the potholes and cracks in the road become ‘a problem.’ The simplest decision rules follow this intuition and involve taking maintenance actions whenever measurements of deterioration processes reach pre-specified critical values. For instance, Flintsch and Zaniewski (1997) advocate requiring pavement preservation activities whenever measurements of roughness or cracking reach threshold values set using expert engineering judgement. Li and Madanat (2002) and Ouyang and Madanat (2004) have formulated and solved mathematical programs for optimal threshold values, the first relating to long-term (infinite-horizon) and the second to short-term infrastructure management. In addition to using critical or threshold values, Flintsch and Zaniewski (1997) and other authors also make use of mathematical programming techniques to find alternate strategies for considering the two objectives of infrastructure management, maximizing facility performance and minimizing maintenance costs.

The two competing objectives of infrastructure management lead to different for-

mulations of the problem. One methodology minimizes management expenditures while forcing public agencies to maintain infrastructure at pre-specified levels. This approach was used in the state of Arizona's Pavement Management System (PMS) (Golabi et al (1982); Wang et al (1994)). The framework is well suited to determining how much budget is necessary to maintain or upgrade the conditions of facilities. However, this approach provides no incentive for raising level of service above minimum standards. An alternate view involves maximizing level of service given a fixed management budget (Liu and Wang (1996); Wang and Liu (1997)). This approach is logical for public agencies with budgets fixed far into the future.

In different situations, different formulations may make more sense. In managing the roadways of an urban area, public agency budgets may be fixed for years to come and clear benefits (in terms of reduced vehicle wear and tear, improved gas mileage, etc.) may accrue to drivers if the quality of roads were maximized. On the other hand, if a cash-strapped public school has to maintain temporary classroom space, it may be best to minimize expenditures ensuring facilities are adequate until they are replaced. Generally speaking, the best approach allows MR&R expenditures and level of service considerations to be traded-off. This can be achieved by minimizing costs that include both agency costs and user costs related to the condition levels of facilities. This approach provides a framework where it is possible to find the 'social-optimal' point that minimizes total costs.

Some of the early infrastructure management system models suggested static management plans. An example of such a plan would be to seal up the cracks on a particular roadway surface in 2 years and then repave the roadway in 6 years. Golabi et al (1982) was

the first to recognize that a dynamic plan, responsive to the actual, realized formation of distresses, might be more appropriate. A dynamic formulation of the plan described above would be to seal up the cracks when severe cracking covers 10% of the roadway surface and repave the roadway when ruts in vehicle wheel paths are more than 3cm deep. Kulkarni (1984) showed that a more dynamic model could lower management costs while raising servicability standards.

Much work has been done in the area of dynamic decision-making for a network of facilities. One of the first (and most influential) papers in this area was Golabi et al (1982), which discussed the management of the pavement network in the state of Arizona. Carnahan et al (1987) formulated the general problem of managing pavement, treating each section of pavement independently. Chan et al (2003) added a layer of complexity, looking at the problem of allocating a fixed budget across a multi-district area with infrastructure management problems in each district in the area.

Extending the infrastructure management problem further, Madanat and Ben-Akiva (1994), Ellis et al (1995), and Jido et al (2004) have considered the costs and (implicit) benefits of inspecting facilities and included inspection frequency flexibility in their infrastructure management system models. The first two of these works dealt with infrastructure management problems where deterioration is represented by jumps between a set of states that summarize condition information. The third work represents the condition of an infrastructure facility via a continuous rating. Management systems that rely on continuous condition ratings typically assume deterministic models of deterioration, while state-spaced models use stochastic models. This distinction will be explored in more detail

in the deterioration modeling section of this work.

For now, let's look at some specific examples involving minimizing user + agency costs and utilizing stochastic, state-space based model of deterioration. Consider a public agency responsible for managing bridge decks. The Federal Highway Administration (FHWA) associates the 'overall condition' of a bridge deck with an integer from 0 to 9. A bridge deck in state 1 will have failed while a bridge deck in state 9 will be like-new FHWA (1979). The 10 point scale may seem limited, but note that the overall condition of a bridge deck is an artificial construct and cannot be measured precisely. The set of possible actions that can be taken on a bridge deck similarly consists of only a handful of options. New asphalt can be overlaid on top of existing bridge decks, decks can be patched, or even rebuilt. Decision makers study bridge deck condition ratings and decide what, if any, maintenance actions to take. It is often desirable to have a simple and transparent rule or policy translating bridge deck condition ratings into maintenance actions. Formulating the problem of deciding which maintenance activities to take as a Markov Decision Problem (MDP) allows for an 'optimal' version of exactly this type of policy.

2.1.1 Single Facility Markov Decision Problem (MDP) Formulation

Markov Decision Problems are ideal descriptions of situations involving limited, discrete sets of possible states of the world (condition states), as well as of actions to take. In addition, a solution to an MDP will consist of a simple map from condition state to action to take. MDP formulations of infrastructure management problems abound in the literature including Golabi et al (1982), Carnahan et al (1987), Gopal and Majidzadeh (1991) and Golabi and Shepard (1997), the last of which relates to the management of bridges.

In order to delve deeper into infrastructure management, it is necessary to study MDP formulations of infrastructure management problems in a bit more detail. Toward that end, let's examine an MDP formulation related to the management of an individual bridge deck.

First of all, let I be the set of condition rating states possible for this bridge deck. This might be the set of FHWA bridge deck condition ratings $\{0, 1, \dots, 9\}$. Similarly, let A be the set of actions that can be taken. This problem will be formulated as one of minimizing costs, both related to management expenditures and the level of service the facility provides. Say $u(i)$ is the user cost associated with the bridge deck being in state i , while $g(i, a)$ is the agency cost of performing maintenance action a on the deck when it is in condition rating state i . Performing maintenance involves incurring agency costs now to lower future user costs. To compare present and future costs, it is necessary to consider discounting. Let α be the discount amount factor we will use. ($\alpha = \frac{1}{1+r}$ where r is the discount rate.) Finally, it is necessary to describe the effects of performing MR&R activities. Let $p(j|i, a)$ be the conditional probability of the bridge deck being in condition rating state j a year after being in state i with action a applied. Given the above definitions, the following Dynamic Programming (DP) recursion holds:

Single Facility Infrastructure Management Problem (SFIMP)

$$v_t(i) = \min_{a_t(i) \in A} \left[g(i, a_t(i)) + u(i) + \alpha \left[\sum_{j \in I} p(j|i, a_t(i)) v_{t+1}(j) \right] \right] \quad (2.1)$$

where $v_t(i)$ is the least expected cost, from year t forward, of managing the deck starting in state i , and $a_t(i)$ is the optimal action to take in year t if the bridge deck is in condition state i .

It would be possible to manage the bridge deck for a planning horizon of T years, given terminal costs ($v_T(i)$ values). Simply work backwards starting at time $T - 1$. For each condition state, examine each action in turn and find the expected costs from year $T - 1$ through year T . In each case, note the action that minimizes expected costs ($a_{T-1}(i)$) and the minimum expected costs incurred ($v_{T-1}(i)$). Repeat this process for year $T - 2$ using the values for $v_{T-1}(i)$ and continue working backwards till reaching year 0, the current time. Optimal actions to take in all condition states possible, at all times, will have been noted.

It would also be possible to manage the facility in question for a long time, using ‘infinite horizon’ dynamic programming. One approach to solve infinite horizon problems is policy iteration: Choose an arbitrary initial management policy a that maps condition states to actions to take, and find the fixed point values for $v(i)$ terms such that $\forall i \in I$ we have $v(i) = g(i, a(i)) + u(i) + \alpha[\sum_{j \in I} p(j|i, a(i))v(j)]$. Next, select a new management policy a' such that $\forall i \in I, a'(i) = \arg \min_{a'(i)} g(i, a'(i)) + u(i) + \alpha[\sum_{j \in I} p(j|i, a'(i))v(j)]$. Using the new management policy a' , recalculate fixed point costs v . Continue along the same lines, until the management policy and costs remain constant from one iteration to the next. The optimal management policy and expected future costs result.

Note that although the discussion was framed in terms of managing a bridge deck, the above formulation was general enough to apply to any infrastructure facility. The specific costs and transition probabilities associated with any facility, or class of homogeneous facilities, could be input and the above mathematical recursion solved. If a heterogeneous set of facilities were to be managed, different instances of the above math program could be solved for each facility and an overall solution found that minimizes the total user +

agency costs across the system. However, this analysis makes the unrealistic assumption of a management budget able to cover the expenses of whatever maintenance actions are necessary to minimize total costs.

2.1.2 System Level MDP Formulation

The state of Arizona's Department of Transportation (ADOT) was reported, in 1994, to manage 12,000 km of pavement divided into 1,765 separate 'sections' (Flintsch and Zaniewski (1997)). The decisions of which actions to take on the separate facilities cannot be made independently, since the funds for all activities are constrained by a fixed budget. Every year, instead of selecting the one action to take on a single facility, it is necessary to select a separate action to take for each facility in the network. The space of possible solutions to the problem of what to do in one year has grown tremendously from the case of managing a single facility. Similarly, the space of possible conditions of all facilities to be managed has grown tremendously.

Note that managing a single facility for a horizon of T years involved going through and comparing all possible actions in all possible condition rating states each and every year / policy iteration. This becomes computationally very challenging if a large number of facilities have to be managed simultaneously. The procedures used to solve the single facility infrastructure management problem are forms of dynamic programming. It is well known that dynamic programming becomes increasingly difficult as the dimensions of a problem increase, the so-called 'curse of dimensionality.'

There are clever procedures for solving a MDP-type network level infrastructure management problem while avoiding the curse of dimensionality. A classic example is

included in the Pavement Management System (PMS) used by the Arizona Department of Transportation (ADOT) (Golabi et al (1982)). This system minimizes the lifecycle costs of the pavement sections managed by the ADOT using convex optimization techniques. Facilities are separated into groups based on their construction, environment, and current condition ratings. ADOT's PMS recommends MR&R actions for fractions of facilities in the different groups. Pavement managers are given leeway in determining which specific facilities to choose to make up the fractions selected by the PMS.

An example of a mathematical programming formulation of a system level infrastructure management problem is presented here. This example is built using the classic ADOT PMS as a template. In this example, it is assumed that there is a network of N facilities to be managed for a period of T years. To simplify the analyses, each of these N facilities is considered identical in the sense that each deteriorates in exactly the same manner. There will be some discussion later of how this assumption may be relaxed. Decisions have to be made regarding all facilities simultaneously because there is a management budget b_t that dictates how much can be spent on MR&R activities in year t .

For decision variables, let $x_{i,t}(a)$ be the fraction of all facilities, in state i in year t , that will have action a applied. Like the original ADOT PMS, the decision variable produces fractions of facilities to which MR&R actions should be applied. This formulation does not offer explicit facility by facility guidance. However, there are benefits to allowing engineers with immediate knowledge of the facilities in question to choose a repair plan out of a set considered equivalent by an asset management system. Furthermore, the formulation shown here has the crucial advantage of allowing the scope of the problem to be independent of

the number of facilities in the network.

Let $f_t(i)$ be the expected fraction of all facilities, in year t , that will be in condition rating state i . Assume at the start of the management exercise, the conditions of all facilities are known. It would be possible to construct $f_0(i)$, the starting fractions of all facilities in different states. Then, given a management policy x and transition probability matrix p , the expected fractions of facilities in different states one year into the management exercise and throughout the planning horizon can be computed. Thus f_0 is a parameter and f_t where $t > 0$ is fixed by x and p .

System Level Infrastructure Management Problem (SLIMP)

$$\min_x \sum_{t=0}^T \alpha^t \left[\sum_i \sum_a [g(i, a) + u(i)] f_t(i) x_{i,t}(a) N \right] \quad (2.2)$$

subject to the following constraints:

$$\begin{aligned} (1) \quad & \sum_i \sum_a g(i, a) f_t(i) x_{i,t}(a) N \leq b_t & \forall t \in \{0, 1, 2, \dots, T\} \\ (2) \quad & \sum_i \sum_a f_t(i) x_{i,t}(a) p(j|i, a) = f_{t+1}(j) & \forall j, t \in \{0, 1, 2, \dots, T\} \\ (3) \quad & x_{i,t}(a) \geq 0 & \forall i, a, t \in \{0, 1, 2, \dots, T\} \end{aligned}$$

The math program shown above minimizes the discounted sum of user and agency costs over the planning horizon. Constraint (1) demands that the expected agency spending does not run over the allowable budget in any given year. Constraint (2) is the key to how deterioration is considered in decision-making. Decisions made in a given year must be based on the expected conditions of facilities, which are obtained from the decisions made the previous year and the transition probability matrix. Constraint (2) says that at a time t the expected number of facilities in a state j is equal to the sum across all states of the expected number of facilities in the state at time $t - 1$ times the respective

probabilities of transitioning into state j . It is worth noting here that constraint (2) and the pre-specified f_0 , taken together, ensure that all facilities are accounted for and make a constraint like $\sum_{i \in I} \sum_{a \in A} f_t(i) = 1$ unnecessary. Finally, constraint (3) ensures negative fractions of facilities are never considered.

If a heterogeneous set of facilities were to be managed, the SLIMP presented above could be modified. Imagine facilities are separated into classes with each facility within a class considered to be homogeneous. Each class could then be associated with a different set of user and agency costs, transition probabilities, and decision variables. The objective function could be modified to sum costs across all facility classes. Similarly, the sum of agency costs across all facility classes could be kept below some maximum budget.

The SLIMP math program shown above is an example of an MDP formulated as a convex optimization program. The word convex is used because the constraints are convex functions of the decision variables, and we are minimizing a convex objective function. There are well known, efficient techniques for solving problems like this that avoid a ‘brute-force’ comparison of each and every potential solution. A review of convex optimization is beyond the scope of this text, but for now note that there is plenty of commercially available software that quickly and efficiently solves problems like the one shown above.

Criticism of MDP formulations, like the ones shown here, often revolve around the fact that they are “memoryless” or “history-independent.” In the MDP formulations here, all facilities in the same condition state are treated as identical. Perhaps not all bridge decks with the same FHWA rating face identical prospects with respect to future deterioration. Perhaps some benefit could be gained by considering more information, like

the *history* of FHWA ratings or MR&R actions taken for individual bridge decks. However, the definition of condition state can be generalized to include histories of FHWA ratings and MR&R actions taken. Through this process of ‘state augmentation,’ MDP formulations can incorporate historical data, though this leads to an increase in complexity.

The more serious restrictions of MDP formulations of infrastructure management problems are:

- that proper definitions have to be found for a discrete sets of condition states so that all facilities in the same condition state will be homogeneous with regards to deterioration over the next year / time step,
- that accurate and precise transition probabilities have to be found to describe the deterioration of all facilities from condition state to condition state, and
- that the spaces of condition states and actions to take have to be small enough to allow optimization of the infrastructure management problem

2.1.3 Conclusion

Single facility and network level MDP infrastructure management problems were presented above. It should be clear now what criteria are important in infrastructure management decision-making, and what the strengths and weaknesses of the MDP approach are. In the models shown, the transition probabilities ($p(j|i, a)$ terms) define deterioration. In the next chapter, we consider how $p(j|i, a)$ terms and other possible parameters of deterioration modeling are set.

Chapter 3

Infrastructure Deterioration

Modeling

The level of service an infrastructure facility provides deteriorates as time passes and ‘distresses’ accumulate. In the case of a simple roadway surface, distresses may take a number of forms including ruts (depressions in vehicle wheel paths) and fatigue or alligator cracks (sets of connected cracks on the surface of the roadway). Mathematical models have been developed to predict how specific distress measures, like the severity of rutting, change over time. Other models summarize information related to multiple distresses into an aggregate overall servicability measure. The artificial construct of an aggregate measure is employed to simplify reporting and suit the requirements of political decision-making (Paterson (1993)). For example, the decision-making problems shown in the preceding section used condition rating state spaces that represent some form of aggregated information related to deterioration. Mathematical models that predict how distress measurements or

overall condition ratings on infrastructure facilities change over time are known as deterioration models.

The first important pavement performance model was created by the American Association of State Highways Officials (AASHO) (AASHTO, 1993), and was based on the results of the AASHO Road Test. The AASHO Road Test involved measuring the performance of various test concrete roadway loops as trucks repeatedly drove around the loops. This provided a large amount of experimental data, data generated from purpose built facilities subject to accelerated usage (and possibly environmental) conditions. Alternate studies have used field data from measurements of in-service facilities to eliminate any potential bias stemming from the use of experimental data. Prozzi and Madanat (2004) and others have even shown how the large bodies of experimental data may be considered in conjunction with field data using joint estimation techniques.

Deterioration models in general focus on three areas when predicting deterioration: the structural composition of the built infrastructure, how the facility will be used, and the environmental conditions the facility will face. The AASHO pavement performance model cited above estimated the overall ‘damage’ done to a pavement section, at a time t , as $\left(\frac{N_t}{\rho}\right)^\omega$ where N_t represents the number of Equivalent Single Axle Loads (ESALs) applied until time t , and ρ and ω are parameters estimated from data that reflect the characteristics of particular roadway surfaces and sets of environmental conditions (AASHTO (1993)). Prozzi and Madanat (2004) extended the analysis of the AASHO model and identified independent variables including:

- asphalt concrete mix characteristics, and strength of the pavement (structural condi-

tions);

- axle loads and configurations (usage conditions);
- temperature, moisture, freezing, and thaw (environmental conditions).

Deterioration models can be classified as being deterministic or stochastic. Deterministic models reduce distress evolution or level of service degradation to a deterministic process that can be precisely quantified. In the AASHO model, once estimates for ρ and ω have been found, a count of ESALs applied on a particular roadway will provide one precise estimate of the damage done to that roadway. However, in Carnahan et al (1987) it was noted that pavement condition quality can actually “vary considerably” even when similar pavement sections are subjected to similar stresses. Stochastic models of deterioration recognize this ‘randomness’ and assign probabilities to a range of different estimates of damage done.

Stochastic models are typically parametric, in the sense that a functional form is assumed to describe the relationship between independent variables and stochastic deterioration. Data are used to estimate the parameters associated with such models. It is worth noting that there has been some research in the areas of non-parametric and semi-parametric deterioration modeling. In particular, DeStefano and Grivas (1998) did not assume any functional form in defining transition probabilities associated with deterioration. Mauch and Madanat (2001) used the Cox proportional hazards model to create a more descriptive model of deterioration without pre-specifying distributions for parameters relating independent variables and deterioration.

The Mauch and Madanat (2001) work cited above is an example of a time-based

deterioration model. In this formulation, probabilities are estimated related to the length of time facilities spend in particular condition rating states. The alternative to time-based models is state-based models that look at the probabilities of condition state changes over fixed time intervals. Time-based models are more flexible than state-based models since they allow consideration of distress progression over different time frames. However, it should be clear how state-based models are sufficient for MDP formulations of the infrastructure management problem like those presented in section 2.1.

In this section, we examine one state-of-the-art, time-based, stochastic deterioration model found in Mishalani and Madanat (2002). Although the model chosen allows for consideration of different time frames, special attention will be paid to explaining how the model could be used to provide inputs to optimization algorithms based on discrete time-step MDP formulations of the infrastructure management problem (as presented in section 2.1). In keeping with the discussions of the preceding section, consider the deterioration of a bridge deck rated by the FHWA on a 10 point scale.

3.1 Survival Analysis

Survival analysis deals with situations where a certain process ‘survives’ for a stochastic length of time and then stops. The length of time, T , a bridge deck spends in a certain FHWA condition state can be thought of as a survival time. Let $f(t)$ be the probability density function of T and $F(t)$ be the cumulative density function of T . One function of interest is the survival function $S(t)$, the probability that the bridge deck remains

in the specified state through at least time t .

$$S(t) = 1 - F(t) \quad (3.1)$$

There is a relationship between the survival function and the probability density function of T .

$$f(t) = \frac{dF(t)}{dt} = -\frac{dS(t)}{dt} \quad (3.2)$$

Another function that might be of interest is the hazard rate function $h(t)$, the instantaneous risk that the bridge deck will leave the condition state at time t :

$$\begin{aligned} h(t) &= \lim_{\Delta \rightarrow 0} \frac{\Pr(t \leq T \leq t + \Delta | t \leq T)}{\Delta} \\ h(t) &= \lim_{\Delta \rightarrow 0} \left(\frac{F(t + \Delta) - F(t)}{S(t)} \right) \left(\frac{1}{\Delta} \right) \\ h(t) &= \frac{f(t)}{S(t)} \end{aligned} \quad (3.3)$$

Equations 2.4 and 2.5, taken together, yield

$$h(t) = -\frac{d \log S(t)}{dt} \quad (3.4)$$

and:

$$S(t) = e^{-\int_0^t h(\tau) d\tau} \quad (3.5)$$

If $S(t)$ and $h(t)$ can be defined, they provide a model for bridge deck deterioration. This model, in turn, can be used to parameterize simpler models that can be used in decision-making. For example, the math programs in the preceding section on this chapter relied on transition probabilities to describe infrastructure deterioration. Notice that the conditional probability a bridge deck will be in a particular condition state i in year $t + 1$ given that it

was in state i in year t with no action applied, $p_t(i|i, none)$ can be expressed as a function of the survival function:

$$p_t(i|i, none) = \frac{S(t+1)}{S(t)} \quad (3.6)$$

Also note that if $p_t(i|i, none)$ does not depend on t , then we have a Markovian transition probability $p(i|i, none)$.

Assume the condition of a facility only worsens over time, and that we know survival functions for states i and $i-1$. It would then be possible to find the probability that a facility in state i at time t is in state $i-1$ at time $t+1$. Simply condition on the time the facility spends in state i :

$$p_t(i-1|i, none) = \int_{\tau=0}^1 (S_{i-1}(1-\tau)) \left(\frac{f_i(t+\tau)}{S_i(t)} \right) d\tau$$

or, in terms of the hazard function:

$$p_t(i-1|i, none) = \int_{\tau=0}^1 (S_{i-1}(1-\tau)) \left(\frac{(h_i(t+\tau))(S_i(t+\tau))}{S_i(t)} \right) d\tau \quad (3.7)$$

Again, assuming $p_t(i-1|i, none)$ does not depend on t , this would represent a transition probability that could be plugged into the decision-making problem formulations shown earlier. Other transition probabilities could be estimated using similar techniques.

So far our use of survival analysis has been fairly vague, there have been no attempts to describe the structure of hazard rate or survival functions. Survival analysis has merely provided a framework which allows for an arbitrarily detailed model of deterioration yet can provide inputs (transition probabilities) to decision-making platforms that rely on simplified models of deterioration.

3.2 The Weibull Distribution

In order to estimate specific values of interesting functions, it is useful to assume the distributions associated with survival time T values adhere to pre-specified form. Clearly, the length of time a bridge deck or other facility spends in a condition state must be positive. It is useful to incorporate a term measuring the degree to which T exhibits positive or negative duration dependence. Markovian transition probabilities assume no duration dependence, i.e. that the length of time a facility has spent in a state is irrelevant to how long it will remain in that state in the future. Measuring duration dependence can provide insight as to how appropriate it is to use Markovian transition probabilities. One useful model can be built by assuming the duration of time a bridge deck spends in a state follows the Weibull distribution. The hazard rate is then:

$$h(t) = p\lambda^p t^{p-1} \quad (3.8)$$

where λ and p are non-negative parameters to be estimated from data. λ is known as the scale parameter because it multiplies and scales t values while p is known as the shape parameter because of its importance in determining the shape of the hazard rate. p is also a measure of duration dependence. If $p < 1$, then there is evidence of negative duration dependence and the longer a facility has spent in a condition state, the less likely it will remain in that state. Likewise, $p > 1$ is a sign of positive duration dependence.

Given the hazard rate from above, the survival function takes the form:

$$S(t) = e^{-(\lambda t)^p} \quad (3.9)$$

Similarly, given these hazard and survival functions, it must be that:

$$f(t) = p\lambda^p t^{p-1} e^{-(\lambda t)^p} \quad (3.10)$$

Clearly, there are important characteristics of bridge decks that are relevant in modeling time-in-state. Some factors that have been found to be statistically significant when looking at reinforced concrete bridge decks in Indiana are presented in Table 3.1 (Mishalani and Madanat, 2002).

Binary Variables	Description of when set to 1
hw class 1	Interstate, rural, open to traffic
hw class 3	Other primary road, rural
hw class 5	Secondary road, rural, state jurisdiction
region	Northern Indiana
wear surf 1	Concrete wearing surface material (no protective system)
type 2	Concrete (not prestressed), continuous deck
Continuous Variable	Description
age	age at beginning of time in state or observation period (years)

Table 3.1: Deterioration model fit: Bridge deck characteristics related to deterioration

In order to include the variables from Table 3.1 in our analysis, say that for bridge deck i , the parameter λ_i that appears in hazard and survival functions depends on the characteristics of bridge deck i . λ_i must be non-negative to ensure the hazard function is non-negative. Say then that:

$$\lambda_i = e^{-\beta X_i} \quad (3.11)$$

where X_i is a column vector of exogeneous variables that describe the characteristics of bridge deck i and β is a row vector of parameters to be estimated.

Data is available from the Indiana Bridge Inventory (IBI) regarding both the characteristics of bridge decks in Indiana and the time these decks spent in state 8. Parameters

p and β introduced in the previous section can be estimated from this data, as was done in Mishalani and Madanat (2002). The data available contains both uncensored and right censored data points measuring the time different bridge decks have spent in state 8. Right censoring refers to situations where it is known that the bridge deck spent *at least* a certain amount of time in state 8.

Maximum likelihood estimation presents a reasonable way to estimate p and β and can incorporate uncensored and right censored data. The likelihood of a transition out of state 8 at a time t_i is $f_i(t_i)$ for bridge deck i . Similarly, the likelihood that bridge deck i was in state 8 through time t_i is $S_i(t_i)$. Assuming each bridge deck i has one observation associated with it, this yields the likelihood function over all observations:

$$L = \prod_i [f_i(t_i)]^{\delta_i} [S_i(t_i)]^{1-\delta_i} \quad (3.12)$$

where δ_i is an indicator variable that takes value 1 if the data associated with bridge deck i is uncensored and 0 if the data is right censored. According to equations 2.11 and 2.12 then

$$L = \prod_i (p\lambda_i^p t_i^{p-1})^{\delta_i} e^{-(\lambda_i t_i)^p} \quad (3.13)$$

Taking the natural log of both sides yields

$$\log(L) = (\log p) \sum_i \delta_i + p \sum_i \delta_i \log \lambda_i + (p-1) \sum_i \delta_i \log t_i - \sum_i (\lambda_i t_i)^p \quad (3.14)$$

or equivalently

$$\log(L) = (\log p) \sum_i \delta_i - p \sum_i \delta_i \beta X_i + (p-1) \sum_i \delta_i \log t_i - \sum_i (t_i e^{-\beta X_i})^p \quad (3.15)$$

Equations 2.16 and 2.17 define the log-likelihood as a function of p and β . Maximizing these functions gives Maximum Likelihood Estimates (MLEs) for p and β . Using the

Indiana Bridge Inventory data, estimated values for p and β matched those in the literature (Mishalani and Madanat (2002)) and are presented in Table 3.2.

Variable	Parameter Value
hardware class 1	-0.639
hardware class 3	-0.553
hardware class 5	-0.574
region	-0.840
age	0.153
wearing surface 1	-0.801
pavement type 2	0.384
constant	2.13
p	1.93

Table 3.2: Deterioration model fit: Estimated values of parameters related to bridge deck deterioration

Given the parameter values from Table 3.2, it is possible to construct the probability and cumulative density functions associated with the time different bridge decks spend in state 8. For example, Figure 3.1 shows the pdf of time in state 8 for a 10 year old bridge deck on a rural interstate in Southern Indiana that is made of prestressed concrete with a protective system.

Figure 3.2 shows the conditional probability of the same type of bridge deck remaining in state 8 for 1 year plotted against the length of time the deck has already spent in state 8. The risk of a bridge deck leaving state 8 is clearly related to how long it's previously been in state 8. This was clear from the fact that the MLE for p is 1.93 (significantly greater than 1).




Figure 3.1: Deterioration model fit: Pdf of the time a bridge deck will spend in state 8




Figure 3.2: Deterioration model fit: Probability of a bridge deck remaining in state 8 for 1 year as a function of the time already spent in state 8

A detailed example of a deterioration model has been described. It is only one of many infrastructure deterioration models in the literature. Some constructive criticisms of deterioration models have been made in the past.

For instance, models sometimes consider maintenance as an explanatory variable when it is in fact usually a function of the condition of managed infrastructure, creating endogeneity bias (Ramaswamy and Ben-Akiva, 1990). Field data often provides a time window during which infrastructure assets failed or distresses appeared. The absence of exact event times, if not properly accounted for, creates censoring bias (Prozzi and Madanat, 2004). In the case of pavements, experimental data suffers from its failure to consider material aging, or varying traffic speeds and paths (Archilla, 2001). Returning to pavements, the original AASHO model referenced previously made use of a large experimental data set but set parameters using only data from sections of pavement that had failed during trials, leading to truncation bias (Prozzi and Madanat, 2004). Statistical techniques can address truncation, censoring, and endogeneity bias, while joint estimation techniques can address the bias experimental data models may have vis a vis field data models.

A state-of-the-art time-based deterioration model was presented in this section. This deterioration model and those used in practice today are quite sophisticated and based upon best statistical practices. The next section will address difficulties related to deterioration modeling that cannot be addressed using statistical techniques. The goal of the next section of this paper is not to critique deterioration models in use today, but to draw attention to their limitations.

3.3 The Curse of Uncertainty

Clearly, the conditional probabilities ($p(j|i, a)$ terms) drive math programs like those shown in section 2.1. It is known that the solution to a MDP “is often quite sensitive to changes in the transition probabilities” (Nilim and El Ghaoui, 2004). This is troubling because it is difficult to say precisely what the conditional probability of a bridge deck (or other facility) being in state j is one year after it was in state i with action a applied. Data regarding deterioration may be limited and accumulate slowly; changes in the condition of a bridge deck may take years before they are noticeable. Condition ratings used to describe bridge decks are artificial constructs and expert judgement regarding how transitions occur between different condition ratings may be quite limited. Finally, there may be factors important to the manner in which a bridge deck deteriorates that are not, and in many cases cannot be, considered in state-of-the-art deterioration models.

3.3.1 Parametric vs. Epistemic Uncertainty

Problems associated with the definition of condition states or, broadly speaking, the choice of deterioration model, can be differentiated from problems associated with setting transition probabilities or fixing the parameters of a deterioration model. The first set of problems is more fundamental and relates to errors in our intuition and understanding of deterioration. Let’s call uncertainty that stems from these type of errors *epistemic uncertainty* since it relates to a fundamental lack of knowledge. In contrast, statistical uncertainty related to the choice of parameters for deterioration models assumed to be valid will be referred to as *parametric uncertainty*.

Epistemic uncertainty may stem from a variety of sources, including uncertainty in how to characterize the condition of an infrastructure facility, how to characterize the independent variables that relate to the condition of that facility, and how to characterize the interaction between the independent variables and the condition of the facility. Some of the independent variables that were identified in the preceding section of this chapter related to environmental conditions. It is famously difficult to describe how environmental conditions change over time, and it is no clearer to describe even the form of how these conditions impact the condition of infrastructure facilities. Similar difficulties arise in describing how a facility will be used in the future. It may even be difficult to describe precisely how a facility was used in the past, or how it was constructed, and how this relates to deterioration.

Parametric uncertainty, in contrast, stems from a small number of sources. It is primarily associated with the limited, and at times biased, data used to calibrate deterioration models. Even if the complex form of the mathematical relation between structural, usage, and environmental conditions and the condition of an infrastructure facility were assumed known, limited data would make it impossible to precisely parameterize the relation.

In the next subsection of this text, evidence of epistemic uncertainty and estimates of the magnitude of parametric uncertainty in the sample deterioration model introduced in the preceding section will be presented.

3.3.2 Quantifying Uncertainty

Let's return to the example deterioration model introduced in the previous section of this paper. Figure 2.2 showed the conditional probability of a bridge deck being in state 8 one year after being in state 8, as a function of the time that facility has already spent in

state 8. As Figure 2.2 makes clear, the risk of a bridge deck leaving state 8 is related to how long that bridge deck has already been in state 8. Any assumption otherwise does not hold and a deterioration model making this type of assumption would incur errors associated with epistemic uncertainty. In this case, probabilities associated with leaving state 8 could range between 0.8 and 1.0 and using one fixed value could result in using a probability that is as much as 0.2 off. It is not typically so easy to quantify the potential magnitude of errors associated with epistemic uncertainty.

However, there are techniques that can be used to estimate the magnitude of statistical uncertainty associated with parameter estimates. For instance, the non-parametric bootstrap method can be used to estimate uncertainty surrounding terms set using a collection of independent and identically distributed (i.i.d.) data points. Clearly, the raw data of time in state for the different bridge decks in the state of Indiana used earlier in this text are not independent and identically distributed. However, it is possible to create data that are more plausibly i.i.d..

Residual data values left over after deterioration model fitting may be assumed to be i.i.d.. However, it is possible that explanatory variables remain related to residual values after fitting. Luckily, focusing attention on one particular highway class and wearing surface at a time negates any potential differences between residual data values associated with different highway classes and wearing surfaces. As Table 3.3 demonstrates, there remain enough data values relating to decks with different highway class, wearing surface combinations to allow this type of analysis.

	hw class 1	3	5	other
wear surface 1	24	64	81	35
not wear surface 1	23	45	35	61

Table 3.3: Deterioration model uncertainty: Counts of bridge deck sample population by hw class and wear surface

Also note that resampling from absolute residual data values and attaching these to expected times in state could create a situation where the sum of the two yields a negative value. This would correspond to an observation of a negative time in state, which is impossible. Say that a bridge deck has characteristics that can be represented by a variable vector X and that this bridge deck has an expected time in state $E[T(X)]$ and an observation of time in state $T_1(X)$. The residuals data value associated with this bridge deck would be $T_1(X) - E[T(X)]$ whereas the relative residual data value would be $\frac{T_1(X) - E[T(X)]}{E[T(X)]}$. Relative residuals are always greater than or equal to -1. Resampling from this data pool, multiplying the results by expected times in state and adding expected times in state cannot yield negative data values.

Using 500 groupwise bootstrapped simulations of resampling relative error values provided distributions of estimated deterioration model parameters as shown in Figure 3.3. Note that the observed distributions appear roughly but not exactly symmetric and normal, neither of which was an assumption of the non-parametric bootstrap analysis performed.

Table 3.4 shows confidence intervals around parameter estimates. The mean of most parameter estimates are close to the original parameter estimates, indicating little bias in these values. However, there does appear to be some bias in the estimation of parameters associated with region and type of bridge deck, as well as with the shape parameter p . Importantly, the confidence intervals surrounding most of the parameters appear sizable.

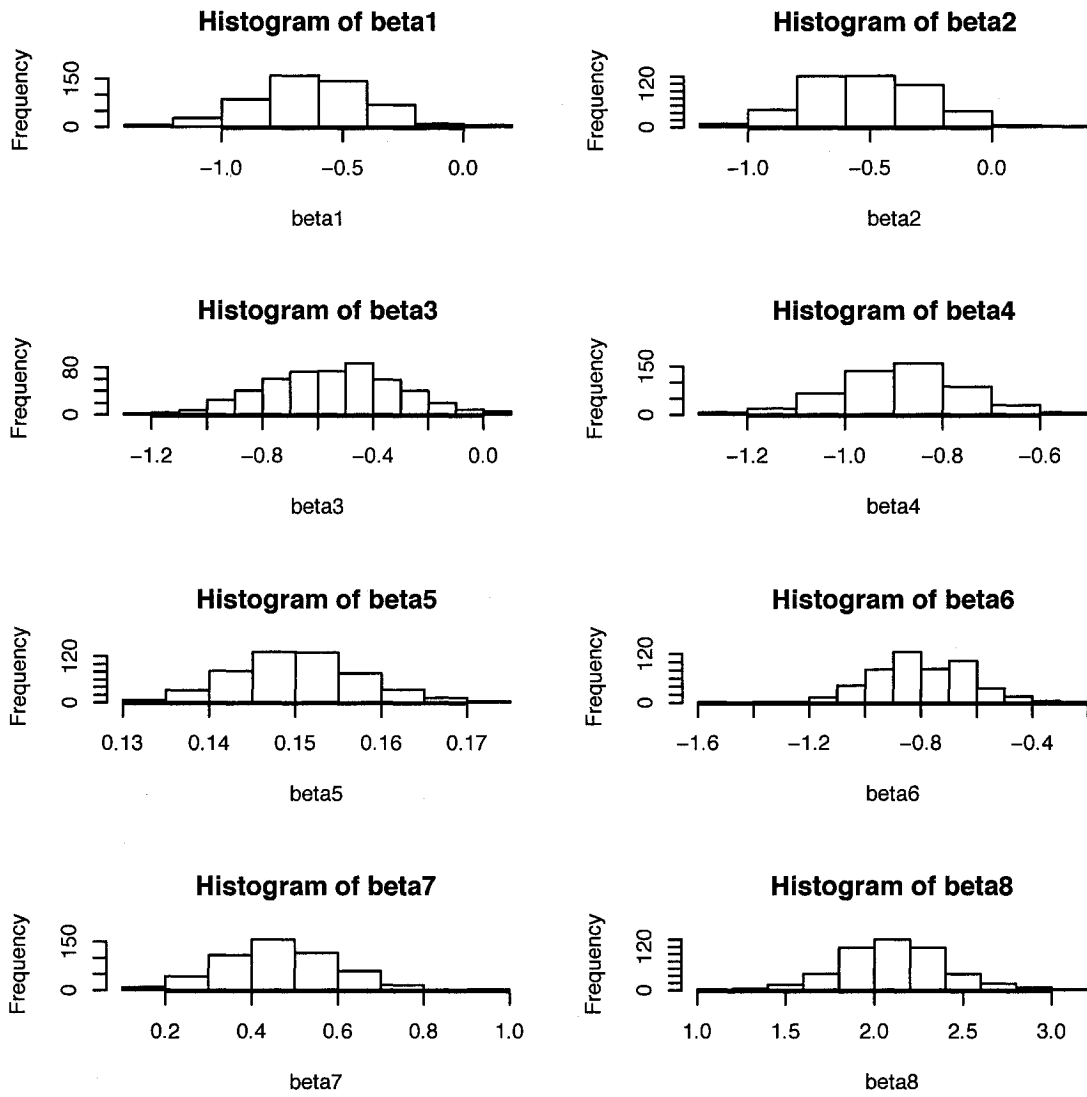


Figure 3.3: Deterioration model uncertainty: Histograms of estimated parameter values

	95% confidence	75%	mean	MLE
hw class 1	(-1.05,-0.29)	(-0.79,-0.48)	-0.64	-0.64
hw class 3	(-0.91,-0.15)	(-0.69,-0.36)	-0.52	-0.55
hw class 5	(-0.93,-0.17)	(-0.71,-0.38)	-0.55	-0.57
region	(-1.10,-0.70)	(-0.96,-0.80)	-0.89	-0.84
age	(0.14,0.16)	(0.15,0.16)	0.15	0.15
wear surf	(-1.07,-0.54)	(-0.92,-0.67)	-0.80	-0.80
type	(0.27,0.66)	(0.37,0.55)	0.46	0.38
constant	(1.65,2.58)	(1.91,2.29)	2.11	2.13
p	(1.81,2.35)	(1.93,2.17)	2.06	1.93

Table 3.4: Deterioration model uncertainty: Confidence intervals of model parameters

Figure 3.4 demonstrates that the uncertainty in the estimation of parameters of bridge deck deterioration leads to uncertainty in the probability density function of how long a bridge deck will spend in state 8. The figure places a 95% confidence interval around the function shown in Figure 3.1, the pdf of time in state 8 for a 10 year old bridge deck on a rural interstate in Southern Indiana that is made of prestressed concrete with a protective system.

Figure 3.5 places a 95% confidence interval around the function shown in Figure 3.2, the estimated conditional probability of remaining in state 8 for 1 year. It was already discussed how difficult it would be to set one value to describe the probability of remaining in state 8 for one year given that this quantity depends on the time the deck has already spent in state 8. The situation is made more complex with the recognition that significant uncertainty surrounds conditional probabilities values even for a fixed t . Given Figure 3.5, the best one can do is to say that the probability of a bridge deck in state 8 remaining in state 8 through the next year is most likely between 0.6 and 1.0.




Figure 3.4: Deterioration model uncertainty: 95% confidence interval pdf of time in state 8




Figure 3.5: Deterioration model uncertainty: 95% confidence interval of probability of remaining in state 8 for 1 year

3.4 Conclusion

This section has provided logical reasons why there might be significant uncertainty in infrastructure deterioration modeling, and has shown empirical evidence of this uncertainty. This uncertainty is ignored in traditional MDP formulations of the infrastructure management problem. (Note how the management problems presented in section 2.1 of this text assumed an accurate and precise Markov chain model of facility deterioration.)

A few questions immediately arise:

- What alternatives exist for considering uncertainty more explicitly within the infrastructure management process?
- What are the trade-offs between these alternatives?
- Would consideration of uncertainty alter optimal infrastructure management policies?
- Would consideration of uncertainty alter optimal infrastructure management costs?

This series of related questions motivates the research that will be presented in the following chapters of this text.

Chapter 4

Robust Infrastructure

Management: At the Single

Facility Level

One strategy for dealing with uncertainty involves using robust optimization to guide decision-making. Robust optimization is a modeling methodology to solve optimization problems in which parameters are not known precisely but known to be in certain ranges, or to belong to certain sets (“uncertainty sets”). The approach is to seek solutions that are not overly sensitive to any realization of uncertainty.

In the field of infrastructure management and elsewhere, stochastic optimization techniques have been used to deal with situations where there is significant uncertainty regarding how a (possibly deterministic) system behaves. Robust stochastic optimization is the more appropriate methodology to use when there is a stochastic system to be managed

and that system cannot be precisely modeled. In robust stochastic programming, no true underlying stochastic model of the data is assumed to be known. A robust feasible solution is one that tolerates changes in the parameters of the problem, up to a given bound known a priori, and a robust optimal solution is a robust feasible solution with the best possible value of the objective function.

4.1 Introduction to Robust Optimization

The origins of robust optimization can be traced back to work done by Soyster in 1973. Soyster considered linear programming problems where the feasible region was defined in terms of set containment. In this model, instead of fixing model parameters used to define the feasible region, Soyster allowed them to vary within pre-specified convex sets. Soyster maximized benefits ensuring the proposed solution was feasible for all possible values of model parameters, a methodology he admitted ‘provides an ultraconservative strategy’ (Soyster, 1973).

Bertsimas and Sim (2004) modified Soyster’s formulation to allow the ‘level of conservatism’ to be adjusted. In the Bertsimas model, uncertainty is measured in terms of the number of linear constraints that may be violated. It is possible to obtain different solutions considering different numbers of constraints that may be violated. This allows for an analysis of how optimal strategies change as more and more initial assumptions are considered possibly erroneous. However, it may not always make sense to define uncertainty in terms of number of initial assumptions violated. For instance, there may be more or less uncertainty surrounding different constraints of a mathematical program.

El Ghaoui has used likelihood regions and entropy bounds to perform robust optimization using a statistically sophisticated representation of uncertainty (El Ghaoui (2003), Nilim and El Ghaoui (2004)). The El Ghaoui work allows for the consideration of different levels of conservatism via adjustable entropy or likelihood bounds. In addition, the El Ghaoui work was one of the first to apply robust optimization to dynamic programming. This is especially important given that the Single Facility Infrastructure Management DP Recursion presented in section 2.1 was set up to be solved via dynamic programming.

4.2 The Maximin Criterion

A sample robust optimization problem will be presented in this section. It represents one way that the Single Facility Infrastructure Management DP Recursion presented in Section 2.1 might be reformulated to consider uncertainty in transition probabilities.

In the Single Facility Infrastructure Management DP Recursion, we used a transition probability matrix p to define deterioration. Given that there may be uncertainty surrounding this matrix, let's now consider a whole "uncertainty set," \mathcal{P} , that includes many transition probability matrices, any one of which may define the system in question.

Consider minimizing costs assuming that nature will act as an opponent. This decision criterion is known as the MAXIMIN criterion and is quite well known in the literature of game theory. Given that the motivation for using robust optimization was a lack of knowledge of how to precisely set transition probabilities, it makes sense to avoid assuming anything about how likely it is that different transition probability matrices within \mathcal{P} describe deterioration. The MAXIMIN decision criterion makes no such assumptions. Ad-

ditionally, the majority of work in the field of robust optimization uses such an approach. Alternatives to the MAXIMIN approach will be discussed shortly. For now, examine a DP recursion based on the Single Facility Infrastructure Management Problem (SFIMP) showing how a robust MAXIMIN formulation might look in practice.

$$v_t(i) = \min_{a_t(i) \in A} \max_{p \in \mathcal{P}} \left[[u(i) + g(i, a_t(i))] + \alpha \left[\sum_{j \in I} p(j|i, a_t(i)) v_{t+1}(j) \right] \right] \quad (4.1)$$

It is possible to say more about the uncertainty set \mathcal{P} . As a set containing transition probability matrices, it must be that any element $p \in \mathcal{P}$ must be in $I \times I \times A \rightarrow [0, 1]$ and must satisfy $\sum_{j \in I} p(j|i, a) = 1$ ($\forall i \in I, a \in A$) and $p(j|i, a) \geq 0$ ($\forall i, j \in I, a \in A$). Additional constraints need to be added to ensure that any matrix p in the set \mathcal{P} is a plausible description of deterioration. Robust optimization will be performed considering, and only considering, transition probability matrices in the uncertainty set, so the definition of uncertainty set becomes crucial to problem formulation.

In the formulation developed here, an ‘‘uncertainty level’’ δ will be employed. δ will be a parameter set ahead of time to a value between 0 and 1, with a higher value corresponding to a higher degree of uncertainty. Setting the uncertainty level to 0 implies no uncertainty, meaning an uncertainty set is defined that includes only the transition probability matrix q given by some initial set of assumptions. Increasing the uncertainty level adds new transition probability matrices to the set. The uncertainty level can be thought of as a measure of level of conservatism, as laid out in Bertsimas and Sim (2004), to be adjusted by decision-makers according to their beliefs and preferences.

In this example, a transition matrix will be included in the uncertainty set if and

only if the difference between any element of the transition matrix and the corresponding element of the matrix q is less than or equal to the uncertainty level. Thus, in addition to the above constraints on elements $p \in \mathcal{P}$, add the constraint that $|p(j|i, a) - q(j|i, a)| \leq \delta$ ($\forall i, j \in I, a \in A$). Seen in this light, the uncertainty level represents how large an error in estimated transition probabilities is considered possible. An uncertainty level of 0 would correspond to assuming a known transition probability matrix accurately and precisely defines an infrastructure asset's decay. On the other hand, an uncertainty level of 1 corresponds to a complete lack of confidence in any given transition probability matrix.

The formulation of this example is admittedly very simple. The same amount of uncertainty is assumed to surround all transition probability estimates. The form of this uncertainty is always the same, real transition probabilities are assumed to fall within a symmetric box around initially assumed transition probability estimates. It is definitely possible to imagine more complex definitions of uncertainty sets. For instance, in Chapter 5 a formulation is introduced where the uncertainty sets surrounding different transition probability estimates are different and related to standard error distributions. However, for now our simple definition of uncertainty is sufficient. The aim here is only to show a sample robust infrastructure management problem, and then to test to see if this formulation yields substantially different results as compared to non-robust formulations. Furthermore, it is worth noting that the formulation shown here requires only the specification of an 'uncertainty level' with a clear intuitive meaning. For practicing engineers and political policy makers, the simplicity of this approach may be a strength rather than a weakness.

Moving forward, it is now possible to build a complete formulation of a Single

Facility Robust Infrastructure Management Problem (SFRIMP).

Single Facility Robust Infrastructure Management Problem (SFRIMP)

$$v_t(i) = \min_{a_t(i) \in A} \max_{p \in \mathcal{P}} \left[[u(i) + g(i, a_t(i))] + \alpha \left[\sum_{j \in I} p(j|i, a_t(i)) v_{t+1}(j) \right] \right] \quad (4.2)$$

where $\mathcal{P} = \{p \in I \times I \times A \rightarrow [0, 1] : |p(j|i, a) - q(j|i, a)| \leq \delta \ (\forall i, j \in I, a \in A), \sum_{j \in I} p(j|i, a) = 1 \ (\forall i \in I, a \in A), p(j|i, a) \geq 0 \ (\forall i, j \in I, a \in A)\}$

The constraints on the set \mathcal{P} can also be written as constraints of the math program. Such a formulation is more typical of the field of Operations Research, and can be directly input into mathematical program solver software.

Alternate Formulation (SFRIMP)

$$v_t(i) = \min_{a_t(i)} \max_p \left[[u(i) + g(i, a_t(i))] + \alpha \left[\sum_{j \in I} p(j|i, a_t(i)) v_{t+1}(j) \right] \right] \quad (4.3)$$

subject to the following constraints:

- (1) $|p(j|i, a) - q(j|i, a)| \leq \delta \quad \forall i, j \in I \text{ and } a \in A$
- (2) $\sum_{j \in I} p(j|i, a) = 1 \quad \forall i \in I \text{ and } a \in A$
- (3) $p(j|i, a) \geq 0 \quad \forall i, j \in I \text{ and } a \in A$

Planning agencies may perceive the MAXIMIN approach being used here to be too conservative. When managing a large network of facilities, it may be too costly, and unrealistic, to manage each one under the assumption that nature is always malevolent. However, it is important to note that by controlling the uncertainty level, decision-makers will be able to control how conservative optimal strategies will be.

There are some alternatives to the MAXIMIN approach. Averbakh (2000) has proposed a methodology for finding minmax regret solutions to problems with uncertainty

objective function coefficients. A minmax regret strategy minimizes the maximum regret, a measure of the opportunity cost of selecting one strategy when a better strategy existed for a particular realization of uncertainty. This approach is still fairly conservative, considering worst-case values of regret. A more optimistic strategy, known as MAXIMAX, is similar to MAXIMIN but involves acting under the assumption that nature will work with decision makers instead of against them. The most realistic point of view would be to recognize that nature will act neither as an adversary nor as an ally, but somewhere in between.

4.3 The Hurwicz Criterion

One attractive alternative involves applying the Hurwicz criterion, as laid out in Hurwicz (1951). The Hurwicz criterion allows a decision maker to set his or her own 'optimism level,' β . The optimism level β must be a number between 0 and 1. The pessimism level is defined as $1 - \beta$. Decisions are then made by selecting actions that maximize benefits obtained by summing the optimism level times the greatest possible benefit level with the pessimism level times the least possible benefit level. An optimism level of 0 would correspond to minimizing costs assuming deterioration is the most severe of all the rates of deterioration considered possible. This decision criterion would be equivalent to MAXIMIN. An optimism level of 1 would correspond to minimizing costs assuming deterioration is the least severe considered possible. All optimism levels between 0 and 1 trade off costs in the most severe case with costs in the least severe case.

Note that consideration of best and worst case outcomes requires using two different transition probability matrices. Let p^b be the transition probability matrix that

corresponds to ‘best case’ conditions and let p^w correspond to ‘worst case’ conditions. In the context of asset management, a robust optimization problem that employs the Hurwicz criterion might be defined as below.

Single Facility Hurwicz Robust Infrastructure Management Problem (SFHRIMP)

$$v_t(i) = \min_{a_t(i)} \max_{p^w \in \mathcal{P}} \min_{p^b \in \mathcal{P}} [u(i) + g(i, a_t(i))] + (\alpha) \left[(\beta) \left(\sum_{j \in I} p^b(j|i, a_t(i)) v_{t+1}(j) \right) + (1 - \beta) \left(\sum_{j \in I} p^w(j|i, a_t(i)) v_{t+1}(j) \right) \right] \quad (4.4)$$

Again, this formulation can be altered to reflect the traditional formulation structure used in Operations Research.

Alternate Formulation (SFHRIMP)

$$v_t(i) = \min_{a_t(i)} \max_{p^w} \min_{p^b} [u(i) + g(i, a_t(i))] + (\alpha) \left[(\beta) \left(\sum_{j \in I} p^b(j|i, a_t(i)) v_{t+1}(j) \right) + (1 - \beta) \left(\sum_{j \in I} p^w(j|i, a_t(i)) v_{t+1}(j) \right) \right] \quad (4.5)$$

subject to the following constraints:

$$(1.1) \quad |p^b(j|i, a) - q(j|i, a)| \leq \delta \quad \forall i, j \in I \text{ and } a \in A$$

$$(1.2) \quad |p^w(j|i, a) - q(j|i, a)| \leq \delta \quad \forall i, j \in I \text{ and } a \in A$$

$$(2.1) \quad \sum_{j \in I} p^b(j|i, a) = 1 \quad \forall i \in I \text{ and } a \in A$$

$$(2.2) \quad \sum_{j \in I} p^w(j|i, a) = 1 \quad \forall i \in I \text{ and } a \in A$$

$$(3.1) \quad p^b(j|i, a) \geq 0 \quad \forall i, j \in I \text{ and } a \in A$$

$$(3.2) \quad p^w(j|i, a) \geq 0 \quad \forall i, j \in I \text{ and } a \in A$$

Hurwicz criterion based robust optimization does require the specification of both an uncertainty and an optimism level. Planning agencies may find it difficult to specify

how much uncertainty they have with regards to infrastructure decay rates, or may find it undesirable to have to place a level of optimism on their management strategies. However, asset management clearly does involve managing systems with some degrees of uncertainty. The more the issue of uncertainty and the decision of how to manage it are discussed, the more informed asset management policies will be.

Uncertainty levels associated with transition probability matrices can be derived from confidence intervals surrounding the statistics provided by deterioration models. Deterioration models that are based on large data sets of infrastructure deterioration over extended periods of time will produce smaller uncertainty sets than less refined models. Optimism levels are more subjective. However it will be shown that the asset management problem can be made robust without making its computational complexity too great, making it is possible to imagine solving a particular asset management problem numerous times with various optimism (and uncertainty) levels to see how performance and reliability guarantees can be traded off.

It still might be possible to characterize the formulations shown above as too conservative on the grounds that certain transitions might be considered, even though such transitions are considered impossible in real life. It is a relatively simple task to ensure that certain 'impossible' transitions are never considered in robust optimization. In the definition of \mathcal{P} , simply fix corresponding transition probabilities to 0.

Alternately, in the language of Operations Research, it is possible to add a decision variable m where $m \in I \times I \times A \rightarrow \{0, 1\}$ ensures whenever an initial model considers transitions impossible, models considered in the optimization do likewise. Additional constraints

are required in the optimization problem:

- $q(j|i, a) + m(j|i, a) > 0$ $\forall i, j \in I$ and $a \in A$
- $p(j|i, a) m(j|i, a) = 0$ $\forall i, j \in I$ and $a \in A$

OR, for Hurwicz robust optimization:

- $q(j|i, a) + m(j|i, a) > 0$ $\forall i, j \in I$ and $a \in A$
- $p^b(j|i, a) m(j|i, a) = 0$ $\forall i, j \in I$ and $a \in A$
- $p^w(j|i, a) m(j|i, a) = 0$ $\forall i, j \in I$ and $a \in A$

The new constraints work as follows. For any $i, j \in I$ and $a \in A$, if $q(j|i, a) = 0$ then $m(j|i, a)$ must be set equal to 1. This forces $p(j|i, a)$ or $p^b(j|i, a)$ and $p^w(j|i, a)$ to 0 to satisfy the remaining constraints. If $q(j|i, a) > 0$, then the first constraint is not binding on $m(j|i, a)$. It will be set equal to 0 so that the remaining constraints are not binding on $p(j|i, a)$ or $p^b(j|i, a)$ and $p^w(j|i, a)$.

The optimization problems described here are not as complex as they might appear at first glance. The Hurwicz criterion asset management problem just combines the costs associated with best case and worst case transition probability matrices. The simplest way to solve this problem is to first solve problems of finding best and worst case transition probabilities and associated cost-to-go functions for all potential actions in all states.

Transition probability matrices are included in the uncertainty set if and only if their every term is within the uncertainty level of some given matrix. Thus the constraint of being in the uncertainty set is “separable” into individual constraints on individual probabilities. Note that different transition probabilities are only linked by the fact that they must

together form a complete transition probability matrix, so the sum of the probabilities from any initial state-action pair must be 1. In any given state i , taking action a , maximizing costs just implies finding a solution to the objective function $\max_{p \in \mathcal{P}} [c(i, a) + \alpha \sum_{j \in I} p(j|i, a) v_{t+1}(j)]$ which can be reduced to $\max_{p \in \mathcal{P}} \sum_{j \in I} [p(j|i, a) v_{t+1}(j)]$. This maximization problem can be solved exactly by altering the initial model transition probability matrix via shifting probability from less costly to more costly states. How much probability can be shifted is determined by the constraints used to define \mathcal{P} . The computational complexity is limited to the size of the transition probability matrix.

4.4 Computational Studies

In order to illustrate the application of robust dynamic programming algorithms to infrastructure management problems, an example is presented here. A one lane-mile segment of highway pavement is managed according to a policy obtained from infinite horizon robust dynamic programming. Previous research provides a ready source of data for how pavement deterioration can be modeled via static transition probabilities. However given the uncertainty in these transition probabilities, potential cost savings can be achieved by applying robust dynamic programming to this problem. It is worth noting that uncertainty is of more concern for infrastructure assets that have less refined deterioration models than pavement sections. Thus robust optimization may actually be better suited to the management of infrastructure assets like underground pipelines and drainage systems, where collecting data regarding deterioration is more problematic.

4.4.1 Problem Specifications

When managing a section of pavement, the decisions to be made include when and how to maintain, overlay, or reconstruct the pavement. In the example presented here, it is assumed that the choices of actions to take in any given year are those presented by Durango and Madanat (2002). These actions include: (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.

The costs of the actions presented here are derived from empirical work done by Carnahan et al (1987) and are included in Table 4.1 and Table 4.2. These costs vary according to the condition state of the pavement. A section of pavement is said to be in state 1 if it is unusable and in state 8 if it is brand new, with the intermediate states representing intermediate condition ratings. In the course of the computational studies done here, condition state 1 is to be avoided at all costs. Alongside agency costs, user costs associated with various pavement condition ratings are included in calculations.

Condition rating	Action to take						
	1	2	3	4	5	6	7
2	0.00	2.00	10.40	12.31	16.11	19.92	25.97
3	0.00	1.40	8.78	10.69	14.49	18.30	25.97
4	0.00	0.83	7.15	9.06	12.86	16.67	25.97
5	0.00	0.65	4.73	6.64	10.43	14.25	25.97
6	0.00	0.31	2.20	4.11	7.91	11.72	25.97
7	0.00	0.15	2.00	3.91	7.71	11.52	25.97
8	0.00	0.04	1.90	3.81	7.61	11.42	25.97

Table 4.1: Agency costs (dollars per lane-yard) of performing different MR&R actions on pavement with different condition ratings.

Condition rating	Cost (dollars per lane-yard)
2	25.00
3	22.00
4	14.00
5	8.00
6	4.00
7	2.00
8	0.00

Table 4.2: User costs (dollars per lane-yard) associated with pavement characterized by different condition ratings.

In their study, Durango and Madanat (2002) present three sets of transition probability matrices. The matrix that describes a section of pavement deteriorating at a “medium” rate is meant to reflect the current best estimate of how a given, random section of pavement will deteriorate. This transition probability matrix is itself derived from normal distributions associated with the performance of various pavement management maintenance actions, (Madanat and Ben-Akiva 1994). The parameters of the normal distributions are presented in Table 4.3. The inclusion of alternate “fast” and “slow” rates of deterioration draws attention to the fact that the “medium” estimate may under or over estimate decay in meaningful ways.

In the present example, the medium decay rate probabilities are used to initialize the robust dynamic programming application. Various uncertainty and optimism levels are considered. The policies obtained by robust optimization are compared to those obtained via non-robust optimization, and resulting expected costs calculated. Cases where the actual probabilities that guide system dynamics are ‘best-case’ and ‘worst-case’ are considered.

Action	Deterioration rate		
	Slow	Medium	Fast
Mean effects			
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
Standard deviation			
	0.30	0.50	0.70

Table 4.3: Mean and standard deviation of effects of actions on the condition rating of a pavement section.

The specifics of pavement management, as described above, will be used to create programs that simulate pavement management. Different decision-making methodologies will be tested, under different scenarios (best or worst-case), and the expected discounted costs of management noted. The ‘control’ strategy in this case will involve taking actions always assuming that the medium deterioration rate is correct. This assumption may or may not hold depending of if there is uncertainty (i.e. error) in initial assumptions or not. MAXIMIN and Hurwicz robust decision-making algorithms that are constrained to fix transition probabilities initially assumed to be impossible to 0 are compared to the control. It will be assumed that these robust algorithms correctly specify system uncertainty. This means that the performance of robust decision-making assuming an uncertainty level of 0.4 will be tested under best and worst-case scenarios given a true uncertainty level of 0.4.

4.4.2 Results

In this particular set of computational studies, actions found to be optimal in a robust sense differed substantially from those optimal in a non-robust sense. Table 4.4 displays MAXIMIN robust optimal policies for varying uncertainty levels. Note that the MAXIMIN robust policy when uncertainty level is 0 is the optimal policy for non-robust optimization.

Uncertainty level	Optimal policy						
	state 8	7	6	5	4	3	2
0.0	2	3	4	4	5	6	7
0.1	2	3	4	5	6	6	7
0.2	2	3	4	5	6	6	7
0.3	2	3	4	5	6	7	7
0.4	3	4	4	5	6	7	7
0.5	4	4	4	5	6	7	7
0.6	4	4	5	6	6	7	7
0.7	4	5	6	6	6	7	7
0.8	4	5	6	6	7	7	7
0.9	4	5	6	6	7	7	7
1.0	4	5	6	6	7	7	7

Table 4.4: Single facility robust optimization: Maximin robust optimal actions when different uncertainty levels are considered

The optimal management policies in MAXIMIN robust optimization are more conservative than those employed in non-robust optimization, especially as uncertainty becomes more significant. However, the actions prescribed by the MAXIMIN robust dynamic programming algorithm can yield significantly lower agency + user costs when compared to that prescribed by non-robust dynamic programming. For example, assuming transition probabilities follow those considered in the MAXIMIN robust algorithm (i.e. worst-case transition probabilities) costs are presented in Table 4.5.

Uncertainty level	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
non-robust cost	13	30	46	62	78	94	110	126	142	154	162
MAXIMIN robust cost	13	30	46	62	72	80	80	80	80	80	80

Table 4.5: Single facility robust optimization: Worst-case future discounted costs of managing a like new pavement

Figure 4.1 depicts the cost savings achieved by using the MAXIMIN robust algorithm rather than the non-robust algorithm by uncertainty level. Note that savings from using robust optimization techniques increase both in size and in relative terms as the uncertainty level increases. Figure 4.2 shows the ranges of potential costs of using the non-robust approach, by uncertainty level. The MAXIMIN robust optimization recognizes the potential for extraordinarily large costs in cases where uncertainty is high and chooses to maintain the pavement on a more regular basis in order to limit the potential maximum costs. Figure 4.3 shows the ranges of potential costs of using the MAXIMIN approach, by uncertainty level.

Clearly MAXIMIN asset management limits the maximum costs, but Figure 4.3 also shows that MAXIMIN is unable to lower costs as much as non-robust asset management systems can in best-case situations. This is one of the shortcomings of the MAXIMIN approach, and one area in which the less conservative Hurwicz robust optimization is able to do substantially better.




Figure 4.1: Single facility robust optimization: Relative and absolute benefit of using max-min robust optimization in worst-case conditions




Figure 4.2: Single facility robust optimization: Cost ranges of non-robust asset management




Figure 4.3: Single facility robust optimization: Cost ranges of Maximin robust asset management

Uncertainty level	Optimism level	Optimal policy		
		in state 8	7	6
0.0	0.00 - 1.00	2	3	4
0.2	0.00 - 1.00	2	3	4
0.4	0.00 - 0.11	3	4	4
	0.12 - 0.41	2	4	4
	0.42 - 0.98	2	3	4
	0.99 - 0.99	2	2	4
	1.00 - 1.00	2	2	3
0.6	0.00 - 0.09	4	4	5
	0.10 - 0.14	4	4	2
	0.15 - 0.21	2	4	2
	0.22 - 0.55	2	3	2
	0.56 - 1.00	2	2	2
0.8	0.00 - 0.08	4	5	6
	0.09 - 0.12	4	5	2
	0.13 - 0.15	4	3	2
	0.16 - 0.38	2	3	2
	0.39 - 0.78	2	2	2
	0.79 - 1.00	1	2	2
1.0	0.00 - 0.07	4	5	6
	0.08 - 0.12	4	5	2
	0.13 - 0.15	4	3	2
	0.16 - 0.51	2	3	2
	0.52 - 0.81	2	2	2
	0.82 - 1.00	1	2	2

Table 4.6: Single facility robust optimization: Hurwicz optimal actions at different uncertainty and optimism levels

Like MAXIMIN robust optimization, Hurwicz criterion based robust optimization can yield management policies that are significantly different from those provided by non-robust asset management. Optimal policies by uncertainty and optimism levels are presented in Table 4.6. For each of the different uncertainty levels, optimism levels of 0.00, 0.01, ... , 1.00 were set. The table groups together ranges of the optimism level that yielded similar optimal policies.

Note that as uncertainty increases, the range of policies that might be optimal increases, so the choice of optimism level becomes more important. Also note that the non-robust optimal policy is not chosen by the Hurwicz robust optimization for uncertainty levels greater than 0.4. This seems to indicate that the policy chosen by the non-robust optimization does a relatively poor job in terms of best and worst case transition probability matrices. This is logical since non-robust asset management does not consider best and worst-case scenarios, it only works with one set of transition probabilities.

Figure 4.4 presents the cost ranges of Hurwicz robust optimization. The Hurwicz robust optimization is able to reap the benefits of best-case transition probabilities, incurring near zero maintenance costs, but also able to limit the worst-case costs. In many ways, the cost ranges observed under this type of asset management offer a suitable compromise between the conservativeness of MAXIMIN style robust optimization and the optimism of MAXIMAX or non-robust schemes.




Figure 4.4: Single facility robust optimization: Cost ranges of Hurwicz (optimism = 0.5) robust asset management.

4.5 Conclusion

Robust optimization offers one way to mitigate against the effects of uncertainty in the deterioration models that underlie infrastructure management systems. Applying robust optimization to the management of a single infrastructure facilities can achieve significant cost savings. These savings can also be thought of as the costs associated with uncertainty in the transition probability matrices used in management systems that ignore uncertainty.

A bootstrap simulation of the computation of bridge deck deterioration transition probabilities, based on the statistical uncertainty surrounding parameters used in this process, revealed standard errors in the range of 0.2 to 0.4. Consideration of these standard errors supports uncertainty levels of 0.5 or greater. The computational study undertaken here reveals the potential for significant benefits associated with considering uncertainty of this magnitude.

It is worth noting that the work presented so far focuses upon the management of a single infrastructure facility, in this case a single lane-yard of pavement. Although the models presented here are benchmarked in a specific pavement management problem, they are generalizable to a wide range of problems regarding the management of different infrastructure assets. Furthermore, the models presented here may be modified to consider the management of a system of infrastructure assets.

Chapter 5

Robust Infrastructure

Management: At the System Level

In the previous chapter, robust optimization was applied to single facility infrastructure management. The results obtained in parametric studies demonstrated that robust optimization was able to limit excess costs associated with errors in infrastructure deterioration modeling. This chapter is an extension of the work of the previous chapter to the more complex system level case.

Recall that in the last chapter, several infrastructure management recursions were shown, all of which could be solved via dynamic programming. Also recall that in Chapter 2, it was discussed how dynamic programming becomes computationally intractable for the large scale problems associated with maintaining a system of infrastructure facilities related by a common management budget. A convex optimization problem was presented that aimed to capture a sophisticated approach to system level infrastructure management. In

this chapter, alternate robust optimization problems will be presented. Solution strategies will be discussed, and computational studies similar to those done in the last chapter of the text will be run.

5.1 The Maximin Criterion

As in the preceding section, first consider a MAXIMIN robust optimization problem. Recall that this involves taking the management actions to minimize costs, given that nature will select a worst-case transition probability matrix to maximize costs. Begin with the System Level Infrastructure Management Problem (SLIMP) of Section 2.1. Now modify this formulation so that the transition probability matrix p becomes a decision variable, with the worst-case (i.e. highest-cost) matrix in the uncertainty set \mathcal{P} considered during optimization.

System Level Robust Infrastructure Management Problem (SLRIMP)

$$\min_x \max_{p \in \mathcal{P}} \sum_{t=0}^T \alpha^t \left[\sum_{i \in I} \sum_{a \in A} [g(i, a) + u(i)] f_t(i) x_{i,t}(a) N \right] \quad (5.1)$$

subject to the following constraints:

- (1) $\sum_i \sum_a g(i, a) f_t(i) x_{i,t}(a) N \leq b_t \quad \forall t \in \{0, 1, 2, \dots, T\}$
- (2) $\sum_i \sum_a f_t(i) x_{i,t}(a) p(j|i, a) = f_{t+1}(j) \quad \forall j, t \in \{0, 1, 2, \dots, T\}$
- (3) $x_{i,t}(a) \geq 0 \quad \forall i, a, t \in \{0, 1, 2, \dots, T\}$

Assuming the uncertainty set \mathcal{P} is defined as before, all terms in any matrix $p \in \mathcal{P}$ must be within the uncertainty level δ of the corresponding terms in some initial estimate matrix q . This leads to the definition of \mathcal{P} used before, namely $\mathcal{P} = \{p \in I \times I \times A \rightarrow [0, 1] :$

$|p(j|i, a) - q(j|i, a)| \leq \delta$ ($\forall i, j \in I, a \in A$), $\sum_{j \in I} p(j|i, a) = 1$ ($\forall i \in I, a \in A$), $p(j|i, a) \geq 0$ ($\forall i, j \in I, a \in A$). It also leads to an alternate formulation of the System Level Robust Infrastructure Management Problem (SLRIMP).

Alternate Formulation(SLRIMP)

$$\min_x \max_p \sum_{t=0}^T \alpha^t \left[\sum_{i \in I} \sum_{a \in A} [g(i, a) + u(i)] f_t(i) x_{i,t}(a) N \right] \quad (5.2)$$

subject to the following constraints:

- (1) $\sum_i \sum_a g(i, a) f_t(i) x_{i,t}(a) N \leq b_t \quad \forall t \in \{0, 1, 2, \dots, T\}$
- (2) $\sum_i \sum_a f_t(i) x_{i,t}(a) p(j|i, a) = f_{t+1}(j) \quad \forall j, t \in \{0, 1, 2, \dots, T\}$
- (3) $x_{i,t}(a) \geq 0 \quad \forall i, a, t \in \{0, 1, 2, \dots, T\}$
- (4) $|p(j|i, a) - q(j|i, a)| \leq \delta \quad \forall j, i, a$
- (5) $\sum_{j \in I} p(j|i, a) = 1 \quad \forall i, a$
- (6) $p(j|i, a) \geq 0 \quad \forall j, i, a$

5.1.1 Solution Strategy

The formulations of the System Level Robust Infrastructure Management Problem shown above may appear quite complex. However, a few inferences can be made that point towards a solution strategy.

First note that transition probabilities capture deterioration and condition states often represent varying levels of facility decay. Seen in this light, it becomes intuitively clear that ‘worst-case’ conditions refer to situations where deterioration proceeds faster than anticipated. To achieve this effect, take the initial probability matrix estimate q and

shift probability from less decayed states to more decayed states in a manner similar to that employed in the single facility case. The amount of probability that can be shifted depends on the constraints of \mathcal{P} . This method typically provides the ‘worst-case’ transition probability matrix p .

Next note that fixing p makes it possible to find the least-cost management policy x , and vice-versa. In fact, the problem of finding the optimal x for a fixed p is identical to the (non-robust) System Level Infrastructure Management Problem of Section 2.1. So it is possible to form a worst-case transition probability matrix p and use this to find the optimal policy x . If it is not clear that the matrix p chosen is in fact the worst-case transition probability matrix, it is possible to fix the chosen policy x and then solve for the worst-case transition probability matrix. Continue iterating until a fixed management policy and transition probability matrix are found.

It is now clear that p represents the worst-case or highest-cost transition probability matrix for the management policy x . Similarly, x represents the optimal or least-cost management policy for the transition probability matrix p . Thus, all policies other than x must have at least as great worst-case costs and x must be the MAXIMIN robust optimal solution to the system level infrastructure management problem.

5.2 The Hurwicz Criterion

Computational studies relating to the management of single facilities demonstrated how using a robust approach based on the Hurwicz criterion to guide infrastructure maintenance decision-making could be worthwhile. This section generalizes the MAXIMIN ap-

proach to system level infrastructure management introduced in the previous section to create a methodology based on the Hurwicz criterion.

Use of the Hurwicz criterion requires consideration of best and worst case transition probability matrices, p^b and p^w respectively. Notice now that in the System Level Infrastructure Management Problem (SLIMP) it is necessary to keep track of $f_t(a)$ terms, the expected fractions of all facilities, at time t , in condition state a . Consider a two time step infrastructure management problem. The present condition of the facilities, as represented by $f_0(a)$ terms, is known. In the following year, the expected condition of the facilities depends upon transition probabilities. The Hurwicz criterion involves considering best and worst case conditions, so let $f_1^b(a)$ and $f_1^w(a)$ describe the expected conditions of the facilities in best and worst cases respectively. Then the management problem can be expressed:

System Level Hurwicz Robust Infrastructure Management Problem (SLHRIMP)

$$\min_x \max_{p^w \in \mathcal{P}} \min_{p^b \in \mathcal{P}} \sum_{i \in I} \sum_{a \in A} [g(i, a) + u(i)] [f_0(i)x_{i,0}(a) + (\beta f_1^b(i) + (1 - \beta)f_1^w(i))x_{i,1}(a)] N \quad (5.3)$$

subject to the following constraints:

$$(1.1) \quad \sum_i \sum_a g(i, a) f_0(i) x_{i,0}(a) N \leq b_0$$

$$(1.2) \quad \sum_i \sum_a g(i, a) f_1^b(i) x_{i,1}(a) N \leq b_1$$

$$(1.3) \quad \sum_i \sum_a g(i, a) f_1^w(i) x_{i,1}(a) N \leq b_1$$

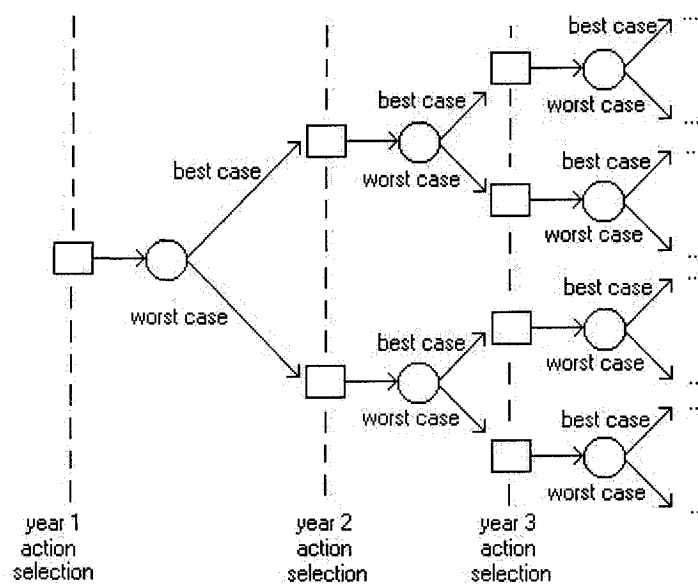
$$(2.1) \quad \sum_i \sum_a f_0(i) x_{i,0}(a) p^b(j|i, a) = f_1^b(j) \quad \forall j$$

$$(2.2) \quad \sum_i \sum_a f_0(i) x_{i,0}(a) p^w(j|i, a) = f_1^w(j) \quad \forall j$$

$$(3.1) \quad x_{i,t}(a) \geq 0 \quad \forall i, a, t \in \{0, 1\}$$

This mathematical program is fairly complex, and it is only for a two year management horizon. Consider a situation where the management horizon is extended to three years. In order to consistently apply the Hurwicz criterion, in the second year best and worst cases must be considered regardless of what happened between years one and two. So, it is necessary to consider best and worst case conditions in year two and in each of these situations best and worst case scenarios for year three. Thus, the number of possible scenarios for year three is 4 and in year t is 2^{t-1} . The complexity of Hurwicz criterion based decision-making increases exponentially with the length of the planning horizon. Figure 5.1 illustrates this point graphically.

Figure 5.1: System level robust optimization: The complexity of Hurwicz criterion based robust optimization



The increasing complexity of using the Hurwicz criterion in robust decision-making was not noticed during single facility infrastructure management because problems were formulated in terms of a recursion representing only a single year of management. In system level infrastructure management, as considered here, it is necessary to make decisions regarding maintenance throughout a (possibly long term) planning horizon at once. Despite its increasing complexity, as shown in Figure 5.1, the Hurwicz criterion remains an interesting approach to robust system level infrastructure management.

5.3 Computational Studies

In order to test the utility of the robust approaches to system level infrastructure management presented above, computational studies were run. As in the preceding chapter, these studies relate to the simulated management and deterioration of pavement. A network of 10,000 lane-yards of pavement will be managed for a period of 5 years. The different sections of pavement are assumed to be homogeneous with the same user and agency costs defined in the preceding chapter (see Table 4.1 and Table 4.2). The facilities to be managed are related by a joint management budget of \$25,000 that can be spent on MR&R activities each year. Again, the unusable condition state, state 1, is to be avoided. As in the case of the computational studies of the last chapter, emphasis is placed on tracking management policies selected by robust and non-robust optimization, as well as the expected costs these policies incur, in trials representing worst and best case conditions given various uncertainty levels.

5.3.1 Results

The actions prescribed by the MAXIMIN robust linear programming algorithm are able to achieve significant (user + agency) cost savings in worst-case conditions, when compared to traditional non-robust optimization. Figure 5.2 shows the accumulated five year management costs accrued to both the users of the system and the planning agency.

Figure 5.2: System level robust optimization: Worst case costs of maximin robust and non-robust asset management




Figure 5.2 shows the worst case costs of asset management, both using an approach that seeks to minimize those costs and the non-robust approach, as a function of uncertainty. If the uncertainty level exceeds 0.8 then the worst case cost of non-robust optimization is undefined. Given the available budget, non-robust optimization is unable to meet the level of service requirement that the worst condition state be avoided. At the same uncertainty levels, robust optimization is able to meet budget and service requirements. The reason is that the non-robust management scheme either chooses not to spend or inefficiently spends its available budget in the first years of asset management. The belief that decay occurs slower than it actually does leads to a crisis in later years when the planning agency suddenly realizes in one year that given its budget it cannot maintain the system above the minimum service requirement within budget.

In order to illustrate the failure of non-robust optimization algorithms in situations where there is great uncertainty, it would be beneficial to see the management policies recommended by the algorithms and the evolution of deterioration over time. Management policies for system level infrastructure management consist of fractions of all facilities, in given condition states at given points in time, to which to apply different management actions. A complete management policy is a matrix of size $|I| \times T \times |A|$ where $|I|$ is the number of condition states, T is the length of the planning horizon, and $|A|$ is the number of management actions possible. It is thus difficult to present numerous complete management policies succinctly.

That being said, it is possible to present information regarding management policy chosen and deterioration progression for one example computational study. Consider the

fractions of all facilities in different condition states and to which different actions will be applied for one example involving an uncertainty level of 0.6 and non-robust optimization.

Table 5.1 presents the results.

Year	Condition	Fraction	Action	Fraction	Agency costs	User costs
1	8	1.00	2	1.00	0.40	0.00
2	8	0.24	2	0.24	2.50	2.72
	7	0.16	3	0.31		
	6	0.60	4	0.45		
3	8	0.20	2	0.20	2.50	3.26
	7	0.15	3	0.52		
	6	0.56	4	0.28		
	5	0.09				
4	8	0.13	2	0.32	2.50	4.35
	7	0.12	3	0.47		
	6	0.47	4	0.21		
	5	0.28				
5	8	0.03	1	0.12	2.50	6.60
	7	0.09	2	0.42		
	6	0.33	3	0.46		
	5	0.49				
	4	0.02				
	3	0.04				

Table 5.1: System level robust optimization: Fractional Variables of Condition State and Actions Taken

Note that Table 5.1 shows a case where a non-robust optimization algorithm recommended not spending the entire available agency budget in year 1. Given an assumption of a medium deterioration rate and a desire to minimize user + agency costs, this is the logical thing to do. However, deterioration proceeds faster than anticipated. The condition of facilities rapidly deteriorates as user costs expand. Similar scenarios play out when

the uncertainty level were set to 0.8 or higher, but in these cases deterioration proceeds so rapidly that non-robust decision-makers are unable to prevent facilities from becoming unusable.

While MAXIMIN robust optimization yields lower costs in worst case conditions, non-robust optimization produces lower costs if decay proceeds as anticipated by the initial model. It is interesting to look at the additional cost of robust optimization in this scenario (shown in Figure 5.3), and in particular to compare this potential extra cost of robust management with the potential savings observed earlier in worst-case conditions (Figure 5.2).

Discontinuities in costs in Figure 5.3 are associated with the fact that at certain threshold levels of uncertainty, robust asset management schemes may decide to alter their maintenance schedule recommendations. The costs of asset management in the expected case were plotted on the same scale as the costs in the worst case shown before. Note that the potential savings achieved by robust optimization in worst case conditions are of a larger scale than the extra costs in expected case conditions. This was by no means guaranteed; MAXIMIN robust optimization does not even consider costs in expected case conditions. Next consider the total cost ranges possible under traditional and MAXIMIN robust asset management, as shown in Figure 5.4.

Figure 5.3: System level robust optimization: Expected case costs of maximin robust and non-robust asset management






Figure 5.4: System level robust optimization: Cost ranges of maximin robust and non-robust asset management



MAXIMIN robust optimization dramatically shrinks the range of potential costs associated with asset management in cases of decay rate uncertainty. Cost uncertainty may be undesirable to planning agencies, particularly given the political environment in which they operate. It is worth noting that the extra costs of MAXIMIN robust management as compared to traditional management in best case conditions are actually slightly greater than those seen earlier in expected case conditions. MAXIMIN robust management is more conservative than traditional asset management schemes and incurs excess costs more or less in proportion to how benevolent conditions are. Thus, consideration of the whole range of costs that are possible may suggest the use of Hurwicz style robust optimization.

The cost ranges of Hurwicz management are presented in Figure 5.5. Note that Hurwicz robust optimization is able to achieve worst-case costs comparable to those of MAXIMIN robust optimization. Furthermore, Hurwicz robust optimization is able to significantly lower best-case costs as compared to MAXIMIN robust optimization. These results are in line with those shown in the preceding chapter related to the management of a single infrastructure facility.

Figure 5.5: System level robust optimization: Cost ranges of hurwicz robust (for $\beta = 0.5$) and non-robust asset management



5.4 Conclusion

As in the case of single facility management, robust optimization was shown to offer potentially significantly lower costs of maintenance than non-robust optimization. In addition, the computational studies presented in this chapter illustrated how ignoring uncertainty and incorrectly estimating deterioration rates can lead to management failures.

The fact that robust optimization recommends significantly different ‘optimal’ management policies than non-robust optimization makes it clear that the policies recommended by infrastructure management systems are sensitive to the (possibly erroneous) assumptions of the systems. The potential cost savings of robust optimization in the computational studies shown here make it clear how important these policy differences are. It should be clear now that it is essential to consider the assumptions of infrastructure management systems before trusting their results.

It should also be clear now that it’s possible to consider uncertainty during decision-making via robust optimization. If traditional robust optimization techniques are viewed as being ‘too conservative,’ it has been shown how an alternate decision criteria, like the Hurwicz criterion, can be used. Furthermore, it is worth noting that the degree of conservatism of robust optimization approaches depends upon the manner in which uncertainty sets are defined. Robust infrastructure management using relatively small uncertainty sets provides results more or less analogous to non-robust infrastructure management.

One thing that may seem odd about the computational studies performed so far is that during simulated infrastructure management facilities are observed deteriorating, but this information is not used to update initial estimates regarding deterioration processes.

The work done so far may have left the reader under the impression that all infrastructure management systems in use today ignore the issue of uncertainty in deterioration modeling and do not keep track of how deterioration actually proceeds. This is untrue. Many modern infrastructure management systems use adaptive control approaches to reduce the uncertainty associated with imprecise initial models of facility deterioration. The following chapter describes state-of-the-art adaptive control formulations. Comparisons will be made to the robust optimization approaches discussed to this point, including computational studies and qualitative analyses comparing the performance of both approaches.

Chapter 6

Comparing Adaptive and Robust Infrastructure Management

6.1 Adaptive Control in Infrastructure Management

The developers of modern infrastructure management systems have recognized the presence of uncertainty in the deterioration models used in practice. Therefore, 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. For example, Harper et al (1990) use Bayesian methods to update the parameters of their deterioration models. Likewise, the popular bridge management system Pontis updates transition probability matrices over time (Golabi and Shepard, 1997). Durango and Madanat (2002) have proposed a decision support system where the uncertainty in the deterioration model is represented by a probability mass function of deterioration rates. Rather than updating

the parameters of the deterioration model, it is this probability mass function of deterioration rates that is updated in light of inspection data in their system. Irrespective of these differences, the common element in these three systems is that they are adaptive, in that they combine model updating and optimization.

Broadly speaking, there are two types of adaptive optimization routines. Open Loop Feedback Control methods alternate between updating model parameters and optimizing decision making with respect to the most recent estimates of parameters. Closed Loop Control improves on this methodology by explicitly considering the future updating of deterioration model parameters within present time MR&R optimization. Unfortunately, consideration of the many ways a network of facilities may deteriorate and how this will lead to different updated deterioration model parameters is not always possible within the framework of a solvable optimization problem. For this reason, adaptive infrastructure management systems that deal with a network of related facilities are typically based on Open Loop Feedback Control (OLFC) approaches.

In adaptive control, model uncertainty is often modeled by treating deterioration model parameters as continuous random variables. The successive updating of these parameters will improve the representation of the actual deterioration process only if the variables converge to their true values. For this to happen, a large number of observations have to be made for every combination of MR&R action performed and infrastructure facility condition state. This means that all MR&R activities must be applied to infrastructure facilities in every possible condition state a sufficient number of times. This may not happen in adaptive infrastructure management systems because the optimization process will tend

to select only a subset of MR&R activities to apply to each condition state. As a result, the deterioration model parameters relating to state-action pairs that are not selected a sufficient number of times may converge to incorrect values. This is a limitation of all OLFC-based adaptive optimization approaches.

However, the most serious limitation to the effectiveness of the adaptive control approach is probably that while managing a network of infrastructure facilities, data on condition, deterioration, and the effectiveness of different MR&R actions accumulates slowly. Thus, adaptive control approaches require a long time to improve the precision of the transition matrices and, during this time, will incur high costs associated with transition matrix uncertainty.

6.1.1 An Example Adaptive Control Approach

An adaptive control approach is here presented using a MDP formulation like those presented in section 2.1 of this paper. New condition information is used to recalculate estimates of transition probabilities (call these $\pi(j|i, a)$ terms). New condition information takes the form of a multinomial random variable; facilities are being observed as being in one of a discrete and mutually exclusive set of condition rating states. It is assumed that transition probability estimates follow the fairly general Dirichlet distribution, the conjugate distribution to the multinomial (as in Madanat et al (2004)). Then, updating transition probabilities becomes straightforward.

Let $n(j|i, a)$ represent the number of facilities that have been observed transitioning from state i to state j in the time step after having action a applied. Further, let $N(i, a)$ represent the number of facilities that have been observed in state i with action a applied

($N(i, a) = \sum_j n(j|i, a)$). Every time step, it would be possible to update the counts, the $n(j|i, a)$ and $N(i, a)$ terms.

Given these counts, the maximum likelihood estimates of transition probabilities are quite simple to calculate.

$$\pi(j|i, a) = \frac{n(j|i, a)}{N(i, a)} \quad (6.1)$$

These transition probability estimates can then be used in infrastructure management decision-making. For example, plug the $\pi(j|i, a)$ terms as the $p(j|i, a)$ terms in the Single Facility Infrastructure Management Problem (SFIMP). As information becomes available, $n(j|i, a)$ and $N(i, a)$ counts can be updated and new estimates calculated for $\pi(j|i, a)$ terms.

Figure 6.1: OLFC Adaptive control: Agency responsibilities over time.

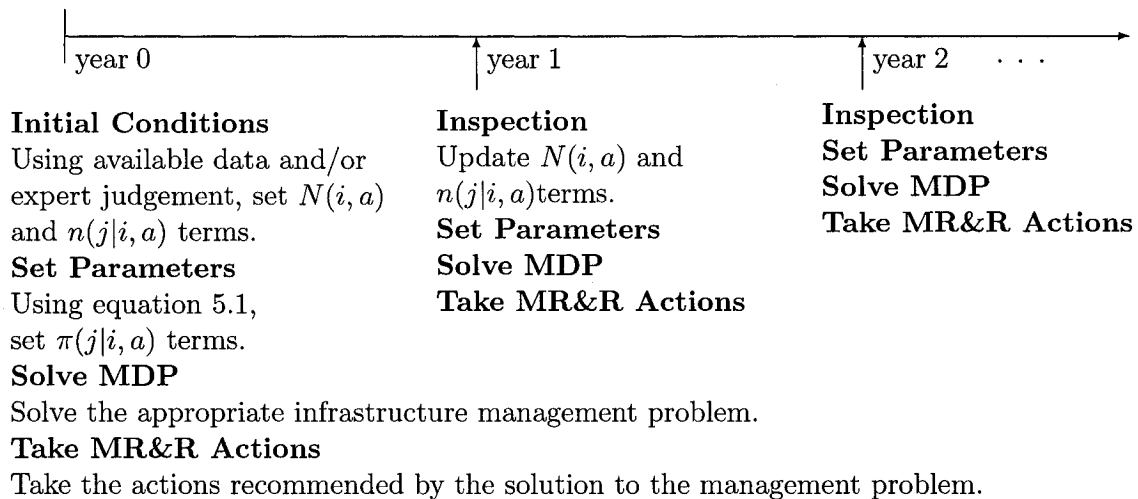


Figure 6.1 shows the responsibilities, over time, of a planning agency using an Open Loop Feedback Control (OLFC) adaptive control approach to infrastructure management. Planning agencies are responsible for taking MR&R actions, and often times are required by law to perform condition surveys that involve inspecting facilities. Many agencies already

use MDP formulations to determine which maintenance actions to perform. Thus, the only step that might be considered an additional responsibility is that of setting parameters given up-to-date counts of condition state transitions. It was already shown that this additional responsibility is quite simple computationally. This step really involves making use of the most up-to-date information available to help guide decision-making, and should not be skipped over.

Note that in Figure 6.1 the possibility of setting $N(i, a)$ and $n(j|i, a)$ terms using ‘expert judgement’ was raised. Harper et al (1990) have written about how expert judgement may be a valuable source of information to use in setting transition probabilities of MDP formulations of infrastructure management problems. As data becomes available, it makes perfect sense to seek some mechanism for combining expert judgement and empirical data. The adaptive control approach shown above leaves open this possibility.

Consider initially setting $N(i, a)$ terms to reflect how much weight is given to the initial assumptions of expert engineers. Larger values for these terms will mean a set amount of new data will have less impact on estimated transition probabilities. Smaller values for $N(i, a)$ terms will give less weight to expert judgement vis a vis new empirical data. Next, consider setting $n(j|i, a)$ terms so that initial estimates of transition probabilities, $\pi(j|i, a) = \frac{n(j|i, a)}{N(i, a)}$ terms reflect expert judgement regarding the likelihood of different transitions occurring.

6.1.2 Systematic Probing

The approach described above learns about infrastructure deterioration by observing facility condition information and relating this to the information regarding maintenance

actions performed the previous time step. No information will be gathered with regards to maintenance actions perceived to be sub-optimal during decision-making, and estimates of the costs and benefits of these actions will remain unchanged. This, in turn, will keep these alternate maintenance actions unattractive during decision-making.

In order to learn about alternate policies not initially favored by robust optimization, it is necessary to ‘probe’ the system. This implies taking actions believed to be sub-optimal given current information in order to learn more about the system in question. For example, it may be beneficial to make decisions using a ϵ -greedy policy in combination with the adaptive control approach presented previously. This involves taking actions consistent with the results of a management problem $100-\epsilon$ percent of the time and selecting a random action (all possible actions equally likely to be selected) ϵ percent of the time where ϵ is some small number. Alternate strategies are also possible; ϵ -greedy represents only one (simple) strategy of systematic probing. Systematic probing involves balancing exploitation, the desire to perform actions currently believed to be best, with exploration, the desire to test the efficacy of various actions and determine which actions truly are best.

6.2 Computational Studies

Adaptive control and robust optimization both represent strategies for mitigating the impacts of deterioration model uncertainty during infrastructure management. It therefore makes sense to study and compare the performance of each in computational studies simulating infrastructure management.

The studies of the preceding two chapters investigated the performance of robust

optimization assuming it was possible to correctly ascertain the uncertainty level surrounding estimated transition probabilities. This assumption may be unrealistic and likely produces results that favor robust optimization over alternate techniques unable to make use of this assumed accurate knowledge of levels of uncertainty. In the computational studies that are presented here, robust optimization strategies with various assumed levels of uncertainty, none fully correct, are studied.

Similarly, adaptive control methodologies with various ϵ -greedy strategies are studied. It should be remembered that adaptive control methodologies need to consider the problem of balancing the desire to explore with the desire to exploit, and need to formulate a strategy for doing so. ϵ -greedy strategies of the type used here represent only one type of strategy out of many.

The computational studies presented here once again involve pavement maintenance. In Chapter 3, Tables 4.1 and 4.2 identified the costs of pavement maintenance activities and pavement deterioration. Table 4.3 outlined three rates of pavement deterioration, described only as being fast, medium, and slow. For the computational studies done here, four scenarios are considered. In the first scenario, the 'real' deterioration rate used to simulate decay is the fast rate, but the initial assumptions of the 'decision-maker' are that deterioration will be slow. Similarly, scenario two involves a real deterioration rate characterized as slow and an assumption that deterioration will be fast. Scenarios three and four involve initial assumptions that deterioration will be medium coupled with real deterioration that is fast and slow, respectively. The different scenarios are chosen because they represent a mix of over and under estimation of deterioration rates, as well as of different

degrees of error in initial assumptions.

In all cases, a pavement network is managed and user + agency costs are minimized in the absence of any strict agency budgets. 60% of the network is assumed to start the planning exercise in (best) condition state 8, 20% in state 7, 10% in state 6, 5% in state 5, 3% in state 4, and 2% in state 3. Simulations of management are run, for each of several adaptive control and robust optimization methodologies. Simulations are necessary because different realizations of deterioration lead to different assumptions used in, and results of, decision-making algorithms, which in turn impact future observations of deterioration. The overall process is complex and stochastic and the true expected costs of different management strategies cannot be calculated using a closed-form expression.

6.2.1 Results

The results, in terms of costs of infrastructure management, of the computational studies are presented on the following few pages. Graphs of the costs of adaptive control management mechanisms, with various values for ϵ , the probing fraction, are presented above graphs of the costs of robust optimization management, with various assumed uncertainty levels. These sets of graphs are presented for the four scenarios outlined in the preceding section. A solid line on each the graphs indicates the (low) expected cost of maintenance given perfect information (a model of deterioration known to be correct at the start of the planning exercise).

Figure 6.2: Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = fast, assumed deterioration = slow




Figure 6.3: Adaptive vs. Robust: Cost ranges of robust management, deterioration = fast, assumed deterioration = slow




Figure 6.4: Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = slow, assumed deterioration = fast

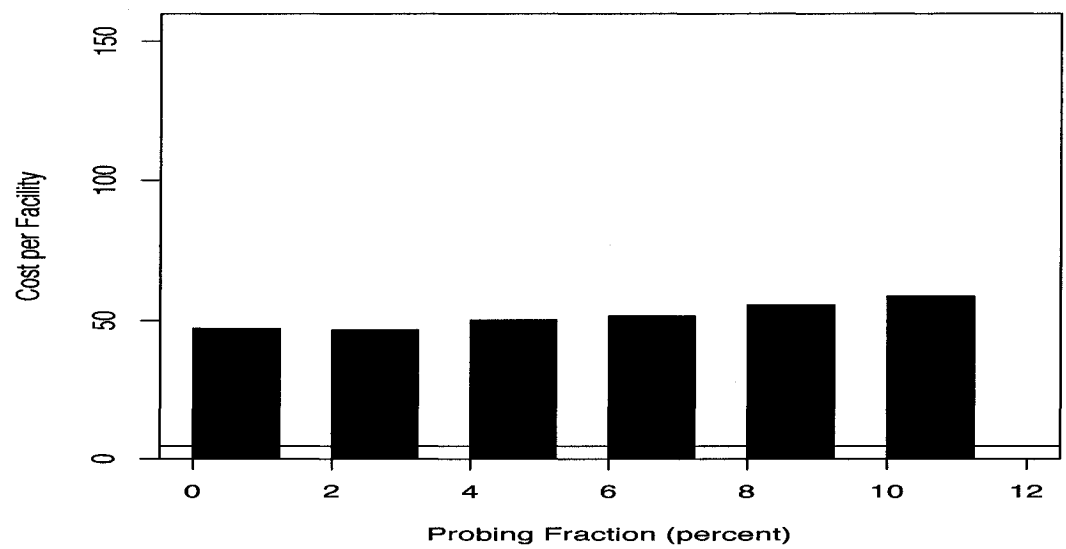


Figure 6.5: Adaptive vs. Robust: Cost ranges of robust management, deterioration = slow, assumed deterioration = fast




Figure 6.6: Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = fast, assumed deterioration = medium

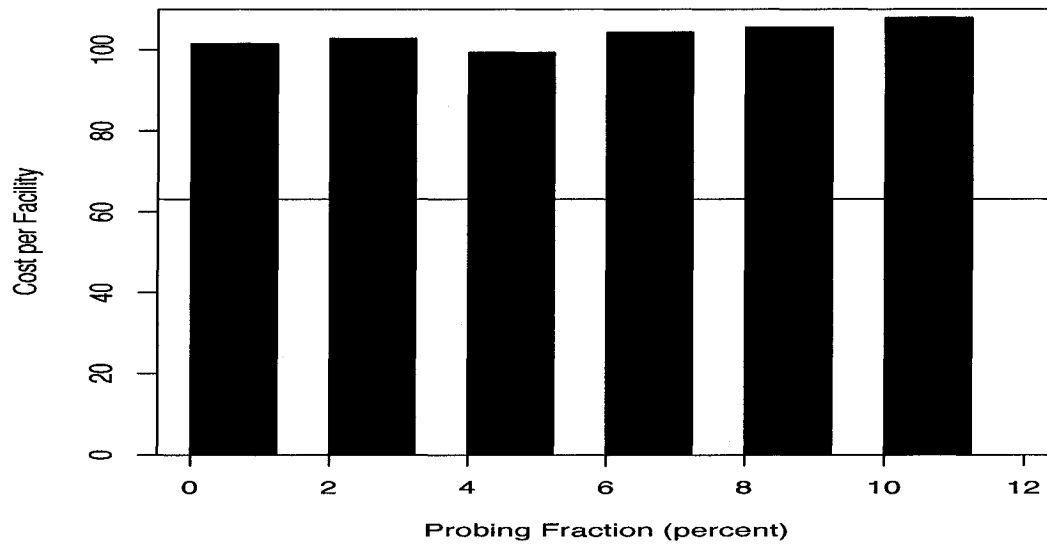


Figure 6.7: Adaptive vs. Robust: Cost ranges of robust management, deterioration = fast, assumed deterioration = medium

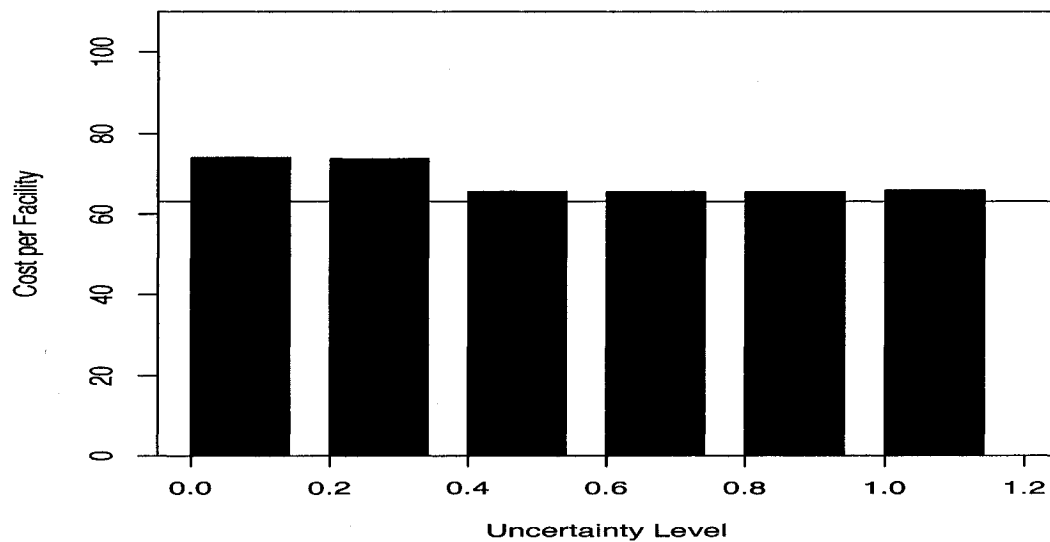


Figure 6.8: Adaptive vs. Robust: Cost ranges of adaptive control management, deterioration = slow, assumed deterioration = medium





Figure 6.9: Adaptive vs. Robust: Cost ranges of robust management, deterioration = slow, assumed deterioration = medium



Approaches based on both adaptive control and robust optimization do not perform as well as those that correctly estimate deterioration. Clearly the issue of adaptive control versus robust optimization is irrelevant if deterioration processes are well understood.

It is not immediately clear which probing fractions produce the best results in adaptive control methodologies. Setting the probing fraction, and indeed the systematic probing strategy, is an art and not a science. On the other hand, some clear trends are observable in relating the performance of robust optimization methodologies to the values for uncertainty level. In cases where deterioration proceeds faster than anticipated, Graphs 6.3 and 6.7, strategies involving larger uncertainty levels outperform those involving smaller uncertainty levels. On the other hand, in cases where deterioration is slower than anticipated, Graphs 6.5 and 6.9, the opposite is true.

Note that the costs of robust optimization, using a fairly large uncertainty level, are significantly less than those of adaptive control in comparing Graphs 6.2 and 6.3 or Graphs 6.6 and 6.7, when real deterioration is characterized as fast but initial assumptions were that deterioration would be slow or medium. As was mentioned previously, robust optimization leads to fairly conservative management strategies that involve significantly more regular maintenance of facilities than non-robust optimization might recommend. This approach saves a lot in user costs when deterioration is faster than initially anticipated.

The flip side to this argument is that robust optimization wastes money in excess agency costs when deterioration is slower than initially anticipated. Note how robust optimization, particularly when using a fairly large uncertainty level, can cost significantly

more than adaptive control as in comparing Graphs 6.4 and 6.5.

It is interesting to note that acting under the naive assumption that initial assumptions are correct often produces lower management costs as compared to adaptive control. (Acting naively is here shown as using robust optimization with a 0% uncertainty level.) In the case studies shown here, initial assumptions were discarded once data was obtained regarding deterioration. It may be better to incorporate initial assumptions and weight them as equivalent to a certain number of empirical observations, in a manner similar to that discussed after Figure 6.1.

Taken together, the graphs presented on the preceding few pages do not uniformly favor adaptive control over robust optimization or vice-versa. Apart from the costs incurred, there are important differences between the two approaches.

6.3 Qualitative Comparisons

Adaptive control is explicitly learning about the system it is managing, while robust optimization is not. Thus, at the end of the simulated management exercises presented above adaptive control will have a more accurate representation of the underlying true transition probability matrix than will robust optimization.

However, it is important to recognize that in the computational studies there was always a true underlying transition probability matrix that defined deterioration. There was no epistemic uncertainty of the form that robust optimization is well suited to considering. Adaptive control methodologies necessarily assume a functional form to define deterioration, and then assume another functional form for the distribution of estimates of parameters

of deterioration models. In the computational studies examined in the preceding section, the assumptions of the adaptive control methodologies were more or less correct. One of the benefits of robust optimization based approaches is that it is typically not necessary to make assumptions regarding the distribution of estimates of deterioration model parameters. Furthermore, the assumed functional form of the deterioration model itself becomes less important since parameters are not considered fixed but rather as belonging to (possibly large) uncertainty sets.

While the assumptions made by adaptive decision-makers in the preceding section were more or less correct, the primary assumptions made by robust decision-makers regarding uncertainty levels and uncertainty sets were incorrect. The uncertainty surrounding an initial assumption of a fast deterioration rate when deterioration could actually be slow, or vice-versa, is not the same as saying there is an uncertainty level of 0.2 or 0.4.

It is interesting to compare the necessity of specifying an uncertainty level in robust optimization and that of possibly specifying a weight to give to initial assumptions in adaptive control. In both cases, this relates to noting how much faith should be placed in expert judgement or initially available data regarding deterioration. It is clear that in problems of infrastructure management involving significant uncertainty, it is necessary, at the start, to estimate how much weight should be given to initial estimates of deterioration model parameters. The agreement between robust optimization and adaptive control based approaches on this matter can lead to speculation on the possibility of combining the two approaches.

6.4 Towards A Hybrid Management System

The best approach to dealing with situations where there is significant uncertainty regarding facility deterioration may be to combine adaptive deterioration modeling with robust decision-making. In this way, new data is used to re-evaluate deterioration model parameters and limit the future magnitude of uncertainty while decision-making is done recognizing the current extent of uncertainty. Such a hybrid approach is possible given that robust control methodologies focus on changing decision-making algorithms, while adaptive control approaches are concerned with updating deterioration models after decisions have been made.

The basic idea of the hybrid approach is simple. Whenever information regarding the deterioration of infrastructure facilities is collected, update the statistical models of deterioration used in decision-making. Particular attention must be paid to quantifying the uncertainty surrounding the terms of the deterioration models. Then this information can be passed to a robust optimization routine, which will provide infrastructure management strategies that are optimal given current levels of uncertainty.

6.4.1 An Example Hybrid Approach

Return now to the example adaptive control approach introduced in this chapter where it was assumed that estimated rows of a transition probability matrix follow the Dirichlet distribution. Given this assumption, standard errors associated with transition probability estimates may be calculated.

$$\sigma(j|i, a) = \sqrt{\frac{n(j|i, a)(N(i, a) - n(j|i, a))}{(N(i, a))^2(N(i, a) + 1)}} \quad (6.2)$$

Standard error terms can be used to describe statistical or parametric uncertainty, uncertainty related to the size of the data sets used to come up with parameter estimates. If it is believed that epistemic uncertainty is sizable, it would be wise to sum terms representing this more fundamental, underlying uncertainty with terms related to standard errors. For example, let $\gamma(j|i, a)$ be a term estimating the epistemic uncertainty surrounding the transition probability estimate $\pi(j|i, a)$. Then one robust mathematical program that could be used in a hybrid approach to infrastructure management would be:

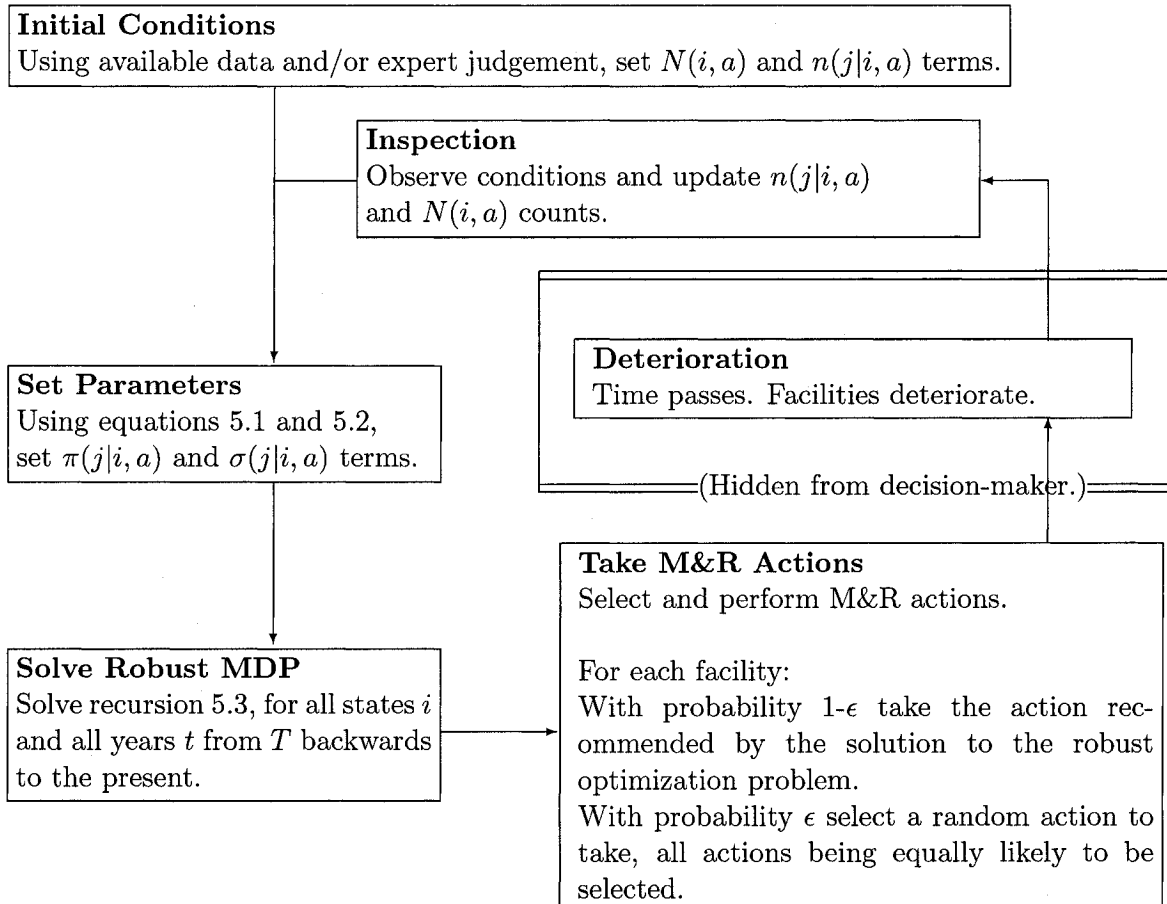
$$v_t(i) = \min_{a_i^t} \max_p \left[g(i, a_i^t) + u(i) + \alpha \left[\sum_j p(j|i, a_i^t) v_{t+1}(j) \right] \right] \quad (6.3)$$

with the added constraints that:

- (1) $|p(j|i, a) - \pi(j|i, a)| \leq 2\sigma(j|i, a) + \gamma(j|i, a) \quad \forall j, a$
- (2) $p(j|i, a) \geq 0 \quad \forall j, a$
- (3) $\sum_j p(j|i, a) = 1 \quad \forall a$

In flowchart form, the hybrid approach that has been outlined above, including systematic probing, could be represented as in Figure 6.10. Note that parameters are updated whenever inspections take place; between condition surveys, robust infrastructure management would be used without adaptive parameter estimates.

Figure 6.10: Adaptive and robust infrastructure management: Flowchart of agency responsibilities over time.



A hybrid approach to infrastructure management, like the one presented above, appears to provide an excellent way to deal with situations where there is significant deterioration rate uncertainty. As an adaptive approach, this methodology will allow an infrastructure management system to learn about deterioration and how fast it progresses. As a robust approach, this methodology will be less sensitive to errors in initial assumptions regarding deterioration processes and to errors in assumptions related to the distribution of estimates of deterioration model parameters. The extent of uncertainty will be reduced over time, while decision-making is always done recognizing current levels of uncertainty.

This is not to say that no concerns arise in combining adaptive control and robust optimization. Recall that in adaptive control there was some concern about the fact that information is only gathered on the effectiveness of actions that have actually been performed. This may be even more of a problem in cases where robust optimization is combined with adaptive control. Worst case costs are used to guide management decisions. The sizes of uncertainty sets used to come up with worst case scenarios relate directly to how much information is available. Thus, the impacts of actions not taken will be highly uncertain, leading to high estimates of worst-case costs, and sub-optimality in the eyes of a robust decision-maker. One effect of using robust decision-making in the context of adaptive control may be to strengthen the chain involving the (possibly incorrect) assumption that an action is sub-optimal, the failure to take that action, the failure to gather more information about that action, and the continued assumed sub-optimality of that action.

It's not clear if the concern raised above represents a serious detriment to the ability to combine adaptive control and robust optimization based methodologies for infrastructure management. It may be that the interaction between the two methodologies is insignificant compared to the basic problem of balancing exploitation and exploration in an adaptive control problem. Further study in this area would be worthwhile. It may be that a hybrid approach is the best way to manage infrastructure in situations where there is significant deterioration model uncertainty.

Chapter 7

Conclusion

The past few chapters of this dissertation have presented results obtained during several different research efforts. It is worthwhile to briefly summarize the findings of these chapters and re-examine their relation to one another.

This work began with an introduction to the field of infrastructure management and specifically to the problems of infrastructure deterioration modeling and infrastructure maintenance decision-making. Examples of mathematical formulations of both deterioration modeling and maintenance decision problems were presented. Next came a section detailing the problem of uncertainty in infrastructure deterioration modeling. Explanations were provided as to why there is reason to believe that there is uncertainty in deterioration modeling. Evidence was brought forward to show that, in fact, there is some proof of this uncertainty. It was hypothesized that this issue of uncertainty might be important in infrastructure maintenance decision-making, despite the fact that it was often ignored. This issue then became the problem that motivated the analyses and research efforts presented

in the remaining chapters of this dissertation.

The next two chapters of the work showed that considering uncertainty during decision-making can lead to vastly different recommended maintenance actions and management costs as compared to those that result from ignoring uncertainty. Single facility and system level problems were formulated and computational studies relating to the simulated management of both single facilities and systems of facilities were performed. Many of the results presented in this part of the text can be interpreted as showing the maximum potential costs of ignoring uncertainty. The form of consideration of uncertainty used, robust optimization, was also presented as a potential mechanism for infrastructure management in the presence of uncertainty. It is worth noting that this is the first time that robust optimization has been studied in the context of infrastructure management. A Hurwicz criterion based form of robust optimization was offered as an alternative to the possibly conservative MAXIMIN criterion form most commonly used in the field of Operations Research.

The chapter directly before this conclusion compared robust optimization to the method currently favored for use in infrastructure management in the presence of uncertainty, adaptive control. A sample adaptive control formulation was presented and computational studies were run comparing the performance of adaptive control and robust optimization. Pains were taken to ensure computational studies run did not favor either approach. It was found that neither robust optimization nor adaptive control based methodologies consistently outperform each other. There are interesting qualitative differences between the two approaches, and these were discussed. Finally, the possibility of combining the two methodologies was investigated and a hybrid approach was formulated. There was some

discussion of what the strengths and weaknesses of such a hybrid approach might be.

7.1 Future Directions

The research that was presented here recommends several areas for future study. A few of the most promising and pragmatic areas of study are briefly outlined below.

Statistical Testing of Deterioration Model Uncertainty

The research that was presented in this dissertation included a section labeled ‘Quantifying Uncertainty’ where some statistical testing was undertaken in an attempt to estimate the magnitude of the uncertainty surrounding transition probabilities used in infrastructure management systems. The most memorable data from this section resulted from a bootstrap analysis that placed uncertainty bounds around one particular estimated transition probability. This analysis was sufficient to prove that even in state of the art deterioration models, there exists significant uncertainty. However, if a robust optimization scheme were really to be used in practice, it would be necessary to come up with more sophisticated descriptions of uncertainty.

Bootstrap analyses, similar to the one presented in this dissertation, could be performed on the parameters of all sorts of different deterioration models relating to different infrastructure facilities. Alternate non-parametric methods, like the jackknife could be investigated, as could more traditional parametric techniques. Emphasis could be placed on the relationship between the uncertainty surrounding different but related deterioration model parameters.

Applying Different Robust Optimization Formulations to Infrastructure Management

Assuming a detailed description of deterioration model uncertainty were to be found, it would be useful to have a robust optimization formulation of the infrastructure management problem that could take advantage of this increased resolution. The robust optimization formulations presented in this dissertation assumed a ‘box’ model of uncertainty where true transition probabilities were assumed to lie in a closed range of possible values. The most advanced robust optimization models, like those found in El Ghaoui (2003) allow for consideration of how the uncertainty surrounding different transition probability estimates are related.

It should be noted at this point that it remains to be seen whether engineers working at planning agencies would be comfortable with the added complexity of more detailed descriptions of uncertainty and robust optimization formulations. This issue too is worthy of further investigation. It is important to know what planning agencies perceive as the strengths and weaknesses of infrastructure management systems in use today, in order to best plan how to modify these systems to improve performance.

Testing the Hybrid Approach

A hybrid, adaptive and robust, approach to infrastructure management was presented in the latter sections of this text. Issues were raised as to whether this approach would offer any material benefit over alternate adaptive or robust approaches. In order to gain a better understanding of the performance of hybrid style approaches, it would be

beneficial to test different hybrid formulations against simulated infrastructure management problems. Special attention should be paid to how well the approaches estimate uncertainty surrounding deterioration model parameters. It is imperative to the success of any robust methodology that uncertainty be estimated accurately. Special attention should also be paid to how well the approaches learn how deterioration progresses. The promise of the adaptive approach is that it learns, but this learning is predicated on the fulfillment of certain assumptions that may or may not hold.

7.2 Final Thoughts

Statistician Leo Breiman once noted that he was “deeply troubled by the current and past use of data models in applications, where quantitative conclusions are drawn and perhaps policy decision are made” since “conclusions are about the model’s mechanism, and not about nature’ mechanism” (Breiman, 2001). This dissertation has shown that there are indeed errors in even state of the art models used to define infrastructure deterioration. This is worrying because these models are used to provide quantitative conclusions regarding infrastructure maintenance. However, alternate methodologies are available for guiding infrastructure maintenance decision-making that are less sensitive to uncertainty or error in model formulation.

It is hoped that this work convinces those in the field of infrastructure management, as well as in other fields, to test the assumptions of the models they use to simplify complex real world phenomenon. Even in situations where little or no empirical data is available, robust methodologies can provide logical mechanisms to aid decision-making. Adaptive

methodologies can provide frameworks where new data is used as soon as it is collected and the consideration of the desire to 'explore' the effectiveness of alternate actions can be brought into the discussion of which actions are 'optimal' to take. A combined hybrid and robust approach may yet provide the most logical mechanism for making decisions in the presence of uncertainty. The number of important decisions made in the presence of uncertainty, and the promise of this research, is nearly limitless.

Bibliography

American Association of State Highways and Transportation Officials (1993) *AASHTO Guide for Design of Pavement Structures*, AASHTO, Washington D.C.

Archilla, A., and Madanat, S. (2001) Estimation of Rutting Models by Combining Data from Different Sources. *Journal of Transportation Engineering*, 127(5), 379-389.

Averbakh, I. (2000) Minmax Regret Solutions for Minimax Optimization Problems with Uncertainty. *Operations Research Letters*, 27(2), 57-65.

Bertsimas, D. and Sim, M. (2004) The Price of Robustness. *Operations Research*, 52(1), 35-53.

Breiman, L. (2001) Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3), 199-231.

Carnahan, J.V., Davis, W.J., Shahin, M.Y., Kean, P.L. and Wu, M.I. (1987) Optimal Maintenance Decisions for Pavement Management, *Journal of Transportation Engineering*, **113**(5), 554-572.

- Chan, W.T., Fwa, T.F., and Tan, J.Y. (2003) Optimal Fund-Allocation Analysis for Multidistrict Highway Agencies. *Journal of Infrastructure Systems*, 9(4), 167-175.
- Congressional Budget Office (1999) Trends in Public Infrastructure Spending. Washington, DC. Available at <http://www.cbo.gov/showdoc.cfm?index=1256&sequence=0>.
- DeStefano, P. and Grivas, D. (1998) Method for Estimating Transition Probability in Bridge Deterioration Models. *Journal of Infrastructure Systems*, 4(2), 56-62.
- Durango, P.L. and Madanat, S. (2002) Optimal Maintenance and Repair Policies in Asset Management Under Uncertain Facility Deterioration Rates: An Adaptive Control Approach. *Transportation Research*, 36A, 763-778.
- El Ghaoui, L. (2003) *A Tutorial on Robust Optimization*. presentation at the Institute of Mathematics and Its Applications, March 11, 2003. Available at <http://robotics.eecs.berkeley.edu/elghaoui>.
- Ellis, H., Jiang, M., and Corotis, R., (1995) Inspection, Maintenance, and Repair with Partial Observability. *Journal of Infrastructure Systems*, 1(2), 92-99.
- Ferreira, A., Antunes, A., and Picado-Santos, L., (2002) Probabilistic Segment-linked Pavement Management Optimization Model. *Journal of Transportation Engineering*, 128(6), 568-577.
- Federal Highway Administration (1979) Recording and Coding Guide for Structure Inventory and Appraisal of the Nation's Bridges. US Department of Transportation. Washington, DC.

- Flintsch, G. and Zaniewski, J. (1997) Expert Project Recommendation Procedure for Arizona Department of Transportation's Pavement Management System. *Transportation Research Record*, 1592, 26-34.
- General Accounting Office (2001) U.S. Infrastructure: Funding Trends and Federal Agencies' Investment, Estimates. *GAO-01-986T*. Available at <http://www.gao.gov/new.items/d01986t.pdf>.
- Golabi, K., Kulkarni, R., and Way, G. (1982) A Statewide Pavement Management System. *Interfaces*, 12(6), 5-21.
- Golabi, K. and Shepard, R. (1997) Pontis: A System for Maintenance Optimization and Improvement of US Bridge Network. *Interfaces*, 27 (1), 71-78.
- Gopal, S. and Majidzadeh, K. (1991) Application of Markov Decision Process to Level-of-Service-Based Systems. *Transportation Research Record*, 1304, 12-17.
- Harper, W., Majidzadeh, K. and Hurst (1990) Application of Expert Opinion in Customized Pavement Management Systems. *Highways and Data Processing*, 1304, 549-557.
- Hurwicz, L. (1951) Optimality Criteria for Decision Making Under Ignorance. *Discussion Paper: Statistics*, 370, Cowles Commission.
- Jido, M., Otazawa, T., and Kobayashi, K. (2004) Optimal Repair and Inspection Rules under Uncertainty. working paper, Kyoto University.
- Kulkarni, R. (1984) Dynamic Decision Model for a Pavement Management System. *Transportation Research Record*, 997, 11-18.

- Li, Y., and Madanat, S. (2002) A Steady-State Solution for the Optimal Pavement Resurfacing Problem. *Transportation Research Part A*, 36(6).
- Liu, F., and Wang, K. (1996) Pavement Performance-Oriented Network Optimization System. *Transportation Research Record*, 1524, 86-93.
- Madanat, S. and Ben-Akiva, M. (1994) Optimal Inspection and Repair Policies for Infrastructure Facilities. *Transportation Science*, 28(1), 55-62.
- Madanat, S., Park, S. and Kuhn, K. (2004) Adaptive Optimization and Systematic Optimization of Infrastructure System Maintenance Policies under Model Uncertainty. accepted for publication in *Journal of Infrastructure Systems*.
- Mauch, M. and Madanat, S. (2001) Semiparametric Hazard Rate Models of Reinforced Concrete Bridge Deck Deterioration. *Journal of Infrastructure Systems*, 7(2), 49-57.
- Mishalani, R. and Madanat, S. (2002) Computation of Infrastructure Transition Probabilities Using Stochastic Duration Models. *Journal of Infrastructure Systems*, 8(4), 139-148.
- Nilim, A., and El Ghaoui, L. (2004) Robust Solutions to Markov Decision Problems with Uncertain Transition Matrices. *EECS ERL Memo MO 04/26*.
- Ouyang, Y., and Madanat, S. (2004) Optimal Scheduling of Rehabilitation Activities for Multiple Pavement Facilities: Exact and Approximate Solutions. *Transportation Research Part A*, 38, 347-365.
- Paterson, W.D.O., (1993) A Standard Surface Integrity Index of Pavement Condition: Definition and Measurement Procedure. World Bank internal paper.

- Prozzi, J.A. and Madanat, S. (2004) Development of Pavement Performance Models by Combining Experimental and Field Data. *Journal of Infrastructure Systems*, 10(1), 9-22.
- Ramaswamy, R. and Ben-Akiva, M. (1900) Estimation of Highway Pavement Deterioration from In-Service Pavement Data. /it Transportation Research Record, 1272, 96-106.
- Soyster, A.L. (1973) Convex Programming with Set-Inclusive Constraints and Application to Inexact Linear Programming. *Operations Research*, 21, 1154-1157.
- Wang, K., Zaniwski, J., and Delton, J. (1994) Analysis of Arizona Department of Transportation's New Pavement Network Optimization System. *Transportation Research Record*, 1455, 91-100.
- Wang, K., and Liu, F. (1997) A Fuzzy Set Based and Performance Oriented Pavement Network Optimization System. *Journal of Infrastructure Systems*, 3(4), 154-159.