OPTIMAL REPLACEMENT STRATEGIES FOR WIND ENERGY SYSTEMS

by

John A. Flory

B.S., Colorado School of Mines, 2000

M.S., Air Force Institute of Technology, 2006

M.S., Rochester Institute of Technology, 2007

Submitted to the Graduate Faculty of

the Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2013

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

John A. Flory

It was defended on

June 24, 2013

and submitted to

Jeffrey P. Kharoufeh, Ph.D., Associate Professor, Department of Industrial Engineering

Lisa M. Maillart, Ph.D., Associate Professor, Department of Industrial Engineering

Bryan A. Norman, Ph.D., Associate Professor, Department of Industrial Engineering

Nagi Gebraeel, Ph.D., Associate Professor, School of Industrial and Systems Engineering

Dissertation Director: Jeffrey P. Kharoufeh, Ph.D., Associate Professor, Department of Industrial Engineering

OPTIMAL REPLACEMENT STRATEGIES FOR WIND ENERGY SYSTEMS

John A. Flory, Ph.D.

University of Pittsburgh, 2013

Motivated by rising energy prices, global climate change, escalating demand for electricity and global energy supply uncertainties, the U.S. government has established an ambitious goal of generating 80% of its electricity supply from clean, renewable sources by 2035. Wind energy is poised to play a prominent role in achieving this goal as it is estimated that 20% of the total domestic electricity supply can be reliably generated by land-based and offshore wind turbines by 2030. However, the cost of producing wind energy remains a significant barrier with operating and maintenance (O&M) costs contributing 20 to 47.5% of the total cost of energy. Given the urgent need for clean, renewable energy sources, and the widespread appeal of wind energy as a viable alternative, it is imperative to develop effective techniques to reduce the O&M costs of wind energy.

This dissertation presents a framework within which real-time, condition-based data can be exploited to optimally time the replacement of critical wind turbine components. First, hybrid analytical-statistical tools are developed to estimate the current health of the component and approximate the expected time at which it will fail by observing a surrogate signal of degradation. The signal is assumed to evolve as a switching diffusion process, and its parameters are estimated via a novel Markov chain Monte Carlo procedure. Next, the problem of optimally replacing a critical component that resides in a partially-observable environment is addressed. Two models are formulated using a partially-observed Markov decision process (POMDP) framework. The first model ignores the cost of turbine downtime, while the second includes this cost explicitly. For both models, it is shown that a threshold replacement policy is optimal with respect to the cumulative level of component degradation. A third model is presented that considers cases in which the environment is partially observed and degradation measurements are uncertain. A threshold policy is shown to be optimal for a special case of this model. Several numerical examples will illustrate the main results and the value of including environmental observations in the wind energy setting.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

I am indebted to many people who helped bring this doctoral dissertation to fruition. First and foremost, I owe much gratitude to my Ph.D. advisor, Dr. Jeff Kharoufeh, for his careful attention to my research during the past four years. His guidance and feedback were instrumental in my development as a doctoral student, and his numerous comments on this document are greatly appreciated. Additionally, I would like to thank the members of my dissertation committee, Drs. Lisa Maillart and Bryan Norman of the University of Pittsburgh and Dr. Nagi Gebraeel of the School of Industrial and Systems Engineering at Georgia Tech, whose careful reading and constructive comments have significantly improved this dissertation. Special thanks are due to Dr. Gebraeel for graciously sharing his bearing vibration data with me. This work would not have been possible without the support of a grant from the National Science Foundation (CMMI-0856702).

For moral support and friendship, I am grateful to all my fellow Ph.D. students whom I have met during the course of my studies. Getting to know such a diverse group of people was truly enriching. Finally, I am most indebted to my parents, John and Alma Flory, for their support, not only during the Ph.D. program, but throughout the various twists and turns of my life.

1.0 INTRODUCTION

1.1 BACKGROUND AND MOTIVATION

A *wind turbine* is a device designed to convert wind (kinetic energy) into mechanical energy in order to generate electricity. Although the concept of wind-generated electricity has a long history, only recently have economic and political conditions emerged that foster widespread efforts to establish competitive parity between wind energy and energy generated using fossil fuels. For example, the United States (U.S.) government established an ambitious goal of generating 80% of its electricity supply from clean, renewable sources (e.g., wind, solar, nuclear, geo-thermal, and others) by the year 2035 to ameliorate the negative effects of rising energy prices, greenhouse gases, escalating demand, and global energy supply uncertainty [5].

Wind energy is poised to play a prominent role in any future renewable energy portfolio due to the relatively advanced state and scalability of wind-turbine technology. The U.S. Department of Energy (DOE) estimates that 20% of the total domestic energy supply can be reliability generated by land-based and offshore wind turbines [1], and some international organizations estimate that wind energy can contribute as much as 12% of the targeted reductions in worldwide greenhouse gas emissions by 2050 (cf. [2, 3]). Yet the overall cost of producing wind energy, as measured in dollars per kilowatt-hour generated, is a critical issue that currently impedes the development of new, large-scale wind energy systems. New land-based wind farm installations are very costly, requiring large capital investments, supporting infrastructure, and connection services; and these costs are substantially higher for offshore systems. Inefficiencies in manufacturing and in the wind turbine supply chain also contribute substantially to the overall cost of operating and maintaining wind turbine farms. Furthermore, it is estimated that operating and maintenance (O&M) costs stemming from scheduled and unscheduled turbine maintenance activities may contribute at least 20%, and as much as 47.5%, of the cost of producing wind energy [112]. Given the urgent need

for clean, renewable energy sources, and the widespread appeal of wind energy as a viable alternative, the search for more effective techniques to reduce the total cost of energy (COE) for wind is of critical importance. While all efforts to reduce the COE are necessary, establishing the most effective strategies for O&M activities is likely to confer the most immediate benefit. Wind farm maintenance planners usually have real-time, condition-based data at their disposal when making decisions regarding inspections, repairs, or replacements. However, they currently lack the necessary techniques to effectively exploit these valuable data, since the information emitted from a single component (or multiple components) does not include optimal maintenance decisions. The objective of this dissertation is to fill this critical gap and reduce wind turbine O&M costs by developing a comprehensive wind turbine maintenance optimization framework that converts observed condition- or environment-based data into cost-optimal replacement actions for individual wind turbine components.

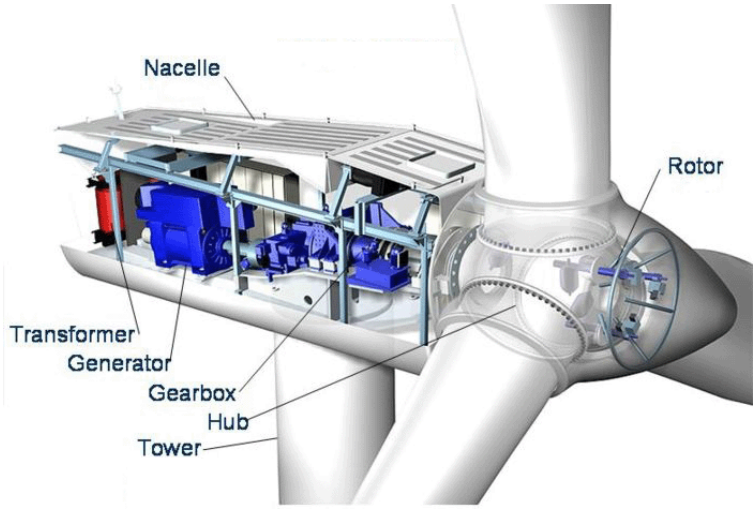

Figure 1: Components and subsystems of a horizontal axis wind turbine.

Any maintenance framework for wind turbine systems must consider the unique array of structural, mechanical, and electrical components that compose them, as well as the reliability issues that can stem from their use in exposed and dynamic environments. Wind turbines are typically categorized as either horizontal axis (the rotor shaft is perpendicular to the tower) or vertical axis (the rotor shaft is parallel to the tower). Irrespective of type, wind turbines all consist of a structural tower that supports a rotor blade and a *nacelle* that houses most of the mechanical and electrical systems, including the generator (see Figure 1). The wind turbine tower and its blades are designed to be structurally sound so as to withstand strong wind gusts and other harsh

weather conditions. Towers can range in height from 50 feet to well over 400 feet (as measured from ground level to the hub level). The rotational energy of the turbine blades is transferred via a low-speed shaft into the nacelle, and the rotor blades are connected to the shaft via a rotor hub which contains mechanisms to adjust blade pitch. A yaw system is mounted between the nacelle and tower to orient the rotor blades into the prevailing wind. Traditionally, wind turbines use a gearbox that transfers the rotation of the low-speed shaft at roughly 25 revolutions per minute (rpm) to a high-speed shaft that rotates at roughly $1,500$ rpm and is used to power the generator [113]. Some wind turbines do not use a gearbox but rather drive the generator directly with the low-speed shaft; however, such turbines are not necessarily more reliable than those with gearboxes [94, 111].

Wind turbines are prone to a variety of failure types. Recent field studies in Europe (cf. [94, 111]) suggest that failure frequencies can be ordered (from high to low) as follows: electrical systems, rotor blades and hub, converter, generator, hydraulics, and the gearbox. Corrosion, vibration fatigue, and mechanical overload (or shocks) are the main contributors to component and subassembly failures. Temporal and seasonal effects of varying wind patterns also have a significant impact on wind turbine reliability and maintainability. In fact, there is now a substantial body of literature that suggests wind speed, wind turbulence, and location are good indicators of wind turbine failure rates. Some of the studies corroborating this assertion include [15, 16, 17, 18, 33, 35, 36, 74, 78, 85, 98, 97, 101]. Numerous studies on the reliability of both land-based and offshore wind turbines in Europe have recently appeared, and a sampling of these includes [26, 35, 36, 43, 93, 99, 105]. Particularly informative are the studies by Faulstich et al. [35, 36] which show that minor failures (those requiring about four hours of downtime) comprise roughly 75% of all failures but account for only 5% of the turbine's downtime, whereas major failures (those requiring 6 to 12.5 days of downtime) are far less frequent but comprise 95% of downtime. For example, failures of the electrical system occur most frequently but are typically repaired in less than a single work day. Gearbox failures are less frequent, but require days (or even weeks) to repair or replace.

All wind turbine maintenance activities (e.g., inspections, repairs, or replacements) are very costly and complicated because wind turbines are typically installed in difficult to access, remote locations. These complications are magnified when repairs or replacements of larger critical components or subsystems are needed, or when the wind farm is located offshore. Although their failure rates are comparatively small, turbine blade, generator, and gearbox failures constitute

75% of O&M costs [84, 107] due to (1) component procurement costs; (2) transportation costs; (3) specialized maintenance equipment, including a large crane to remove failed components and lift new components, personnel, and other resources into the nacelle; and (4) specialized (possibly third-party) work crews. Obviously, wind turbine availability is impacted significantly by the availability of components and cranes, safety concerns for workers (due to harsh weather conditions), and minimum staffing levels (at least two crew members are required for most tasks). It goes without saying that maintenance costs for offshore turbines are nearly double those of their land-based counterparts due to accessibility issues (e.g., distance from shore, weather conditions, wave height) and availability of transportation and lifting equipment (e.g., sea vessels, cranes, and personnel) [35, 103]. The risks of unanticipated wind turbine component failures are often mitigated by scheduled maintenance (SM). The SM regimen generally practiced by wind farm operators involves biannual inspections of major components, assessing their functionality, and preforming minor maintenance activities. Tasks include visually inspecting turbine blades, gearboxes and shafts, replacing the gearbox oil and hydraulic system fluid, and testing major systems, such as pitch control [48]. Despite the fact that some studies estimate that 20–70% of turbine parts by cost must be replaced during the first 10 years of a wind turbine's operation [48], operators are not mandated to replace major components such as blades, generators, and gearboxes prior to failure. In other words, SM for wind turbines effectively assumes that the components will operate without failure for the 20–30 year design life of the turbine. Consequently, the current SM practice is clearly not only based on unrealistic assumptions about the useful lifetimes of individual components but also provides no allowance for preventively replacing major components.

Condition-based maintenance (CBM) techniques are currently the focus of many efforts to improve wind turbine reliability and availability. CBM techniques can be off-line (those requiring a turbine shutdown) or online (those not requiring a turbine shutdown). Basic online condition monitoring can track important indicators (e.g., temperatures, speeds, fluid levels, line-phase imbalance, voltages, and tower vibration). However, sophisticated sensing systems can monitor vibrations in bearing housings, detect discontinuities in gears or rolling surfaces, and monitor fluids for contamination [28, 67, 100, 107, 113]. As noted by Walford [107], although CBM's benefits are manifold, it is rather costly (as compared to SM).

Nilsson and Bertling [75] analyzed a CBM policy via a life-cycle cost analysis in two case studies using real turbine data in Sweden and the U.K. Using both land-based and offshore scenarios, they showed that, despite the elevated costs of CBM, implementing CBM strategies will reduce total

turbine life cycle costs (which include capital investment costs, O&M costs, and lost revenue due to downtime). So there is ample evidence to suggest that scientifically grounded CBM strategies will reduce O&M costs.

Many of the condition monitoring techniques proposed for wind turbines are experimental and not used in practice by the wind industry. Those condition monitoring systems that are currently utilized focus primarily on the timely detection of component faults. An example of such a system is the Bently Nevada ADAPT Wind Condition Monitoring Solution developed by General Electric (GE) [4]. This system monitors vibration frequencies to detect and trend particle contamination in the planetary gearbox. By monitoring the signals for an extended time period, the system establishes a baseline of vibration frequencies for healthy gearbox operation. As the gearbox degrades and the frequencies deviate from the baseline, the system indicates a gearbox fault to the wind farm operator so that repair or replacement actions can be taken. Although the condition monitoring systems used by GE and other manufacturers detect early-stage component faults and provide wind farm operators with a forewarning of component failures, they do not provide information about the component's residual lifetime distribution, nor do they provide a comprehensive optimal policy for replacing components.

1.2 RELEVANT LITERATURE

Attempts to characterize and model the reliability of wind energy systems have burgeoned during the last decade. Most relevant to this dissertation are simulation and analytical models of CBM for wind energy systems. As noted by Byon et al. [18], modeling wind turbine reliability is difficult due to cascading failures and the complexity of multiple failure modes within a single turbine (or across turbines in a single wind farm); therefore, many researchers have proposed simulation models of reliability and/or maintenance activities. Rademakers et al. [84] developed a Monte Carlo simulation model (with an analytical stochastic model) describing the long-run average O&M costs for an offshore wind farm and concluded that 55% of total costs stem from lost revenues due to downtime. McMillan and Ault [72] developed a discrete-time Markov chain (DTMC)-based Monte Carlo simulation to analyze the cost efficiency of condition monitoring systems with perfect sensor information. They compared a periodic SM policy to a CBM policy and found that CBM dominates with respect to maximizing availability and minimizing costs. Using a sequential Monte Carlo

simulation model, Kim and Singh [59] considered the impact of aging components (i.e., those with increasing hazard rate functions) and concluded that aging effects are significant when estimating performance losses (e.g., expected energy not supplied). Byon et al. [18] developed a discrete-event simulation model that considers wind turbine state changes as a result of stochastic events. They developed a power generation model, a wind speed model, a DTMC degradation model, and a maintenance model. A hidden Markov model (HMM) was used to account for a partially observed degradation state of the component. They found that CBM dominates periodic SM with respect to performance and cost metrics. In fact, using a CBM strategy, the maintenance costs per turbine were reduced by 24%.

Of particular importance to this dissertation are those studies using stochastic (namely Markov) models of maintenance strategies for wind turbine systems. Among these, most assume a Markov degradation model (i.e., one in which the component's level of degradation transitions between a finite set of discrete states). McMillan and Ault [72] assume that generator, gearbox, and turbine blades can assume one of three states (functioning fully, degraded, or failed) while the electrical system's status is binary (functioning or failed). The transition probability matrix was parameterized using historical data and used as the basis of a Monte Carlo simulation model. Arabian et al. [7] used a two-state continuous-time Markov chain (CTMC) model to compare the reliability of geared generators and direct-drive wind turbines. Dobakhshari and Fotuhi-Firuzabad [31] developed a CTMC model for power output from a single turbine and from an entire wind farm. They included time-varying patterns of wind speed and seasonal effects to show that these have a significant impact on reliability indices. Leite et al. [64] developed a probabilistic model of wind farm power generation that includes both a wind model and a turbine model. The status of the wind turbine is modeled by a two-state CTMC, and the mean time to repair the turbine is assumed to be a function of the weather conditions, the type of the failed component, the maintenance strategy, and wind speed at the time of failure. A finite-state Markov process was used to characterize the wind speed. Their results suggest that the fidelity of the wind speed model contributes significantly to the prediction accuracy when compared with real data obtained at three Brazilian wind energy centers.

An especially useful stochastic maintenance model for capturing complex features in wind turbine systems and prescribing optimal maintenance strategies is a partially observable Markov decision process (POMDP) model. In these models, a decision maker does not have the benefit of perfect state information (e.g., the exact degradation status of a wind turbine component) and

must make maintenance decisions based on partial information. The information might be incomplete due to imperfect models linking condition monitoring signals to specific faults and/or noise in the sensor signal itself (cf. [72]). POMDP models have found wide applicability in multi-state maintenance optimization problems (cf. [49, 69, 70, 114] and references therein). Of particular relevance to this dissertation is a model by Byon et al. [16] of wind turbine maintenance subject to stochastic weather conditions. The authors developed a POMDP model that includes such factors as adverse weather conditions, repair interruptions and delays, and the lead time required to procure replacement parts and assemble ground crews. The objective is to minimize the expected per period cost, where costs are tied to both maintenance activities and downtime. The wind turbine's degradation level is assumed to evolve as a finite-state DTMC, and in each period, three actions are available: no action, preventive maintenance (PM), or exactly observing the degradation level. If adverse weather conditions occur during either preventive or corrective maintenance, the activity must be suspended until the weather is favorable, resulting in downtime costs. The authors obtained closed-form expressions for the optimal policy, including control limits for PM.

Byon et al. [15] extend the model in [16] to include multiple components and season-dependent (or temporally-nonhomogeneous) weather conditions. Their model is a discounted POMDP whose objective is to minimize the total expected discounted cost, and the degradation process is a finite-state DTMC with multiple failure states. A PM action can restore the wind turbine component to any improved degradation level with cost dictated by the type of PM. Corrective maintenance actions have costs, lead times, repair times, and weather requirements that are dependent on the type of failure that occurs. As in [16], maintenance activities are suspended during adverse weather conditions; however, the revenue loss in each period is dynamic. Using a backward dynamic programming algorithm, they solve the POMDP model and numerically illustrate optimal policies for a wind turbine gearbox application.

Maintenance models using a POMDP framework need not only consider discrete degradation processes. Though not focused on wind turbine maintenance applications, one study considers a partially-degraded machine whose degradation status assumes values on a continuous state space. Zhou et al. [117] developed a partially observable semi-Markov process (POSMDP) model which is continuous in both time and state. Suppose $\Lambda(t)$ is the *true* health status of the machine at time t, and let $X(t)$ be the *observed* status which is assumed to be noise-corrupted. The model assumes that $\{\Lambda(t) : t \geq 0\}$ is a Gamma-based state space model, i.e., $\Lambda(t + \Delta t) - \Lambda(t)$ has a Gamma$(a\Delta t, \xi)$ distribution where $a\Delta t$ and ξ are the shape and scale parameters, respectively. The

model also assumes that $X(t)$ is normally distributed with mean $\Lambda(t)$ and standard deviation σ_ξ. Three actions are permissible: inspections, replacements (preventive or corrective), and imperfect maintenance. Because the state space of $\Lambda(t)$ is continuous, the *belief space*, or set of all probability distributions over the uncertain states, is infinite dimensional. To reduce the dimensionality of the belief space, Monte Carlo-based density projection (see Zhou et al. [116]) is used to project the set of possible health distributions onto a (smaller) space of parametric distributions. Using three different maintenance strategies, they consider two objectives: minimizing the long-run cost rate and maximizing the long-run availability of the machine. The models are solved numerically using a policy iteration algorithm.

The simulation and analytical models described in this section are extremely useful for elucidating the benefits of a sound CBM strategy. However, nearly all the CBM approaches are *ad hoc*; that is, they assume a stochastic model for the degradation process that cannot incorporate new information acquired about the operating environment. Although some models, such as those by Byon et al. [16, 15], incorporate the impact of weather conditions on maintenance activities, they do not allow for weather conditions to affect degradation. In addition, nearly all of the CBM models assume that degradation evolves as a discrete-state process when in reality degradation is often continuous. The approach taken in this dissertation is distinguished from all existing approaches in that it will explicitly account for the impact of the specific operating environment on a wind turbine component's degradation and prescribe optimal replacement strategies for the component in that operating environment.

1.3 PROBLEM STATEMENT AND RESEARCH OBJECTIVES

This dissertation addresses three important problems related to reducing the operating and maintenance costs of a single wind turbine. The primary objectives of this research include the development of analytical and statistical tools to assess the current and future health of critical components, and to prescribe optimal policies for replacing them. The specific objectives are as follows:

1. To develop a hybrid analytical-statistical framework for transforming observed condition- or environment-based data into cost-optimal replacement actions for a single wind turbine component;

2. To develop statistical procedures, based on Markov chain Monte Carlo methods, for estimating the parameters of the operating environment and the environment's relationship to degradation by observing component degradation or a surrogate signal;
3. To create and solve stochastic optimization models that prescribe optimal replacement actions based on partially observable environment conditions and/or imperfectly observable degradation observations.

1.4 DISSERTATION OUTLINE AND CONTRIBUTIONS

Chapter 2 presents a general framework to approximate the stochastic, environment-driven degradation process of a wind turbine component by observing a signal of the component's degradation. The degradation signal is assumed to evolve as a switching diffusion process, and a Markov chain Monte Carlo (MCMC) statistical procedure is adapted to estimate the process parameters. The results of numerical experiments are presented to evaluate the framework's ability to characterize both simulated and real degradation processes. Performance is evaluated by comparing component lifetime estimates obtained using the estimated parameters to the actual component lifetime, and by comparing a mean signal computed via a forward-filtering-backward-smoothing algorithm to the observed signal. The results indicate that the framework can be used to estimate a diverse set of degradation processes and is effective even in cases where the signal deviates significantly from a switching diffusion model. The major research contributions of Chapter 2 are the following:

1. A general framework, based on a switching diffusion process model, to approximate the stochastic, environment-driven degradation process of a wind turbine component by observing a signal of component degradation;
2. The development of performance measures to assess the quality of the approximation based on component lifetime and mean signal estimates.

Chapter 3 introduces two models for optimally replacing a single wind turbine component using a partially-observed Markov decision process (POMDP) model with the objective of minimizing the long-run average replacement cost per unit time. Both models assume that the component's degradation is observed perfectly at each decision epoch but that the environment is only partially observable. The two models differ in that the first assumes replacements occur instantaneously

with fixed, deterministic costs, whereas the second model assumes replacements require a fixed, deterministic time period for completion during which downtime costs are accrued as a function of the prevailing environment conditions. For each model, we prove the existence and optimality of a threshold-type replacement policy with respect to the cumulative degradation level whose thresholds depend on some assessment of the environment's state. Numerical solution techniques are introduced to compute optimal policies using the estimated environment parameters, and numerical examples are presented to illustrate the optimal policies. The main contributions of Chapter 3 are as follows:

1. The formulation of realistic, POMDP replacement models for a wind turbine component that assumes degradation is a function of a randomly-evolving, partially-observed environment;
2. The development of an optimal replacement model that explicitly incorporates the environment-dependent downtime costs accrued during wind turbine component replacements;
3. Structural results that characterize the optimal replacement policy in each model (including the existence of optimal replacement thresholds with respect to the component's cumulative degradation level).

Chapter 4 considers an optimal replacement model for a single wind turbine component whose cumulative degradation is imperfectly observed according to a parameterized probability density function (p.d.f.). Similar to the models in Chapter 3, the replacement problem is formulated using a POMDP model with the objective of minimizing the long-run average replacement cost per unit time. Optimal replacement thresholds, with respect to the observed degradation level, are shown to exist under special conditions. The problem is solved numerically using a projection-filtering algorithm that employs belief projection and particle filtering techniques. A modification of the algorithm is developed for the case when the environment state is known with certainty. The main contributions of Chapter 4 are as follows:

1. The formulation of a realistic, POMDP replacement model for a wind turbine component whose degradation is a function of a randomly-evolving, partially-observed environment and is imperfectly observed;
2. The derivation of conditions under which optimal replacement thresholds exist with respect to the degradation observation;
3. An extension of the particle-filtering algorithm for the case of a mixed state-space that is induced by both partially-observable and observable random environments.

In the next chapter, we begin by addressing the problem of modeling the environment-driven degradation of a wind turbine component. It is assumed that the environment evolves as a continuous-time Markov chain (CTMC) on a finite state space, and an observed, surrogate signal of component degradation is modeled as a switching diffusion process. The model parameters are estimated from the signal by adapting a MCMC-based statistical procedure. Though the assumptions may seem restrictive, the flexibility of this approach will become apparent in the numerical illustrations.

2.0 DEGRADATION IN A RANDOM ENVIRONMENT

In this chapter, a mathematical model for the environment-driven degradation of a wind turbine component is presented, as well as a model to characterize a signal of the component's degradation. A Markov chain Monte Carlo (MCMC) inference procedure is introduced to estimate the environment parameters from observations of the degradation signal. The effectiveness of the degradation model and inference procedure for estimating the lifetimes of components is illustrated by way of several numerical examples.

2.1 MODEL OF DEGRADATION IN A RANDOM ENVIRONMENT

For this discussion, all random variables are defined on a common, complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\mathcal{F}_t = \sigma((X(t), Z(t)) : 0 \leq u \leq t)$ be the filtration at time t, where $\mathcal{F}_t \subseteq \mathcal{F}$ and $\mathcal{F}_t \subseteq \mathcal{F}_u$ if $t < u$. For any $t \geq 0$, $\mathcal{F}_t$ can be thought to contain all of the available information up to time t. For an event $A \in \mathcal{F}$, denote by $\mathbb{I}(A)$ the indicator function, where $\mathbb{I}(A) = 1$ if A occurs and $\mathbb{I}(A) = 0$ otherwise, and for $a, b \in \mathbb{R}$ let $a \wedge b \equiv \min(a, b)$ and $a \vee b \equiv \max(a, b)$. Consider a wind turbine component that is placed into service at time zero with no degradation. Over time, the component degrades due to normal usage and the influence of its operating environment. Once the component's cumulative degradation level exceeds a (deterministic) critical threshold, x_c ($x_c > 0$), it is declared to be failed. Let $X(t)$ be an $\mathcal{F}_t$-measurable random variable that denotes the cumulative degradation at time t and $T(x_c) \equiv \inf\{t \geq 0 : X(t) \geq x_c\}$ be the random time for cumulative degradation to exceed the critical threshold. It is assumed that $X(t)$ is a function of the state of the operating environment. Let $\mathcal{S} \subseteq \mathbb{N}$ be the state space of the environment, and denote the *stochastic environment process* by $\mathcal{Z} \equiv \{Z(t) : t \geq 0\}$, where $Z(t) \in \mathcal{S}$ is the state of the environment at time t.

If $r : \mathcal{S} \to \mathbb{R}$ is a one-to-one mapping of the degradation rate associated with each environment state, then $X(t)$ is given by

$$X(t) = X(0) + \int_0^t r_{Z(u)} \mathrm{d}u, \tag{2.1}$$

where $X(0) = 0$ with probability 1 (w.p. 1). Degradation models, such as the one in equation (2.1), have been analyzed using many different approaches. One class of approaches consists of statistically-based techniques. For example, regression models are used to characterize the random effects of degradation in [68, 88] but assume a static environment. In time-varying environments, proportional hazard models (PHM) have been used to estimate the residual lifetime distributions (RLD) of systems in [73, 50, 38, 80, 10, 62, 106, 95, 65, 66, 115]. Other statistical approaches attempt to estimate the RLD by modeling degradation using Weiner processes, gamma proceses, or random coefficients models [32, 110, 40, 39, 41, 109]. These statistical approaches have been employed with great success to model various degradation processes; however, the resulting models are typically limited by restrictive assumptions (e.g. normality) and are often difficult to generalize beyond the specific system considered.

Another approach is to develop stochastic failure models based on various assumptions of $\mathcal{Z}$ and r. Foundational efforts include the results of Esary [34], who examined failure distributions for a component subject to a general degradation process and random shocks. Çinlar [22, 23] considered $\{(X(t), Z(t)) : t \geq 0\}$ as a more general Markov additive process (MAP) and also examined two models in which $\mathcal{Z}$ is a finite Markov process. In the first model, degradation corresponds to a Lévy process, and the component is subject to random shocks at environment transitions, whereas in the second model, degradation is a continuous additive functional of $\mathcal{Z}$. Results when the degradation of multiple components is influenced by a common environment are presented by Çinlar, *et al.* [24, 25] including conditions under which the component lifetimes are associated. In [52, 53] Kharoufeh considers a single component and assumes $\mathcal{Z}$ is a continuous-time Markov chain (CTMC) on a finite state space and that component degradation increases at a constant, positive rate in each environment state. Under identical degradation dynamics, Kiessler [58] computes the limiting average availability of a periodically inspected component. In [55] Kharoufeh *et al.* extend the model of [58] by including random shocks and derive the component's lifetime distribution, mean time to failure, and limiting availability. In the case when $\mathcal{Z}$ is a semi-Markov process, Kharoufeh *et al.* [57] provide an approximation of the component's lifetime using phase-type distributions and in [54] derive results for the lifetime distribution when $\mathcal{Z}$ is a nonhomogeneous CTMC or a semi-Markov process. Strictly discrete processes modulated by an environment have also been

considered. Özeckici and Soyer [76] study a Markov modulated Bernoulli process where the success probability, the probability that a periodically inspected component survives a given inspection period, is determined by a Markov process $\mathcal{Z}$. They apply Bayesian inference techniques to obtain the model parameters depending on the observability of $\mathcal{Z}$. A network reliability application of this model is presented by Özekici and Soyer in [77], as well as a software reliability application of the models in [24, 25].

In this research, assumptions are imposed on $\mathcal{Z}$ and r that result in a tractable stochastic model for the degradation of a wind turbine component. It is assumed that component degradation is influenced by (i) the actual ambient environment in which it resides and operates (e.g., wind speed, ambient air temperature or humidity can induce degradation); and/or (ii) the various operational settings of the equipment. Potential interactions between environmental conditions and operational settings are not precluded from influencing degradation; therefore, the set of governing environment states, denoted by $\mathcal{S}$, is a finite set consisting of all *unique* combinations of the environmental conditions and operational settings. Denote the total number of environment states by $\ell \equiv |\mathcal{S}|$. For example, if a unit is subject to two different temperature regimes, denoted t_- and t_+, and two different operating speeds, denoted v_- and v_+, the state space of the governing environment is the Cartesian product of $\mathcal{S}_1 = \{t_-, t_+\}$ and $\mathcal{S}_2 = \{v_-, v_+\}$; that is, $\mathcal{S} \equiv \mathcal{S}_1 \times \mathcal{S}_2 = \{(t_-, v_-), (t_-, v_+), (t_+, v_-), (t_+, v_+)\}$, and $\ell = 4$. In reality, it is often the case that the environmental conditions, or even the operational settings experienced by a unit, are unknown or unobservable. Therefore, ℓ must be inferred from real data.

The present model assumes that $\mathcal{Z}$ is a finite irreducible, temporally-homogeneous CTMC with infinitesimal generator matrix $\boldsymbol{Q} = [q_{ij}]$, $i, j \in \mathcal{S}$. Advantages of this assumption are that under various forms of r, the Laplace-Stieltjes transform (LST) of the lifetime distribution is known, and CTMCs can be constructed to approximate the stochastic behavior of many different forms of degradation processes. It is important to note that the CTMC assumption can be relaxed to analyze non-Markovian environments [57]; however, the simplest case is presented here to elucidate the main concepts. Let $\{Z_n : n \geq 0\}$ be the discrete-time Markov chain (DTMC) embedded at transition epochs of $\mathcal{Z}$, i.e., Z_n is the state of $\mathcal{Z}$ just after the nth transition, $n \geq 1$, and let $\boldsymbol{P} = [p_{ij}]$ be the transition probability matrix. Assume the existence of a measurable, one-to-one function $r : \mathcal{S} \to \mathbb{R}_+$ such that whenever $Z(t) = i \in \mathcal{S}$, the component degrades at a *unique* constant rate $r_i > 0$. Because the degradation rates are distinguishable, the states of $\mathcal{S}$ can be ordered such that $i < j$ if $r_i < r_j$, $i, j \in \mathcal{S}$.

Let the row vector $\boldsymbol{r} = [r_1, r_2, \ldots, r_\ell]$ contain these ℓ ordered rates, and set $\boldsymbol{R}_d = \text{diag}(\boldsymbol{r})$. So that $X(t)$ is well defined for each $t \geq 0$, assume that

$$\int_0^t |r_{Z(u)}| \mathrm{d}u < \infty \quad \text{w.p. 1.}$$

The strict positivity of the degradation rates r_i, $i \in \mathcal{S}$, ensures that the sample paths of $\{X(t) : t \geq 0\}$ are almost surely piecewise linear and monotone increasing; therefore, for $x_c > 0$ and $t \geq 0$, the events $\{X(t) \leq x_c\}$ and $\{T(x_c) > t\}$ are equivalent.

Let $F(x_c, t) \equiv \mathbb{P}(T(x_c) \leq t) = 1 - \mathbb{P}(X(t) > x_c)$ denote the cumulative distribution function (c.d.f.) of the component's lifetime. As proved in [55, 56], the Laplace-Stieltjes transform (LST) of $F(x_c, t)$, with respect to the spatial variable x_c, is

$$\widetilde{F}(u, t) \equiv \int_0^\infty e^{-ux_c} F(\mathrm{d}x_c, t) = 1 - \boldsymbol{\alpha} \exp\left[(\boldsymbol{Q} - u\,\boldsymbol{R}_d)t\right] \boldsymbol{e}, \quad u > 0, \tag{2.2}$$

where $\boldsymbol{\alpha} = [\mathbb{P}(Z(0) = i)]$, $i \in \mathcal{S}$, is the environment's initial distribution, $\exp(\boldsymbol{A})$ denotes matrix exponentiation of the square matrix $\boldsymbol{A}$, and $\boldsymbol{e}$ is an $\ell \times 1$ vector of ones. Furthermore, if $m^n(x_c) \equiv \mathbb{E}(T^n(x_c))$ denotes the nth moment of the lifetime distribution, $n \geq 1$, then $\widetilde{m}^n(u)$, the LST of $m^n(x_c)$ with respect to the spatial variable x_c, is given by

$$\widetilde{m}^n(u) \equiv \int_0^\infty e^{-ux_c} \mathrm{d}m^n(x_c) = n!\, \boldsymbol{\alpha}(u\boldsymbol{R}_d - \boldsymbol{Q})^{-n} \boldsymbol{e}. \tag{2.3}$$

Now, if $\boldsymbol{\alpha}$, $\boldsymbol{r}$, and $\boldsymbol{Q}$ are known, or if they can be estimated from observed data, the lifetime distribution and its moments can be obtained by inverting the Laplace transforms of (2.2) and (2.3), respectively. That is,

$$F(x, t) = \mathcal{L}^{-1}\{u^{-1}\widetilde{F}(u, t)\}, \tag{2.4}$$

$$m^n(x) = \mathcal{L}^{-1}\{u^{-1}\widetilde{m}^n(u)\}, \tag{2.5}$$

where $\mathcal{L}^{-1}$ denotes the inverse Laplace transform operator. Often the inversion operations (2.4) and (2.5) are not tractable analytically, in which case the inversion can be done numerically using stable techniques (see for example [6]). In the asymptotic regime (as $x_c \to \infty$), the expected lifetime can be approximated using a relatively simple expression due to Kharoufeh *et al.* [56] as follows:

Theorem 2.1. *As* $x_c \to \infty$,

$$\frac{\mathbb{E}[T(x_c)]}{x_c} \to \frac{1}{\pi_s \boldsymbol{r}'}, \tag{2.6}$$

where π_s *is the stationary distribution of* $\boldsymbol{Q}$.

Ultimately, for the degradation model to be useful, it must be possible to estimate its parameters by observing the degradation process. Kharoufeh and Cox [53] proposed a clustering technique for estimating $(\boldsymbol{Q}, \boldsymbol{r}, \ell)$ from degradation observations. However, this approach tends to overestimate the number of environment states, leading to over-specified models. Instead, an alternative Markov chain Monte Carlo (MCMC) inference procedure is proposed to estimate $(\boldsymbol{Q}, \boldsymbol{r}, \ell)$.

2.2 REVIEW OF MARKOV CHAIN MONTE CARLO (MCMC) ESTIMATION

Markov chain Monte Carlo (MCMC) refers to a class of techniques that approximate sampling from an arbitrary distribution function G. The fundamental concept underlying these techniques is that random samples of G are realizations of a Markov chain $\mathcal{X} = \{X_n \in S_X : n \geq 0\}$ that converges in distribution to G. In this section only, it is assumed that $S_X \subseteq \mathbb{R}$ and that G is absolutely continuous, possessing probability density function (p.d.f.) g.

An early MCMC-based technique is the *Metropolis-Hastings* (MH) algorithm [47], which obtains a sequence of random sample $X_1, X_2, \ldots$ for g by randomly accepting or rejecting samples from a candidate density $q(y|x)$, where x is the most-recently accepted candidate. The algorithm is as follows:

1. Given X_n, randomly sample $Y \sim q(y|X_n)$.
2. Randomly sample $U \sim \mathrm{U}(0,1)$ and set

$$X_{n+1} = \begin{cases} Y, & \text{if } U \leq \alpha(X_n, Y), \\ X_n, & \text{otherwise,} \end{cases}$$

   where

$$\alpha(x,y) = 1 \wedge \frac{g(y)q(x|y)}{g(x)q(y|x)}$$

   is the probability of accepting the proposed sample.

Special cases of the MH algorithm include the *independence sampler*, where $q(y|x)$ is independent of x, the *random walk sampler*, where $q(x|y) = q(y|x)$, and the *Gibbs sampler* [21]. The *Gibbs sampler* generates a sequence of m-dimensional random vectors $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots$, where $\boldsymbol{X}_n = [X_1^{(n)}, X_2^{(n)}, \ldots, X_m^{(n)}] \in \mathbb{R}^m$, using a sequence of conditional distributions from the distribution G with support $K \subseteq \mathcal{B}(\mathbb{R}^m)$, where $\mathcal{B}(\mathbb{R}^m)$ denotes the Borel sets of $\mathbb{R}^m$.

Let $g(z_i|z_1,\ldots,z_{i-1},z_{i+1},\ldots,z_m)$ be the conditional density of the ith component of $\boldsymbol{X}$, given components $(z_1,\ldots,z_{i-1},z_{i+1},\ldots,z_m)$. The Gibbs sampler generates random samples from G as follows:

1. Given X_n,

   Randomly sample $X_1 \sim g(x_1|X_2,\ldots,X_m)$

   Randomly sample $X_2 \sim g(x_2|Y_1,X_2,\ldots,X_m)$

   $\vdots$

   Randomly sample $X_m \sim g(x_m|X_1,\ldots,X_{m-1})$

2. Set $X_{n+1} = [X_1, X_2, \ldots, X_m]$.

There are many other specialized MCMC sampling techniques such as *Metropolis-Gibbs hybrids*, the *multiple-try Metropolis-Hastings method*, the *hit-and-run sampler*, and *auxiliary variable samplers* [61].

Another application of MCMC-based techniques is for Bayesian parameter inference, where the objective is to estimate the parameters of a stochastic process by assuming each parameter has a prior distribution. The parameters are sampled in an iterative fashion, and each sample of the entire parameter set corresponds to a realization of an underlying Markov chain that converges to a stationary distribution [42]. In principle, these procedures resemble a Gibbs sampler, where each parameter is sampled iteratively from its prior distribution, conditioned on the estimates of the other parameters. However, generating each parameter may require one or more additional MCMC procedures of varying complexity. Often MCMC-based inference is an attractive alternative to maximum-likelihood or psuedo-likelihood approaches that may not be analytically tractable or are excessively cumbersome computationally. Applications of MCMC-based interference include those for hidden-Markov models [86, 9], diffusion processes [82, 63, 87, 13, 30, 45, 44], and finite mixtures [37].

Of particular interest to this research are the MCMC inference procedures developed for a special class of diffusion models known as *switching diffusion processes* [63, 13, 45, 44]. A switching diffusion process satisfies the general stochastic differential equation (SDE)

$$\mathrm{d}\boldsymbol{Y}(t) = \mu(\boldsymbol{Y}(t), Z(t), t) + \boldsymbol{\sigma}(\boldsymbol{Y}(t), Z(t), t)\,\mathrm{d}\boldsymbol{B}(t), \tag{2.7}$$

where $\boldsymbol{Y}(t) = [Y_1(t), Y_2(t), \ldots, Y_n(t)] \in \mathbb{R}^n$, $\mu(\boldsymbol{Y}(t), Z(t), t)$ is a n-dimensional drift function, $\boldsymbol{\sigma}(\boldsymbol{Y}(t), Z(t), t)$ is a $n \times n$ diffusion matrix with strictly positive elements, $\boldsymbol{B}(t)$ is standard n-dimensional Brownian motion, and $\{Z(t) : t \geq 0\}$ is CTMC on finite S_X. Although (2.7) can admit

many different classes of functions for $\mu(\boldsymbol{Y}(t), Z(t), t)$ and $\boldsymbol{\sigma}(\boldsymbol{Y}(t), Z(t), t)$, the MCMC-based inference procedures presented here restrict these functions to be constants for each state in $\mathcal{S}$; that is,

$$\begin{aligned} \mu(\boldsymbol{Y}(t), Z(t), t) &\equiv [\mu_i(Z(t))], \quad i = 1, 2, \ldots, n, \\ \boldsymbol{\sigma}(\boldsymbol{Y}(t), Z(t), t) &\equiv [\sigma_{ij}(Z(t))], \quad i, j = 1, 2, \ldots, n. \end{aligned}$$

The objective is to estimate $\boldsymbol{Q}$, $\mu_i(k)$, and $\sigma_{ij}(k)$ by observing $\{\boldsymbol{Y}(t) : 0 \leq t \leq T\}$ for $i, j = 1, 2, \ldots, n$, $k \in \mathcal{S}$, and $0 < T < \infty$.

In the next section, a special case of (2.7) is considered where $n = 1$ and the diffusion coefficient is independent of $(\boldsymbol{Y}(t), Z(t), t)$. The relevance of this case to modeling the degradation signal is elucidated, and an MCMC-based inference procedure is described to estimate $(\boldsymbol{Q}, \boldsymbol{r}, \ell)$ from the degradation signal.

2.3 MCMC INFERENCE PROCEDURE FOR THE DEGRADATION SIGNAL

In this section, a special case of the switching diffusion model is considered to characterize an observed degradation signal from the model in Section 2.1. Then a MCMC-based inference procedure is presented, originally developed by Leichty and Roberts [63], to estimate the environment parameters by observing the signal. Consider a special case of the SDE (2.7) when $n = 1$ and the diffusion coefficient is independent of $(\boldsymbol{Y}(t), Z(t), t)$. Stating (2.7) in the notation of Section 2.1, where $\mu(Z(t)) \equiv r_{Z(t)}$, gives

$$\mathrm{d}Y(t) = r_{Z(t)}\mathrm{d}t + \sigma \mathrm{d}B(t), \quad \sigma > 0. \tag{2.8}$$

To illuminate the relevance of (2.8) to the degradation model, consider the fact that a wind turbine's degradation, and even its environment conditions, are often not directly or precisely observed. For example, measuring gear degradation requires a costly shutdown of the wind turbine and an invasive inspection of the gearbox.

To avoid costly inspections, wind farm operators instead use sensors that measure proxies of gear degradation, such as lubricant contamination and vibrations. Even measurements of basic environmental conditions, such as wind speed by anemometers placed on a turbine nacelle, can be corrupted by disruption of local air currents caused by the turbine blades. The fact that

degradation and the environment are seldom fully, or even directly, observed in practice motivates an approach to estimate the environment parameters using a signal that serves as a proxy for the true degradation.

Let the random variable $Y(t)$ be the signal level at time t and call $\{Y(t) : t \geq 0\}$ the degradation signal process. Simply using $Y(t)$ in place of $X(t)$ in the degradation model of (2.1) is problematic since the sample paths of $\{Y(t) : t \geq 0\}$ are not monotone increasing w.p. 1. Instead, note that a solution to (2.8), with initial condition $Y(0) = 0$ w.p. 1, satisfies

$$Y(t) = \int_0^t r_{Z(u)} du + \sigma \int_0^t \mathrm{d}B(t) = X(t) + \sigma B(t),$$

so that sample paths of $\{Y(t) : t \geq 0\}$ correspond to a superposition of the true degradation process $\{X(t) : t \geq 0\}$ with a Brownian process having diffusion coefficient σ. Recalling that the degradation model defines failure when $X(t) \geq x_c$, the first-passage time $T(x_c)$ cannot be directly ascertained since $X(t)$ is not observable. Instead, the first passage time $T'(x_c)$, defined as

$$T'(x_c) = \inf\{t > 0 : Y(t) \geq x_c\},$$

will be observed.

In practice degradation signals are observed at discrete times, not continuously. Let $\mathcal{T} = \{t_0, t_1, t_2, \ldots, t_N\}$ denote a set of N discrete signal observation times, where $t_0 \equiv 0$ and $t_N \equiv T < T(x_c)$, and define $\mathcal{Y} = \{Y(0), Y(t_1), Y(t_2), \ldots, Y(T)\}$ as the set of signal observations at times in $\mathcal{T}$, where $Y(0) = 0$ w.p. 1. Define a piecewise-linear function $Y_c(t)$ to approximate $Y(t)$ on $[0, T]$ such that: (i) $Y_c(t_j) = Y(t_j)$, $j = 0, 1, \ldots, N$, and (ii) $\mathrm{d}Y_c(t)/\mathrm{d}t$ is defined for all t except $t_j \in \mathcal{T}$, $j = 1, 2, \ldots, N$, and is constant within each interval $(t_{j-1}, t_j]$, $j = 1, 2, \ldots, N$. The objective is to estimate the environment parameters $(\boldsymbol{Q}, \boldsymbol{r}, \ell)$ from $Y_c(t)$ and obtain an estimate of the component's expected lifetime. Achieving this objective requires first obtaining parameter estimates using a MCMC inference procedure and then using the approximations of (2.3) and (2.5) to obtain an estimate of the expected lifetime.

Now a detailed description of the MCMC-based procedure originally developed by Leichty and Roberts [63] is provided. Initially, it is assumed that $\mathcal{Z}$ has ℓ states. Let $\widehat{Z}(t)$ be the estimated state of $Z(t)$ at time t, $t \in [0, T]$, and $\widehat{z} = \{\widehat{Z}(t) : t \in [0, T]\}$ be the estimated sample path of $\mathcal{Z}$. The MCMC procedure consists of sequentially sampling from and updating the conditional densities of $\boldsymbol{r}|\boldsymbol{Q}, \widehat{z}, \mathcal{Y}$, $\boldsymbol{Q}|\widehat{z}, \boldsymbol{r}, \mathcal{Y}$, and $\widehat{z}|\boldsymbol{r}, \boldsymbol{Q}, \mathcal{Y}$ denoted by $f_{\boldsymbol{r}}(\boldsymbol{r}|\boldsymbol{Q}, \widehat{z}, \mathcal{Y})$, $f_{\boldsymbol{Q}}(\boldsymbol{Q}|\widehat{z}, \boldsymbol{r}, \mathcal{Y})$, and $f_{\widehat{z}}(\widehat{z}|\boldsymbol{r}, \boldsymbol{Q}, \mathcal{Y})$, respectively. A single iteration of the algorithm consists of the following:

1. Randomly sample $\widehat{\boldsymbol{r}}_{v+1} \sim f_{\boldsymbol{r}}(\boldsymbol{r}|\widehat{\boldsymbol{Q}}_v, \widehat{z}_v, \mathcal{Y})$,
2. Randomly sample $\widehat{\boldsymbol{Q}}_{v+1} \sim f_{\boldsymbol{Q}}(\boldsymbol{Q}|\widehat{z}_v, \widehat{\boldsymbol{r}}_{v+1}, \mathcal{Y})$,
3. Randomly sample $\widehat{z}_{v+1} \sim f_{\widehat{z}}(\widehat{z}|\widehat{\boldsymbol{r}}_{v+1}, \widehat{\boldsymbol{Q}}_{v+1}, \mathcal{Y})$,

where $\widehat{\boldsymbol{r}}_v$, $\widehat{\boldsymbol{Q}}_v$, and $\widehat{z}_v$ denote the values of $\widehat{\boldsymbol{Q}}$, $\widehat{\boldsymbol{r}}$, and $\widehat{z}$ respectively, at the vth iteration, $v = 0, 1, 2, \ldots$. The estimate of the diffusion coefficient σ^2, denoted by $\widehat{\sigma}^2$, is computed prior to initializing the procedure using the following estimator for the quadratic variation of a diffusion process:

$$\widehat{\sigma}^2 = \frac{1}{N} \sum_{j=1}^{N} \frac{(Y_c(t_j) - Y_c(t_{j-1}))^2}{t_j - t_{j-1}}.$$

Prior densities and hyperparameters for $\boldsymbol{r}_0$, $\boldsymbol{Q}_0$, and $\widehat{z}_0$, the initial values of $\boldsymbol{r}$, $\boldsymbol{Q}$, and $\widehat{z}$, respectively, must be specified. Assume the prior density for $\boldsymbol{r}_0$ is a constrained, multivariate normal density and obtain $\widehat{\boldsymbol{r}}_0 = [\hat{r}_1^{(0)}, \hat{r}_2^{(0)}, \ldots, \hat{r}_\ell^{(0)}]$ by randomly sampling

$$\widehat{\boldsymbol{r}}_0 \sim N(\boldsymbol{0}, \boldsymbol{\delta}) \quad \text{s.t.} \quad \hat{r}_1^{(0)} < \hat{r}_2^{(0)} < \cdots < \hat{r}_\ell^{(0)},$$

where $\boldsymbol{0}$ is an $\ell \times 1$ matrix of zeros and hyperparameter $\boldsymbol{\delta} \equiv \text{diag}(\delta_1, \delta_2, \ldots, \delta_\ell)$. Typically $\delta_i = 3\widehat{\sigma}^2$, $i \in \mathcal{S}$, so that the prior distribution of $\widehat{\boldsymbol{r}}_0$ is relatively diffuse. The prior distribution of $\boldsymbol{Q}_0$ assumes that the off-diagonal elements are independent and exponentially distributed with a rate determined by a hyperparameter β. To initialize $\widehat{\boldsymbol{Q}}_0 = [\widehat{q}_{ij}^{(0)}]$, sample $\widehat{q}_{ij}^{(0)} \sim \text{Exp}(\beta)$, $j \neq i$, and compute $q_{ii}^{(0)}$ by normalizing the ith row of $\widehat{\boldsymbol{Q}}_0$, $i, j \in \mathcal{S}$, where $\beta = T/3$. The estimated sample path $\widehat{z}$ is initialized by generating $\widehat{z}_0$ via simulation using $\widehat{\boldsymbol{Q}}_0$, where the initial state $\widehat{Z}(0)$ is sampled from a uniform-discrete distribution with support $\mathcal{S}$.

Now the conditional densities $f_{\boldsymbol{r}}(\boldsymbol{r}|\boldsymbol{Q}, \widehat{z}, \mathcal{Y})$, $f_{\boldsymbol{Q}}(\boldsymbol{Q}|\widehat{z}, \boldsymbol{r}, \mathcal{Y})$, and $f_{\widehat{z}}(\widehat{z}|\boldsymbol{r}, \boldsymbol{Q}, \mathcal{Y})$ are described. Define for each $i \in \mathcal{S}$ the constants

$$\begin{aligned} a_i &= \frac{1}{\widehat{\sigma}^2} \int_0^T \mathbb{I}(Z(u) = i) \mathrm{d}Y_c(u), \\ b_i &= \frac{1}{\widehat{\sigma}^2} \int_0^T \mathbb{I}(Z(u) = i) \mathrm{d}u + \frac{1}{\delta_i^2}. \end{aligned}$$

Intuitively, a_i is a proxy for the cumulative change of $Y(t)$ while the environment is in state i, whereas, b_i is a proxy for the cumulative amount of time the environment is in state i during [0,T]. Define an $\ell \times 1$ vector $\boldsymbol{\mu} = [a_1 b_1^{-1}, a_2 b_2^{-1}, \ldots, a_\ell b_\ell^{-1}]$ and an $\ell \times \ell$ matrix $\boldsymbol{\Sigma} = \text{diag}(b_1^{-1}, b_2^{-1}, \ldots, b_\ell^{-1})$. The conditional density $f_{\boldsymbol{r}}(\boldsymbol{r}|\boldsymbol{Q}, \widehat{z}, \mathcal{Y})$ is a constrained bivariate normal, where

$$f_{\boldsymbol{r}}(\boldsymbol{r}|\boldsymbol{Q}, \widehat{z}, \mathcal{Y}) \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad \text{s.t.} \quad r_1 < r_2 < \cdots < r_\ell. \tag{2.9}$$

Note that the elements of $\boldsymbol{\mu}$, a_i/b_i, are statistical estimates of the degradation rate in each state i. For the conditional density of $\boldsymbol{Q}$, it is assumed that the off-diagonal elements are independent, Gamma-distributed random variables. Denoting ϑ_{ij}, $i, j \in \mathcal{S}$, as the number of transitions that occur from state i to j in $\widehat{z}$, define $\alpha_{ij} = \vartheta_{ij} + 1$ and

$$\gamma_i^{-1} \equiv \int_0^T \mathbb{I}(Z(u) = i)\mathrm{d}u + \beta, \quad i, j \in \mathcal{S}.$$

Let

$$f_{\boldsymbol{Q}}(\boldsymbol{Q}|\widehat{z}, \boldsymbol{r}, \mathcal{Y}) \sim \prod_{j \neq i} \text{Gamma}(\alpha_{ij}, \gamma_i^{-1}) \tag{2.10}$$

$$\text{subject to} \quad q_{ii} = -\sum_{j \neq i} q_{ij} \quad \text{w.p. } 1, \quad i, j \in \mathcal{S}.$$

Intuitively, α_{ij} is a proxy for the total number of transitions from state i to j, and γ_i^{-1} is a proxy for the total cumulative time spent in state i. For $j \neq i$, the mean of q_{ij} is $\alpha_{ij}\gamma_i^{-1}$, which is approximately equal to the maximum likelihood estimate (MLE) of q_{ij} given by $(\alpha_{ij} - 1)/(\gamma_i - \beta)$ (see [12]). Lastly, $f_{\widehat{z}}(\widehat{z}|\boldsymbol{r}, \boldsymbol{Q}, \mathcal{Y})$ is the nonstandard density

$$\begin{aligned} f_{\widehat{z}}(\widehat{z}|\boldsymbol{r}, \boldsymbol{Q}, \mathcal{Y}) \quad \propto \quad & \prod_{j \neq i} \vartheta_{ij} \times \exp\left[\frac{1}{\widehat{\sigma}^2} \int_0^T \sum_{i=1}^{\ell} \mathbb{I}(Z(u) = i)\mathrm{d}Y_c(u)\right] \\ & \times \exp\left[\int_0^T \sum_{i=1}^{\ell} \left[\mathbb{I}(Z(u) = i)\left(-q_{ii} + \frac{r_i^2}{2\widehat{\sigma}^2}\right)\right]\mathrm{d}u\right]. \end{aligned} \tag{2.11}$$

In a given iteration, proposed values $\widehat{\boldsymbol{r}}'$ and $\widehat{\boldsymbol{Q}}'$ for $\widehat{\boldsymbol{r}}$ and $\widehat{\boldsymbol{Q}}$ are generated from (2.9) and (2.10), respectively. The candidate $\widehat{\boldsymbol{r}}'$ is always accepted with probability 1, but $\widehat{\boldsymbol{Q}}'$ is accepted with probability

$$\alpha(\widehat{\boldsymbol{Q}}, \widehat{\boldsymbol{Q}}') = 1 \wedge \frac{\pi_{\widehat{\boldsymbol{Q}}'}(\hat{Z}(0))}{\pi_{\widehat{\boldsymbol{Q}}}(\hat{Z}(0))},$$

where $\pi_{\widehat{\boldsymbol{Q}}'}(i)$ and $\pi_{\widehat{\boldsymbol{Q}}}(i)$, $i \in \mathcal{S}$, are the ith elements of the stationary distributions of CTMCs with respective generator matrices $\widehat{\boldsymbol{Q}}'$ and $\widehat{\boldsymbol{Q}}$. While obtaining $\widehat{\boldsymbol{r}}$ and $\widehat{\boldsymbol{Q}}$ from (2.9) and (2.10) is relatively straightforward, obtaining $\widehat{z}$ from (2.11) is not. Instead $\widehat{z}$ is modified using one of three procedures. The procedures are dubbed "the independence sampler", "the refinement sampler", and "the birth-death sampler", where p_{is}, p_{rs}, and p_{bd} denote the samplers' respective selection probabilities and $p_{is} + p_{rs} + p_{bd} = 1$. The process of updating $\widehat{z}$ is a Metropolis-Hastings (MH) step. In particular, a sampler is randomly selected to propose a realization of the estimated CTMC

sample path, denoted $\widehat{z}'$, and accept $\widehat{z}'$ with probability $\alpha_{is}(\widehat{z},\widehat{z}')$, $\alpha_{rs}(\widehat{z},\widehat{z}')$, or $\alpha_{bd}(\widehat{z},\widehat{z}')$ for the independence, refinement, and birth-death samplers, respectively. If $\widehat{z}'$ is accepted, $\widehat{z}_{v+1} = \widehat{z}'$; otherwise, $\widehat{z}_{v+1} = \widehat{z}$. In what follows, a description of each sampler is provided.

The independence sampler generates $\widehat{z}'$ independently of $\widehat{z}$ via simulation using $\widehat{\boldsymbol{Q}}$, where $\widehat{Z}'(0)$ is drawn from a uniform-discrete distribution with support $\mathcal{S}$. The acceptance probability is simply a ratio of the likelihoods of $\widehat{z}$ and $\widehat{z}'$, where for a given $\widehat{z}$, the likelihood, denoted $L(Y|\boldsymbol{r},\boldsymbol{Q},\widehat{z})$, is

$$L(\mathcal{Y}|\boldsymbol{r},\boldsymbol{Q},\widehat{z}) = \exp\left[\frac{1}{\hat{\sigma}^2}\int_0^T \sum_{i=1}^{\ell} r_i \mathbb{I}(Z(u)=i)\mathrm{d}Y_c(u) - \frac{1}{2\widehat{\sigma}^2}\int_0^T \sum_{i=1}^{\ell} r_i^2 \mathbb{I}(Z(u)=i)\mathrm{d}u\right],$$

so that

$$\alpha_{is}(\widehat{z},\widehat{z}') = 1 \wedge \frac{L(\mathcal{Y}|\boldsymbol{r},\boldsymbol{Q},\widehat{z}')}{L(\mathcal{Y}|\boldsymbol{r},\boldsymbol{Q},\widehat{z})}.$$

Before describing the refinement sampler, let M denote the total number of transitions in $\widehat{z}$, and ordered sets $\mathcal{I} = \{i_0,i_1,\ldots,i_M\}$ and $\tau = \{s_1, s_2, \ldots, s_M\}$ denote the state sequence and transition epochs, respectively, of $\widehat{z}$, where $i_0, i_1, \ldots, i_M \in \mathcal{S}$ and $s_1, s_2, \ldots s_M \in (0,T)$. The refinement sampler begins by randomly selecting a time $t' \sim \mathrm{U}[0,T]$. Assume that $t' \in [s_{m-1}, s_m]$ so that $\widehat{Z}(t') = i_m$. If the time interval is *internal*, that is $t' \in (s_1, s_M)$, the refinement sampler sets $s_{m-1} \to t'$ or $s_m \to t'$ with probability (w.p.) $1/2$. If $t' \in [0, s_1]$ or $t' \in [s_M, T]$, that is the interval is *external*, the sampler sets $s_1 \to t'$ or $s_M \to t'$ w.p. 1, respectively. Letting $\mathcal{I}'$ and τ' denote the state sequence and transition times of $\widehat{z}'$, the refinement sampler obtains $\widehat{z}'$ from $\widehat{z}$ by setting $\mathcal{I}' = \mathcal{I}$ and

$$\tau' = \begin{cases} \tau \setminus \{s_{m-1}\} \cup \{t'\}, & \text{if } s_{m-1} \to t', \\ \tau \setminus \{s_m\} \cup \{t'\}, & \text{if } s_m \to t'. \end{cases}$$

The refinement sampler changes the lengths of two adjacent sojourn intervals leaving all other intervals unchanged. Define a constant c, where

$$c = \begin{cases} \frac{1}{2}, & \text{if } t' \in [0, s_1] \cup [s_M, T] \text{ and } M > 1, \\ 2, & \text{if } t' \in (s_1, s_M), \text{ and either } s_1 \to t' \text{ or } s_M \to t', \\ 1, & \text{otherwise.} \end{cases}$$

The acceptance probability of the refinement sampler is

$$\alpha_{rs}(\widehat{z},\widehat{z}') = 1 \wedge c\frac{f_{\widehat{z}}(\widehat{z}'|\boldsymbol{r},\boldsymbol{Q},\mathcal{Y})}{f_{\widehat{z}}(\widehat{z}|\boldsymbol{r},\boldsymbol{Q},\mathcal{Y})},$$

where $f_{\widehat{z}}(.|\boldsymbol{r},\boldsymbol{Q},\mathcal{Y})$ is given in (2.11). The birth-death sampler constructs $\widehat{z}'$ by randomly choosing to add or delete an element of $\mathcal{I}$ along with the elements corresponding its transition times in τ.

If $M > 0$, a birth or death event occurs w.p. 1/2; otherwise, a death event occurs w.p. 1. For both birth and death events, select a random time $t' \sim \mathrm{U}[0, T]$. Assume that $t' \in [s_{m-1}, s_m]$, so $\widehat{Z}(t') = i_m$. For death events, set $\mathcal{I}' = \mathcal{I} \setminus \{i_m\}$ and

$$\tau' = \begin{cases} \tau \setminus \{s_{m-1}\}, & \text{w.p. } 1/2, \text{ if } m \neq 0 \text{ and } m \neq M, \\ \tau \setminus \{s_m\}, & \text{w.p. } 1/2, \text{ if } m \neq 0 \text{ and } m \neq M, \\ \tau \setminus \{s_{M-1}\}, & \text{w.p. } 1, \text{ if } m = M, \\ \tau \setminus \{s_1\}, & \text{w.p. } 1, \text{ if } m = 0, \end{cases}$$

If $i_{m-1} \neq i_{m+1}$, then operations $t \setminus \{s_{m-1}\}$ and $t \setminus \{s_m\}$ correspond to "left" and "right" deaths, respectively, otherwise they correspond to a "middle" death. In the case of a birth, both the new state and its sojourn time must be randomly selected. Births consist of "left", "middle", and "right" types, and each birth type is selected with equal probability. Let $t'' \sim \mathrm{U}[s_{m-1}, s_m]$, and i' denote the new birth state (for now not specified). For left or right births, $\tau' = \tau \cup \{t'\}$; otherwise, $\tau' = \tau \cup \{t'\} \cup \{t''\}$ for middle births. Insert i' into $\mathcal{I}$ such that its position corresponds to the interval $[s_{m-1}, t']$, $[t' \wedge t'', t' \vee t'']$, or $[t', s_m]$, for left, middle, or right births, respectively. Letting $\mathcal{S}^* \subset \mathcal{S}$ denote the subset of all states that if assigned to i' would result in the equivalence of the birth to a refinement sampler update, i' is randomly selected from the set $\mathcal{S} \setminus (\mathcal{S}^* \cup \{i_m\})$. The acceptance probability for the birth-death sampler is

$$\alpha_{bd}(\widehat{z}, \widehat{z}') = 1 \wedge \frac{q_p(\widehat{z}', \widehat{z})}{q_p(\widehat{z}, \widehat{z}')} \times \frac{p(\widehat{z}'|\widehat{z})}{p(\widehat{z}|\widehat{z}')} \times \frac{L(\mathcal{Y}|\boldsymbol{r}, \boldsymbol{Q}, \widehat{z}')}{L(\mathcal{Y}|\boldsymbol{r}, \boldsymbol{Q}, \widehat{z})}, \tag{2.12}$$

where for $\widehat{z}_1$ and $\widehat{z}_2$, $q_p(\widehat{z}_1, \widehat{z}_2)$ is the joint proposal density of $\widehat{z}_1$ and $\widehat{z}_2$ and $p(\widehat{z}_1|\widehat{z}_2)$ is the prior density of $\widehat{z}_1$, given $\widehat{z}_2$. The expressions for the ratios in (2.12) are complicated and depend on M, the birth/death type, and whether the birth/death interval is internal or external. Full details are given in [63].

To obtain the final estimates of $\boldsymbol{r}$ and $\boldsymbol{Q}$, recall the MCMC convergence result [42] that $(\widehat{\boldsymbol{r}}_v, \widehat{\boldsymbol{Q}}_v, \widehat{z}_v) \Rightarrow \pi(\boldsymbol{r}, \boldsymbol{Q}, \widehat{z})$ as $v \to \infty$, where $\pi(\boldsymbol{r}, \boldsymbol{Q}, \widehat{z})$ is the stationary distribution of $\mathcal{X}$. Therefore, to obtain $\widehat{\boldsymbol{r}}$ and $\widehat{\boldsymbol{Q}}$, the sampled vectors $\widehat{\boldsymbol{r}}_v$ and off-diagonal $\widehat{\boldsymbol{Q}}_v$ elements are averaged after a "burn in" period. Although implementing the MCMC procedure requires that the *order* of $\mathcal{Z}$, or unknown number of environment states ℓ, be known, it is possible to estimate the order by using the MCMC procedure in concert with the Bayesian information criterion (BIC) statistic [90]. Let $\widehat{\ell}$ denote the BIC estimate of the order ℓ, and define $\widehat{\boldsymbol{r}}^{(\ell)}$, $\widehat{\boldsymbol{Q}}^{(\ell)}$, and $\widehat{z}^{(\ell)}$ as the estimates of $\boldsymbol{r}$, $\boldsymbol{Q}$, and $\widehat{z}$, respectively, obtained using the MCMC procedure assuming $\mathcal{Z}$ is order ℓ.

The BIC is a *penalized maximum likelihood estimator* statistic [19], and the estimated order is computed as follows:

$$\widehat{\ell} = \underset{\ell \in \mathbb{N}}{\operatorname{argmin}} \left\{ -2 \ln L(\mathcal{Y} | \widehat{\boldsymbol{r}}^{(\ell)}, \widehat{\boldsymbol{Q}}^{(\ell)}, \widehat{z}^{(\ell)}) + \ell^2 \ln(N+1) \right\}. \tag{2.13}$$

The ℓ^2 in the second term of (2.13) corresponds to the total number of parameters estimated by the MCMC procedure (ℓ elements of $\boldsymbol{r}$ and $\ell^2 - \ell$ off-diagonal elements of $\boldsymbol{Q}$), and $N+1$ is the total number of signal observations.

Finally, a procedure for estimating $\widehat{Y}(t) \equiv \mathbb{E}[Y(t)]$, the expected value of $Y(t)$, using $\widehat{\boldsymbol{r}}$, $\widehat{\boldsymbol{Q}}$, and $\widehat{z}$ is described. This estimate is useful as an additional way to assess the quality of parameter estimates by comparing $\widehat{Y}(t)$ with the actual signal path. Let $\Delta\widehat{Y}_m \equiv \mathbb{E}[Y(s_m) - Y(s_{m-1})]$ be the expected change in Y on $[s_{m-1}, s_m]$, where

$$\Delta\widehat{Y}_m = \sum_{i=1}^{\widehat{\ell}} r_i (s_m - s_{m-1}) \mathbb{P}(Z_{m-1} = i | \mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}), \quad m = 1, 2, \ldots, M.$$

Assuming $\Delta\widehat{Y}_m$ for $m = 1, 2, \ldots, M$ can be obtained, then it is possible to compute

$$\widehat{Y}(t) = (t - s_\kappa)\Delta\widehat{Y}_{\kappa+1} + \sum_{i=1}^{\kappa} \Delta\widehat{Y}_m, \quad t \in [0, T],$$

where $\kappa \equiv \max\{m : s_m \le t\}$. Computing $\Delta\widehat{Y}_m$ requires that $\mathbb{P}(Z_m = i | \mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$ be determined for all $i = 1, 2, \ldots, \widehat{\ell}$ and $m = 0, 1, \ldots, M$. These conditional probabilities are estimated by applying a *forward-filtering-backward-smoothing* (FFBS) algorithm to $\widehat{z}$ using $\widehat{\boldsymbol{Q}}$ and $\widehat{\boldsymbol{r}}$. Let $\mathcal{Y}^{(m)} \equiv \{Y(t_j) : t_j \le s_m\}$ be the set of data observed up to time s_m, $m = 1, 2, \ldots, M$, and define a piecewise-linear function $Y_c^{(m)}(t)$ on $[0, s_m]$ such that $Y_c^{(m)}(t) = Y_c(t)$ for $t \in [0, s_m]$ and is zero otherwise. Lastly, for $m = 1, 2, \ldots, M$, let

$$\begin{aligned} f(Y_c(s_m) | Z_m = i, \mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) &= \exp\left[\frac{\widehat{r}_i}{\widehat{\sigma}^2} \int_{s_{m-1}}^{s_m} \mathrm{d}Y_c^{(m)}(u) - \frac{\widehat{r}_i^2}{2\widehat{\sigma}^2} \int_{s_{m-1}}^{s_m} \mathrm{d}u \right] \\ &= \exp\left[\frac{\widehat{r}_i}{\widehat{\sigma}^2} \left(Y_c^{(m)}(s_m) - Y_c^{(m)}(s_{m-1}) \right) - \frac{\widehat{r}_i^2}{2\widehat{\sigma}^2} (s_m - s_{m-1}) \right] \end{aligned}$$

be the conditional density of the observation $Y_c(s_m)$ at time s_m, given observations $\mathcal{Y}^{(m-1)}$, $Z(m) = i$, and estimated parameters $\widehat{\boldsymbol{r}}$ and $\widehat{\boldsymbol{Q}}$. Next the filtering and smoothing procedures are described.

The objective of the filtering procedure, originally developed in [11], is to compute $\mathbb{P}(Z_m = i | \mathcal{Y}^{(m)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$, for $i = 1, 2, \ldots, \widehat{\ell}$ and $m = 1, 2, \ldots, M$. This probability is computed recursively for $m = 1, 2, \ldots, M$ as follows:

1. For $i = 1, 2, \ldots, \widehat{\ell}$, compute $\mathbb{P}(Z_m = i|\mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$, the one-step ahead prediction probabilities for Z_m, where

$$\mathbb{P}(Z_m = i|\mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) = \sum_{j=1}^{\widehat{\ell}} \hat{p}_{ji} \mathbb{P}(Z_{m-1} = j|\mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}), \tag{2.14}$$

$$\hat{p}_{ji} = \begin{cases} -\hat{q}_{ji}/\hat{q}_{jj}, & j \neq i, \\ 0, & j = i, \end{cases}$$

and $\mathbb{P}(Z_0 = i|\mathcal{Y}^{(0)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) = 1/\widehat{\ell}$.

2. Compute the filtered probabilities $\mathbb{P}(Z_m = i|\mathcal{Y}^{(m)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$ for $i = 1, 2, \ldots, \widehat{\ell}$, using (2.14), where

$$\mathbb{P}(Z_m = i|\mathcal{Y}^{(m)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) = \frac{f(Y(s_m)|Z_m = i, \mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) \mathbb{P}(Z_m = i|\mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})}{\sum_{j=1}^{\widehat{\ell}} f(Y(s_m)|Z_m = j, \mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) \mathbb{P}(Z_m = j|\mathcal{Y}^{(m-1)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})}. \tag{2.15}$$

The filtered probabilities $\mathbb{P}(Z_m = i|\mathcal{Y}^{(m)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$ from (2.15) are then used in the smoothing procedure, to compute $\mathbb{P}(Z_m = i|\mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$, $i \in \mathcal{S}$, iteratively backward in time for $m = M-1, M-2, \ldots, 1$ as follows:

$$\mathbb{P}(Z_m = i|\mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) = \sum_{j=1}^{\widehat{\ell}} \frac{\hat{p}_{ij}\, \mathbb{P}(Z_m = i|\mathcal{Y}^{(m)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}) \mathbb{P}(Z_{m+1} = j|\mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})}{\sum_{k=1}^{\widehat{\ell}} \hat{p}_{kj} \mathbb{P}(Z_m = k|\mathcal{Y}^{(m)}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})}, \tag{2.16}$$

where $\mathbb{P}(Z_M = i|\mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$ is obtained from (2.15) with $m = M$ (see [27]). The initial distribution, $\mathbb{P}(Z_0 = i|\mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$ can be obtained by continuing computation of (2.16) to $m = 0$, and set $\widehat{\boldsymbol{\alpha}} = \mathbb{P}(Z_0 = i|\mathcal{Y}, \widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}})$.

In the next section, numerical examples that illustrate the performance of the MCMC procedure are presented for both simulated and real degradation processes. These examples show that the proposed signal estimation procedure is viable for wind turbine components operating in time-varying environments.

2.4 EXAMPLES OF COMPONENT LIFETIME ESTIMATION

In this section, the degradation model and inference procedure for estimating component lifetimes are illustrated on data obtained from both simulated and real degradation processes. To evaluate

performance, estimated environment parameters are obtained by sampling signals up to various percentages of their threshold crossing-times, and the parameters are used to compute expected lifetime estimates. These lifetime estimates are then compared to the actual component failure times (first passage times to the critical threshold, x_c).

First the effectiveness of the BIC statistic is illustrated in estimating ℓ when the environment experiences a large number of transitions prior to a component's failure. Consider a CTMC environment with $\ell = 3$ states and the following generator matrix and degradation rates:

$$\boldsymbol{Q} = \begin{bmatrix} -1.0 & 0.5 & 0.5 \\ 0.75 & -1.5 & 0.75 \\ 1.0 & 1.0 & -2.0 \end{bmatrix}, \quad \boldsymbol{r} = \begin{bmatrix} 1 & 5 & 10 \end{bmatrix}. \tag{2.17}$$

Degradation signals are generated according to equation (2.8) using a diffusion coefficient $\sigma = 0.316$ and observed at equidistant times $t_j - t_{j-1} = 0.01$, $j = 1, 2, \ldots, N$, where N is the random number of signal observations obtained prior to the first-passage time. Consider a simulated signal path for a component with threshold $x_c = 4000$. At $t = 963.67$ the signal reached the threshold at which time the real environment had transitioned $N = 1,337$ times. Computed values of the BIC statistic for $\ell = 2$, $\ell = 3$, and $\ell = 4$ are shown in Table 1. The BIC estimated order $\widehat{\ell} = 3$ corresponds to the actual environment order.

Table 1: BIC values.

ℓ	BIC value ($\times 10^4$)
2	-6.0940
3	-6.4869
4	-6.4418

Example 1: Now consider a component in the environment defined by (2.17) but for which failure is assumed to occur at threshold $x_c = 1000$. Although the environment does not evolve for sufficient time to reliably estimate the order, the component's expected lifetime can still be reasonably estimated. For what follows, the BIC estimated order $\hat{\ell} = 3$ is used from Table 1. Figure 2 shows a randomly generated path with $x_c = 1000$. The path was observed at $N = 23,466$ discrete times, and the estimates of the environment parameters are

$$\widehat{\boldsymbol{Q}} = \begin{bmatrix} -0.8190 & 0.4784 & 0.3406 \\ 0.7550 & -1.5112 & 0.7561 \\ 0.7037 & 0.8709 & -1.5745 \end{bmatrix}, \quad \widehat{\boldsymbol{r}} = \begin{bmatrix} 1.1046 & 4.8863 & 9.6783 \end{bmatrix}.$$

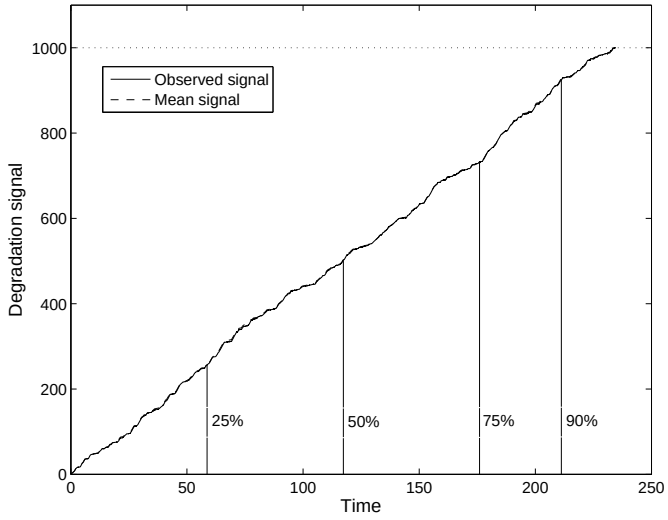

Figure 2: Observed and mean signal paths (CTMC environment).

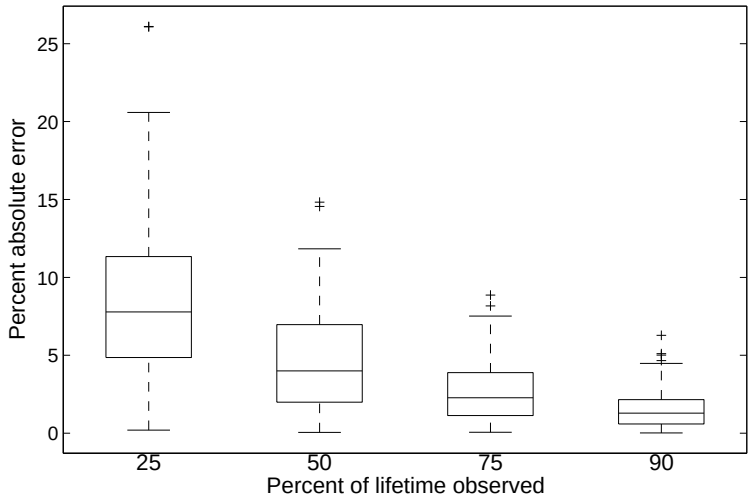

Figure 3: Box-and-whisker plot of percent absolute error (CTMC environment).

To validate the estimates, the mean signal path was computed using the FFBS algorithm and is also shown in Figure 2. The close correspondence of the observed and mean signal paths suggests that the estimates $(\widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}, \widehat{z})$ characterize the signal process well.

Next 100 random signal paths were generated for $x_c = 1000$ up to times $T^{(n)}(x_c)$, where $T^{(n)}(x_c)$ denotes the crossing time of the nth path, $n = 1, 2, \ldots, 100$. Each path was observed up to p-percent of its crossing time, where $p \in \{25, 50, 75, 90\}$. Let $t_p^{(n)} \equiv p \times T^{(n)}(x_c)/100$

denote the time corresponding to p percent of the nth component's lifetime. For each (n, p), set $\mathcal{T}_p^{(n)} = \{t_j : t_j \leq t_p^{(n)}\}$ and $\mathcal{Y}_p^{(n)} = \{Y(t_j) : t_j \in \mathcal{T}_p^{(n)}\}$ to obtain $(\widehat{\boldsymbol{r}}, \widehat{\boldsymbol{Q}}, \widehat{z})$ using the MCMC procedure. The expected lifetime is then computed using (2.3) and (2.5). The quality of the estimates is assessed by computing the percent absolute error, denoted $\epsilon_p^{(n)}$, where

$$\epsilon_p^{(n)} = \frac{|\tau_p^{(n)} - T^{(n)}(x_c)|}{T^{(n)}(x_c)} \times 100.$$

Figure 3 shows box-and-whisker plots of the percent absolute error of estimated lifetimes for each p. The plots indicate that both the median and variance of absolute error tend to decrease as p increases. Although the absolute error can be large (as evidenced by an outlier above 25%), the fact that the median absolute error is well below 10% indicates the procedure provides reasonable estimates of expected lifetime when the environment process is a CTMC.

Example 2: Next, a semi-Markov environment, whose inter-transition times are not exponentially distributed, is assumed to affect degradation. Signal paths are generated for three different semi-Markov environments with $\ell = 5$, $\ell = 10$, and $\ell = 20$ states, respectively. The transition probabilities, as well as the holding time distribution and degradation rates for each environment state, are randomly generated according to Table 2, where, $[\nu_1, \nu_2, \nu_3] \sim \mathrm{U}(0,1)^3$ such that $\nu_1 < \nu_2 < \nu_3$, $\nu_4 = 0.1 + 3.0\eta_4$, $\nu_5 = 0.1 + 3.0\eta_5$, and $\eta_4, \eta_5 \sim \mathrm{U}(0,1)$.

Table 2: Summary of holding-time distributions and degradation rates.

State index	Holding Time Distribution	Degradation Rate
1,4,7,9,12,15,19,20	Uniform$(0,1)$	Uniform$(0,2)$
2,10,13,16	Triangle(ν_1, ν_2, ν_3)	Uniform$(0,2)$
3,6,14,17	Gamma$(0.5, 0.5)$	20
5,8,11,19	Beta(ν_4, ν_5)	Uniform$(0,2)$

A key feature of the environments is that signal paths tend to have sudden increases in growth due to the relatively high degradation rate and potentially long holding times associated with the Gamma-distributed states. Each signal path is simulated with a diffusion coefficient of $\sigma = 0.707$ and observed at equidistant times $t_j - t_{j-1} = 0.01$, $j = 1, 2, \ldots, N$. Failure is assumed when the signal reaches a threshold $x_c = 3000$. The values of the BIC statistic obtained from each path are shown in Table 3.

Table 3: BIC values ($\times 10^3$) (semi-Markov environments).

ℓ	5-state	10-state	20-state
2	-3.6097	-3.7346	-3.5176
3	-0.6504	-2.6752	-0.9471
4	0.0221	-1.5648	-0.8553

Figures 4, 6, and 8 show simulated signal paths observed at $N = 62,389$, $N = 88,050$ and $N = 99,863$ discrete times for $\ell = 5$, $\ell = 10$, and $\ell = 20$, respectively.

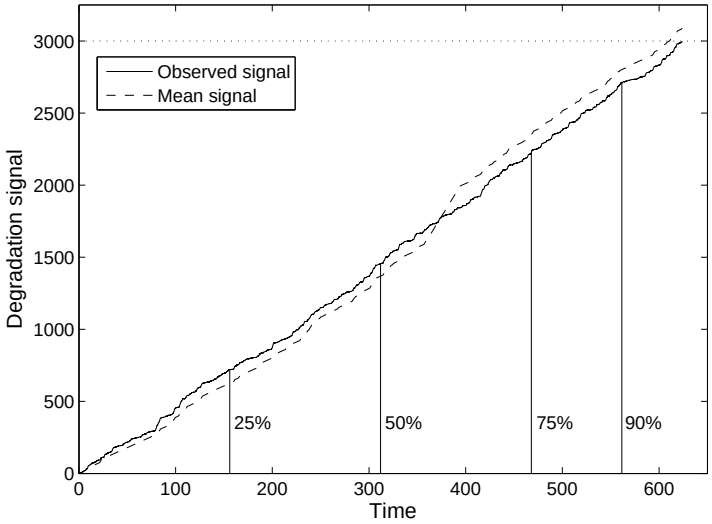

Figure 4: Observed and mean signal paths ($\ell = 5$ semi-Markov environment).

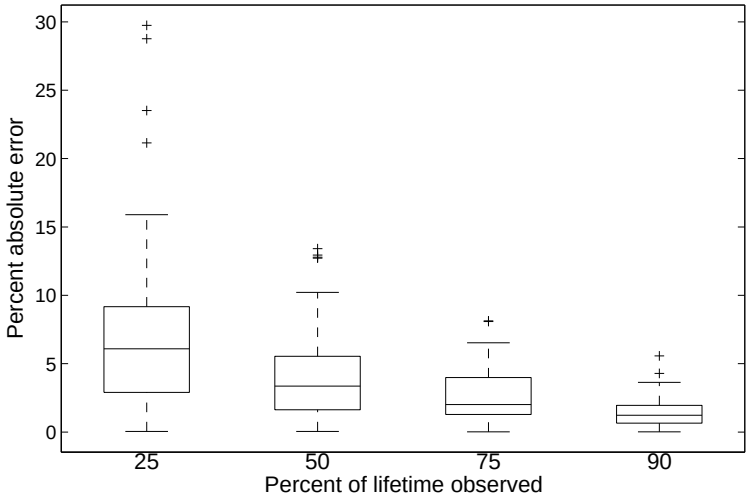

Figure 5: Box-and-whisker plot of percent absolute error ($\ell = 5$ semi-Markov environment).

For each environment, $\hat{\ell} = 2$ is the BIC-estimated order. However, $\hat{\ell} = 3$ will be used for the 10-state environment to provide more modeling fidelity since the BIC-values for $\ell = 2$ and $\ell = 3$ are relatively close. Estimates of the environment parameters obtained from the respective signal paths are as follows:

$$\widehat{\boldsymbol{Q}}_5 = \begin{bmatrix} -0.1846 & 0.1846 \\ 0.6280 & -0.6280 \end{bmatrix}, \quad \widehat{\boldsymbol{r}}_5 = \begin{bmatrix} 3.0489 & 10.8933 \end{bmatrix},$$

$$\widehat{\boldsymbol{Q}}_{10} = \begin{bmatrix} -0.0795 & 0.0511 & 0.0284 \\ 0.1585 & -0.2661 & 0.1076 \\ 0.2322 & 0.1906 & -0.4229 \end{bmatrix}, \quad \widehat{\boldsymbol{r}}_{10} = \begin{bmatrix} 2.1408 & 4.8024 & 9.4680 \end{bmatrix},$$

$$\widehat{\boldsymbol{Q}}_{20} = \begin{bmatrix} -0.0841 & 0.0841 \\ 0.4275 & -0.4275 \end{bmatrix}, \quad \widehat{\boldsymbol{r}}_{20} = \begin{bmatrix} 1.9587 & 8.4067 \end{bmatrix}.$$

Comparing the mean signal paths with the observed paths in Figures 4, 6, and 8 indicates that the parameters obtained based on the relatively small BIC-estimated orders adequately characterize the signal processes. One hundred simulated signal paths were generated for each environment and observed up to p percent of their respective failure times, $p \in \{25, 50, 75, 90\}$. Figures 5, 7, and 9 show box-and-whisker plots of percent absolute error for each p in the $\ell = 5$, $\ell = 10$, and $\ell = 20$ environments, respectively. In contrast to the CTMC example, the box-and-whisker plots show a strong propensity for large outlying errors. Such large errors are consistent with the

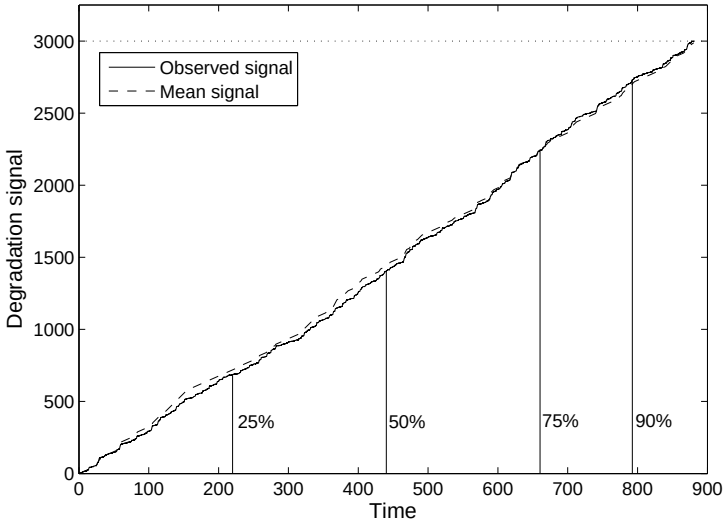

Figure 6: Observed and mean signal paths ($\ell = 10$ semi-Markov environment).

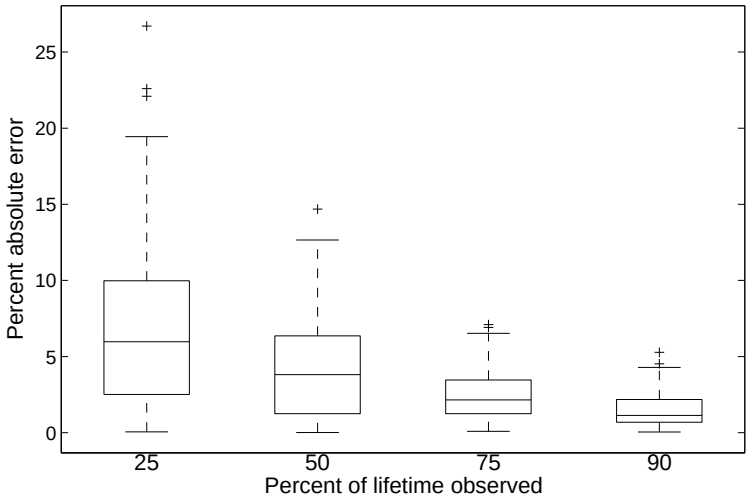

Figure 7: Box-and-whisker plot of percent absolute error ($\ell = 10$ semi-Markov environment).

irregularity induced by the Gamma-distributed states. However, despite these large outliers, the median absolute error and variance decrease rapidly with increasing p. As the median absolute error is below 5% for each environment by $p = 50$, the procedure seems quite suitable at estimating expected lifetimes in certain types of non-Markovian environments.

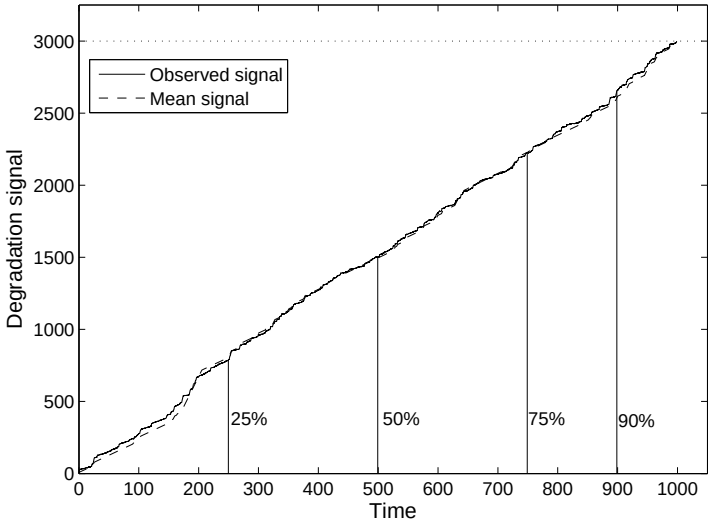

Figure 8: Observed and mean signal paths ($\ell = 20$ semi-Markov environment).

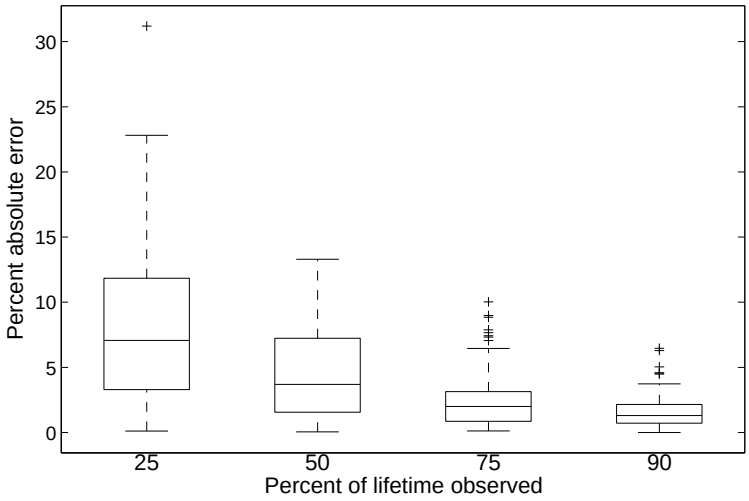

Figure 9: Box-and-whisker plot of percent absolute error ($\ell = 20$ semi-Markov environment).

Example 3: Now the performance of the inference procedure is illustrated on real degradation data originally obtained by Virkler *et al.* [104]. The data contain measurements of fatigue crack length over time (as measured by the number of load cycles) for 68 specimens for 2024-T3 aluminum alloy. Representative sample paths are shown in Figure 10, and specimen failure is assumed when the crack length exceeds 45 mm.

The curvature of the sample paths suggests crack length increases exponentially with time and is not consistent with the assumption that sample paths evolve according to (2.8). The MCMC

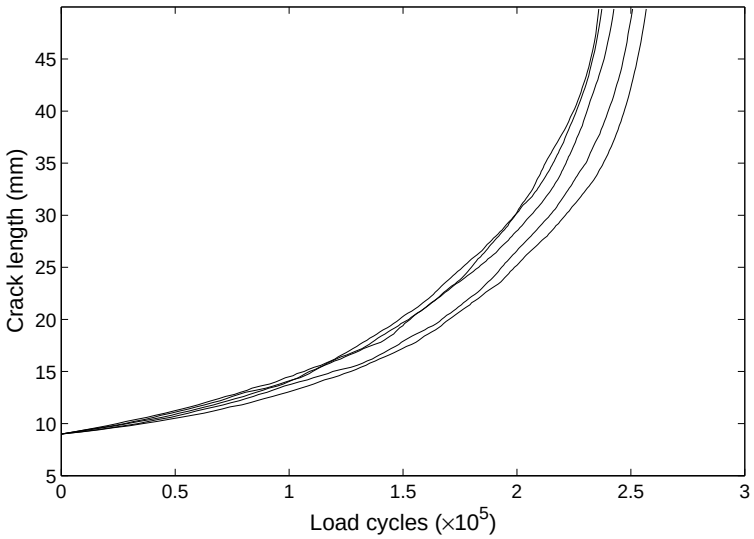

Figure 10: Sample paths of fatigue-crack propagation.

procedure will be applied to paths of log-transformed crack lengths which exhibit a milder curvature than the original paths. Since the MCMC procedure often performs poorly, or even completely fails, for signals that deviate strongly from the behavior of a switching diffusion model when $\sigma^2 \leq 10$, the transformed crack-lengths (and the associated transformed threshold) are multiplied by a factor of 100 before applying the MCMC procedure to increase σ^2 to an appropriate magnitude. All transformed paths are translated to begin at the origin.

To evaluate performance, path observations are restricted to $p \in \{75, 90, 100\}$ as the curvature is not strongly apparent until $p > 50$. Furthermore, it is often the case that elements of $\widehat{\boldsymbol{r}}$ obtained from the MCMC procedure are negative due to the relatively long regime of near-zero crack growth. When $\widehat{\boldsymbol{r}} \notin \mathbb{R}_+^\ell$, approximating the expected lifetimes via equations (2.3) and (2.5) is not possible. Instead, expected lifetimes are computed using the limiting result (2.6). The BIC values for a single sample path of log-transformed crack length are shown in Table 4, and the estimated number of environment states is $\hat{\ell} = 2$. Figure 11 shows one of the transformed sample paths and its corresponding mean path. Although the evolution of the true degradation sample path deviates significantly from that of a switching diffusion model, it can still be well characterized using the estimated model parameters. A box-and-whisker plot of the percent absolute error is shown in Figure 12 for all 68 sample paths observed up to each respective p. The exponential curvature

Table 4: BIC values (crack-length data).

ℓ	BIC value
2	-486.0313
3	15.5098
4	15.5098

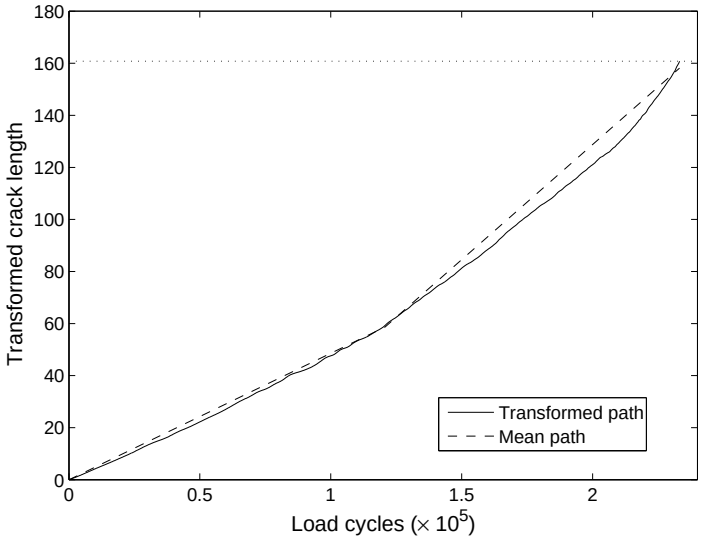

Figure 11: Observed and mean signal paths (crack-length data).

of the paths leads the procedure to consistently underestimate the expected lifetime. For $p = 75$, $p = 90$ and $p = 100$, the median absolute error is 24.4%, 12.1%, and 4.2% respectively, while the error variance tends to decrease slightly as p increases. The plots indicate that the MCMC procedure can provide reasonable estimates of lifetime in a fatigue-crack application provided a sufficiently large percentage of total lifetime is observed.

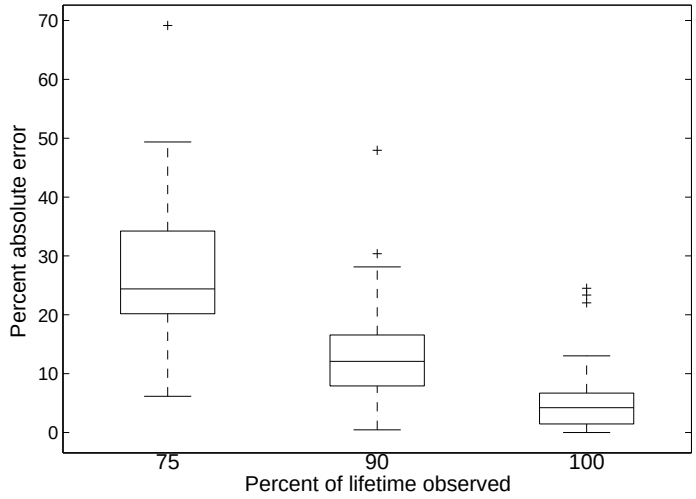

Figure 12: Box-and-whisker plot of percent absolute error (crack-length data).

Example 4: In this example, the procedure is shown to be useful at predicting bearing lifetimes from vibration data. The vibration data were collected by Gebraeel and Pan [41] for 25 bearings tested until failure under constant loading and rotational conditions. Figure 13 shows the vibration signal for a single bearing. The vibration signals are collected at 2 minute intervals and represent

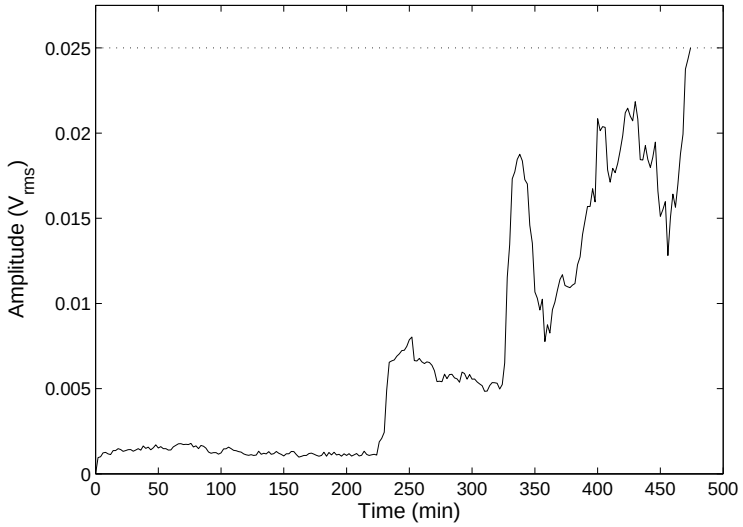

Figure 13: A vibration signal path.

the average amplitude over time of the defective bearing frequency and its first six harmonics. The threshold for bearing failure is based on the root means square (rms) of the vibration amplitude,

defined here as 0.025 V_{rms}. This vibration level corresponds to 2.2 Gs of acceleration and is consistent with the "danger-level" specification of $2.0 - 2.2$ Gs in the industrial standards for machinery vibration ISO 2372. Of the 25 bearings tested, one signal did not reach the threshold and will not be considered.

The vibration signals are characterized by an initial period at which they remain at a relatively constant level with only small oscillations due to random noise. After some time, the signals have a sudden, upward jump after which they undergo large oscillations and exhibit a general increasing trend until failure. The vibration signals clearly do not satisfy the assumption of (2.8), which requires a time-homogeneous diffusion coefficient. For applicability to our model, the bearing lifetimes are only predicted from the instant at which a given vibration signal exhibits its first large jump. Therefore, the entire bearing lifetimes are not predicted but rather the remaining lifetimes from the instant they first enter their failure regimes. The vibration signals are transformed in a similar manner to the Virker data. To "dampen" the signal oscillations, inference is performed on log-transformed amplitudes. The transformed amplitudes are also multiplied by a factor of 10^3 to increase $\hat{\sigma}$ to an appropriate magnitude for adequate performance of the MCMC procedure. All paths are translated to begin at the origin. For each path, at least one element of $\widehat{\boldsymbol{r}}$ is negative, so remaining lifetimes are again estimated using the limiting result (2.6).

The BIC values for a single transformed vibration signal are shown in Table 5, and the estimated order is $\hat{\ell} = 3$. Figure 14 shows a plot of one of the vibration paths (untransformed) and the

Table 5: BIC values (vibration data).

ℓ	BIC value ($\times 10^3$)
2	-0.2687
3	-1.4136
4	-1.3906

corresponding mean signal path. The close correspondence of the paths indicates that parameter estimates are obtained that characterize the vibration process well. Box-and-whisker plots of percent absolute error for all 24 paths observed at $p \in \{25, 50, 75, 90\}$ are shown in Figure 15. While the error variance remains relatively constant as p increases, the median absolute error decreases as p increases; however, it is not below 10% until $p = 90$. That the error variance is relatively constant and the absolute error is relatively large is likely attributable to the large oscillations of the vibration signal which make extremely precise lifetime prediction difficult.

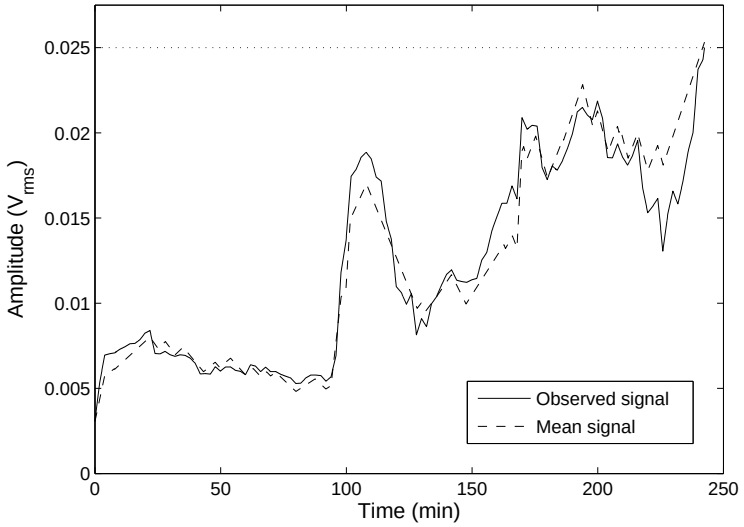

Figure 14: Observed and mean signal paths (vibration data).

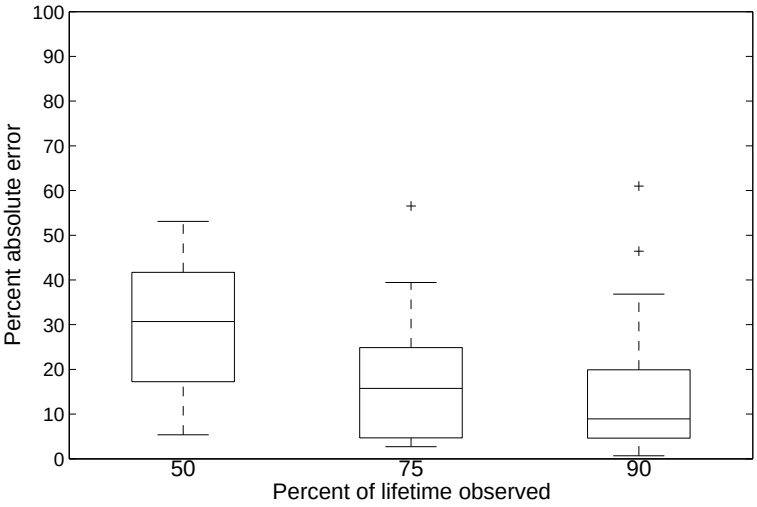

Figure 15: Box-and-whisker plot of percent absolute error (vibration data).

Example 5: The procedure is now used to predict the lifetime of a gear that operates in a wind turbine's drivetrain. The signal for gear wear corresponds to the effective number of load cycles and is computed from measurements of shaft torque and rotor speed. The measurements are from a dataset collected for a single wind turbine over a 19-month period that consists of 10-minute summary statistics. That is, for each 10-minute interval, a set of descriptive statistics (including the mean, minimum, maximum, standard deviation, etc.) is obtained for each turbine

parameter. After discarding intervals where data is missing due to downtime and other causes, summary statistics of turbine torque and rotor speed are available for 12,067 time periods. The failure threshold is defined here as the effective number of load cycles at which a gear tooth will fail. Let ζ_i and ω_i denote the mean torque (foot-pounds) and rotor speed (revolutions per minute) of the wind turbine over the ith time period, $i = 1, 2, \ldots, N = 12,067$.

To relate these parameters to the effective number of load cycles imposed on the gear tooth, first note that the relationship between torque and gear stress, denoted by ξ, can be approximated as

$$\xi = h\zeta,$$

where $h > 0$ is a constant [96]. Let ξ_w denote the stress amplitude corresponding to w cycles on a stress-life (S-N) curve of the gear material, and $(w', \xi_{w'})$ denote another known point on the S-N curve. The function $n(\xi)$, defined as the number of cycles to failure under stress amplitude ξ, is approximated by a power law as follows [89]:

$$\begin{aligned} n(\xi) &= w(\xi/\xi_w)^{1/b}, \\ b &= -\frac{\ln(\xi_w/\xi_{w'})}{\ln(w'/w)}. \end{aligned}$$

For stress amplitude ξ, define $\tilde{n}(\xi) \equiv w/n(\xi)$ as the effective number of load cycles imposed on a gear tooth during one cycle of amplitude ξ. Assuming a gear tooth experiences c load cycles per revolution, the total effective number of load cycles during the ith 10-minute period, denoted η_i, is approximately

$$\eta_i = 10\, c\omega_i\, \tilde{n}(\xi_i).$$

The signal value at t_i is computed in a straight-forward way as

$$Y(t_i) = \sum_{j=1}^{i} \eta_j, \quad i = 0, 1, 2, \ldots, N, \tag{2.18}$$

where $Y(0) \equiv 0$.

In this example, $h = c = 1$, $w = 10^7$, and $\xi_w = 63.0$ (a value close to the overall mean torque observed in the dataset). Three S-N curves are considered based on different values of $(w', \xi_{w'})$: Case (i) $(100.0, 0.3 \times 10^7)$, Case (ii) $(100.0, 0.5 \times 10^6)$, and Case (iii) $(100.0, 0.1 \times 10^6)$. Note that in progressing from Case (i) to Case (iii) the total effective number of load cycles imposed on material increases for each fixed $\xi > \xi_w$. Signal paths are simulated to a failure threshold $x_c = 10^7$ using a bootstrap technique; whereby, the torque and rotor speed values are generated by randomly

sampling four-hour, continuous blocks of torque and rotor speed data. The BIC statistic values computed using a simulated path for each case are shown in Table 6, and the estimated orders are $\hat{\ell} = 2$, $\hat{\ell} = 2$, and $\hat{\ell} = 3$ for Cases (i), (ii), and (iii), respectively. Figures 16, 18, and 20 show the

Table 6: BIC values (drivetrain gear example).

ℓ	Case (i) ($\times 10^3$)	Case (ii)	Case (iii)
2	-3.5863	-381.9225	2.0317
3	-2.1527	-134.1971	-16.9701
4	-0.8514	-27.2337	15.4704

observed signals used to compute the BIC statistic for Cases (i), (ii), and (iii), respectively, along with the mean signal paths obtained from the estimated parameters corresponding to each path. That the mean and observed signal paths correspond relatively closely for all three cases indicates that the degradation process for the gear tooth is well characterized by the estimated parameters.

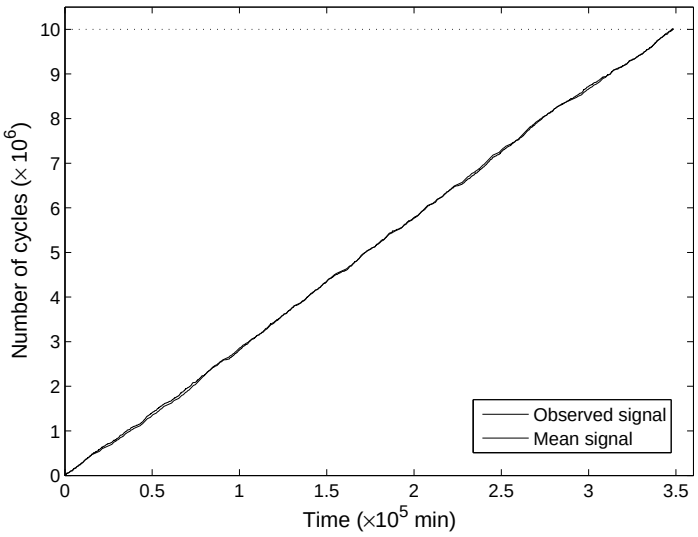

Figure 16: Observed and mean signal paths (gear tooth degradation Case (i)).

Figures 17, 19, and 21 show box-and-whisker plots of the percent absolute error for Cases (i), (ii), and (iii), respectively with 100 randomly generated signal paths observed up to p percent of their respective failure times, $p \in \{25, 50, 75, 90\}$. For fixed p, the median and variance of the absolute error increases from Case (i) to Case (iii)—an intuitive result as Cases (ii) and (iii) can have relatively large degradation rates leading to more irregular sample paths. For Case (i), the

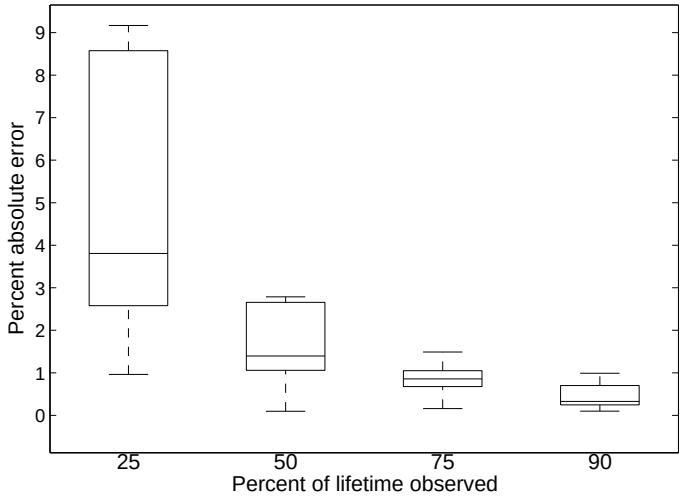

Figure 17: Box-and-whisker plot of percent absolute error (gear tooth degradation Case (i)).

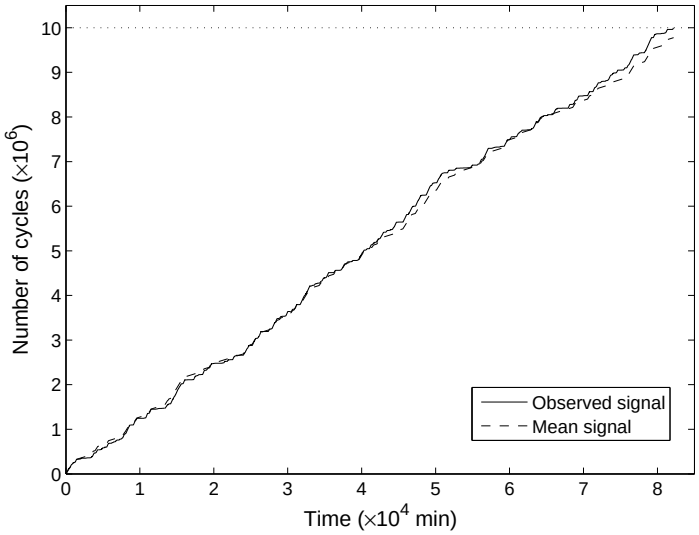

Figure 18: Observed and mean signal paths (gear tooth degradation Case (ii)).

most realistic case in a practical application, the median absolute error is less than 5% for all p, and for Cases (ii) and (iii), the median absolute error is below 10% and 20%, respectively, for all p. That the inference procedure also performs well in the more extreme Cases (ii) and (iii) suggests that the procedure is well-suited to estimate the lifetimes of gears that operate in relatively extreme degradation environments.

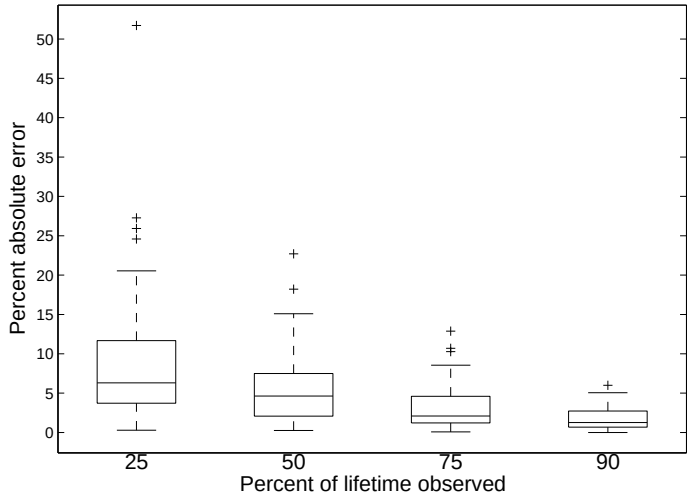

Figure 19: Box-and-whisker plot of percent absolute error (gear tooth degradation Case (ii)).

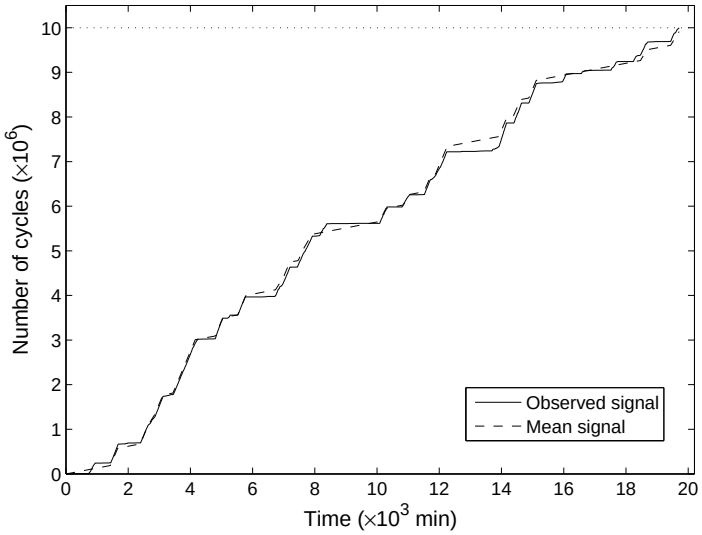

Figure 20: Observed and mean signal paths (gear tooth degradation Case (iii)).

Example 6: Using the same dataset as in Example 5, consider a wind turbine shaft bearing. The degradation is again defined as the effective number of load cycles but is computed based on rotor speed and bearing temperature measurements. For this example, it is assumed that the primary determinant of the degradation rate at a given time is the viscosity of the bearing lubricant. Let $\rho = \nu_0/\nu_1$ be the relative lubricant viscosity, where ν_0 and ν_1 denote the specified and actual lubricant viscosity, and $a(\rho)$ be a life adjustment factor that is a function of ρ. If T_b denotes the

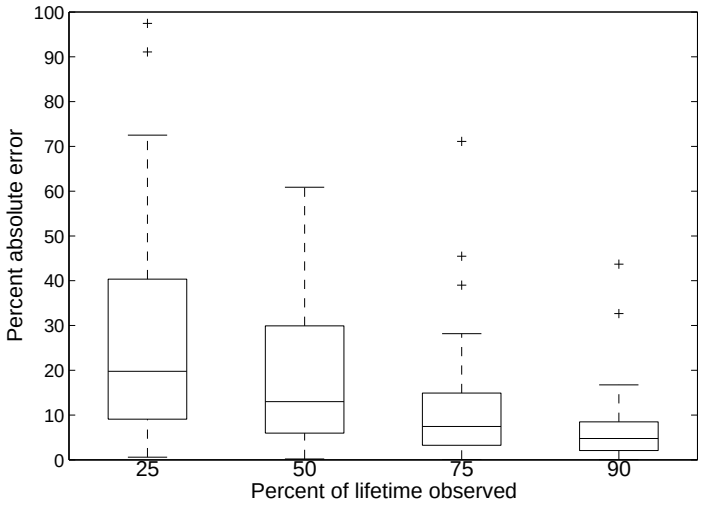

Figure 21: Box-and-whisker plot of percent absolute error (gear tooth degradation Case (iii)).

bearing's specified lifetime, defined as the time at which the bearing is expected to survive with 90% probability given operational loading conditions, then it is known [46] that

$$T_b \propto a(\rho).$$

The relationship between lubricant viscosity (in centistokes (cSt)) and bearing temperature (in Celsius (C)), denoted θ, can be characterized using the Ubbelohde-Walther equation [8] as follows:

$$\ln\left[\ln(\nu_1 + 0.7)\right] = u + v\ln(\theta + 273.15), \tag{2.19}$$

where u and v are constants that depend on the lubricant type.

Summary data for rotor speed and bearing temperature are available for 25,421 10-minute periods after disregarding periods with spurious or missing data. To generate degradation signals, rotor speed and bearing temperature data are bootstrapped by randomly sampling from four-hour, contiguous blocks of wind turbine data. For the ith 10-minute interval, ρ_i is computed as a function of the bearing temperature using (2.19) and assuming $v_0 = 150$ cSt. To obtain the constants u and v in (2.19), it is assumed that a commercially available wind turbine lubricant is used with viscosities of 150 cSt at 40° C and 20.7 cSt at 100° C, so that $u = 17.77$ and $v = 2.81$. Two different cases are considered for the function $a(\rho)$. In the first case, the life adjustment factor, denoted $a_1(\rho)$, is estimated based on data for standard steel bearings [29]. In the second case, the life adjustment factor, denoted $a_2(\rho)$, is selected to be significantly larger and smaller than $a_1(\rho)$

for relatively small and large values of ρ, respectively, and is intended to represent a material with more exaggerated degradation characteristics. Plots of $a_1(\rho)$ and $a_2(\rho)$ are shown in Figure 22(a) and Figure 22(b), respectively. The total number of effective load cycles in the ith period for Case

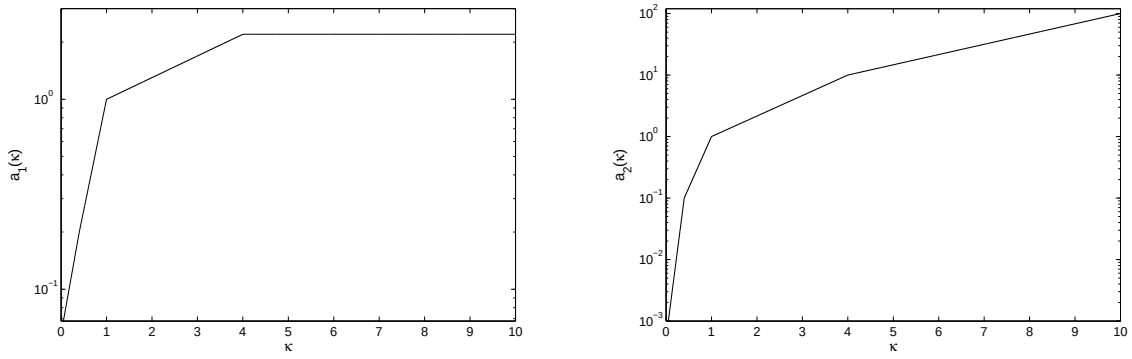

(a) Plot of $a_1(\rho)$ for a steel bearing (Case (i)).

(b) Plot of $a_2(\rho)$ (Case (ii)).

Figure 22: Assessing degradation signal estimation (Case (i)).

$j \in \{1, 2\}$ is denoted η'_{ij}, where

$$\eta'_{ij} = 10\, a_j(\rho_i)\, \omega_i, \quad i = 1, 2, \ldots,$$

and the cumulative degradation signal at a given time is computed as in equation (2.18) but with η'_{ij} in place of η.

In this example, $x_c = 10^7$. Simulated degradation paths are shown in Figures 23 and 25 for Cases (i) and (ii), respectively, and the BIC values computed by observing each degradation path are shown in Table 7. For both cases, $\hat{\ell} = 2$. The estimated parameters appear to characterize the

Table 7: BIC values (shaft bearing example).

ℓ	Case (i) ($\times 10^3$)	Case (ii) ($\times 10^3$)
2	-3.8946	-1.7135
3	-2.4338	-1.2094
4	-1.3038	0.0182

environments well as the mean signal paths for each case closely correspond with the observed signal paths. Figures 24 and 26 show box-and-whisker plots of the percent absolute error of estimated lifetime for 100 randomly generated sample paths in Cases (i) and (ii), respectively, observed up

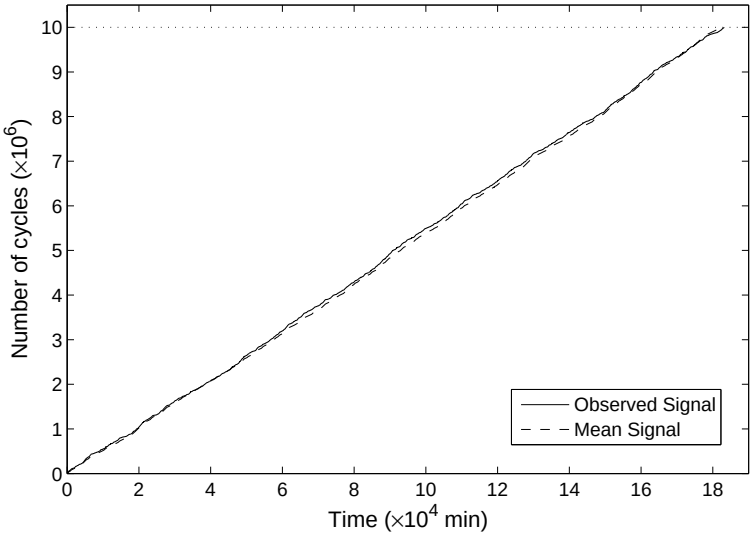

Figure 23: Observed and mean signal paths (shaft bearing degradation Case (i)).

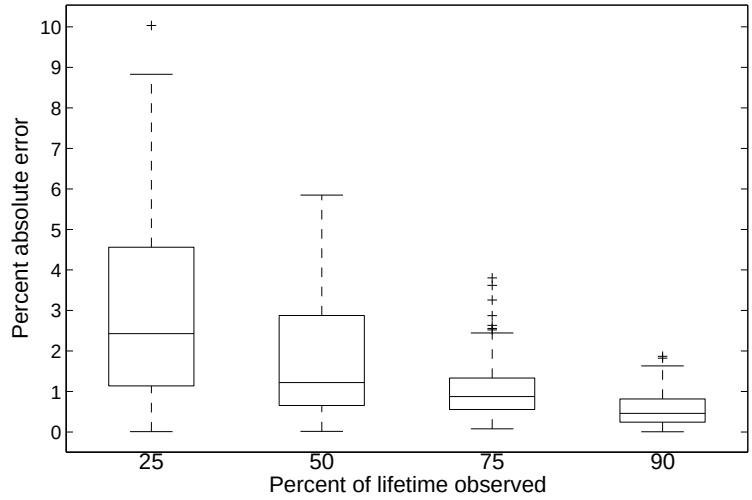

Figure 24: Box-and-whisker plot of percent absolute error (shaft bearing degradation Case (i)).

to p-percent of lifetime, $p \in \{25, 50, 75, 90\}$. The absolute error in each case is well below 10%, and both the median and variance of absolute error decrease as p increases. It is intuitive that for identical p-values both the median and variance of the error in Case (i) is smaller than in Case (ii) due to Case (i) having a more limited range of possible degradation rates. The relatively small absolute errors indicates that the procedure provides an effective way to estimate the lifetimes of wind turbine shaft bearings.

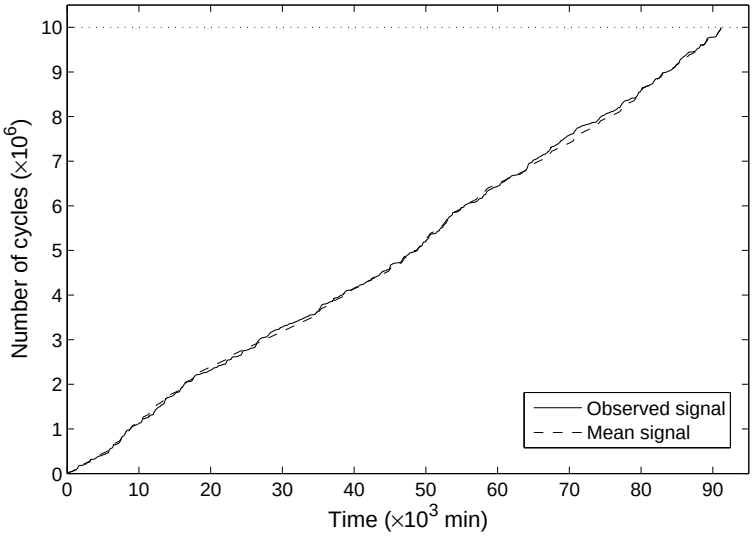

Figure 25: Observed and mean signal paths (shaft bearing degradation Case (ii)).

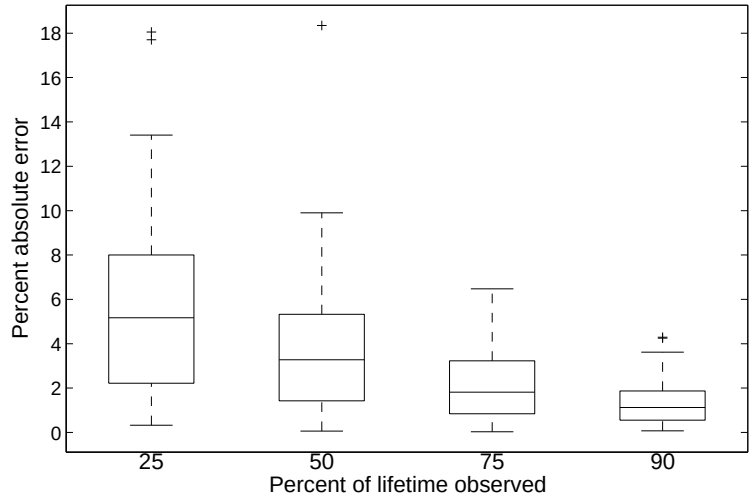

Figure 26: Box-and-whisker plot of percent absolute error (shaft bearing degradation Case (ii)).

In this chapter, a mathematical model for the environment-driven degradation of a wind turbine component was presented. An MCMC-based statistical procedure was adapted to estimate the environment parameters from an observed signal of component degradation. In the next chapter, two models are presented for optimally replacing a wind turbine component based on the estimated environment parameters.

3.0 REPLACEMENT FOR OBSERVABLE DEGRADATION

In this chapter, two models for optimally replacing a single wind turbine component are presented using a partially observed Markov decision process (POMDP) framework. Both models assume that the environment is partially-observable and degradation is observable, but only one model explicitly accounts for stochastic downtime costs incurred during component replacements. The POMDP models are formulated on a continuous belief space, and structural results are derived for the optimal policies. Subsequently, a numerical solution procedure and numerical examples are presented to illustrate the optimal policies of both models.

3.1 SUMMARY OF RELEVANT LITERATURE

The optimal replacement or repair of a stochastically-degrading system is a classical operations research problem that has been considered by many researchers. Some important surveys are given in [79, 92, 108]. A number of recent efforts have focused on replacement models where degradation is partially or imperfectly observed. For instance, Maillart [69] considers the optimal maintenance of a component whose partially-observed degradation process evolves as a discrete-time Markov chain (DTMC). Structural policy results are proven for a model that minimizes the long-run average cost of maintenance and then modified for the case of imperfect observations using a heuristic. Maillart [70] considers a similar model in an adaptive maintenance scheduling application with silent failures. Ivy [49] considers a component that degrades in a stepwise manner through increasing degradation levels until a silent failure occurs. The component's degradation state is observed indirectly by a continuous random variable having a probability distribution conditioned on the actual degradation state. Byon *et al.* [16] present a wind turbine maintenance application that uses a DTMC-based degradation model and includes such factors as adverse weather conditions, repair interruptions and

delays, and maintenance lead times. Other replacement models incorporate degradation processes that are continuous in time. Makis and Jiang [71] consider optimally replacing a component whose degradation evolves as a CTMC on a finite state space and is observed imperfectly at discrete times. Associated with each observation is a probability mass function conditioned on the true degradation state. The replacement model is formulated as an optimal stopping problem, and the long-run expected replacement cost per unit time is minimized. This work is extended in [60] to include multivariate observations that are normally distributed with mean and covariance matrices determined by the degradation state. A similar model is used by Jiang, *et al.* [51] to maximize the component's long-run expected availability per unit time.

Typically replacement models assume that the component can attain only a finite number of degradation levels; however, some researchers have considered cases where the underlying degradation is continuous. In [117], Zhou *et al.* develop a partially observable semi-Markov process (POSMDP) model which is continuous in both time and state, where component degradation evolves according to a Gamma-based state space model. A Monte Carlo-based density projection procedure (see [116]) is used to reduce the infinite-dimensional belief space to a finite dimension so that the problem can be formulated as a Markov decision process (MDP) model on the lower-dimensional belief space. The MDP is solved numerically using policy iteration for both long-run average cost and availability objectives.

Most existing replacement models view the component's degradation state as the only relevant observational variable for making replacement decisions. Often, degradation can be affected by other factors such as the component's external operating environment. Despite its importance, relatively few replacement models consider replacement decisions based on knowledge of the environment state, in addition to knowledge of the cumulative degradation. In [102], Ulukus *et al.* consider optimally replacing a component whose operating environment evolves as a CTMC with state space $\mathcal{S}$. When the environment is in state i, the component accumulates degradation at a constant rate $r_i > 0$, where $i \in \mathcal{S}$. At fixed times, both the environment state and cumulative degradation are observed, and a decision is made whether to preventively replace the component. The model is formulated as an infinite-horizon MDP with the objective of minimizing the total expected discounted replacement cost, and for each environment state, an associated threshold of cumulative degradation is shown to exist such that preventive replacement is optimal in the given environment state when degradation exceeds the threshold.

The replacement models that currently exist in the literature omit one or more important features required when considering the replacement of a wind turbine component. Realistic replacement models require incorporating not only the notion that a component's degradation is determined largely by the exogenous, operating environment but also the potential uncertainty in discerning the exact environment state. Models that seek to minimize the long-run average replacement cost per unit time, as opposed to the total discounted cost, are more suitable for wind turbine components, as replacement decisions can potentially be made frequently in response to electronic sensor readings. In addition, a realistic model must consider the potential revenue losses incurred during replacement operations when the turbine cannot produce electricity. These revenue losses are stochastic and depend on the prevailing weather conditions (i.e. the environment state) during the replacement period. In the remainder of this chapter, two replacement models are presented that utilize the same degradation dynamics as considered by Ulukus *et al.* [102] but assume the environment is partially observable. Moreover, while one model assumes the replacement costs are deterministic, the other model incorporates stochastic downtime costs. Both models seek to minimize the long-run average replacement cost per unit time.

3.2 REPLACEMENT WITH DETERMINISTIC COSTS

In this section, we present a replacement model that assumes deterministic costs for preventive and reactive replacements and minimizes the long-run average cost of replacements per unit time. To account for an uncertain environment state, the replacement problem will be formulated using a partially-observed Markov decision process (POMDP) model. Let $I = \{1, 2, \ldots\}$ be the set of decision epochs, and define a *period* as the time between two successive epochs, where each period is of length $\delta > 0$. The component's cumulative degradation at decision epoch n, X_n, is perfectly observed (i.e. known with certainty) for each $n \geq 1$. After each observation, a decision is immediately made whether to preventively replace the component, and the set of feasible actions at each epoch is $\mathcal{A} \equiv \{0, 1\}$, where actions 0 and 1 signify "do nothing" and "preventively replace," respectively. By contrast, the environment is assumed to be unobservable; that is, the environment state is inferred from an observed degradation increment and represented using the concept of a "belief state." The belief state of the environment at the nth decision epoch is the probability distribution $\pi_n = [\pi_n^{(1)}, \pi_n^{(2)}, \ldots, \pi_n^{(\ell)}]$, where $\pi_n^{(i)} \equiv \mathbb{P}(Z_n = i | \Delta X_n, \pi_{n-1})$ and $\Delta X_n \equiv X_n - X_{n-1}$ is

the random degradation increment during the nth period. The belief space of the environment is then the ℓ-dimensional probability simplex $\Pi \equiv \left\{[\pi^{(1)}, \pi^{(2)}, \ldots, \pi^{(\ell)}] : \sum_{i=1}^{\ell} \pi^{(i)} = 1\right\}$. Immediately following the observation of a degradation increment, the belief state is recursively updated. Let $W_{ij}(x,t) = \mathbb{P}(X(t) \leq x, Z(t) = j | Z(0) = i)$ for $i, j \in \mathcal{S}$ and $t \geq 0$, and define the density

$$p(u, j|i) \equiv \left. \frac{\partial W_{ij}(x, \delta)}{\partial x} \right|_{x=u},$$

where $p(u, j|i)$ is defined for all $u \in (r_1\,\delta, r_\ell\,\delta)$ such that $u \neq r_k\,\delta$, $k \in \mathcal{S}$. For $j \in \mathcal{S}$, define

$$T_j(u, \pi) \equiv \mathbb{P}(Z_{n+1} = j | \Delta X_n = u, \pi_n = \pi) = \sum_{i \in \mathcal{S}} \mathbb{P}(Z_{n+1} = j | \Delta X_n = u, Z_n = i) \pi^{(i)},$$

where

$$\mathbb{P}(Z_{n+1} = j | \Delta X_n = u, Z_n = i) = \begin{cases} \dfrac{p(u, j|i)}{\sum_{k \in \mathcal{S}} p(u, k|i)}, & u \in (r_1\,\delta, r_\ell\,\delta) \text{ s.t. } u \neq r_l\,\delta,\ l \in \mathcal{S}, \\ \mathbb{I}(j = k), & u = r_k\,\delta,\ k \in \mathcal{S}. \end{cases}$$

The quantity $T_j(u, \pi)$ is the probability that the environment is in state j at the next decision epoch given an initial belief state π and an observed increase in degradation over the period of u. Given $\pi_n = \pi$, and a realized degradation increment $u = x_{n+1} - x_n$, it follows that $\pi_{n+1} = \boldsymbol{T}(u, \pi)$, where $\boldsymbol{T}(u, \pi) \equiv [T_1(u, \pi), T_2(u, \pi), \ldots, T_\ell(u, \pi)]$.

The replacement problem is now formulated as a POMDP model on the belief space $B \equiv [0, x_c] \times \Pi$. Let the belief state of the POMDP at the nth decision epoch be denoted denoted by $b^{(n)} = (x_n, \pi_n)$. For $(x, \pi) \in B$ and $(x', \pi') \in B$, define the conditional probability $K((x, \pi), (x', \pi')) \equiv \mathbb{P}(\Delta X_{n+1} \leq x' - x, \pi_{n+1} = \pi' | \pi_n = \pi)$. Note that $K((x, \pi), (x', \pi')) = K((0, \pi), (x' - x, \pi'))$, where for $u \geq 0$,

$$K((0, \pi), (u, \pi')) = \begin{cases} \displaystyle\int_0^u \sum_{i \in \mathcal{S}} q(v|i) \pi^{(i)} \mathrm{d}v + \sum_{i \in \mathcal{S}} \mathbb{I}(u \geq r_i \delta) \pi^{(i)} \exp(q_{ii}\,\delta), & \text{if } \pi' = \boldsymbol{T}(u, \pi), \\ 0, & \text{if } \pi' \neq \boldsymbol{T}(u, \pi), \end{cases}$$

and $q(v|i) \equiv \sum_{j \in \mathcal{S}} p(v, j|i)$. Denote the transition kernel density between $(x, \pi) \in B$ and $(x', \pi') \in B$ as

$$k(x' - x, \pi) \equiv \left. \frac{\partial}{\partial u} \left[K((0, \pi), (u, \pi')) \right] \right|_{u = x' - x},$$

where for $u \geq 0$,

$$k(u,\pi) = \begin{cases} \pi^{(i)} \exp(q_{ii}\,\delta), & u = r_i\,\delta, \quad i \in \mathcal{S}, \\ \displaystyle\sum_{j\in\mathcal{S}} q(u|j)\pi^{(j)}, & u \neq r_i\,\delta, \quad i \in \mathcal{S}. \end{cases}$$

Degradation is assumed to be observed instantaneously at the beginning of each period with a cost $c_0 > 0$. Preventive replacements (if chosen) occur immediately following the degradation observation with a cost c_1 ($0 < c_0 < c_1 < \infty$) and are assumed to be instantaneous. If the component fails between two decision epochs, a reactive replacement is instantly performed with costs $c_1 + c_2$, where $c_2 > 0$ is a penalty. All components are assumed to begin operation in belief state $(0, \pi_s)$, where π_s is the stationary distribution of the environment. Define a *policy* as a function $a : B \to \mathcal{A}$, where $a(X_n, \pi_n)$ is the action taken in $(X_n, \pi_n) \in B$, and let $\mathcal{P}$ denote the set of all possible policies. The objective is to find the policy that minimizes the long-run average cost of replacements per unit time, denoted by γ, where

$$\gamma = \inf_{a\in\mathcal{P}} \mathbb{E}_a \left\{ \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} c_0 + c_1\,\mathbb{I}\{a(X_n,\pi_n) = 1\} + (c_1 + c_2)\,\mathbb{I}\{a(X_n,\pi_n) = 0, \widehat{H}(X_n,\pi_n) = 1\} \right\}$$

and $\widehat{H}(X_n, \pi_n)$ is the event that the component fails between decision epochs n and $n+1$, given $(X_n, \pi_n) \in B$.

The optimality equations are now provided. Let $V(x,\pi)$ be the minimum relative cost per unit time, given that a component starts operation in $(x,\pi) \in B$, and define $V_0(x,\pi)$ and $V_1(x,\pi)$ as the relative costs if either no action or preventive replacement, respectively, are taken in $(x,\pi) \in B$. The expected survival time of the component in the next period given $(x,\pi) \in B$ is denoted $\tau(x,\pi)$, where

$$\tau(x,\pi) = \sum_{i\in\mathcal{S}} \left[\sum_{j\in\mathcal{S}} \int_0^\delta W_{ij}(x_c - x, t)\,\mathrm{d}t \right] \pi^{(i)}. \tag{3.1}$$

The optimality equations are as follows:

$$V(x,\pi) = \min\{V_1(x,\pi), V_0(x,\pi)\}, \quad (x,\pi) \in B \tag{3.2}$$

where for $\mathbb{I}^+(x,u) \equiv \mathbb{I}(x+u \geq x_c)$, $\mathbb{I}^-(x,u) \equiv \mathbb{I}(x+u < x_c)$, and $V_\pi(x,u) \equiv V(x+u, \boldsymbol{T}(u,\pi))$,

$$\begin{aligned} V_1(x,\pi) &= c_0 + c_1 + V(0,\pi_s) \\ V_0(x,\pi) &= c_0 + \int_0^\infty \mathbb{I}^+(x,u)\,[c_1 + c_2 + V(0,\pi_s)]\,k(u,\pi)\mathrm{d}u \\ &\quad + \int_0^\infty \mathbb{I}^-(x,u) V_\pi(x,u) k(u,\pi)\mathrm{d}u - \gamma\tau(x,\pi). \end{aligned}$$

3.2.1 Structural Results

An exact analytical solution for the optimality equations (3.2) is not known except in trivial cases. Nonetheless, it is possible to prove some basic structural properties of the resulting optimal policy. To begin, a nonnegative lower bound for γ is established.

Definition 3.1. *For $(x, \pi) \in B$, let*

$$H(x, \pi) \equiv \int_0^{x_c} \mathbb{I}^+(x, u) k(u, \pi)\, \mathrm{d}u, \quad x \geq 0, \tag{3.3}$$

be the probability that the component fails in the next period, given $(x, \pi) \in B$.

The function $H(x, \pi)$ gives the probability that the component's cumulative degradation exceeds the critical threshold before the next decision epoch, given the current belief state is (x, π).

Lemma 3.1. *The average cost of an optimal policy is bounded below as follows:*

$$\gamma > \frac{c_0}{\delta}. \tag{3.4}$$

Proof. The lower bound can be established by considering the average cost of a policy for the case when $c_1 = c_2 = 0$. To establish this bound rigorously, note that immediately following replacement it is optimal to do nothing; otherwise, $\gamma = \infty$. Set $V(0, \pi_s) = 0$ and observe that $V(0, \pi_s) = V_0(0, \pi_s) = 0$. This implies that

$$c_0 + (c_1 + c_2) H(0, \pi_s) + \int_0^{\infty} I^-(x, u) V_\pi(0, u) k(u, \pi) \mathrm{d}u - \gamma \tau(0, \pi_s) = 0. \tag{3.5}$$

Solving (3.5) for γ gives

$$\gamma \;=\; \frac{c_0 + (c_1 + c_2) H(0, \pi_s) + \displaystyle\int_0^{\infty} I^-(0, u) V_\pi(0, u) k(u, \pi) \mathrm{d}u}{\tau(0, \pi_s)} > \frac{c_0}{\delta}. \tag{3.6}$$

□

The lower bound (3.4) is quite loose; nonetheless, the existence of this bound is useful to establish some basic properties of the optimal policy.

Lemma 3.2. *Let $C_0 = c_0 + c_1 + V(0, \pi_s)$. Then for all $(x, \pi) \in B$, $V(x, \pi) \leq C_0$.*

Proof. For $(x, \pi) \in B$, $V(x, \pi) = \min\{V_0(x, \pi), V_1(x, \pi)\} \leq V_1(x, \pi) = C_0$. □

Definition 3.2. *Let the subset of the belief space where preventive replacement is optimal be denoted by $D = \{(x, \pi) \in B : V_1(x, \pi) \leq V_0(x, \pi)\}$ and D^c denote the complement of D. Define the sets $D_\pi \equiv \{x : (x, \pi) \in D\}$ and $D_x \equiv \{\pi : (x, \pi) \in D\}$.*

Characterizing the optimal policy is equivalent to providing a complete description of the region D, which generally is difficult to obtain. In fact, it is not known whether D is convex or even connected. Nonetheless, the next theorem and subsequent corollary place some bounds on the region.

Theorem 3.1. *If the component survives the next decision period w.p. 1 for $(x, \pi) \in B$, then $(x, \pi) \in D^c$.*

Proof. If the component survives in the next period w.p. 1 for $(x, \pi) \in B$, then

$$\begin{aligned} V_0(x,\pi) &= c_0 + \int_0^\infty I^-(x,u) V_\pi(x,u) k(u,\pi) \mathrm{d}u - \gamma\delta \\ &\leq c_0 + \int_0^\infty I^-(x,u)[c_0 + c_1 + V(0,\pi_s)] k(u,\pi) \mathrm{d}u - \gamma\delta \quad \text{(Lemma 3.2)} \\ &= 2c_0 + c_1 + V(0,\pi_s) - \gamma\delta \\ &< c_0 + c_1 + V(0,\pi_s) \\ &= V_1(x,\pi). \end{aligned}$$

Therefore, $(x, \pi) \in D^c$. □

Simply put, preventive replacement is never optimal if the probability of the component failing before the next decision epoch is zero. Theorem 3.1 leads immediately to the following corollary:

Corollary 3.1. $D \subseteq \{(x, \pi) \in B : x \in (x_c - r_\ell\, \delta, x_c]\}$.

Proof. For all $(x, \pi) \in B$ such that $x \in [0, x_c - r_\ell\, \delta]$, the component will survive the next period w.p. 1 and $\{(x, \pi) \in B : x \in [0, x_c - r_\ell\, \delta]\} \subseteq D^c$. Taking the complements of both sets yields the result. □

Corollary 3.1 places a significant boundary on D. Namely, D is confined to exist only in the subset of B where $x \geq x_c - r_\ell\, \delta$ so that the probability of failure before the next decision epoch is nonzero. Figure 27 provides a schematic that depicts this bound in the case of a three-state environment. (Note that the third component of the environment state is determined by the values of the first two components.)

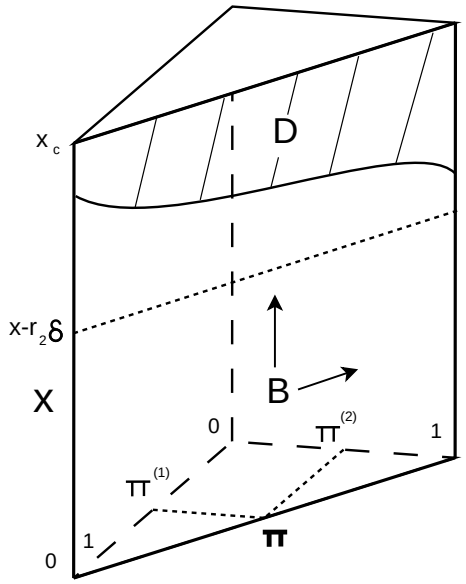

Figure 27: Depiction of boundary for preventive replacement region D.

Theorem 3.2. *If* $\tau(x,\pi) = 0$*, then* $(x,\pi) \in D$*.*

Proof. If $\tau(x,\pi) = 0$, it follows that failure occurs in the next period w.p. 1, and

$$\begin{aligned} V_0(x,\pi) &= c_0 + \int_0^\infty \mathbb{I}^+(x,u)\left[c_1 + c_2 + V(0,\pi_s)\right] k(u,\pi)\mathrm{d}u \\ &\qquad + \int_0^\infty \mathbb{I}^-(x,u) V_\pi(x,u) k(u,\pi)\mathrm{d}u - \gamma\tau(x,\pi) \\ &= c_0 + c_1 + c_2 + V(0,\pi_s) \\ &> c_0 + c_1 + V(0,\pi_s) \\ &= V_1(x,\pi). \end{aligned}$$

Therefore, $(x,\pi) \in D$. □

If $\tau(x,\pi) = 0$, it is not possible to reduce the average cost by delaying preventive replacement because a costlier, reactive replacement would be immediately required. So it is quite intuitive that the optimal action is an immediate, preventive replacement.

Corollary 3.2. $(x_c,\pi) \in D$ *for all* $\pi \in \Pi$*.*

Proof. Since $\tau(x_c,\pi) = 0$ for $\pi \in \Pi$, the result follows from Theorem 3.2. □

It is tempting to conclude that a sufficient condition for the optimality of preventive replacement in a given belief state is that component failure occurs in the next period w.p. 1; however, a stronger condition is required to account for pathological instances of replacement problems where, for example, the inspection interval is longer than the component lifetime. The sufficient condition is based on the following lemma.

Lemma 3.3. *An upper bound for γ is*

$$\gamma \leq c_0 + (c_1 + c_2)\frac{r_\ell}{x_c}. \tag{3.7}$$

Proof. It is known that $T(x_c) \geq x_c/r_\ell$ w.p. 1 [55]. Therefore, an upper bound for the policy cost can be obtained by considering the average cost of allowing a component to continuously operate to failure at this rate, without preventive replacements.

$$\gamma \leq \frac{\left\lfloor \dfrac{x_c}{r_\ell} \right\rfloor c_0 + c_1 + c_2}{\dfrac{x_c}{r_\ell}} \leq c_0 + (c_1 + c_2)\frac{r_\ell}{x_c}.$$

□

As for the lower bound (3.4), the upper bound (3.7) is not particularly tight; however, its existence illuminates some additional properties of the optimal policy in the case of a particular class of replacement instances.

Theorem 3.3. *If the component fails w.p. 1 in the next period for $(x, \pi) \in B$, then it is optimal to preventively replace if*

$$\tau(x, \pi) \leq \frac{c_2\, x_c}{c_0 + (c_1 + c_2)r_\ell}. \tag{3.8}$$

Proof. If the component fails w.p. 1 in the next period, then

$$\begin{aligned} V_0(x, \pi) &= c_0 + \int_0^\infty \mathbb{I}^+(x, u)\left[c_1 + c_2 + V(0, \pi_s)\right] k(u, \pi)\mathrm{d}u \\ &\quad + \int_0^\infty \mathbb{I}^-(x, u) V_\pi(x, u) k(u, \pi)\mathrm{d}u - \gamma\tau(x, \pi) \\ &= c_0 + c_1 + c_2 + V(0, \pi_s) - \gamma\tau(x, \pi). \end{aligned}$$

For preventive replacement to be optimal, it must be the case that $V_0(x, \pi) \geq V_1(x, \pi)$. Therefore, it follows that

$$c_0 + c_1 + c_2 + V(0, \pi_s) - \gamma\tau(x, \pi) \geq c_0 + c_1 + V(0, \pi_s),$$

which implies $c_2 - \gamma\tau(x,\pi) \geq 0$. From Lemma 3.3, it follows that

$$c_2 - \gamma\tau(x,\pi) \geq c_2 - \left[c_0 + (c_1 + c_2)\frac{r_\ell}{x_c}\right]\tau(x,\pi).$$

So a sufficient condition for preventive replacement is

$$c_2 - \left[c_0 + (c_1 + c_2)\frac{r_\ell}{x_c}\right]\tau(x,\pi) \geq 0,$$

or

$$\tau(x,\pi) \leq \frac{c_2\, x_c}{c_0 + (c_1 + c_2)r_\ell}.$$

□

The inequality (3.8) can now be used to define a class of replacement instances where preventive replacement is *always* optimal when component failure is certain before the next decision epoch.

Corollary 3.3. *Preventive replacement is always optimal in state $(x,\pi) \in B$ when component failure occurs in the next decision period w.p. 1 if*

$$c_2 - \delta\left[c_0 + (c_1 + c_2)\frac{r_\ell}{x_c}\right] \geq 0. \tag{3.9}$$

Proof. If failure occurs w.p. 1 in the next decision period for $(x,\pi) \in B$, then $\tau(x,\pi) \leq \delta$. Therefore, if the left-hand side of inequality (3.8) at least as large as δ, the condition will always be satisfied. Rearranging terms yields condition (3.9). □

It should be noted that condition (3.9) will typically be satisfied in "practical" replacement problems where the degradation threshold is much larger than the degradation rates and the inspection interval is not too long. A *practical* replacement problem is assumed to have the following properties:

(i) $x_c \gg r_\ell\, \delta$,
(ii) $\delta \ll \dfrac{c_2\, x_c}{(c_1 + c_2)r_\ell}$.

3.2.2 Threshold Policy with Respect to Degradation

In this subsection, it is proved that the optimal replacement policy has a threshold-type structure. That is, for each environment belief state $\pi \in \Pi$, there exists a threshold cumulative degradation level above which it is always optimal to preventively replace. To establish this result, the next lemma bounds $V_0(x,\pi)$.

Lemma 3.4. *For each $x \in [0, x_c]$ and $\pi \in \Pi$, $V_0(x,\pi)$ is bounded as follows:*

$$\underline{V}_0(x,\pi) \leq V_0(x,\pi) \leq \bar{V}_0(x,\pi), \tag{3.10}$$

where,

$$\begin{aligned} \underline{V}_0(x,\pi) &= (C_0 - c_1) + (c_1 + c_2)H(x,\pi) - \gamma\tau(x,\pi), \\ \bar{V}_0(x,\pi) &= c_0 + C_0 + (c_2 - c_0)H(x,\pi) - \gamma\tau(x,\pi). \end{aligned} \tag{3.11}$$

Proof. By Lemma 3.2, $V_\pi(x,u) \leq C_0$ for all $u \geq 0$. Let $M \equiv c_1 + c_2 + V(0,\pi_s)$, then

$$\begin{aligned} V_0(x,\pi) &= c_0 + \int_0^\infty \mathbb{I}^+(x,u)[c_1 + c_2 + V(0,\pi_s)]k(u,\pi)\mathrm{d}u \\ &\qquad + \int_0^\infty \mathbb{I}(x+u<x_c)V_\pi(x,u)k(u,\pi)\mathrm{d}u - \gamma\tau(x,\pi) \\ &= c_0 + M\int_0^\infty \mathbb{I}^+(x,u)k(u,\pi)\mathrm{d}u \\ &\qquad + \int_0^\infty \mathbb{I}(x+u<x_c)V_\pi(x,u)k(u,\pi)\mathrm{d}u - \gamma\tau(x,\pi) \\ &\leq c_0 + MH(x,\pi) + C_0[1 - H(x,\pi)] - \gamma\tau(x,\pi) \\ &= c_0 + C_0 + (M - C_0)H(x,\pi) - \gamma\tau(x,\pi) \\ &= c_0 + C_0 + (c_2 - c_0)H(x,\pi) - \gamma\tau(x,\pi). \end{aligned}$$

Likewise, $V(x,\pi) \geq V(0,\pi_s)$ for all $\pi \in \Pi$. Therefore,

$$\begin{aligned} V_0(x,\pi) &= c_0 + M\int_0^\infty \mathbb{I}^+(x,u)k(u,\pi)\mathrm{d}u \\ &\qquad + \int_0^\infty \mathbb{I}(x+u<x_c)V_\pi(x,u)k(u,\pi)\mathrm{d}u - \gamma\tau(x,\pi) \\ &\geq c_0 + M\int_0^\infty \mathbb{I}^+(x,u)k(u,\pi)\mathrm{d}u \\ &\qquad + V(0,\pi_s)\int_0^\infty \mathbb{I}(x+u<x_c)k(u,\pi)\mathrm{d}u - \gamma\tau(x,\pi) \\ &= c_0 + MH(x,\pi) + (C_0 - c_0 - c_1)[1 - H(x,\pi)] - \gamma\tau(x,\pi) \\ &= (C_0 - c_1) + [M - V(0,\pi_s)]H(x,\pi) - \gamma\tau(x,\pi) \\ &= (C_0 - c_1) + (c_1 + c_2)H(x,\pi) - \gamma\tau(x,\pi). \end{aligned}$$

□

Although the monotonicity of $V_0(x,\pi)$ in x, and by extension $V(x,\pi)$, cannot be established, both the lower and upper bounds in (3.10) are monotone increasing in x. As x increases, it can

be shown that $\underline{V}(x,\pi)$ will eventually exceed $V_1(x,\pi)$, which is constant in x, so that there is a threshold of cumulative degradation above which preventive replacement is optimal. A depiction of the relationship between the value functions and the bounds of Theorem 3.4 is shown in Figure 28. The next lemma states that if preventive replacement is optimal for some degradation level in a

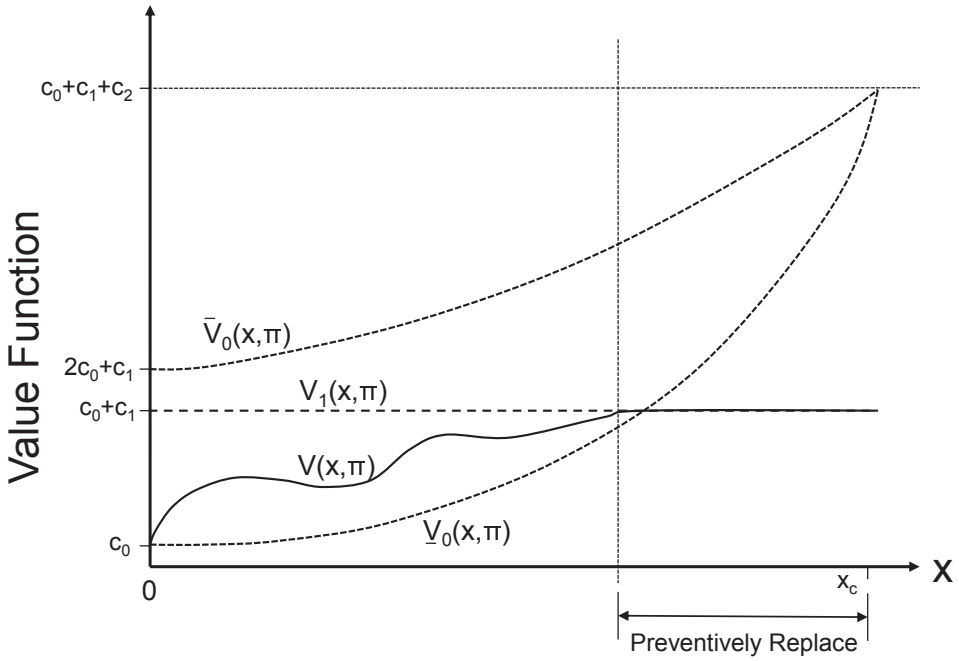

Figure 28: Depiction of the value functions and their bounds.

fixed environment state, then preventive replacement is also also optimal for all higher degradation levels. The lemma is particularly useful because it establishes that D_π consists of only a single interval of degradation levels.

Lemma 3.5. *For each $\pi \in \Pi$, if $(x,\pi) \in D$, then $(x',\pi) \in D$ for all $x' \geq x$.*

Proof. Assume that $(x,\pi) \in D$ and there exists $x^* > x$ such that $(x^*,\pi) \in D^c$. Let $B^D_{(x^*,\pi)}$ as the set of all $(x,\pi) \in B$ that are reachable from (x^*,π) under policy D prior to a replacement event. Consider a policy defined by D^* that is obtained by modifying D as follows:

(i) Set $D^* = D \setminus \{(x,\pi)\}$

(ii) For all $(y,\pi') \in B^D_{(x^*,\pi)}$, set $D^* = D^* \cup \{(y+x-x^*,\pi')\}$ if $(y,\pi') \in D$; otherwise, set $D^* = D^* \setminus \{(y+x-x^*,\pi')\}$.

Note that prior to a replacement event the evolution of two sample paths starting from (x,π) and (x^*,π) under policy D^* differ only through a translation of degradation levels by $x^* - x$. For any

sample path starting at (x^*, π) that exceeds the threshold at some time t, the corresponding sample path starting from (x, π) will be less than the threshold at time t. Therefore, $\gamma_{D'} < \gamma_D$, and D cannot be optimal. This implies that $(x', \pi) \in D$ for all $x' \geq x$. □

Using the lower bound of (3.10) and Lemma 3.5, the following theorem establishes the existence of a threshold policy with respect to x for fixed π.

Theorem 3.4 (Threshold policy for fixed π)**.** *There exists, for each $\pi \in \Pi$, a threshold $x_\pi < x_c$ at which it is optimal to replace for all $x \geq x_\pi$.*

Proof. A sufficient condition to replace in $(x, \pi) \in B$ is that $\underline{V}_0(x, \pi) \geq V_1(x, \pi)$. Therefore, $\underline{V}_0(x, \pi) \geq V_1(x, \pi)$ implies that

$$(C_0 - c_1) + (c_1 + c_2)H(x, \pi) - \gamma\tau(x, \pi) \geq C_0,$$

which implies

$$(c_1 + c_2)H(x, \pi) - \gamma\tau(x, \pi) \geq c_1. \tag{3.12}$$

Note that the left side of inequality (3.12) is increasing in x as $H(x, \pi)$ is increasing in x and $\gamma\tau(x, \pi)$ is decreasing in x. Furthermore, the inequality is strict at the point (x_c, π) for all $\pi \in \Pi$ since $H(x_c, \pi) = 1$ and $\tau(x_c, \pi) = 0$. Let $x' = \inf\{x : (c_1 + c_2)H(x, \pi) - \gamma\tau(x, \pi) \geq c_1\} < x_c$, then $x' < x_c$ and (3.12) is satisfied for all $x \geq x'$. Therefore, $\{x : x \geq x'\} \subseteq D_\pi$. In addition, it follows by Lemma 3.5 that $D_\pi = \{x : x \geq x_\pi\}$ for some $x_\pi \leq x_c$. However, since $\{x : x \geq x'\} \subseteq D_\pi$, it follows that $x_\pi < x_c$. □

Theorem 3.4 establishes that, for each environment belief state, there is an associated degradation threshold such that it is optimal to preventively replace the wind turbine component when its cumulative degradation exceeds the threshold. The fact that replacement thresholds are dependent on the environment state is critical because it allows information about the dynamic operating environment to affect replacement decisions. Such environment information is constantly monitored at wind farms and is readily available for utilization in decision-making.

3.2.3 Threshold Policy with Respect to the Environment State

Ideally, it is desirable to obtain an analogous threshold policy to Theorem 3.4 with respect to the belief state π for fixed x. However, such a result requires characterization of the set D_x, which

is difficult in general. But the following corollary provides some insights into the properties of D_x—namely that D_x is an increasing set in x.

Corollary 3.4. *For all x and x' such that $x \leq x'$, $D_x \subseteq D_{x'}$.*

Proof. Consider $\pi \in D_x$ for $x \geq x_\pi$. Then for $x' \geq x$, $(x', \pi) \in D$ by Theorem 3.4 so that $\pi \in D_{x'}$, and $D_x \subseteq D_{x'}$. □

Therefore, the subset of Π in which preventive replacement is optimal for a given degradation level x will tend to grow larger as x increases. A more specific characterization can only be established for the case $\ell = 2$ when $\boldsymbol{Q}$ and $\boldsymbol{r}$ satisfy certain conditions. Under those assumptions, it is possible to show that the optimal policy has a threshold-type structure with respect to π. Assume the environment has infinitesimal generator $\boldsymbol{Q}$ and degradation rates given by

$$\boldsymbol{Q} = \begin{bmatrix} -\alpha & \alpha \\ \beta & -\beta \end{bmatrix}, \quad \boldsymbol{r} = \begin{bmatrix} r_0 & r_1 \end{bmatrix}, \tag{3.13}$$

where $\alpha > \beta > 0$ and $r_1 > r_0 > 0$. Define $\lambda = \beta/\alpha$ and $\widehat{\boldsymbol{P}} = \boldsymbol{I} + \boldsymbol{Q}/\alpha$, where $\widehat{\boldsymbol{P}}$ is the transition probability matrix of the uniformized environment process given by

$$\widehat{\boldsymbol{P}} = \begin{bmatrix} 0 & 1 \\ \lambda & 1-\lambda \end{bmatrix}.$$

Sericola [91] derived a formula to compute the reward distribution, which can be written as follows when $\ell = 2$:

$$\boldsymbol{W}(x,t) = \sum_{n=0}^{\infty} e^{-\alpha t} \frac{(\alpha t)^n}{n!} \sum_{k=0}^{n} \binom{n}{k} x^k (1-x)^{n-k} \boldsymbol{C}(n,k), \tag{3.14}$$

where $\boldsymbol{W}(x,t) = [W_{ij}(x,t)]$, $i, j \in \{1,2\}$, and $\boldsymbol{C}(n,k) = [c_{ij}(n,k)]$ such that

$$c_{ij}(n,k) = \begin{cases} c_{2j}(n-1,k), & i=1,\ j \in \{1,2\},\ k \in \{0,1,\ldots,n-1\}, \\ \left(\widehat{\boldsymbol{P}}^n\right)_{1j}, & i=1,\ j \in \{1,2\},\ k=n, \\ 0, & i=2,\ j \in \{1,2\},\ k=0, \\ \lambda c_{1j}(n-1,k-1) + (1-\lambda) c_{2j}(n-1,k-1), & i=2,\ j \in \{1,2\},\ k \in \{1,2,\ldots,n\}. \end{cases}$$

For notational brevity, let $c_{i\cdot}(n,k) \equiv c_{i1}(n,k) + c_{i2}(n,k)$, $i \in \{1,2\}$. In what follows, it is shown that a threshold policy exists with respect to π for fixed x if $c_{1\cdot}(n,k) \leq c_{2\cdot}(n,k)$ for all $k \leq n$ and $n \geq 0$. To prove this condition holds for $\boldsymbol{C}(n,k)$, an induction argument is used that relies on several preliminary lemmas.

Lemma 3.6. *For $j \in \{1, 2\}$,*

$$c_{2j}(n,k) = \begin{cases} \lambda^{n-k} c_{2j}(2k-n, 2k-n) \\ +(1-\lambda)\left[\sum_{s=1}^{n-k} \lambda^{n-k-s} c_{2j}(2k-n+2s-1, 2k-n+s-1)\right], \quad k \geq n/2, \\ 0, \quad k < n/2. \end{cases}$$

Proof. When $k \geq n/2$,

$$\begin{aligned} c_{2j}(n,k) &= \lambda c_{21}(n-2,k-1) + (1-\lambda) c_{2j}(n-1,k-1) \\ &= \lambda^2 c_{21}(n-4,k-2) + \lambda(1-\lambda) c_{2j}(n-3,k-2) + (1-\lambda) c_{2j}(n-1,k-1) \\ &\quad \vdots \\ &= \lambda^{n-k} c_{2j}(2k-n,2k-n) + \lambda^{n-k-1}(1-\lambda) c_{21}(2k-n+1,2k-n) \\ &\qquad + \lambda^{n-k-2}(1-\lambda) c_{2j}(2k-n+3,2k-n+1) + \cdots + \lambda(1-\lambda) c_{2j}(n-3,k-2) \\ &\qquad + (1-\lambda) c_{21}(n-1,k-1) \\ &= \lambda^{n-k} c_{2j}(2k-n,2k-n) + (1-\lambda)\left[\sum_{s=1}^{n-k} \lambda^{n-k-s} c_{2j}(2k-n+2s-1,2k-n+s-1)\right], \end{aligned}$$

and when $k < n/2$,

$$\begin{aligned}
c_{2j}(n,k) &= \lambda^k c_{2j}(n-2k,0) + \lambda^{k-1}(1-\lambda)c_{2j}(n-2k+1,0) \\
&\qquad +\lambda^{k-2}(1-\lambda)c_{2j}(n-2k+3,n-k+1)+\cdots \\
&\qquad \cdots+\lambda(1-\lambda)c_{2j}(n-3,k-2)+(1-\lambda)c_{2j}(n-1,k-1) \\
&= (1-\lambda)\left[\sum_{s_0=1}^{k}\lambda^{k-s_0}c_{2j}(n-2k+2s_0-1,s_0-1)\right] \\
&= (1-\lambda)^2\left[\sum_{s_0=1}^{k}\lambda^{k-s_0}\sum_{s_1=1}^{s_0-1}\lambda^{s_0-1-s_1}c_{2j}(n-2k+2s_1,s_1-1)\right] \\
&= (1-\lambda)^3\left[\sum_{s_0=1}^{k}\lambda^{k-s_0}\sum_{s_1=1}^{s_0-1}\lambda^{s_0-1-s_1}\sum_{s_2=1}^{s_1-1}\lambda^{s_1-1-s_2}c_{2j}(n-2k+2s_2+1,s_2-1)\right] \\
&= \vdots \\
&= (1-\lambda)^k\left[\sum_{s_0=1}^{k}\lambda^{k-s_0}\sum_{s_1=1}^{s_0-1}\lambda^{s_0-1-s_1}\cdots\right. \\
&\qquad \left.\cdots\sum_{s_{k-1}=1}^{1}\lambda^{s_{k-2}-1-s_{k-1}}\underbrace{c_{2j}(n-k+2s_{k-1}-1,0)}_{=0}\right]=0
\end{aligned}$$

since $c_{2j}(r,0)=0$ for all $r \geq 0$ and $j \in \{1,2\}$. □

Lemma 3.7. *For $n \geq 1$, the n-step transition probability matrix of the uniformized chain is*

$$\widehat{\boldsymbol{P}}^n = \begin{bmatrix} 1-\sum_{k=0}^{n-1}(-1)^k\lambda^k & \sum_{k=0}^{n-1}(-1)^k\lambda^k \\ \\ 1-\sum_{k=0}^{n}(-1)^k\lambda^k & \sum_{k=0}^{n}(-1)^k\lambda^k \end{bmatrix}.$$

Proof. This result can be easily proved using an induction argument. □

Lemma 3.8. *For $n \geq 1$, $c_{2\cdot}(n,n) = \lambda\sum_{k=0}^{n-1}(1-\lambda)^k$.*

Proof. By Lemmas 3.6 and 3.7, it is seen that

$$\begin{aligned}
c_{21}(n,n) &= \lambda\left(\widehat{\boldsymbol{P}}^{n-1}\right)_{11}+\lambda(1-\lambda)\left(\widehat{\boldsymbol{P}}^{n-2}\right)_{11}+\lambda(1-\lambda)^2\left(\widehat{\boldsymbol{P}}^{n-3}\right)_{11}+\cdots \\
&\qquad \cdots+\lambda(1-\lambda)^{n-2}\left(\widehat{\boldsymbol{P}}^{1}\right)_{11}+\lambda(1-\lambda)^{n-1}(\boldsymbol{I})_{11}+(1-\lambda)^n c_{21}(0,0), \\
c_{22}(n,n) &= \lambda\left(\widehat{\boldsymbol{P}}^{n-1}\right)_{12}+\lambda(1-\lambda)\left(\widehat{\boldsymbol{P}}^{n-2}\right)_{12}+\lambda(1-\lambda)^2\left(\widehat{\boldsymbol{P}}^{n-3}\right)_{12}+\cdots \\
&\qquad \cdots+\lambda(1-\lambda)^{n-2}\left(\widehat{\boldsymbol{P}}^{1}\right)_{12}+\lambda(1-\lambda)^{n-1}(\boldsymbol{I})_{12}+(1-\lambda)^n c_{22}(0,0).
\end{aligned}$$

Since $\boldsymbol{C}(0,0) = \mathbf{0}$ and $\left(\widehat{\boldsymbol{P}}^{n-1}\right)_{11} + \left(\widehat{\boldsymbol{P}}^{n-1}\right)_{12} = 1$ for all $n \geq 0$,

$$\begin{aligned} c_{2\cdot}(n,n) &= \lambda\left[\left(\widehat{\boldsymbol{P}}^{n-1}\right)_{11} + \left(\widehat{\boldsymbol{P}}^{n-1}\right)_{12}\right] + \lambda(1-\lambda)\left[\left(\widehat{\boldsymbol{P}}^{n-2}\right)_{11} + \left(\widehat{\boldsymbol{P}}^{n-2}\right)_{12}\right] \\ &\quad +\lambda(1-\lambda)^2\left[\left(\widehat{\boldsymbol{P}}^{n-3}\right)_{11} + \left(\widehat{\boldsymbol{P}}^{n-3}\right)_{12}\right] + \cdots + \lambda(1-\lambda)^{n-1}\left[(\boldsymbol{I})_{11} + (\boldsymbol{I})_{12}\right] \\ &\quad +(1-\lambda)^n c_{2\cdot}(0,0) \\ &= \lambda + \lambda(1-\lambda) + \lambda(1-\lambda)^2 + \cdots + \lambda(1-\lambda)^{n-1} \\ &= \lambda\sum_{k=0}^{n-1}(1-\lambda)^k. \end{aligned}$$

□

Lemma 3.9. *For all $n \geq 0$, $c_{2\cdot}(n,n) \leq c_{1\cdot}(n,n)$.*

Proof. By Lemma 3.9,

$$c_{2\cdot}(n,n) = \lambda\sum_{k=0}^{n-1}(1-\lambda)^k \leq 1 = c_{1\cdot}(n,n).$$

□

Lemma 3.10. *For all $n \geq 1$, $c_{2\cdot}(n,n-1) \leq c_{1\cdot}(n,n-1)$.*

Proof. For $n \geq 1$ and $j \in \{1,2\}$, Lemma 3.6 gives the following:

$$\begin{aligned} c_{2j}(n,n-1) &= \lambda c_{11}(n-1,n-2) + (1-\lambda)c_{2j}(n-1,n-2) \\ &= \lambda c_{2j}(n-2,n-2) + (1-\lambda)c_{2j}(n-1,n-2) \\ &= \lambda c_{2j}(n-2,n-2) + \lambda(1-\lambda)c_{2j}(n-3,n-3) \\ &\quad +\lambda(1-\lambda)^2 c_{2j}(n-4,n-4) + \cdots + \lambda(1-\lambda)^{n-3}c_{2j}(1,1) \\ &\quad +\lambda(1-\lambda)^{n-2}c_{2j}(0,0) + (1-\lambda)^{n-1}c_{2j}(1,0). \end{aligned}$$

The result is shown as follows:

$$
\begin{aligned}
c_{2.}(n, n-1) &= \lambda c_{2.}(n-2, n-2) + (1-\lambda) c_{2.}(n-1, n-2) \\
&= \lambda c_{2.}(n-2, n-2) + \lambda(1-\lambda) c_{2.}(n-3, n-3) + \lambda(1-\lambda)^2 c_{2.}(n-4, n-4) + \cdots \\
&\qquad \cdots + \lambda(1-\lambda)^{n-3} c_{2.}(1,1) + \lambda(1-\lambda)^{n-2} c_{2.}(0,0) + (1-\lambda)^{n-1} c_{2.}(1,0) \\
&= \lambda \left[\lambda \sum_{k=0}^{n-3} (1-\lambda)^k \right] + \lambda(1-\lambda) \left[\lambda \sum_{k=0}^{n-4} (1-\lambda)^k \right] + \lambda(1-\lambda)^2 \left[\lambda \sum_{k=0}^{n-5} (1-\lambda)^k \right] + \cdots \\
&\qquad \cdots + \lambda(1-\lambda)^{n-3} \\
&= \lambda^2 \left[\sum_{k=0}^{n-3} (1-\lambda)^k + (1-\lambda) \sum_{k=0}^{n-4} (1-\lambda)^k + (1-\lambda)^2 \sum_{k=0}^{n-5} (1-\lambda)^k + \cdots + (1-\lambda)^{n-3} \right] \\
&\leq \lambda^2 \left[\frac{1}{\lambda} + (1-\lambda)\frac{1}{\lambda} + (1-\lambda)^2 \frac{1}{\lambda} + \cdots + (1-\lambda)^{n-3} \frac{1}{\lambda} \right] \\
&= \lambda \left[1 + (1-\lambda) + (1-\lambda)^2 + \cdots + (1-\lambda)^{n-3} \right] \\
&= \lambda \sum_{k=0}^{n-3} (1-\lambda)^k \\
&\leq \lambda \sum_{k=0}^{n-2} (1-\lambda)^k \\
&= c_{2.}(n-1, n-1) \\
&= c_{1.}(n, n-1).
\end{aligned}
$$

□

Lemma 3.11. *For all $n \geq 0$ and $k \leq n$, $c_{2.}(n,k) \leq c_{2.}(n-1,k)$.*

Proof. By Lemma 3.6,

$$\begin{aligned} c_{2.}(n,k) &= \underbrace{\lambda^{n-k}c_{2.}(2k-n,2k-n)}_{1} + \underbrace{\lambda^{n-k-1}(1-\lambda)c_{2.}(2k-n+1,2k-n))}_{2} \\ &\quad + \underbrace{\lambda^{n-k-2}(1-\lambda)c_{2.}(2k-n+3,2k-n+1)}_{3} + \cdots + \underbrace{\lambda^{2}(1-\lambda)c_{2.}(n-5,k-3)}_{\text{n-k-1}} \\ &\quad + \underbrace{\lambda(1-\lambda)c_{2.}(n-3,k-2) + (1-\lambda)c_{2.}(n-1,k-1)}_{\text{n-k}} \end{aligned}$$

$$\begin{aligned} c_{2.}(n-1,k) &= \underbrace{\lambda^{n-k-1}c_{2.}(2k-n+1,2k-n+1)}_{1} + \underbrace{\lambda^{n-k-2}(1-\lambda)c_{2.}(2k-n+2,2k-n+1)}_{2} \\ &\quad + \underbrace{\lambda^{n-k-3}(1-\lambda)c_{2.}(2k-n+4,2k-n+2)}_{3} + \cdots + \underbrace{\lambda^{2}(1-\lambda)c_{2.}(n-6,k-3)}_{\text{n-k-2}} \\ &\quad + \underbrace{\lambda(1-\lambda)c_{2.}(n-4,k-2) + (1-\lambda)c_{2.}(n-2,k-1)}_{\text{n-k-1}} \end{aligned}$$

A one-to-one comparison of the terms $\{3,4,\ldots,n-k\}$ of $c_{2.}(n,k)$ with the corresponding terms $\{2,3,\ldots,n-k-1\}$ of $c_{2.}(n-1,k)$ indicates that each pair is of the form $\alpha c_{2.}(w,v)$ and $\alpha c_{2.}(w-1,v)$, respectively, where $w \leq n$, $v \geq w/2$, and $\alpha \in \mathbb{R}_{+}$. Therefore, the paired terms in each set have the same relationship as the original expressions that generated them, for which $(\alpha,w,v) = (1,n,k)$. It follows that if the recursion is true for $(\alpha,w,v) = (1,n,k)$, it holds for all the generated terms by extension. Now proving the result is reduced to just showing that the sum of the first and second terms of $c_{2.}(n,k)$ are less than or equal to the first term of $c_{2.}(n-1,k)$. To prove this is true, first note a result from [91] that states $c_{ij}(n,k) \leq c_{ij}(n,k+1)$ for all $i,j \in \{1,2\}$, $n \geq 0$, and $k \leq n$. Now observe that the sum of the first and second terms of $c_{2.}(n,k)$ can be bounded above as follows:

$$\begin{aligned} &\lambda^{n-k-1}(1-\lambda)c_{2.}(2k-n+1,2k-n) \\ &\qquad +\lambda^{n-k}c_{2.}(2k-n,2k-n) \quad \leq \quad \lambda^{n-k-1}(1-\lambda)c_{2.}(2k-n+1,2k-n+1) \\ &\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad +\lambda^{n-k}c_{2.}(2k-n,2k-n) \\ &\qquad\qquad\qquad\qquad = \quad \lambda^{n-k}(1-\lambda)\sum_{r=0}^{2k-n}(1-\lambda)^{r} + \lambda^{n-k+1}\sum_{r=0}^{2k-n-1}(1-\lambda)^{r} \\ &\qquad\qquad\qquad\qquad = \quad \lambda^{n-k} - \lambda^{n-k+1}\sum_{r=0}^{2k-n}(1-\lambda)^{r} + \lambda^{n-k+1}\sum_{r=0}^{2k-n-1}(1-\lambda)^{r} \\ &\qquad\qquad\qquad\qquad = \quad \lambda^{n-k} - \lambda^{n-k+1}(1-\lambda)^{2k-n} \\ &\qquad\qquad\qquad\qquad \leq \quad \lambda^{n-k}, \end{aligned}$$

and the first term of $c_{2.}(n-1,k)$ is bounded below as follows:

$$\begin{aligned}\lambda^{n-k-1}c_{2.}(2k-n+1,2k-n+1) &= \lambda^{n-k}\sum_{r=0}^{2k-n}(1-\lambda)^r \\ &= \lambda^{n-k}\left(1+\sum_{r=1}^{2k-n}(1-\lambda)^r\right) \\ &\geq \lambda^{n-k}.\end{aligned}$$

□

Lemma 3.12. *For all $n \geq 0$ and $k \leq n$, $c_{2.}(n,k) \leq c_{1.}(n,k)$.*

Proof. In the case where $k \leq n/2$, $c_{2.}(n,k) = 0 \leq c_{1.}(n,k)$ by Lemma 3.6. For $k \geq n/2$, Lemmas 3.9 and 3.10 prove that the result holds for all $n \geq 0$ and $k \in \{n-1, n\}$. Now assume that $c_{2.}(n,k) \leq c_{1.}(n,k)$ for all $n \geq 0$ and $k \in \{n-(w-1), n-(w-2), \ldots, n\}$ for some integer $w \geq n/2$. Then by Lemma 3.6,

$$\begin{aligned}c_{2.}(n,n-w) &= \underbrace{\lambda^w c_{2.}(n-2w,n-2w)}_{1} + \underbrace{\lambda^{w-1}(1-\lambda)c_{2.}(n-2w+1,n-2w)}_{2} \\ &\quad +\lambda^{w-2}(1-\lambda)c_{2.}(n-2w+3,n-2w+1) \\ &\quad +\lambda^{w-3}(1-\lambda)c_{2.}(n-2w+5,n-2w+3) \\ &\quad +\cdots+(1-\lambda)c_{2.}(n-1,n-w-1),\end{aligned} \tag{3.15}$$

and

$$\begin{aligned}c_{1.}(n,n-w) &= c_{2.}(n-1,n-w) \\ &= \underbrace{\lambda^{w-1}c_{2.}(n-2w+1,n-2w+1)}_{1} + \underbrace{\lambda^{w-2}(1-\lambda)c_{2.}(n-2w+2,n-2w+1)}_{2} \\ &\quad +\lambda^{w-3}(1-\lambda)c_{2.}(n-2w+4,n-2w+2) \\ &\quad +\cdots+(1-\lambda)c_{2.}(n-2,n-w-1).\end{aligned} \tag{3.16}$$

Equations (3.15) and (3.16) are compared on a term-by-term basis. For the first terms,

$$\begin{aligned}\lambda^w c_{2.}(n-2w,n-2w) &= \lambda^{w+1}\sum_{k=0}^{n-2w-1}(1-\lambda)^k \\ &\leq \lambda^w\sum_{k=0}^{n-2w}(1-\lambda)^k \\ &= \lambda^{w-1}c_{2.}(n-2w+1,n-2w+1).\end{aligned}$$

Next the second term of (3.15) is shown to be less than the corresponding second term of equation (3.16) as follows:

$$\begin{aligned}
\lambda^{w-2}(1-\lambda)c_{2.}(n-2w+1,n-2w) &= \lambda^{w-2}(1-\lambda)[\lambda c_{2.}(n-2w+1,n-2w) \\
&\qquad +(1-\lambda)c_{2.}(n-2w+1,n-2w)] \\
&\leq \lambda^{w-2}(1-\lambda)[\lambda c_{2.}(n-2w,n-2w) \\
&\qquad +(1-\lambda)c_{2.}(n-2w+1,n-2w)] \quad \text{(Lemma 3.11)} \\
&= \lambda^{w-2}(1-\lambda)[\lambda c_{1.}(n-2w+1,n-2w) \\
&\qquad +(1-\lambda)c_{2.}(n-2w+1,n-2w)] \\
&= \lambda^{w-2}(1-\lambda)c_{2.}(n-2w+2,n-2w+1).
\end{aligned}$$

The preceding argument can be made for the relationship between each corresponding pair of terms in $\{3,4,\ldots,w-1\}$ so that

$$\begin{aligned}
c_{1.}(n,n-w) &= K_1+(1-\lambda)c_{2.}(n-2,n-w-1) \\
c_{2.}(n,n-w) &= K_2+(1-\lambda)c_{2.}(n-1,n-w-1),
\end{aligned}$$

where $K_1 > K_2 > 0$. Now by Lemma 3.11, $c_{2.}(n-1,n-w-1) \leq c_{2.}(n-2,n-2-1)$ so that $c_{2.}(n,n-w) \leq c_{1.}(n,n-w)$, and the induction holds for w. □

Lemma 3.13. *Let $\pi(a) = [1-a,a]$ for $a \in [0,1]$. Then for fixed $x \in [0,x_c]$, $H(x,\pi(a))$ is monotone increasing in a, and $\tau(x,\pi(a))$ is monotone decreasing in a.*

Proof. For $a \in [0,1]$ and $(x,\pi(a)) \in B$,

$$\begin{aligned}
H(x,\pi(a)) &= 1-\sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{S}} W_{ij}(x,\delta)\,\pi^{(i)}(a) \\
&= 1-(1-a)\sum_{j\in\mathcal{S}} W_{1j}(x,\delta) - a\sum_{j\in\mathcal{S}} W_{2j} \\
&= a\left[\sum_{j\in\mathcal{S}} W_{1j}(x,\delta) - W_{2j}(x,\delta)\right].
\end{aligned}$$

Therefore, it is sufficient to show that

$$\sum_{j\in\mathcal{S}} W_{1j}(x,\delta) - W_{2j}(x,\delta) \geq 0 \tag{3.17}$$

for all $x \in [0, x_c]$ and $\delta \geq 0$. By Lemma 3.12, $c_{1\cdot}(n,k) \geq c_{2\cdot}(n,k)$ for all $n \geq 0$ and $k \leq n$. Therefore, follows immediately from (3.14) that (3.17) holds, and $H(x,\pi)$ is monotone increasing in a. Likewise,

$$\begin{aligned}
\tau(x,\pi(a)) &= \sum_{i\in\mathcal{S}} \left[\sum_{j\in\mathcal{S}} \int_0^\delta W_{ij}(x_c - x, \delta)\, \mathrm{d}\delta \right] \pi^{(i)}(a) \\
&= \int_0^\delta \left[(1-a) \sum_{j\in\mathcal{S}} W_{1j}(x,\delta) + a \sum_{j\in\mathcal{S}} W_{2j}(x,\delta) \right] \mathrm{d}\delta \\
&= \int_0^\delta \left\{ 1 - a \sum_{j\in\mathcal{S}} [W_{1j}(x,\delta) - W_{2j}(x,\delta)] \right\} \mathrm{d}\delta,
\end{aligned}$$

where it follows that the integrand and, therefore, $\tau(x,\pi(a))$ are decreasing in a. □

Using Lemmas 3.4 and 3.13, the existence and optimality of a threshold policy are shown with respect to the parameterized belief state.

Theorem 3.5. *If for fixed $x \in [0, x_c]$, it is optimal to replace for some $\pi(a) \in \Pi$, then it is optimal to replace for all $\pi(a')$, $a' \in [a, 1]$.*

Proof. By Lemma 3.4, $V(x,\pi(a)) \geq \underline{V}_0(x,\pi(a)) = (C_0 - c_1) + (c_1 + c_2)H(x,\pi(a)) - \gamma\tau(x,\pi(a))$. Since $\underline{V}_0(x,\pi(a))$ is increasing in a by Lemma 3.13, it remains optimal to replace for all $a' \in [a,1]$. □

In the special case of a two-state environment, Theorem 3.5 proves the existence of a threshold-type policy with respect to the environment state for a fixed level of degradation. The replacement threshold is equivalent to the probability that the environment is in the highest-degradation state, and as this probability increases, it remains optimal to preventively replace. This result is especially useful in a wind turbine application where it may be the case that information for the environment state is suddenly updated before another degradation observation can be acquired. In such an instance, a wind farm operator may elect to conduct a preventive replacement if the probability that the environment is in the highest-degradation state now exceeds the threshold associated with the last observed degradation level.

3.3 REPLACEMENT WITH STOCHASTIC DOWNTIME COSTS

In this section, an alternate replacement model is considered that incorporates stochastic downtime costs during replacements. The observability assumptions for the environment and degradation, as well as the underlying model dynamics, are unchanged from Section 3.2; however, preventive and reactive replacements are no longer assumed to be completed instantaneously at a fixed cost. Instead, the replacements require a fixed, deterministic time period for completion during which downtime cost accrues as a function of the random environment state.

As before, degradation is observed instantaneously at each decision epoch with cost c_0 ($c_0 > 0$), and if a decision is made to replace the component or it fails, the respective preventive or reactive replacement starts immediately at the beginning of the period or immediately upon failure. However, now it is assumed that the preventive and reactive replacements require deterministic time periods of $\delta_1 > 0$ and $\delta_2 > 0$, respectively, for completion. The total cost of a replacement now consists of a fixed capital cost, denoted c_1 ($0 < c_0 < c_1 < \infty$) and a downtime cost that is a function of the random environment. Let d_i denote the downtime cost per unit time when the environment is in state i, $i \in \mathcal{S}$, and $\boldsymbol{d} = [d_1, d_2, \ldots, d_\ell]$. Define the cumulative downtime cost accrued during a replacement up to time t by $X_d(t)$, where

$$X_d(t) = \int_0^t d_i \, \mathbb{I}(Z(u) = i) \mathrm{d}u.$$

Let $W_{ij}^{(d)}(t) \equiv \mathbb{P}(X_d(t) \leq x, Z(t) = j | Z(0) = i)$, for $i, j \in \mathcal{S}$ and ≥ 0, and

$$p_d^{(t)}(u, j|i) \equiv \left. \frac{\partial W_{ij}^{(d)}(x,t)}{\partial x} \right|_{(u,t)}.$$

The expected total downtime cost accrued during a replacement of time t, given it starts with the environment in state i is denoted $C_i^{(t)}$, where

$$\begin{aligned} C_i^{(t)} &\equiv \mathbb{E}[X_d(t) | X_d(0) = 0, Z(0) = i] \\ &= \sum_{j \in \mathcal{S}} \int_{d_1 \, t}^{d_\ell \, t} u \, p_d^{(t)}(u, j|i) \, \mathrm{d}u + t \, d_i \, \exp(q_{ii} \, t), \end{aligned}$$

and it follows that the total expected downtime cost of a preventive replacement in $(x, \pi) \in B$, denoted $C^{(p)}(\pi)$, is

$$\begin{aligned} C^{(p)}(\pi) &= \mathbb{E}[X_d(\delta_1) | X_d(0) = 0, \pi] \\ &= \sum_{i \in \mathcal{S}} C_i^{(\delta_1)} \pi^{(i)}. \end{aligned}$$

Let $L(x,\pi)$ be the random remaining lifetime of the component given $(x,\pi) \in B$, and define $\boldsymbol{T}^*(u,\pi) \equiv [T_1^*(u,\pi), T_2^*(u,\pi), \ldots, T_\ell^*(u,\pi)]$, where

$$\begin{aligned} T_i^*(u,\pi) &\equiv \mathbb{P}(Z(u) = i | L(x,\pi) = u, \pi) \\ &= \frac{\sum_{j\in\mathcal{S}} p^{(u)}(x_c - x, i|j)\pi^{(j)}}{\sum_{j\in\mathcal{S}}\sum_{k\in\mathcal{S}} p^{(u)}(x_c - x, j|k)\pi^{(k)}}. \end{aligned}$$

Then the expected total downtime cost of a reactive failure in the next period given $(x,\pi) \in B$ is denoted $C^{(r)}(x,\pi)$, where

$$\begin{aligned} C^{(r)}(x,\pi) &= \mathbb{E}[X_d(L+\delta_2) | L \le \delta, X_d(L) = 0, X(L) - X(0) = x_c - x, \pi] \\ &= \frac{1}{\mathbb{P}(L(x,\pi) \le \delta)} \int_{(x_c-x)/r_\ell \wedge \delta}^{(x_c-x)/r_1 \wedge \delta} \sum_{i\in\mathcal{S}} C_i^{(\delta_2)} T_i^*(u,\pi) \mathrm{d}\mathbb{P}(L(x,\pi) \le u). \end{aligned}$$

All components are assumed to begin operation in $(0,\pi_s)$, and the objective is to minimize the long-run average replacement cost per unit time, where

$$\gamma = \inf_{a\in\mathcal{P}} \mathbb{E}_a \left\{ \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} c_0 + c_1 \, \mathbb{I}\{a(X_n,\pi_n) = 1\} \right.$$
$$\left. + [c_1 + X_d(L+\delta_2) - X_d(L)] \, \mathbb{I}\{a(X_n,\pi_n) = 0, L(x,\pi) \le \delta\} \right\}$$

Let $V(x,\pi)$ be the minimum relative cost per unit time given a component starts operation in $(x,\pi) \in B$, and $V_0(x,\pi)$ and $V_1(x,\pi)$ be the relative costs if either no action or preventive replacement, respectively, are taken in $(x,\pi) \in B$. The optimality equations are as follows:

$$V(x,\pi) = \min\{V_1(x,\pi), V_0(x,\pi)\}, \quad (x,\pi) \in B, \tag{3.18}$$

where

$$\begin{aligned} V_1(x,\pi) &= c_0 + c_1 + C^{(p)}(\pi) + V(0,\pi_s) - \delta_1\gamma, \\ V_0(x,\pi) &= c_0 + \left[c_1 + C^{(r)}(x,\pi) - \delta_2\gamma + V(0,\pi_s)\right] \int_0^\infty \mathbb{I}^+(x,u) k(u,\pi) \mathrm{d}u \\ &\quad + \int_0^\infty \mathbb{I}^-(x,u) V_\pi(x,u) k(u,\pi) \mathrm{d}u - \gamma\tau(x,\pi). \end{aligned}$$

Obtaining structural results for (3.18) is considerably more difficult than for (3.2). Nonetheless, some basic results for the optimal policies can be obtained under relatively restrictive conditions. The following lemma establishes a lower bound for the optimal policy cost:

Lemma 3.14. *The average cost of an optimal policy can be bounded below as follows:*

$$\gamma > \frac{c_0}{\delta + \delta_2}.$$

Proof. Set $V(0, \pi_s) = 0$, and note that the optimal action immediately following replacement is to do nothing so that $V(0, \pi_s) = V_0(0, \pi_s) = 0$. Substituting equation (3.19) for $V_0(0, \pi_s)$ and solving for γ yields

$$\gamma = \frac{c_0 + \left[c_1 + C^{(r)}(0, \pi_s)\right] H(0, \pi_s) + \int_0^\infty \mathbb{I}^-(0, u) V_\pi(0, u) k(u, \pi)\, \mathrm{d}u}{\tau(0, \pi_s) + \delta_2 H(0, \pi_s)},$$

which gives

$$\gamma > \frac{c_0 + \left[c_1 + C^{(r)}(0, \pi_s)\right] H(0, \pi_s)}{\tau(0, \pi_s) + \delta_2 H(0, \pi_s)} > \frac{c_0}{\delta + \delta_2}.$$

□

Unlike the model of Section 3.2, it is not necessarily optimal to delay preventive replacement when the component will survive the next period w.p. 1. In fact, one can consider situations where initiating a preventive replacement immediately could take advantage of the current environment state having a relatively low downtime cost rate. However, when failure is imminent in the next period, preventive replacement is always optimal under certain conditions.

Lemma 3.15. *If the component fails w.p. 1 in the next period for $(x, \pi) \in B$, then a sufficient condition for the optimality of preventive replacement is*

$$C^{(r)}(x, \pi) - C^{(p)}(\pi) \geq c_0 \left(\frac{\delta_2 - \delta_1}{\delta_2 + \delta} \right).$$

Proof. For $(x, \pi) \in B$, preventive replacement is optimal if $V_1(x, \pi) \leq V_0(x, \pi)$. Assuming failure in the next decision period w.p. 1, $V_1(x, \pi) \leq V_0(x, \pi)$, which implies

$$c_0 + c_1 + C^{(p)}(\pi) + V(0, \pi_s) - \delta_1 \gamma \leq c_0 + c_1 + C^{(r)}(x, \pi) - \delta_2 \gamma + V(0, \pi_s) - \gamma \tau(x, \pi). \tag{3.19}$$

Rearranging (3.19) and applying Lemma 3.14 gives

$$C^{(r)}(x, \pi) - C^{(p)}(\pi) \geq \gamma[\delta_2 - \delta_1 + \tau(x, \pi)] > c_0 \left(\frac{\delta_2 - \delta_1}{\delta_2 + \delta} \right).$$

□

The next three lemmas are required to prove the existence of a replacement threshold with respect to x for fixed π. The first lemma establishes a lower bound for $V_0(x, \pi)$ and the two following lemmas establish some properties of $C^{(r)}(x, \pi)$.

Lemma 3.16. *For each* $\pi \in \Pi$,

$$\begin{aligned} V_0(x,\pi) \geq \underline{V}_0(x,\pi) \quad &\equiv \quad c_0 + \left[c_1 + C^{(r)}(x,\pi)\right] H(x,\pi) \\ &\quad -\gamma\left[\delta_2 H(x,\pi) + \tau(x,\pi)\right] + V(0,\pi_s). \end{aligned} \tag{3.20}$$

Proof. For $(x,\pi) \in B$ and $u \geq 0$, $V_\pi(x,u) \geq V(0,\pi_s)$. Therefore,

$$\begin{aligned} V_0(x,\pi) &= c_0 + \left[c_1 + C^{(r)}(x,\pi) - \delta_2\gamma + V(0,\pi_s)\right] \int_0^\infty \mathbb{I}^+(x,u)k(u,\pi)\mathrm{d}u \\ &\qquad + \int_0^\infty \mathbb{I}^-(x,u)V_\pi(x,u)k(u,\pi)\mathrm{d}u - \gamma\tau(x,\pi) \\ &\geq c_0 + \left[c_1 + C^{(r)}(x,\pi)\right] H(x,\pi) - \gamma\delta_2 H(x,\pi) \\ &\qquad + V(0,\pi_s)H(x,\pi) + V(0,\pi_s)[1 - H(x,\pi)] - \gamma\tau(x,\pi) \\ &= c_0 + \left[c_1 + C^{(r)}(x,\pi)\right] H(x,\pi) - \gamma\left[\delta_2 H(x,\pi) + \tau(x,\pi)\right] + V(0,\pi_s). \end{aligned}$$

□

Lemma 3.16 establishes a lower bound for $V_0(x,\pi)$. The lower bound (3.20) includes a term for the expected downtime cost of a reactive replacement, which is not necessarily monotone in x. Therefore, in contrast to the simple properties of (3.11), the behavior of (3.20) is considerably more difficult to characterize.

Lemma 3.17. *For all* $\pi \in \Pi$,

$$C^{(r)}(x_c,\pi) = \sum_{i\in\mathcal{S}} C_i^{(\delta_2)}\,\pi^{(i)}.$$

Proof. When $x = x_c$, the remaining lifetime is zero w.p. 1, and $\mathrm{d}\mathbb{P}(L(x_c,\pi) \leq u) = \omega(u)$, where $\omega(u)$ is the Dirac delta function. Therefore,

$$\begin{aligned} C^{(r)}(x_c,\pi) &= \int_{\mathbb{R}_+} \sum_{i\in\mathcal{S}} C_i^{(\omega_2)} T_i^*(u,\pi)\omega(u) \\ &= \sum_{i\in\mathcal{S}} C_i^{(\omega_2)} T_i^*(0,\pi) \\ &= \sum_{i\in\mathcal{S}} C_i^{(\omega_2)} \mathbb{P}(Z(u) = i | L(x,\pi) = 0, \pi) \\ &= \sum_{i\in\mathcal{S}} C_i^{(\omega_2)}\,\pi^{(i)}. \end{aligned}$$

□

Lemma 3.18. *For $\epsilon > 0$ and $(x, \pi) \in B$ such that $x_c - x < \epsilon$,*

$$C^{(r)}(x, \pi) \geq \sum_{i \in \mathcal{S}} C_i^{(\delta_2)} \pi^{(i)} \exp(-q_{ii}\epsilon / r_i).$$

Proof. Pick $\epsilon < \delta r_1$ and x such that $x_c - x < \epsilon$. Then the remaining lifetime of the component is $(x_c - x)/r_1 < \epsilon/r_1 < \delta$ so that $\mathbb{P}(L(x, \pi) \leq \delta) = 1$. Then

$$\begin{aligned}
C^{(r)}(x, \pi) &= \frac{1}{\mathbb{P}(L(x, \pi) \leq \delta)} \int_{\epsilon/r_\ell}^{\epsilon/r_1} \sum_{j \in \mathcal{S}} C_j^{(\delta_2)} T_j^*(u, \pi) \mathrm{d}\mathbb{P}(L(x, \pi) \leq u) \\
&= \int_{\epsilon/r_\ell}^{\epsilon/r_1} \sum_{j \in \mathcal{S}} C_j^{(\delta_2)} \mathbb{P}(Z(u) = j | L(x, \pi) = u, \pi) \mathrm{d}\mathbb{P}(L(x, \pi) \leq u) \\
&= \int_{\epsilon/r_\ell}^{\epsilon/r_1} \sum_{j \in \mathcal{S}} C_j^{(\delta_2)} \mathrm{d}\mathbb{P}(Z(u) = j, L(x, \pi) \leq u | \pi) \\
&\geq \int_{\epsilon/r_\ell}^{\epsilon/r_1} \mathbb{I}(u \in \{\epsilon/r_i : i \in \mathcal{S}\}) \sum_{j \in \mathcal{S}} C_j^{(\delta_2)} \mathrm{d}\mathbb{P}(Z(u) = j, L(x, \pi) \leq u | \pi) \\
&= \sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{S}} C_j^{(\delta_2)} \mathrm{d}\mathbb{P}(Z(u) = j, L(x, \pi) \leq u | \pi)|_{u = \epsilon/r_i} \\
&\geq \sum_{i \in \mathcal{S}} C_i^{(\delta_2)} \mathrm{d}\mathbb{P}(Z(u) = i, L(x, \pi) \leq u | \pi)|_{u = \epsilon/r_i} \\
&= \sum_{i \in \mathcal{S}} \sum_{k \in \mathcal{S}} C_i^{(\delta_2)} \mathrm{d}\mathbb{P}(Z(u) = i, L(x, \pi) \leq u | Z(0) = k)|_{u = \epsilon/r_i} \pi^{(k)} \\
&= \sum_{i \in \mathcal{S}} \sum_{k \in \mathcal{S}} C_i^{(\delta_2)} \mathbb{P}(Z(\epsilon/r_i) = i, X(\epsilon/r_i) = \epsilon | Z(0) = k) \pi^{(k)} \\
&\geq \sum_{i \in \mathcal{S}} C_i^{(\delta_2)} \mathbb{P}(Z(\epsilon/r_i) = i, X(\epsilon/r_i) = \epsilon | Z(0) = i) \pi^{(i)} \\
&= \sum_{i \in \mathcal{S}} C_i^{(\delta_2)} \pi^{(i)} \exp(-q_{ii}\epsilon/r_i),
\end{aligned}$$

where the last equality follows from the fact that $\mathbb{P}(X(t) = r_i t, Z(t) = j | Z(0) = i) = \mathbb{I}(i = j) \exp(-q_{ii} t)$ for all $t \geq 0$ and $i \in \mathcal{S}$. □

A structural result for the optimal policy can now be inferred from Lemmas 3.14-3.18. Unlike in Theorem 3.4, a replacement threshold with respect to x is not guaranteed to exist for a given $\pi \in \Pi$ unless a necessary condition holds that involves the inspection and replacement times, as well as the downtime costs. Furthermore, even provided this necessary condition holds, the optimality of preventive replacement can only be shown for all x sufficiently close to the failure threshold.

Theorem 3.6. *If for some* $\pi \in \Pi$,

$$c_0 \left(\frac{\delta_2 - \delta_1}{\delta + \delta_2} \right) < \sum_{i \in \mathcal{S}} \left(C_i^{(\delta_2)} - C_i^{(\delta_1)} \right) \pi^{(i)}, \tag{3.21}$$

then there exists an $\epsilon > 0$ *such that it is optimal to preventively replace for all* (x, π) *such that* $x \in (x_c - \epsilon, x_c]$.

Proof. Using Lemma 3.16, a sufficient condition for the optimality of preventive replacement at $(x, \pi) \in B$ is $\underline{V}_0(x, \pi) \geq V_1(x, \pi)$, which implies

$$c_0 + c_1 + C^{(r)}(x, \pi) - \delta_2 \gamma + V(0, \pi_s) \geq c_0 + c_1 + C^{(p)}(\pi) + V(0, \pi_s) - \delta_1 \gamma. \tag{3.22}$$

Rearranging (3.22) and applying Lemma 3.16 gives that

$$C^{(r)}(x, \pi) - \sum_{i \in \mathcal{S}} C_i^{(\delta_1)} \pi^{(i)} \geq \gamma(\delta_2 - \delta_1) > c_0 \left(\frac{\delta_2 - \delta_1}{\delta + \delta_2} \right),$$

and further rearrangement yeilds

$$C^{(r)}(x, \pi) \geq \sum_{i \in \mathcal{S}} C_i^{(\delta_1)} \pi^{(i)} + c_0 \left(\frac{\delta_2 - \delta_1}{\delta + \delta_2} \right) \equiv K. \tag{3.23}$$

Evaluating (3.23) at (x_c, π) by applying Lemma 3.17 and rearranging terms gives condition (3.21). Now assuming (3.21) is satisfied at x_c, then it can also be shown to be satisfied for $x \in (x_c - \epsilon, x_c]$ for some $\epsilon > 0$ as follows. Assume condition (3.21) is satisfied and rearrange terms so that

$$\sum_{i \in \mathcal{S}} C_i^{(\delta_2)} \pi_i > \sum_{i \in \mathcal{S}} C_i^{(\delta_1)} \pi_i + c_0 \left(\frac{\delta_2 - \delta_1}{\delta + \delta_2} \right) = K,$$

which implies

$$\sum_{i \in \mathcal{S}} C_i^{(\delta_2)} \pi_i - K \equiv K_1 > 0. \tag{3.24}$$

By Lemma 3.18, for $\epsilon > 0$ and all x such that $x_c - x < \epsilon > 0$,

$$C^{(r)}(x, \pi) \geq \sum_{i \in \mathcal{S}} C_i^{(\delta_2)} \pi^{(i)} \exp(-q_{ii} \epsilon / r_i).$$

Picking ϵ such that $1-\exp(-q_{ii}\epsilon/r_i) < K_1/\sum_{j\in\mathcal{S}} C_j^{(\delta_2)}$ for all $i \in \mathcal{S}$ gives

$$\begin{aligned} C^{(r)}(x,\pi) - K &\geq \sum_{i\in\mathcal{S}} C_i^{(\delta_2)}\pi^{(i)}\exp(-q_{ii}\epsilon/r_i) - K \\ &> \left(\sum_{i\in\mathcal{S}} C_i^{(\delta_2)}\pi^{(i)}\right)\left(1 - \frac{K_1}{\sum_{i\in\mathcal{S}} C_i^{(\delta_2)}}\right) - K \\ &= \sum_{i\in\mathcal{S}} C_i^{(\delta_2)}\pi^{(i)} - \left(\frac{\sum_{i\in\mathcal{S}} C_i^{(\delta_2)}\pi^{(i)}}{\sum_{i\in\mathcal{S}} C_i^{(\delta_2)}}\right) K_1 - K \\ &> 0, \end{aligned}$$

where the last inequality holds by (3.24). Since $C^{(r)}(x,\pi) - K > 0$ implies (3.23) holds, it is optimal to replace for all $x \in (x_c - \epsilon, x_c]$. □

That a replacement threshold does not necessarily exist for each environment belief state in Theorem 3.6 indicates that environment information is especially critical to consider when stochastic downtime costs are incurred during replacement operations. For example, a particular environment state could represent ideal wind conditions for electricity production, during which significant revenues are generated. In such an environment state, it may be optimal to continue operating the wind turbine, regardless of the component's degradation level, rather than shutting it down to perform a preventive replacement. Likewise, in an environment state that represents insufficient wind conditions, it may be optimal to preventively replace a slightly-degraded component to minimize the revenue lost during future periods of ideal wind conditions.

3.4 NUMERICAL SOLUTION TECHNIQUES

In this section, numerical solutions techniques are described for solving the optimality equations of the models in Sections 3.2 and 3.3. These procedures require discretization of the belief space and the use of either a policy iteration or linear programming (LP) approach. The key idea is to approximate the integrals in the optimality equations as Riemann sums and formulate the problem as a Markov decision process (MDP) model on a finite state space.

Let $\dot{\Pi} = \{\dot{\pi}_1, \dot{\pi}_2, \ldots, \dot{\pi}_{L_1}\}$ be a discretization of Π, where $\dot{\pi}_i \in \Pi$ and $L_1 \geq 1$. For $L_2 \geq 2$, let the discretization interval between the realizable levels of degradation on $\mathcal{X} \setminus \{0\}$ be $b \equiv$

$(x_c + (r_i - r_1)\delta)/L_2$ so that the ordered set $\dot{\mathcal{X}} = \{0, \dot{x}_1, \dot{x}_2, \ldots, \dot{x}_{L_2+1}\}$ is the discretization of $\mathcal{X}$, where $\dot{x}_i = r_1\delta + (i-1)b$, $i = 1, 2, \ldots, L_2 + 1$. The complete discretized belief space is denoted $\dot{B} \equiv (\dot{\Pi} \times \dot{\mathcal{X}}) \cup \{(0, \pi_s)\}$, where $L \equiv |\dot{B}| - 1 = L_1(L_2 + 2)$, and denote the ith discretized belief state in $\dot{B}$ by $\dot{b}_i \equiv (\dot{x}_i, \dot{\pi}_i)$, where $(\dot{x}_i, \dot{\pi}_i) \in \dot{\Pi} \times \dot{\mathcal{X}}$, $i = 1, 2, \ldots, L$ and $\dot{b}_0 \equiv (0, \pi_s)$. The set of discretized cumulative degradation levels attainable from a given $\dot{b}_i$ is $\bar{\mathcal{X}}_i = \{\bar{x}_{i1}, \bar{x}_{i2}, \ldots, \bar{x}_{i,C+1}\}$, where $\bar{x}_{ij} = x_i + r_1\delta + (j-1)b$ and $C \equiv |\bar{\mathcal{X}}_i| = \delta(r_i - r_1)/b$, $l = 1, 2 \ldots, L$, $j = 1, 2, \ldots, C + 1$. Estimating the transition probabilities between discretized belief states consists of three steps. In the first step for a given $\dot{b}_i$, compute $\bar{\Pi}_i = \{\bar{\pi}_{i1}, \bar{\pi}_{i2}, \ldots, \bar{\pi}_{i,C+1}\}$, where $\bar{\pi}_{ij} = \boldsymbol{T}(\bar{x}_{ij}, \dot{\pi}_i)$ is the updated belief state obtained from $\dot{b}_i$ by observing an increase in degradation of $\bar{x}_{ij}$ in the next decision epoch, $i = 1, 2, \ldots, L$, $j = 1, 2, \ldots, C + 1$. Next approximate each $\bar{\pi}_{ij}$ with a discretized $\pi^*_{ij} \in \dot{\Pi}$, where

$$\pi^*_{ij} = \min_{\dot{\pi} \in \dot{\Pi}} \|\bar{\pi}_{ij} - \dot{\pi}\|, \quad i = 1, 2, \ldots, L, \quad j = 1, 2, \ldots, C + 1,$$

and let $\Pi^*_i \equiv \{\pi^*_{i1}, \pi^*_{i2}, \ldots, \pi^*_{i,C+1}\}$. Finally, compute the row vector $\dot{P}_i \equiv \{\dot{p}_{i1}, \dot{p}_{i2}, \ldots, \dot{p}_{i,C+1}\}$, where $\dot{p}_{ij}$ is the transition probability between $\dot{b}_i$ and updated belief state $(\dot{x}_i + \bar{x}_{ij}, \pi^*_{ij})$. To compute $\dot{p}_{ij}$, first define $R(x; \pi_n, x_n) \equiv \mathbb{P}(X_{n+1} \leq x | x_n, \pi_n)$ as the probability that the cumulative degradation in epoch $n + 1$ does not exceed x, given $b^{(n)} = (x_n, \pi_n)$, where

$$\begin{aligned}
R(x; \pi_n, x_n) &= \mathbb{P}(X_{n+1} \leq x | x_n, \pi_n) \\
&= \sum_{i \in \mathcal{S}} \mathbb{P}(X_{n+1} \leq x | x_n, z_n = i) \pi^{(i)}_n \\
&= \sum_{i \in \mathcal{S}} \mathbb{P}(X_{n+1} \leq x - x_n | 0, z_n = i) \pi^{(i)}_n \\
&= \sum_{i \in \mathcal{S}} \mathbb{P}(X_1 \leq x - x_n | 0, z_0 = j) \pi^{(i)}_n \\
&= \sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{S}} \mathbb{P}(X_1 \leq x - x_n, z_1 = j | z_0 = i) \pi^{(i)}_n \\
&= \sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{S}} W_{ij}(x - x_n, \delta) \pi^{(i)}_n.
\end{aligned}$$

Then $\dot{p}_{ij}$ is computed as follows:

$$\dot{p}_{ij} = R_i(\bar{x}_{ij} + b/2 \wedge x_{i,C+1}) - R_i(x_{i1} \vee x_{ij} - b/2), \quad i = 0, 1, \ldots, L, \quad j = 1, 2, \ldots, L,$$

where $R_i(x) \equiv R(x; \dot{\pi}_i, \dot{x}_i)$, $i = 0, 1, \ldots, L$. The function $R_i(x)$ is approximated using numerical results by Sericola [91] and Bladt, *et al.* [14]. Computing $(\Pi^*_i, \bar{\mathcal{X}}_i, \dot{P}_i)$ for all $\dot{b}_i \in \dot{B}$ establishes the transition probabilities between all pairs of discretized belief states.

Now the optimality equations are stated on $\dot{B}$ for the model of Section 3.2. Let $\dot{V}(i) \equiv \dot{V}(\dot{b}_i)$ denote the relative cost, given that the component starts operation in $\dot{b}_i$, $i = 0, 1, \ldots, L$. Then

$$\dot{V}(i) = \min\{\dot{V}_0(i), \dot{V}_1(i)\}, \quad i = 0, 1, \ldots, L. \tag{3.25}$$

where $\dot{V}_0(i) \equiv \dot{V}_0(\dot{b}_i)$ and $\dot{V}_1(i) \equiv \dot{V}_1(\dot{b}_i)$ denote the relative costs of doing nothing and preventive replacement, respectively, in $\dot{b}_i$. The relative cost of preventive replacement in $\dot{b}_i$ is

$$\dot{V}_1(i) = c_0 + c_1 + \dot{V}(0), \quad i = 0, 1, \ldots, L.$$

In the case of doing nothing, the relative cost is computed by conditioning on the event that the component survives the current decision period. Denote the single period transition probabilities between $\dot{b}_i$ and $\dot{b}_j$, given the component survives or fails by q_{ij}, where

$$\begin{aligned} q_{ij} &= \mathbb{P}(b^{(n+1)} = \dot{b}_j | b^{(n)} = \dot{b}_i, \dot{x}_j \leq x_c) \\ &= \dot{p}_{ij}\, \mathbb{I}(\dot{x}_j \leq x_c)/R_i(x_c), \quad i = 0, 1, \ldots, L, \quad j = 1, 2, \ldots, L. \end{aligned}$$

Let τ_i be the expected survival time in the decision period, given $\dot{b}_i$, where

$$\tau_i = \sum_{j \in \mathcal{S}} \sum_{k \in \mathcal{S}} \int_0^{\delta} W_{jk}(x_c - \dot{x}_i, t) \mathrm{d}t, \quad i = 0, 1, \ldots, L.$$

Then the relative cost of doing nothing in $\dot{b}_i$ is then given as follows:

$$\dot{V}_0(i) = c_0 + \left[c_1 + c_2 + \dot{V}(0)\right] [1 - R_i(x_c)] + R_i(x_c) \sum_{j=1}^{L} \dot{V}(j) q_{ij} - \gamma \tau_i, \quad i = 0, 1, \ldots, L.$$

The optimality equations (3.25) can be solved using policy iteration; however, for large problems, a linear programming formulation (LP) can be advantageous in that it can be easily solved using commercial optimization software. By convention, the solution of the LP primal formulation corresponds to the optimal relative cost of each belief state; whereas, the solution of the LP dual formulation corresponds to the optimal action for each belief state. The values of both the LP primal and dual objective functions are equal to the long-run average cost per unit time under the optimal replacement policy. Before presenting the primal formulation, observe that setting $\dot{V}(0) \equiv 0$ leads to the following condition for each discretized belief state $i = 0, 1, \ldots, L$:

$$\dot{V}(i) = \min \left\{ c_0 + c_1, \quad c_0 + (c_1 + c_2) [1 - R_i(x_c)] + R_i(x_c) \sum_{j=1}^{L} \dot{V}(j) q_{ij} - \gamma \tau_i \right\}. \tag{3.26}$$

However, condition (3.26) is equivalent to the following pair of linear constraints for states $i = 0, 1, \ldots, L$:

$$\begin{aligned} &\dot{V}(i) \leq c_0 + c_1, \\ &\dot{V}(i) - R_i(x_c) \sum_{j=1}^{L} \dot{V}(j) q_{ij} + \gamma \tau_i + \leq c_0 + (c_1 + c_2)[1 - R_i(x_c)]. \end{aligned} \tag{3.27}$$

Using the constraints (3.27), the primal LP formulation is as follows:

$$\text{maximize} \quad \gamma \tag{3.28a}$$

$$\begin{aligned} \text{subject to} \quad & \dot{V}(i) \leq c_0 + c_1, \quad \forall i = 0, 1, \ldots, L \\ & \dot{V}(i) - R_i(x_c) \sum_{j=0}^{L} \dot{V}(j) q_{ij} + \gamma \tau_i + \leq c_0 + (c_1 + c_2)[1 - R_i(x_c)], \quad \forall i = 0, 1, \ldots, L, \\ & \gamma \in \mathbb{R}. \end{aligned}$$

The dual LP is formulated directly from the primal. Let x_{ia} denote the steady state probability of being in belief state i and taking action a, $i = 0, 1, \ldots, L$, $a \in \mathcal{A}$. The dual formulation is

$$\begin{aligned} \text{minimize} \quad & \sum_{i=1}^{L} \{c_0 + (c_1 + c_2)[1 - R_i(x_c)]\} x_{i0} + (c_0 + c_1) \sum_{i=1}^{L} x_{i1} \\ \text{subject to} \quad & x_{i0} + x_{i1} - \sum_{j=0}^{L} q_{ji} \, R_j(x_c) \, x_{j0} = 0, \quad \forall \, i = 1, 2, \ldots, L, \\ & \sum_{i=1}^{L} \tau_i \, x_{i0} = 1, \\ & x_{ia} \geq 0, \quad i = 0, 1, \ldots, L, \quad a \in \{0, 1\}. \end{aligned}$$

In practice, obtaining the optimal policy by solving the dual formulation is problematic. The problem arises because a large subset of the belief space for both POMDP models consists of transient states, and the dual formulation can only assign optimal actions to recurrent states [83]. This limitation combined with effects due to discretizing the belief space often causes the dual solution to provide an ambiguous policy. As an alternative to solving the dual, the optimal policy is obtained directly from the primal solution. Consider a component that begins operation in $\dot{b}_i \in \dot{B}$. Preventive replacement is only optimal in $\dot{b}_i$ if the total expected bias incurred for immediate preventive replacement followed by resumption of the optimal policy is less than the total expected bias for doing nothing for a single period; that is, if

$$c_0 + (c_1 + c_2)[1 - R_i(x_c)] - \gamma \tau_i + R_i(x_c) \sum_{j=1}^{L} q_{ij} \dot{V}(j) \geq c_0 + c_1,$$

so that

$$c_2 + R_i(x_c)\left[\sum_{j=1}^{L} q_{ij}\dot{V}(j) - (c_1 + c_2)\right] - \gamma\tau_i \geq 0. \tag{3.30}$$

Therefore, the optimal policy will be obtained by first solving the primal LP formulation (3.28a) to obtain γ and $\dot{V}(i)$, for $i = 1, 2, \ldots, L$. Then condition (3.30) will be checked for each discretized belief state to obtain the corresponding optimal action.

For the replacement model of Section 3.3, the optimality equations on $\dot{B}$ follow in a similar manner. Let $\ddot{V}(i) \equiv \ddot{V}(\dot{b}_i)$ denote the relative cost given that the component starts operation in i, $i = 1, 2, \ldots, L$, and let $\ddot{V}_0(i) \equiv \ddot{V}_0(\dot{b}_i)$ and $\ddot{V}_1(i) \equiv \ddot{V}_1(\dot{b}_i)$ denote the relative costs given no action and replacement, respectively, in state i, $i = 1, 2, \ldots, L$. Then

$$\ddot{V}(i) = \min\{\ddot{V}_0(i), \ddot{V}_1(i)\}, \quad i = 0, 1, \ldots, L, \tag{3.31}$$

where

$$\begin{aligned}
\ddot{V}_0(i) &= c_0 + \left[c_1 + C_i^{(r)} - \delta_2\gamma + \ddot{V}(0)\right]\bar{R}_i(x_c) + R_i(x_c)\sum_{j=1}^{L}\ddot{V}(j)q_{ij} - \gamma\tau_i, \\
\ddot{V}_1(i) &= c_0 + c_1 + C_i^{(p)} - \delta_1\gamma + \ddot{V}(0).
\end{aligned}$$

To obtain the LP formulation of (3.31), observe that setting $\ddot{V}(0) \equiv 0$ leads to the following condition for each discretized belief state $i = 0, 1, \ldots, L$:

$$\ddot{V}(i) = \min\left\{c_0 + c_1 + C_i^{(p)} - \delta_1\gamma, \quad c_0 + \left(c_1 + C_i^{(r)} - \delta_2\gamma\right)\bar{R}_i(x_c) + R_i(x_c)\sum_{j=1}^{L}\ddot{V}(j)q_{ij} - \gamma\tau_i\right\},$$

which is equivalent to the following pair of linear constraints for states $i = 0, 1, \ldots, L$:

$$\begin{aligned}
&\ddot{V}(i) + \delta_1\gamma \leq c_0 + c_1 + C_i^{(p)}, \\
&\ddot{V}(i) - R_i(x_c)\sum_{j=1}^{L}\ddot{V}(j)q_{ij} + \gamma\left[\tau_i + \delta_2\bar{R}_i(x_c)\right] \leq c_0 + \left(c_1 + C_i^{(r)}\right)\bar{R}_i(x_c).
\end{aligned} \tag{3.32}$$

Using the constraints (3.32), the primal LP formulation is as follows:

$$\text{maximize} \quad \gamma \tag{3.33a}$$

$$\begin{aligned}
\text{subject to} \quad & \ddot{V}(i) + \delta_1\gamma \leq c_0 + c_1 + C_i^{(p)}, \quad \forall i = 0, 1, \ldots, L \\
& \ddot{V}(i) - R_i(x_c)\sum_{j=1}^{L}\ddot{V}(j)q_{ij} + \gamma\left[\tau_i + \delta_2\bar{R}_i(x_c)\right] \leq c_0 + \left(c_1 + C_i^{(r)}\right)\bar{R}_i(x_c), \quad \forall i = 0, 1, \ldots, L, \\
& \gamma \in \mathbb{R}.
\end{aligned}$$

The dual LP, formulated directly from the primal, is as follows:

$$\begin{aligned}
\text{minimize} \quad & \sum_{i=1}^{L} \left[c_0 + \left(c_1 + C_i^{(r)} \right) \bar{R}_i(x_c) \right] x_{i0} + \left(c_0 + c_1 + C_i^{(p)} \right) \sum_{i=1}^{L} x_{i1} \\
\text{subject to} \quad & x_{i0} + x_{i1} - \sum_{j=0}^{L} q_{ji} \, R_j(x_c) \, x_{j0} = 0, \quad \forall \, i = 1, 2, \ldots, L, \\
& \sum_{i=1}^{L} \left[\tau_i + \delta_2 \bar{R}_i(x_c) \right] x_{i0} + \delta_1 \sum_{i=1}^{L} x_{i1} = 1, \\
& x_{ia} \geq 0, \quad i = 0, 1, \ldots, L, \quad a \in \{0, 1\}.
\end{aligned}$$

Analogous to (3.30), a condition that must hold in a given belief state for preventive replacement to be optimal is

$$c_0 + \left(c_1 + C_i^{(r)} - \delta_2 \gamma \right) \bar{R}_i(x_c) + R_i(x_c) \sum_{j=1}^{L} \ddot{V}(j) q_{ij} - \gamma \tau_i \geq c_0 + c_1 + C_i^{(p)} - \delta_1 \gamma$$

or

$$\gamma \left(\delta_1 - \tau_i \right) + \left(C_i^{(r)} - \delta_2 \gamma \right) \bar{R}_i(x_c) + R_i(x_c) \left[\sum_{j=1}^{L} \ddot{V}(j) q_{ij} - c_1 \right] - C_i^{(p)} \geq 0. \tag{3.35}$$

Therefore, the optimal policy is obtained by solving the primal LP formulation (3.33a) to obtain γ and $\ddot{V}(i)$, for $i = 1, 2, \ldots, L$. Then condition (3.35) is checked for each discretized belief state to assign optimal actions.

3.5 POLICIES FOR OBSERVABLE DEGRADATION

In this section, the results of numerical experiments are provided in which optimal policies of both replacement models are obtained under various environment, cost, or replacement assumptions. First considered are experiments for the replacement model of Section 3.2 in which replacement policies are obtained for a single component that is subjected to different operating environments. It is assumed that the observation interval is $\delta = 1.0$, and the component fails when its cumulative degradation reaches $x_c = 40.0$. The operating environments are assumed to consist of $\ell = 3$ states with the following parameterized generator matrix and degradation rates:

$$\boldsymbol{Q}(q) = \begin{bmatrix} -1.0 & 0.9 & 0.1 \\ 0.9 & -1.0 & 0.1 \\ 0.5 + q/2 & 0.5 + q/2 & -1.0 - q \end{bmatrix}, \qquad \boldsymbol{r}(w) = \begin{bmatrix} 1.0 & 2.0 & w \end{bmatrix},$$

where $q \in \{-0.5, 0.0, 1.0\}$ and $w \in \{3.0, 5.0, 7.0\}$. In what follows, optimal replacement policies are obtained for each pair (q, w) when $c_0 = 1.0$, $c_1 = 10.0$, and $c_2 = 2.5$.

To solve each example numerically, a discretization interval of 0.2 is used for both Π and $\mathcal{X}\backslash\{0\}$, which corresponds to fixing $L_1 = 21$ and setting $L_2 = 210$, $L_2 = 220$ and $L_2 = 230$ when $w = 3.0$, $w = 5.0$, and $w = 7.0$, respectively. The respective total number of states when $w = 3.0$, $w = 4.0$, and $w = 5.0$ are $L = 4,453$, $L = 4,663$, and $L = 4,875$. The optimal policy costs associated with each environment are shown in Table 8. The results show that γ increases as q decreases and w

Table 8: Policy cost (γ) for various (q, w).

w \ q	-0.5	0.0	1.0
3.0	1.4797	1.4557	1.4422
5.0	1.5664	1.5084	1.4709
7.0	1.6433	1.5599	1.5004

increases, which is quite intuitive. As q decreases, the environment spends a greater proportion of time in $\ell = 3$, and the component's expected lifetime is reduced. The resulting increase in replacement frequency leads to a higher policy cost. Likewise, as w increases, the component's expected lifetime is reduced as degradation proceeds more rapidly, and the resulting replacement frequency increases.

For each (q, w) and $\dot{\pi}$, Table 9 shows the corresponding replacement threshold $x_{\dot{\pi}}$. For a fixed $\dot{\pi}$, the threshold $x_{\dot{\pi}}$ tends to decrease as q decreases and w increases, which is opposite to the behavior of γ. Although this trend is weak when the probability of being in a low-degradation state is relatively high, the trend becomes markedly stronger as the probably of being in a high-degradation state increases. The rationale is that when the components have shorter expected lifetimes, it is advantageous to be more conservative by replacing at a lower cumulative degradation level in order to avoid incurring costly reactive replacement penalties. For this same reason, preventive replacement also tends to occur earlier for fixed (q, w) when $\dot{\pi}$ indicates a higher likelihood that the environment is in a state with a relatively high degradation rate.

The performance of the POMDP policies is evaluated by comparing their costs with the costs of an age-replacement policy and a policy consisting of only reactive replacements. The age-replacement policy is defined to be such that the component is replaced either at 90% of the component's expected lifetime or upon failure, whichever occurs first. Performance is evaluated for

Table 9: Optimal replacement thresholds ($x_{\dot{\pi}}$) for (q, w) and $\dot{\pi}$.

	$w = 3.0$			$w = 5.0$			$w = 7.0$		
$\dot{\pi}$	$q = -0.5$	$q = 0.0$	$q = 1.0$	$q = -0.5$	$q = 0.0$	$q = 1.0$	$q = -0.5$	$q = 0.0$	$q = 1.0$
[1.0 0.0 0.0]	38.6	38.6	38.6	38.6	38.6	38.6	38.6	38.6	38.6
[0.8 0.2 0.0]	38.4	38.4	38.4	38.4	38.4	38.4	38.6	38.4	38.4
[0.8 0.0 0.2]	38.2	38.2	38.4	38.2	38.2	38.2	38.2	38.0	38.0
[0.6 0.4 0.0]	38.2	38.2	38.2	38.4	38.2	38.2	38.4	38.2	38.2
[0.6 0.2 0.2]	38.0	38.0	38.2	38.0	38.0	38.0	38.0	37.8	38.0
[0.6 0.0 0.4]	37.4	37.8	38.0	37.0	37.2	37.4	36.6	36.8	37.2
[0.4 0.6 0.0]	38.2	38.2	38.2	38.2	38.2	38.2	38.2	38.2	38.2
[0.4 0.4 0.2]	38.0	38.0	38.0	38.0	38.0	38.0	38.0	37.8	38.0
[0.4 0.2 0.4]	37.4	37.8	38.0	36.8	37.0	37.4	36.6	36.6	37.0
[0.4 0.0 0.6]	37.2	37.4	37.8	36.2	36.4	37.0	35.4	35.8	36.4
[0.2 0.8 0.0]	38.0	38.0	38.0	38.0	38.0	38.0	38.2	38.0	38.0
[0.2 0.6 0.2]	38.0	38.0	38.0	38.0	38.0	38.0	38.0	37.8	38.0
[0.2 0.4 0.4]	37.4	37.8	38.0	36.8	37.0	37.4	36.4	36.6	37.0
[0.2 0.2 0.6]	37.2	37.4	37.8	36.0	36.4	37.0	35.4	35.8	36.4
[0.2 0.0 0.8]	37.2	37.2	37.6	35.2	36.0	36.6	34.2	34.8	36.0
[0.0 1.0 0.0]	38.0	38.0	38.0	38.0	38.0	38.0	38.0	38.0	38.0
[0.0 0.8 0.2]	38.0	38.0	38.0	38.0	38.0	38.0	38.0	37.8	38.0
[0.0 0.6 0.4]	37.4	37.8	38.0	36.6	36.8	37.4	36.2	36.4	36.8
[0.0 0.4 0.6]	37.2	37.4	37.8	36.0	36.2	37.0	35.2	35.8	36.2
[0.0 0.2 0.8]	37.2	37.2	37.6	35.2	35.8	36.6	34.0	34.8	35.8
[0.0 0.0 1.0]	37.2	37.2	37.6	35.2	35.4	36.4	33.2	34.2	35.4

$(q, w) \in \{(0, 3.0), (-0.5, 7.0)\}$, $c_0 = 1.0$, $c_1 = 10.0$, and $c_2 \in \{2.5, 5.0, 10.0, 20.0, 30.0, 40.0\}$. Define

$$c_r \equiv \frac{c_1 + c_2}{c_1}$$

as the ratio of the reactive replacement cost to the preventive replacement cost. Figures 29(a) and 29(b) show a comparison of the policy costs for environments $(q, w) = (0, 3.0)$ and $(q, w) = (-0.5, 7)$, respectively, as a function of c_r. In both environments, the performance of the POMDP policies is superior to those of the age- and reactive-replacement policies. As c_r increases and reactive failures become more punitive, the cost of both the age- and reactive-replacement policies increases at a much faster rate than the cost of the POMDP policy. The performance disparity is greater in environment $(q, w) = (-0.5, 7.0)$ which indicates that environment-state information is more valuable in cases when the environment has a state with relatively high degradation rate and long expected holding time.

In the next experiment, the replacement model of Section 3.3 is solved numerically under different assumptions for the operating environment, downtime costs, and replacement times. Here it is assumed that $x_c = 40.0$, $\delta = 1.0$, $\delta_1 = 0.5$, $c_0 = 1.0$, and $c_1 = 1.5$. All degradation environ-

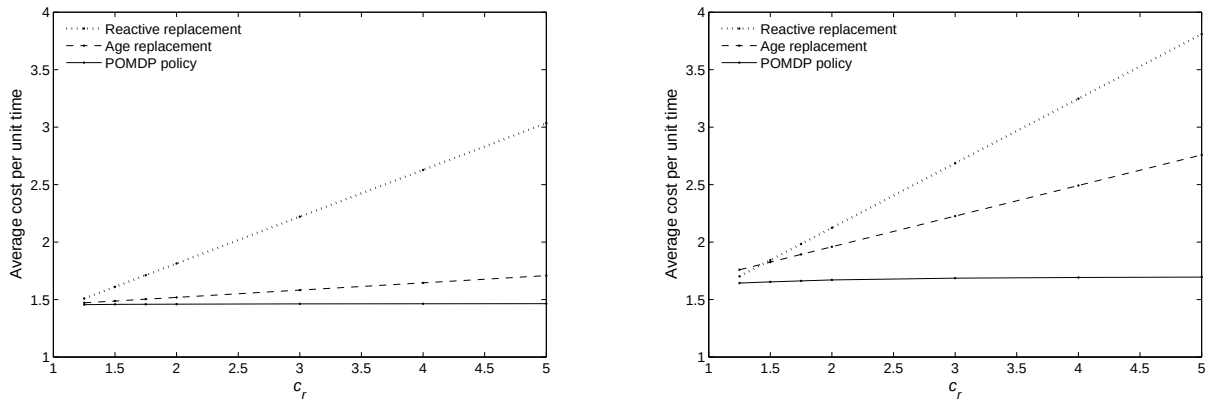

(a) Policy costs for $(q, w) = (0, 3.0)$.

(b) Policy costs for $(q, w) = (-0.5, 7.0)$.

Figure 29: Comparison of replacement policy performance.

ments consist of $\ell = 3$ states with degradation rates $\boldsymbol{r} = [1.0\ 2.0\ 5.0]$; however, the environments' stochastic behavior corresponds to one of three generator matrices as follows:

$$\boldsymbol{Q}_1 = \begin{bmatrix} -1.0 & 1.0 & 0 \\ 0 & -1.0 & 1.0 \\ 1.0 & 0 & -1.0 \end{bmatrix}, \quad \boldsymbol{Q}_2 = \begin{bmatrix} -1.0 & 0.5 & 0.5 \\ 0.5 & -1.0 & 0.5 \\ 0.5 & 0.5 & -1.0 \end{bmatrix},$$

$$\boldsymbol{Q}_3 = \begin{bmatrix} -1.0 & 0 & 1.0 \\ 1.0 & -1.0 & 0 \\ 0 & 1.0 & -1.0 \end{bmatrix}.$$

Note that $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_3$ are cyclic environments in which transitions occur in the order $(1 \rightarrow 2 \rightarrow 3 \rightarrow 1)$ and $(3 \rightarrow 2 \rightarrow 1 \rightarrow 3)$, respectively. Environments defined by $\boldsymbol{Q}_2$ are acyclic and transition randomly to a higher or lower environment state index. To observe the effect of different reactive replacement times and downtime cost vectors, optimal polices are derived for $\delta_2 \in \Delta_2 \equiv \{0.75, 1.0, 1.5, 2.0, 5.0\}$ and

$$\boldsymbol{d}_1 = \begin{bmatrix} 1.0 & 1.1 & 8.0 \end{bmatrix}, \quad \boldsymbol{d}_2 = \begin{bmatrix} 1.0 & 4.0 & 8.0 \end{bmatrix}.$$

Let $\mathcal{P}(i, j, \delta_2)$ denote the optimal policy obtained by solving the optimality equations assuming $\boldsymbol{Q}_i$, $\boldsymbol{d}_j$, and δ_2, where $i \in \{1, 2, 3\}$ and $j \in \{1, 2\}$. Let $\mathcal{D}_j \equiv \{\mathcal{P}(i, j, \delta_2) : i \in \{1, 2, 3\}, \delta_2 \in \Delta_2\}$ denote the set of optimal policies obtained under $\boldsymbol{d}_j$, $j \in \{1, 2\}$.

The optimality equations are solved numerically by discretizing Π and $\mathcal{X} \setminus \{0\}$ over intervals of length 0.2 so that $L_1 = 21$, $L_2 = 220$, and $L = 4,663$. Figures 30(a) and 30(b), 31(a) and

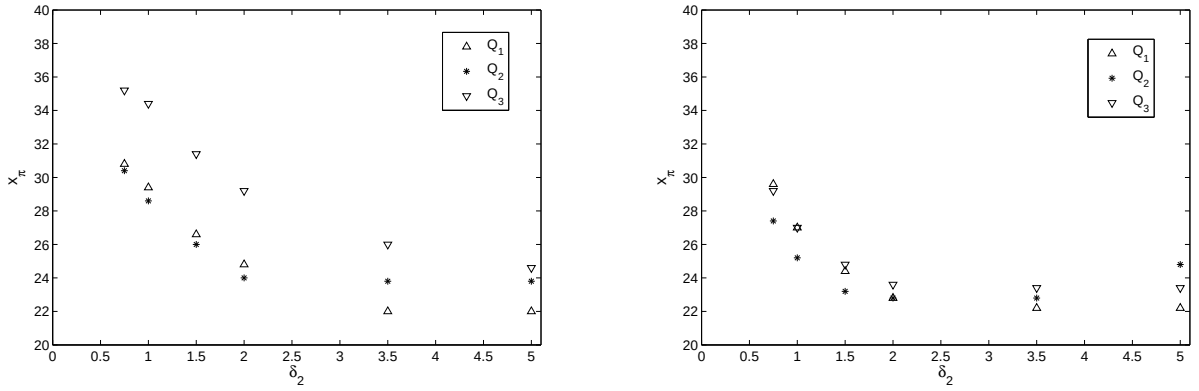

(a) Replacement thresholds policies in $\mathcal{D}_1$.

(b) Replacement thresholds policies in $\mathcal{D}_2$.

Figure 30: Replacement thresholds for policies in $\mathcal{D}_1$ and $\mathcal{D}_2$ when $\pi = [1\ 0\ 0]$.

31(b), and 32(a) and 32(b) show the cumulative degradation thresholds that correspond to being in environment states 1, 2, and 3, respectively, w.p. 1 for $\mathcal{D}_1$- and $\mathcal{D}_2$-policies. (Note that for cases in which no degradation threshold exists for a given environment belief state, the threshold is assigned the value $x_c = 40.0$.) The relatively high threshold values for the optimal policies when $\pi = [0\ 0\ 1]$ as compared to when $\pi \in \{[1\ 0\ 0], [0\ 1\ 0]\}$ indicate that policies tend to avoid initiating replacements in state 3, due to the state's relatively high downtime cost rate. In fact, when $\pi = [0\ 0\ 1]$ and $\delta_2 \leq 1.5$, preventive replacement is never optimal. For $\mathcal{D}_1$-policies, there is a markedly different preference for the optimal environment state to initiate a preventive replacement between $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_3$ environments. For $\pi = [1\ 0\ 0]$, the replacement thresholds for the $\boldsymbol{Q}_1$ environment in $\mathcal{D}_1$ are less than those of the corresponding thresholds for the $\boldsymbol{Q}_3$ environment. In other words, $\mathcal{D}_1$-policies have a stronger preference to initiate a preventive replacement when the environment is in state 1 for a $\boldsymbol{Q}_1$ environment than for a $\boldsymbol{Q}_3$ environment. That the opposite relationship exists between the replacement thresholds in $\pi = [0\ 1\ 0]$ indicates that $\mathcal{D}_1$-policies for $\boldsymbol{Q}_3$ tend to favor initiating preventive replacements in state 2.

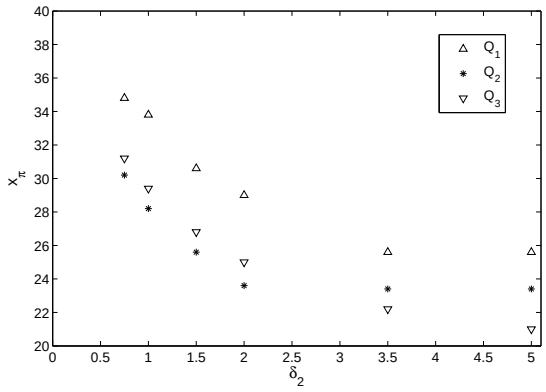

(a) Replacement thresholds policies in $\mathcal{D}_1$.

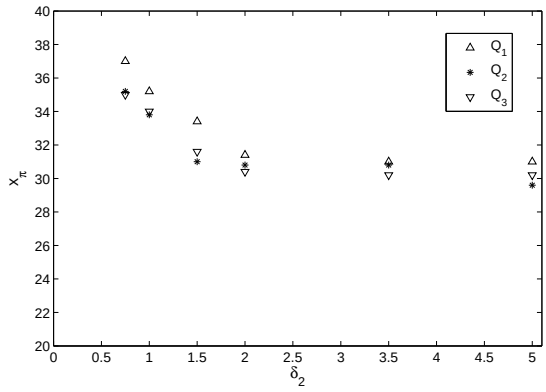

(b) Replacement thresholds policies in $\mathcal{D}_2$.

Figure 31: Replacement thresholds for policies in $\mathcal{D}_1$ and $\mathcal{D}_2$ when $\pi = [0\ 1\ 0]$.

The difference of environment state preferences arises because the optimal policies seek to provide the longest amount of time for a preventive replacement to complete before their respective environments reach state 3. Since the downtime costs between states 1 and 2 in $\boldsymbol{d}_1$ are nearly identical, a cost-minimizing strategy mainly consists of avoiding replacements during periods when the environment is in state 3. For $\mathcal{D}_2$-policies, in which state 2 has a moderately high downtime cost, this preference is relatively weak and does not hold for every value of δ_2.

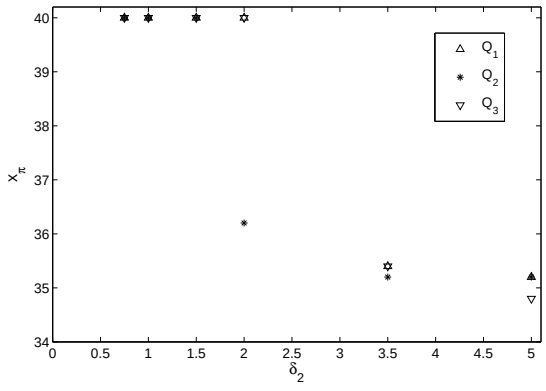

(a) Replacement thresholds policies in $\mathcal{D}_1$.

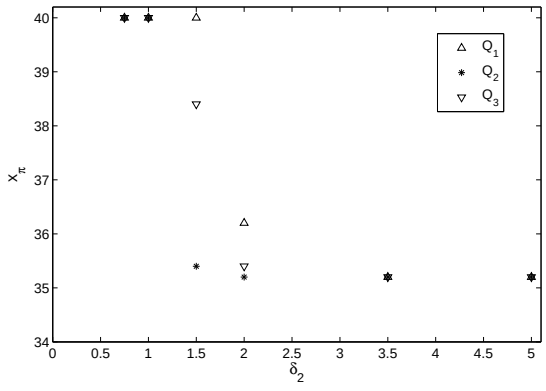

(b) Replacement thresholds policies in $\mathcal{D}_2$.

Figure 32: Replacement thresholds for policies in $\mathcal{D}_1$ and $\mathcal{D}_2$ when $\pi = [0\ 0\ 1]$.

Figures 33(a) and 33(b) show plots of the costs of $\mathcal{D}_1$- and $\mathcal{D}_2$-policies, respectively. The costs of the $\mathcal{D}_2$-policies for $\boldsymbol{Q}_3$ are significantly higher than those of the respective $\mathcal{D}_2$-policies for $\boldsymbol{Q}_1$. That the $\boldsymbol{Q}_3$ policies tend to have higher costs is due to the fact that their cost-minimizing strategy of avoiding replacements while the environment is in state 3 requires initiating replacements in state 2, which has a moderately high downtime cost. In the case of $\boldsymbol{Q}_2$, the cost is almost always higher than those of the corresponding $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_3$ environment policies. This relatively high cost arises because it is not possible to guarantee avoidance of state 3 for at least one period by initiating a preventive replacement in state 1 or 2.

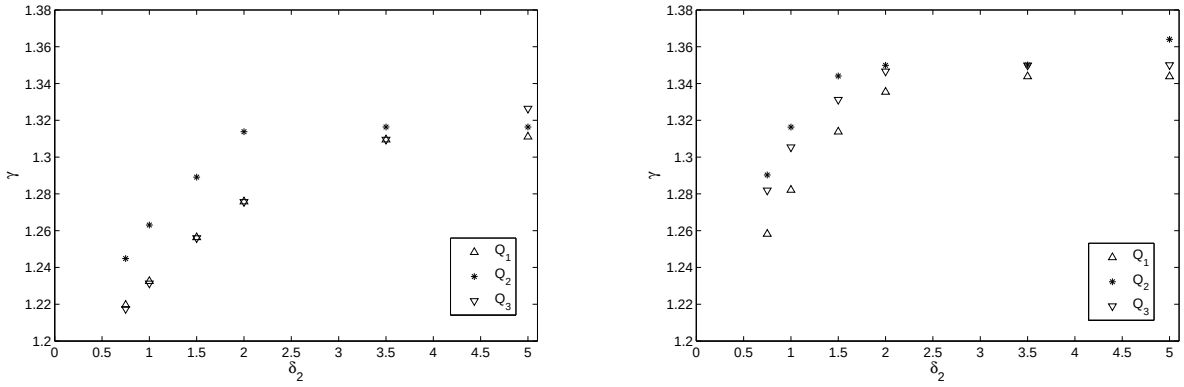

(a) Average cost of policies in $\mathcal{D}_1$.

(b) Average cost of policies in $\mathcal{D}_2$.

Figure 33: Average costs of policies in $\mathcal{D}_1$ and $\mathcal{D}_2$ as a function of δ_2.

Two POMDP models for optimally replacing a single wind turbine component were presented in this chapter that assumed perfect observations of the component's degradation level. In the next chapter, replacement of the wind turbine component is considered for imperfect observations, where a surrogate signal of component degradation is observed.

4.0 REPLACEMENT FOR IMPERFECTLY-OBSERVED DEGRADATION

In this chapter, the replacement model of Section 3.2 is extended to cases in which degradation is observed imperfectly through readings from a noisy sensor. Structural results are presented to characterize the optimal policy, and a projection-filtering algorithm that employs a belief-projection technique to reduce the dimensionality of the belief space, is used to solve the optimality equations numericaly. A modification of the algorithm is developed for the case when the environment state is known with certainty. To illustrate the optimal policies of the replacement models, and to compare their performance with the case where the environment is perfectly observed, the results of several numerical experiments are presented.

4.1 PARTIALLY-OBSERVED MARKOV DECISION PROCESS MODEL

In this section, the model of Section 3.2 is extended to the case of imperfectly-observed degradation. Here, it is assumed that the degradation level is assessed by observing a random signal Y_n that has a probability density function (p.d.f.) parameterized by X_n and $\theta \in \Theta$, where $\Theta \subseteq \mathbb{R}^m$. Let

$$\psi_n(x) \equiv \frac{\partial}{\partial u}\mathbb{P}(X_n \leq u | y_n, y_{n-1}, \ldots, y_0, a_{n-1}, a_{n-2}, \ldots, a_0)\Big|_{u=x}$$

be defined as the belief state of degradation at the beginning of the nth epoch, and Ψ be the belief space of degradation, where Ψ is an infinite-dimensional space of probability densities. To formulate the replacement problem as a POMDP, denote the belief state in the nth decision epoch by $b^{(n)} = (\psi_n, \pi_n)$. Given all prior signal observations up to epoch n, $\{y_0, y_1, \ldots, y_n\}$, and actions $\{a_0, a_1, \ldots, a_{n-1}\}$, the density $\psi_n(x)$ is expressed as a recursive function of a p.d.f. p as follows:

$$\begin{aligned}\psi_n(x) &\equiv p(x|y_0, \ldots, y_n, a_0, \ldots, a_{n-1}, \pi_{n-1}) \\ &\propto p(y_n|x_n = x, a_{n-1}, \pi_{n-1}) \int_0^{x_c} p(x|a_{n-1}, x_{n-1}, \pi_{n-2})\psi_{n-1}(x_{n-1})\mathrm{d}x_{n-1}. \qquad (4.1)\end{aligned}$$

Computing the updated belief state using (4.1), which is conditioned on all past signal observations and actions, is not analytically tractable. Instead, a simplified POMDP model will first be described in which the belief state is recursively updated using only the most recent signal observation and action. Let $f(y_n, x_n; \theta)$ denote the joint p.d.f. of (Y_n, X_n) under parameter θ, where $(x_n, \theta) \in [0, x_c] \times \Theta$, and let the conditional density of X_n, given observation y_n be denoted $h(x_n|y_n; \theta)$, where

$$h(x_n|y_n;\theta) = \frac{f(y_n, x_n; \theta)}{\displaystyle\int_0^{x_c} f(y_n, u; \theta)\,\mathrm{d}u}, \quad (x_n, y_n) \in [0, x_c] \times \mathbb{R}, \quad \theta \in \Theta.$$

Given no replacement, the updated belief state element $\pi_n^{(i)}$ is computed in this simplified model as $\pi_n^{(i)} = \boldsymbol{T}^*(y_n, y_{n-1}, \pi_{n-1})$, where for $(y', y, \pi) \in \mathbb{R}^2 \times \Pi$ and $T_i(u, \pi)$ as defined in (3.1),

$$\begin{aligned}
\boldsymbol{T}^*(y', y, \pi) &\equiv [T_1^*(y', y, \pi), T_2^*(y', y, \pi), \ldots, T_\ell^*(y', y, \pi)], \\
T_i^*(y', y, \pi) &\equiv \mathbb{P}(Z_n = j | y', y, \pi) \\
&= \int_0^\infty \mathbb{P}(Z_{n+1} = j | \Delta X_n = u, \pi) g(u|y', y)\,\mathrm{d}u \\
&= \int_0^\infty T_i(u, \pi) g(u|y', y)\,\mathrm{d}u,
\end{aligned}$$

and

$$\begin{aligned}
g(u|y', y) &= \frac{\partial}{\partial u'} \mathbb{P}(X_n - X_{n-1} \le u' | y', y)\Big|_{u'=u} \\
&= \frac{\partial}{\partial u'} \int_{-\infty}^{\infty} \int_{-\infty}^{u'+v} h(w|y', \theta) h(v|y, \theta) \mathrm{d}w\,\mathrm{d}v \Big|_{u'=u} \\
&= \int_{-\infty}^{\infty} h(u + v|y', \theta) h(v|y, \theta) \mathrm{d}v.
\end{aligned}$$

Because the density $h(x|y;\theta)$ is determined by y, the belief state of the simplified POMDP is denoted by $b^{(n)} = (y_n, \pi_n)$, and the belief space of the POMDP is $\mathcal{H}(\Theta) \times \Pi$, where $\mathcal{H}(\Theta) = \{h(x|y;\theta) : \theta \in \Theta\}$ is the set of all possible densities $h(x|y;\theta)$ for $\theta \in \Theta$. Letting $k(u, \pi)$ be defined as in (3.1), the transition kernel density between $(y, \pi) \in B$ and $(y', \pi') \in B$ conditioned on the non-occurrence of a failure in the next period is denoted by $k^*((y, \pi), (y', \pi'))$, where

$$k^*((y, \pi), (y', \pi')) = \int_0^{x_c} \int_0^\infty f(y'|x + u, \theta) k(u, \pi) h(x|y, \theta) \mathrm{d}u\,\mathrm{d}x \tag{4.2}$$

when $\pi' = \boldsymbol{T}^*(y', y, \pi)$ and is zero otherwise. It is important to note in (4.2) that the transition kernel density is only non-zero at a single element $\pi' \in \Pi$, which is obtained directly as a function of (y', y, π).

As in Section 3.2, all degradation observations and preventive replacements (if chosen) occur instantaneously at the beginning of a period and cost c_0 and c_1, respectively ($0 < c_0 < c_1 < \infty$). If the component fails between two observation epochs, a reactive replacement is instantly performed and costs $c_1 + c_2$, where $c_2 > 0$ is a penalty. The components are assumed to begin operation in belief state $(0, \pi_s)$, where π_s is the stationary distribution of the environment. The objective is to minimize the long-run average cost of replacements per unit time, denoted γ, where

$$\gamma = \inf_{a \in \mathcal{P}} \mathbb{E}_a \left\{ \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} c_0 + c_1 \, \mathbb{I}\{a(Y_n, \pi_n) = 1\} + (c_1 + c_2) \, \mathbb{I}\{a(Y_n, \pi_n) = 0, \widehat{G}(Y_n, \pi_n) = 1\} \right\},$$

and $\widehat{G}(Y_n, \pi_n)$ is the event the component fails between decision epochs n and $n+1$ given $(Y_n, \pi_n) \in B$.

The optimality equations are now presented. Let $V(y, \pi)$ be the minimum relative cost per unit time given a component starts operation in $(y, \pi) \in B$, and define $V_0(y, \pi)$ and $V_1(y, \pi)$ as the relative costs if either no action or preventive replacement, respectively, are taken in $(y, \pi) \in B$. For a new component that enters operation, the relative cost is denoted $V(0)$. The expected survival time of the component in the next period given $(y, \pi) \in B$ is denoted $\tau^*(y, \pi)$, where

$$\tau^*(y, \pi) = \int_0^{x_c} \tau(x, \pi) h(x|y, \theta) \, \mathrm{d}x,$$

and $\tau(x, \pi)$ defined by (3.1). The optimality equations are as follows:

$$V(y, \pi) = \min\{V_1(y, \pi), V_0(y, \pi)\}, \tag{4.3}$$

where $V_\pi(y', y) \equiv V(y', \boldsymbol{T}^*(y', y, \pi))$,

$$\begin{aligned} V_1(y, \pi) &= c_0 + c_1 + V(0) \\ V_0(y, \pi) &= c_0 + \int_0^{x_c} \int_0^{\infty} \mathbb{I}^+(x, u)[c_1 + c_2 + V(0)] k(u, \pi) h(x|y, \theta) \, \mathrm{d}u \, \mathrm{d}x - \gamma \tau^*(y, \pi) \\ &\quad + \int_0^{x_c} \int_0^{\infty} \int_{-\infty}^{\infty} \mathbb{I}^-(x, u) V_\pi(y', y) f(y'|x + u, \theta) k(u, \pi) h(x|y, \theta) \, \mathrm{d}y' \, \mathrm{d}u \;\; \mathrm{d}x. \end{aligned}$$

An analytical solution for the optimality equations (4.3) is not known except in trivial cases. Nonetheless, it is still possible to show some basic structural properties of the resulting optimal policy. In what follows, it is proved that under certain conditions the optimal policy has a threshold structure with respect to y for fixed $\pi \in \Pi$.

Lemma 4.1. *Let* $C_0 \equiv c_0 + c_1 + V(0)$*, then*

$$V_0(y,\pi) \geq C_0 - c_1 + (c_1 + c_2)G(y,\pi) - \gamma\tau^*(y,\pi), \tag{4.4}$$

where

$$G(y,\pi) \equiv \int_0^{x_c} \int_0^{\infty} H(x,\pi)h(x|y,\theta)\,\mathrm{d}x$$

is the failure probability in the next period, given $(y,\pi) \in B$ *and* $H(x,\pi)$ *is defined in* (3.3)*.*

Proof. To establish the lower bound, note that $V_\pi(y,\pi) \geq V(0)$ for $(y,\pi) \in B$. Defining $C_1 \equiv c_1 + c_2 + V(0)$, it follows that

$$\begin{aligned}
V_0(y,\pi) &= c_0 + \int_0^{x_c}\int_0^{\infty} \mathbb{I}^+(x,u)[c_1 + c_2 + V(0)]k(u,\pi)h(x|y,\theta)\,\mathrm{d}u\,\mathrm{d}x - \gamma\tau^*(y,\pi) \\
&\quad + \int_0^{x_c}\int_0^{\infty}\int_{-\infty}^{\infty} \mathbb{I}^-(x,u)V_\pi(y',y)f(y'|x+u,\theta)k(u,\pi)h(x|y,\theta)\,\mathrm{d}y'\,\mathrm{d}u\,\mathrm{d}x \\
&\geq c_0 + C_1\int_0^{x_c}\int_0^{\infty} \mathbb{I}^+(x,u)k(u,\pi)h(x|y,\theta)\,\mathrm{d}u\,\mathrm{d}x - \gamma\tau^*(y,\pi) \\
&\quad + V(0)\int_0^{x_c}\int_0^{\infty} \mathbb{I}^-(x,u)\left(\int_{-\infty}^{\infty} f(y'|x+u,\theta)\,\mathrm{d}y'\right)k(u,\pi)h(x|y,\theta)\,\mathrm{d}u\,\mathrm{d}x \\
&= c_0 + C_1G(y,\pi) + V(0)(1 - G(y,\pi)) - \gamma\tau^*(y,\pi) \\
&= C_0 - c_1 + (c_1 + c_2)G(y,\pi) - \gamma\tau^*(y,\pi).
\end{aligned}$$

□

Establishing simple properties of $V(y,\pi)$, such as monotonicity, is difficult and is strongly dependent on the choice of $f(x,y;\theta)$. However, it is possible to utilize the lower bound (4.4), which is a relatively simple function, to obtain some basic structural results for the optimal policy. In order to establish these results, the conditional density $h(y|x;\theta)$ must satisfy some conditions which are based on the following definitions:

Definition 4.1. *For* $t \geq 0$ *and* $(x,\pi) \in [0,x_c] \times \Pi$*, let*

$$L(x,\pi,t) \equiv \sum_{i\in\mathcal{S}} \mathbb{P}(X(t) \geq x_c - x|X(0) = 0, Z(0) = i)\pi^{(i)}.$$

That is, $L(x,\pi,t)$ is the probability that the component fails during an observation interval of length t, given belief state (x,π).

Definition 4.2. *For any real function* $\phi(x,\pi)$ *and* $(x,\pi) \in [0,x_c] \times \Pi$*, define*

$$\mathbb{E}_y[\phi(x,\pi)] \equiv \int_0^{x_c} \phi(x,\pi) h(x|y,\theta)\,\mathrm{d}x.$$

Based on Definitions 4.1 and 4.2, a lemma can be stated that characterizes the functions $G(y,\pi)$ and $\tau^*(y,\pi)$.

Lemma 4.2. *If* $\mathbb{E}_y[L(x,\pi,t)]$ *is monotone increasing in* $y \in \mathbb{R}$*, then* $G(y,\pi)$ *is monotone increasing in* $y \in \mathbb{R}$*, and* $\tau^*(y,\pi)$ *is monotone decreasing in* $y \in \mathbb{R}$*.*

Proof. Observe that

$$\begin{aligned} H(x,\pi) &= \mathbb{P}(X_{n+1} \geq x_c | X_n = x, \pi_n = \pi) \\ &= \sum_{i\in\mathcal{S}} \mathbb{P}(X_1 \geq x_c - x | X_0 = 0, Z_0 = i)\pi^{(i)} \\ &= \sum_{i\in\mathcal{S}} \mathbb{P}(X(\delta) \geq x_c - x | X(0) = 0, Z(0) = i)\pi^{(i)} \\ &= L(x,\pi,\delta). \end{aligned}$$

Therefore, it follows directly from the hypothesis that $\mathbb{E}_y[H(x,\pi)] = G(y,\pi)$ is monotone increasing in $y \in \mathbb{R}$. Also, note that

$$\begin{aligned} \tau(x,\pi) &= \int_0^\delta \sum_{i\in\mathcal{S}} \mathbb{P}(X(t) \leq x_c | X(0) = x, Z(0) = i)\pi^{(i)}\,\mathrm{d}t \\ &= \sum_{i\in\mathcal{S}} \left\{ \int_0^\delta [1 - \mathbb{P}(X(t) \geq x_c - x | X(0) = 0, Z(0) = i)]\,\mathrm{d}t \right\} \pi^{(i)} \\ &= \sum_{i\in\mathcal{S}} \left[\delta - \int_0^\delta \mathbb{P}(X(t) \geq x_c - x | X(0) = 0, Z(0) = i)\,\mathrm{d}t \right] \pi^{(i)} \\ &= \delta - \sum_{i\in\mathcal{S}} \left[\int_0^\delta \mathbb{P}(X(t) \geq x_c - x | X(0) = 0, Z(0) = i)\,\mathrm{d}t \right] \pi^{(i)} \\ &= \delta - \int_0^\delta \left[\sum_{i\in\mathcal{S}} \mathbb{P}(X(t) \geq x_c - x | X(0) = 0, Z(0) = i)\pi^{(i)} \right] \mathrm{d}t, \end{aligned}$$

so

$$\begin{aligned}
\tau^*(y,\pi) &= \mathbb{E}_y[\tau(x,\pi)] \\
&= \mathbb{E}_y\left[\delta - \int_0^\delta \left(\sum_{i\in\mathcal{S}} \mathbb{P}(X(t) \geq x_c - x | X(0)=0, Z(0)=i)\pi^{(i)}\right) \mathrm{d}t\right] \\
&= \delta - \mathbb{E}_y\left[\int_0^\delta \left(\sum_{i\in\mathcal{S}} \mathbb{P}(X(t) \geq x_c - x | X(0)=0, Z(0)=i)\pi^{(i)}\right) \mathrm{d}t\right] \\
&= \delta - \int_0^{x_c}\int_0^\delta \left(\sum_{i\in\mathcal{S}} \mathbb{P}(X(t) \geq x_c - x | X(0)=0, Z(0)=i)\pi^{(i)}\right) h(x|y,\theta)\,\mathrm{d}t\,\mathrm{d}x \\
&= \delta - \int_0^\delta \int_0^{x_c} L(x,\pi,t)h(x|y,\theta)\,\mathrm{d}x\,\mathrm{d}t \quad \text{(Fubini)} \\
&= \delta - \int_0^\delta \mathbb{E}_y[L(x,\pi,t)]\,\mathrm{d}t.
\end{aligned}$$

Then $\int_0^\delta \mathbb{E}_y[L(x,\pi,t)]\,\mathrm{d}t$ is monotone increasing in y under the hypothesis, and $\tau^*(y,\pi)$ is monotone decreasing. □

Lemma 4.2 is based on the rather technical condition that $\mathbb{E}_y[L(x,\pi,t)]$ is increasing in y. However, this condition can be understood intuitively. The condition essentially requires that the conditional density $h(x|y;\theta)$ shift more probability mass to larger values of x as y increases. Since the probability $L(x,\pi,t)$ is clearly increasing in x, the integral $\mathbb{E}_y[L(x,\pi,t)]$ will increase as more probability mass is shifted to larger values of x. An illustration of this condition is shown in Figure 34.

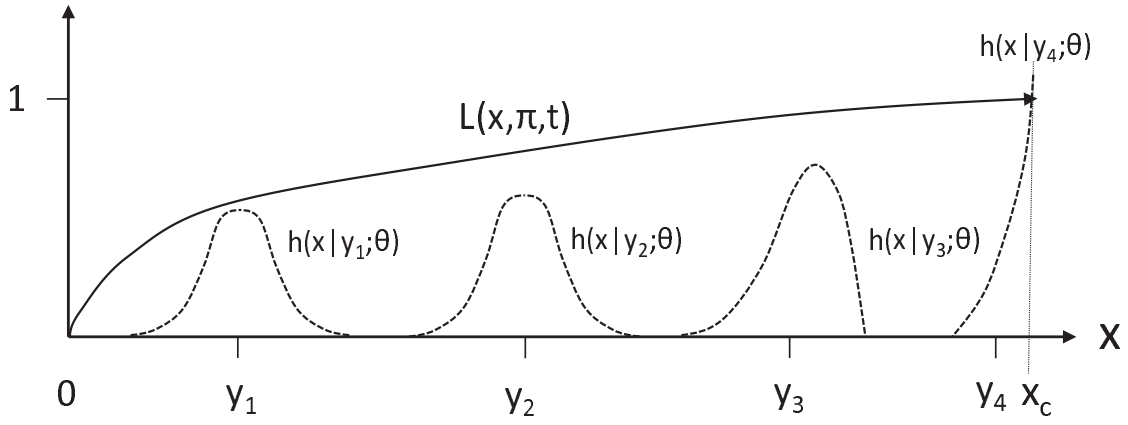

Figure 34: Depiction of monotonicity condition.

Theorem 4.1. *If* $\mathbb{E}_y[L(x,\pi,t)]$ *is monotone increasing in* $y \in \mathbb{R}$, $\lim_{y\to\infty} G(y,\pi) = 1$, *and* $\lim_{y\to\infty} \tau^*(y,\pi) = 0$, *then for each* $\pi \in \Pi$, *there exists a threshold* y_π *such that it is optimal to replace for any* $y' \geq y_\pi$.

Proof. By Lemma 4.1, a sufficient condition for the optimality of preventive replacement is

$$V_1(y,\pi) \leq C_0 - c_1 + (c_1 + c_2)G(y,\pi) - \gamma\tau^*(y,\pi),$$

which implies

$$0 \leq -c_1 + (c_1 + c_2)G(y,\pi) - \gamma\tau^*(y,\pi). \tag{4.5}$$

Under the hypothesis, the right side of (4.5) is monotone increasing by Lemma 4.2, and the inequality is satisfied in the limit as $y \to \infty$. Therefore, there exists $y_\pi \in \mathbb{R}$ such that $V_0(y,\pi) \geq V_1(y_\pi,\pi)$ for all $y \geq y_\pi$, and preventive replacement is optimal. □

Theorem 4.1 establishes that it is possible to obtain replacement thresholds for each environment state, not with respect to the cumulative degradation level, but with respect to a signal serving as an indirect measure of the true degradation. The existence of such thresholds is significant because the optimal policy can be easily implemented in wind turbine applications where the true degradation is typically not directly observed. As replacement decisions are not only based on the explicit value of the signal but also on the environment state, replacement decisions can take into account all information known about the environment.

4.2 NUMERICAL SOLUTION TECHNIQUES

Solving for the optimal policy on the belief space $B = \Psi \times \Pi$ is not analytically tractable as closed-form expressions for the updated belief state cannot be obtained in general. This section describes a numerical solution technique that uses particle-filtering and belief-projection techniques to approximate the optimal policy.

Standard approaches for solving POMDPs numerically often entail discretizing the belief space and then applying value- or policy-iteration algorithms. Because B is infinite dimensional, these standard solution approaches are not computationally tractable [116]. Instead a technique known as belief (or density) projection is employed that projects Ψ onto a family of parameterized densities

$\mathcal{G} = \{g(.;\theta) : \theta \in \Theta\}$, where Θ is the parameter space of the density family. The key idea is that for any given belief state $\psi \in \Psi$, there exists a $g^* \in \mathcal{G}$ with parameterization $\theta^* \in \Theta$ such that a measure of distance between the densities ψ and g^* (e.g., the Kullback-Leibler (KL) divergence) is minimized. In this way, elements of Ψ can be approximated by distributions parameterized on Θ, thereby, yielding a "parameterized" belief space $B_\Theta \equiv \Theta \times \Pi$. Assuming the dimensionality of Θ is low, the POMDP can be approximately solved using a policy iteration algorithm on the set $\dot{B}_\Theta \equiv \dot{\Theta} \times \dot{\Pi}$, where $\dot{\Pi}$ and $\dot{\Theta}$ are discretizations of Π and Θ, respectively. Table 10 summarizes the attributes and notation associated with each type of belief space.

Table 10: Summary of belief space attributes.

Type	State space	Belief state
True belief space	$B = \Psi \times \Pi$	$\psi \times \pi; \quad \psi \in \Psi, \quad \pi \in \Pi$
$\Downarrow$	$\Downarrow$	$\Downarrow$
Parameterized belief space	$B_\Theta = \Theta \times \Pi$	$g(\cdot;\theta) \times \pi; \quad g(\cdot;\theta) \in \mathcal{G}, \quad \pi \in \Pi$
$\Downarrow$	$\Downarrow$	$\Downarrow$
Discretized belief state	$\dot{B}_\Theta = \dot{\Theta} \times \dot{\Pi}$	$(\dot{\theta}_i, \dot{\pi}_i); \quad \dot{\theta}_i \in \dot{\Theta}, \quad \dot{\pi}_i \in \dot{\Pi}, \quad i = 1, 2, \ldots, L$

Ψ: Infinite-dimensional space of probability densities

Π: Probability simplex

$\mathcal{G} \equiv \{g(\cdot;\theta), \theta \in \Theta\}$: Family of densities parameterized on $\Theta \subseteq \mathbb{R}^m$

$\dot{\Theta}, \dot{\Pi}$: Discretizations of Θ and Π, respectively, $L \equiv |\dot{B}_\Theta|$

The numerical solution approach is now described in detail. Obtaining an approximate solution to the MDP on $\dot{B}_\Theta$ requires estimating the transition probabilities between all pairs of discretized belief states. Let $\dot{b}^{(n)} \in \dot{B}_\Theta$ denote the discretized belief state at the nth decision epoch, and $L \equiv |\dot{B}_\Theta|$ denote the total number of discretized states. For notational convenience, write $(\dot{b}^{(n)} = i) \equiv (\dot{b}^{(n)} = (\dot{\theta}_i, \dot{\pi}_i))$, where $\dot{\theta}_i \in \dot{\Theta}$, $\dot{\pi}_i \in \dot{\Pi}$, $i = 1, 2, \ldots, L$. We seek to estimate the matrix $\dot{\boldsymbol{P}} = [\dot{p}_{ij}]$, where $\dot{p}_{ij} = \mathbb{P}(\dot{b}^{(n+1)} = j | \dot{b}^{(n)} = i, x_{n+1} < x_c)$ as the transition probability between discretized belief states i and j in the nth decision epoch, given degradation has not exceeded x_c, $i, j = 1, 2, \ldots, L$. Let $\upsilon = (x, i)$, where $(x, i) \in \Upsilon = [0, x_c] \times \mathcal{S}$. This estimate is obtained in three major steps: (1) Estimating $b^{(n)}$ given $(b^{(n-1)}, \upsilon_{n-1}, a_{n-1})$, (2) Projecting $b^{(n)}$ onto B_Θ, and (3) Approximating B_Θ with $\dot{B}_\Theta$. In what follows, each step is described in detail.

Computing $b^{(n)}$ using (4.1) requires *filtering*. With the exception of a few instances (e.g. linear Gaussian systems), filtering cannot be solved analytically [20], so an approximate *particle filtering*

technique is used instead. The notion of particle filtering is to approximate a density with a probability mass function (p.m.f.) defined by a finite set of weighted particles [81]. Particle filtering confers the advantage that potentially intractable computations can be avoided by simulating the particles' probabilistic evolutions. In what follows, it is assumed that $Y_n \sim N(X_n, \sigma_\epsilon^2)$, where $\sigma_\epsilon > 0$. Denote the total number of particles by $\widehat{N}$, and let $\upsilon_n^{(i)} = (x_n^{(i)}, z_n^{(i)}) \in \Upsilon$ and $w_n^{(i)} \in (0, 1]$ be the core state and weight, respectively, of the ith particle at the nth decision epoch, $i = 1, 2, \ldots, \widehat{N}$, where $\sum_{i=1}^{\widehat{N}} w_n^{(i)} = 1$. Denoting by $\upsilon_{n|n-1}^{(i)}$ the state of particle i in decision epoch n, given it was in state $\upsilon_{n-1}^{(i)}$ in state $n-1$, the set $\hat{B}^{(n)} = \left\{(\upsilon_{n|n-1}^{(i)}, w_n^{(i)}) : i = 1, 2, \ldots, \widehat{N}\right\}$ defines a p.m.f. $\hat{b}^{(n)}(\upsilon)$ that approximates $b^{(n)}$ on the support Υ, where

$$\begin{aligned} \hat{b}^{(n)}(\upsilon) &= \sum_{i=1}^{\widehat{N}} w_n^{(i)} \mathbb{I}\left(\upsilon = \upsilon_{n|n-1}^{(i)}\right), \\ w_n^{(i)} &\propto p\left(y_n | \upsilon_{n|n-1}^{(i)}, a_{n-1}\right). \end{aligned}$$

The p.m.f. $\hat{b}^{(n)} \in B$ must be projected onto an element $\bar{b}^{(n)} \equiv (\bar{\theta}_n, \bar{\pi}_n) \in B_\Theta$. The vector $\bar{\pi}_n$ is obtained in a straightforward manner by averaging the environment states in $\hat{B}^{(n)}$. That is

$$\bar{\pi}_n = \widehat{N}^{-1} \sum_{i=1}^{\widehat{N}} \sum_{j=1}^{\ell} \mathbb{I}\left(z_{n|n-1}^{(i)} = j\right) \cdot \boldsymbol{e}_j, \tag{4.6}$$

where $\boldsymbol{e}_j$ is the jth unit vector. Obtaining $\bar{\theta}_n$ requires minimizing some measure of distance between $\hat{b}^{(n)}$ and $g(.;\theta)$. Typically the KL-divergence, denoted D_{KL}, is minimized where for distributions F_1 and F_2 with respective densities f_1 and f_2,

$$D_{KL}(F_1 || F_2) \equiv \int_{-\infty}^{\infty} f_1(x) \ln \frac{f_1(x)}{f_2(x)} \mathrm{d}x. \tag{4.7}$$

However, minimizing (4.7) can be quite arduous computationally, so for practical reasons, the maximum likelihood estimate (MLE) is used as a proxy for D_{KL}. Given $\hat{b}^{(n)}$ and $\mathcal{G}$, we set $\bar{\theta}_n$ equal to its MLE; that is,

$$\bar{\theta}_n = \underset{\theta \in \Theta}{\operatorname{argmin}} \sum_{i=1}^{\widehat{N}} w_n^{(i)} g\left(x_{n|n-1}^{(i)}; \theta\right). \tag{4.8}$$

The discretization $\dot{B}_\Theta$ of Θ is then obtained by mapping the projected belief state $\bar{b}^{(n)} = (\bar{\theta}_n, \bar{\pi}_n)$ to a discrete state with index $\bar{v} \in \{1, 2, \ldots, L\}$ by minimizing the norms between the projected and discretized parameters as follows:

$$\bar{v} = \underset{i=1,2,\ldots,L}{\operatorname{argmin}} \|\dot{\theta}_i - \bar{\theta}_n\| + \|\dot{\pi}_i - \bar{\pi}_n\|. \tag{4.9}$$

A simulation-based *projection particle filtering* algorithm, based on the approach of Zhou, *et al.* [117], is employed that computes each row of $\dot{\boldsymbol{P}}$ by performing both particle filtering and density projection. For $i = 1, 2, \ldots, L$, the projection-filtering algorithm computes row $[\dot{p}_{l\cdot}]$ as follows:

1. Sample particles $\{\upsilon_{n-1}^{(1)}, \upsilon_{n-1}^{(2)}, \ldots, \upsilon_{n-1}^{(2\widehat{N})}\}$, where $\upsilon_{n-1}^{(i)} \sim g(\cdot; \dot{\theta}_l) \times \dot{\pi}_l$, $i = 1, 2, \ldots, 2\widehat{N}$.

2. Simulate $\{\upsilon_{n|n-1}^{(1)}, \upsilon_{n|n-1}^{(2)}, \ldots, \upsilon_{n|n-1}^{(2\widehat{N})}\}$ according to (2.1) for time δ.

3. Sample $\{\upsilon_{n|n-1}^{'(1)}, \upsilon_{n|n-1}^{'(2)}, \ldots, \upsilon_{n|n-1}^{'(2\widehat{N})}\}$, where $\upsilon_{n|n-1}^{'(i)}$ is randomly drawn with replacement from $\{\upsilon_{n|n-1}^{(j)} : \upsilon_{n|n-1}^{(j)} < x_c, j = 1, 2, \ldots, 2\widehat{N}\}$.

4. Simulate cumulative degradation observations $\{y_1, y_2, \ldots, y_{\widehat{N}}\}$, where $y_i \sim \mathrm{N}(x_{n|n-1}^{'(i)}, \sigma_\epsilon^2)$.

5. Compute $w_n^{(ij)}$, the weight of the jth particle, given the ith observation, where

$$w_n^{(ij)} = \frac{\phi\left(\dfrac{y_i - \upsilon_{n|n-1}^{'(\widehat{N}+j)}}{\sigma_\epsilon}\right)}{\displaystyle\sum_{k=1}^{\widehat{N}} \phi\left(\frac{y_i - \upsilon_{n|n-1}^{'(\widehat{N}+k)}}{\sigma_\epsilon}\right)}, \quad i, j = 1, 2, \ldots, \widehat{N},$$

   and

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}x^2\right)$$

   is the standard normal density.

6. For $i = 1, 2, \ldots, \widehat{N}$, set $\hat{B}_i^{(n)} = \{\upsilon_{n|n-1}^{(j)}, w_n^{(ij)}, j = 1, 2, \ldots, \widehat{N}\}$ and compute $\bar{b}^{(n)}$ using (4.6) and (4.8). Determine the closest discretized belief state $\bar{v}_i$ to $\bar{b}_i^{(n)}$ using (4.9).

7. For $m = 1, 2, \ldots, L$, compute

$$p_{lm} = \widehat{N}^{-1} \sum_{i=1}^{\widehat{N}} \mathbb{I}(\bar{v}_i = m)$$

Now the MDP formulation on $\dot{B}_\Theta$ is presented. Denote the relative cost of starting in $\dot{b}^{(0)} = i$ by $V(i)$ and the relative costs of doing nothing and replacing a component in $\dot{b}^{(0)} = i$, respectively, by $V_0(i)$ and $V_1(i)$. The optimality equation is

$$V(i) = \min\{V_0(i), V_1(i)\}, \quad i = 1, 2, \ldots, L. \tag{4.10}$$

When a preventive replacement is performed in $\dot{b}^{(n)} = i$, the total relative cost consists of the inspection and replacement costs as well as the relative cost of starting in a "good-as-new" condition, denoted $V(0)$. It is assumed that after a replacement the cumulative degradation is zero w.p. 1, and the environment is distributed as the stationary distribution π_s; that is, $(x_n, z_n) \sim \delta(u) \times \pi_s$, where $\delta(u)$ is the dirac delta function. Then

$$V_1(i) = c_0 + c_1 + V(0), \quad i = 1, 2, \ldots, L.$$

If no replacement is made in a decision epoch, computing $V_0(i)$ requires conditioning on whether a failure occurs in the subsequent period. Abusing notation slightly, let $R_i(t) = \mathbb{P}(X(s+t) < x_c | b^{(s)} = i)$, $s \in \mathbb{R}_+$, be the probability of survival in time interval $(s, s+t)$, given $b^{(s)} = i$, where

$$R_i(t) = \frac{\int_0^{x_c} F(x_c - u, \delta; \dot{\pi}_i) g(u; \theta_i) \mathrm{d}u}{\int_0^{x_c} g(u; \theta_i) \mathrm{d}u} \tag{4.11}$$

and

$$F(x, t; \pi) = \sum_{i \in \mathcal{S}} \mathbb{P}(X(t) \leq x | X(0) = 0, Z(0) = i) \pi^{(i)}, \tag{4.12}$$

is the c.d.f. of the degradation increment on $[0, t]$, given the initial environment distribution $\pi \in \Pi$. The expected survival time during the next period, given $b^{(n)} = i$, denoted τ_i, is computed using (4.11) as

$$\tau_i = \int_0^{\delta} R_i(u) \mathrm{d}u, \quad i = 1, 2, \ldots, L.$$

The value function for doing nothing in $\dot{b}^{(n)} = i$ is

$$V_0(i) = c_0 + [c_1 + c_2 + V(0)] [1 - R_i(\delta)] + \sum_{j=1}^{L} p_{ij} R_i(\delta) V(j) - \gamma \tau_i,$$

where γ is the minimum long-run expected cost per unit time. The optimal policy to minimize long-run average cost can be obtained from the optimality equation (4.10) in a straightforward manner using policy iteration or an LP-based approach.

The replacement model and projection-filtering algorithm can be easily modified for the case where degradation is imperfectly observed, but the environment state is known with certainty. The belief space of the modified POMDP is $B_- = \Psi \times \Pi_-$, where $\Pi_- = \{\boldsymbol{e}_i : i = 1, 2, \ldots, \ell\}$ is simply the set of extreme vectors in the simplex Π. The POMDP is solved using the same general approach by formulating it as an MDP on B_-, which is then approximated using density projection as $B_{\Theta-} = \Theta \times \Pi_-$. However, the fact that the environment is observable actually complicates

the particle filtering slightly as resampling must be performed in a way that captures the extra information known about the environment. The transition probability matrix $\dot{\boldsymbol{P}}$ is estimated using a modification of the projection-filtering algorithm that partitions the particles by their environment state before resampling to account for observability of the environment state at each decision epoch. The modified projection-filtering algorithm to compute row $[\dot{p}_{l\cdot}]$ of $\dot{\boldsymbol{P}}$ for $l = 1, 2, \ldots, L$ is as follows:

1. Sample particles $\{\upsilon_{n-1}^{(1)}, \upsilon_{n-1}^{(2)}, \ldots, \upsilon_{n-1}^{(2\widehat{N})}\}$ from $\dot{b}_u \in \dot{B}_\Theta$, where $\upsilon_{n-1}^{(i)} \sim g(\cdot;\theta_l)\times\pi_l$, $i = 1, 2, \ldots, 2\widehat{N}$.

2. Simulate $\{\upsilon_{n|n-1}^{(1)}, \upsilon_{n|n-1}^{(2)}, \ldots, \upsilon_{n|n-1}^{(2\widehat{N})}\}$ according to (2.1) for time ΔI.

3. For $k \in \mathcal{S}$, compute $\widehat{N}^{(k)}$ the number of particles in environment state k, where

$$\widehat{N}^{(k)} = \sum_{i=1}^{2\widehat{N}} \mathbb{I}\left(z_{n|n-1}^{(i)} = k\right),$$

   and set

$$n(k) = \begin{cases} (\widehat{N}^{(k)} - 1)/2, & \text{if } \widehat{N}^{(k)} \text{ is even,} \\ \widehat{N}^{(k)}/2, & \text{if } \widehat{N}^{(k)} \text{ is odd.} \end{cases}$$

4. For $k \in \mathcal{S}$,

   a. Sample $\{\upsilon_{n|n-1}^{'(1)}, \upsilon_{n|n-1}^{'(2)}, \ldots, \upsilon_{n|n-1}^{'(n(k))}\}$, where $\upsilon_{n|n-1}^{'(i)}$ is randomly sampled with replacement from $\{\upsilon_{n|n-1}^{(j)} : \upsilon_{n|n-1}^{(j)} < x_c, z_{n|n-1}^{(j)} = k, j = 1, 2, \ldots, n(k)\}$.

   b. Simulate cumulative degradation observations $\{y_1, y_2, \ldots, y_{n(k)}\}$, where $y_i \sim N\left(x_{n|n-1}^{'(i)}, \sigma_\epsilon^2\right)$.

   c. Compute $w_n^{(ij)}$, the weight of the jth particle, given the ith observation, where

$$w_n^{(ij)} = \frac{\phi\left(\dfrac{y_i - \upsilon_{n|n-1}^{'(n(k)+j)}}{\sigma_\epsilon}\right)}{\displaystyle\sum_{l=1}^{\widehat{N}} \phi\left(\frac{y_i - \upsilon_{n|n-1}^{'(n(k)+l)}}{\sigma_\epsilon}\right)}, \quad i, j = 1, 2, \ldots, n(k).$$

d. For $i = 1, 2, \ldots, \widehat{N}$ set $\hat{B}_i^{(n)} = \{\upsilon_{n|n-1}^{(j)}, w_n^{(ij)}, j = 1, 2, \ldots, n(k)\}$ and compute $\bar{b}_i^{(n)}$ using (4.6) and (4.8). Determine the closest discretized belief state $\bar{v}_i$ to $\bar{b}_i^{(n)}$ using (4.9).

e. For $m = 1, 2, \ldots, L$, compute

$$p_{lm} = \eta_k \sum_{i=1}^{n(k)} \mathbb{I}(\bar{v}_i = m),$$

where the correction factor η_k is

$$\eta_k = \frac{\widehat{N}^{(k)}}{2\widehat{N}n(k)}.$$

The optimality equations are identical to (4.10).

4.3 POLICIES FOR IMPERFECTLY-OBSERVED DEGRADATION

In this section, three numerical examples are presented that illustrate optimal policies obtained from the two replacement models of Section 4.2 in which (i) the environment is partially-observable and degradation is imperfectly observed (PEID) and (ii) the environment is observable and degradation is imperfectly observed (OEID). To provide a basis for comparison in order to assess the cost of uncertainty in degradation observations, the performance of PEID and OEID policies is compared with policies obtained from the replacement model of Section 3.2 that assumes a partially-observed environment and observable degradation (PEOD). The performance of all three models will also be compared with those of the age- and reactive-replacement policies described in Section 3.5. The first numerical example illustrates optimal policies for a relatively simple degradation environment; whereas, the second and third examples consider replacement policies for wind turbine components based on real data.

4.3.1 An Illustrative Example

Consider a component that operates in an $\ell = 3$ state CTMC environment with the following parameters:

$$\boldsymbol{Q} = \begin{bmatrix} -1.0 & 0.9 & 0.1 \\ 0.9 & -1.0 & 0.1 \\ 0.5 & 0.5 & -1.0 \end{bmatrix}, \qquad \boldsymbol{r} = \begin{bmatrix} 1.0 & 2.0 & 5.0 \end{bmatrix}.$$

The component's cumulative degradation is measured at a constant inspection interval $\delta = 1.0$, and failure is assumed to occur when the cumulative degradation exceeds the critical threshold $x_c = 40.0$. The inspection, preventive replacement, and reactive replacement penalty costs are $c_0 = 1.0$, $c_1 = 10.0$ and $c_2 = 100.0$, respectively. In what follows, optimal replacement policies for the component are derived based on the observability of the environment states for values of $\sigma_\epsilon \in \{0.25, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0\}$. To solve each instance numerically, the number of particles used is $\widehat{N} = 500$, and we let $\mathcal{G}$ be the family of truncated normal densities on $[0, x_c]$, where $\Theta = \{(\mu, \sigma) \in (-\infty, \infty) \times (0, \infty)\}$. (It is important to not confuse σ, a component in the parameterized belief state of degradation, with σ_ϵ, the degree of "noise" in degradation observations.) For $\theta \in \Theta$, the p.d.f. is given by

$$g(x; \theta) = \frac{\sigma^{-1}\phi\left(\dfrac{x-\mu}{\sigma}\right)}{\phi\left(\dfrac{x_c-\mu}{\sigma}\right) - \phi\left(\dfrac{\mu}{\sigma}\right)}, \tag{4.13}$$

where ϕ is the standard normal density. In the case of the PEID model, policies are obtained on a discretized belief space $\dot{B}_\Theta = \dot{\Theta} \times \dot{\Pi}$, where

$$\begin{aligned}
\dot{\Theta} &= \begin{cases}
\{(\dot{\mu}, \dot{\sigma}) \in \{0.2, 0.4, \ldots, 38.6\} \times \{0.04, 0.08, \ldots, 0.4\}\}, & \sigma_\epsilon = 0.25, \\
\{(\dot{\mu}, \dot{\sigma}) \in \{0.2, 0.4, \ldots, 38.6\} \times \{0.1, 0.2, \ldots, 1.0\}\}, & \sigma_\epsilon = 0.5, \\
\{(\dot{\mu}, \dot{\sigma}) \in \{0.2, 0.4, \ldots, 38.6\} \times \{0.1, 0.4, \ldots, 2.8\}\}, & \sigma_\epsilon \in \{1.0, 2.0\}, \\
\{(\dot{\mu}, \dot{\sigma}) \in \{0.2, 0.4, \ldots, 38.6\} \times \{1.0, 1.3, \ldots, 3.7\}\}, & \sigma_\epsilon \in \{3.0, 4.0, 5.0\},
\end{cases} \\
\dot{\Pi} &= \{\pi_s, [1/2\ 1/2\ 0], [0\ 1/2\ 1/2], [1/2\ 0\ 1/2]\},
\end{aligned}$$

and $L = 7,720$. The OEID model is solved numerically using the same parameters for σ_ϵ and $\dot{\Theta}$ as in the PEID case but with $\dot{\Pi} \in \Pi_-$ so that $L = 5,790$.

Table 11: Policy cost per unit time.

$\sigma_\epsilon \rightarrow$	0.25	0.5	1.0	2.0	3.0	4.0	5.0
γ (PEID)	1.5503	1.5586	1.5673	1.5844	1.5968	1.6058	1.6124
γ (OEID)	1.5328	1.5378	1.5434	1.5489	1.5570	1.5620	1.5677

The optimal policy costs of the PEID model for each value of σ_ϵ are shown in Table 11, and Figures 35, 36, and 37 display the optimal policies for $\sigma_\epsilon \in \{1.0, 3.0, 5.0\}$, respectively, where each policy is partitioned into separate plots by $\dot{\pi} \in \dot{\Pi}$. As σ_ϵ increases, the PEID policy cost

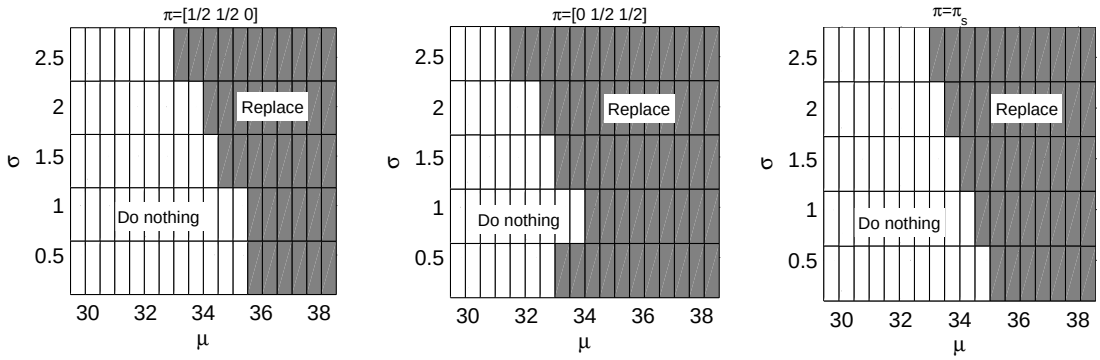

Figure 35: Optimal policy (PEID model, $\sigma_\epsilon = 1.0$)).

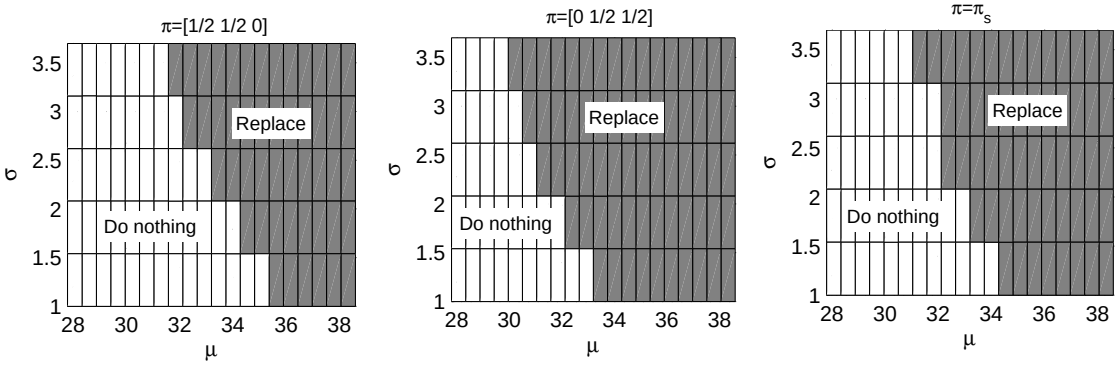

Figure 36: Optimal policy (PEID model, $\sigma_\epsilon = 3.0$)).

increases as the true degradation level of the component becomes more difficult to discern, and the policies react by favoring preventive replacement at smaller values of μ. In addition, when it is likely the environment is in a high-degradation state, the policies tend to replace for smaller μ due to the relatively high likelihood that cumulative degradation will exceed the critical threshold in the next decision period. The policy structure with respect to σ is more complex due to the

characteristics of the truncated normal distribution used in the density projection. In fact, for a fixed μ, the component's survival probability and expected survival time over the next period are not necessarily monotone with respect to σ. Therefore, it is possible that the preventive replacement boundary with respect to μ is not monotone in σ. This behavior is especially pronounced when the probability of being in a state with a high degradation rate is large. Next we consider the case when

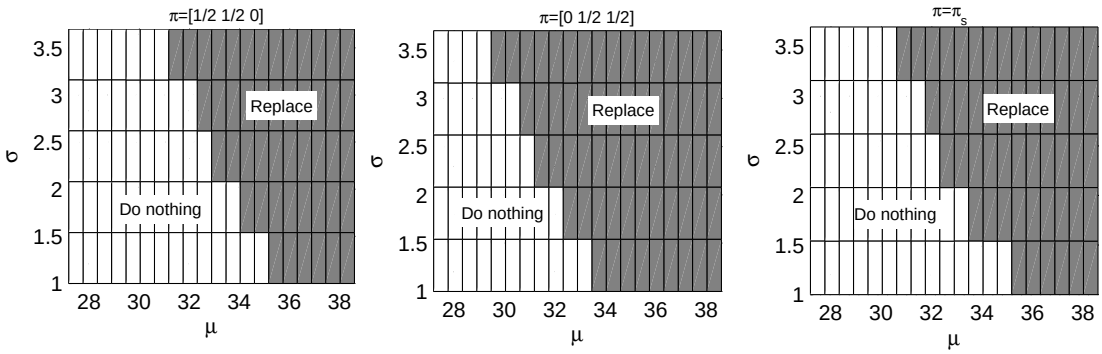

Figure 37: Optimal policy (PEID model, $\sigma_\epsilon = 5.0$)).

the environment is observable (OEID). The cost of each optimal policy is shown in Table 11. The fact that the average cost of each OEID policy for a given σ_ϵ is less than that of the corresponding PEID policy is consistent with the notion that having perfect information about the environment leads to better replacement decisions. The optimal policies for $\sigma_\epsilon \in \{1.0, 3.0, 5.0\}$ are shown in Figures 38, 39, and 40, respectively. Each policy exhibits a strong structural similarity to those of the PEID model. The trends to replace for smaller μ when σ_ϵ is relatively large or when the degradation rate of the current environment state is high are still present as well as the tendency to preventively replace for a fixed μ only for large or small values of σ. Plots of the the average policy costs of the PEID and OEID models as a function of σ_ϵ are shown together in Figure 41. For purposes of comparison, the optimal policy cost ($\gamma = 1.5389$) for a replacement model in a PEOD environment is also shown on the plot. As σ_ϵ approaches zero, the cost of the PEID policy decreases toward that of the PEOD policy; whereas, the cost of the OEID policy is already less than that of the PEOD policy for $\sigma_\epsilon \leq 0.5$. Therefore, perfect knowledge of the operating environment state does in fact confer a noticeable benefit to the resulting replacement policy. The age- and

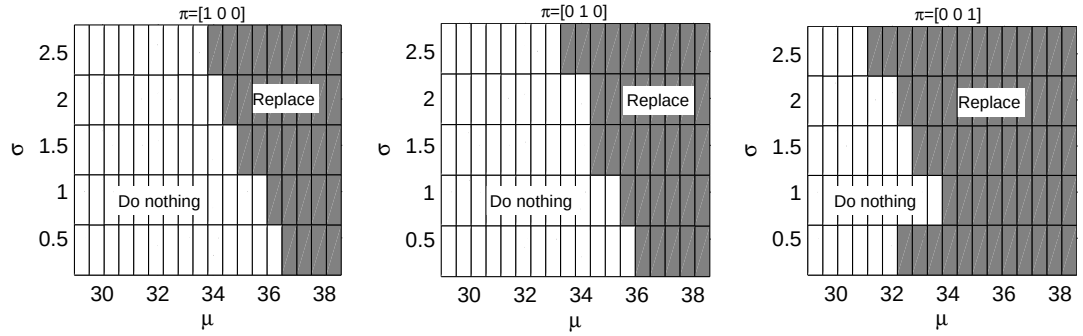

Figure 38: Optimal policy (OEID model, $\sigma_\epsilon = 1.0$)).

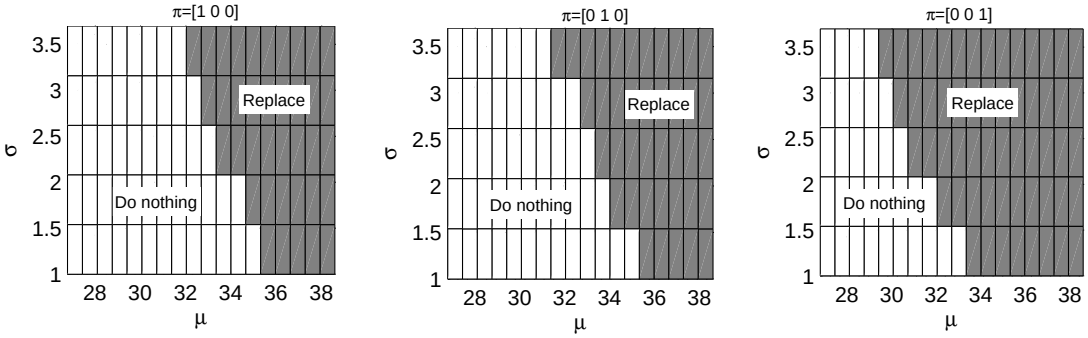

Figure 39: Optimal policy (OEID model, $\sigma_\epsilon = 3.0$)).

reactive-replacement policies have costs of 2.8058 and 5.9225, respectively. The costs of these two policies do not depend on σ_ϵ as the component's lifetime distribution is determined by $(\boldsymbol{Q}, \boldsymbol{r})$ only. All three POMDP policies outperform the age- and reactive replacement policies by at least 42.5% and 72.8%, respectively, over the range of σ_ϵ.

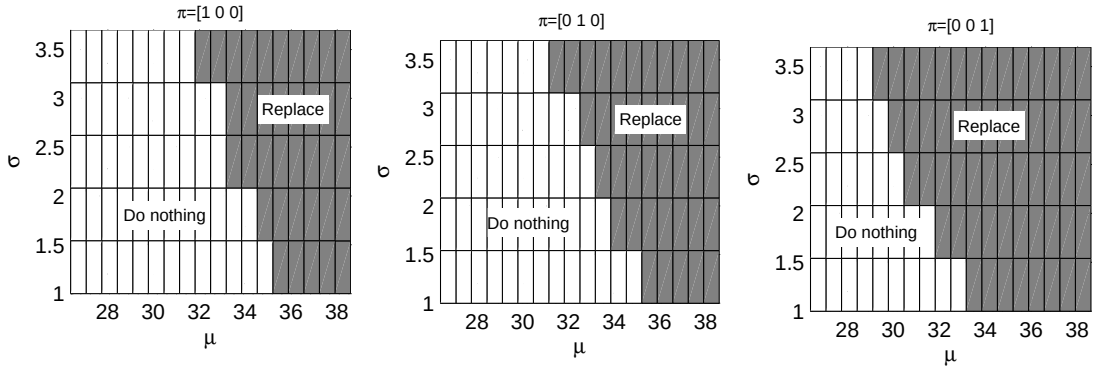

Figure 40: Optimal policy (OEID model, $\sigma_\epsilon = 5.0$)).

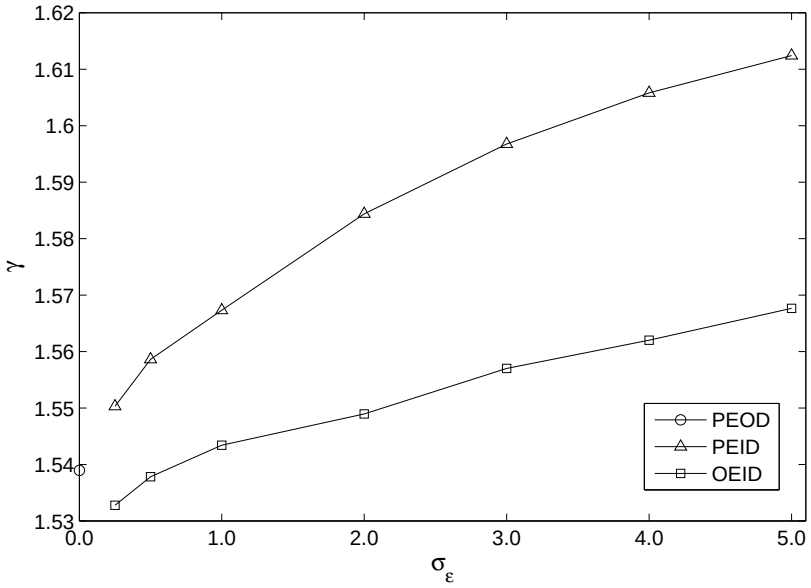

Figure 41: Average policy costs (PEOD, PEID, and OEID models)).

4.3.2 A Wind Turbine Gear Replacement Example

This example considers the optimal replacement of a gear operating in a wind turbine's drive train. Gear tooth degradation is measured as the effective number of load cycles, which is computed as a function of shaft torque and rotor speed. Failure is assumed when the effective number of

load cycles exceeds $x_c = 10^7$. To estimate the environment parameters, a degradation path is simulated to failure using the approach outlined in Example 4 of Section 2.4, where the S-N curve is parameterized by the values in Case (iii). The environment parameters $\boldsymbol{r}$, $\boldsymbol{Q}$, and ℓ are estimated using the MCMC-based inference procedure of Section 2.3 to be as follows:

$$\boldsymbol{Q} = \begin{bmatrix} -0.0131 & 0.0081 & 0.0050 \\ 0.0443 & -0.0836 & 0.0393 \\ 0.0276 & 0.0159 & -0.0435 \end{bmatrix}, \qquad \boldsymbol{r} = 10^4 \times \begin{bmatrix} 0.1327 & 1.1231 & 1.6746 \end{bmatrix}.$$

Due to errors in both the observed data as well as deviations from the physics-based wear model, it is assumed that there are various degrees of uncertainty in assessing the true number of effective load cycles imposed on a gear tooth. To account for this uncertainty, the computed number of load cycles is assumed to be normally distributed with a mean corresponding to the actual number of cycles with variance σ_ϵ^2. The normality of error is typically a valid assumption in practice where the error can be considered as the sum of a large number of independent error terms which converge in distribution to a normal random variable. The inspection interval is $\delta = 100.0$ minutes, and the inspection, preventive replacement, and reactive replacement penalty costs are assumed to be $c_0 = 1.0$, $c_1 = 100.0$ and $c_2 = 200.0$, respectively. The large penalty for reactive replacements is chosen to reflect the extra expenses incurred in the logistics and lost production associated with a sudden replacement.

Table 12: Policy cost per unit time (gear replacement example).

$\sigma_\epsilon \rightarrow$	2.5×10^4	7.5×10^4	2.5×10^5
γ (PEID)	0.07075	0.07885	0.08896
γ (OEID)	0.06894	0.07634	0.08235

Optimal policies are obtained for both the PEID and OEID models when $\sigma_\epsilon \in \{2.5 \times 10^4, 7.5 \times 10^4, 2.5 \times 10^5\}$ by projecting the belief space onto the truncated family of truncated normal densities (4.13). For the PEID model, $\widehat{N} = 500$ and the discretized belief space is $\dot{B}_\Theta = \dot{\Theta} \times \dot{\Pi}$, where

$$\begin{aligned} \dot{\Theta} &= \left\{(\dot{\mu}, \dot{\sigma}) \in \{1.0 \times 10^5, 6.0 \times 10^5, \ldots, 98.6 \times 10^5\} \times \{0.1\,\sigma_\epsilon, 0.2\,\sigma_\epsilon, \ldots, \sigma_\epsilon\}\right\}, \\ \dot{\Pi} &= \{\pi_s, [1/2\ 1/2\ 0], [0\ 1/2\ 1/2]\}, \end{aligned}$$

and $L = 5,940$. The costs for optimal policies of the PEID model are shown in Table 12 for each value of σ_ϵ. As in the previous example, the policy cost increases with σ_ϵ as the true degradation

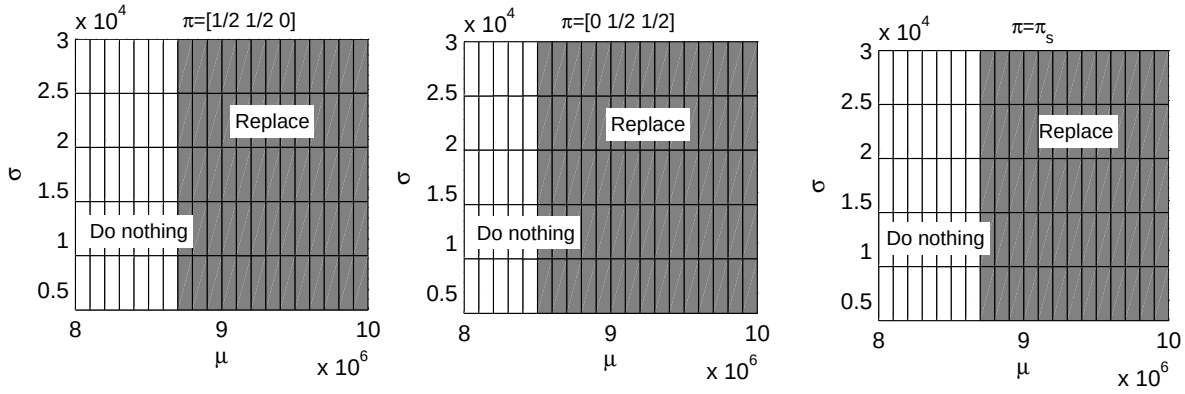

Figure 42: Optimal policy (PEID model, $\sigma_\epsilon = 2.5 \times 10^4$)).

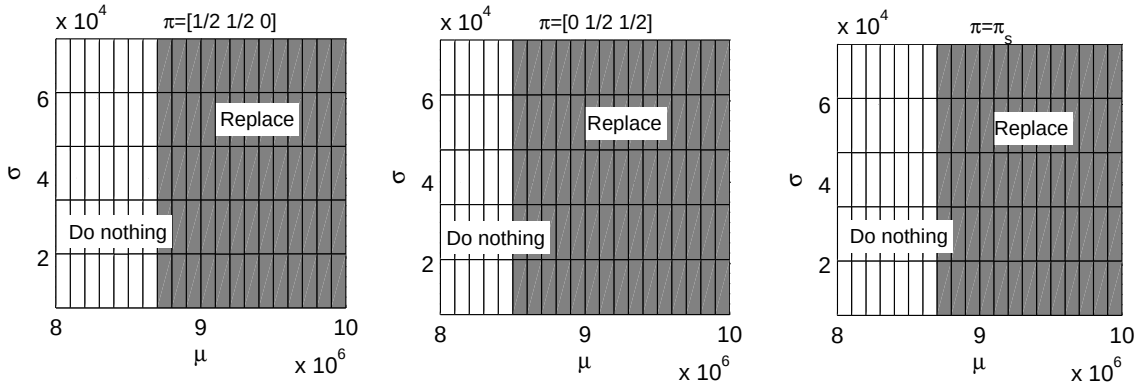

Figure 43: Optimal policy (PEID model, $\sigma_\epsilon = 7.5 \times 10^4$)).

level of the gear becomes less discernable. The associated optimal policies for $\sigma_\epsilon \in \{1.0, 3.0, 5.0\}$ are shown in Figures 42, 43, and 44, respectively. The figures indicate the same general characteristics of policy structure observed in the previous example. The policies favor replacing for smaller μ when it is likely the environment is in a high-degradation state due to the corresponding higher likelihood that cumulative degradation will exceed the critical threshold in the next decision period. Also, as

σ_ϵ increases, the optimal policies tend to favor preventive replacement for smaller μ reflecting that increased uncertainty leads to more conservative replacement decisions. Some of the same complex behavior that results from effects of the truncated normal distribution are also observed as in the last example, though the policies are relatively insensitive to σ.

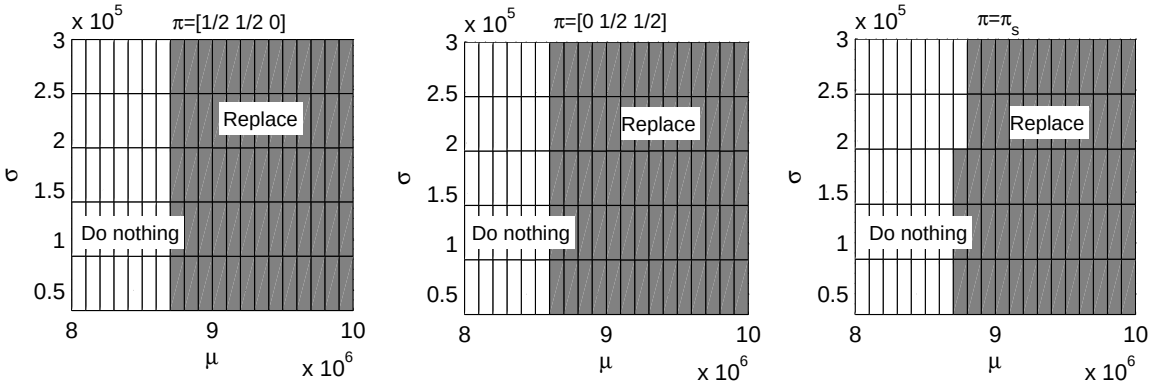

Figure 44: Optimal policy (PEID model, $\sigma_\epsilon = 2.5 \times 10^5$)).

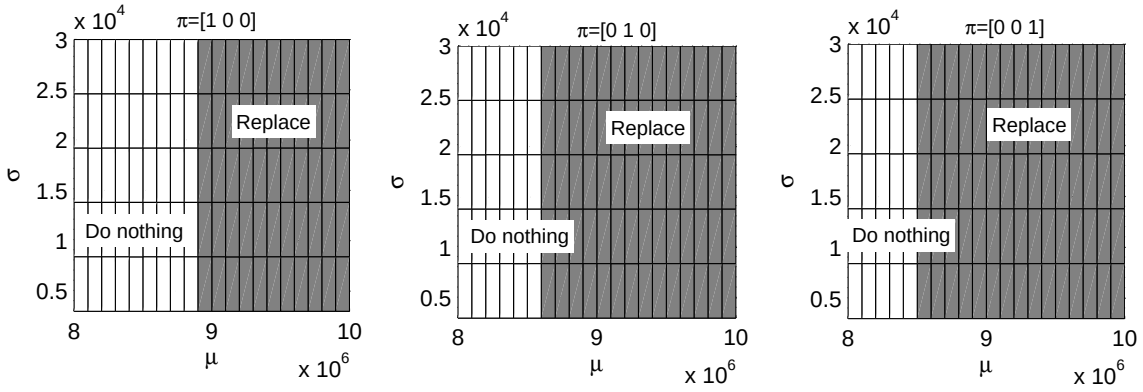

Figure 45: Optimal policy (OEID model, $\sigma_\epsilon = 2.5 \times 10^4$)).

The OEID model is solved numerically using $\widehat{N} = 750$ particles on $\dot{B}_{\Theta\_} = \dot{\Theta} \times \dot{\Pi}_{\_}$, where $L = 4,455$. Policy costs are shown in Table 12 and are lower than those of the PEID model for corresponding σ_ϵ. Figures 45, 46, and 47 show the optimal policies for $\sigma_\epsilon \in \{2.5 \times 10^4, 7.5 \times 10^4, 2.5 \times 10^5\}$, respectively. A comparison plot of the average costs of the PEID and OEID policies

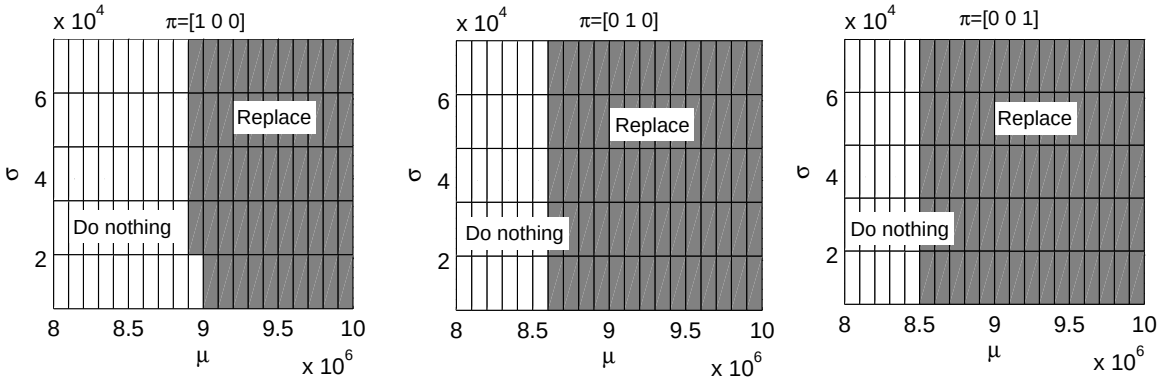

Figure 46: Optimal policy (OEID model, $\sigma_\epsilon = 7.5 \times 10^4$)).

for various σ_ϵ are shown together in Figure 48 along with the cost of a policy under the PEOD assumption ($\gamma = 0.06649$). As in the previous example, the PEID policy tends to outperform the OEID policy, particularly for large values of σ_ϵ. This result suggests that environment information becomes more valuable as the uncertainty in the degradation level increases.

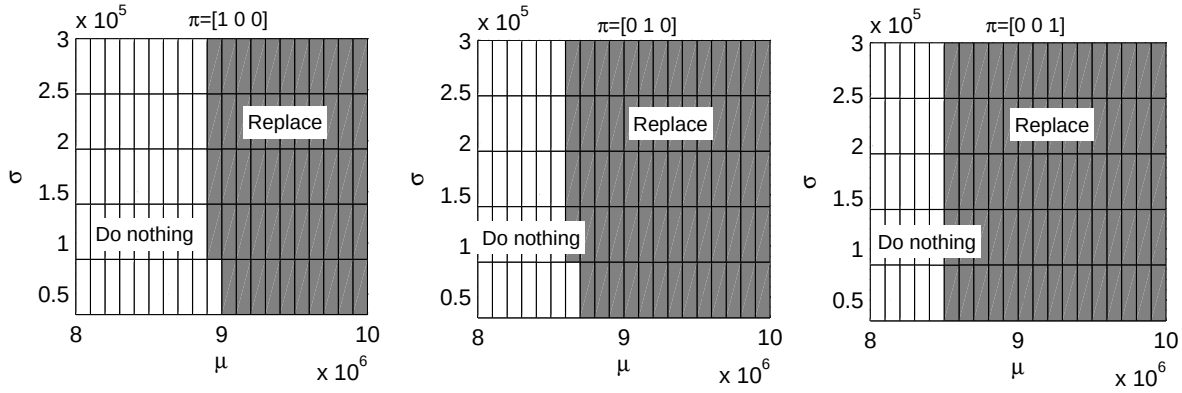

Figure 47: Optimal policy (OEID model, $\sigma_\epsilon = 2.5 \times 10^5$)).

The costs of the three POMDP models are much lower than the costs of the age- and reactive-replacement policies, which are 0.1022 and 0.1589, respectively, for moderate values of σ_ϵ. However, when the degradation uncertainty is high, the age-replacement policy is more competitive with the POMDP policies, though the POMDP policies still have lower costs. The POMDP policies outperform the age- and reactive-replacement policies by at least 22.9% and 50.4%, respectively, for $\sigma_\epsilon = 7.5 \times 10^4$ and by at least 13.0% and 44.0%, respectively, for $\sigma_\epsilon = 2.5 \times 10^5$.

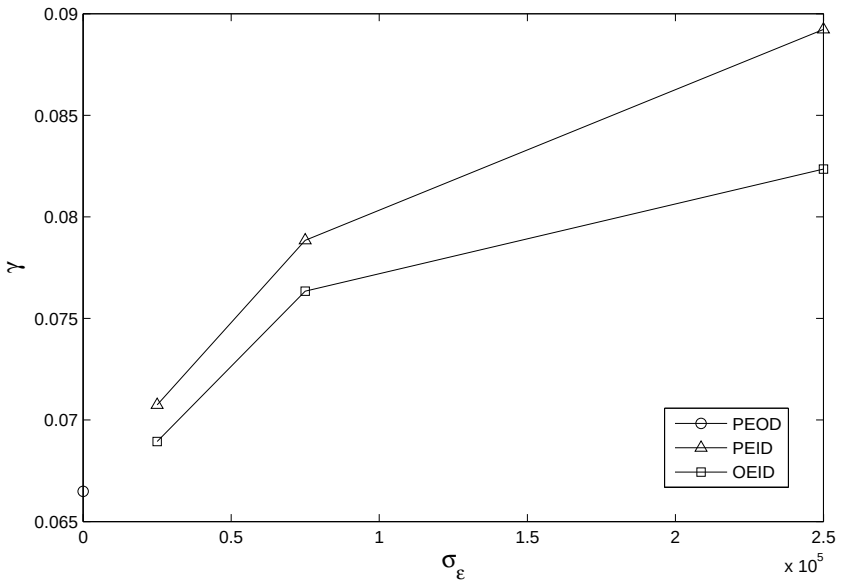

Figure 48: Average gear replacement policy costs (PEOD, PEID, and OEID models).

4.3.3 A Wind Turbine Shaft Bearing Replacement Example

Now replacement policies are obtained for a wind turbine's shaft bearing where degradation is again assessed as the effective number of load cycles. As in the previous example, it is assumed that there is uncertainty in both assessing the state of the operating environment as well as computing the actual number of cycles imposed on the bearing. This estimation error is similarly assumed to be the sum of a large number of independent error terms and will be treated as unbiased and normally distributed with variance σ_ϵ^2. A degradation signal was simulated from temperature and rotor speed data using the approach of Example 5 in Section 2.4 with the life adjustment factor $a_2(\rho)$. Using the MCMC-based inference procedure of Section 2.3, the operating environment is estimated to have $\ell = 2$ states and the following parameters:

$$\boldsymbol{Q} = \begin{bmatrix} -0.0089 & 0.0089 \\ 0.0304 & -0.0304 \end{bmatrix}, \qquad \boldsymbol{r} = 10^3 \times \begin{bmatrix} 0.7580 & 2.2585 \end{bmatrix}.$$

Replacement policies are obtained assuming that $x_c = 10^7$ cycles, $\delta = 300$ minutes, $c_0 = 1.0$, $c_1 = 100.0$, $c_2 = 200.0$, and $\sigma_\epsilon \in \{5.0 \times 10^3, 1.75 \times 10^4, 7.5 \times 10^4\}$.

The belief space is projected onto the truncated family of normal densities (4.13). To obtain policies for the PEID instances, $\hat{N} = 100$ particles and $\dot{B}_\Theta = \dot{\Theta} \times \dot{\Pi}$, where

$$\begin{aligned} \dot{\Theta} &= \left\{(\dot{\mu}, \dot{\sigma}) \in \{5.0 \times 10^4, 10.0 \times 10^4, \ldots, 9.75 \times 10^6\} \times \{0.1\,\sigma_\epsilon, 0.2\,\sigma_\epsilon, \ldots, \sigma_\epsilon\}\right\}, \\ \dot{\Pi} &= \{\pi_s, [1\ 0], [0\ 1]\}, \end{aligned}$$

and $L = 5,850$. The OEID policies are obtained by setting $\hat{N} = 150$ particles and $\dot{B}_{\Theta\_} = \dot{\Theta} \times \dot{\Pi}_-$, where $L = 3,900$.

Table 13: Policy cost per unit time (shaft bearing replacement example).

$\sigma_\epsilon \rightarrow$	7.5×10^3	1.75×10^4	7.5×10^4
γ (PEID)	0.01525	0.01579	0.01623
γ (OEID)	0.01525	0.01567	0.01600

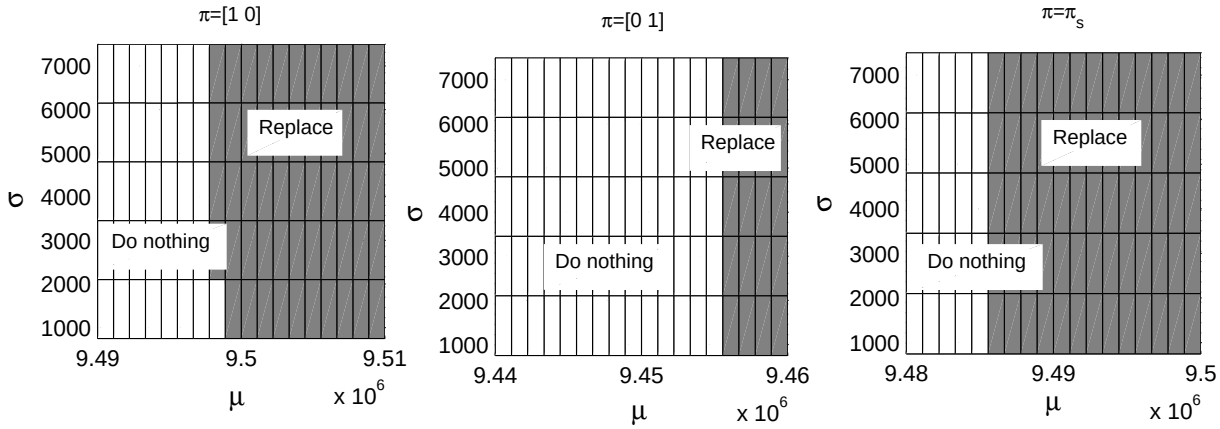

Figure 49: Optimal policy (PEID model, $\sigma_\epsilon = 7.5 \times 10^3$)).

Table 13 shows the optimal policy costs for the PEID and OEID models for each σ_ϵ, and the same relationships between the policy costs can be observed as in the previous examples. For a given σ_ϵ, the PEID policy cost is greater than that of the corresponding OEID policy, and the cost of both models' policies increase with σ_ϵ. The optimal policies for the PEID models for $\sigma_\epsilon \in \{7.5 \times 10^3, 1.75 \times 10^4, 7.5 \times 10^4\}$ are shown in Figures 49, 50, and 51, respectively. In contrast to the gear replacement example, the boundary of the preventive replacement region clearly

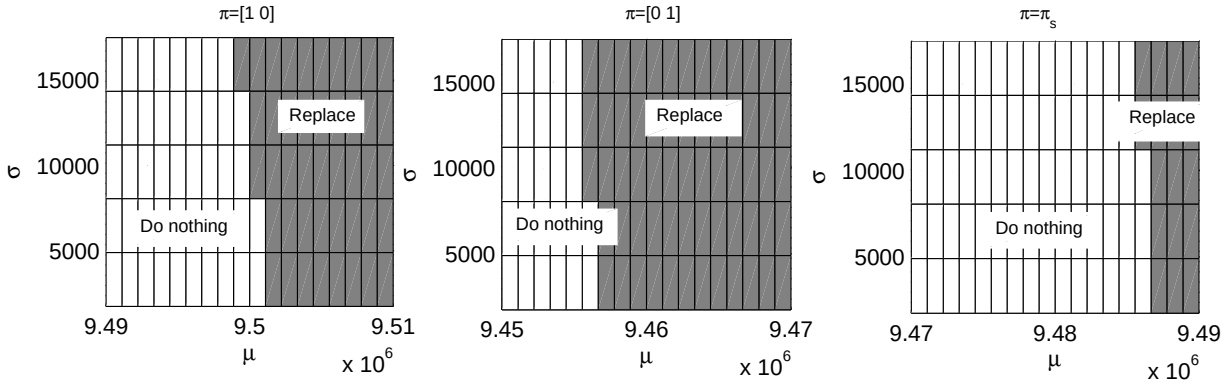

Figure 50: Optimal policy (PEID model, $\sigma_\epsilon = 1.75 \times 10^4$)).

has a strong dependence on both μ and σ within each environment belief state, which becomes especially pronounced when σ_ϵ is large. Specifically, when there is significant uncertainty in the degradation level, the policies behave quite conservatively by preventively replacing for relatively small μ. Similar structural behavior as in the previous examples, such as the tendency to replace for smaller μ when the current degradation rate is relatively high, is also present. The OEID policies shown in Figures 52, 53, and 54 for $\sigma_\epsilon \in \{7.5 \times 10^3, 1.75 \times 10^4, 7.5 \times 10^4\}$, respectively, are nearly identical to the PEID policies for corresponding environment states.

Figure 55 is a plot of the average cost of the PEID and OEID policies for various σ_ϵ as well as the cost of the PEOD policy ($\gamma = 0.014896$). In contrast to the other examples, the average cost of the PEID and OEID policies are approximately equal for small σ_ϵ indicating that perfect knowledge of the environment state does not improve replacement decisions markedly when the uncertainty of the cumulative degradation level is small. The reduced value of environment information is likely due to the relatively long inspection interval coupled with the fact that environment consists of only two states. As a consequence, regardless of the initial environment state when a decision is made, the environment process will be "well mixed" prior to the next decision, and the degradation accrued during the period will be relatively insensitive to the initial environment state. However, in cases where there is a large degree of uncertainty in the degradation level, perfect knowledge of the exact environment state could lead to avoidance of a reactive failure or premature

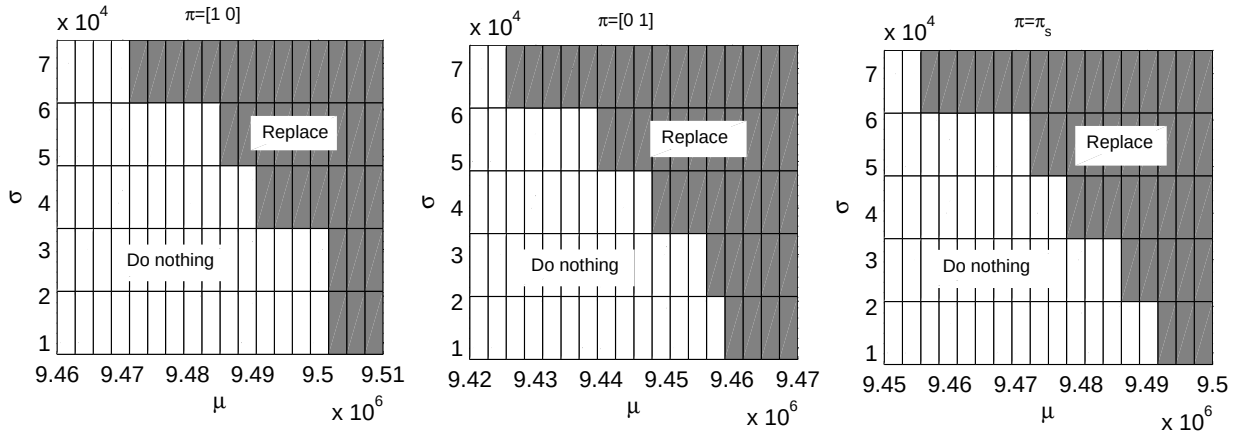

Figure 51: Optimal policy (PEID model, $\sigma_\epsilon = 7.5 \times 10^4$)).

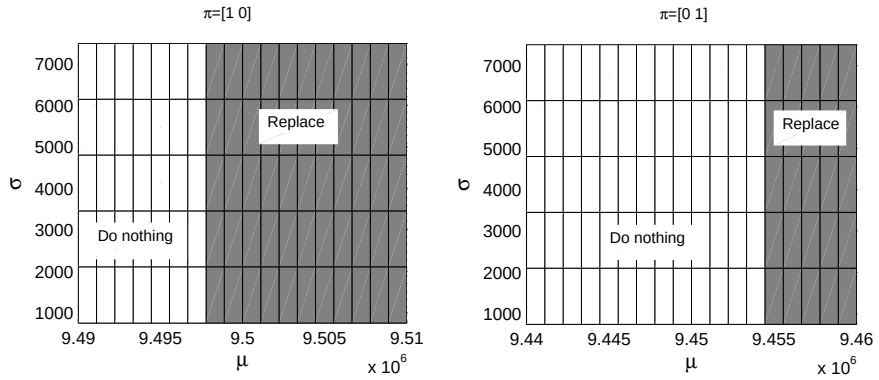

Figure 52: Optimal policy (OEID model, $\sigma_\epsilon = 7.5 \times 10^3$)).

preventive replacement. The costs of the age- and reactive-replacement policies are 0.0158 and 0.0362, respectively. All three POMDP policies greatly outperform the reactive policy; however, for larger values of σ_ϵ, the age-replacement policy slightly outperforms both the PEID and OEID policies. The diminished value of environment information in this replacement example causes the performance of the POMDP policy to be more dependent on accurate information about the component's degradation level. As σ_ϵ increases and the degradation level becomes relatively un-

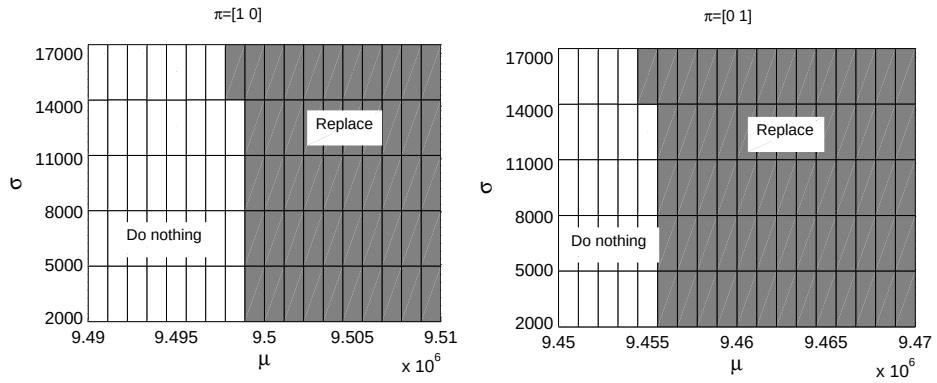

Figure 53: Optimal policy (OEID model, $\sigma_\epsilon = 1.75 \times 10^4$)).

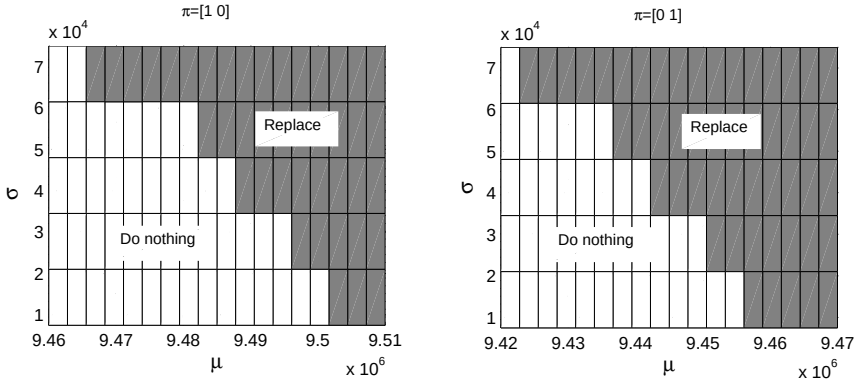

Figure 54: Optimal policy (OEID model, $\sigma_\epsilon = 7.5 \times 10^4$)).

certain, the POMDP policies become less competitive with the age-replacement policy, and when degradation is relatively certain, the POMDP policies are more competitive. For $\sigma_\epsilon = 7.5 \times 10^3$, the POMDP policies outperform the age- and reactive-replacement policies by at least 3.5% and 57.9%, respectively.

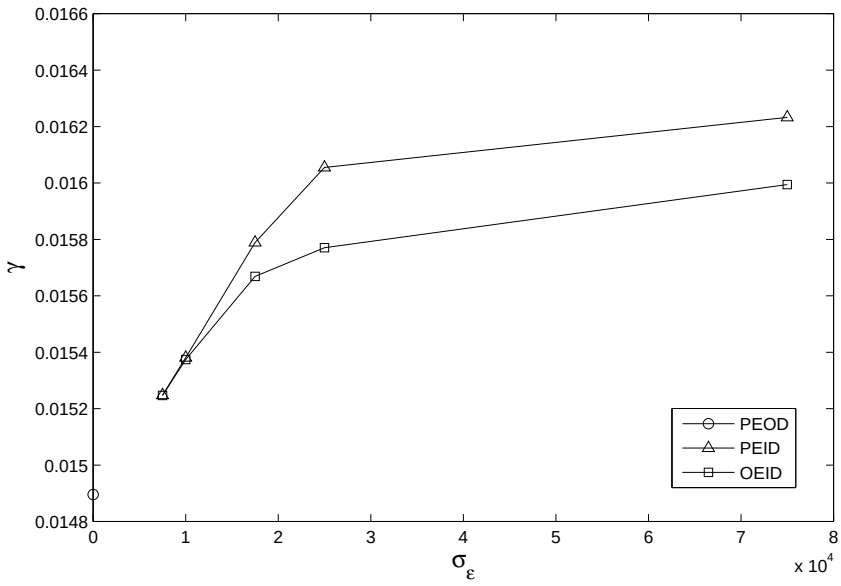

Figure 55: Average shaft bearing replacement policy costs (PEOD, PEID, and OEID models).

In this chapter, a POMDP model to optimally replace a single wind turbine component was described that incorporates the uncertainty in determining the component's degradation level from a surrogate signal. The replacement problem was solved numerically using particle filtering and belief projection techniques to obtain approximate, optimal policies. In the next chapter, we summarize the results of this dissertation and propose directions for future research.

5.0 CONCLUSIONS AND FUTURE RESEARCH

Wind energy is poised to play a crucial role in the future energy portfolio of the United States and the world at large. This dissertation has focused on reducing the operating and maintenance costs of wind energy by developing a decision-framework for optimally replacing a wind turbine component. This framework required development of a mathematical model of the environment-driven degradation of a wind turbine component whose parameters can be estimated from a surrogate signal of the component's degradation. The estimated parameters were then used in POMDP models to optimally replace a single wind turbine component under different observability assumptions for the environment and the component's degradation level.

In Chapter 2, a hybrid analytical-statistical framework was presented to approximate the stochastic, environment-driven, degradation process of a wind turbine component from an observed signal of degradation. The signal was assumed to evolve as a switching diffusion process, and the environment parameters were estimated using a Markov chain Monte Carlo (MCMC) statistical procedure. The framework's performance was evaluated by comparing the expected component lifetime and mean degradation signal computed from the estimated parameters with the actual component lifetime and the observed degradation signal, respectively. Numerical results indicated that the framework was robust against signals that departed significantly from the stochastic behavior of a switching diffusion process. For degradation processes driven by non-Markovian environments, non-constant degradation rates, or time-inhomogeneous variance, the framework was still able to characterize the degradation environment well, provided the signal is observed for a sufficiently long time period.

Although the framework in Chapter 2 provides a viable approach to approximate the stochastic behavior of an arbitrary degradation process, it is limited by assumptions required for the switching diffusion model. In particular, when the variance of the degradation signal varies significantly over non-overlapping time intervals, substantial preprocessing of the signal is required to estimate the

environment parameters. To more effectively handle such cases, the switching diffusion model must incorporate a time-inhomogeneous diffusion coefficient, and an appropriate inference procedure is required for the new model. Similarly, the estimation procedure does not perform as well in applications such as crack propagation, where degradation signals evolve as nearly-deterministic, nonlinear functions of time and violate the assumption that the environment process is time-homogeneous. Although signal preprocessing can be helpful, a more realistic model should include degradation rates that are functions of time and/or the cumulative degradation level. Lastly, the MCMC statistical procedure is not ideal for real-time updating of the environment parameters as the signal is observed. In fact, there are no known filtering- or simulation-based procedures that are able to infer the parameters of a switching diffusion process in a real-time, iterative manner. Developing such an inference procedure would significantly reduce the computational burden and allow near-continuous updating of the model parameters as sensor data are collected.

Using the environment-driven degradation model of Chapter 2, two wind turbine component replacement models were presented in Chapter 3 that assume the environment state is partially-observable. The models were formulated using a POMDP framework with the objective of minimizing the long-run average replacement cost per unit time. Different assumptions for the replacement of a component were considered in each model. The first model assumed instantaneous replacements with fixed, deterministic costs; whereas, the second model assumed that replacements require a fixed, deterministic time period for completion during which downtime costs accrue as a function of the environment. The optimal policy of the first model was shown to have a threshold-type structure with respect to the component's cumulative degradation level and, in special cases, with respect to the environment belief state. However, optimal replacement thresholds for the second model only exist with respect to the component's cumulative degradation level when special conditions are satisfied. The results of numerical experiments indicated that in both models the replacement thresholds can differ substantially between environment belief states. Therefore, condition monitoring is critical so that the most informative belief state, and its associated degradation threshold, can be discerned prior to a replacement decision. The value of this environment information became evident when the optimal policy of the first model was shown to outperform both age- and reactive-replacement policies.

The results obtained for the replacement models in Chapter 3 indicate several avenues for future research. The first is obtaining more detailed analytical results for each model to include a precise description of the optimal policy structure and tighter bounds on the optimal policy cost

as well as determining whether the preventive replacement regions are convex and/or connected. Other avenues include extensions of the replacement models that enhance realism. Although the notion of stochastic downtime costs is explicitly considered in this research, the total time to complete a replacement is assumed to be deterministic. In practice there are significant safety limits on the wind speed and weather conditions that must be satisfied in order for work crews to conduct replacement operations. These limits are particularly critical for offshore wind turbines, which cannot be accessed during periods when the significant wave height of the surrounding water exceeds a certain level. Therefore, including replacement interruptions and/or delays that depend on the environment state would greatly enhance realism.

In Chapter 4, the replacement problem was extended to the case where cumulative degradation is imperfectly observed according to a parameterized probability density function (p.d.f.). It was assumed that the environment is partially-observable and that component replacements are instantaneous with fixed, deterministic costs. The problem was formulated using a POMDP model to minimize the long-run average replacement cost per unit time. The POMDP was shown to be analytically intractable, and a simplified instance of the model was formulated in which the belief state is updated using only the previous signal observation. Optimal thresholds were shown to exist with respect to the observed signal level in the simplified model if the p.d.f. of the observed signal satisfies a technical condition. To solve the POMDP numerically, a known projection-filtering algorithm was extended to the mixed state space induced by both partially-observable and observable environments. Comprehensive numerical experiments illustrated optimal replacement policies for both a wind turbine drivetrain gear and a shaft bearing. The results show that the policy cost decreases with the level of uncertainty in both the environment state and the cumulative degradation level so that utilizing a condition-based information is clearly valuable. However, in environments whose distributions are nearly stationary before each decision epoch, the environment information becomes less valuable as the uncertainty in the degradation level is reduced. For such environments, the POMDP policies outperform age-replacement policies when degradation is relatively certain. However, as the uncertainty in degradation increases, the POMDP policies become less competitive with age-replacement policies. In such instances, it may be advantageous for wind farm operators to invest more heavily in sensors that discern the degradation level as opposed to the environment state.

Despite realistically incorporating uncertainty in the degradation level of a wind turbine component, the models of Chapter 4 must be extended in a similar manner to those of Chapter 3 to

incorporate stochastic downtime costs and replacement times as well as replacement interruptions and/or delays that depend on the environment. Obtaining structural results for those cases is beneficial from a computational point of view and in practice, as the policies can be communicated to wind farm operators as a set of elementary rules. Since the replacement models of Chapter 4 are based on sensor readings, which are obtained for multiple wind turbines and component types, it will be useful to consider a group replacement policy, in which multiple components are simultaneously monitored and replaced. A group replacement policy is particularly advantageous for large wind farms where there may be a substantial fixed cost associated with transporting the necessary equipment and work crews to the turbine location.

The models proposed in this dissertation constitute a first step in developing a comprehensive decision-making framework for replacing a wind turbine component. Although the assumptions for the environment-driven degradation model seem restrictive, the dissertation results highlight the flexibility of the model at estimating various degradation processes. For the framework to be practicable, the computational burden of estimating the degradation model's parameters must be reduced. The POMDP replacement models incorporate some realistic features for the uncertainty in discerning both the environment state and degradation level, as well as the opportunity cost of lost energy production. However, in practice there are many other considerations for replacement decisions that are not featured in the models (e.g. repair interruptions and/or delays, component lead times, and group replacement strategies). By incorporating more realistic features into the degradation and replacement models and reducing their computational burden, operators will have a viable tool for making effective turbine maintenance decisions.

BIBLIOGRAPHY

[1] 20 percent wind energy by 2030: Increasing wind energy's contribution to U.S. electricity supply. Technical Report DOE/GO-102008-2567, U.S. Department of Energy (DOE), Washington, D.C., July 2008.

[2] Energy technology perspectives: Scenarios and strategies to 2020. Technical report, International Energy Agency (IEA), Paris, France, 2008.

[3] Technology roadmap: Wind energy. Technical report, International Energy Agency (IEA), Paris, France, 2009.

[4] ADAPT.wind cumulative impulse fact sheet. Technical Report GEA19503, General Electric Corporation, Fairfield, CT, October 2011.

[5] A national offshore wind strategy: Creating an offshore wind energy industry in the united states. Technical Report DOE/EE-0798, U.S. Department of Energy (DOE), Washington, D.C., February 2011.

[6] J. Abate and W. Whitt. Numerical inversion of Laplace transforms of probability distributions. *ORSA Journal on Computing*, 7(1):36–43, 1995.

[7] H. Arabian-Hoseynabadi, P. Tavner, and H. Oraee. Reliability comparison of direct-drive and geared-drive wind turbine concepts. *Wind Energy*, 13(1):62–73, 2010.

[8] M. B. B. Jakoby, M. Scherer and H. Eisenschmid. An automotive engine oil viscosity sensor. *IEEE Sensors Journal*, 3(5):562–568, 2003.

[9] F. Ball, Y. Cai, and A. O'Hagen. MCMC for hidden, continuous-time Markov chains. Technical report, University of Nottingham, Nottingham, UK, November 1997.

[10] D. Banjevic, A. Jardine, V. Makis, and M. Ennis. A control-limit policy and software for condition-based maintenance optimization. *INFOR*, 39:32–50, 2001.

[11] E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics*, 37(6):1554–1563, 1966.

[12] P. Billingsley. Statistical methods in Markov chains. *The Annals of Mathematical Statistics*, 32(1):12–40, 1961.

[13] P. Blackwell. Bayesian inference for Markov processes with diffusion and discrete components. *Biometrika*, 90(3):613–627, 2003.

[14] M. Bladt, B. Meini, M. Neuts, and B. Sericola. Distributions of reward functions on continuous-time Markov chains. In *Proceedings of the Fourth International Conference on Matrix Analytic Methods in Stochastic Models*, Adelaide, Australia, 2002.

[15] E. Byon and Y. Ding. Season-dependent condition-based maintenance for a wind turbine using a partially observed Markov decision process. *IEEE Transactions on Power Systems*, 25(4):1823–1834, 2010.

[16] E. Byon, L. Ntaimo, and Y. Ding. Optimal maintenance strategies for wind turbine systems under stochastic weather conditions. *IEEE Transactions on Reliability*, 59(2):393–404, 2010.

[17] E. Byon, L. Ntaimo, C. Singh, and Y. Ding. Wind energy facility reliability and maintenance. In Pardalos, Pereira, Rebennack, and Boyko, editors, *Handbook of Wind Power Systems: Optimization, Modeling, Simulation and Economic Aspects*. In press.

[18] E. Byon, E. Perez, Y. Ding, and L. Ntaimo. Simulation of wind farm operations and maintenance using discrete event system specification. *Simulation*, 87(12):1091–1115, 2011.

[19] O. Cappé, E. Moulines, and R. Rydén. *Inference in Hidden Markov Models*. Springer, New York, NY, 2005.

[20] J. Carpenter, P. Clifford, and P. Fearnhead. Improved particle filter for nonlinear problems. *IEE Proceedings on Radar, Sonar, and Navigation*, 146(1):2–7, 1999.

[21] G. Casella and E. George. Explaining the Gibbs sampler. *The American Statistician*, 46(3):167–174, 1992.

[22] E. Çinlar. Shock and wear models and Markov additive processes. In I. Shimi and C. Tsokos, editors, *The Theory and Applications of Reliability*, pages 193–214. Academic Press, New York, NY, 1977.

[23] E. Çinlar. Markov and semi-Markov models of deterioration. In M. Abdel-Hameed, E. Çinlar, and J. Quinn, editors, *Reliability Theory and Models*, pages 3–41, Orlando, FL, 1984. Academic Press.

[24] E. Çinlar and S. Özekici. Reliability of complex devices in random environments. *Probability in the Engineering and Informational Sciences*, 1:97–115, 1987.

[25] E. Çinlar, M. Shaked, and J. Shanthikumar. On lifetimes influenced by a common environment. *Stochastic Processes and Their Applications*, 33:347–359, 1989.

[26] P. Cheng, G. van Bussel, G. van Kuik, and J. H. Vugts. Reliability-based design methods to determine the extreme response distribution of offshore wind turbines. *Wind Energy*, 6(1):1–22, 2003.

[27] S. Chib. Calculating posterior distributions and model estimates in Markov mixuture models. *Journal of Econometrics*, 75(1):79–97, 1996.

[28] C. Ciang, J. Lee, and H. Bang. Structural health monitoring for a wind turbine system: A review of damage detection methods. *Measurement Science and Technology*, 19(12):1–20, 2008.

[29] M. Cutler. Bearing life prediction in paper making machinery, including environmental effects. In *TAPPI Engineering Conference*, Toronto, Ontario, 1990.

[30] P. Dellaportas. Bayesian model selection for partially observed diffusion proceses. *Biometrika*, 93(4):809–825, 2006.

[31] A. Dobakhshari and M. Fotuhi-Firuzabad. A reliability model of large wind farms for power system adequacy studies. *IEEE Transactions on Energy Conversion*, 24(3):792–801, 2009.

[32] K. Doksum and A. Hóyland. Models for variable-stress accelerated life testing experiments based on Wiener processes and the inverse Gaussian distribution. *Technometrics*, 34(1):74–82, 1992.

[33] L. Dueas-Osorio and B. Basu. Unavailability of wind turbines due to wind-induced accelerations. *Engineering Structures*, 30(4):885–893, 2008.

[34] J. Esary, A. Marshall, and F. Prochan. Shock models and wear processes. *Annals of Probability*, 1(4):627–649, 1973.

[35] S. Faulstich, B. Hahn, and P. Tavner. Wind turbine downtime and its importance for offshore deployment. *Wind Energy*, 14(3):327–337, 2011.

[36] S. Faulstich, P. Lyding, and P. Tavner. Effects of wind speed on wind turbine availability. In *Proceedings of the European Wind Energy Conference and Exhibition*, Brussels, Belgium, 2011.

[37] S. Frühwirth-Schnatter. Markov chain Monte Carlo estimation of classical dynamic and switching mixture models. *Journal of the Americal Statistical Association*, 96(453):194–209, 2001.

[38] D. Gamerman. Dynamic Bayesian models for survival data. *Applied Statistics*, 40(1):63–79, 1991.

[39] N. Gebraeel. Sensory-updated residual life distributions for components with exponential degradation patterns. *IEEE Transactions on Automation Science and Engineering*, 3(4):382–393, 2006.

[40] N. Gebraeel, M. Lawley, R. Li, and J. Ryan. Residual-life distributions from component degradation signals: A Bayesian approach. *IIE Transactions*, 37(6):543–557, 2005.

[41] N. Gebraeel and J. Pan. Prognostic degradation models for computing and updating residual life distributions in a time-varying environment. *IEEE Transactions on Reliability*, 57(4):539–550, 2008.

[42] A. Gelgrand and A. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409, 1990.

[43] C. Gray and S. Watson. Physics of failure approach to wind turbine condition based maintenance. *Wind Energy*, 13(5):395–405, 2010.

[44] M. Hahn, S. Frühwirth-Schnatter, and J. Sass. Markov chain Monte Carlo methods for parameter estimation in multidimension continuous time switching diffusion models. *Journal of Financial Econometrics*, 8(1):88–121, 2010.

[45] M. Hahn and J. Sass. Parameter estimation in continuous-time Markov switching models: A semi-continuous Markov chain Monte Carlo approach. *Bayesian Analysis*, 4(1):63–84, 2009.

[46] T. Harris and R. Barnsby. Life rating for ball and roller bearings. *Proceedings of the Institution of Mechanical Engineers, Part J: Journal of Engineering Tribology*, 215(6):577–595, 2001.

[47] W. Hastings. Monte Carlo sampling methods using Markov chains and their applicaitons. *Biometrika*, 57(1):97–109, 1970.

[48] E. Hau. *Wind Turbines: Fundamentals, Technologies, Application, Economics*. Springer, New York, NY, 2013.

[49] J. Ivy and S. Pollock. Marginally monotonic maintenance policies for a multi-state deteriorating machine with probabilistic monitoring and silent failures. *IEEE Transactions on Reliability*, 54(3):489–497, 2005.

[50] A. Jardine, P. Anderson, and D. Mann. Application of the Weibull proportional hazards model to aircraft and marine engine failure data. *Quality and Reliability Engineering International*, 3(2):77–82, 1987.

[51] R. Jiang, M. Kim, and V. Makis. Availability maximization under partial observations. *OR Spectrum*, DOI: 10.1007/s00291-012-0294-3, 2012.

[52] J. Kharoufeh. Explicit results for wear processes in a Markovian environment. *Operations Research Letters*, 31(3):237–244, 2003.

[53] J. Kharoufeh and S. Cox. Stochastic models for degradation-based reliability. *IIE Transactions*, 37(6):533–542, 2005.

[54] J. Kharoufeh, S. Cox, and M. Oxley. Reliability of manufacturing equipment in complex environments. *Annals of Operations Research*, DOI 10.1007/s10479-011-0839-x.

[55] J. Kharoufeh, D. Finkelstein, and D. Mixon. Availability of periodically inspected systems with Markovian wear and shocks. *Journal of Applied Probability*, 43(2):303–317, 2006.

[56] J. Kharoufeh and D. Mixon. On a Markov-modulated shock and wear process. *Naval Research Logistics*, 56(6):563–576, 2009.

[57] J. Kharoufeh, C. Solo, and M. Ulukus. Semi-Markov models for degradation-based reliability. *IIE Transactions*, 42(8):599–612, 2010.

[58] P. Kiessler, G. Klutke, and Y. Yang. Availability of periodically inspected systems subject to Markovian degradation. *Journal of Applied Probability*, 39(1):700–711, 2012.

[59] H. Kim and C. Singh. Reliability modeling and simulation in power systems with aging characteristics. *IEEE Transactions on Power Systems*, 25(1):21–28, 2010.

[60] M. Kim and V. Makis. Optimal control of a partially observable failing system with costly multivariate observations. *Stochastic Models*, 28(4):584–608, 2012.

[61] D. Kroese, T. Taimre, and Z. Botev. *Handbok of Monte Carlo Methods*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2011.

[62] M. Lee, G. Whitmore, F. Laden, J. Hart, and E. Garshick. Assessing lung cancer risk in railroad workers using a first hitting time regression model. *Environmetrics*, 15(5):501–512, 2004.

[63] J. Leichty and G. Roberts. Markov chain Monte Carlo methods for switcing diffusion models. *Biometrika*, 88(2):299–315, 2001.

[64] A. Leite, C. Borges, and D. Falcao. Probabilistic wind farms generation model for reliability studies applied to Brazilian sites. *IEEE Transactions on Power Systems*, 21(4):1493–1501, 2006.

[65] C. Liao and S. Tseng. Optimal design for step-stress accelerated degradation tests. *IEEE Transactions on Reliability*, 55(1):59–66, 2006.

[66] H. Liao, W. Zhao, and H. Guo. Predicting remaining useful life of an individual unit using proportional hazards model and logistic regression model. In *Annual Reliability and Maintainability Symposium*, pages 127–132, 2006.

[67] B. Lu, Y. Li, X. Wu, and Z. Yang. A review of recent advances in wind turbine condition monitoring and fault diagnosis. In *Power Electronics and Machines in Wind Applications (PEMWA)*, pages 1–7. IEEE, 2009.

[68] C. Lu and W. Meeker. Using degradation measures to estimate a time-to-failure distribution. *Technometrics*, 35(2):161–174, 1993.

[69] L. Maillart. Maintenance policies for systems with condition monitoring and obvious failures. *IIE Transactions*, 38(6):463–475, 2006.

[70] L. Maillart and L. Zheltova. Structured maintenance policies on interior sample paths. *Naval Research Logistics*, 54(6):645–655, 2007.

[71] V. Makis and X. Jiang. Optimal replacement under partial observations. *Mathematics of Operations Research*, 28(2):382–394, 2003.

[72] D. McMillan and G. Ault. Condition monitoring benefit for onshore wind turbines: Sensitivity to operational parameters. *IET Renewable Power Generation*, 2(1):60–72, 2008.

[73] L. Myers. Survival functions induced by stochastic covariate processes. *Journal of Applied Probability*, 18(2):523–529, 1981.

[74] L. Nelson, L. Manuel, H. Sutherland, and P. Veers. Statistical analysis of wind turbine inflow and structural response data from the LIST program. *Journal of Solar Energy Engineering*, 125(4):541–550, 2003.

[75] J. Nilsson and L. Bertling. Maintenance management of wind power systems using condition monitoring systems-life cycle cost analysis for two case studies. *IEEE Transactions on Energy Conversion*, 22(1):223–229, 2007.

[76] S. Özekici and R. Soyer. Bayesian analysis of Markov modulated Bernoulli processes. *Mathemathical Methods of Operations Research*, 57(1):125–140, 2003.

[77] S. Özekici and R. Soyer. Mathematical reliability: An expository perspective. *International Series in Operations Research and Management Science*, 67(1):249–273, 2004.

[78] J. Peinke, S. Barth, F. Böttcher, D. Heinemann, and B. Lange. Turbulance, a challenging problem for wind energy. *Physica A: Statistical Mechanics and its Applications*, 338(1-2):187–193, 2004.

[79] W. Pierskalla and J. Voelker. A survey of maintenance models: The control and surveillance of deteriorating systems. *Naval Research Logistics Quarterly*, 23(3):353–388, 2006.

[80] M. Pijnenburg. Additive hazards models in repairable systems reliability. *Reliability Engineering and System Safety*, 31(3):369–390.

[81] M. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599, 1999.

[82] G. Polson and O. Roberts. Bayes factors for discrete observations from diffusion processes. *Biometrika*, 81(1):11–26, 1994.

[83] M. Puterman. *Markov Decision Processes*. John Wiley and Sons Inc., New York, NY, 2005.

[84] L. Rademakers, H. Braam, M. Zaaijer, and G. van Bussel. Assessment and optimisation of operation and maintenance of offshore wind turbines. In *Proceedings of the European Wind Energy Conference*, Madrid, Spain, 2003.

[85] V. A. Riziotis and S. G. Voutsinas. Fatigue loads on wind turbines of different control strategies operating in complex terrain. *Journal of Wind Engineering and Industrial Aerodynamics*, 85(3):211–240, 2000.

[86] C. Robert, G. Celeux, and J. Diebolt. Bayesian estimation of hidden Markov chains: A stochastic implementation. *Statistics and Probability Letters*, 16(1):77–83, 1993.

[87] G. Roberts and O. Stramer. On inference for partially observed nonlinear diffusion models using the Metropolis-Hastings algorithm. *Biometrika*, 88(3):603–622, 2001.

[88] M. Robinson and M. Crowder. Bayesian methods for a growth-curve degradation model with repeated measures. *Lifetime Data Analysis*, 6(4):357–374, 2000.

[89] J. Schijve. *Fatigue of Structures and Materials*. Springer, New York, NY, 2009.

[90] G. Schwartz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.

[91] B. Sericola. Occupation times in Markov chains. *Communications in Statistics–Stochastic Models*, 16(5):479–510, 2000.

[92] Y. Sherif and M. Smith. Optimal maintenance models for systems subject to failure. *Naval Research Logistics Quarterly*, 28(1):47–74, 1981.

[93] J. D. Sørensen. Framework for risk-based planning of operation and maintenance for offshore wind turbines. *Wind Energy*, 12(5):493–506, 2009.

[94] F. Spinato, P. J. Tavner, G. J. W. van Bussel, and E. Koutoulakos. Reliability of wind turbine subassemblies. *Renewable Power Generation, IET*, 3(4):387–401, 2009.

[95] Y. Sun, L. Ma, J. Mathew, W. Wang, and S. Zhang. Mechanical systems hazard estimation using condition monitoring. *Mechanical Systems and Signal Processing*, 20(5):1189–1201, 2006.

[96] H. Sutherland and D. Burwinkle. The spectral content of the torque loads on a turbine gear tooth. *Wind Energy*, 16(1):91–97, 1995.

[97] P. Tavner, C. Edwards, A. Brinkman, and F. Spinato. Influence of wind speed on wind turbine reliability. *Wind Engineering*, 30(1):55–72, 2006.

[98] P. Tavner, Y. Qiu, A. Korogiannos, and Y. Feng. The correlation between wind turbine turbulance and pitch failure. In *Proceedings of the European Wind Energy Conference and Exhibition*, Brussels, Belgium, 2011.

[99] P. Tavner, J. Xiang, and F. Spinato. Reliablity analysis for wind turbines. *Wind Energy*, 10(1):1–18, 2007.

[100] P. J. Tavner. Review of condition monitoring of rotating electrical machines. *Electric Power Applications, IET*, 2(4):215–247, 2008.

[101] P. J. Tavner, R. Gindele, and S. Faulstich. Study of effects of weather and location on wind turbine failure rates. In *European Wind Energy Conference (EWEC 2010)*, Warsaw, Poland, 2010.

[102] M. Ulukus, J. Kharoufeh, and L. Maillart. Optimal replacement policies under environment-driven degradation. *Probability in the Engineering and Informational Sciences*, 26(3):405–424, 2012.

[103] G. J. W. van Bussel and A. R. Henderson. State of art and technology trends for offshore wind energy: Operation and maintenance issues. In *Offshore Wind Energy, EWEA Special Topic Conference*, Brussels, Belgium, 2001.

[104] D. Virkler, B. Hilberry, and P. Goel. The statistical nature of fatigue crack propagation. *ASME Journal of Engineering Materials and Technology*, 101(2):148–153, 1979.

[105] S. Vittal and M. Teboul. Performance and reliability analysis of wind turbines using Monte Carlo methods based on system transport theory. In *Proceedings of the 46th AIAA Structures, Structural Dynamics and Materials Conference*, pages 1–7, Austin, Texas, 2005.

[106] P. Vlok, M. Wnek, and M. Zygmunt. Utilising statistical residual life estimates of bearings to quantify the influence of preventive maintenance actions. *Mechanical Systems and Signal Processing*, 18(4):833–847, 2004.

[107] C. Walford. Wind turbine reliability: Understanding and minimizing wind turbine operation and maintenance costs. Technical Report SAND2006–1100, Sandia National Laboratories, Albuquerque, NM, March 2006.

[108] H. Wang. A survey of maintenance policies of deteriorating systems. *European Journal of Operations Research*, 139(3):469–489, 2002.

[109] X. Wang. Wiener processes with random effects for degradation data. *Journal of Multivariate Analysis*, 101(2):340–351, 2010.

[110] G. Whitmore and F. Schenkelberg. Modelling accelerated degradation data using Wiener diffusion with a time scale transformation. *Lifetime Data Analysis*, 3(1):27–45, 1997.

[111] M. Wilkinson, F. Spianto, and M. Knowles. Towards the zero maintenance wind turbine. In *Proceedings of the 41st International Universities Power Engineering Conference*, pages 74–78, Newcastle Upon Tyne, UK, 2006.

[112] R. Wiser and M. Bolinger. 2010 wind technologies market report. Technical report, U.S. Department of Energy, Washington, D.C., June 2011.

[113] R. J. Wood, A. S. Bahaj, S. R. Turnock, L. Wang, and M. Evans. Tribological design constraints of marine renewable energy systems. *Philosophical Transactions of the Royal Society A*, 368(1929):4807–4827, 2010.

[114] E. Yaylali and J. Ivy. Partially observable MDPs (POMDPs): An introduction. In J. Cochran, L. Cox, P. Keskinocak, J. Kharoufeh, and J. Smith, editors, *Wiley Encyclopedia of Operations Research and Management Science*, pages 4017–4036, New York, NY, 2011. John Wiley and Sons, Inc.

[115] M. You, L. Li, G. Meng, and J. Ni. Two-zone proportional hazard model for equipment remaining useful life prediction. *Journal of Manufacturing Science and Engineering*, 132(4):041008–1–041008–6, 2010.

[116] E. Zhou, M. Fu, and S. Marcus. Solving continuous-state POMDPs via density projection. *IEEE Transactions on Automatic Control*, 55(5):1101–1116, 2010.

[117] Y. Zhou, L. Ma, J. Mathew, Y. Sun, and R. Wolff. Maintenance strategy optimization using a continuous-state partially observable semi-Markov decision process. *Microelectronics Reliability*, 51(2):300–309, 2011.