



City Research Online

City St George's, University of London

Citation: Dagg, R.A. (1999). Optimal inspection and maintenance for stochastically deteriorating systems. (Unpublished Doctoral thesis, City University London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/8227/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Optimal Inspection and Maintenance for Stochastically Deteriorating Systems

Richard Andrew Dagg

A Thesis submitted for the degree of Doctor of Philosophy in the Department of
Actuarial Science and Statistics, The City University, London

December 1999

Table of Contents

	List of Figures	6
	List of Tables	7
	Acknowledgements & Declaration	8
	Abstract	9
1	Reliability and Maintenance of Deteriorating Systems	10
1.1	Introduction	10
1.2	Modelling system degradation	12
1.3	Modelling Failure as a result of degradation	14
1.4	Modelling maintenance and inspection	16
1.5	Optimisation Criteria	19
1.6	Summary of Thesis	20
2	Modelling Systems subject to Stochastic Degradation	22
2.1	Introduction	22
2.2	Technical Properties of degradation processes	24
2.3	The Wiener Process as a Degradation Model	28
2.3.1	Relevant Properties of Wiener Processes	28
2.3.2	Degradation models based on the Wiener process	32
2.4	The Gamma Process as a Degradation Model	35
2.5	Models of degradation in the presence of covariates	39
2.6	Maintenance and inspection Models	44
2.7	Summary	48

3	Optimal Perfect Inspection	50
3.1	Introduction	50
3.2	Underlying Inspection and replacement model	51
3.2.1	Model for periodic inspection	54
3.2.2	Model for non-periodic inspections	55
3.2.3	Model for no inspections and continuous condition monitoring	55
3.3	Optimal Periodic Inspection : Average cost criterions	56
3.3.1	Derivation of Expected Cost per Unit Time	57
3.3.2	Special Case of Monotone Degradation Process Model	62
3.3.3	Costs of no inspections and condition monitoring	64
3.4	Optimal Periodic Inspection: Discounted Cost Criterion	66
3.4.1	Derivation of Discounted total cost	66
3.4.2	Special Case of Monotone Degradation Process Model	68
3.4.3	Costs of no inspection and condition monitoring	70
3.5	Optimal Non-Periodic Inspection: Discounted Cost Criteria	73
3.5.1	Optimality Equations in General	73
3.5.2	Special Case of Monotone Degradation Process Model	76
3.6	Example : Gamma Process Degradation	77
3.6.1	Periodic Inspection : Average cost per unit time	78
3.6.2	Periodic Inspection : Discounted total cost	80
3.6.3	Non-Periodic Inspection : Discounted Total cost	81
3.6.4	Numerical Results and Comments	81
3.7	Conclusion	92
4	Optimal Imperfect Inspection of Lévy process models	94
4.1	Introduction	94
4.2	Underlying Inspection and Replacement Model.	95
4.3	Optimal Periodic Inspection of a covariate Process: Average Cost Criterion.	97
4.3.1	Derivation of Expected Cost Per unit time	97

4.3.2 Special case of Monotone Degradation process model	103
4.3.3 Costs of no inspections and condition monitoring	104
4.4 Optimal periodic inspection of a covariate process: Discounted cost criterion	105
4.4.1 Derivation of disocunted total cost	106
4.4.2 Special case of monotone degradation process model	108
4.4.3 Costs of no inspection and condition monitoring	109
4.5 Optimal Non-Periodic Inspection of a covariate process: Discounted cost criterion.	110
4.6 Models for imperfect inspection	112
4.7 Example : Imperfect inspection of a Gamma Process	114
4.7.1 Model Assumptions	114
4.7.2 Distributional Results for Gamma process with error.	115
4.7.3 Solution of the Integral Equation.	117
4.7.4 Numerical Results	119
4.8 Conclusion	122

5 Optimal Maintenance Policies 124

5.1 Introduction	124
5.2 Underlying Inspection and Repair Model.	125
5.2.1 System model	125
5.2.2 Models for inspection and maintenance actions	127
5.3 Optimal inspection in the presence of maintenance	129
5.3.1 Optimal Inspection under deterministic maintenance.	130
5.3.2 Optimal Inspection of randomly Maintained Systems.	135
5.3.3 Optimal Non-periodic inspections.	137
5.4 Optimal Maintenance for a given Inspection Policy	140
5.4.1 Optimal Deterministic Maintenance	141
5.4.2 Optimal General Maintenance	143
5.4.3 Joint Optimisation of Maintenance and Inspections	144
5.5 Example : Gamma Process Degradation	146
5.5.1 Optimal Inspection of a deterministically maintained system	147

5.5.2 Optimal Periodic Inspection and Maintenance Policies	152
5.6 Conclusions	155
6 Summary and Conclusions	157
6.1 Summary	157
6.2 Possibilities for further research	159
A1 Approximate Solution of Integral Equations	161
A1.1 Approximate solution of Fredholm integral equations	161
A1.2 Fredholm equations with diagonally singular kernel	163
A2 Results in the case of Wiener process Degradation	165
A2.1 Introduction.	165
A2.2 Joint Density of X_t and M_t	166
A2.3 Distribution function of X_t and M_t	167
A2.4 Probability Density function of Y_t and M_t	168
A3 Derivation of the Dynamic Programming Equation	172
A3.1 Derivation of the dynamic programming equation	172
References	175

List of Figures

- Fig 1.1 Graphical Representation of a Degradation Model
- Fig 1.2 Example of Threshold Failure Model
- Fig 1.3 Graphical Representation of Degradation Models Incorporating inspection and Maintenance
- Fig 2.1 Example of a Markov Additive Process Degradation Model
- Fig 2.2 Example of degradation with a control limit replacement policy and threshold failure model
- Fig 3.2 Possible forms for cost functions for periodic inspection policies

List of Tables

Table 3.1	Optimal Inspection interval and corresponding discounted total cost – Average cost criterion
Table 3.2	Optimal Inspection interval and corresponding discounted total cost – Discounted total cost
Table 3.3	Minimum discounted total cost of optimal non-periodic inspection policy
Table 3.4	Optimal Non-periodic inspection policies
Table 4.1	Optimal inspection policy, corresponding discounted total cost and percentage of optimal DTC with perfect inspection: $v = 0.1$
Table 4.2	Optimal inspection policy and corresponding discounted total cost for a system with perfect inspection
Table 5.1	Table showing effect of maintenance schedule with $\alpha = \beta = 7$, on discounted total cost and optimal inspection interval
Table 5.2	Table showing effect of maintenance schedule with $\alpha = \beta = 25$, on discounted total cost and optimal inspection interval
Table 5.3	Optimal Non-periodic Inspection schedule for system subject to maintenance. Table shows optimal inspection policy and corresponding discounted total cost, alongside optimal periodic policy and cost.
Table 5.4	Jointly Optimal Maintenance and Periodic Inspection schedule. Table shows optimal policy and corresponding discounted total cost

Acknowledgements

Firstly I would like to thank my supervisor, Professor Martin Newby for his help and encouragement throughout the period of my research. He was always available for consultation, discussion and help whenever it was required, and provided valuable insight into many of the problems and difficulties I encountered. Further, I would like to thank the Department of Actuarial Science and Statistics for funding my research through a departmental scholarship.

At a personal level, I must thank my family and friends for supporting my research throughout the last three years. In particular my greatest thanks must go to Qadar Asad. Without her friendship, help and support this thesis would certainly not have been completed.

Declaration of Powers of Discretion

I grant powers of discretion to the University Librarian to allow this thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgement.

Abstract

This thesis concerns the optimisation of maintenance and inspection for stochastically deteriorating systems. The motivation for this thesis is the problem of determining condition based maintenance policies, for systems whose degradation may be modelled by a continuous time stochastic process. Our emphasis is mainly on using the information gained from inspecting the degradation to determine efficient maintenance and inspection policies.

The system we shall consider is one in which the degradation is modelled by a Lévy process, and in which failure is defined to occur when the degradation reaches a critical level. It is assumed that the system may be inspected or repaired at any time, and that the costs of inspections and repairs may depend on the level of system degradation.

Initially we look at determining optimal inspection policies for systems whose degradation may be directly and perfectly observed, before extending this analysis to the case where the degradation is unobservable, and a related covariate process is used to determine maintenance decisions. In both cases it is assumed the replacement policy is fixed and known in advance. Finally we consider the case of joint optimisation of maintenance and inspection, for cases in which the maintenance action has either deterministic or random effect on the degradation level.

In all of these cases we use the properties of the Lévy process degradation model to form a recursive relationship which allows us to determine integral and functional equations for the maintenance cost of the system. Solutions to these determine optimal periodic and non-periodic inspection and maintenance policies.

Throughout the thesis we use the gamma process degradation model as an example. For this model we determine optimal perfect inspection policies for the cases when inspections are periodic and non-periodic. As a special case of a covariate process we consider the optimal imperfect periodic inspection policy. Finally we obtain jointly optimal deterministic-maintenance and periodic-inspection policies.

Chapter 1

Introduction: Maintenance and Inspection of Deteriorating Systems

1.1 Introduction

Engineering advances in recent years have meant that the complexity of many systems has increased greatly. Systems are now being designed to carry out ever more complex and difficult tasks. This has resulted in engineers seeking more effective methods to improve and monitor the reliability of these systems.

There is a limit, however, to how far the reliability of a system may be improved by improvements in the quality of components and construction. When this limit has been reached other methods must be found to ensure that systems are reliable enough to safely complete their required task. An important consequence of this is that maintenance and inspections of these systems is now of much greater importance. However, the improvements in reliability achieved by quality improvement methods mean that traditional methods of reliability centred maintenance are much less effective. The main reason for this is that, in the case of a highly reliable system, it is difficult to estimate the failure characteristics of the system, when failures occur very rarely. This means that the distribution of the time to failure is unavailable, so that other methods of analysis must be found.

Since the event of system failure cannot be observed, an alternative measure of system performance must be found. It is natural, as a proxy for failure, to consider the degradation of the system, where degradation is a measure of performance, capability, quality or damage. For example, the degradation of a tyre is measured by tread depth. The tread depth clearly affects the safety, performance and capability of the tyre. Using crack length as a degradation measure is natural for many mechanical systems under cyclical loading. In this case, the length of the crack does not affect the performance of the system, until the crack becomes large enough for the material to break, in which case the system fails. Clearly, the degradation is related to failure, and so measurements of degradation allow us to indirectly determine the failure characteristics of a system.

The main focus of this thesis is the analysis of deteriorating systems, with the aim of determining optimal maintenance and inspection policies. We are mainly interested in systems whose degradation may be observed by the system user. On the basis of observed degradation, decisions can then be made regarding the inspection and maintenance of the system.

In this introductory chapter we will consider the practical problems associated with the maintenance of deteriorating systems, and look in detail at the aspects of these systems which are important when developing mathematical models. The presentation is largely non-mathematical, since a survey of mathematical models for maintenance and inspection is given in Chapter 2.

We shall divide our discussion of maintenance into four sections:

1. Modelling System Degradation
2. Failure Characteristics
3. Inspection and Maintenance Policies
4. Optimisation Criteria

Each section presents an important practical aspect of a system or its management, and discusses problems which arise when considering mathematical models for deteriorating systems and their inspection and maintenance.

1.2 Modelling System Degradation

In cases when it is not possible or economical to test a system to failure, it is usual to measure the degradation of a system, to gain insight into the system failure characteristics. In what follows we shall assume that degradation of a system is measured by a physical observation of a characteristic of the system.

A classical example of this type of model is that of Crack Growth (Sobczyk, 1987). In this case the 'system' fails when the crack length becomes too large. It is natural, then, to use crack length as a measure of system degradation. By observing the crack length, and by using knowledge of system characteristics, we are able to determine how much longer the system is likely to function successfully. Using this information we can formulate a maintenance and inspection policy.

Another classical example of degradation is in erosion or wear processes. In this case the system has a specific characteristic which is eroded or wears out over time. For example, in the case of tyres, the tread depth is reduced by tyre wear, and eventually the tread is worn away, making the tyre useless. Another example is given by Van Noortwijk (1996), who considers the problems of maintaining coastal flood barriers. In this case the sea gradually erodes the flood barrier, and the barrier is deemed to have failed when it is no longer able to withstand the pressure of the sea. In this case by observing the height and width of the barrier, a prediction can be made regarding the expected failure time of the barrier. Using this information maintenance can be planned to prevent the barrier from failing.

In the two preceding examples, the degradation process clearly has a direct influence on the failure of the system. In many systems however, the degradation process which should be considered is not obviously apparent. In this section we shall consider the problems in modelling degradation, and how they may be overcome.

Before considering a model for degradation, it is necessary to determine which system characteristics are of use in 'predicting' failure. It is clear in the crack growth example that crack length is of prime importance. It may also be the case however that other variables influence the failure time of the system, and so should be incorporated into any degradation model. In general, we assume that there are a number of system variables which directly describe the failure characteristics of the system. This implies that our general degradation model should be in the form of a vector valued stochastic

process. We shall assume that an analysis of the system has been carried out and that the system variables influencing failure have been identified.

Having identified which variables affect system failure, we must address the problem of observability. If a system variable cannot be observed, for technical reasons say, then it cannot be incorporated directly into a degradation model. In this case it may be necessary to consider a ‘covariate’ of the system variable. By a covariate, we mean a system variable which is conditionally independent of system failure, given the degradation process. This means that the covariate does not directly influence system failure, but provides statistical information about failure, when the underlying degradation process is unobservable. Clearly there may be more than one covariate relating to any single unobservable system variable, so that the covariate model may be a vector-valued stochastic process.

There is a possibility that some of the covariates are themselves unobservable, in which case we could consider covariates of the initial covariate under consideration. Carrying on in this fashion we obtain a hierarchy of processes, each level being less important to failure than the previous level. To avoid this type of hierarchical model we shall assume that all covariate processes are observable. Therefore the model of degradation is restricted to three levels, as illustrated in the graphical model shown in figure 1.1 below. This is the most general type of model we shall consider.




Figure 1.1 – Graphical Representation of a degradation model

In this model, the main point is that the covariate influences only the degradation of the system. If the degradation process is observable, then the covariate process and the event of failure are independent. From a modelling perspective this is important since it allows us to greatly simplify the analysis of a system.

Having determined the degradation variables and covariates, it is necessary to define a mathematical model defining the relationship between (i) The covariate process and the degradation process; and (ii) The degradation process and failure. In the first case it is

assumed by definition of our terminology that a known relationship between the covariate and the degradation process exists. For example, it could be the case that the average rate of deterioration of the system is a function of the value of the covariate process. In the crack growth example this might correspond to the average rate of crack growth being a function of the operating temperature of the system. In the coastal flood protection example, the level of rainfall is an obvious environmental factor influencing the failure of the system. Regarding the second case, it is often much more difficult to adequately define a relationship between degradation and physical failure. This problem is discussed further in the following section.

In both cases however, it can be seen that for systems with vector valued degradation and covariate processes, the system models become very complex. The choice of model for each of these processes plays a very important role, since it determines how useful and accurate the model is. Much work has been done on the estimation of reliability for systems subject to random covariates. However, most of this literature is related to medical statistics, and analysis of survival data (for example see Jewell and Kalbfleisch, 1996). This means that the models considered are generally regression models, which attempt to explain the effect of various covariates/factors on the lifetime of a system/individual. These models are not generally suited to optimisation of maintenance and inspection. We consider some of these models in greater detail in the following chapter.

Further aspects of degradation modelling are more mathematical in nature and are discussed in the following chapter.

1.3 Modelling Failure as a Result of Degradation

While the degradation of a system is often relatively easy to model, the actual failure is not as easy to predict. Clearly, modelling degradation can only provide us with partial information regarding the failure of the system.

Given that this is the case, the question arises as to how failure can be modelled as a function of the degradation of the system. In general we are interested in preventing a physical failure of a system. Physical failure is generally defined as ‘the termination of the ability of an entity to perform a required function’ (IEC 1991). Clearly, in the

general model we are using we must define a relationship between the degradation process and failure. There are two common assumptions which have been made.

The first, and perhaps simplest method, is to define a threshold failure model in terms of technical failure. More precisely, if X represents the degradation process of the system we are considering, then a technical failure occurs at the first time $X_t \in C$, where C is some critical failure set. Under this definition the failure time represents the hitting time of the set C . The set C is then chosen so that the system is unlikely to physically fail for values of degradation outside the critical set (with a specified probability).

The most common example of this model is shown in figure 1.2 below.




Figure 1.2 – Example of Threshold Failure Model

In the example of figure 1.2, failure is assumed to occur when the degradation process first reaches the barrier C . In this particular case the degradation reaches C at time T , and this is defined as the failure time.

In some senses this definition is unsatisfactory, since failure is now assumed to occur when the degradation reaches a specific value; in reality this is not the case. While this model can be criticised as being overly simplistic, its simplicity allows computations to be carried out, which would be impossible under more realistic assumptions. With a good definition of technical failure this model can be very useful, and has often been used in the literature.

One difficulty with threshold failure models is that they take no account of external factors on the system. For example, it may be the case that a system will fail due to incompetence of a system operator. This cannot be incorporated into a threshold failure

model. In general we must be careful to ensure that a technical failure, defined by first entrance to a critical set C , must occur for the system to fail (within reason, and subject to external factors). Clearly, these external factors are much more difficult to model than the degradation of the system.

A simplification of this type of threshold failure model may be defined when we are dealing with a degradation process with discrete state space. In this case it is common to view failure as the culmination of the degradation process, and to define failure as a state of the process. Examples of models using this approach are given by Milioni and Pliska (1988a,b) and Dagg and Newby (1998). These models, however, are based on systems with precise properties and cannot be applied to more general models.

The second main failure model which has been proposed is the doubly stochastic Poisson Process (See, for example, Cox and Isham, 1980). The idea of this model is to assume that failures occur according to a Poisson Process, whose intensity is given by a function of the stochastic covariate process. While this assumption seems more realistic there are some difficulties, principally in determining the nature of the failure intensity. If this can be obtained, however, the assumptions result in a realistic model, allowing failures to occur at any level of degradation, and to allow the level of degradation to influence the failure rate of the system.

A comprehensive review covering much of the material of sections 1.2 and 1.3 may be found in Singpurwalla (1995) or Yashin and Manton (1997). We consider some particular failure models in the following chapter.

1.4 Modelling Maintenance and Inspection

Having outlined how the system degradation and failure may be modelled, we now consider the important topic of how this failure can be prevented. In general, to prevent a system from failing it must be maintained in some way. To decide on effective maintenance, it is necessary to obtain information about the system degradation, which is done by inspecting the system. Clearly, the introduction of inspection and maintenance complicates the degradation models somewhat, as it is now assumed that the maintenance of the system affects the level of degradation.

Graphically, the model we are dealing with is shown below




Figure 1.3 – Graphical Representation of Degradation Models Incorporating Inspection and Maintenance

The arrows in the diagram are interpreted as meaning ‘directly influences’. Thus, as before covariates affect degradation, which in turn affects failure. Additionally, we now have that the level of degradation affects the observed value of inspection, which in turn determines the maintenance action, which affects degradation and possibly also affects the value of covariates.

It is important, given what we have said above, to define inspections and maintenance actions as precisely as possible. In general there are two main forms of inspection model, namely perfect and imperfect inspection. Under perfect inspection we assume that the true value of the degradation process is observed at each inspection. Similarly imperfect inspection assumes that the degradation process is observed with some error. The exact way these are defined depends on the underlying system model. There are commonly two assumption made in regard to imperfect inspection:

1. Measurement Error
2. Classification Error

As the name suggests, measurement error occurs when the true level of degradation cannot be measured because of inaccuracies or faults with equipment. A common model for measurement error assumes that the degradation process is observed with an additive error term (See for example Whitmore, 1995). This is not the only model available, but its structure makes the analysis of this type of model quite straightforward.

Classification error occurs usually when the system has a discrete state space, and the system is categorised as being in one state, when in fact it is in another. In this case imperfect inspection may be defined by allowing a probability of observing degradation is state x when in fact degradation is in state y , and this model fits well with standard models of systems with discrete state space.

As with inspection, there are many possible types of maintenance action. The first distinction which can be made is between complete and partial maintenance. Complete maintenance is defined as any action which returns the system to a good-as-new condition. The most common example of such maintenance is replacement, but others do exist. In contrast to complete maintenance we define incomplete or general maintenance. In this case the maintenance action results in the degradation of the system being reduced to a specified level. This is clearly more realistic in terms of maintenance than complete repair, although problems still exist. Principally it is often difficult to determine the exact effect maintenance will have on the degradation and/or covariate process of the system.

We may also distinguish between perfect and imperfect maintenance. As with inspection, perfect maintenance occurs when the affect of maintenance on the system is known deterministically. For example, with replacement, we know the degradation of a system will be reduced to the level of a new system. With imperfect maintenance, the effect of maintenance is in some way random. This randomness may be intrinsic to the maintenance action itself, or may be the result of errors on the part of those carrying out maintenance. We may also define hazardous maintenance, in a similar way to hazardous inspection, but we do not consider this here.

Having defined the types of maintenance action which may be available, it is also important to consider the purpose of maintenance. Broadly, maintenance actions are described as being either preventive or corrective. Preventive maintenance takes place prior to system failure, and aims to prevent failure occurring by improving the condition of the system. The actual level and quality of improvement depends on the maintenance action applied. Corrective maintenance, however, takes place after a system has failed, and aims to fix the system so that it can be re-used. In many cases, however, it is impossible to fix a failed system, and a replacement will be carried out.

Having defined the maintenance and inspection actions available, we must now consider the timing of maintenance and inspection. Broadly, both of these may be characterised in terms of periodic and non-periodic policies. As the name suggests, periodic policies involve making inspections or carrying out maintenance actions at periodic intervals. Generally non-periodic policies are in some way adaptive, defining the next maintenance action or inspection in terms of the results of previous inspections. Generally non-periodic policies result in lower costs than periodic policies for the same system, although periodic policies are often considered easier to implement by maintenance managers

It can be seen from the above discussion that incorporating inspection and maintenance results in considerable complication of the degradation model. As a result, there are few papers in the literature which deal with the maintenance-inspection problem as a single entity. Most of the results in the literature tend to be focused on either maintenance or inspection, while ignoring the other.

It seems clear that maintenance and inspection should be jointly modelled, since both actions are very much inter-dependent, and we proceed with this in mind. The survey paper by Pham and Wang (1996) considers the literature on imperfect maintenance in more detail.

1.5 Optimisation Criteria

The idea of maintenance optimisation is that, based on observations of state variable and or covariates, we choose a maintenance policy (which may include inspections, repairs, and replacements and any other relevant action which may be taken to affect the state of a system) which best meets our objectives, however they are defined. Therefore we are dealing with a constrained optimisation problem which is time dependent.

At outset, even during system design, it is important to consider the objectives of the system and the criteria on which the system will be judged a success or not. These criteria will be defined by the system-user, based on his requirements and responsibilities. Assuming such criteria are set out, we assume that we wish to optimise

inspection and maintenance, such that we achieve the least possible cost of system operation and maintenance, subject to the system satisfying its objective criteria.

The procedure of optimisation then, is a mathematical problem, and the possibility of computing optimal policies depends upon the underlying models being used. In this thesis we shall assume that our objective in maintenance optimisation is to minimise the ‘costs’ of inspection and maintenance. Generally we shall define the costs as either

1. Expected Average cost per unit time; or
2. Expected Discounted total cost (over an infinite horizon)

These two criteria are the most commonly found in the literature on maintenance.

In some applications it is important to consider optimisation of quantities other than cost. For example, we may be required to find the maintenance policy which reduces costs, but does not allow the probability of failure to be greater than 1% at any time. We do not consider this type of problem, assuming that the costs of failure in the model will fully reflect the consequences of failure, so that the results of failure occurring are taken into account by the model and optimisation process.

1.6 Summary of Thesis

We have, in the preceding sections considered the most important problems in the area of maintenance modelling, from a practical point of view. The main outcome of this analysis was the need for coherent models incorporating degradation, covariates, inspection and maintenance. It is the aim of this thesis to go some way in presenting a model of degradation which allows these various factors to be incorporated, and hence to obtain corresponding optimal inspection and maintenance policies.

In Chapter 2 we consider in detail some of the most important mathematical degradation models which have been recently considered. Broadly we look at the whole range of models, but we pay special attention to those models which are related to those we shall consider later in this thesis. We also briefly look at maintenance models which have been considered for various types of systems, to set the scene for the analysis of the following chapters,

In Chapters 3, 4 and 5 we propose and extend a model for optimisation of inspection and maintenance, based on Lévy process degradation. Chapter 3 presents the basic derivation of a model for optimal markovian non-periodic inspection of a deteriorating system, based on methods of dynamic programming theory. These results are extended in chapter 4 to the case in which the degradation process is unobservable, and covariates are observed. Chapter 5, then considers optimal markovian maintenance policies for a similar system, allowing one to obtain jointly optimal (markovian) non-periodic inspection and maintenance policies.

Throughout the thesis the example of Gamma process degradation is used as an example. In chapter 3 and 5 optimal maintenance and inspection policies are found for gamma process degradation, and in chapter 4 the special case of imperfect inspection of a gamma process is considered.

We conclude in chapter 6 by summarising the main findings of the thesis, and considering possible extensions of the model proposed, and other future work based on our findings.

Chapter 2

Modelling Systems subject to Stochastic Degradation

2.1 Introduction

Having looked at reliability and maintenance in general, let us now consider how we can best encapsulate system properties in a mathematical model. As we have said, there are a great many types of system in use, and no single model is adequate for all systems. A great many models have been considered in the literature, but many of these are used for modelling a specific system. We shall mainly consider here models for systems whose degradation is a continuous time stochastic process, and look at related models for optimisation of maintenance and inspection of these systems.

From a modelling point of view, we are now interested in translating the system properties into a mathematical model. The aspects of reliability and maintenance described in the previous chapter must each be considered to find an appropriate model, so the system as a whole can be described. For the type of system we are interested in most previous work has been done in the field of degradation modelling, and the review which follows necessarily reflects this.

In what follows we let X_t be a random variable representing the accumulated amount of degradation of a system. In general, X_t may be a multivariate process, but we assume for the moment that it is a univariate stochastic process. It is possible to deal with the case of multivariate degradation case, by considering the multivariate process as set of correlated univariate processes.

Further, we assume that the degradation of the system is measured on an increasing scale. That is, if $X_s < X_t$ then the system is in a better condition at time s than at time t . If $U(x)$ represents the utility function of having a system in state x and if $\{X_t : t > 0\}$ is a univariate Stochastic Process representing the degradation of the system, we assume that

$$X_s < X_t \Rightarrow U(X_s) > U(X_t) \quad (2.1)$$

This assumption is not universal in the literature, since some authors assume that a decrease in the level of degradation X decreases the utility of the system. In certain applications it is more logical to consider the degradation process as a decreasing measure of the system condition. This is particularly the case when the state of the system is given by some measurement of system quality. In particular the models of Doksum and Normand (1995) and Van Noortwijk (1996) are of this type. In the former the system state is taken to be the number of Healthy cells in a blood sample, and in the latter the system state is, for example, the amount of sand remaining on a coastal flood barrier. However, by defining the degradation as decreasing measure, we avoid any confusion between the cases of monotonic increasing and decreasing processes separately, and allow ourselves to be completely consistent throughout. Clearly, if X_t is an increasing process, then the process $X'_t = K - X_t$ is a decreasing process, for any K . Transforming the process in this way always allows us to use an increasing measure of degradation.

In the case of a multivariate stochastic degradation process, we can assume that all of the component processes follow the above convention, i.e. smaller levels of degradation are preferred to larger levels of degradation. This, however, does not provide an ordering on the state space of the multivariate process, and so there may be problems when dealing with multivariate processes, and decisions regarding which states are preferable to others.

When discussing degradation process above we did not preclude the possibility that X is not-monotonically increasing. In most applications, we observe that the degradation of a non-maintained system is monotonically increasing, and it is preferable to incorporate this feature into a degradation model. Under the further assumption that X is a monotonically increasing process, i.e. for all $s < t$, $X_s \leq X_t$, it is clear that the degradation process then has the property that

$$s < t \Rightarrow U(X_s) \geq U(X_t) \quad (2.2)$$

This is clearly an intuitively desirable property for a degradation process to have, saying that as a system ages, its condition deteriorates. A degradation process satisfying the above condition will be called a monotone degradation process. We assume a multivariate process X is a monotone degradation process if all of the components are monotone degradation processes. This avoids the partial ordering problems described above. In this case equation (2.2) applies to a stochastic process in any number of dimensions.

With these basic definitions in mind we now consider some desirable properties of degradation processes, before looking at three common degradation models: the Wiener process, the compound Poisson process and the Gamma Process. We then consider some particular models which are useful in the case of a system in which covariates are observed. Finally we consider models for optimisation of maintenance and inspection of deteriorating systems.

2.2 Technical Properties of degradation processes

Let us now consider the properties of system degradation which must be encapsulated in our mathematical model of degradation. From these system properties we are able to decide which stochastic model is the most suitable for any given system.

Firstly, let us consider how the degradation of a system occurs. Various models have been considered, but most of these are in some sense based on the concept of accumulated damage. This is a natural assumption in many situations, where the degradation and resulting failure of a system may be viewed as the results of the

accumulation of wear and tear over the lifetime of the system. Van Noortwijk (1996) points out that in many systems which are subject to shocks, the order in which the damage (i.e. the shocks) occur is immaterial. This suggests that the random deterioration incurred in equal time intervals form a set of exchangeable random variables (Bernardo and Smith, 1992). This also implicitly assumes that the distribution of the degradation incurred is independent of the time scale, i.e. the process has stationary increments. The properties of exchangeable and stationary increments are very close indeed to the stronger properties of stationary and independent increments, which suggests Lévy processes may be a suitable candidate for our model. See Breiman (1968) for a discussion of the properties of general Lévy processes.

This reasoning however may be criticised in at least one important respect. While the independent increments property seems intuitively correct for many applications, the stationary increments property does not seem so intuitively appealing. In some applications (as described by Van Noortwijk 1996) it may be possible that this assumption holds, but it is reasonable to expect that under some circumstances, that the degradation accumulated by the system was in some way related to the age and/or the state of the system. We shall consider this problem further below.

This criticism however is overshadowed by the advantages that are to be gained in using certain Lévy processes as degradation models. The class of Lévy processes contains certain processes which are extremely tractable and for which there are many results readily available; namely the Compound Poisson Process, the Wiener process and the Gamma Process. In this respect, the simplification made in assuming independent and stationary increments is justifiable simply on the basis of tractability.

An important result which is proven by Breiman (1968, Ch 9,14) regarding Lévy processes is Lévy-Khinchine Decomposition. This states that that the characteristic function of a Lévy process $X(t)$ may be decomposed as $\varphi(\theta) = E(e^{i\theta X(t)}) = e^{\psi(\theta)t}$ where

$$\psi(\theta) = ic\theta - \frac{1}{2}\sigma^2\theta^2 + \int_{\mathbf{R}\setminus\{0\}} \left\{ e^{i\theta b} - 1 - \frac{i\theta b}{(1+b^2)} \right\} \nu(db)$$

where $c \in \mathbf{R}$, $\sigma > 0$, and ν is a measure on $\mathbf{R}\setminus\{0\}$ such that $\int_{\mathbf{R}\setminus\{0\}} (b^2 \wedge 1) \nu(db) < \infty$. This

result is extremely important because it shows that any Lévy process can be expressed as the sum of a Wiener process and a jump process. This has the consequence that any

degradation model based on Lévy processes must be either a Wiener process, a jump process (such as a Compound Poisson process or shot noise process) or the sum of these two processes. This decomposition also shows that the Wiener process is the only continuous Lévy process, and so if we want to model a system which degrades continuously (such as a fatigue crack) using a Lévy process, we are theoretically limited to a Wiener Process.

As we have said above, monotonicity may be a crucial property of degradation processes in many applications. It can be seen from the above decomposition that only processes which do not have a Wiener process component can be monotonic. It is clear from the above discussion then, that any monotonic increasing Lévy process must be a jump process. Rogers and Williams (1994) give a decomposition of the Laplace transform for such processes, known as subordinators: The distribution F on $\mathbf{R}^+$ is infinitely divisible if and only if there is a representation

$$\int_{[0,\infty)} e^{-\lambda x} F(dx) = \exp \left[-c\lambda - \int_0^\infty (1 - e^{-\lambda x}) \mu(dx) \right]$$

for some $c > 0$, and measure μ on $(0, \infty)$ satisfying the integrability condition $\int_0^\infty (x \wedge 1) \mu(dx) < \infty$. This condition is useful in determining whether or not a Lévy process is monotonic or not.

We have therefore something of a paradox. In many applications we expect degradation to be both continuous and monotonic, but using Lévy processes, we cannot model both of these conditions simultaneously. In such an application it is necessary to decide which of these properties is more important, and choose a model on this basis.

Within the class of Lévy processes the two processes which are most important are the Wiener process and the Compound Poisson process. These are the basic building blocks from which all Lévy processes are constructed. In reliability modelling, both of these processes have been used extensively in the literature. There are however other Lévy processes which have been used such as the Gamma process applied by Van-Noortwijk (1996) among others and Stable processes used by Hougaard (1986). In terms of degradation modelling we shall focus mainly on the Wiener process and its extensions and the Gamma process, although the compound Poisson process has formed the basis for many degradation processes of the Shock model type.

Firstly, the Wiener process is ubiquitous in the literature on probabilistic modelling. The reason being that it has many useful properties and has been studied extensively. Secondly, the Gamma process, which is less commonly found, but perhaps more important for our purpose, since it addresses some of the shortcomings of the Brownian motion process (see below). The gamma process captures the global features of the physical process and provides enough structure to enable calculations of interest to be carried out.

As was mentioned above, of all of the assumptions underlying Lévy processes, for our applications that of stationary increments seems most restrictive. In certain applications we may decide that to provide a suitable model it is necessary to drop this assumption. Doing this gives us further possibilities for our degradation model.

The most obvious generalisation is to make the assumption that the parameters of a Wiener process may be time and/or state dependent, which defines the Class of Diffusion processes. These processes inherit many of the nice properties of Wiener processes, but are much more general and flexible for modelling. The analysis of such processes is more difficult than for a simple Wiener process, but much work has been carried out and many results are available. Karlin and Taylor (1981, Ch 15) provide an extensive survey of the elementary properties of such processes.

A more general approach however, is to allow the parameters of a Lévy process be dependent on another stochastic process. In this way the resulting process evolves locally as a Lévy process, but the distribution of increments is dependent on another stochastic process. In the particular case where this second process is a Markov process, the resulting bivariate process is known as a Markov Additive process. This allows a Markovian covariate process to be incorporated into a Lévy process degradation model. These processes however are less tractable than the preceding examples, and much less work has been done on their application to reliability.

In what follows we survey the properties of Wiener processes and gamma processes before considering their application as univariate degradation processes. We conclude this chapter by looking at how the models can be extended to the case where a covariate process is involved and how useful the models are in terms of maintenance optimisation.

2.3 The Wiener Process as a Degradation Model

2.3.1 Relevant Properties of Wiener processes

The Brownian motion process is perhaps the most important stochastic process in all of probability theory. It falls into all of the main classes of stochastic processes, namely: Markov processes, Lévy Processes, Martingales, Right Processes (Rogers and Williams, 1994). This means that there are a great many results available so that many useful calculations can be made.

The definition of a Brownian motion is given by

Definition 2.3.1 (Standard Brownian Motion)

A Stochastic Process $\{B_t : t \geq 0\}$ is said to be a Brownian motion if

- (i) $B_0 = 0$ a.s.
- (ii) The map $t \rightarrow B_t$ is a continuous function of $t \in R_+$
- (iii) For every $t, h \geq 0$, $B_{t+h} - B_t$ is independent of $\{B_u : 0 \leq u \leq t\}$, and has a Gaussian distribution with mean zero and variance h . $\square$

However, the Brownian motion clearly has some undesirable properties for use as a degradation process. The main problem is that the expected value of the process at any time is zero. This means we cannot use it as a degradation process in general, since to do so would implicitly assume that a system neither degrades or improves on average as time passes. While it is plausible such a system could exist (i.e. when the failure of a system is subject to a random environmental factor, and the system otherwise does not deteriorate), it is not common. Another difficulty, which is inherited by most models based on the Brownian motion process, is that the sample paths are non-monotonic.

The Wiener process is a transformation of the Brownian motion that is a more suitable as a degradation model. If B_t is a Brownian motion, as defined above, we define a Wiener process W_t as the following function of a Brownian motion

$$W_t = \mu t + \sigma B_t$$

This process is essentially a scaled Brownian motion about the line μt , which represents the mean value of the process at time t . This process is clearly a much more flexible modelling tool than a standard Brownian motion. In fact this process is the most general continuous Lévy process available. It is clear however that the Wiener process inherits the non-monotonicity property of the Brownian motion. While this is not ideal for degradation modelling, it has been found that the Wiener process is a suitable model in many cases (See for example Newby, 1991 and 1998)

In addition it is clear from the definition that this process is a Lévy process. The fact that increments are normally distributed means that any estimation and inference based on data from this process will benefit from Normal distribution theory, and so will take place in the standard setting.

Then, if we use the Wiener process as a degradation model, there are various quantities of interest to us. Firstly, if we consider a threshold failure model, failure will occur at the hitting time of the threshold. If we assume the Wiener degradation process is defined as above, and the failure threshold is c , then the failure time distribution is an inverse Gaussian distribution with density

$$g(t) = \frac{c}{\sqrt{2\pi\sigma^2 t^3}} \exp\left[-\frac{(c - \mu t)^2}{2\sigma^2 t}\right]$$

This density is relatively simple, and provides a flexible model for many. A detailed discussion of the properties of the inverse Gaussian distribution may be found in the book of Chhikara and Folks (1989).

The definition of the Wiener process makes it clear that, given an initial value $W_0 = x$, the distribution of a future value of the process, W_t is a normal distribution with mean $x + \mu t$ and variance $\sigma^2 t$. Hence, for a known initial value of degradation we can determine the probability the system has reached a given level of degradation as

$$P(X_t < c | X_0 = x) = \Phi\left(\frac{c - x - \mu t}{\sigma\sqrt{t}}\right)$$

which is a special case of the more general results, which is easily shown as a consequence of the independent and stationary increments property of the Wiener process:

$$P(W_t - W_s < \Delta | W_s = w_s) = \Phi\left(\frac{\Delta - \mu(t-s)}{\sigma\sqrt{t-s}}\right)$$

There are many other useful results that are of interest when dealing with the Wiener process degradation model. In particular the joint density of W_t and its maximum variable defined by $M_t = \sup\{X_s : s \leq t\}$, given $W_0 = x$, is given by

$$f_{W_t, M_t}(y, m | x) = \frac{2(2m - y - x)}{\sqrt{2\pi\sigma^6 t^3}} \exp\left\{-\frac{(y - x - \mu t)^2}{2\sigma^2 t}\right\} \exp\left\{-\frac{2(m - x)(m - y)}{\sigma^2 t}\right\}$$

where y represents the future level of degradation, and m represents the maximum level of degradation attained during the period, for an initial level of degradation x . The proof of this, and other, results can be found in appendix A2.

It is clear from these results that the Wiener process model is very tractable. Most of the important results we might require are readily available, or can be easily derived. This is the main reason the Wiener process model is so commonly used.

As we have said the Wiener process has some problems which can make it unsuitable as a degradation model. In particular, being a Lévy process, it has independent and stationary increments, the problems associated with which, have been described above. A natural generalisation of the Wiener process is to the class of processes known as diffusions.

Diffusion processes are a general class of continuous processes which satisfy the strong Markov property. They are generally not Lévy processes, since the distribution of increments is not usually stationary. However, because they are Markov processes, they are tractable and a large theory has built up around them giving many results of use in reliability and degradation modelling. (Karlin and Taylor, 1991)

Manipulation of the above general definition of a diffusion, results in a practical definition of a diffusion process which is most easily expressed in terms of a stochastic differential equation. So, in general a diffusion with drift parameter $\mu(t, x)$ and

diffusion parameter $\sigma(t, x)$ is a stochastic process X which is a weak solution of the Stochastic Differential Equation (SDE) (Rogers and Williams, 1987)

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dB_t$$

Where B is a standard Brownian motion process. This process is the obvious generalisation of the Wiener process. The difference being that the parameters of the process are now permitted to depend on the time parameter t and the current state of the process X . It is clear that the local properties of the diffusion process are similar to those of a Wiener process.

Further information on the definition and application of Stochastic differential equations may be found in, among others, Rogers and Williams (1987). We shall however not use the theory of SDE's in this thesis, as we will not be dealing further with general diffusion processes. A classical treatment of diffusion processes (which does not use the SDE approach) may be found in Karlin and Taylor (1981). We do not use the theory of SDEs in this thesis, and Karlin and Taylor (1981) is more than adequate for our purposes.

For degradation modelling we are more likely to be concerned with the case in which the diffusion parameters are dependent on the current level of degradation. In this case we have a time-homogeneous SDE

$$dX_t = b(X_t)dt + \sigma(X_t)dB_t \tag{2.3}$$

This simplification makes the analysis of these processes slightly easier. There are of course some examples of systems whose degradation will exhibit a time dependent component, but the analysis in this case is similar to the homogeneous case.

The properties of diffusion processes are well known, so we will concentrate on the specific features of diffusions that are of interest to failure modelling. As with the Wiener process it is often possible to determine the distribution of the hitting time of a point in a diffusion process. In general this is found to be a generalised inverse Gaussian distribution. Further information on this and similar results may be found in the work of Barndorff-Neilson et al (1978).

In addition to distributional results much work has been done in computing functionals of Diffusion processes. In particular, if $X(t)$ is a diffusion process, and T represents the Hitting time of the set $\mathbf{R} \setminus (a, b)$, the functional

$$w(x) = E \left[\int_0^T g(X(s)) ds \mid X(0) = x \right]$$

is found to be a solution of the differential equation

$$\mu(x) \frac{dw}{dx} + \sigma(x) \frac{d^2w}{dx^2} = -g(x) \quad \text{for } a < x < b$$

with boundary conditions $w(a) = w(b) = 0$. This is the classical Fokker-Planck equation, and can generally be solved using analytical methods. Solving and letting $b \rightarrow -\infty$ gives expected value of the integral up to the hitting time of a . This functional is of importance, since if we let $g(x)$ be the instantaneous cost of being in state x , this functional represents the total cost incurred up to failure, which we assume occurs when the process hits point a . Clearly, in many cases this differential equation will have an explicit analytical solution. In more complex cases it is a relatively simple matter to obtain a numerical solution to this type of equation.

Similar results to those described above may be obtained for many other functionals. Some of these results may be found in the book of Karlin and Taylor (1981), which also described how the results may be derived.

2.3.2 Degradation models based on the Wiener Process

The diffusion model has been used extensively in the literature as a degradation model. The main models which have been used are described below.

Non-monotonicity is the main drawback of using a diffusion process as a degradation model. The problem can be reduced if the drift of the process is increased relative to the variance, but in the case where we are fitting a model to data the drawback is obvious. If the degradation process is monotonic increasing (which is a very reasonable assumption), all of the increments which we observe will be positive. If we try to fit a normal distribution to this data we will clearly not be proceeding according to best statistical practise, instead we are forcing the data to fit the model.

Having said this however, there are certain cases where these drawbacks do not apply. Sobczyk (1987) provides justification for using a diffusion process in the case of Crack

Growth governed by the Paris-Erdogan law, while Newby (1998) and Whitmore(1995) provide details of estimation and inference for crack growth processes under this model. In addition, there are cases, particularly with regard to deterioration of biological systems (such as the human body), where there is the possibility of non-monotonic degradation processes occurring. This can be seen for example in the model of CD4 counts of Doksum and Normand (1995). Further, it is often the case that the degradation properties of a system are not particularly well known, in which case a Wiener process, being the limiting process of the random walk model, can be used as an approximation.

The Wiener Process $W_t = \mu t + \sigma B_t$, which retains most of the structural properties of the Brownian motion process, is the most common diffusion process found in the literature on reliability and failure modelling. This process is the basis for most of the work done on estimation and inference on degradation processes, and its tractability mean that useful results are relatively easy to come by. We shall now consider some important examples of such models.

As we have mentioned above Doksum and Normand (1995) present a model based on the Wiener process to model the CD4 marker for a HIV infected individual. Under their assumptions the marker is assumed to follow a geometric Brownian motion, so that $\log X_t$ is a Wiener process, namely

$$\log\left(\frac{X_t}{X_0}\right) = \mu t + \sigma B_t$$

where X_t represents the CD4 count, which measures degradation, with X_0 being the initial CD4 count. Using this model, standard estimation and inference procedures are carried out and predictions are made of latency time and time until development of AIDS (Time to Failure). However, as with many applications outside the field of maintenance, no attempt is made to consider how we may intervene to affect the degradation process.

A similar model is proposed by Whitmore (1995) who assumes the degradation process follows a Wiener process, but makes the additional assumption that inference is based on observations which are subject to some measurement error. So if X represents the true degradation process, and Y represents the observed degradation we have

$$\begin{aligned}X_t &= \mu t + \sigma B_t \\Y_t &= X_t + \varepsilon_t\end{aligned}$$

where $\varepsilon_t \sim N(0, \nu^2)$. Clearly, the assumption of normal errors and normally distributed increments simplifies the analysis of this problem to a great extent. Whitmore (1995) gives results concerning the estimation and inference of degradation from such processes. In this model, and that of Doksum and Normand (1995) failure is defined according to a threshold failure model, and so in each case it is possible to derive the time to failure distribution under the model.

A more complex model which uses a more general diffusion process was presented by Lemoine and Wenocur (1985). They assume that the degradation of a system is modelled by a diffusion process satisfying the time homogeneous SDE (2.3). The system is also assumed to be subject to failure according to killing at rate $k(x)$ (Karlin and Taylor, 1981) when the system state is x , and is also defined to fail when the degradation reaches a predetermined threshold value. The survival function of the system is then given by

$$P^x(T > t) = w(x, t) = E^x \left[\exp \left\{ - \int_0^t k(X(s)) ds \right\} 1_{\{T > t\}} \right]$$

Since X follows a diffusion process, it is shown in Karlin and Taylor (1981, Ch. 15) that $w(x, t)$ is a solution to the partial differential equation

$$\frac{\partial w}{\partial t} = -k(x)w + b(x)\frac{\partial w}{\partial x} + \frac{1}{2}\sigma^2(x)\frac{\partial^2 w}{\partial x^2}$$

which can be solved numerically (and in some cases analytically) to give the required survival function. Lemoine and Wenocur (1985) then go on to consider various special cases of this model and show that many common reliability models can be derived from this kind of analysis. This model has an advantage over the previous models in that it can incorporate state dependent behaviour in the degradation model, making it more flexible, if more difficult to work with.

Another generalisation of the Wiener process is given by Whitmore and Schenkelberg (1997). This model assumes that degradation follows a Wiener process, but that the time scale may be accelerated so that, to a certain extent, a time dependent factor may be taken into account. While this is less flexible than the model of Lemoine and

Wenocur, it is easier to work with, and for most standard transformations of the time scale, inference and estimation procedures carry over easily from the Wiener case. (See Lee and Whitmore (1993) or Feller (1971, Ch. X.7) for a general description of stochastic process under a randomised time transformation).

There are other models based on diffusion processes, some of which are similar to the above, and others that deal with either maintenance or incorporate covariates in the analysis, which are considered below.

2.4 The Gamma Process as a Degradation Model

As we noted above, a general Lévy process contains two components: a Wiener component and a pure jump process component. Intuitively, if we seek a degradation model which has only positive increments, we should be able to define a process in the same way as Brownian motion is defined, but replacing the normally distributed increments, with increments from a distribution which takes only positive values. If we assume that the distribution of the increment $X_t - X_s$ follows a gamma distribution with density

$$f(x) = \frac{\beta^{\alpha(t-s)} x^{\alpha(t-s)-1} e^{-\beta x}}{\Gamma(\alpha(t-s))}$$

we obtain a gamma process, which is defined as

Definition 2.4.1 (Gamma Process)

A Stochastic Process $\{X_t : t \geq 0\}$ is said to be a Gamma process if

- (i) $X_0 = 0$ a.s
- (ii) X has independent and stationary increments
- (iii) The distribution of $X_t - X_s$ is Gamma with parameters $\alpha(t-s)$ and β .

In the case where $\alpha = 1$ and $\beta = 1$ we call X a standard Gamma process.

This process as it is defined here is not, however, a Lévy process, since it is not generally time homogeneous. For the process X to be a Lévy process we must assume that $\alpha(t) = \alpha \cdot t$, for some constant α . Henceforth we assume that we are dealing with this form of Gamma process.

Let us consider some of the main properties of this process. Firstly it is clearly monotone increasing, which is the behaviour observed in most physical degradation processes. Secondly, from the properties of Lévy processes it is discontinuous, unlike the diffusion models considered above. The gamma process can be thought of therefore as the accumulation of an infinite number of small shocks. This interpretation gives credence to the model, since this is often how degradation occurs. Thirdly, the independence and stationarity of increments of this process mean that this model assumes future degradation is independent of the current level of degradation, and depends only on the period over which the system will be allowed to deteriorate. This is in some ways unrealistic, although models based on the Gamma process should be at least as suitable as those based on the Wiener process since both have the flexibility of two parameters, while the Gamma process has the advantage of monotonic sample paths.

However, just as Wiener processes can be extended to a Diffusion processes, the gamma process can be extended in a similar way. The results of these extensions in this case are not as tractable as in the case of diffusions, since there is not a well developed framework for stochastic differential equations involving Gamma processes.

Using the gamma process as a model for degradation, it is straightforward to compute the hitting time distribution of a point. Let H_x be the hitting time of point x , and assume that $X(t)$ is a gamma process started at zero. Then it is clear that $P(H_x \leq t) = P(X(t) > x)$, since the process is monotonic. Then it is simply found that

$$P(H_x \leq t) = P(X(t) > x) = \int_x^\infty \frac{\beta^{\alpha t} y^{\alpha t - 1} e^{-\beta y}}{\Gamma(\alpha t)} dy = \frac{\Gamma(\alpha t; \beta x)}{\Gamma(\alpha t)} = F(t | x)$$

where $\Gamma(\alpha; x) = \int_x^\infty t^{\alpha-1} e^{-t} dt$ is the incomplete Gamma function. So, if we assume that X_t measures degradation, and the system fails when degradation first reaches the point c ,

it is clear that the expression above represents the failure time distribution function for the system. It is possible to obtain the density of the hitting time distribution as

$$f(t) = (\log(\beta(c-x)) - \alpha \Psi(\alpha t)) \cdot F(t, c-x) + \frac{\beta(c-x)G}{\Gamma(\alpha t)}$$

Where G is a Meijer G function (See Erdelyi, 1954b) and Ψ is the derivative of the natural logarithm of the Gamma function. In this particular case we have

$$G = G_{2,3}^{3,0} \left(\beta(c-x) \left| \begin{array}{ccc} 0 & 0 & 0 \\ -1 & \alpha t - 1 & -1 \end{array} \right. \right)$$

This density was found using the symbolic manipulation package MAPLE. This density is not suitable for carrying out computations regarding the gamma process, and it is better to work with the much simpler distribution function.

The first failure model to use the gamma process was that of Abdel-Hameed (1975). However, this model is based on a slightly different model in which the wear at time t , is defined to be a random variable $X(t)$, where $X(t) \sim \Gamma(\alpha(t), \beta)$. The idea behind the model is that there is a known probabilistic relationship between the degradation of the system and its failure. This is expressed simply by the survival function $\bar{G}(x) = P(W > x)$, where W is the wear at failure. The Gamma process is then hypothesised as the relationship between the age of the system and the degradation of the system. Then if we let $f_t(x)$ be the probability density function of $X(t)$, it is clear that the survival function of the time to failure T is given by

$$\bar{F}(t) = \int_0^{\infty} f_t(x) \bar{G}(x) dx \quad (2.4)$$

It is clear that (2.4) is simply a mixture of the $G(x)$ distributions, where the mixing distribution is Gamma. Abdel-Hameed (1975) then provided conditions under which the lifetime distribution properties of $G(x)$ are inherited by $F(t)$.

The model is intuitively very appealing, since it is often the case that the relationship between degradation and failure are well known, but the time until a critical amount of deterioration has occurred is stochastic. However, the model has several shortcomings. Firstly, it assumes that the degradation process is continuously monitored, and that the degradation process can be observed without error. Both of these assumptions may be

unrealistic for certain applications. Secondly. The underlying gamma process has stationary increments, so that the process evolves according to the same probabilistic laws, independently of the current level of degradation and the current age of the system. This assumption is very restrictive.

Another model based on a gamma process is that of Wenocur (1988). This model is a generalisation of diffusion processes, but with the Brownian motion replaced by a standard gamma process. In stochastic differential equation form the model may be written as

$$dX_t = b(X_t)dt + \sigma(X_t)d\gamma_t$$

where γ_t follows a gamma process with parameters $\alpha=\beta=1$. This model is a generalisation of the gamma process defined above, and appears to be a much more flexible modelling tool. There are some problems however, since the ‘drift’ and ‘variance’ of the process X_t are possibly dependent on X_t , thus losing the property of stationary increments, and any results based on the theory of Lévy processes. The process X_t is still a Markov process however, and various results have been obtained by Wenocur (1988) who derived the backward equation for the Kac Functional

$$w(x, t) = E^x \left[\exp \left\{ \int_0^t -k(X(s))ds \right\} \cdot f(X(t)) \right]$$

This is important, since many quantities of interest can be put into this form, and may be able to be derived by solving the associated integro-differential equation. It was found however, that the generator of the process does not take a nice form, and so many calculations which can be carried out easily in the diffusion case, are impossible here.

It remains to be seen if this process offers any advantage over the original gamma process defined above. The diffusion process defined by a SDE can be interpreted intuitively as a deterministic process with an additive ‘error’ component, so that the process fluctuates around its expected value (drift component) in a random manner. In the case of the SDE driven by a standard gamma process, the drift component cannot be interpreted as the average value of the process, and is instead interpreted as the ‘baseline’ rate of change of deterioration at level x . This gives perhaps a more flexible model than the basic Gamma process, but at the cost of increased complexity. It is possible that a similar effect can be achieved by adjusting the shape parameter $\alpha(t)$, so

as to increase the rate of increase of the process by increasing the average size of the jumps. It is clear however that making the shape parameter of the basic Gamma process dependent on the system state would result in a model more complex than that of Wenocur(1988). The fact that no papers using this variant of the Gamma process model have been published since its origination may be indicative of its lack of tractability and practical application.

2.5 Models of Degradation in the Presence of Covariates

As was noted in Chapter 1, it is common to find that a system characteristic which provides direct information on the failure characteristics of the system is not observable. In such cases it may be possible to consider a covariate process as an alternative to observing degradation. To make allowance for this it is often desirable to jointly model the degradation of the system (which is unobservable) along with a (possibly multi-dimensional) covariate process.

In many circumstances a similar problem arises in that the degradation of the system is only observable subject to some kind of measurement error. This case may be included under the structure of covariates outlined above, if the measurement taken from the system is assumed to be the covariate, while the true degradation is represented by the unobservable degradation process. In this case the relationship between the covariate and system variable is determined by the observation error structure.

The above factors make the modelling of degradation in the presence of covariates a much more complex task. We can no longer assume that the process involved is stationary, and we must make allowance for the evolution of the state and the covariate over time. The resulting modelling becomes complex even under the most simple assumptions.

There has been a great deal of interest recently in this subject because of its relevance to the subject of AIDS epidemiology. In that field, covariate processes are known as marker processes and generally refer to measurements taken from the human body which are of some relevance to the unknown condition of the AIDS sufferer (because time since infection is unknown). In comparison there has been relatively little work

done specifically in the area of maintenance modelling. (Desmond, 1985 and Singpurwalla, 1995 give detailed overviews of this area)

Most of the work in this area, from both bio-statistical and reliability perspectives has involved the use of hazard rate modelling. This is unsurprising, since much of the work can be seen as a stochastic generalisation of the Cox Proportional Hazards model, of Cox (1972) in which regression is carried out in the presence of covariates using a hazard function of the form

$$h(t; \mathbf{z}) = \lambda(t) \exp(\boldsymbol{\beta}^T \mathbf{z})$$

where $\lambda(t)$ represents the baseline hazard function, $\boldsymbol{\beta}$ is a vector of parameters and $\mathbf{z}$ is a vector of covariates.

Jewell and Kalbfleisch (1996) consider the case where the hazard rate is related to the covariates by the formula

$$h(t | H_t^X) = h_0(t) + \beta X(t)$$

where $H_t^X = \{X(s) : 0 \leq s \leq t\}$ represents the history of the covariate process X , β is an unknown parameter and $h_0(t)$ represents a baseline hazard function. They consider some special cases including the case of a Poisson Marker process and a Multivariate zero-one marker process, and prove some results relating to residual and past life.

A similar model is given by Fusaro et al (1992) and Shi et al (1996) who consider a hazard function of the form

$$h(t | H_t^X) = h(X(t))$$

so that the survival of the item is completely dependent on the value of the covariate, with no other time varying factors allowed. Shi et al (1996) consider various marker processes $X(t)$ including the processes (which are applicable to the case of CD4 count modelling for AIDS survival data)

$$\begin{aligned} \text{(i)} \quad X(t) &= [a + bt + B(t)]^4 \\ \text{(ii)} \quad X(t) &= [a + bt + U(t)]^4 \end{aligned}$$

Where $B(t)$ is a Brownian motion, $U(t)$ is an integrated Ornstein-Uhlenbeck Process (Karlin and Taylor, 1981) and $[a, b]^T$ follows a bivariate normal distribution. It is

unlikely that these covariate models (i) and (ii) above have any application to this area of reliability.

A similar approach is taken by Myers (1981) who considers the case in which the hazard rate of the system is a quadratic form, dependent on a multidimensional diffusion covariate process.

$$h(t) = Z_t^T B Z_t + 2C^T Z_t + d$$

where Z is a vector of covariates, B is a positive definite matrix and C is a parameter vector. Myers makes the further assumption that the covariate Z follows the m -dimensional diffusion process defined by

$$dZ_t = a(t)Z_t dt + b(t)dB_t$$

where B_t is an m -dimensional Brownian motion and a and b are continuous functions of t . In this case it is found that the Survival function for the system may be written in terms of the solution of a Matrix Ricatti equation.

A review of this type of hazard modelling may be found in Yashin and Manton (1997), who also consider a model similar to that of Myers (1981) described above.

A model which deals more directly with this type of problem is that of Whitmore, Crowder and Lawless (1998). The underlying structure of the model is based on a bivariate Wiener process $(X(t), Y(t))$. The components of the Wiener process $(X(t), Y(t))$ represent the (unobservable) degradation process and the covariate (or marker) process, respectively. To simplify the modelling, it is further assumed that the system is observed for a fixed time t , at which time it is inspected and found to be either failed or operating. Using standard results from the theory of Wiener process, a likelihood is created and this allows for estimation and inference on the process.

The model of Whitmore, Crowder and Lawless (1998) explicitly assumes that the covariate is correlated to the degradation process, but the simplification of a fixed time frame makes the model unsuitable for considering systems over more than one time period or operating cycle. In chapter 4, we shall propose a model which extends the results of Whitmore, Crowder and Lawless, by allowing optimal inspection policies to be computed.

A more indirect method of dealing with covariates is given by Rishel (1991) and extended by Lefebvre and Gaspo (1996a,b). Rishel (1991) proposes a model in which the wear of a system is related to a stochastic process (which represents environmental variables) by an ordinary differential equation which represents the evolution of system degradation in the presence of the environmental factor. This system is described by the stochastic differential equations

$$\begin{aligned}dX_t &= \rho(X_t, Z_t) dt \\dZ_t &= f(X_t, Z_t) dt + \sigma(X_t, Z_t) dB_t\end{aligned}$$

Where $\rho(x, z)$ is a nonnegative continuous function of x and z and Z_t is a vector valued diffusion process. Under the assumption that the environmental process Z and the wear process X are dependent on a control parameter u , Rishel(1991) and Lefebvre and Gaspo (1996b) go on to consider optimal control policies which minimise the wear of the system. This assumes that the system can be controlled continuously, which is not always the case. In the models we shall consider and our only allowable actions are those which come under the scope of inspection and maintenance, and which therefore take place at discrete time points. For models which are monitored and may be maintained continuously, this control theoretic approach has many advantages. We shall not consider this approach further in this thesis.

A recent paper by Lim (1998) applies regime switching techniques to a related problem in which the failure time distribution of an entity is dependent upon the state of a Markov chain. While this is not directly comparable with the degradation models we have been considering, the underlying structure appears to allow covariates to be taken into account in certain cases. The technique of regime switching models has been applied in the financial literature to the prediction of time series. See Hamilton (1994, Ch. 21) for further analysis of such problems

The regime switching model described above allows the state of a degradation process over the next ‘period’ to depend on the current state of the covariate process. While this partly achieves our aims, in many cases this would not adequately describe the problem at hand. To generalise further we must consider Markov Additive Process of Cinlar (1972b).

Markov Additive processes may be regarded as generalisations of Lévy processes, the Brownian motion and Gamma processes both being special cases. We do not give a definition here, as we shall not consider these processes outside of this section.

The diagram below (figure 2.5.1), which is modified from Cinlar (1977) shows the type of situation that Markov Additive processes may be used to model:




Figure 2.1 – Example of a Markov Additive Process Degradation Model

This situation may be thought of (for example) as a degradation process which evolves according to a Gamma process when the covariate is in state 1, evolves according to a Compound Poisson Process while the covariate is in state 2, and remains constant when the covariate is in state 3 (which may correspond to the system being under repair, say). While this situation may seem rather unlikely, it shows the potential flexibility of Markov Additive processes as modelling tools.

As we remarked above, Markov Additive processes are a general and flexible modelling tool in that they allow for the presence of one or more covariates in the model. This model therefore, if fully developed, should allow us to take into account the effect of other factors on the failure mechanisms, and therefore gain better understanding of system behaviour so that maintenance effort is better applied. However, although the model appears to be suitable for more general applications than those discussed above, the modelling process and any subsequent optimisation is made more difficult by the increased structural complexity of the process.

This complexity however has not been a barrier to the application of these processes in other fields. In particular Markov Additive processes have been widely used as models for the arrival and departure processes in Storage process (see for example the papers of Kaspi (1984) and Asmussen and Kella (1998)).

It appears, however, that there has been minimal application of MAPs in the field of failure time modelling. Some of the reasons for this are clear. For the most part degradation models have been constructed by those seeking to model the failure time characteristics of a system or a component, without looking at inspection or maintenance of the system. In such cases there has been little interest in modelling a covariate process alongside the degradation process, possibly because the increased cost of data collection and testing is prohibitive. Another possible reason is that it was felt the added complexity of the model was too much to justify the added information it may provide. As has been pointed out by Singpurwalla (1995), there may be some problems in obtaining results about these processes because of an apparent lack of tractability.

2.6 Maintenance and inspection Models

These failure models are constructed to describe systems, so that we can compute the optimal maintenance strategy for the system. While much work has been carried out on optimisation of maintenance and inspection, relatively little of this seems to be applied to these types of degradation model.

For the most part, classical models of inspection and replacement are based on the distribution of the inter-failure times of a system. Most of these may be regarded as extensions of the classical models of replacement such as age replacement and block replacement (See Barlow and Proschan, 1965). By assuming that the system is replaced by an identical system at failure or preventive replacement, the act of replacement is seen as a renewal event. This assumption allows the theory of renewal processes to be applied to replacement problems. In particular the renewal reward theorem (Ross, 1970) allows us to show that the long run average cost per unit time is equivalent to the average cost per cycle. This result is perhaps the most often used in maintenance

modelling, since it reduces the general problem to one in which only the expected cost and expected length of a single cycle needs to be found. As we have said, the literature on this type of modelling is vast, and since it does not directly influence the remainder of this thesis we shall not expand on it here. The reader is referred to the book of Jardine (1973) or the excellent review paper of Valdez-Flores and Feldman (1988).

While the literature on the models described above is vast, the literature on models which directly apply degradation processes is much less commonly found. As before, the models may be broadly split into inspection models, replacement models and maintenance models. (Replacement models are of course a special case of maintenance model, but we differentiate them here as their properties make modelling much simpler).

By inspection models we generally mean models designed to address the question ‘At what times should the system be inspected’. In answering this question using a mathematical model, it is generally necessary to make assumptions about the maintenance or replacement policy being applied to the system, and the relative costs of these policies. For example, the inspection policy instructs us when to inspect the system, and hence obtain information on the level of degradation. Based on this level a rule must be applied determining what maintenance actions will be taken. If no action is to be taken, the inspection is unnecessary and there is no problem.

An example is given by Anderson and Friedman (1977). The basic model proposed assumes that the degradation of a system is modelled as a Brownian motion process B_t , and that costs are incurred continuously at rate $f(B_t)$ per unit time. If, at an inspection time τ , it is found that $B_\tau \in A$ the system operation continues, otherwise it is shut down at cost $\gamma(B_\tau)$. Further, each inspection is assumed to incur a cost $\beta(B_\tau)$. Extensions of this model also consider the case when the inspection is subject to a Gaussian measurement error and when the system degradation is reduced, rather than shut down when $B_\tau \notin A$. The model was extended by Anderson and Friedman (1978) to consider the case when the underlying degradation is a bivariate Brownian motion, one component of which is continuously observable, while the other may only be observed via costly inspections.

Using techniques of quasi-variational inequalities, Anderson and Friedman (1977,78) are able to obtain an optimal sequence of inspections which minimises the discounted

total cost of operating the system. While this analysis is useful, there are problems with applying it in real situations. The work is largely theoretical in nature, and so ignores some important practical factors. Primarily, the system cannot ‘fail’ between inspections, which clearly is unreasonable for most systems. Further, the underlying degradation process is a Brownian motion, which, as we have said, is generally not suitable as a degradation model.

Yeh (1996) also considers the case of optimal inspections. The underlying degradation process is assumed to be a semi-Markov process. By transforming this process into an equivalent Markov process, techniques from Markov decision theory are used to determine optimal inspection intervals. The model is extended to allow for the case of imperfect inspection. This model has a much more application oriented perspective than that of Anderson and Friedman (1977,78), and it would seem that it is general enough to be applied in many situations.

In general optimal replacement policies for deteriorating systems are specified in terms of a control limit policy, with a pre-specified inspection policy. Such policies define a control limit, say r , and assume that the system is replaced if an inspection reveals a level of degradation greater than the control limit, r . In many replacement problems, it is possible to show that the control limit type policy is optimal.

Park (1988a,b) gives an example of this type of model. The models given in Park (1988a,b) differ only in their failure assumptions. The first assumes threshold failure, while the second assumes a wear dependent hazard rate. In both cases the degradation level at time t is assumed to have a known probability distribution. In the first model (Park 1988a) it is assumed that the system is replaced if inspection reveals a level of degradation greater than r , but fails if the level of degradation reaches level c . (See Figure 2.2 overleaf). In this case, failure or replacement define a regeneration point of the degradation process, allowing the renewal reward theorem to be applied. Having computed the expected cost and length of a cycle, an expression for the average cost per unit time is obtained, so that optimal control limit r may be computed. In Park (1988b), similar calculations are made, but the system may now fail at any time, according to a degradation dependent hazard rate. The model of Park (1988a) is particularly important for what follows in the remainder of this thesis. In Park (1988a) the problem of determining a control limit for a given inspection policy is considered. We shall begin by considering a ‘dual’ of this problem (using different methods) of obtaining the

optimal inspection schedule for a given control limit. We further extend these results obtaining optimal non-periodic inspection policies, along with policies for imperfect inspection and maintenance. In particular, our methods allow computation of jointly optimal control limit and inspection policies as a special case of more general maintenance policies.




Figure 2.2 – Example of degradation with a control limit replacement policy and threshold failure model

Another model which is similar to that we shall consider is that of Antelman and Savage (1965). Here two models are considered, one with continuous and one with discrete inspections. Both of these are applied to industrial process control making them unsuitable for degradation modelling. As with Anderson and Friedman's (1977,78) work, the underlying process considered is the Brownian motion and failures are not considered. In each case expressions are found for the costs of operating the system. However, in the case of discrete inspections the resulting functional equation is not solved, and no optimal policies are obtained.

Other models for optimal replacement, not based on inspection policies, have also been considered. In particular Zuckerman (1978) considers optimal replacement times for a system with degradation modelled by a one-sided Lévy process (subordinator). The model incorporates threshold failure, but the threshold itself is modelled by a random variable with known distribution. This is particularly applicable to stress-strength models for loading of structures. In this case, failure is not only determined by the degradation of the structure, but also by the loading applied to it.

As one might expect, maintenance models are much less commonly found, than models for replacement of corresponding systems. The basic reason for this lies in the added complexity of maintenance. In the case of replacement models, intervention in degradation behaviour is characterised by returning the level of degradation to zero, which naturally defines a regeneration point for the process. When more general maintenance actions are allowed, it may be the case that the level of degradation is reduced, but this does not provide a regeneration point for the degradation process.

There are few maintenance models based on ‘degradation’ models. One example is that of Hontelez et al (1996). In this case the degradation of the system is modelled by a Wiener process. The state space is then discretised so that the discretised system is represented by a discrete time Markov chain., and Markov decision theory can then be applied to determine optimal maintenance policies.

Another common model is the so-called ‘Virtual Age’ model, based on Kijima et al (1988). The model assumes that the ‘degradation’ of the system is measured by its ‘virtual age’, which represents the effective age of the system, taking maintenance into account. In this way, each maintenance action corresponds to a reduction in virtual age. Applying techniques from dynamic programming theory integral equations are derived, which can be solved to obtain the costs of operating the system. While this model uses the lifetime distribution of the system rather than a degradation process, it is included here since the methods used are similar to those we shall use in the following chapters.

Stadje and Zuckerman (1992) extend their earlier results by considering two different maintenance models, one allowing only minimal or perfect repair and one allowing a general degree of repair. A similar model is proposed by Dagpunar (1997), who considers the effect of different types of maintenance on a system.

2.7 Summary

The choice of degradation model in any particular situation depends largely on the properties of the system being considered. It is clear however that the choice of model will greatly affect the techniques which must be applied to determine optimal maintenance and inspection policies.

Many degradation models have been proposed whose main aim is to determine the reliability characteristics of a system in terms of that system's degradation. In this application it is clear that a large amount of model complexity is useful, the main purpose of the analysis being to encapsulate as much information as possible into the model. However, the complexity of these models often means they are unsuitable for maintenance modelling.

On the other hand, many models of inspection and maintenance have been proposed which are based on simplified degradation models, for the sake of tractability. In addition, many of the models proposed treat maintenance and inspection as separate entities, when it is clear that they are inter-related.

These observations suggest that there is a need for models which can adequately combine maintenance and inspection with a realistic degradation model. In what follows we consider various models which we believe go some way in addressing these problems. We begin by considering optimal inspection of a general Lévy process degradation model, before extending this to the case of inspection of a system with covariates. Finally we look at the case of joint optimisation of maintenance and inspection. The models considered are similar to those of Park (1988a) and apply methods similar to those of Stadje and Zuckerman (1991).

Chapter 3

Optimal Perfect Inspection Policies

3.1 Introduction

In the following chapter we shall consider optimal inspection policies for systems which can be modelled by a Lévy Process, i.e. a process with stationary and independent increments. We shall also consider a particular example of this type of process: the gamma process.

The general set-up of the system we shall look at is as follows: we consider a system or component whose degradation may be modelled by a Lévy Process. We assume that the level of degradation is not continuously observable, but may be found by inspecting the system/component, at any time of the decision-maker's choosing. On making an inspection, the Decision-maker may decide to replace the system/component or allow it to continue operating until the next inspection. The system/component may however fail if the level of degradation reaches a high enough level. We assume that no other action may be taken and that the replacement policy is predetermined. We further assume that the income from the system is the same irrespective of its level of degradation (and therefore disregard it in what follows) and that replacement is instantaneous. Any loss of income incurred due to failure or replacement is assumed to

be incorporated into the cost of replacement or failure. The general set-up is similar to that considered by Antelman and Savage (1965) and Park (1988a,b).

We begin by looking at the case of periodic inspection and in particular the case of systems with degradation modelled by Wiener and gamma processes. We consider the case of perfect inspection and consider both average cost per unit time and discounted total cost criteria. In each case integral equations are developed giving the expected costs and expected time to replacement for a general degradation model. The development of these equations is similar to that of Stadje and Zuckerman (1991), who consider optimal maintenance strategies using a virtual age model.

Following this, the case of non-periodic inspection is considered in general. Again we consider perfect inspection, but consider only discounted total cost criterion in this case. The main reason for this simply being that the average cost criterion is much more complicated. Applying the theory of Semi-Markov Decision processes optimality equations are derived, and solved numerically, obtaining the optimal inspection policy. The gamma process provides the basis of our examples. In addition, we consider briefly the case of no inspections and of continuous condition monitoring, for comparison purposes. Finally we examine the numerical results from this model and discuss practical issues relating to the model.

3.2 Underlying Inspection and Replacement model

In order to proceed we must make precise the definition of the system we are considering and how inspections and replacements affect the system. The basis for our system is that of a threshold failure model. We assume that the level of degradation is measured at inspection, and on the basis of this a decision is made whether or not the system is to be replaced. The cost of any such replacement may depend upon the state of the system, and the system is replaced immediately on failure.

We make the following assumptions regarding the nature of the system and the replacement policy. The policy is assumed to be fixed in all respects, other than the inspection interval. Throughout this chapter we assume that the degradation process is denoted by the stochastic process $X = \{X_t : t \geq 0\}$, and assume the initial level of

degradation is given by $X_0 = 0$. This implies a new system is assumed to have zero level of degradation. If this is not the case a simple transformation of the model results in a model for which the initial level of degradation is indeed zero.

MODEL ASSUMPTIONS M1 (perfect inspection) :

1. We assume the state-space of the system $S = (-\infty, \infty)$ is partitioned into intervals $A_0, A_1, A_2, \dots, A_n$ such that $A_0 = (-\infty, s_0)$, and $A_k = [s_k, s_{k+1})$ for all $k = 0, 1, \dots, n-1$ with $s_k < s_{k+1}$ and $s_{n+1} = \infty$.
2. Each inspection reveals the true state of the system, and the state can only be determined by carrying out an inspection.
3. If, at an inspection $X_t \in A_0$, the system is not replaced and is allowed to continue operating until the next inspection. Each inspection incurs cost C_0 , which may be regarded as the cost of inspection and any loss caused by the system being unavailable during inspection.
4. If, at an inspection, $X_t \in A_i$ for $i = 1, 2, \dots, n-1$ then, the system is replaced at cost C_i . This may be regarded as the cost of replacing a system with degradation level in A_i and any loss incurred due to the system being unavailable. It is assumed the replacement system is identical to the original system, with zero level of degradation.
5. The system is deemed to have 'failed' at the first moment the process hits the set A_n and this failure is immediately observed resulting in an immediate replacement of the system.
6. The cost of replacing a failed system is C_n , and is assumed that $C_i < C_j$ for $i < j$. (That is, the cost of replacement does not decrease as degradation increases)

We assume throughout that the degradation of the system is modelled by a Lévy process X , having increments $X_{t+\tau} - X_t$ with probability density function $f_\tau(z)$, which are independent of t by the Lévy property. When $X_t = x$, we denote the density of $X_{t+\tau}$ given $X_t = x$ as $f_\tau(z|x)$.

Let us now consider in detail the implications of these assumptions carefully. The definition of the replacement policy implies that there are a number of distinct phases which the process goes through, before reaching the final failure level.

For example, we may have a system that has four different levels for costs. We may consider A_0 to represent the situation in which the system operates effectively, A_1 represents a situation in which the system operates effectively, but is close to a possible failure and so should be repaired at low cost, to restore it to a good-as-new position. A_2 represents the situation in which the system is very close to failure and can no longer be repaired, but must be replaced. A_3 represents failure, in which case the system must be replaced, but additional costs are incurred, perhaps for loss of production, compensation etc. It is clear that this general situation covers many different possible applications. It is probable that the most common case would be the case where $n = 2$ and there are three regions, one representing effective operation (A_0), one for planned replacement (A_1) and a third for failure (A_2). We shall consider this case in a numerical example below. As we have said, this is similar to the case considered by Park (1988a) who considers the optimal wear limit replacement policy for a general degradation process, under the assumption of known periodic inspection times. Figure 3.1 below, graphically shows the operation of this type of inspection policy. We shall consider this model in our gamma process example of section 3.6.




Figure 3.1 – Example of a three region replacement policy

In addition, assumption 4 is also important because it assures us that the process is replaced immediately the process hits set A_n . In the case of non-monotonic degradation processes there may be the possibility of failure having occurred, but being ignored because the process had left state A_n before the time of the next inspection. This implies

that in the case of a general degradation model, we must not only consider the state of the process at inspection, but also the values of the process between inspections. The simplest way to do this is to introduce a supplementary process M_t , where M_t represents the maximum value of the process X_t in the period $[0, t)$, formally $M_t = \sup\{X_u : 0 \leq u \leq t\}$. Having done this, the actions taken after an inspection are completely determined by the state of the bivariate process (X_t, M_t) . Alternatively, depending on the circumstances, it may be more convenient to introduce the supplementary variable $H_{A_n}^x$, which is the hitting time of the critical set starting from state x . (See Rogers and Williams, 1994). The derivations which follow emphasise these points in more detail.

3.2.1 Model for periodic inspection

We begin by considering the optimal periodic inspection policy for a system whose degradation can be perfectly observed, at any time. By periodic inspection we mean only policies which are of the form ‘Inspect the system every τ time units’. This means that there is an inspection schedule $\sigma = \{\tau, 2\tau, 3\tau, \dots\}$, so that inspections continue at this fixed time interval, until the system fails or is replaced, at which point we reset time to zero, and continue with this inspection schedule, so that an inspection occurs τ time units after replacement/failure.

This policy is simple to understand and apply. Because there is no adaptive component, inspections can be planned well in advance. In some cases it may be advantageous that the timing of maintenance and inspection is the same for all components, because fixed costs may be reduced. Additional to these practical benefits, a static policy also avoids the complications of Dynamic programming arguments, making the solutions to these problems easier to obtain.

The disadvantage of such a policy is that it is fixed. Once the system is running it does not matter what we observe, the inspection intervals remain the same. This policy is generally sub-optimal when compared to a more general state dependent policy, which takes into account the level of degradation of the system. The policy may result in too frequent inspections in the early life of the system, and too infrequent inspections as the

system nears failure. We consider the difference in costs of these static and dynamic policies later in this chapter.

3.2.2 Model for non-periodic inspections

In the more general case of non-periodic inspections, we still apply model assumptions M1 as above. We amend the inspection policy, so that we now consider a stationary state-dependent policy. This policy is of the form:

‘We observe the state of the degradation process to be x . If $x \in A$ replace the system, but if $x \notin A$, allow the system to continue operating and inspect again after $\tau(x)$ time units ‘

This policy is similar to that used by Antelman and Savage (1965), who consider optimal inspection and replacement of a system modelled as a Brownian Motion with zero drift. This policy is adaptive, in the sense that the inspection intervals change as the degradation level changes. It is perhaps more suitable for systems which are safety critical. In addition, the policy should be less costly than the periodic policy considered above, because the inspections take into account the current level of degradation, and so are optimal at each stage rather than only at the beginning of the operation period. It is expected therefore that savings are available using this policy over periodic alternatives.

The optimal policy in this case is defined completely by the function $\tau^*(x)$. We apply the standard approach of Dynamic programming, and consider the discounted total cost of applying the policy $\tau(x)$. Applying a standard policy improvement routine to an arbitrary discretisation of the problem, easily provides an optimal policy.

3.2.3 Model for No Inspections and Continuous Condition Monitoring

The model described above has been defined for the case when the system is inspected by periodic perfect inspections. In some cases however, periodic inspections may not be applicable, and a limiting case of these inspections may be optimal. The two such limiting cases are the case where no inspections occur (the system being simply

replaced on failure), and the case where there is some form of continuous condition monitoring.

The first of these cases is likely to occur for systems that have low replacement cost in relation to the cost of making an inspection. An example of this would be a component of a system, which is cheap and easy to replace, but inspection required the system to be shutdown, so that the state of degradation of the component could be ascertained. The condition monitoring case is likely to be applicable when the replacement cost of the system is extremely high, and inspections are thus relatively inexpensive. We treat these two cases using similar methodology to those of the more general cases already discussed.

3.3 Optimal Periodic Inspection: Average Cost Criterion

To compute the optimal inspection policy, which in this case is defined completely by the number τ , we apply the Renewal Reward Theorem (Ross, 1970). This is done using standard arguments by considering the expected cost and expected length of a cycle, where a cycle is defined as the time between replacements of the system (either preventive or corrective). Defining the start of each new cycle as a regeneration point, we can apply the renewal reward theorem to calculate the average cost per unit time over an infinite horizon.

Application of this technique implies we must assume that the lengths of cycles are independently and identically distributed. This assumption holds only if each replacement system has initially zero degradation, and is identical to the original system in its degradation characteristics. In this chapter we assume this is the case throughout. In addition assumptions 4 and 5 of M1 ensure that the replacement system is identical to the original system, with zero level of degradation. This implies that replacement may be regarded as a regeneration point, so that we may regard the process as being a renewal reward process with renewals occurring at replacement of the system. Following Ross (1970), if $C(t)$ represents the cost incurred up to time t , then the limiting average cost per unit time is given by

$$E\left[\frac{C(t)}{t}\right] \rightarrow \frac{E[X]}{E[Y]} \quad \text{as } t \rightarrow \infty$$

Where $E[X]$ and $E[Y]$ represent the expected cost per cycle and expected length of a cycle, respectively.

3.3.1 Derivation of Expected Cost per Unit Time

We define $V(x, \tau)$ to be the (random) cost per cycle for a system with initial level of degradation level x and inspection interval τ . Similarly, let $L(x, \tau)$ be the (random) length of a cycle for a system under the same conditions. Although, at this time, we are only really interested in obtaining the average cost and cycle length for a system which starts with zero level of degradation, we find it advantageous to use a formulation which is state dependent. This is for computational reasons and because it allows us to generalise the results easily to a state-dependent non-periodic inspection policy. In addition, we note that the terms derived above represent an abuse of notation. The random variables V and L are not functions of x and τ , and the functional notation is used simply for convenience in further development, and to simplify the following derivations.

Following the above discussion we define $v(x, \tau) = E[V(x, \tau) \mid X_0 = x]$ to be the expected cost until replacement, and $l(x, \tau) = E[L(x, \tau) \mid X_0 = x]$ to be the average time until replacement for a system with initial degradation level x , assuming that we apply inspection policy τ . By the renewal reward theorem we may then find the expected total cost per unit time as

$$\mathbf{C}(0, \tau) = \frac{v(0, \tau)}{l(0, \tau)}$$

We now develop integral equations for the expected cost and length of a cycle, enabling the optimal inspection policy to be obtained, under the average cost criteria.

Expected Cost per Cycle

Using a standard dynamic programming argument, we can express the cost until replacement as the sum of the cost incurred at the next inspection and the cost incurred

thereafter. Let us assume that X_t represents the degradation process, and M_t represents the maximum value of that process over the period $[0, t]$. Since we are dealing only with Lévy processes, the results are independent of the current time, and so, without loss of generality, we assume that the current time is $t = 0$, and condition all variables on the event $X_0 = x$. (1_A is the indicator function of the set A)

$$V(x, \tau) | X_\tau, M_\tau = [C_0 + V(X_\tau, \tau)] 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \sum_{i=1}^{n-1} C_i 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + C_n 1_{\{M_\tau \in A_n\}} \quad (3.1)$$

where we have expressed the cost function as the sum of the cost in the event of continuing (i.e. not replacing the system) and the costs for each possible replacement event. The expected cost is thus

$$\begin{aligned} E(V(x, \tau) | X_\tau, M_\tau) &= E \left(V(x, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \sum_{i=0}^{n-1} C_i 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + C_n 1_{\{M_\tau \in A_n\}} \right) \\ &= E(V(x, \tau) | X_\tau, M_\tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \sum_{i=0}^{n-1} C_i 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + C_n 1_{\{M_\tau \in A_n\}} \quad (3.2) \\ &= v(X_\tau, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \sum_{i=0}^{n-1} C_i 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + C_n 1_{\{M_\tau \in A_n\}} \end{aligned}$$

Define $f_\tau(y, m | x)$ to be the joint density of X_τ and M_τ , conditional upon $X_0 = x$. Taking expectations of the above expression with respect to the joint density of X_τ and M_τ we get

$$\begin{aligned} v(x, \tau) &= E(E(V | X_\tau, M_\tau)) \\ &= E \left(v(X_\tau, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \sum_{i=0}^{n-1} C_i 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + C_n 1_{\{M_\tau \in A_n\}} \right) \quad (3.3) \\ &= \int \int_{A_0 \cup S \setminus A_n} v(y, \tau) f_\tau(y, m | x) dm dy + \sum_{i=0}^{n-1} C_i P^x(X_\tau \in A_i, M_\tau \notin A_n) + C_n P^x(M_\tau \in A_n) \end{aligned}$$

Where S represents the whole state space, and superscript x indicates probabilities are conditional on $X_0 = x$.

As we have mentioned in previous chapters, a failing (in certain circumstances) of many non-monotonic degradation models is that they allow for some unrealistic consequences. The fact that the process can decrease means that, however unlikely, it is possible to have a negative level of degradation, which would correspond to an improvement beyond what we would regard as a new system. This is a shortcoming of the degradation model that we would like to remove from our model for inspections. To

this end, we assume that, for $y < 0$, $v(y, \tau) = v(0, \tau)$, so that we assume a negative level of degradation is equivalent to zero degradation. Of course in some applications it may be that one would like to consider the case where improvement beyond a new system is allowed, but we do not consider this case here.

From our initial assumptions we know that the set A_0 has the form $(-\infty, s_0)$, so we can rewrite the integral in the previous expression as

$$\begin{aligned}
 & \int_{A_0 \setminus A_n} \int v(y, \tau) f_\tau(y, m | x) dm dy \\
 &= \int_{-\infty \setminus A_n}^0 \int v(y, \tau) f_\tau(y, m | x) dm dy + \int_0^{s_0} \int v(y, \tau) f_\tau(y, m | x) dm dy \\
 &= v(0, \tau) P^x(X_\tau < 0, M_\tau \notin A_n) + \int_0^{s_0} \int v(y, \tau) f_\tau(y, m | x) dm dy
 \end{aligned} \tag{3.4}$$

Now define $c(x, \tau)$ to be expected cost incurred at the next inspection, which corresponds to the non-integral terms in the above expression, giving

$$c(x, \tau) = \sum_{i=0}^{n-1} C_i P^x(X_\tau \in A_i, M_\tau \notin A_n) + C_n P^x(M_\tau \in A_n) \tag{3.5}$$

And define $u(x, \tau)$ to be the probability that there is a negative level of degradation, and the system does not fail, so that

$$u(x, \tau) = P^x(X_\tau < 0, M_\tau \notin A_n) \tag{3.6}$$

Substituting these functions into the integral equation and interchanging the order of integration, we obtain:

$$\begin{aligned}
 v(x, \tau) &= c(x, \tau) + v(0, \tau) u(x, \tau) + \int_0^{s_0} \int v(y, \tau) f_\tau(y, m | x) dm dy \\
 &= c(x, \tau) + v(0, \tau) u(x, \tau) + \int_0^{s_0} v(y, \tau) K_\tau(y | x) dy
 \end{aligned} \tag{3.7}$$

Where $K_\tau(y | x)$ is given by

$$K_\tau(y | x) = \int_{S \setminus A_n} f_\tau(y, m | x) dm \tag{3.8}$$

This function represents the probability density that the process is in state y at the next inspection, and the process does not hit the set A_n before the next inspection, given the

current state is x and the inspection interval is τ . Whether or not this integral can be computed in closed form depends largely on the nature of the underlying degradation process. It is known that using the Wiener process degradation model, this integral can be computed in closed form. (See appendix A2)

The integral equation above has general form

$$v(x) = p(x) + q(x)v(0) + \int_0^s v(y)K(y, x)dy \quad (3.9)$$

In general this is a Fredholm integral equation of the second kind. The nature of the functions c , u and K mean, however, that the equation is unlikely to have a closed form solution. In general numerical methods have to be used, but fortunately there are simple effective methods available.

Expected Length of a cycle

To calculate the average cost per unit time to replacement it is necessary to calculate the expected time to replacement defined by $l(x, \tau) = E[L(x, \tau) | X_0 = x]$. We calculate this in a similar fashion to the calculation of average cost, the only further complication being that on absorption in set A_n , a full inspection period τ is not completed, and so we must incorporate the hitting time of the set A_n . However, since A_n is an interval, we need only consider the hitting time of the least point in that interval, which is s_n according to our definition above. Define $H_{s_n}^x$ to be the hitting time of this point starting from $X_0 = x$. Therefore when considering the time to failure we must take into account the distribution of the failure time, which is just the hitting time distribution of the set A_n . Then proceeding as above we find the (random) time until replacement, starting in state x . So conditioning on the values of X_τ , M_τ and $H_{s_n}^x$ we find

$$\begin{aligned} L(x, \tau) | X_\tau, M_\tau, H_{s_n}^x &= [\tau + L(X_\tau, \tau)] 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \tau 1_{\{X_\tau \in A_1 \cup \dots \cup A_{n-1}, M_\tau \notin A_n\}} + H_{s_n}^x 1_{\{M_\tau \in A_n\}} \\ &= [\tau + L(X_\tau, \tau)] 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \tau 1_{\{X_\tau \in A_1 \cup \dots \cup A_{n-1}, H_{s_n}^x > \tau\}} + H_{s_n}^x 1_{\{H_{s_n}^x < \tau\}} \end{aligned} \quad (3.10)$$

Now, applying the same reasoning as in the case of the cost per cycle, the conditional expectation of $L(x, \tau)$ given $(X_\tau, M_\tau, H_{s_n}^x)$ is simply given by

$$E(L(x, \tau) | X_\tau, M_\tau, H_{s_n}^x) = l(X_\tau, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \tau 1_{\{H_{s_n}^x > \tau\}} + H_{s_n}^x 1_{\{H_{s_n}^x < \tau\}} \quad (3.11)$$

Taking expectations of the above expression with respect to the joint density of H , X and M we get

$$\begin{aligned} E(E(L(x) | X_\tau, M_\tau, H_{s_n}^x)) &= E\left(l(X_\tau, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \tau 1_{\{H_{s_n}^x > \tau\}} + H_{s_n}^x 1_{\{H_{s_n}^x < \tau\}}\right) \\ &= E\left(l(X_\tau, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}}\right) + \tau P(H_{s_n}^x > \tau) + \int_0^\tau h g(h | x) dh \\ &= \int_0^\tau [1 - G(h | x)] dh + \int_{A_0 S \setminus A_n} \int l(y, \tau) f_\tau(y, m | x) dm dy \end{aligned} \quad (3.12)$$

Where g and G represent the density and cumulative distribution functions of $H_{s_n}^x$, conditional on the current state x . As in the case of the equation for cost per cycle, we wish to prohibit negative levels of degradation, and to this end we define, as in the above case,

$$u(x, \tau) = P^x(X_\tau < 0, M_\tau \notin A_n) \quad (3.13)$$

and using the same reasoning as in the case of cost per cycle, we find the integral equation for the expected length of a cycle to be

$$l(x, \tau) = \int_0^\tau [1 - G_\tau(h | x)] dh + l(0, \tau) u(x, \tau) + \int_0^{s_0} l(y, \tau) K_\tau(y | x) dy \quad (3.14)$$

where K_τ is the same function as in the previous case, namely

$$K_\tau(y | x) = \int_{S \setminus A_n} f_\tau(y, m | x) dm \quad (3.15)$$

It is easily seen that this integral equation has the same general form as the equation for the expected cost per cycle function, and hence the same considerations will apply in solving this equation.

Optimal Inspection Policies

Having computed the average cost and length of a cycle, we can now directly apply the renewal reward theorem to compute the average cost. So the expected average cost per

unit time over an infinite time horizon, for a new system with inspection policy τ is given by

$$\mathbb{C}(0, \tau) = \frac{v(0, \tau)}{l(0, \tau)} \quad (3.16)$$

Since we are assuming that each new system has degradation level zero, the policy used for inspection of the system should be that which minimises $\mathbb{C}(0, \tau)$, that is

$$\tau^* = \arg \inf_{\tau > 0} \{ \mathbb{C}(0, \tau) \} \quad (3.17)$$

Finding the optimal periodic inspection policy is now a matter of optimising this function. An example of this optimisation is given in sections 3.6.1 and 3.6.4 where the case of a gamma process degradation model is considered.

3.3.2 Special Case of Monotone Degradation Process Model

In the case of a monotonic increasing degradation process we can simplify the above equation since X_t and M_t are identical. As above we define $v(x, \tau) = E[V(x, \tau) | X_0 = x]$ to be the average cost until replacement/failure a system with degradation level x . Then applying the same reasoning as on the case of the general model we find

$$\begin{aligned} V(x, \tau) | X_\tau &= [C_0 + V(X_\tau, \tau)] 1_{\{X_\tau \in A_0\}} + C_1 1_{\{X_\tau \in A_1\}} + \dots + C_{n-1} 1_{\{X_\tau \in A_{n-1}\}} + C_n 1_{\{X_\tau \in A_n\}} \\ &= V(X_\tau, \tau) 1_{\{X_\tau \in A_0\}} + \sum_{i=0}^n C_i 1_{\{X_\tau \in A_i\}} \end{aligned} \quad (3.18)$$

Where we assume $A_0 = B$. Taking expectations we find

$$\begin{aligned} E(V(x, \tau) | X_\tau) &= E\left(V(X_\tau, \tau) 1_{\{X_\tau \in A_0\}} + \sum_{i=0}^n C_i 1_{\{X_\tau \in A_i\}} \mid X_\tau\right) \\ &= v(X_\tau, \tau) 1_{\{X_\tau \in A_0\}} + \sum_{i=0}^n C_i 1_{\{X_\tau \in A_i\}} \end{aligned} \quad (3.19)$$

Thus, taking expectations with respect to X , we get

$$\begin{aligned} v(x, \tau) &= E(E(V(x, \tau) | X_\tau)) \\ &= E\left(v(X_\tau, \tau) 1_{\{X_\tau \in A_0\}} + \sum_{i=0}^n C_i 1_{\{X_\tau \in A_i\}}\right) \\ &= \int_{A_0} v(y, \tau) f_\tau(y | x) dy + \sum_{i=0}^n C_i P^x(X_\tau \in A_i) \end{aligned} \quad (3.20)$$

In this case we can define $c(x, \tau)$ as

$$c(x, \tau) = \sum_{i=0}^n C_i P^x(X_\tau \in A_i) \quad (3.21)$$

which again represents the average cost incurred at the next inspection. In this case $f_\tau(y|x)$ is the density of X_τ conditional on $X_0 = x$. So rewriting the above expression we find that $v(x)$ satisfies the integral equation

$$v(x) = c(x, \tau) + \int_0^{s_0} v(y) f_\tau(y | x) dy \quad (3.22)$$

So that, in this case, we need only consider the transition density for period τ , rather than the joint probability of the state after period τ , and the probability that the process does not reach the set A_n in that period. In addition, because the process is assumed to be monotonic increasing, there are no difficulties here with the problem of negative levels of degradation.

Applying identical reasoning to the derivation of the equation for average length of a cycle, we find that

$$L(x, \tau) | X_\tau, H_{s_n}^x = [\tau + L(X_\tau, \tau)] 1_{\{X_\tau \in A_0\}} + \tau 1_{\{X_\tau \in A_1 \cup \dots \cup A_{n-1}, H_{s_n}^x > \tau\}} + H_{s_n}^x 1_{\{H_{s_n}^x < \tau\}} \quad (3.23)$$

so that the expectation over L becomes

$$E(L(x, \tau) | X_\tau, H_{s_n}^x) = l(X_\tau, \tau) 1_{\{X_\tau \in A_0\}} + \tau 1_{\{H_{s_n}^x > \tau\}} + H_{s_n}^x 1_{\{H_{s_n}^x < \tau\}} \quad (3.24)$$

so that, upon taking expectations with respect to the density of X_t and H_s , we get

$$l(x, \tau) = \int_0^\tau [1 - G_\tau(h | x)] dh + \int_{A_0} l(y, \tau) f_\tau(y | x) dy \quad (3.25)$$

The structure of the monotonic degradation process implies that these integral equations are of the Volterra type. An example of this is shown in section 3.6.1 below where a Gamma process is used as the degradation model. The details of the numerical solution of this form of equation is given in Appendix A1.

As in the general case, the average cost per unit time is given by

$$C(0, \tau) = \frac{v(0, \tau)}{l(0, \tau)} \quad (3.26)$$

which is minimised to determine the optimal inspection interval.

3.3.3 Costs of No Inspections and Condition Monitoring

Under the average cost per unit time criteria, the cases are particularly simple. We consider first the case of no inspections, before looking at CCM.

No Inspections (NI)

Let us assume that we are dealing with the same system as described above. Then the expected length of a cycle, starting in state zero, is simply the expected failure time of the system which is simply the hitting time of S_n . Likewise the only cost incurred during a cycle is the failure replacement cost of the system C_F .

$$\mathbb{C}(0, \infty) = \frac{C_F}{E[H_{S_n}]} = \frac{C_F}{\int_0^{\infty} [1 - G_{S_n}(h|0)] dh} \quad (3.27)$$

Which represents the average cost per unit time of allowing the system to run without making any inspection, and simply replacing on failure.

The limiting value for the cost of replacement given by equation (3.27) can also be easily derived from the integral equations (3.7) and (3.14) in section 3.3.1, by letting $\tau \rightarrow \infty$, and rearranging the resulting simplification.

Continuous Condition Monitoring (CCM)

Suppose now that condition monitoring is available, and incurs monitoring cost ρ per unit time. Now assume that the CCM replacement limit is given by r , and the cost of replacement is given by C_R . Then, denoting $G_r(h|0)$ as the distribution of the hitting time of r from zero, we have simply

$$\mathbb{C}_{CCM} = \rho + \frac{C_F}{E[H_r]} = \rho + \frac{C_R}{\int_0^{\infty} [1 - G_r(h|0)] dh} \quad (3.28)$$

These simple results can now be used to compute the limiting cases, which are not covered directly by the model described above.

These results can be used to give an indication of when condition monitoring is more appropriate than a traditional inspection regime.

Relationship to the Age Replacement model

While we are considering alternative policies based on our model, we make the following comment regarding an age replacement model. By letting $s_0 = 0$ in our model, the integral equations for v and l become

$$v(x, \tau) = c(x, \tau) + v(0, \tau)u(x, \tau) \quad (3.29a)$$

$$l(x, \tau) = \int_0^\tau [1 - G_{s_n}(h | x)] dh + l(0, \tau)u(x, \tau) \quad (3.29b)$$

Letting $x = 0$, and rearranging, it is clear that the average cost per unit time becomes

$$C(0, \tau) = \frac{c(0, \tau)}{\int_0^\tau [1 - G_{s_n}(h | 0)] dh} \quad (3.30)$$

since the u function is cancelled when the ratio is formed. In this expression c may be regarded as the expected cost incurred on replacement, given by equation (3.5). As in the usual age replacement model, the denominator is simply the expected length of a cycle under policy τ .

It is clear then that this is exactly the same as the simple age replacement policy, except that the replacement cost at the end of a cycle is a function of the level of degradation. If we define a simple two region replacement policy with cost C_R incurred if degradation is below the failure limit, and C_F incurred otherwise then we have exactly the simple age replacement policy, as defined in, for example, Cox (1962). This generalisation however allows degradation dependent age replacement policies to be considered, and seems suitable for modelling perfect repair, with repair costs dependent on the level of degradation.

3.4 Optimal Periodic Inspection: Discounted Cost Criterion

3.4.1 Derivation of Discounted Total Cost

Let us now consider the optimal perfect inspection of Lévy degradation processes under the discounted total cost criteria. This criteria is perhaps simpler than the previous case, and also allows future costs to be discounted, which is a significant factor in real applications. In addition, this criterion is not dependent on the renewal reward theorem for its validity and can be extended easily to the case of non-periodic inspections.

The system we shall consider is identical to that considered above, and we again use model assumptions M1 as given in section 3.2. We apply the same reasoning as in the previous case.

We define $V_\delta(x, \tau)$ to be the discounted total cost for a system which has current level of degradation x . So, if we define t_i to be the inter-event times (an event being either an inspection or failure), then define

$$v_\delta(x, \tau) = E(V_\delta(x, \tau) | X_0 = x) = E\left(\sum_{n=1}^{\infty} e^{-\delta(t_1 + \dots + t_n)} C(X_{t_{n-1}}, X_{t_n}; \tau) | X_0 = x\right) \quad (3.31)$$

to be the expected discounted total cost, where $C(x, y; \tau)$ represents the (random) cost incurred if the system is in state y at time t_n when it was in state x at time t_{n-1} , and the inspection policy is τ .

Using standard dynamic programming arguments, we can express this cost function as the sum of the cost incurred at the next inspection and the cost incurred thereafter. This case differs only slightly from the calculation of expected cost per cycle in the average cost criterion case above. The differences are simply that we must now discount all of the costs and include the cost of operating the new system, after a replacement or failure occurs.

Again we assume that X_t represents the degradation process, and M_t represents the maximum value of that process over the period $[0, t]$. Let δ represent the discount rate which applies to our calculations, and assume $\delta > 0$. The cost function may be written

$$\begin{aligned} I'_\delta(x, \tau) | X_\tau, M_\tau, H_{s_n}^x &= e^{-\delta\tau} (C_0 + V_\delta(X_\tau, \tau)) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} \\ &+ \sum_{i=1}^{n-1} e^{-\delta\tau} (C_i + I'_\delta(0, \tau)) 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + e^{-\delta H_s^x} (C_n + I'_\delta(0, \tau)) 1_{\{H_s^x < \tau\}} \end{aligned} \quad (3.32)$$

where we have expressed the cost function as the sum of the cost in the event of continuing (i.e. not replacing the system) and the costs for each possible replacement event. Equation (3.32) is equivalent to (3.1) in the average cost criterion case. We proceed similarly

$$\begin{aligned}
& E(V_\delta(x, \tau) | X_\tau, M_\tau, H_{s_n}^x) \\
&= E \left(e^{-\delta\tau} (C_0 + V_\delta(X_\tau, \tau)) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \right. \\
&\quad \left. \sum_{i=1}^{n-1} e^{-\delta\tau} (C_i + V_\delta(0, \tau)) 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + e^{-\delta H_s^x} (C_n + V_\delta(0, \tau)) 1_{\{H_s^x < \tau\}} \right) \quad (3.33) \\
&= e^{-\delta\tau} v(X_\tau, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \sum_{i=0}^{n-1} e^{-\delta\tau} C_i 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + e^{-\delta H_s^x} C_n 1_{\{H_s^x < \tau\}} \\
&\quad + \sum_{i=1}^{n-1} e^{-\delta\tau} v_\delta(0, \tau) 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + e^{-\delta H_s^x} v_\delta(0, \tau) 1_{\{H_s^x < \tau\}}
\end{aligned}$$

We define $f_\tau(y, m | x)$ to be the joint density of X_τ and M_τ , conditional upon $X_0 = x$. Taking expectations of the above expression with respect to the joint density of M_τ and X_τ we get

$$\begin{aligned}
v_\delta(x, \tau) &= E(E(V_\delta(x, \tau) | X_\tau, M_\tau)) \\
&= E \left(e^{-\delta\tau} v(X_\tau, \tau) 1_{\{X_\tau \in A_0, M_\tau \notin A_n\}} + \sum_{i=0}^{n-1} e^{-\delta\tau} C_i 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} \right. \\
&\quad \left. + e^{-\delta H_s^x} C_n 1_{\{H_s^x < \tau\}} + \sum_{i=1}^{n-1} e^{-\delta\tau} v_\delta(0, \tau) 1_{\{X_\tau \in A_i, M_\tau \notin A_n\}} + e^{-\delta H_s^x} v_\delta(0, \tau) 1_{\{H_s^x < \tau\}} \right) \quad (3.34) \\
&= \int \int_{A_0 \setminus S \setminus A_n} e^{-\delta\tau} v(y, \tau) f_\tau(y, m | x) dm dy + e^{-\delta\tau} C_0 P^x(X_\tau \in A_0, M_\tau \notin A_n) \\
&\quad + \sum_{i=1}^{n-1} e^{-\delta\tau} (C_i + v_\delta(0, \tau)) P^x(X_\tau \in A_i, M_\tau \notin A_n) + (C_n + v_\delta(0, \tau)) \int_0^\tau e^{-\delta h} g_\tau(h | x) dh
\end{aligned}$$

where $g(h | x)$ is the probability density function of $H_{s_n}^x$, the hitting time of the critical set from x .

As in the case of the discounted total cost criterion, we make the assumption that $v_\delta(x, \tau) = v_\delta(0, \tau)$ for $x < 0$. The integral term in the above expression becomes

$$\begin{aligned}
& \int \int_{A_0 \setminus S \setminus A_n} e^{-\delta\tau} v_\delta(y, \tau) f_\tau(y, m | x) dm dy = \\
& e^{-\delta\tau} v_\delta(0, \tau) P^x(X_\tau < 0, M_\tau \notin A_n) + \int_0^{s_0} \int_{S \setminus A_n} e^{-\delta\tau} v_\delta(y, \tau) f_\tau(y, m | x) dm dy \quad (3.35)
\end{aligned}$$

the details following from the average cost case. Substituting equation (3.35) into equation (3.34) we obtain

$$v_\delta(x, \tau) = c_\delta(x, \tau) + v_\delta(0, \tau)u_\delta(x, \tau) + \int_0^{s_0} v_\delta(y, \tau) K_\tau(y | x) dy \quad (3.36)$$

where

$$c_\delta(x, \tau) = \sum_{i=0}^{n-1} e^{-\delta\tau} C_i P^x(X_\tau \in A_i, M_\tau \notin A_n) + C_n \int_0^\tau e^{-\delta h} g_\tau(h | x) dh \quad (3.37a)$$

$$u_\delta(x, \tau) = e^{-\delta\tau} P^x(X_\tau \notin A_0^+ \cup A_n, M_\tau \notin A_n) + \int_0^\tau e^{-\delta h} g_\tau(h | x) dh \quad (3.37b)$$

$$K_\tau(y | x) = \int_{S \setminus A_n} e^{-\delta\tau} f_\tau(y, m | x) dm \quad (3.37c)$$

A_0^+ denotes the set of positive states which are in A_0 , i.e. $A_0^+ = A_0 \setminus (-\infty, 0)$.

The optimal policy in this case is given by

$$\tau_\delta^* = \arg \inf_{\tau > 0} \{ v(0, \tau) \}$$

which in all but degenerate cases must be obtained by numerical methods.

We note that the forms of the functions c , u and K are more complicated than is the case for the average cost criteria. It can be seen however, that the integral equation above has the same general form as in the previous case, and can be solved using the same methods.

As in the average cost case, this equation is simplified if we are dealing with a monotonic degradation process. It is this simplification that we consider in the next section.

3.4.2 Special Case of Monotone Degradation Process Model

As before, in the case of monotonic degradation process, X_t and M_t are identical. We therefore find that the total discounted cost may be expressed as

$$V_\delta(x, \tau) | X_\tau, H_{s_n}^x = e^{-\delta\tau} (C_0 + V_\delta(X_\tau, \tau)) 1_{\{X_\tau \in A_0\}} + \sum_{i=1}^{n-1} e^{-\delta\tau} (C_i + V_\delta(0, \tau)) 1_{\{X_\tau \in A_i\}} + e^{-\delta H_\tau^x} (C_n + V_\delta(0, \tau)) 1_{\{H_\tau^x < \tau\}} \quad (3.38)$$

So that the expected cost conditional on X and H is given by

$$E(V_\delta(x, \tau) | X_\tau, H_{s_n}^x) = e^{-\delta\tau} v_\delta(X_\tau, \tau) 1_{\{X_\tau \in A_0\}} + \sum_{i=0}^{n-1} e^{-\delta\tau} C_i 1_{\{X_\tau \in A_i\}} + e^{-\delta H_\tau^x} C_n 1_{\{H_\tau^x < \tau\}} + \sum_{i=1}^{n-1} e^{-\delta\tau} v_\delta(0, \tau) 1_{\{X_\tau \in A_i\}} + e^{-\delta H_\tau^x} v_\delta(0, \tau) 1_{\{H_\tau^x < \tau\}} \quad (3.39)$$

Thus, taking expectations with respect to X and H , we get

$$\begin{aligned} v_\delta(x, \tau) &= E(E(V_\delta(x, \tau) | X_\tau, H_{s_n}^x)) \\ &= \int_{A_0} e^{-\delta\tau} v_\delta(y, \tau) f_\tau(y | x) dy + e^{-\delta\tau} C_0 P^x(X_\tau \in A_0) \\ &\quad + \sum_{i=1}^{n-1} e^{-\delta\tau} (C_i + v_\delta(0, \tau)) P^x(X_\tau \in A_i) + (C_n + v_\delta(0, \tau)) \int_0^\tau e^{-\delta h} g_\tau(h | x) dh \end{aligned} \quad (3.40)$$

In this case we define $c(x, \tau)$ as

$$c_\delta(x, \tau) = \sum_{i=0}^{n-1} e^{-\delta\tau} C_i P^x(X_\tau \in A_i) + C_n \int_0^\tau e^{-\delta h} g_\tau(h | x) dh \quad (3.41)$$

Also, because we have a monotonic (increasing) degradation process there is no possibility that the level of degradation will become negative, so $u_\delta(x, \tau)$ is given by

$$u_\delta(x, \tau) = e^{-\delta\tau} P^x(X_\tau \notin A_0 \cup A_n) + \int_0^\tau e^{-\delta h} g_\tau(h | x) dh \quad (3.42)$$

Where A_0^+ is defined as in the previous section. Then, making the appropriate substitutions we get the integral equation

$$v_\delta(x, \tau) = c_\delta(x, \tau) + v_\delta(0, \tau) u_\delta(x, \tau) + \int_0^{s_0} v_\delta(y, \tau) K_\tau(y | x) dy \quad (3.43)$$

with c and u defined as above and

$$K_\tau(y | x) = e^{-\delta\tau} f_\tau(y | x) \quad (3.44)$$

is the discounted transition function. Once again, the integral equation is in the same standard form, and may be solved by the methods given in appendix A1.

3.4.3 Costs of No Inspections and Continuous Condition monitoring

The case of no inspections and condition Monitoring are again much easier than in the general case. In the discounted total cost criteria however, a recursive definition must be used

No Inspections

As in section 3.3.3, we assume that the system is allowed to run until failure, so that the only costs incurred are the costs of replacing the system at each failure. We define C_F to be the cost of replacing the system on failure, and assume failure occurs when the degradation level hits the critical set S_n . Letting $V_\delta(0, \infty)$ denote the total discounted costs with discount rate δ , starting in state zero. Then

$$V_\delta(x, \infty) | H_{S_n}^x = [C_F + V_\delta(0, \infty)] e^{-\delta H_{S_n}^x} \quad (3.45)$$

so that

$$E(V_\delta(x, \infty) | H_{S_n}^x) = E([C_F + V_\delta(0, \infty)] e^{-\delta H_{S_n}^x}) = [C_F + v_\delta(0, \infty)] e^{-\delta H_{S_n}^x} \quad (3.46)$$

hence, in the particular case when $x = 0$, which is the case we are most concerned with, we have

$$\begin{aligned} v_\delta(0, \infty) &= [C_F + v_\delta(0, \infty)] \int_0^\infty e^{-\delta h} g_{S_n}(h|0) dh \\ \Rightarrow v_\delta(0, \infty) &= \frac{C_F \int_0^\infty e^{-\delta h} g_{S_n}(h|0) dh}{1 - \int_0^\infty e^{-\delta h} g_{S_n}(h|0) dh} \end{aligned} \quad (3.47)$$

Where $g_{S_n}(h|x)$ represents the density of the hitting time of the critical set S_n from x . If we define $\tilde{g}_{S_n}(\theta|x)$ to be the Laplace transform of g , it is clear that the discounted total cost in this case may be written in terms of the Laplace transform of g as

$$v_\delta(0, \infty) = \frac{C_F \tilde{g}_{S_n}(\delta|0)}{1 - \tilde{g}_{S_n}(\delta|0)} \quad (3.48)$$

It can be easily shown that the solution of the integral equation above gives values consistent with those found directly using the above formula.

In this case the integral equation is given by equation (3.36) and (3.37a,b,c). It is clear then, that as $\tau \rightarrow \infty$, the summations of probabilities in c and u tend to zero because of the presence of the exponential function. Thus we find from equations (3.35a) and (3.36b)

$$c_\delta(x, \tau) \rightarrow C_n \int_0^\infty e^{-\delta h} g_{s_n}(h | x) dh \quad u_\delta(x, \tau) \rightarrow \int_0^\infty e^{-\delta h} g_{s_n}(h | x) dh \quad (3.49)$$

$K_\tau^\delta(y | x)$ represents the probability that the system is in state y , for some $y < r$, at time τ (without having first failed). Clearly, for degradation modelled by a Lévy process with positive drift, K must tend to zero as $\tau \rightarrow \infty$. Thus the integral equation may be re-written as

$$v_\delta(x, \infty) = C_n \int_0^\infty e^{-\delta h} g_{s_n}(h | x) dh + v_\delta(0, \tau) \int_0^\infty e^{-\delta h} g_{s_n}(h | x) dh \quad (3.50)$$

So, letting $x = 0$, and rearranging gives

$$v_\delta(0, \infty) = \frac{C_F \int_0^\infty e^{-\delta h} g_{s_n}(h | 0) dh}{1 - \int_0^\infty e^{-\delta h} g_{s_n}(h | 0) dh} \quad (3.51)$$

As is found by direct argument above.

Continuous Condition Monitoring

In the case of CCM we use the same approach, the only difference being the cost of monitoring. We assume that the cost of CCM is incurred continuously at rate ρ per unit time, and that the system is to be replaced with cost C_R , when the degradation reaches level r . Then starting in state x we have

$$\begin{aligned} V_\delta(x) | H_x^r &= \int_0^{H_x^r} \rho e^{-\delta t} dt + [C_R + V_\delta(0)] e^{-\delta H_x^r} \\ &= \frac{\rho}{\delta} + \left[C_R + V_\delta(0) - \frac{\rho}{\delta} \right] e^{-\delta H_x^r} \end{aligned} \quad (3.52)$$

so that, taking expectations, we get,

$$\begin{aligned}
 v_\delta(x) &= E[E[V_\delta(x) | H'_x]] \\
 &= E_H \left[\frac{\rho}{\delta} + \left[C_R + v_\delta(0) - \frac{\rho}{\delta} \right] e^{-\delta H'_1} \right] \\
 &= \int_0^\infty \left\{ \frac{\rho}{\delta} + \left[C_R + v_\delta(0) - \frac{\rho}{\delta} \right] e^{-\delta h} \right\} g_r(h|x) dh \\
 &= \frac{\rho}{\delta} + \left[C_R + v_\delta(0) - \frac{\rho}{\delta} \right] \int_0^\infty e^{-\delta h} g_r(h|x) dh \\
 &= \frac{\rho}{\delta} + \left[C_R + v_\delta(0) - \frac{\rho}{\delta} \right] \tilde{g}_r(\delta|x)
 \end{aligned} \tag{3.53}$$

so that, in the particular case when $x = 0$, we have

$$\begin{aligned}
 v_\delta(0) &= \frac{\rho}{\delta} + \left[C_R + v_\delta(0) - \frac{\rho}{\delta} \right] \tilde{g}_r(\delta|0) \\
 \Rightarrow v_\delta(0) &= \frac{\frac{\rho}{\delta} + \left[C_R - \frac{\rho}{\delta} \right] \tilde{g}_r(\delta|0)}{1 - \tilde{g}_r(\delta|0)}
 \end{aligned} \tag{3.54}$$

Depending on the nature of the degradation process, the Laplace transforms of the hitting time densities may be computed directly, or by numerical integration. In both cases however, the computations are simple enough to give a good idea of the costs involved in condition monitoring.

Age Replacement Model

As for the average cost case above, by letting the replacement limit tend to zero, we can replicate the simple age replacement model, with a degradation dependent cost of replacement.

In this case, equation (3.36) becomes simply

$$v_\delta(x, \tau) = c_\delta(x, \tau) + v_\delta(0, \tau)u_\delta(x, \tau) \tag{3.55}$$

With c and u defined by (3.37a,b). So that upon letting $x = 0$, we get

$$v(0, \tau) = \frac{c(0, \tau)}{1 - u(0, \tau)} \tag{3.56}$$

On calculating these functions with $s_0 = 0$, we find that the discounted total cost is simply

$$v_\delta(0, \tau) = \frac{e^{-\delta\tau} (C_1 P(X_\tau \in A_1, M_\tau \notin A_n) + \dots + C_{n-1} P(X_\tau \in A_{n-1}, M_\tau \notin A_n)) + C_n \int_0^\tau e^{-\delta h} g(h|x) dh}{1 - e^{-\delta\tau} P(M_\tau \in A_n) - \int_0^\tau e^{-\delta h} g(h|x) dh} \quad (3.57)$$

Which, in the case of a two-region replacement policy gives an expression for the discounted total cost of an age replacement policy.

3.5 Optimal Non-Periodic Inspection: Discounted Cost Criteria

3.5.1 Optimality Equations in General

Let us now consider the case of optimal non-periodic inspection. This case is more general and often more useful than the previously considered case of periodic inspection, since it generally results in policies with lower costs. This is because each inspection interval is in some way determined by the state of the system at the previous inspection. This case however, is more complex and requires the use of dynamic programming to obtain the optimal policies.

As in the case of periodic inspection under the discounted cost criterion we assume a total discounted cost function of the form

$$v_\pi(x) = E(V_\pi(x) | X_0 = x) = E\left(\sum_{n=1}^{\infty} e^{-\delta(t_1 + \dots + t_n)} C(X_{t_{n-1}}, X_{t_n}; \pi(X_{t_{n-1}})) | X_0 = x\right) \quad (3.58)$$

where π represents the inspection policy, and other notation is as in section 3.4.1. We are only concerned with deterministic stationary policies, since the Lévy property of our degradation process implies the current time has no bearing on future levels of degradation.

We define the function v_δ to be the value function for the δ -optimal policy, i.e.

$$v_\delta(x) = \inf_{\pi} v_\pi(x) \quad (3.59)$$

For the general cost function given by equation (3.57) above, we find that the optimality equation is given by

$$v_\delta(x) = \inf_{\tau > 0} \left\{ c(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau+h)} v_\delta(y) f_\tau(y, h | x) dh dy \right\} \quad (3.60)$$

where c represents the discounted costs incurred at the next inspection, when in inspection interval τ is chosen. A derivation of this based on Ross (1970) is given in appendix A3. In the particular case we are dealing with, the results of section 3.4.1 show that the optimality equation is given by

$$v_\delta(x) = \inf_{\tau > 0} \left\{ c_\delta(x, \tau) + v_\delta(0)u_\delta(x, \tau) + \int_0^{s_0} v_\delta(y) K_\tau^\delta(y | x) dy \right\} \quad (3.61)$$

where c , u and K are given by equations (3.37a,b,c) of section 3.4

Now let us consider the convergence of this dynamic programming equation. We follow the standard method, given by, among others, Ross (1970).

We now assume that c is continuous and bounded, as it will be in reasonable cases, and define $T_\delta : C^b[0, \infty) \rightarrow C^b[0, \infty)$ as

$$(T_\delta v)(x) = \inf_{\tau > 0} \left\{ c_\delta(x, \tau) + v(0)u_\delta(x, \tau) + \int_0^{s_0} v(y) K_\tau^\delta(y | x) dy \right\} \quad (3.62)$$

Now, assume that $v_1, v_2 \in C^b[0, \infty)$ and $x \geq 0$. Then let $\tau_0 = \pi^*(x, v_2)$ be a minimising point of the function

$$\tau \rightarrow c_\delta(x, \tau) + v_2(0)u_\delta(x, \tau) + \int_0^{s_0} v_2(y) K_\tau(y | x) dy \quad (3.63)$$

over the range $[0, x]$. Then

$$\begin{aligned} & (T_\delta v_1)(x) - (T_\delta v_2)(x) \\ &= \inf_{\tau > 0} \left\{ c_\delta(x, \tau) + v_1(0)u_\delta(x, \tau) + \int_0^{s_0} v_1(y) K_\tau(y | x) dy \right\} \\ & \quad - c_\delta(x, \tau_0) + v_2(0)u_\delta(x, \tau_0) + \int_0^{s_0} v_2(y) K_{\tau_0}(y | x) dy \end{aligned} \quad (3.64)$$

$$\begin{aligned}
&\leq u_\delta(x, \tau_0) \cdot (v_1(0) - v_2(0)) + \int_0^{s_0} (v_1(y) - v_2(y)) K_\tau(y|x) dy \\
&\leq \|v_1 - v_2\| \left\{ u_\delta(x, \tau_0) + \int_0^{s_0} K_\tau(y|x) dy \right\} \\
&\leq \kappa \|v_1 - v_2\|
\end{aligned}$$

We note that κ is less than unity for $\delta > 0$, since it represents the discounted total transition probability out of state x . The function u in the penultimate line of the derivation is the discounted probability that the system will be replaced (or have negative degradation) and the integral term is the discounted probability that the system will not fail or be preventively replaced. It is clear from equation (3.64), and a similar derivation with v_1 and v_2 interchanged that

$$\|T_\delta v_1 - T_\delta v_2\| \leq \kappa \|v_1 - v_2\| \quad (3.65)$$

so that T defines a contraction mapping on $C^b[0, \infty)$. We deduce therefore, under the assumptions given that $T^k v$ converges to v_δ uniformly on $[0, \infty)$ as $k \rightarrow \infty$.

We can apply a standard policy improvement algorithm to the discretisation of this problem, where the discretisation is carried out by applying a quadrature rule to the integral. The policy improvement algorithm based on Puterman (1994) is

1. Set $k = 0$, and select an arbitrary inspection rule $\pi_0 = \{\tau_0^0, \dots, \tau_n^0\}$
2. (Policy Evaluation) Obtain $\mathbf{v}^k$ by solving $\mathbf{v}^k = \mathbf{M}_k^{-1} \mathbf{c}_k$, where the subscript k indicates the matrix is to be evaluated with policy π_k .

3. (Policy Improvement) Choose $\pi_{k+1} = \{\tau_0^{k+1}, \dots, \tau_n^{k+1}\}$ to satisfy

$$\pi_{k+1} = \arg \min_{\pi \in \Pi} \{T_\pi(\mathbf{v}^k)\}$$

4. If $\pi_{k+1} = \pi_k$, stop and set $\pi^* = \pi_k$. Otherwise, increment k by 1 and return to step 2.

The exact definition of the matrix $\mathbf{M}$ depends on the situation, but the methods of appendix A1 may be used to determine its form.

See the example of section 3.6, where this algorithm is used to compute optimal non-periodic inspection policies.

3.5.2 Special Case of Monotone Degradation Process Model

Applying the same reasoning as above to the monotone degradation case in section 3.4.2, gives us the following optimality equation for the case of a monotonic degradation process

$$v_\delta(x) = \inf_{\tau > 0} \left\{ c_\delta(x, \tau) + v_\delta(0)u_\delta(x, \tau) + \int_0^{s_0} v_\delta(y) K_\tau(y | x) dy \right\} \quad (3.66)$$

with c , v and K defined as in equations (3.41), (3.42) and (3.44), namely:

$$c_\delta(x, \tau) = \sum_{i=0}^{n-1} e^{-\delta\tau} C_i P^x(X_\tau \in A_i) + C_n \int_0^\tau e^{-\delta h} g_\tau(h | x) dh$$

$$u_\delta(x, \tau) = e^{-\delta\tau} P^x(X_\tau \notin A_0 \cup A_n) + \int_0^\tau e^{-\delta h} g_\tau(h | x) dh$$

$$K_\tau(y | x) = e^{-\delta\tau} f_\tau(y | x)$$

The argument showing that the mapping

$$(T_\delta v)(x) = \inf_{\tau > 0} \left\{ c_\delta(x, \tau) + v(0)u_\delta(x, \tau) + \int_0^{s_0} v(y) K_\tau(y | x) dy \right\} \quad (3.67)$$

is a contraction mapping, is identical to that above for the non-monotonic degradation process.

In both cases, of general and monotonic degradation processes these equations must be solved numerically, since their complexity makes closed form analytical solutions unavailable. However, using the methods given above we may determine the cost of any particular policy. Using a standard policy improvement routine, we can derive the optimal policy. An example is shown in section 3.6, immediately below.

In this case, it is clear that the costs of no-inspections and of continuous condition monitoring are identical to those given in section 3.4.3, namely

$$v_{\delta}(0, \infty) = \frac{C_F \tilde{g}_c(\delta | 0)}{1 - \tilde{g}_c(\delta | 0)}$$

$$v_{\delta}^{CCM}(0) = \frac{\frac{\rho}{\delta} + \left[C_R - \frac{\rho}{\delta} \right] \tilde{g}_r(\delta | 0)}{1 - \tilde{g}_r(\delta | 0)}$$

where $\tilde{g}$ represents the Laplace transform of the density of a hitting time, the suffix denoting the point which we are considering.

3.6 Example : Gamma Process Degradation

As an example we assume that we are dealing with a Gamma degradation process. Therefore the increments of the degradation process have a gamma distribution

$$X_t - X_s \sim \text{Gamma}(\alpha(t - s), \beta)$$

Under our notation this means that the transition density function is given by

$$f_{\tau}(y | x) = \frac{\beta^{\alpha\tau} (y - x)^{\alpha\tau-1} e^{-\beta(y-x)}}{\Gamma(\alpha\tau)} \quad (3.68)$$

Let us in addition simplify the problem by assuming that the replacement policy is determined by only three sets. We assume that the system fails if the degradation reaches a certain level c , and we replace the system with a new system, if an inspection reveals a level of degradation greater than amount $r < c$. Figure 3.1 shows how the state space is subdivided under these assumptions.

We further assume the following cost structure

- (i) Each inspection incurs a cost C_I .
- (ii) Each preventive replacement incurs cost C_R , in addition to the cost of inspection.
- (iii) Replacement on failure incurs cost C_F , with no inspection cost incurred.

This simplified model, then, corresponds to the more general model previously considered with

$$A_0 = [0, r), \quad A_1 = [r, c), \quad A_2 = [c, \infty), \quad C_0 = C_I, \quad C_1 = C_I + C_R, \quad C_2 = C_F$$

Now let us consider the three cases of Periodic inspection (average and discounted costs) and non-periodic discounted costs separately.

3.6.1 Periodic Inspection : Average Cost per unit time

From the results of section 3.2, for the case of a monotonic degradation process we find the cost and length per cycle are the solution of the equations (3.7) and (3.14), which are

$$v(x, \tau) = c(x, \tau) + \int_{A_0} v(y, \tau) f_{\tau}(y | x) dy$$

$$l(x, \tau) = \int_0^{\tau} [1 - G_{\tau}(h | x)] dh + \int_{A_0} l(y, \tau) f_{\tau}(y | x) dy$$

where c is given by equation (3.5), and u is zero since we are dealing with a monotone process.

$$c(x, \tau) = \sum_{i=0}^n C_i P^x(X_{\tau} \in A_i)$$

and as before, $f_{\tau}(y | x)$ is the probability density function of X_{τ} given $X_0 = x$, and $G_c(h | x)$ is the cumulative distribution function of H_c^x , the hitting time of c starting from $X_0 = x < c$. These are as follows:

$$f_{\tau}(y | x) = \frac{\beta^{\alpha\tau} (y - x)^{\alpha\tau-1} e^{-\beta(y-x)}}{\Gamma(\alpha\tau)}$$

and it is simple to compute the CDF of H_c^x in terms of the incomplete gamma function

$$\begin{aligned} G_c(h | x) &= P^x(H_c^x \leq h) = P^x(X_h > c) \\ &= \int_{c-x}^{\infty} \frac{\beta^{\alpha h} y^{\alpha h-1} e^{-\beta y}}{\Gamma(\alpha h)} dy \\ &= \frac{\Gamma(\alpha h; \beta(c-x))}{\Gamma(\alpha h)} \\ &= Q(\alpha h, \beta(c-x)) \end{aligned} \tag{3.69}$$

Where the incomplete Gamma function is defined as $\Gamma(\alpha; x) := \int_x^{\infty} t^{\alpha-1} e^{-t} dt$, and Q is an incomplete Gamma Ratio defined by

$$Q(\alpha; x) := \frac{\Gamma(\alpha; x)}{\Gamma(\alpha)}$$

The function $c(x, \tau)$ is found by computing the probabilities shown in the general case above. In this particular case we have that

$$\begin{aligned} c(x, \tau) &= \sum_{i=0}^n C_i P^x(X_\tau \in A_i) = C_I P(X_\tau < c) + C_R P(r \leq X_\tau < c) + C_F P(X_\tau \geq c) \\ &= C_I \left(1 - \frac{\Gamma(\alpha \tau; \beta(c-x))}{\Gamma(\alpha \tau)} \right) \\ &\quad + C_R \left(\frac{\Gamma(\alpha \tau; \beta(r-x)) - \Gamma(\alpha \tau; \beta(c-x))}{\Gamma(\alpha \tau)} \right) + C_F \left(\frac{\Gamma(\alpha \tau; \beta(c-x))}{\Gamma(\alpha \tau)} \right) \\ &= C_I + C_R Q(\alpha \tau, \beta(r-x)) + (C_F - C_R - C_I) Q(\alpha \tau, \beta(c-x)) \end{aligned}$$

And in addition we find that

$$\int_0^\tau [1 - G_\tau(h | x)] dh = \tau - \int_0^\tau \frac{\Gamma(\alpha h; \beta(c-x))}{\Gamma(\alpha h)} dh = \tau - \int_0^\tau Q(\alpha h, \beta(c-x)) dh$$

Substituting these functions into the integral equations for average cost and average time per cycle we get:

$$\begin{aligned} v(x, \tau) &= C_I + C_R Q(\alpha \tau, \beta(r-x)) + (C_F - C_R - C_I) Q(\alpha \tau, \beta(c-x)) \\ &\quad + \int_x^r v(y, \tau) \frac{\beta^{\alpha \tau} (y-x)^{\alpha \tau - 1} e^{-\beta(y-x)}}{\Gamma(\alpha \tau)} dy \\ l(x, \tau) &= \tau - \int_0^\tau Q(\alpha h, \beta(c-x)) dh + \int_x^r l(y, \tau) \frac{\beta^{\alpha \tau} (y-x)^{\alpha \tau - 1} e^{-\beta(y-x)}}{\Gamma(\alpha \tau)} dy \end{aligned}$$

Given the form of these integral equations we use a numerical approximation to solve them. Details of the approximation used, and other methods used later in this chapter are given in Appendix A1. Using this approximation we compute values of $v(x)$ and $l(x)$ for various values of x , thus obtaining the cost function C . Section 3.6.4 below summarises the numerical results found for this model, and compares them to alternative policies.

3.6.2 Periodic Inspection: Discounted Total Cost

In the case of periodic inspection under the discounted total cost criterion we must solve the integral equation

$$v_\delta(x, \tau) = c_\delta(x, \tau) + v_\delta(0, \tau)u_\delta(x, \tau) + \int_0^{s_0} v_\delta(y, \tau) K_\tau(y | x) dy$$

where the functions c , u and K in this case are given by equations (3.41), (3.42) and (3.44) of section 3.3. Expanding these functions we find them to be

$$K_\tau(y | x) = e^{-\delta\tau} \frac{\beta^{\alpha\tau} (y-x)^{\alpha\tau-1} e^{-\beta(y-x)}}{\Gamma(\alpha\tau)} \quad y \geq x$$

$$c_\delta(x, \tau) = e^{-\delta\tau} \{C_I + C_R Q(\alpha\tau, \beta(r-x)) + (C_F - C_R - C_I) Q(\alpha\tau, \beta(c-x))\} \\ + \delta C_F \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh$$

$$u_\delta(x, \tau) = e^{-\delta\tau} Q(\alpha\tau, \beta(r-x)) + \delta \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh$$

where we have used the elementary relation

$$\begin{aligned} \int_0^\tau e^{-\delta h} g(h | x) dh &= \int_0^\tau e^{-\delta h} \frac{\partial}{\partial h} G(h | x) dh \\ &= \left[e^{-\delta h} G(h | x) \right]_{h=0}^\tau + \delta \int_0^\tau e^{-\delta h} G(h | x) dh \\ &= e^{-\delta\tau} G(\tau | x) + \delta \int_0^\tau e^{-\delta h} G(h | x) dh \end{aligned} \quad (3.70)$$

This transformation is necessary, since the density of the hitting time does not have a useful closed form expression. Substituting these values into equation (3.45) above we have the integral equation

$$v_\delta(x, \tau) = e^{-\delta\tau} \{C_I + C_R Q(\alpha\tau, \beta(r-x)) + (C_F - C_R - C_I) Q(\alpha\tau, \beta(c-x))\} \\ + \delta C_F \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh + v_\delta(0, \tau) \left[e^{-\delta\tau} Q(\alpha\tau, \beta(r-x)) + \delta \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh \right] \\ + \int_x^{s_0} v_\delta(y, \tau) e^{-\delta\tau} \frac{\beta^{\alpha\tau} (y-x)^{\alpha\tau-1} e^{-\beta(y-x)}}{\Gamma(\alpha\tau)} dy$$

Which we solve by the Modified Nystrom Method (Press et. al., 1992). The details of the numerical solution are given in Appendix A1.

3.6.3 Non-Periodic Inspection : Discounted Total Cost

Let us now consider the case of non-periodic inspection. From the analysis of section 3.6.2 it is clear that the dynamic programming equation we must solve is:

$$v_\delta(x) = \inf_{\tau > 0} \left\{ \begin{aligned} & e^{-\delta\tau} \left[C_I + C_R Q(\alpha\tau, \beta(r-x)) + (C_F - C_R - C_I) Q(\alpha\tau, \beta(c-x)) \right] \\ & + \delta C_F \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh \\ & + v_\delta(0) \left[e^{-\delta\tau} Q(\alpha\tau, \beta(r-x)) + \delta \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh \right] \\ & + \int_x^{s_0} v_\delta(y) e^{-\delta\tau} \frac{\beta^{\alpha\tau} (y-x)^{\alpha\tau-1} e^{-\beta(y-x)}}{\Gamma(\alpha\tau)} dy \end{aligned} \right\}$$

To solve this, we discretise the integral equation and apply a standard policy improvement algorithm, described above. Because we use the discretisation, we can find optimal policies for arbitrarily fine meshes over the state-space and time domain. While these do not give the exact optimal policy, we may find a policy which is arbitrarily close to that policy from the discretised problem.

3.6.4 Numerical Results and Comments on the model

Let us now consider the numerical results found by solving the equations of section 3.6.1 to 3.6.3. In all three cases, our underlying assumptions are as follows:

- (i) The underlying degradation process is a Gamma process with $\alpha/\beta = 1$. We vary the parameter α to look at the effects of changing the variability of the underlying process, for a given average rate of degradation
- (ii) The failure limit of the process $c = 1$. We vary the replacement limit r to see the effect of changing this limit.
- (iii) The cost of an inspection C_I is 1 unit. We vary the replacement cost C_R and failure cost C_F , to look at changes in the optimal policies.

For the numerical work we assume that the state space is subdivided into 10 steps, and the time domain is divided into steps of length 0.05. We choose the state and time steps simply because they are small enough to show the nature of the optimal policies, and still not be too computationally demanding. Firstly, we consider the optimal periodic policies found in the average cost and discounted total cost case. We consider them together since their solutions are very similar, in that, for all cases considered, they result in the same optimal policy. This seems to be because of the high probability that each cycle will have the same length and costs under periodic inspection.

The optimal periodic policy in these cases is determined by the inspection period. The optimal policies for the average cost and discounted total cost are shown in tables 3.1 and 3.2 on pages 87 and 88 respectively. These tables, however, conceal many interesting facts about the cost functions.

Properties of the cost function for Periodic Policies

We now consider the form of the cost function under periodic inspection. For the model we are using we have found three distinct shapes for the cost function.

- (i) Single Local Minimum (Fig 3.2a)
- (ii) Double Local Minimum (Fig 3.2b)
- (iii) Monotone Decreasing (Fig 3.2c)




Figure3.2 (a)




Figure 3.2 (b)




Figure 3.2 (c)

The first of these, Single Local Minimum, is the standard cost function found in most of the literature. As the name suggests it corresponds to the case in which the costs are a decreasing function of the inspection interval up to the optimal point, and thereafter the cost is an increasing function of the inspection interval.

The second case of a double local minimum is more interesting. As can be seen from figure 3.2b, the cost function has two local minima. For the model we are using this unusual behaviour is intuitively correct, and results from the nature of the periodic inspection policy.

The third case shown in figure 3.2c is also a standard one, and occurs when the cost of failure is not sufficiently high to outweigh the costs of inspection. The resulting curve is thus monotone decreasing, and the optimal policy is not to carry out inspections, and

replace the system on failure. We do not consider this case in detail, since we are assuming the system is important enough that failure has large consequences.

Let us now consider how variations in the parameter values affect the shape of the cost function. We do not give details of the precise numbers used, since we are only looking at how the parameters effect the shape of the cost function. Later we consider how these variations effect the optimal policy itself.

Firstly let us consider the effect of the choice of replacement level on the shape of the cost function. We have found that low replacement limits result in single minimum cost functions, which become double minimum as the replacement limit gets closer to the failure limit. When the replacement limit is low, it has very little effect on the optimal inspection policy. The lower the replacement policy, the more likely that a preventive replacement will occur at the next inspection. So, for low replacement limits, we are simply assuming that every inspection is in fact a replacement and are effectively applying an age replacement policy. The optimal policy then is close to that found for an age replacement policy. It is clear then that the cost function will have a single minimum, as in the age replacement case, for low values of r .

As the replacement level increases however, the cost function tends towards a double minimum form, especially in the presence of large failure cost as a proportion of replacement cost. The reason for the double minimum is clear. If the replacement limit is high, the difference between the failure and replacement limits is reduced, whereas the difference between the replacement limit and zero has increased. Because the distance between the two limits has fallen, this results in an increased chance that the system will fail between two inspections. This would imply a shorter inspection interval, but this would mean more frequent inspections during the time when the system is below the replacement level.

Thus, at the first minimum, inspections are often enough to reduce the chance that a system may fail before the next inspection interval. At the second minimum, the inspection period is long enough so that the first inspection is likely to find the degradation level between the replacement and failure limits. Beyond this minimum, the cost function increases, tending toward the limit, which is equivalent to the case where no inspections take place. Between these two minima, we must have a local maximum. This represents the worst of both positions, as the increased inspection

interval is too large to be able to avoid failures occurring, but too small for the first inspection to reveal a level of degradation between the replacement and failure limits.

Now we have seen why these shapes of curve occur, it is important to see how the parameters effect the cost functions. Firstly let us consider the effect of increasing the process variance, while holding the rate of degradation constant. This has a very predictable effect. Low process variability results in predictable behaviour, which causes the problems described above to become more prevalent. Since we know with more certainty the level of degradation at a specific point in time, we can more accurately choose the inspection intervals, hence reducing costs. However, the cost of making the wrong decision (as in the case of the local maximum described above) is increased. This emphasises the local minima and maximum in the cost function.

In the case of high process variance, the increased uncertainty makes the problem described above less important, and the shape of the curve becomes smoother, so that the local maximum and minima are less pronounced. This is exactly as we would expect.

It is also very important to consider which of these two local minima will in fact represent the global minimum of the cost function. This is largely determined by the costs of replacement and cost of failure. Essentially there are two cases: those of the single and double minima.

In the case of a single minimum, the results are largely as expected in, say, an age replacement model. Increasing failure cost relative to replacement cost results in optimal inspection period becoming closer to the origin.

In the case of a double minimum, the effects are more interesting and practically more important. It is the relative values of the replacement and failure cost which determines which of the local minima is the global minimum. Essentially, for large failure cost, the first local minimum is emphasised and the global minimum tends to be located at this minimum point. This is also intuitively clear, since for larger failure costs it is more important that we prevent failure, rather than have our first inspection occur at a time when we expect the degradation to be between the replacement and failure levels. For smaller values of the failure cost, emphasis is placed on the second minimum, which may even disappear, so that the cost function becomes monotonic decreasing beyond

the local maximum. In this case it is the second local minimum which becomes the global minimum.

The question of these double local minimum values is also important practically. As we have seen, it is possible that small changes in the costs, replacement levels and process variability, can have large changes on the optimal policy. It is perhaps also the case (although this has not been investigated) that the optimal policy is not a continuous function of the parameters of the model. This is particularly the case in the double minimum case, when small changes in the cost of failure can result in the global minimum jumping from one local minimum to the other.

Properties of Optimal Inspection Policies

Having looked at the properties of the cost functions, let us now consider the optimal inspection policies themselves. Tables 3.1 to 3.3 show respectively, the average cost optimal periodic policy, the discounted total cost optimal periodic policy, and the discounted total cost for the optimal non-periodic policy, expressed as a percentage of the costs of the periodic policy.

As we have noted above the average and discounted total cost periodic policies have largely the same properties, so we consider them together. These periodic policies behave largely as expected, and as described in the above discussion. Firstly, we consider the effect of the costs of replacement on the optimal policy, basing our comments on the results shown in Table 3.1. It can be seen that there are two effects on the optimal inspection interval, based on movements in the cost of failure and replacement. The first of these relates to the absolute value of these costs, and the second to the ratio of failure cost to replacement cost. Increasing the absolute values of either cost tends to make the inspection interval become shorter. This is simply because the cost of inspection becomes relatively cheaper, and so it is optimal to inspect more often. Similarly, if the costs are increased (or decreased) so that the ratio of failure to replacement costs increases, this has the effect of decreasing the optimal inspection interval. The reverse of this situation is that an increase in the costs which reduces the ratio of failure to replacement costs, may increase the inspection interval. The exact effect on the inspection interval depends upon the interaction of these two factors. In

some cases the increasing costs will cause the inspection interval to fall, but this depends upon the change in the cost ratio.

		Optimal Inspection Interval and Average Cost Per unit Time								
		Low Variance ($a=20$)			Medium Variance ($a=10$)			High Variance ($a=5$)		
		$C_F=10$	30	100	$C_F=10$	30	100	$C_F=10$	30	100
Low Rep. Limit $r=0.3$	$C_R=5$	0.9 8.346	0.65 10.746	0.5 13.249	0.95 8.706	0.5 12.640	0.3 15.726	1.2 8.754	0.4 14.404	0.2 20.734
	20		0.95 27.322	0.65 37.259		1.15 27.724	0.5 43.010		1.55 27.076	0.35 47.416
	50			0.8 74.855			0.85 80.001			0.95 82.620
Med. Rep. Limit $r=0.6$	$C_R=5$	0.9 8.424	0.3 11.689	0.2 13.931	0.95 8.729	0.25 13.133	0.15 17.919	1.1 8.710	0.25 15.172	0.15 25.777
	20		0.95 27.223	0.25 34.187		1.05 27.506	0.2 37.633		1.25 6.904	0.15 42.951
	50			0.35 71.786			0.3 73.816			0.3 75.503
High Rep. Limit $r=0.9$	$C_R=5$	1.1 9.487	0.1 22.039	0.05 34.366	1.25 9.408	0.15 22.986	0.05 55.717	1.65 9.057	0.2 23.591	0.1 54.892
	20		1.0 28.464	0.05 44.8514		1.1 28.150	0.05 53.576		1.35 27.106	0.1 62.124
	50			0.05 65.822			0.05 72.884			0.1 76.588

Table 3.1 - This figure shows the optimal inspection interval, and corresponding discounted total cost subject to variations in replacement limit r , degradation process variance, cost of failure and cost of replacement.

We can certainly say however that, from these results, an increase in failure cost, holding all other costs fixed, will result in a decreased inspection interval. Likewise, decreasing the cost of replacement, holding other costs constant will result in an increased inspection interval. Also, increasing both costs, and holding them in the same ratio, will result in a decreased inspection interval. These results show the importance of using an appropriate model for choosing an inspection interval. It is difficult to tell in any particular case how the costs, and other variables that have an effect, will affect the optimal policy.

		Optimal Inspection Interval and Discounted Total Cost								
		Low Variance ($a=20$)			Medium Variance ($a=10$)			High Variance ($a=5$)		
		$C_F=10$	30	100	$C_F=10$	30	100	$C_F=10$	30	100
Low Rep. Limit $r=0.3$	$C_R=5$	0.9 83457	0.65 107462	0.5 132488	0.95 87054	0.5 126398	0.3 157259	1.2 87538	0.4 144035	0.2 207333
	20		0.95 273206	0.65 372577		1.15 277220	0.5 430083		1.55 270735	0.35 474144
	50			0.8 740515			0.85 799975			0.95 826155
Med. Rep. Limit $r=0.6$	$C_R=5$	0.9 84231	0.3 116884	0.2 139309	0.95 87283	0.25 131324	0.15 179187	1.1 87093	0.25 151719	0.15 257761
	20		0.95 272217	0.25 341855		1.05 275039	0.2 376316		1.25 269019	0.15 429498
	50			0.35 717830			0.3 738130			0.3 754999
High Rep. Limit $r=0.9$	$C_R=5$	1.1 94860	0.1 220381	0.05 343653	1.25 94075	0.15 229850	0.05 439203	1.65 90561	0.2 235899	0.1 548898
	20		1.4 284641	0.05 448500		1.1 281479	0.05 535738		1.35 271028	0.1 621216
	50			0.05 658193			0.05 728809			0.1 765853

Table 3.2 - This figure shows the optimal inspection interval, and corresponding discounted total cost subject to variations in replacement limit r , degradation process variance, cost of failure and cost of replacement.

It is also clear from the table that the variability of the degradation process plays an important role, but seems less clear cut than in the above case. Increasing the variability of the process has various effects depending on the level of the costs involved. The general pattern is that the inspection interval will increase for systems whose failure and replacement costs are not large relative to the inspection cost, while it will decrease for systems that have high costs relative to the cost of inspection. The reason for this is clear. Under a periodic policy, it is important that the replacement region is not missed by inspections, so increasing the variability of the process makes this more likely to happen. When costs are small, the inspection interval increases so that the inspection is unlikely to take place before the replacement limit has been reached. On the other hand with high costs it is more important that the system does not ever fail, so that a much lower inspection interval occurs.

		Discounted Total Cost as percentage of Periodic Optimal DTC								
		Low Variance ($\sigma=20$)			Medium Variance ($\sigma=10$)			High Variance ($\sigma=5$)		
		$C_F=10$	30	100	$C_F=10$	30	100	$C_F=10$	30	100
Low Rep. Limit $r=0.3$	$C_R=5$	100% 83457	99.5% 106860	94.5% 125259	100% 87046	98% 123917	97.5% 153513	100% 87532	99% 142844	98% 203768
	20		100% 273206	99.5% 370651		100% 277220	98.5% 423975		100% 270736	100% 473395
	50			100% 748453			100% 799519			100% 825597
Med. Rep. Limit $r=0.6$	$C_R=5$	98% 82310	85% 99463	81.5% 113706	98% 85367	88% 116003	84.5% 151239	99% 86282	92% 139871	88.5% 228721
	20		99% 270451	94% 320533		99.5% 273305	94.5% 355392		100% 268731	95% 408925
	50			98% 702130			98.5% 728081			99% 747815
High Rep. Limit $r=0.9$	$C_R=5$	91.5% 94860	65.2% 143825	74% 254144	96% 90311	76% 172898	83.5% 267128	98.5% 89340	84.5% 199920	89.5% 490960
	20		91% 259020	85% 382543		94.5% 266366	89% 475919		98% 265006	92.5% 575126
	50			96.5% 635076			95% 692967			97% 741044

Table 3.3 - This figure shows the minimum discounted total cost of the optimal non-periodic Inspection Policy subject to variations in the replacement limit r , degradation process variance, cost of failure and cost of replacement.

The other parameter that we have considered is the replacement level r . As we have said above, lower replacement levels imply a simple age replacement model. For small r it is clear that the aim of the policy is to stop the process as close as possible to, but not above, the failure level. It is clear from the table of costs that this level of r is less costly than higher levels of r , for low replacement and failure costs. For higher levels of failure and replacement costs, it is clear that the low level of r results in higher costs. In this case it is often much better to consider a medium or high level of r . Which of these gives the least cost is uncertain, but it appears that the process variability plays a big role.

In some cases, the parameter r may be chosen at the discretion of the decision-maker. In this case, it is clear that the optimal value of r may be found by considering the joint optimisation of r and the inspection interval. This is easily done using our model. We note further that fixing the inspection interval allows us to use the model to find the optimal replacement limit, as in, for example, Park (1988a).

Now let us consider the optimal non-periodic policies. Table 3.3 on page 89 shows the minimum costs for various choices of parameters, and compares them to the corresponding non-periodic policies. Table 3.4 on page 91 shows some examples of the optimal policy, and corresponding costs, along with the optimal periodic policy in the same case. In the table, values of 100% indicate that the cost was greater than 99% of the periodic policy.

Firstly, table 3.3 shows how the minimum costs for the periodic and non-periodic cases compare. We note, and emphasise, that in all but one extreme case, the optimal general policy is non-periodic. In only one of the cases shown, did the optimal general policy coincide with the optimal periodic policy. It is clear then that in many cases the periodic policy is very nearly optimal. This is seen particularly for the case of a low replacement limit, with exception of the extreme case in which the failure cost is 20 times greater than the replacement cost. The reason for this is simply that low replacement limit is equivalent to an age replacement model and so we are looking for a replacement time rather than an inspection time.

It can be further seen that the benefits of having a non-periodic structured policy is seen most when we have a high ratio of failure to replacement cost, with a high replacement limit, and with a low process variance. This is largely as we would expect. The non-periodic policy will always result in a larger number of inspections, but is more likely to prevent failure occurring. Thus the benefit is seen most when the consequences of failure are greatest, and this is when the ratio of failure to replacement costs is high, and both of these are high relative to the inspection cost.

These findings are again found in Table 3.3, which gives the optimal inspection policies for various values of C_F , C_R and r . The policies themselves are as we would expect, being decreasing functions of the level of degradation.

Therefore, we have seen in this section that the model produces policies that are sensible and useful in the examples we have considered. The model produces results that are consistent with alternative models, such as the simple age replacement model, and with common sense. Also, it shows new and interesting behaviour of cost functions for degrading systems, which has important consequences for application to real life systems.

			Optimal Non-periodic Policy under Discounted Cost Criterion																Cost of Optimal Policy	Optimal Periodic Policy	
			Discretised State x																		
Rep. Limit	C_R	C_F	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.5	0.65	0.7	0.75		τ^*	DTC
$r=0.25$	5	10	0.9	0.85	0.75	0.7	0.65	0.6											83458	0.9	83458
	10	50	0.65	0.6	0.55	0.55	0.5	0.45											193329	0.65	193512
	20	200	0.6	0.55	0.5	0.45	0.4	0.4											413721	0.6	415688
$r=0.5$	5	10	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.5	0.45	0.4	0.35						83029	0.9	83529
	10	50	0.55	0.55	0.5	0.5	0.45	0.4	0.4	0.35	0.3	0.25	0.25						180262	0.35	191100
	20	200	0.45	0.45	0.4	0.35	0.35	0.3	0.3	0.3	0.25	0.2	0.2						357592	0.3	379530
$r=0.75$	5	10	0.8	0.8	0.75	0.7	0.65	0.6	0.55	0.5	0.45	0.4	0.35	0.3	0.25	0.25	0.2	0.05	82003	0.95	87833
	10	50	0.55	0.5	0.5	0.45	0.4	0.4	0.35	0.3	0.3	0.25	0.2	0.2	0.15	0.1	0.1	0.05	170443	0.15	212081
	20	200	0.45	0.4	0.4	0.35	0.3	0.3	0.25	0.25	0.2	0.2	0.15	0.15	0.1	0.1	0.05	0.05	326842	0.1	394383

Table 3.4 . This table shows the optimal policies for various combinations of the parameters r , C_R and C_F . The optimal periodic policy and corresponding costs are shown for comparison

3.7 Conclusion

In this chapter we have considered the problem of computing optimal inspection intervals for systems whose degradation follows a Lévy process, and fail when this degradation reaches a given threshold level. We have applied arguments from dynamic programming theory to derive integral equations and optimality equations to determine the costs of maintenance for such systems. The solutions of these equations can then be used to determine the optimal inspection policy for the system.

The main example considered in the chapter is that of a Gamma Process. The results of the model provide sensible and realistic inspection policies for such systems, and gives insight into the behaviour of the system and the effect of applying various inspection policies.

There are many extensions of this model that could be considered. Firstly, we could consider the case of imperfect inspection, and this case is treated as special case of chapter 4. Another important extension is to the case where the system is not replaced, but instead is imperfectly repaired. This is considered in chapter 5.

Now let us consider some possible extensions that we do not consider in this thesis. Firstly, it is important to note that we have used a stylised example to show the properties of the model. There can be no substitute however for considering a real life system, so that the results can be compared with the observed reality. To do this it would be useful to extend this model to the case of a multivariate degradation process, so that more than one indicator of failure could be considered. While this extension is theoretically possible, we have found that the solutions of the integral equations and optimality equations are more difficult to find. This is because the equations then involve multi-dimensional integrals, and the probabilities are much more difficult to compute, especially when the processes are correlated.

As we noted in chapter 2, the threshold failure model can be criticised since it implies that the system cannot fail unless the degradation reaches a specific point. As we have suggested, we believe that this criticism is unfounded for many modern complex systems, since failure may only be rarely or never observed, and is often itself extrinsic to the measure of degradation being considered. However, for systems where this is not the case it would be interesting to consider the case when the system may fail at any level, according to a degradation dependent hazard rate.

Also, it would be interesting to consider a model in which we are not only trying to reduce the costs of inspection. For example, we may wish to find the least cost subject to a safety constraint etc.

We note that the model as it stands only applies to a simple Lévy process model of degradation. The model can be used however for systems modelled by a generalised gamma process as defined by Van-Noortwijk (1996). If we assume that the results above are derived conditional on the value of the parameter α of the gamma process, we may compute the overall expected cost under a generalised gamma process, taking expectations with respect to the distribution of α .

The model as it stands however, seems to provide a useful first step in the analysis of optimal inspection and maintenance of systems which degrade stochastically. We believe these models will become increasingly important, as degradation type models gain more emphasis over the traditional failure time model.

Chapter 4

Optimal Inspection Policies in the Presence of Covariates

4.1 Introduction

In the previous chapter, we obtained approximate optimal inspection policies for systems whose degradation is perfectly observable. In many cases however, it is not possible to observe the system in this way, and a proxy for the true degradation process must be used. With this in mind, we now extend the analysis of the previous chapter to the case of observation of a covariate process. The system we shall consider is identical to that considered in chapter 3, in all respects, except that inspections now reveal the state of the covariate process, rather than the degradation process of the system. Once again we apply ideas from dynamic programming and Markov Decision Processes to obtain integral equations and optimality equations which we solve numerically to obtain the (approximate) optimal policies. As in chapter 3 we consider both average cost and discounted total cost criteria, for periodic inspections and discounted total cost criterion for non-periodic inspections.

We begin by looking at the general case of a Lévy degradation process, and a Lévy covariate process. We then look at the case of imperfect inspection of the degradation process, which is a special case of a covariate process. As in the previous chapter we

then consider the specific case of Gamma process degradation, subject to imperfect inspection. The model proposed is quite general in nature, and can be applied to any case of Lévy process degradation with threshold failure. The exact nature of the degradation and covariate processes determines the extent to which useful results can be obtained.

4.2 Underlying Inspection and Replacement Model

We now consider the underlying inspection and replacement model, which is largely the same as that considered in Chapter 3. The list below makes all of the assumptions explicit. Throughout this chapter the process $X = \{X_t : t \geq 0\}$ denoted the degradation process of the system, and the process $Y = \{Y_t : t \geq 0\}$ denotes the corresponding covariate process. We make, for now, no assumptions about the relationship between the degradation and covariate process.

MODEL ASSUMPTIONS M2 (Inspection of a covariate process):

1. We assume the state-space of the covariate process S_c is partitioned into intervals $A_0, A_1, A_2, \dots, A_n$, such that $A_0 = [0, s_0)$, and $A_k = [s_k, s_{k+1})$ for all $k = 0, 1, \dots, n-1$ with $s_k < s_{k+1}$ and $s_{n+1} = \infty$.
2. We assume the degradation process has state space S_d partitioned into a set $B = [b, \infty)$ and its complement $[0, b)$. The system is deemed to have failed when the degradation process hits the critical set B , corresponding to level of degradation b .
3. Each inspection reveals the true state of the covariate process Y_t .
4. It is assumed that a new system has covariate level $Y_0 = \lambda \in A_0$. A new system is assumed to have degradation level 0.
5. If, at an inspection $Y \in A_0$, the system is not replaced and is allowed to continue operating until the next inspection. Each inspection incurs a cost C_0 which is regarded as the cost of inspection and cost of any loss incurred by the system being unavailable during inspection.

6. If, at an inspection, $Y \in A_i$ for $i=1,2,\dots,n$ then the system is replaced at a cost C_i . This may be regarded as the total cost of inspection, of replacing the system and of costs incurred by having the system unavailable during replacement. The replaced system is assumed to have covariate level starting at level λ .
7. The system is deemed to have 'failed' at the first moment the true level of degradation process X_t hits the set B and this failure is immediately observed resulting in an immediate replacement of the system at cost C^F . Otherwise the state of the system cannot be revealed.

The main difference between these assumptions and those for perfect inspection given in Chapter 3 is that all of the replacement decisions are now based on the covariate level Y_t rather than the level of degradation.

The fact that measurements are made of a covariate means that much greater care is required with definitions. The most important point concerns the nature of the covariate process. We make the assumption that the covariate process is a time-homogeneous Markov process. It may be possible to consider more general Markov processes, but we do not consider this here. We note that all Levy processes fall into this category, as well as many diffusion processes.

Secondly we assume that the covariate process is bounded, in the following sense. If $X = \{X_t : t \geq 0\}$ represents the theoretical covariate or degradation process we are using, we define a truncated process $X' = \{X'_t : t \geq 0\}$ such that

$$X'_t = \begin{cases} A & X_t < A \\ X_t & A \leq X_t \leq B \\ B & X_t > B \end{cases} \quad \text{for } A, B \in \mathbf{R}$$

We shall assume throughout that the process we are using is truncated in this form. In terms of degradation process with a threshold failure model it is natural to let $A = 0$, and to let B equal to the failure threshold level. In the case of a covariate process such natural truncation points may not arise. The point of this truncation is to remove the possibility of infinite covariate or degradation values. In practical terms, the level of degradation and covariate will have a known range of possible values, and so the imposition of upper and lower limits will not cause any problems. It should be noted

that we have already informally made an assumption that $X'_t = 0$ for $X_t < 0$ in the analysis of non-monotonic degradation processes given in Chapter 3. We formalise this here as it is more important in the case of a covariate process, to avoid the theoretical possibility of infinite levels of the covariate.

The assumptions described above have largely been made for technical reasons, and so that an integral equation with infinite limit may be avoided. We comment now on the implications for practical application of these methods. Firstly, the nature of the covariate process is that it must be in some way dependent on the degradation process, or at least correlated with it. Otherwise there would be no benefit in observing the covariate as a proxy of the degradation level. This means that, to some extent the two processes measure the same effects, and so we should expect them to have similar values and patterns (perhaps subject to scaling). It is thus reasonable to assume that the covariate process cannot take negative values, since a similar assumption is made regarding the degradation process.

4.3 Periodic Inspection of a Covariate Process: Average Cost Criterion

Following the methodology outlined in Chapter 3, we now consider the optimisation of periodic inspections of a covariate process. As in chapter 3, we derive integral equations for the average cost per cycle, and average time of a cycle, and appeal to the renewal reward theorem to obtain the average cost per unit time. In this case, and throughout this chapter, we restrict attention to policies which are dependent only on the current observation of the covariate process. The reason for this is to simplify the problem, and avoid problems with full history dependence.

4.3.1 Derivation of Expected Cost per Unit Time

As we have described in the section 4.2, the main difference between this model and that of Chapter 3 is that the decision variable is now the value of the covariate Y . The analysis may be simplified, however, by simply conditioning on the current level of

degradation. By doing this, we can apply the results of the previous chapter, before integrating over the density of the degradation level, conditional on the covariate level.

With the previous remarks in mind we define $V(y, \tau)$ to be the random cost per cycle for a system currently in state y , that is with covariate level y , and likewise let $L(y, \tau)$ be the remaining length of a cycle for a system currently with covariate level y . Now we define the expected cost and length per cycle as

$$\begin{aligned} v(y, \tau) &= E[V(y, \tau) | Y_0 = y] \\ l(y, \tau) &= E[L(y, \tau) | Y_0 = y] \end{aligned}$$

Now, applying the same methodology as in chapter 3 we compute these by deriving an integral equation based on a recursive relationship, proved using dynamic programming arguments.

Expected Cost per Cycle

Let us assume that X_t and Y_t represent the degradation process at time t , and the value of the covariate process at time t , respectively. In addition we define the process M_t to be the maximum of the process X_t over the period $[0, t)$, and H_x^B to be the hitting time of the critical set B , by the degradation process X_t started at x . Then, conditional on the initial values of these processes, and on the values of M_τ and Y_τ we can write

$$\begin{aligned} V(y, \tau) | Y_\tau, M_\tau, X_0 = x, Y_0 = y &= [C_0 + V(Y_\tau, \tau)] 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} + C_1 1_{\{Y_\tau \in A_1, M_\tau \notin B\}} + \dots \\ &\dots + C_{n-1} 1_{\{Y_\tau \in A_{n-1}, M_\tau \notin B\}} + C_n 1_{\{Y_\tau \in A_n, M_\tau \notin B\}} + C_n^{(F)} 1_{\{M_\tau \in B\}} \\ &= V(Y_\tau, \tau) 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} + \sum_{i=0}^n C_i 1_{\{Y_\tau \in A_i, M_\tau \notin B\}} + C^F 1_{\{M_\tau \in B\}} \end{aligned} \quad (4.1)$$

Which is essentially the same as the perfect inspection case except that X has been replaced by Y . Then, the conditional expectation of $V(y)$ given Y_τ is simply

$$\begin{aligned} E(V(y, \tau) | Y_\tau, M_\tau, X_0 = x, Y_0 = y) &= E \left(V(Y_\tau, \tau) 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} + \sum_{i=0}^n C_i 1_{\{Y_\tau \in A_i, M_\tau \notin B\}} + C^F 1_{\{M_\tau \in B\}} \middle| Y_\tau, M_\tau, X_0 = x, Y_0 = y \right) \\ &= E(V(Y_\tau, \tau) | Y_\tau, M_\tau, X_0 = x, Y_0 = y) 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} + \sum_{i=0}^n C_i 1_{\{Y_\tau \in A_i, M_\tau \notin B\}} + C^F 1_{\{M_\tau \in B\}} \end{aligned} \quad (4.2)$$

Now, we define $v^x(y, \tau) = E(V(y, \tau) | X_0 = x, Y_0 = y)$ to be the expected cost per cycle conditional on the current state of the degradation process.

Then, writing the density of Y_τ, M_τ conditional on $X_0 = x$ as $f_\tau(u, m | x)$, we have

$$\begin{aligned} v^x(y, \tau) &= E(V(y, \tau) | X_0 = x, Y_0 = y) \\ &= \int \int_{A_0 S_d \setminus B} v(u, \tau) f_\tau(u, m | x) dm du + \sum_{i=0}^n C_i P^x(Y_\tau \in A_i, M_\tau \notin B) + C^F P^x(M_\tau \in B) \end{aligned} \quad (4.3)$$

Since, given Y_τ , $V(Y_\tau, \tau)$ is independent of X_0 and Y_0 . (Superscript x on these probability functions signifies that the probabilities are conditional on $X_0 = x$.) Now, we constrain the covariate process to have only positive values by assuming $v(y, \tau) = v(0, \tau)$ for all $y < 0$. This assumption is not restrictive, since we have already assumed that the process is truncated below, and by re-scaling the process we may make this lower limit equal to zero.

As in section 3.3.1, this assumption implies that the integral in equation (4.3) above becomes

$$\begin{aligned} &\int \int_{A_0 S_d \setminus B} v(u, \tau) f_\tau(u, m | x) dm du \\ &= v(0, \tau) P^x(Y_\tau < 0, M_\tau \notin B) + \int_0^{s_0} \int_0^b v(u, \tau) f_\tau(u, m | x) dm du \end{aligned} \quad (4.4)$$

Now, we define analogously to the results of chapter 3, the functions

$$c^x(y, \tau) = \sum_{i=0}^n C_i P^x(Y_\tau \in A_i, M_\tau \notin B) + C^F P^x(M_\tau \in B) \quad (4.5a)$$

$$u^x(y, \tau) = P^x(Y_\tau < 0, M_\tau \notin B) \quad (4.5b)$$

$$K_\tau^x(u | y) = \int_0^b f_\tau(u, m | x) dm \quad (4.5c)$$

so that we can rewrite equation (4.3) as

$$v^x(y, \tau) = c^x(y, \tau) + u^x(y, \tau) v(0, \tau) + \int_0^{s_0} v(u, \tau) K_\tau^x(u | y) du \quad (4.6)$$

Now, the expected cost per cycle, conditional on the level of the degradation process is clearly given by

$$v(y, \tau) = E(v^x(y, \tau)) = \int_0^b v^x(y, \tau) f(x | y) dx \quad (4.7)$$

where $f(x | y)$ is the density of $X_* | Y_* = y$, $0 < X_* < b$. The density is conditional on the fact that the level of degradation must be less than b , since otherwise the system would have failed already. The fact we are making an inspection shows that failure cannot have occurred.

Substituting the right hand side of equation (4.6) into equation (4.7) we obtain

$$\begin{aligned} v(y, \tau) = & \left\{ \int_0^b c^x(y, \tau) f(x | y) dx \right\} + v(0, \tau) \left\{ \int_0^b u^x(y, \tau) f(x | y) dx \right\} \\ & + \int_0^{s_0} v(u, \tau) \left\{ \int_0^b K_\tau^x(u | y) f(x | y) dx \right\} du \end{aligned}$$

so that

$$v(y, \tau) = c(y, \tau) + u(y, \tau)v(0, \tau) + \int_0^{s_0} v(u, \tau) K_\tau(u | y) du \quad (4.8)$$

The functions c , u and K are defined by

$$c(y, \tau) = \int_0^b \left\{ \sum_{i=0}^n C_i P^x(Y_\tau \in A_i, M_\tau \notin B) + C^F P^x(M_\tau \in B) \right\} f(x | y) dx \quad (4.9a)$$

$$u(y, \tau) = \int_0^b P^x(Y_\tau < 0, M_\tau \notin B) f(x | y) dx \quad (4.9b)$$

$$K_\tau(u | y) = \int_0^b \int_0^b f_\tau(u, m | x) f(x | y) dm dx \quad (4.9c)$$

A difficulty which remains hidden in this analysis is computation of the joint density of the maximum variable of the process M , and the value of the covariate process Y . It is possible that the density can be obtained by conditioning on the true level of degradation at time τ . Thus, using f to denote density functions, we have

$$\begin{aligned}
f_{Y_\tau, M_\tau}(u, m) &= \int_0^\infty f_{Y_\tau, M_\tau | X_\tau}(u, m | z) f_{X_\tau}(z) dz \\
&= \int_0^\infty f_{Y_\tau | X_\tau}(u | z) f_{M_\tau | Y_\tau, X_\tau}(m | u, z) f_{X_\tau}(z) dz \\
&= \int_0^\infty f_{Y_\tau | X_\tau}(u | z) f_{M_\tau, X_\tau}(m, z) dz
\end{aligned} \tag{4.10}$$

with all densities conditional on $X_0 = x$. So that the density is obtained as an integral of the joint density of X and M and the density of the covariate process conditional upon the degradation process. In the case of a Wiener process, this density can be obtained in terms of the standard Normal density function.

As in chapter 3, some simplification is possible if we are dealing with a monotone degradation process. This is considered in the next section.

Expected Length of a Cycle

As above we consider the value of $L(y, \tau)$, conditional on $M_\tau, Y_\tau, H_x^B, X_0 = x$ and $Y_0 = y$.

Then, applying the same method as before we have

$$\begin{aligned}
L(y, \tau) | Y_\tau, M_\tau, H_x^B, X_0 = x, Y_0 = y \\
= L(Y_\tau, \tau) 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} + \tau 1_{\{Y_\tau \notin A_0, M_\tau \notin B\}} + H_x^B 1_{\{H_x^B < \tau\}}
\end{aligned} \tag{4.11}$$

So that, upon taking expectations over L we get

$$\begin{aligned}
E(L(y, \tau) | Y_\tau, M_\tau, H_x^B, X_0 = x, Y_0 = y) \\
= E(L(Y_\tau, \tau) 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} | Y_\tau, M_\tau, X_0 = x, Y_0 = y) + \tau 1_{\{Y_\tau \notin A_0, M_\tau \notin B\}} + H_x^B 1_{\{H_x^B < \tau\}}
\end{aligned}$$

Now, taking expectations with respect to Y_τ and M_τ . Then, using the notation of the previous subsection we have

$$\begin{aligned}
I^x(y, \tau) &= E(L(y, \tau) | X_0 = x, Y_0 = y) \\
&= \int_{A_0 \setminus B} \int l(u, \tau) f_\tau(u, m | x) dm du + \tau P^{y, x}(Y_\tau \notin A_0, M_\tau \notin B) + \int_0^\tau h g_B(h | x) dh \\
&= \left\{ \tau P^{y, x}(Y_\tau \notin A_0, M_\tau \notin B) + \int_0^\tau h g_B(h | x) dh \right\} + l(0, \tau) P^{y, x}(Y_\tau < 0, M_\tau \notin B) \\
&\quad + \int_0^{s_0} \int_0^b l(u, \tau) f_\tau(u, m | x) dm du
\end{aligned} \tag{4.12}$$

Where $g_B(h|x)$ is the density function of the hitting time of the critical set B starting at x . So, unconditionally we have

$$l(y, \tau) = \int_{-\infty}^{\infty} l^x(y, \tau) f(x|y) dx \quad (4.13)$$

where, as before, $f(x|y)$ is the density of $X_\bullet | Y_\bullet = y$, $0 < X_\bullet < b$.

$$\begin{aligned} l(y, \tau) = & \int_0^b \left\{ \tau P^{y,x}(Y_\tau \notin A_0, M_\tau \notin B) + \int_0^\tau h g_B(h|x) dh \right\} f(x|y) dx \\ & + l(0, \tau) \int_0^b P^{y,x}(Y_\tau < 0, M_\tau \notin B) f(x|y) dx + \int_0^{s_0} \int_0^b l(u, \tau) f_\tau(u, m|x) f(x|y) dm du dx \end{aligned} \quad (4.14)$$

Now we define

$$d(y, \tau) = \int_0^b \left\{ \tau P^{y,x}(Y_\tau \notin A_0, M_\tau \notin B) + \int_0^\tau h g_B(h|x) dh \right\} f(x|y) dx \quad (4.15a)$$

$$u(y, \tau) = \int_0^b P^{y,x}(Y_\tau < 0, M_\tau \notin B) f(x|y) dx \quad (4.15b)$$

$$K_\tau(u|x) = \int_0^b \int_0^b f_\tau(u, m|x) f(x|y) dm dx \quad (4.15c)$$

So that we can rewrite the equation in our standard form as

$$l(y, \tau) = d(y, \tau) + l(0, \tau)u(y, \tau) + \int_0^{s_0} l(u, \tau)K_\tau(u|x) du \quad (4.16)$$

If all functions can be efficiently computed, this equation can now be solved by numerical means, giving the expected length of a cycle. It can be seen however that these equations involve probabilities of the form $P^x(Y_\tau \notin A_0, M_\tau \notin B)$. Whether or not this can be computed depends on the nature of the degradation and covariate processes.

Thus, as in the case of perfect observations, the average cost per unit time can be obtained and optimised as described in section 3.3.1. The only difference in this case is that the covariate process may have non-zero starting level. Hence we seek the inspection interval τ which minimises the expected average cost per unit time given by

$$C(\lambda, \tau) = \frac{v(\lambda, \tau)}{l(\lambda, \tau)} \quad (4.17)$$

This can be done, by simply enumerating the costs for possible values of τ and choosing the one which results in lowest cost. Alternatively some form of search algorithm may be used to find the optimal inspection interval.

4.3.2 Special Case of a Monotone Degradation Process

As in the case of perfect inspection, things can be simplified if we are dealing with a monotone degradation process. In this case however, the maximum of the process does not completely disappear from the solution, instead it is replaced by X_τ , so that we must now consider the joint distribution of the observation Y_i and the true level of degradation X_τ . Thus we have

$$\begin{aligned}
 V(y, \tau) &| Y_\tau, X_\tau, X_0 = x, Y_0 = y \\
 &= [C_0 + V(Y_\tau, \tau)] 1_{\{Y_\tau \in A_0, X_\tau \notin B\}} + C_1 1_{\{Y_\tau \in A_1, X_\tau \notin B\}} + \dots \\
 &\quad \dots + C_{n-1} 1_{\{Y_\tau \in A_{n-1}, X_\tau \notin B\}} + C_n 1_{\{Y_\tau \in A_n, X_\tau \notin B\}} + C^F 1_{\{X_\tau \in B\}} \quad (4.18) \\
 &= V(Y_\tau, \tau) 1_{\{Y_\tau \in A_0, X_\tau \notin B\}} + \sum_{i=0}^n C_i 1_{\{Y_\tau \in A_i, X_\tau \notin B\}} + C^F 1_{\{X_\tau \in B\}}
 \end{aligned}$$

We define, as in all of this sub-section, $f_\tau(u, z | x)$ as the joint density of Y_τ and X_τ , conditional on $X_0 = x$. So, expanding as before we have

$$v^x(y, \tau) = \int_0^{s_0} \int_0^b v(u, \tau) f_\tau(u, z | x) dz du + \sum_{i=0}^n C_i P^{y,x}(Y_\tau \in A_i, X_\tau \notin B) + C^F P^{y,x}(X_\tau \in B) \quad (4.19)$$

So that the integral equation is once again of the form (4.8)

$$v(y, \tau) = c(y, \tau) + u(y, \tau)v(0, \tau) + \int_0^{s_0} v(u, \tau) K_\tau(u | y) du$$

with functions defined by

$$c(y, \tau) = \int_0^b \left\{ \sum_{i=0}^n C_i P^x(Y_\tau \in A_i, X_\tau \notin B) + C^F P^x(X_\tau \in B) \right\} f(x | y) dx \quad (4.20a)$$

$$u(y, \tau) = \int_0^b P^x(Y_\tau < 0, X_\tau \notin B) f(x | y) dx \quad (4.20b)$$

$$K_{\tau}(u | y) = \int_0^b \int_0^b f_{\tau}(u, z | x) f(x | y) dz dx \quad (4.20c)$$

Likewise, the equation for the expected length of a cycle is given by

$$l(y, \tau) = d(y, \tau) + l(0, \tau)u(y, \tau) + \int_0^{s_0} l(u, \tau)K_{\tau}(u | x) du \quad (4.21)$$

Where

$$d(y, t) = \int_0^b \left\{ \tau P^{y,x}(Y_{\tau} \notin A_0, X_{\tau} \notin B) + \int_0^{\tau} h g_B(h | x) dh \right\} f(x | y) dx \quad (4.22a)$$

$$u(y, \tau) = \int_{-\infty}^{\infty} P^{y,x}(Y_{\tau} < 0, X_{\tau} \notin B) f(x | y) dx \quad (4.22b)$$

$$K_{\tau}(u | x) = \int_0^b \int_0^b f_{\tau}(u, z | x) f(x | y) dz dx \quad (4.22c)$$

So that these equations may be solved in the usual way, giving expected cost and length of a cycle, allowing the long run average cost per unit time to be calculated.

4.3.3 Cost of No inspections and condition monitoring

When no inspections are carried out, the problem reduces to that of the previous chapter, and the average cost per unit time is given by

$$\mathbb{C}(0, \infty) = \frac{C_F}{\int_0^{\infty} [1 - G_B(h | 0)] dh} \quad (4.23)$$

Where G is the distribution function of the hitting time of the critical set starting at zero and C_F is the cost of replacement on failure.

In the case of continuous condition monitoring the most consistent generalisation of the above model is to assume that the covariate process is observed continuously and the system is replaced whenever the covariate reaches a specified level, or at failure whichever occurs first. Let us assume that the system is replaced when the covariate reaches a replacement level r , or when the true level of degradation reaches the failure limit b . The cost of replacement and failure are defined as C_R and C_F respectively. Then

assume the cost of CCM per unit time is given by ρ , so that average cost is given by the expected cost to replacement divided by the expected time to replacement, thus

$$\bar{C}_{CCM} = \frac{E(\rho T + C_R 1_{\{H_r^Y < H_b^X\}} + C_F 1_{\{H_r^Y > H_b^X\}})}{E(T)} \quad (4.23)$$

which upon simplification gives

$$\bar{C}_{CCM} = \rho + \frac{C_R P(H_r^Y < H_b^X) + C_F P(H_r^Y > H_b^X)}{\int_0^\infty t g_T(t) dt} \quad (4.24)$$

where H_x^Z represents the hitting time of the point x by the process Z . The random variable T represents the failure/replacement time of the system given by the minimum of the hitting times for the two processes: $T = H_r^Y \wedge H_b^X$. Clearly, to compute the average cost we require the joint density function of the random variables H_r^Y and H_b^X . Since the processes X and Y are not independent, it is likely that obtaining this distribution is not trivial. We do not consider the case of condition monitoring further in this thesis, as it is outside the main path of our analysis.

If the joint density described in the previous paragraph is available, or may be estimated by simulation, then the model should provide a reasonable model for continuous condition monitoring of systems with Lévy process degradation and Markov covariate processes.

4.4 Optimal Periodic Inspection of a Covariate Process : Discounted Cost Criterion

Having looked at the problem of optimal inspections of a covariate process under the expected average cost criterion, we now consider the case of the discounted total cost criterion. As in chapter 3, we expect that the results under this criterion will be very similar to those under the average cost criterion. As before, we apply dynamic programming arguments to obtain integral equations for the cost functions involved.

4.4.1 Derivation of Discounted Total Cost

The arguments here are similar to those given in section 3.4.1, the only difference being the additional term allowing for the non-zero initial value of the covariate process. We also make the additional assumption, as described previously, that the policy is dependent on the history of observed covariates, only through the current observation. We define $V_\delta(y, \tau)$ to be the discounted total cost incurred for a system currently in observed state y . Then in the usual way we may write

$$\begin{aligned} V_\delta(y, \tau) | (X_\tau, M_\tau, H_x^B, X_0 = x, Y_0 = y) = & e^{-\delta\tau} (C_0 + V_\delta(Y_\tau, \tau)) 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} \\ & + \sum_{i=1}^n e^{-\delta\tau} (C_i + V_\delta(\lambda, \tau)) 1_{\{X_\tau \in A_i, M_\tau \notin B\}} + e^{-\delta H_x^B} (C^F + V_\delta(\lambda, \tau)) 1_{\{H_x^B < \tau\}} \end{aligned} \quad (4.25)$$

Where λ is the assumed covariate level for a system with zero level of degradation.

Then, the conditional expectation of V is given by

$$\begin{aligned} E(V_\delta(y, \tau) | X_\tau, M_\tau, H_x^B, X_0 = x, Y_0 = y) = & e^{-\delta\tau} (C_0 + v_\delta(Y_\tau, \tau)) 1_{\{Y_\tau \in A_0, M_\tau \notin B\}} \\ & + \sum_{i=1}^n e^{-\delta\tau} (C_i + v_\delta(\lambda, \tau)) 1_{\{X_\tau \in A_i, M_\tau \notin B\}} + e^{-\delta H_x^B} (C^F + v_\delta(\lambda, \tau)) 1_{\{H_x^B < \tau\}} \end{aligned} \quad (4.26)$$

Therefore, defining $f_\tau(u, m | x)$ to be the joint density of Y_τ and M_τ , given $X_0 = x$. Then, as in the perfect inspection case we consider the conditional expectation of V given $X_0 = x$. Thus we define $v_\delta^x(y, \tau) = E(V_\delta(y, \tau) | X_0 = x, Y_0 = y)$ so that

$$\begin{aligned} v_\delta^x(y, \tau) = & E(E(V_\delta(y, \tau) | X_\tau, M_\tau, H_x^B)) \\ = & \int \int_{A_0 \setminus B} e^{-\delta\tau} v(u, \tau) f_\tau(u, m | x) dm du + e^{-\delta\tau} C_0 P^x(Y_\tau \in A_0, M_\tau \notin B) \\ & + \sum_{i=1}^n e^{-\delta\tau} (C_i + v_\delta(\lambda, \tau)) P^x(Y_\tau \in A_i, M_\tau \notin B) + (C^F + v_\delta(\lambda, \tau)) \int_0^\tau e^{-\delta h} g_B(h | x) dh \end{aligned} \quad (4.27)$$

with $g_B(h | x)$ being the density of the hitting time of the critical set. Making the usual assumption that $v(y, \tau) = v(0, \tau)$ for all $y < 0$, the integral in the above expression becomes

$$\begin{aligned} & \int \int_{A_0 \setminus B} e^{-\delta\tau} v(u, \tau) f_\tau(u, m | x) dm du \\ & = e^{-\delta\tau} v(0, \tau) P^x(Y_\tau < 0, M_\tau \notin B) + \int_0^{s_0} \int_0^b e^{-\delta\tau} v(u, \tau) f_\tau(u, m | x) dm du \end{aligned} \quad (4.28)$$

So, we may rewrite the integral equation as

$$\begin{aligned}
v_\delta^x(y, \tau) = & \left\{ \sum_{i=0}^n e^{-\delta \tau} C_i P^x(Y_\tau \in A_i, M_\tau \notin B) + C^F \int_0^\tau e^{-\delta h} g_B(h | x) dh \right\} \\
& + v_\delta(\lambda, \tau) \left\{ \sum_{i=1}^n e^{-\delta \tau} P^x(Y_\tau \in A_i, M_\tau \notin B) + \int_0^\tau e^{-\delta h} g_B(h | x) dh \right\} \\
& + v_\delta(0, \tau) e^{-\delta \tau} P^x(Y_\tau < 0, M_\tau \notin B) + \int_0^{s_0} \int_0^b e^{-\delta \tau} v(u, \tau) f_\tau(u, m | x) dm du
\end{aligned} \tag{4.29}$$

So that, the integral equation in this case is

$$v_\delta^x(y, \tau) = c^x(y, \tau) + u^x(y, \tau) v_\delta(0, \tau) + w^x(y, \tau) v_\delta(\lambda, \tau) + \int_0^{s_0} v_\delta(u, \tau) K_\tau^x(u | x) du \tag{4.29}$$

with functions defined as

$$c^x(y, \tau) = \sum_{i=0}^n e^{-\delta \tau} C_i P^x(Y_\tau \in A_i, M_\tau \notin B) + C^F \int_0^\tau e^{-\delta h} g_B(h | x) dh \tag{4.30a}$$

$$w^x(y, \tau) = \sum_{i=1}^n e^{-\delta \tau} P^x(Y_\tau \in A_i, M_\tau \notin B) + \int_0^\tau e^{-\delta h} g_B(h | x) dh \tag{4.30b}$$

$$u^x(y, \tau) = e^{-\delta \tau} P^x(Y_\tau < 0, M_\tau \notin B) \tag{4.30c}$$

$$K_\tau^x(u | y) = \int_0^b e^{-\delta \tau} f_\tau(u, m | x) dm \tag{4.30d}$$

Thus, the unconditional average cost per unit time may be obtained by averaging over the true state of degradation, given the observed level of degradation. Thus, the integral equation becomes

$$v_\delta(y, \tau) = c(y, \tau) + u(y, \tau) v_\delta(0, \tau) + w(y, \tau) v_\delta(\lambda, \tau) + \int_0^{s_0} v(u, \tau) K_\tau(u | y) du \tag{4.31}$$

with functions given by

$$c(y, \tau) = \int_0^b \left\{ \sum_{i=0}^n e^{-\delta \tau} C_i P^x(Y_\tau \in A_i, M_\tau \notin B) + C^F \int_0^\tau e^{-\delta h} g_B(h | x) dh \right\} f(x | y) dx \tag{4.32a}$$

$$w(y, \tau) = \int_0^b \left\{ \sum_{i=1}^n e^{-\delta \tau} P^x(Y_\tau \in A_i, M_\tau \notin B) + \int_0^\tau e^{-\delta h} g_B(h | x) dh \right\} f(x | y) dx \tag{4.32b}$$

$$u(y, \tau) = \int_0^b e^{-\delta\tau} P^x(Y_\tau < 0, M_\tau \notin B) f(x | y) dx \quad (4.32c)$$

$$K_\tau(u | y) = \int_0^b \int_0^b e^{-\delta\tau} f_\tau(u, m | x) f(x | y) dm dx \quad (4.32d)$$

As before these must be evaluated numerically, so that using standard methods the expected discounted total cost may be computed. Whether or not solutions can be obtained depends on being able to find the joint distribution of the maximum variable and the covariate.

4.4.2 Special Case of Monotone Degradation Process

Making the same observations as in section 4.3.2, it is clear that in the case of a monotonic degradation process, the results are simplified by replacing M_τ by X_τ wherever it occurs. This implies we must solve equation (4.31) with functions defined by

$$c(y, \tau) = \int_0^b \left\{ \sum_{i=0}^n e^{-\delta\tau} C_i P^x(Y_\tau \in A_i, X_\tau \notin B) + C^F \int_0^\tau e^{-\delta h} g_B(h | x) dh \right\} f(x | y) dx \quad (4.33a)$$

$$w(y, \tau) = \int_0^b \left\{ \sum_{i=1}^n e^{-\delta\tau} P^x(Y_\tau \in A_i, X_\tau \notin B) + \int_0^\tau e^{-\delta h} g_B(h | x) dh \right\} f(x | y) dx \quad (4.33b)$$

$$u^x(y, \tau) = \int_0^b e^{-\delta\tau} P^x(Y_\tau < 0, X_\tau \notin B) f(x | y) dx \quad (4.33c)$$

$$K_\tau^x(u | y) = \int_0^b \int_0^b e^{-\delta\tau} f_\tau(u, z | x) f(x | y) dz dx \quad (4.33d)$$

it is clear that in this case there is little difference between the case of monotonic and non-monotonic degradation processes. However, it is likely that the joint density of the covariate and degradation process is more easily found than that of the covariate and the maximum of the degradation process.

4.4.3 Costs of No Inspection and Condition Monitoring

As we remarked in the case of the average cost per unit time criterion, in the case of no-inspections, the imperfect inspection may be disregarded and the results of section 3.4.3 may be applied. Thus the cost of not carrying out inspections and allowing the system to fail is given by

$$v_\delta(0, \infty) = \frac{C_F \int_0^\infty e^{-\delta h} g_B(h|0) dh}{1 - \int_0^\infty e^{-\delta h} g_B(h|0) dh} \quad (4.34)$$

In the case of continuous condition monitoring, we can easily derive the equation for the discounted total cost, but its solution depends largely on the availability of the joint density of the hitting times of the degradation and covariate processes.

Let H_x^Z represent the hitting time of the point x by the process Z . As before let T represents the failure/replacement time of the system given by $T = H_r^Y \wedge H_b^X$, and let $U = H_r^Y - H_b^X$. Then, we can express the discounted total cost recursively as

$$V(0) | U, T = \int_0^T \rho e^{-\delta t} dt + e^{-\delta t} (C_R 1_{\{U>0\}} + C_F 1_{\{U<0\}} + V(0)) \quad (4.35)$$

So that

$$v(0) = \int_{-\infty}^{\infty} \int_0^{\infty} \left\{ \frac{\rho}{\delta} (1 - e^{-\delta T}) + e^{-\delta t} (C_R 1_{\{U>0\}} + C_F 1_{\{U<0\}} + v(0)) \right\} f(u, t) dt du \quad (4.36)$$

Expanding this we get

$$\begin{aligned} v(0) = & \frac{\rho}{\delta} \int_0^{\infty} (1 - e^{-\delta t}) f(t) dt + v(0) \int_0^{\infty} e^{-\delta t} f(t) dt \\ & + C_F \int_0^{\infty} e^{-\delta t} f(t, U < 0) dt + C_R \int_0^{\infty} e^{-\delta t} f(t, U > 0) dt \end{aligned} \quad (4.37)$$

so that the discounted total cost is given by

$$v(0) = \frac{\frac{\rho}{\delta} \int_0^{\infty} (1 - e^{-\delta t}) f(t) dt + C_F \int_0^{\infty} e^{-\delta t} f(t, U < 0) dt + C_R \int_0^{\infty} e^{-\delta t} f(t, U > 0) dt}{1 - \int_0^{\infty} e^{-\delta t} f(t) dt} \quad (4.39)$$

To compute the discounted total cost of condition monitoring we require the joint distribution of U and T , which is obtained from the joint distribution of H_r^Y and H_b^X . As we have said, computation of this density is not a trivial task. We have considered the simplest case, in which both the degradation and covariate processes are modelled by (correlated) Wiener processes, and are unable to find the density required. Of course, in the case of uncorrelated processes, the hitting times are independent, and so the densities are easily found. This is of little use for modelling purposes, however.

4.5 Optimal Non-Periodic Inspection of a Covariate Process: Discounted Cost Criteria

We now briefly consider the case of non-periodic inspections. From the analysis of the previous section it is clear that these results may be easily extended to the case of non-periodic inspections. Appendix A3 gives the theoretical derivation of the dynamic programming equation, in the general setting of a semi-Markov decision process, and this justifies our use of the previous section's results. As before we make the simplifying assumption that the policy is dependent only on the current observation simplifies the development. We note that this assumption also implies that the degradation process depends on the covariate process only through the current observation. This may or may not be a reasonable assumption, depending on the nature of the underlying system being modelled.

Our simplifying assumption that the policy is Markov, that is dependent only on the current observation, simplifies matters and means we do not have to consider the general partially observed semi-Markov decision process.

As in the case of non-periodic perfect inspection we assume a total discounted cost function of the form

$$v_{\pi}(y) = E(V_{\pi}(y) | Y_0 = y) = E\left(\sum_{n=1}^{\infty} e^{-\delta(t_1 + \dots + t_n)} C(Y_{t_{n-1}}, Y_{t_n}; \pi(Y_{t_{n-1}})) | Y_0 = y\right) \quad (4.40)$$

where π represents the inspection policy. We are only concerned with deterministic stationary policies, since the Lévy property of our degradation process implies the current time has no bearing on future levels of degradation.

We define the function v_{δ} to be the value function for the δ -optimal policy, i.e.

$$v_{\delta}(y) = \inf_{\pi} v_{\pi}(y)$$

For the general cost function given by equation (4.40) above, we find that the optimality equation is given by

$$v_{\delta}(y) = \inf_{\tau > 0} \left\{ c(y, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_{\delta}(u) f_{\tau}(u, h | y) dh du \right\} \quad (4.41)$$

where c represents the discounted costs incurred at the next inspection, when in inspection interval τ is chosen. A derivation of this based on Ross (1970) is given in appendix A3. We note that this derivation is identical to the perfect inspection case, since we make all decisions on the basis of the covariate process Y . The only remaining point of note is that the hitting time h in the above equation is for the process X .

In the particular case we are dealing with, the results of section 4.4.1 show that the optimality equation is given by

$$v_{\delta}(y) = \inf_{\tau > 0} \left\{ c_{\delta}(y, \tau) + v_{\delta}(0)u_{\delta}(y, \tau) + v_{\delta}(\lambda)w_{\delta}(y, \tau) + \int_0^{s_0} v_{\delta}(u)K_{\tau}(u | y) du \right\} \quad (4.42)$$

where c , u , w and K are given by equations (4.32a-d) of section 4.4.1 (or equations (4.33a-d) in the case of a monotonic degradation process). It is clear from the proof of convergence in section 3.5.1, that this equation produces a unique solution, subject to conditions on the functions c , u , w and K described in section 3.5.1.

We shall not consider an example of non-periodic inspections in this chapter. The main reason, as we shall see, is that the computation of the required functions is extremely time consuming, making policy evaluation very slow. The periodic inspection examples given below utilise the same calculations, but require less function evaluations to determine the optimal policy

4.6 Models for Imperfect Inspection

A special case of the above analysis is that of imperfect inspection. In this case the covariate process is simply the observed level of degradation subject to error. The nature of the imperfection in the inspection process often allows us to derive a relationship between the observed and true processes, so that numerical results may be obtained.

The most common type of imperfect inspection in the degradation case is that the system degradation is observable, but with error. This is the case we are most interested in. The usual assumptions are that, if X_t is the degradation process, then we observe, at inspection times, the random variable

$$Y_t = X_t + \varepsilon_t \quad (4.43)$$

Where ε represents the error in making inspections. This structure is used by Whitmore (1995), who considers the case where X_t is a Wiener process and the errors ε are IID normally distributed with mean zero. In what follows we shall mainly consider this structure, under the assumption that X_t is a Lévy process. Before doing that however, we consider some alternative inspection models, and look at possible problems in their analysis.

Instead of considering additive models, in certain circumstances it may be plausible to consider using a multiplicative model of the form $Y_t = U_t \cdot X_t$. It is possible that for certain models this would be more realistic.

A third possible model could be that the inspection is imperfect in that we observe some discretised variable, which is based on the degradation process, for example this may take the form

$$Y_t = \begin{cases} 1 & X_t \in A_1 \\ 2 & X_t \in A_2 \\ \vdots & \\ n & X_t \in A_n \end{cases}$$

for given intervals $A_1, \dots, A_n$. In this case it is probably better to consider a discrete state space model, and assume that the discrete degradation model may be observed perfectly. In a sense, the cases we have considered in this and the previous chapter,

since the cost structure is dependent only on the structure of the sets A_i rather than the actual level of degradation.

In the case of a Lévy process with additive normal errors, we are basically dealing with a filtering problem. The observed process has a component of noise which we would like to remove to find the value of the underlying degradation process. Determining the relevant probability distributions is largely dependent on the nature of the underlying process and noise process, and the interdependencies between them. In this general setting we can, however, obtain results about the relationship between the underlying and observed degradation processes.

We will now give some properties of general Lévy process models in the presence of inspection errors which are additive. This structure is given by equation (4.43) above, and additionally we assume that the inspection error ε_i has mean μ_ε and variance σ_ε^2 .

As we have said, this model has been proposed by Whitmore(1995), when X is assumed to be a Wiener process and ε follows a Normal distribution with zero mean. In this article, Whitmore (1995) fully discusses the estimation and inference of system parameters, for the Wiener process case.

We shall assume for the moment that ε may have any type of statistical distribution. Based on these assumptions it is an elementary fact that

$$\text{Cov}(X_t, Y_t) = \text{Cov}(X_t, X_t + \varepsilon_t) = V(X_t) + \text{Cov}(X_t, \varepsilon_t)$$

Hence, the correlation between the true degradation X_t and the observation Y_t is given by

$$\begin{aligned} \rho(X_t, Y_t) &= \frac{\text{Cov}(X_t, Y_t)}{\sqrt{V(X_t)V(Y_t)}} \\ &= \frac{V(X_t) + \text{Cov}(X_t, \varepsilon_t)}{\sqrt{V(X_t)V(Y_t)}} \\ &= \frac{\text{sd}(X_t) + \rho(X_t, \varepsilon_t) \cdot \text{sd}(\varepsilon_t)}{\text{sd}(Y_t)} \end{aligned} \tag{4.44}$$

As we would expect, when the errors are perfectly positively correlated with the underlying degradation process, so to is the observation process Y . However, in the case when the errors and the degradation process are uncorrelated, the correlation between the actual and the observed degradation is just the ratio of standard deviations

of the process value and the observed value. In this case we may rewrite the equation in terms of the variances of the errors and the process, thus

$$\rho(X_t, Y_t) = \frac{sd(X_t)}{sd(Y_t)} = \left[\frac{V(X_t)}{V(X_t) + V(\varepsilon_t)} \right]^{\frac{1}{2}} = \left[1 + \frac{V(\varepsilon_t)}{V(X_t)} \right]^{-\frac{1}{2}} \quad (4.45)$$

which shows clearly that the correlation is always positive and decreases as the error variance increases. In addition, we note that the correlation is unaffected by the mean value of the error μ_ε . This is important because it implies that any consistent bias on the part of the inspector or monitoring system is irrelevant in terms of information provided about the degradation of the system. Of course, in the case where there is inconsistent bias this will not be the case.

To make further progress from this point we must make some assumptions regarding the nature of the degradation process and the errors. In the following section we consider the case of gamma process degradation.

4.7 Example : Imperfect Inspection of a Gamma Process

In this section we consider the case of gamma process degradation, with imperfect inspection. As we have said this represents a special case of the observation of covariate process. We use this special case since the structure of the inspection errors allows us to easily determine the stochastic relationship between the degradation process and covariate process (or observed degradation process).

4.7.1 Model Assumptions

In the following example we consider a system almost identical to that considered in the previous chapter. The system is deemed to have failed if the true level of degradation reaches a critical level c . Upon failure, the system is immediately replaced (corrective replacement). If, at an inspection, the 'observed' degradation is greater than the replacement level r , the system is preventively replaced. If the 'observed' level of degradation is found to be less than the replacement level r , the system continues

operating, and is inspected again according to an appropriate inspection schedule. We note that the decision as to whether to preventively replace or not is only dependent on the ‘observed’ degradation process, while the event of failure is only dependent on the true degradation process.

In addition the cost assumptions are identical:

- (i) Each Inspection incurs a cost C_I
- (ii) Each Preventive replacement incurs cost C_R , in addition to inspection costs
- (iii) Replacement on failure incurs cost C_F

The only difference between these assumptions and those of section 3.6.1 concerns the inspection process. Under these assumptions, it is possible to observe levels of degradation less than zero, and greater than the critical level c . It is now that our remarks about truncation of the degradation process (section 4.2) are applied. In this case, we assume that any observed level of degradation less than zero, is equivalent to a degradation level of zero, and assume that any observed level of degradation greater than c , is equivalent to degradation level c . The first assumption here is important, but the second is not, since when $r < c$, an observed level of degradation greater than c will result in immediate preventive replacement.

4.7.2 Distributional Results for the Gamma Process with error

We assume here that the inspection error is modelled by an additive normal random variable, with mean zero and variance ν^2 . We make the assumption that the Gamma process is defined as in Chapter 2, and that the errors are normally distributed with mean zero and variance ν^2 . Thus

$$\begin{aligned} X_{t_j} - X_{t_i} &\sim Ga(\alpha(t_j - t_i), \beta) \\ \varepsilon_t &\sim N(0, \nu^2) \quad \forall t \end{aligned} \tag{4.46}$$

where $\varepsilon_s \perp \varepsilon_t$ for $s \neq t$.

Let us consider the conditional distribution of the true level of degradation given the observation at that time. Let us initially consider the joint distribution of the observed and true values of the degradation process. Then it is elementary that

$$P(Y_* \leq u | X_* = z) = P(X_* + \varepsilon_* \leq u | X_* = z) = P(\varepsilon_* \leq u - z) = \Phi\left(\frac{u - z}{\nu}\right) \quad (4.47)$$

So that $Y_* | X_* \sim N(X_*, \nu^2)$ as we would expect. In addition, we are also interested in the distribution of the true level of degradation at a particular instant, given the observed degradation level. In this case more care must be taken, since our assumptions implicitly assume that the level of degradation is between level 0 and level c . Thus we have

$$\begin{aligned} f(x | y) dx &= \Pr(X_* = x | Y_* = y, 0 \leq X_* \leq c) dx \\ &= \frac{\Pr(X_* = x | Y_* = y)}{\Pr(0 \leq X_* \leq c | Y_* = y)} dx \\ &= \frac{\exp\left\{-\frac{1}{2\nu^2}(y - x)^2\right\}}{\nu\sqrt{2\pi}\left[\Phi\left(\frac{y}{\nu}\right) - \Phi\left(\frac{y - c}{\nu}\right)\right]} dx \end{aligned} \quad (4.48)$$

Where we have used the facts that $Y_* = X_* + \varepsilon_*$, and $\varepsilon_* \sim N(0, \nu^2)$.

We have by assumption that the density of $X_t | X_0 = x$ is given by

$$f(X_t = z | X_0 = x) = \frac{\beta^{\alpha t} (z - x)^{\alpha t - 1} e^{-\beta(z - x)}}{\Gamma(\alpha t)} \quad (4.49)$$

From these basic results we can compute the joint density of the future observed and true levels of degradation using

$$f_{Y_t, X_t | X_0}(u, z | x) = f_{Y_t | X_0, X_t}(u | x, z) f_{X_t | X_0}(z | x) = f_{Y_t | X_t}(u | z) f_{X_t | X_0}(z | x)$$

So that substituting the appropriate densities from equations (4.47) and (4.49), we obtain the joint density of the observed and true level of degradation, conditional on the true initial level of degradation,

$$f_\tau(u, z | x) = \frac{1}{\nu\sqrt{2\pi}} \exp\left\{-\frac{1}{2\nu^2}(u - z)^2\right\} \frac{\beta^{\alpha\tau} (z - x)^{\alpha\tau - 1} \exp\{-\beta(z - x)\}}{\Gamma(\alpha\tau)} \quad (4.50)$$

As in the example of chapter 3, we require the distribution of the hitting time of the critical set from an initial level of degradation x . This is defined in terms of a distribution function as

$$P(H_x^c < h) = \frac{\Gamma(\alpha h, \beta(c-x))}{\Gamma(\alpha h)} \quad (4.51)$$

which is derived in section 3.6.1, as equation (3.69).

4.7.3 Solution of Integral Equation

As an example, let us consider the optimisation of periodic inspections under a discounted total cost criterion. The case of non-periodic inspections applies the methods given here, but additionally requires a policy improvement algorithm to compute the optimal policy. We shall not consider this further here.

Given the information about the gamma process above, we can now proceed to solve the general integral equation given below

$$v_\delta(y, \tau) = c(y, \tau) + u(y, \tau)v_\delta(0, \tau) + w(y, \tau)v_\delta(\lambda, \tau) + \int_0^{s_0} v(u, \tau)K_\tau(u | y)du \quad (4.52)$$

Where c , u , w , and K are defined by equations (4.33a-d). From the assumptions given in section 4.7.1, the equations (4.33a-d) may be simplified, and we obtain the equations

$$c(y, \tau) = \int_0^c \left\{ e^{-\delta\tau} [C_I P^x(X_\tau < c) + C_R P^x(Y_\tau > r, X_\tau < c)] + C^F \int_0^\tau e^{-\delta h} g_c(h | x) dh \right\} f(x | y) dx \quad (4.53a)$$

$$u(y, \tau) = \int_0^c \left\{ e^{-\delta\tau} P^x(Y_\tau \notin (0, r), X_\tau < c) + \int_0^\tau e^{-\delta h} g_c(h | x) dh \right\} f(x | y) dx \quad (4.53b)$$

$$K_\tau(u | y) = \int_0^c \int_x^c e^{-\delta\tau} f_\tau(u, z | x) f(x | y) dz dx \quad (4.53c)$$

where the equations for u and w have been combined, since the initial level of the observed degradation (covariate) process is assumed to be equal to zero.

The expressions for $f_\tau(u, z | x)$, $f(x | y)$ and $G_c(h | x)$ have been computed above and are given by equations (4.50), (4.48) and (4.51) respectively. It now remains to compute the probabilities in equations (4.53a) and (4.53b). Firstly

$$\begin{aligned}
P^x(Y_\tau > r, X_\tau < c) &= \int_r^\infty \int_x^c f_\tau(u, z | x) dz dx \\
&= \int_x^c \int_r^\infty \frac{1}{\nu\sqrt{2\pi}} \exp\left\{-\frac{1}{2\nu^2}(u-z)^2\right\} \frac{\beta^{\alpha\tau}(z-x)^{\alpha\tau-1} e^{-\beta(z-x)}}{\Gamma(\alpha\tau)} du dz \\
&= \int_x^c \frac{\beta^{\alpha\tau}(z-x)^{\alpha\tau-1} e^{-\beta(z-x)}}{\Gamma(\alpha\tau)} \left\{1 - \Phi\left(\frac{r-z}{\nu}\right)\right\} dz \\
&= 1 - Q(\alpha\tau, \beta(c-x)) - \int_x^c \frac{\beta^{\alpha\tau}(z-x)^{\alpha\tau-1} e^{-\beta(z-x)}}{\Gamma(\alpha\tau)} \Phi\left(\frac{r-z}{\nu}\right) dz
\end{aligned} \tag{4.54}$$

since the gamma density in the integrand is independent of u . The function Q is the incomplete gamma ratio, defined by equation (3.70).

Similarly, the second probability of interest may be calculated as

$$\begin{aligned}
P^x(Y_\tau \notin (0, r), X_\tau < c) &= \\
1 - Q(\alpha\tau, \beta(c-x)) - \int_x^c \frac{\beta^{\alpha\tau}(z-x)^{\alpha\tau-1} e^{-\beta(z-x)}}{\Gamma(\alpha\tau)} \left\{\Phi\left(\frac{r-z}{\nu}\right) - \Phi\left(-\frac{z}{\nu}\right)\right\} dz
\end{aligned} \tag{4.55}$$

The integral in both of the functions c and u may be calculated as in chapter 3.

Equation (3.71) which gives

$$\int_0^\tau e^{-\delta h} g(h | x) dh = e^{-\delta\tau} G(\tau | x) + \delta \int_0^\tau e^{-\delta h} G(h | x) dh \tag{4.56}$$

with the distribution function of the hitting time being given by equation (4.51) above.

Substituting these functions into equations (4.53a-c) we arrive at the functions

$$c^x(y, \tau) = e^{-\delta\tau} C_I + e^{-\delta\tau} C_R \left\{1 - \int_x^c \frac{\beta^{\alpha\tau}(z-x)^{\alpha\tau-1} e^{-\beta(z-x)}}{\Gamma(\alpha\tau)} \Phi\left(\frac{r-z}{\nu}\right) dz\right\} \tag{4.57a}$$

$$+ e^{-\delta\tau} (C_F - C_R - C_I) Q(\alpha\tau, \beta(c-x)) + \delta C^F \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh$$

$$\begin{aligned}
u^x(y, \tau) &= e^{-\delta\tau} \left\{1 - \int_x^c \frac{\beta^{\alpha\tau}(z-x)^{\alpha\tau-1} e^{-\beta(z-x)}}{\Gamma(\alpha\tau)} \left\{\Phi\left(\frac{r-z}{\nu}\right) - \Phi\left(-\frac{z}{\nu}\right)\right\} dz\right\} \\
&\quad + \int_0^\tau e^{-\delta h} Q(\alpha h, \beta(c-x)) dh
\end{aligned} \tag{4.57b}$$

where the functions c and u are now given by

$$c(y, \tau) = \int_0^c c^x(y, \tau) f(x | y) dx \quad u(y, \tau) = \int_0^c u^x(y, \tau) f(x | y) dx$$

and the function K is computed directly as

$$K_\tau(u | y) = \frac{e^{-\delta\tau} \beta^{\alpha\tau}}{2\pi\nu^2 \Gamma(\alpha\tau) \left\{ \Phi\left(\frac{y}{\nu}\right) - \Phi\left(\frac{y-c}{\nu}\right) \right\}} \times \int_0^c \int_x^c \exp\left\{ -\frac{1}{2\nu^2} [(u-z)^2 + (y-x)^2] - \beta(z-x) \right\} (z-x)^{\alpha\tau-1} dz dx \quad (4.57c)$$

The functions in this form have been simplified as much as possible. All of the integrals above must be computed numerically. Clearly, computation of these functions requires us to calculate a number of double integrals, which is extremely computationally demanding. This means that the time required to compute optimal inspection policies in the case of imperfect inspection (and more generally in the case of a covariate process) is much longer than in the case of perfect inspection policies.

As an example, applying the Nystrom method of appendix A1, requires us to compute an $n \times n$ matrix with each entry being an evaluation of the function $K_\tau(u | y)$. If we apply a trapezium rule to the double integral in expression (4.57c), with n points, each function evaluation requires n^2 evaluations of the integrand. Thus to compute the whole matrix requires n^4 function evaluations. This compares to n^2 function evaluations required for the perfect inspection case.

4.7.4 Numerical Results

The following tables give the results found in the case of a gamma process degradation model. Table 4.1 shows the optimal periodic inspection policy under the discounted total cost criterion. In addition the table shows what percentage of this optimal cost is given by the cost of the optimal periodic policy with perfect inspection. The comparison is not direct, since one would have to assume that perfect inspections are more expensive than imperfect inspections.

The table is generated by solving the equation (4.52) with functions given by (4.57a-c). The equation is solved by the Nystrom method of appendix A1, and it is assumed that $n = 5$, so that the integral in equation (4.52) is subdivided into 5 intervals. Clearly, this

is not ideal, and greater accuracy could be obtained by increasing the number of subdivisions. The reason for using such a small number of sub-divisions is simply one of time. As we described above (section 3.5.3) the number of computations required in the case of imperfect inspection is much greater than under perfect inspection, and for illustrative purposes we feel it is adequate to reduce the accuracy of the calculation.

In chapter 3, 10 subdivisions were applied, and these figures (some of which are given in Table 4.2) are used for comparison purposes. Clearly, there is some error in these figures, but for comparison purposes we believe they adequately show the properties of the model. In particular, the starred entry in Table 4.1 (below) would appear to show that the cost of imperfect inspections is (slightly) less than the cost of perfect inspections at the same cost. Clearly this aberration is caused by computational error.

In terms of the results from the model, the effect of variation in the costs and process variability is the same as that found in chapter 3 for the case of perfect inspection. In general, increasing costs and increasing process variability tend to reduce the optimal inspection interval. In addition, it is clear from Table 4.1 that for systems with large corrective and preventive replacement costs, the proportionate increase in optimal DTC caused by having imperfect inspection is greater than for systems with relatively smaller replacement costs. (Table 4.2 shows optimal policy for a perfectly observed system)

The imperfect inspection has an unexpected effect on the optimal inspection policy. While we might expect that imperfect inspection would reduce the inspection intervals, it is found in many cases that the inspection interval is greater than in the case of perfect inspections. The reason for this is probably that it is important never to observe the system in a state slightly less than the replacement limit. If this occurs, the system will not be inspected again for a long period and so failure is certain to occur. Therefore, to increase the inspection interval reduces the probability that the observed degradation is less than the replacement limit at inspection. While the probability of failure is increased by increasing the inspection interval, it is likely that this increased cost is more than compensated for by the reduction in the number of cases in which degradation is observed to be less than the replacement limit.

		DTC and Optimal Inspection Policy for $\nu = 0.1$					
		Low Variance ($a=20$)			High Variance ($a=5$)		
		$C_F=10$	30	100	$C_F=10$	30	100
Low Rep. Limit $r = 0.3$	$C_R=5$	0.9 89739 93%	0.6 121007 89%	0.45 147578 89%	1.35 90116 97%	0.45 159846 90%	0.25 535216 39%
	20		1.0 286869 95%	0.55 417751 89%		∞ 271485 99%	0.4 535216 88%
	50			0.75 814338 90%			1.0 868044 95%
High Rep. Limit $r = 0.9$	$C_R=5$	1.15 97137 97%	0.25 247882 89%	0.1 567560 60%	∞ 90495 >100%*	0.35 254184 92%	0.15 701353 78%
	20		∞ 291298 98%	0.15 724764 62%		∞ 271485 99%	0.3 793832 78%
	50			0.25 894616 73%			0.4 888093 86%

Table 4.1 – Optimal inspection policy, corresponding discounted total cost and percentage of optimal DTC with perfect inspection, for a system with $\nu = 0.1$

The table above (Table 4.1) shows the case when the inspection error ε has standard deviation $\nu = 0.1$. Clearly, with our state space defined over the interval $[0,1]$, this represents quite a large inspection error. This value was chosen, so that the effects of the inspection error could be easily seen. As the standard deviation of the inspection error is reduced, the effects described above are less pronounced, and tend to the perfect inspection case as the standard deviation approaches zero (subject to computational error). The effects on the optimal policy for very large errors, depends largely on the replacement level r . For small r , the inspection error has much less effect, and the optimal policy tends towards that for the perfect inspection case. When the replacement level is close to the failure level c , the large inspection error means that inspections have little effect in preventing failures, and so the optimal policy tends to be one of carrying out no inspections, and simply replacing on failure.

		DTC and Optimal Inspection Policy for $\nu = 0$					
		Low Variance ($a=20$)			High Variance ($a=5$)		
		$C_F=10$	30	100	$C_F=10$	30	100
Low Rep. Limit $r = 0.3$	$C_R=5$	0.9 83457	0.65 107462	0.5 132488	1.2 87538	0.4 144035	0.2 207333
	20		0.95 273206	0.65 372577		1.55 270735	0.35 474144
	50			0.8 740515			0.95 826155
High Rep. Limit $r = 0.9$	$C_R=5$	1.1 94860	0.1 220381	0.05 343653	1.65 90561	0.2 235899	0.1 548898
	20		1.4 284641	0.05 448500		1.35 271028	0.1 621216
	50			0.05 658193			0.1 765853

Table 4.2 – Optimal inspection policy and corresponding discounted total cost for a system with perfect inspection (given by Table 3.2)

It is clear from the results given here, that the model provides a useful and reasonable policy for the inspection of systems subject to error (and more generally the inspection of covariate processes). The results are largely what would have been expected, the only problem with the model being the computational complexity, which means that the time taken to compute the optimal policy is very large.

4.8 Conclusion

In this chapter we have looked at the general problem of obtaining optimal markovian inspection policies for systems whose degradation is unobservable, but which have an observable, associated covariate process. We have assumed the system is modelled by a Lévy process degradation model, with threshold failure, and a Lévy covariate process. Using techniques from dynamic programming theory, integral and functional equations have been derived, allowing the optimal periodic and non-periodic inspection policies to

be obtained. As a special case, imperfect inspection of a Lévy process degradation model is considered.

As in the previous chapter, the main example used is that of the gamma process, with additive normal inspection error. The results provide sensible and realistic inspection policies, and give insight into the system behaviour, and into the effect of inspection error on the optimal inspection policy of a system. It is found however, that obtaining optimal inspection policies with inspection error is extremely computationally demanding, due to the large number of numerical multiple integrals which must be evaluated. This may provide a barrier in extending this model further.

Generally, the extensions which could be considered are the same as those described in chapter 3. The main extension of this model which would be interesting to consider is that of Bayesian methods. In this model, we have considered a system with a known degradation process, but which cannot be observed. Clearly this is somewhat paradoxical: if the system cannot be observed, how can we know what the underlying degradation process is. If we assume that the underlying degradation process is of known form, but with unknown parameters, it may be possible to apply Bayesian techniques, so that the sequential estimation of the underlying degradation process may be based upon observation of the covariate process. Whether or not this can be done depends on the underlying form of the degradation process and covariate process. Future work will be carried out in this area.

There are few models in the literature which explicitly deal with optimisation of inspections for systems in the presence of covariates, and as such this model provides a useful addition in the field of degradation modelling.

Chapter 5

Optimal Maintenance for Deteriorating Systems

5.1 Introduction

In the following chapter, we extend the results of the two previous chapters to the case of maintenance. This is a more complex problem than those considered in the previous two chapters, since we now assume that the system may be repaired, and so the level of degradation may be changed by the system operator.

We consider firstly the case in which the optimal inspection interval is found for a system with known maintenance policy. The basis of the model is that of the previous two chapters, but instead of replacement (or perfect repair) we assume that in each state, the system may be repaired, reducing the level of degradation by a deterministic or random amount.

Following this we look at the optimisation of maintenance policies for a system with a given inspection schedule. This case is quite different to the optimal inspection cases previously considered, but the methods applied are similar. We assume that the maintenance action may have a deterministic or random effect on the system state.

Finally we consider jointly optimal periodic inspection and maintenance policies, for various cost structures.

As in chapters 3 and 4, both periodic and non-periodic inspections are considered, but examples are confined to cases of perfect inspections for computational reasons. The methods applied in this chapter mirror closely those in the previous chapters. In each case integral equations or Dynamic programming equations are found for the cost of the system, as a function of the state of the system. Using the standard methods given previously these may be solved numerically to obtain an approximate numerical solution. Since all of the equations are solved by the same method, we restrict examples to the case of deterministic maintenance, which is less computationally demanding to solve.

5.2 Underlying Inspection and Repair Model

5.2.1 System Model

Bringing maintenance and repair into our model means that some of our assumptions must be changed. As we have noted in chapter 1, there are many different types of maintenance and inspection policies that we could consider. In this section, we describe the model that we shall consider. This model has been chosen since we feel that it most adequately describes the general situation we are considering, while remaining computationally feasible.

In general, we shall consider a system that is very similar to that found in chapter 3, with some minor changes. It is clear that the assumptions of chapter 3 are a special case of these assumptions. We consider two models, one incorporating deterministic maintenance, and another allowing for general (random) maintenance.

MODEL ASSUMPTIONS M3 (Deterministic Maintenance):

1. We assume the state-space of the system is S_d , which is partitioned into a set B and its complement. The system is deemed to have failed when the degradation process X hits the critical set $B = (c, \infty)$.
2. Each inspection reveals the true state of the system, and the state can only be determined by carrying out an inspection.
3. Each inspection incurs a cost C_θ . At an inspection, a repair is attempted, and its cost is given by the function $C(x, y)$, where x is the system state before maintenance, and y is the system state after maintenance.
4. The cost of replacing a failed system is C^F . On failure, it is assumed the system is immediately replaced by a new system, identical to the original.

These assumptions are largely unchanged from those of chapter 3, except that we now do not specify a partition of the state space. This subdivision is now encapsulated by the maintenance policy, and corresponding cost function, which are discussed in the next subsection.

In the model described above, it is assumed that the maintenance has a known effect on the state of the system. It is perhaps more realistic to assume that there is some uncertainty about the effect of maintenance on the level of degradation. In the following set of assumptions, we assume that the uncertainty in the effect of maintenance is modelled by the random variable Θ .

MODEL ASSUMPTIONS M4 (Random Maintenance):

1. We assume the state-space of the system is S_d , which is partitioned into a set B and its complement. The system is deemed to have failed when the degradation process X hits the critical set $B = (c, \infty)$.
2. Each inspection reveals the true state of the system, and the state can only be determined by carrying out an inspection.

3. Each inspection incurs a cost C_0 . At an inspection a repair is attempted, the cost of which is given by the function $C_d(x)$, where x is the system state before maintenance, and d represents the maintenance action used. We assume that the system state after maintenance is given by a function $x \rightarrow d(x, \theta)$, where θ represents a particular realisation of the random vector Θ . (See section 5.2.2)
4. The cost of replacing a failed system is C^F . On failure, it is assumed the system is immediately replaced by a new system, identical to the original.

Based on these assumptions, in the cases of random and deterministic maintenance, the two main problems we would like to consider are

- (a) Computation of optimal inspection intervals for fixed maintenance actions
- (b) Computation of optimal maintenance actions for a fixed inspection schedule

The problem of optimal inspection of a maintained system is similar to those considered in chapters 3 and 4. The main difference being that the transitions to new levels of degradation must now incorporate maintenance. As we shall see, this has very little effect on the derivation of the dynamic programming equations, but does make their numerical solution more difficult. The second case of optimal maintenance is slightly different to the foregoing analysis. Integral equations are used to determine the costs of given strategies but new dynamic programming equations are derived to compute the optimal maintenance strategy.

In the case of optimal inspections, as before, we simply assume a fixed maintenance strategy, and proceed in developing integral equations as before. The difference in model assumptions, however, means that the integral equations are slightly different.

5.2.2 Models for Inspection and Maintenance Actions

Firstly, let us briefly look at possible inspection strategies. As we have said in chapter 1, the main distinctions are between perfect and imperfect inspection, and between periodic and non-periodic inspection. In this chapter, we shall mainly consider the case of perfect inspection, both periodic and non-periodic. The main reason for this is that

they are the computationally easier to solve than the corresponding imperfect inspection case. Secondly, as we have seen in chapter 4, the inspection error must become quite large before the imperfect inspection has an effect on the optimal policy.

Let us now look at some possible maintenance actions. The simplest policy from a computational point of view is that of perfect complete repair. Under this policy, maintenance is equivalent to a replacement, since the maintenance effectively results in the system becoming good as new. It is clear therefore that this policy may be dealt with using exactly the methods of the previous two chapters, where we redefine replacement as complete repair.

Slightly more complex is the case of perfect incomplete repair. By this we mean that repairs are carried out, they affect the degradation of the system in a known and fixed way. An example of such a policy would be a maintenance action that reduces the degradation of the system by a fixed percentage. The costs of such deterministic maintenance can be determined by a simple adaptation to the methods given in chapter three. It is this model which we shall use for deterministic maintenance.

A further modification is to allow imperfect or general repair, in which the level of repair is random. In this case, the repairs carried out have a random effect on the degradation of the system. This randomness may be caused by errors on the part of those carrying out the maintenance, or may be simply inherent in the system maintenance. The case of general repair is more difficult to solve, since the transitions now involve an added degree of randomness.

In the derivations that follow, we shall make a general assumption regarding maintenance. We assume that, if the state of the system at inspection is y , then the state after maintenance is given by a maintenance function $d(y, \theta)$. θ represents a particular realisation of a random vector Θ that in some way encapsulates the randomness found in the maintenance. We assume throughout that, conditional on the system's state y , the density of the parameter θ is given by $f(\theta | y)$. This approach gives a great deal of flexibility, for example when maintenance does not occur at a given value of y , we set $d(y, \theta) = y$. Likewise, if in a certain region replacement is to be carried out, we can set $d(y, \theta) = 0$. It is clear that these cases are the extremes of y , and we assume that

$0 \leq d(y, \theta) \leq y$. In the case of deterministic maintenance, the maintenance function is defined simply as $d(y)$.

For the problem of determining an optimal inspection schedule we assume that the maintenance function d is known in advance, for all values of y . Our task then is to determine the cost of maintenance for each possible inspection interval, and hence determine an optimal inspection schedule.

The problem of computing an optimal maintenance strategy for a given inspection schedule is more complex. To determine the optimal maintenance strategy we must make some assumptions about the maintenance function d . The easiest case is that of deterministic maintenance. In this case we must determine for each x , the maintenance function $d(x)$, which can be achieved using dynamic programming.

In the case of general maintenance, we must be more specific. In the case of general maintenance, we must specify all of the maintenance functions, and the distribution of all random parameters. There may be n possible maintenance actions $d_i(y, \theta)$ for $i = 1, \dots, n$ and corresponding probability distributions for Θ , $f_i(\theta | y)$. A more straightforward assumption is that of a single maintenance action $d(y, \theta)$, with a number of possible distributions of Θ , $f_i(\theta | y)$. In this case, we seek the optimal choice of probability distribution for Θ , in each state. We shall restrict our attention to the second case.

In the following sections we consider the problems described above and obtain optimal maintenance and inspection strategies.

5.3 Optimal Inspection in the Presence of Maintenance

As in chapters 3 and 4, we consider optimal perfect inspection in the periodic and non-periodic cases. We shall consider only the total discounted cost criterion, as the examples of chapter 3 show that the results in the average cost case are almost identical. The main application of the results of this section will come in sections 5.4 and 5.5 below. A particular case for which the results of this section are directly applicable

would be the case in which there is only a single possible maintenance action in each state. For other cases, it is usually preferable to consider the joint optimisation of inspection and maintenance.

We note that in this chapter it is important to distinguish between the state of the system before and after repair has taken place. For this reason we define t^- to represent a time just before t , and t^+ to represents a time just after t .

5.3.1 Optimal Inspection under Deterministic Maintenance

We begin by considering the case of optimal inspection for a system subject to a fixed and known maintenance strategy, as described by assumptions M3 above. In this section, we assume that the inspection policy is periodic. In section 5.3.3 below, we consider the case of non-periodic inspections by direct extension of the results given here.

As in chapter 3 we define $V_\delta(x, \tau)$ to be the discounted total cost for a system which has current level of degradation x . We assume that the system as just been inspected and has level of degradation x , and an appropriate maintenance action will be instantaneously taken, changing the state of the system to $d(x)$. It is possible to assume that the argument of the function V is the state of the system immediately after maintenance. In this case, however, it is more difficult to obtain results. In particular this formulation makes it difficult to obtain optimal maintenance policies using standard methods.

So, if we define t_i to be the inter-event times (an event being either an inspection / repair or failure), we define the expected discounted total cost as

$$v_\delta(x, \tau) = E(V_\delta(x, \tau) | X_0 = x) = E\left(\sum_{n=0}^{\infty} e^{-\delta(t_1 + \dots + t_n)} C(X_{t_n^-}, X_{t_{n+1}^-}; \tau) | X_0 = x\right) \quad (5.1)$$

where $C(x, y; \tau)$ represents the (random) cost incurred if the system is in state y at time t_{n+1}^- when it was in state x at time t_n^+ , and the inspection interval is τ . It is assumed that the maintenance policy is fixed, and so we do not show explicit dependencies, except where necessary to avoid confusion.

Again we assume that X_t represents the degradation process, and M_t represents the maximum value of that process over the period $[0, t]$. The cost function may be expressed recursively. Thus

$$\begin{aligned} V_\delta(x, \tau) &| X_\tau, M_\tau, H_{d(x)}^B \\ &= C(x, d(x)) + V_\delta(X_\tau^-, \tau) e^{-\delta\tau} 1_{\{M_\tau \notin B, X_\tau \in B\}} + \{V_\delta(0, \tau) + C^F\} e^{-\delta H_{d(x)}^B} 1_{\{H_{d(x)}^B < \tau\}} \end{aligned} \quad (5.2)$$

This equation closely resembles that given in chapter 3 for perfect inspection under a replacement policy. The main difference is that we now incorporate the cost of maintenance, which is incurred immediately. In addition, since the state of the system changes to $d(x)$, all probability functions must be adjusted to take this into account. In particular the failure time is now given by the hitting time of B from the point $d(x)$.

Taking expectations of the above expression, and applying the reasoning as in chapter 3 gives the integral equation

$$\begin{aligned} v_\delta(x, \tau) &= C(x, d(x)) + \int_B \int_{S \setminus B} v_\delta(y, \tau) e^{-\delta\tau} f_\tau(y, m | d(x)) dm dy \\ &\quad + (v_\delta(0, \tau) + C^F) \int_0^\tau e^{-\delta h} g_B(h | d(x)) dh \end{aligned} \quad (5.3)$$

In this section, we define $K_\tau(y | d(x)) = \int_{S \setminus B} e^{-\delta\tau} f_\tau(y, m | d(x)) dm$, which is slightly different to the definition of previous chapters. As before $f_\tau(y, m | x)$ and $g_B(h | x)$ are the densities of $X_\tau, M_\tau | X_0 = x$, and the hitting time H of the set B by the process X , given $X_0 = x$, respectively.

We now assume that the function $d(y)$ is known for all y , and in particular, that $d(y) = 0$ for $y \leq 0$, and that $C(y, d(y)) = 0$ for $y \leq 0$. As in chapter 3, we do not want to allow negative degradation levels, treating these as aberrations of the model rather than a description of the physical reality. We thus assume $v(y, \tau) = 0 \forall y < 0$.

We can thus rewrite the integrals in the above expression as

$$\int_B v_\delta(y, \tau) K_\tau(y | d(x)) dy =$$

$$v_\delta(0, \tau) e^{-\delta \tau} P^{d(x)}(X_\tau < 0, M_\tau > c) + \int_0^c v_\delta(y, \tau) K_\tau(y | d(x)) dy$$

where the superscript $d(x)$ now indicates that the probability is conditional upon $X_0 = d(x)$. Then the integral equation may be written as

$$v_\delta(x, \tau) = c(x, \tau) + v_\delta(0, \tau) u(x, \tau) + \int_0^c v_\delta(y, \tau) K_\tau(y | d(x)) dy \quad (5.4)$$

Where the functions in this case are defined as

$$c(x, \tau) = C(x, d(x)) + C^F \int_0^\tau e^{-\delta h} g_B(h | d(x)) dh \quad (5.5a)$$

$$u(x, \tau) = e^{-\delta \tau} P^{d(x)}(X_\tau < 0, M_\tau > c) + \int_0^\tau e^{-\delta h} g_B(h | d(x)) dh \quad (5.5b)$$

$$K_\tau(y | d(x)) = \int_{S \setminus B} e^{-\delta \tau} f_\tau(y, m | d(x)) dm \quad (5.5c)$$

As in chapter 3, we expect that for a suitable degradation model the first term of the function u should be close to zero. Otherwise, the model suggests a large probability of the degradation becoming negative.

In the case where the degradation process X is monotonic, the Maximum variable need not be considered and we obtain the functions c , u and K as

$$c(x, \tau) = C(x, d(x)) + C^F \int_0^\tau e^{-\delta h} g_B(h | d(x)) dh \quad (5.6a)$$

$$u(x, \tau) = \int_0^\tau e^{-\delta h} g_B(h | d(x)) dh \quad (5.6b)$$

$$K_\tau(y | d(x)) = e^{-\delta \tau} f_\tau(y | d(x)) \quad (5.6c)$$

It is obvious from the form of these equations that they can be solved by the methods given in appendix 1. In this case however, there are some added complications, which we now consider.

The numerical solution of the integral equations in chapters 3 and 4 is based on a numerical integration procedure constructed over a suitable discretisation of the region of integration. In the case of a non-singular kernel K , the results apply almost directly. Applying a quadrature rule to the integral equation gives, for each $y_j \in M$

$$v(y_j, \tau) = c_d(y_j, \tau) + u_d(y_j, \tau) v(0, \tau) + \sum_{i=0}^n v(y_i, \tau) K_\tau(y_i | d(y_j)) w_{i,j} \quad (5.7)$$

As in appendix 1, by writing this equation replacing x by values in the mesh, $y_i = ih$, $i = 0, 1, \dots, N$ where $h = r/N$, we can write this as a system of linear equations in matrix form as

$$\mathbf{v} = \mathbf{c} + \mathbf{U}\mathbf{v} + \mathbf{K}_d^T \mathbf{v} \quad (5.8)$$

Where the matrices are defined as follows :

$$\mathbf{v} = \begin{bmatrix} v(y_0) \\ v(y_1) \\ \vdots \\ v(y_m) \end{bmatrix} \quad \mathbf{c} = \begin{bmatrix} c(y_0) \\ c(y_1) \\ \vdots \\ c(y_m) \end{bmatrix}$$

$$\mathbf{K}_d^T = \begin{bmatrix} K(y_0, d(y_0))w_{0,0} & K(y_0, d(y_1))w_{0,1} & \cdots & K(y_0, d(y_m))w_{0,m} \\ K(y_1, d(y_0))w_{1,0} & K(y_1, d(y_1))w_{1,1} & \cdots & K(y_1, d(y_m))w_{1,m} \\ \vdots & \vdots & \ddots & \vdots \\ K(y_m, d(y_0))w_{m,0} & K(y_m, d(y_1))w_{m,1} & \cdots & K(y_m, d(y_m))w_{m,m} \end{bmatrix}$$

$$\mathbf{U} = \begin{bmatrix} u(y_0) & 0 & \cdots & 0 \\ u(y_1) & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ u(y_m) & 0 & \cdots & 0 \end{bmatrix}$$

Clearly, the matrices are defined mainly as in appendix 1, which considers the solution of this integral equation in the case when $d(y) = y$. It is clear that in this case, the above matrices reduce to those of Appendix 1/Chapter 3.

So that the cost vector $\mathbf{v}$ may be found as

$$\mathbf{v} = (\mathbf{I} - \mathbf{U} - \mathbf{K}_d^T)^{-1} \mathbf{c} \quad (5.9)$$

if the inverse exists.

In the case of a diagonally singular kernel, the approach given in appendix 1 may be applied, and a similar result to that given above is found. However, there are now some added complexities. In this case the general integral equation (5.7) is rewritten, to remove the singularity. Thus

$$v_\delta(x, \tau) = c(x, \tau) + v_\delta(0, \tau)u(x, \tau) + \int_0^c [v_\delta(y, \tau) - v_\delta(d(x), \tau)] K_\tau(y | d(x)) dy + v_\delta(d(x), \tau) \int_0^c K_\tau(y | d(x)) dy \quad (5.10)$$

For convenience of notation, we define

$$q(x) = \int_0^c K_\tau(y | d(x)) dy. \quad (5.11)$$

Discretising this equation, we obtain for values of x on the mesh $M = \{y_i | i = 0, \dots, n\}$

$$v_\delta(y_j, \tau) = c_d(y_j, \tau) + v_\delta(0, \tau)u_d(y_j, \tau) + \sum_{i=0}^n v(y_i, \tau) K_\tau(y_i | d(y_j)) w_{i,j} - v(d(y_j)) \sum_{i=0}^n K_\tau(y_i | d(y_j)) w_{i,j} + v(d(y_j)) q_d(y_j) \quad (5.12)$$

it can be seen that this system of equations may be written in matrix form as

$$\mathbf{v} = \mathbf{c} + \mathbf{U}\mathbf{v} + \mathbf{K}_d^T \mathbf{v} + \mathbf{Q}\mathbf{v}_d - \mathbf{W}\mathbf{v}_d \quad (5.13)$$

where the matrices $\mathbf{c}$, $\mathbf{U}$, $\mathbf{K}_d$ and $\mathbf{v}$ are defined as above. The matrix $\mathbf{Q}$ and vector $\mathbf{v}_d$ are defined by

$$\mathbf{Q} = \begin{bmatrix} q_d(y_0) & 0 & \dots & 0 \\ 0 & q_d(y_1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & q_d(y_n) \end{bmatrix} \quad \mathbf{v}_d = \begin{bmatrix} v(d(0)) \\ v(d(1)) \\ \vdots \\ v(d(n)) \end{bmatrix} \quad (5.14)$$

The matrix $\mathbf{W}$ is defined as $\mathbf{W} = \text{diag}(\mathbf{w})$, where the vector $\mathbf{w}$ is given by $\mathbf{w} = \mathbf{K}_d^T \mathbf{e}$ where $\mathbf{e}$ is an $(n+1) \times 1$ vector containing 1 in each entry.

It is clear then, that to solve these equations we must determine a relationship between the vector $\mathbf{v}$ and the vector $\mathbf{v}_d$. The simplest method is to assume that $y \in M \Rightarrow d(y) \in M$. Hence, for any point on our mesh the maintenance action d results in a state that is also in the mesh M . Now, define a matrix $\mathbf{M}_d$ such that

$[\mathbf{M}_d]_{i,j} = 1$ if $d(y_i) = y_j$ and 0 otherwise. Then it is clear that $\mathbf{v}_d = \mathbf{M}_d \mathbf{v}$. Substituting this into the above matrix equation and rearranging gives

$$\mathbf{v} = (\mathbf{I} + \mathbf{W}\mathbf{M}_d - \mathbf{U} - \mathbf{Q}\mathbf{M}_d - \mathbf{K}_d^T)^{-1} \mathbf{c} \quad (5.15)$$

when the inverse exists.

We note that the assumption that the maintenance action results in a point on the mesh is not restrictive. In particular, as the mesh becomes finer the effect of this should be negligible. We note also that this assumption is not required in the case when the Kernel is non-singular. In that case, the model is able to allow the maintenance to change the state of the process in a completely general way.

Then, as in chapter 3, for a given maintenance policy we may now compute the discounted total cost of operating any inspection policy. It is therefore simple to determine the optimal inspection policy. An example of this is given in section 5.7, using the gamma process as the model of system degradation.

5.3.2 Optimal Inspection of Randomly Maintained Systems

We now briefly consider the case described by assumptions M4, given in section 5.2.2 above. Again, we restrict attention to periodic inspection policies, leaving the case of non-periodic policies to section 5.3.3 below.

The main difference between the analysis of this section and that preceding it is that the decision-maker does not know with certainty how the system will be affected by the maintenance action. In terms of derivation of appropriate integral equations, this causes no problems. Problems, however, do arise when considering the numerical solution of these equations. These are discussed following the derivations below.

In the main, we use the notation of section 5.3.1, but we extend the notation to deal with random maintenance. We assume, since the maintenance policy is fixed, that the maintenance function is given by $d(y, \theta)$, where θ has probability density function $f(\theta | y)$ for each $y \in [0, c]$.

Then, proceeding as before, the cost function V may be expressed recursively as

$$\begin{aligned} V_\delta(x, \tau) &| X_\tau, M_\tau, H_{d(x, \Theta)}^B, \Theta \\ &= C_d(x) + V(X_\tau^-, \tau) e^{-\delta \tau} 1_{\{M_\tau \notin B, X_\tau \in B\}} + \left\{ V_\delta(0, \tau) + C^F \right\} e^{-\delta H_{d(x, \Theta)}^B} 1_{\{H_{d(x, \Theta)}^B < \tau\}} \end{aligned} \quad (5.16)$$

which is the same as in the deterministic maintenance case, except for the presence of θ in the maintenance function. Taking expectations as before gives

$$\begin{aligned} v_\delta(x, \tau) &= C_d(y) + \int \int \int v_\delta(y, \tau) e^{-\delta \tau} f(\theta | y) f_\tau(y, m | d(x, \theta)) d\theta dm dy \\ &\quad + \left(v_\delta(0, \tau) + C^F \right) \int_0^\tau \int_\Theta e^{-\delta h} g_B(h | d(x, \theta)) d\theta dh \end{aligned} \quad (5.17)$$

Again, we assume that $d(y, \theta) = 0$ for $y \leq 0$, and that $C_d(y) = 0$ for $y \leq 0$. So that the integral equation becomes

$$v_\delta(x, \tau) = c(x, \tau) + u(x, \tau) v_\delta(0, \tau) + \int_0^c v_\delta(y, \tau) K_\tau(y | x) dy \quad (5.18)$$

with

$$K_\tau(y | x) = \int \int_{S \setminus B \times \Theta} e^{-\delta \tau} f(\theta | x) f_\tau(y, m | d(x, \theta)) d\theta dm \quad (5.19a)$$

$$c(x, \tau) = C_d(x) + C^F \int_0^\tau \int_\Theta e^{-\delta h} g_B(h | d(x, \theta)) f(\theta | x) d\theta dh \quad (5.19b)$$

$$u(x, \tau) = \int_{-\infty}^0 K_\tau(y | x) dy + \int_0^\tau e^{-\delta h} g_B(h | d(x, \theta)) f(\theta | x) dh \quad (5.19c)$$

From these equations, it is clear that the randomness in the maintenance policy is simply averaged out. Essentially then, the case of randomness can be approximated by simply considering the mean effect of maintenance rather than taking into account the effect of the parameter θ . We have not investigated how this simplification affects the results of the models of this chapter. It is natural to expect, however, that the error caused by the simplification will become larger, as the variability of θ gets larger.

As usual, in the case of a monotone degradation process, we can simplify the equation a little. In this case the integral equation is in the same form as equation (5.17) and the functions are given by

$$K_\tau(y|x) = \int_{\Theta} e^{-\delta\tau} f(\theta|x) f_\tau(y|d(x,\theta)) d\theta \quad (5.20a)$$

$$c(x,\tau) = C_d(x) + C^F \int_0^\tau \int_{\Theta} e^{-\delta h} g_B(h|d(x,\theta)) f(\theta|x) d\theta dh \quad (5.20b)$$

$$u(x,\tau) = \int_0^\tau e^{-\delta h} g_B(h|d(x,\theta)) f(\theta|x) dh \quad (5.20c)$$

where $f_\tau(y|x)$ is the density of $X_\tau | X_0 = x$. This equation is of the same form as that given above. In this case, the kernel still involves an integral, and so the usual computational advantages of the monotone process are not present. This is the same effect as was observed in the case of optimal inspection of a covariate process. It would appear from what we have done that any attempt to add an extra degree of randomness into the problem results in a problem which is much more computationally demanding to solve. For these reasons, we shall not consider a numerical example of random maintenance, but we give an outline of the method of solution below.

It is clear from the form of the equation that the method used in the deterministic maintenance case will apply.

5.3.3 Optimal Non-Periodic inspections

It is clear from the derivation of the integral equations above, and the development of dynamic programming equation in Appendix 3, that the results of the previous two sections on deterministic and general maintenance policies may be immediately generalised to the non-periodic case.

Firstly, in the case of deterministic maintenance, it is clear that the underlying equation is almost identical to that shown in section 3.5. Clearly, then we may apply the same policy improvement algorithm as given in that section. Then, from the above results it is clear that the dynamic programming equation is of the form

$$v_{\delta}(x) = \inf_{\tau > 0} \left\{ c(x, \tau) + v_{\delta}(0)u(x, \tau) + \int_0^c v_{\delta}(y) K_{\tau}(y | d(x)) dy \right\} \quad (5.21)$$

with the functions c , u and K being defined appropriately by equations (5.5) or (5.6) depending upon whether the process is monotonic or not.

Likewise, in the case of general (random) maintenance, the dynamic programming equation is given by

$$v_{\delta}(x) = \inf_{\tau > 0} \left\{ c(x, \tau) + v_{\delta}(0)u(x, \tau) + \int_0^c v_{\delta}(y) K_{\tau}(y | x) dy \right\} \quad (5.22)$$

where the functions c , u and K are defined by equations (5.19) or (5.20), depending on whether or not the underlying degradation process is monotonic.

It is shown in Appendix 3 how these equations may be derived. The proof given in Chapter 3 showing that the dynamic programming equation of section 3.5 defines contraction mappings, applies equally well here.

In both cases, assuming efficient numerical methods are available to compute the costs of any given policy, and a standard policy improvement algorithm may be used to compute the optimal inspection strategy.

To apply the policy improvement algorithm, we must evaluate the costs of a specific policy. Thus, in equations (5.4) and (5.18) we define $v(x) = v(x, \tau(x))$, where $\tau(x)$ defines a state dependent inspection policy. These equations may be thus written

$$v(x) = c(x) + v(0)u(x) + \int_0^c v(y) K_{\tau(x)}(y | d(x)) dy \quad (5.23)$$

and

$$v(x) = c(x) + u(x)v(0) + \int_0^c v(y) K_{\tau(x)}(y | x) dy \quad (5.24)$$

To allow concise description in all possible cases, we assume now that the discretised integral equation, whichever one we are using, is written in operator notation $\mathbf{v} = T(\mathbf{v})$.

In either case, let us assume that the integral equation has been discretised, using a mesh of m uniformly spaced points over the interval $[0, c]$. Following the development

of sections 5.3.1 and 5.3.2, the solution may now be expressed in the form $\mathbf{v} = \mathbf{M}^{-1}\mathbf{c}$ for some matrix $\mathbf{M}$ and vector $\mathbf{c}$. The vector $\mathbf{v}$, of course, gives the approximate solution at points on the mesh

$$\mathbf{v} = \begin{bmatrix} v(y_0) \\ v(y_1) \\ \vdots \\ v(y_n) \end{bmatrix}$$

The particular form of the matrix $\mathbf{M}$ depends upon whether or not we are dealing with random or deterministic repairs, and on whether or not the Kernel of the integral equation is singular. In any case, we assume that it is invertible. We have not investigated conditions for the invertibility of these matrices, but in all cases considered we have yet to find a singular matrix.

Then, we may apply the following policy iteration algorithm, which is similar to that given in chapter 3, (The algorithm is based on that given by Puterman, 1994), and is identical to that used in chapter 3:

1. Set $k = 0$, and select an arbitrary inspection rule $\pi_0 = \{ \tau_0^0, \dots, \tau_n^0 \}$
2. (Policy Evaluation) Obtain $\mathbf{v}^k$ by solving $\mathbf{v}^k = \mathbf{M}_k^{-1}\mathbf{c}_k$, where the subscript k indicates the matrix is to be evaluated with policy π_k .
3. (Policy Improvement) Choose $\pi_{k+1} = \{ \tau_0^{k+1}, \dots, \tau_n^{k+1} \}$ to satisfy

$$\pi_{k+1} = \arg \min_{\pi \in \Pi} \{ T_\pi(\mathbf{v}^k) \}$$

4. If $\pi_{k+1} = \pi_k$, stop and set $\pi^* = \pi_k$. Otherwise, increment k by 1 and return to step 2.

Where we emphasise that the minimisation in step 3. is carried out component-wise. We may of course use ‘min’ rather than ‘inf’, since we are dealing with a discretised system, having a finite state space and action space. Again, the effect of the discretisation is small for a fine mesh.

As with most of the work in this thesis, the solutions require large amounts of numerical work, and whether or not an efficient solution can be found depends largely on the underlying process and assumptions made. In particular it is found that the case of general maintenance is computationally demanding. It is doubtful whether or not the increased effort required to solve the general problem is worthwhile. While the general model is perhaps more realistic, the deterministic model has the advantage of being computationally easier and giving more informative policies.

5.4 Optimal Maintenance for a given Inspection Policy

In this section, we shall look at a problem closely related to that of section 5.3. In section 5.3 we considered optimisation of the inspection interval, for a given maintenance policy. We shall now consider the problem of determining which maintenance actions are optimal, for a given inspection policy. In this section, we will only consider the case of periodic inspection, since the results are easily extended to an arbitrary inspection policy. In the latter case, the inspection interval τ is replaced by the function $\tau(\cdot)$ wherever it appears.

This analysis of this section is similar to that of Stadje and Zuckerman (1991), who consider a model in which the virtual age of the system may be reduced by maintenance. The other main difference between what follows and the work of Stadje and Zuckerman, is that Stadje and Zuckerman define the failure mechanism in terms of a hazard function, whereas we consider a threshold failure model, for reasons already outlined. A model related to this is given by Dagpunar (1998), who develops integral equations for the various maintenance policies, in which the virtual age of a system is reduced by maintenance actions.

While our model differs by not considering hazard based failures, we feel that it compensates by allowing us to model an observable effect, namely the effect of maintenance on degradation, as opposed to the unobservable effect of the effect of maintenance on the ‘age’ of a system. The models of course are different, and are appropriate for different types of system.

Our model is clearly focused on systems whose degradation can be affected by a maintenance action. We look at a general case in which the degradation can be reduced by any amount chosen by the system user. Of course, in reality the maintenance may be more restricted than this. These restrictions simplify the model below, and are easily accommodated.

5.4.1 Optimal Deterministic Maintenance

Let us assume that the inspection policy for our system is known and fixed. While an inappropriate assumption for many systems, it is often the case that there are a limited number of times available at which inspection, and subsequent repair may take place. A classical example of this is that of commercial aircraft, which may only be inspected at the end of each flight. In addition, an airline may only have facilities to carry out certain maintenance actions in certain places, so that it is only possible for it to do maintenance after a number of flights.

We assume then, that the system has a fixed inspection schedule, which, without loss of generality, we assume is periodic. We shall briefly consider the case of joint optimisation of periodic inspections and maintenance in section 5.4.3. In that case, the results of this section are combined with a simple search algorithm to obtain the optimal inspection and maintenance policy.

Thus, for a given maintenance function d and inspection interval τ , we can compute the discounted total cost of this policy as described in section 5.3.1 above. It is of course the solution of the integral equation.

$$v_\delta(x, \tau) = c(x, \tau) + v_\delta(0, \tau)u(x, \tau) + \int_0^c v_\delta(y, \tau) K_\tau(y | d(x)) dy \quad (5.25)$$

with the functions c , u and K defined as in equations (5.5) or (5.6), whichever is appropriate for the situation at hand. As we have seen, this equation can usually be solved by simple numerical methods.

From the nature and derivation of this equation, it is obvious that we may obtain a dynamic programming equation of the form

$$v_\delta(x) = \inf_{0 \leq d \leq x} \left\{ c(x, d) + v_\delta(0)u(x, d) + \int_0^\varepsilon v_\delta(y) K_\tau(y | d) dy \right\} \quad (5.26)$$

This functional equation can be solved by similar methods to those described above for optimal non-periodic inspections. Since τ is assumed known, we drop it from our notation in this section.

As usual, we begin by discretising the integral equation to form a set of linear equations. Let us assume that a mesh $M = \{y_i | i = 0, \dots, n\}$ is defined for the discretisation. Then we assume that $y \in M \Rightarrow d(y) \in M$, so that the maintenance action always results in the system being on a state on the mesh M . Thus, in the dynamic programming equation above we must assume that $d \in M$. Thus the problem is reduced to a discrete action space problem, and we can solve the discretised problem in the usual way.

The policy improvement algorithm is defined as above, with slight amendments.

1. Set $k = 0$, and select an arbitrary inspection rule $\pi_0 = \{d_0^0, \dots, d_n^0\}$
2. (Policy Evaluation) Obtain $\mathbf{v}^k$ by solving $\mathbf{v}^k = \mathbf{M}_k^{-1} \mathbf{c}_k$, where the subscript k indicates the matrix is to be evaluated with policy π_k .
3. (Policy Improvement) Choose $\pi_{k+1} = \{d_0^{k+1}, \dots, d_n^{k+1}\}$ to satisfy

$$\pi_{k+1} = \arg \min_{\pi \in \Pi} \{T_\pi(\mathbf{v}^k)\}$$

4. If $\pi_{k+1} = \pi_k$, stop and set $\pi^* = \pi_k$. Otherwise increment k by 1 and return to 2..

This algorithm may be applied as in chapter 3.

The policy found using this algorithm gives the state to which the degradation should be reduced. Thus, it provides us with a full description of the appropriate maintenance action in any given state.

It may also be the case that we are restricted to a small number of available maintenance actions, similar to the situation we have described in the random maintenance case. This restriction can easily be accommodated by simply restricting the available values of d in each state. These values, however, must always be members of the set M . Clearly, as the number of points in the mesh increases, this restriction becomes less important.

5.4.2 Optimal General Maintenance

In the case of general maintenance under perfect inspection, we assume that there is a single maintenance function $d(x, \theta)$, but that the random variable θ may come from a number of distributions, at the choosing of the decision-maker. We assume therefore that there is a set of probability density functions $D = \{ f_1(\theta | \cdot), f_2(\theta | \cdot), \dots, f_k(\theta | \cdot) \}$, and that the aim of our optimisation problem is to decide which of these probability functions should be used in each state of the process.

In particular we have in mind the maintenance function

$$d(x, \theta) = x - \theta$$

In this case, we assume that θ represents the reduction in the degradation. Then, two different distributions of θ correspond to two possible maintenance actions, both of which reduce maintenance, but by a random amount, with different probability distributions. There are of course many other possible forms for the maintenance function, but this one seems the most reasonable.

As in the assumptions M4 at the beginning of this chapter, we assume that the cost of applying a particular maintenance action depends only on the state of the system prior to maintenance and the maintenance action. In this section, this implies that the cost of a maintenance action depends only on the state x and the distribution chosen for θ .

Thus, the policy is now described by $\pi = \{ a_0, \dots, a_n \}$, where $f_{a_i}(\theta | x_i)$ is the distribution chosen in state x_i . Then it is clear that the dynamic programming equation is of the form

$$v_\delta(x) = \inf_{a \in A} \left\{ c(x, \tau) + v_\delta(0)u(x, \tau) + \int_0^c v_\delta(y) K_\tau^a(y | x) dy \right\} \quad (5.27)$$

where in this case the kernel is defined by

$$K_\tau^a(y | x) = \int_0^x e^{-\delta\tau} f_a(\theta | x) f_\tau(y | x - \theta) d\theta \quad (5.28)$$

Therefore, the policy improvement routine described in section 5.4.1 may be applied directly, replacing the policy $\pi_k = \{d_0^k, \dots, d_n^k\}$ with $\pi_k = \{a_0^k, \dots, a_n^k\}$.

As in the deterministic case, the policy gives a complete description of the optimal maintenance action in each state. In this case, it is likely that the number of available maintenance actions is small, and so the policy improvement is not computationally demanding. As in the case of inspection of covariates, it is the policy evaluation that proves to be computationally difficult.

5.4.3 Joint Optimisation of Maintenance and Inspections

It is possible, by a simple extension to the above results to consider optimal periodic inspection and maintenance. Simply, we may find for each τ the optimal maintenance policy, and corresponding discounted total cost. By searching for the minimum optimal maintenance costs over all values of τ , it is simple to obtain the value of τ and corresponding maintenance policy that has the least cost. This approach is taken in section 5.6 below, using the Gamma degradation model as an example.

In order to consider the case of non-periodic inspection of a maintained system, we need to extend the dynamic programming formulation of the problem. As before, we can compute the cost of maintenance and inspection under any given policy. The integral equation can thus be written explicitly in terms of the policies applied. Thus

$$v_\delta(x) = c(x, \tau(x), d(x)) + v_\delta(0)u(x, \tau(x), d(x)) + \int_0^c v_\delta(y) K_{\tau(x)}(y | d(x)) dy \quad (5.29)$$

assuming we are dealing with deterministic maintenance. In the case of random maintenance, we amend the equation as in the above case. We can now apply the

standard policy improvement algorithm described above. This can be done by redefining the policy to be applied as $\pi_k = \{ (d_0^k, \dots, d_n^k), (\tau_0^k, \dots, \tau_m^k) \}$, where n is the number of points in the mesh and m is the number of time points considered.

It is clear from previous results that the optimal policy will exist, at least for the discretised problem. The problem with this method is that there are a large number of policies that have to be evaluated during the policy improvement. In this case we have to compute the cost under the $m \times n$ different combinations of d and τ . Clearly, for large values of m and n this will prove time consuming.

A possible approach to reducing the complexity of this problem is to separately evaluate and improve the maintenance and inspection policies. We assume that the discretised integral equation under policy $\pi_k = (d^k, t^k)$ is written in operator notation as $\mathbf{v} = T_{k,k}(\mathbf{v})$. Then the policy iteration algorithm is given by

1. Set $k = 0$. Choose arbitrary policies given by $\mathbf{d}^0 = (d_0^0, \dots, d_n^0)$ and $\mathbf{t}^0 = (\tau_0^0, \dots, \tau_m^0)$.
2. (Policy Evaluation I) Compute the discounted cost under policy $(\mathbf{d}^k, \mathbf{t}^k)$ by solving $\mathbf{v}^{k,k} = T_{k,k}(\mathbf{v}^{k,k})$
3. (Policy Improvement I) Choose $\mathbf{d}^{k+1}$ component-wise such that

$$\mathbf{d}^{k+1} = \arg \min_{d \in D} \{ T(\mathbf{v}^{k,k}) \}$$
4. (Policy Evaluation II) Compute the discounted cost under policy $(\mathbf{d}^{k+1}, \mathbf{t}^k)$ by solving $\mathbf{v}^{k+1,k} = T_{k+1,k}(\mathbf{v}^{k+1,k})$
5. (Policy Improvement II) Choose $\mathbf{t}^{k+1}$ component-wise such that

$$\mathbf{t}^{k+1} = \arg \min_{\tau > 0} \{ T(\mathbf{v}^{k+1,k}) \}$$
6. If $(\mathbf{d}^{k+1}, \mathbf{t}^{k+1}) = (\mathbf{d}^k, \mathbf{t}^k)$ stop and set $\pi^* = (\mathbf{d}^k, \mathbf{t}^k)$. Otherwise, return to step 2, incrementing k by 1.

This algorithm differs only slightly from the standard algorithm given above. The main difference here is that much fewer evaluations are required in the policy improvement section. The minimisation now requires only $m + n$ evaluations, which for large m and n reduces the computational effort required. On the other hand, this algorithm now requires two policy evaluations in each cycle, and it may take more cycles to converge, since not every combination of inspection and maintenance actions is considered in each cycle. Which of the two algorithms is preferred depends largely on how computationally demanding the policy evaluation is. In most cases, however, we would expect that this algorithm is more efficient.

In the example that follows, we consider only the case of optimal periodic inspection and maintenance. The case of non-periodic inspection is more difficult and adds little to what we have already said.

5.5 Example : Gamma Process Degradation

In this section, we shall consider the case of a gamma process degradation model. We shall mainly be concerned with periodic and non-periodic inspections of systems with deterministic maintenance. The case of random maintenance may be solved by the methods described in the chapter, but these cases are much more computationally demanding than those we shall now look at.

We assume therefore the maintenance is deterministic, so that the maintenance function is defined by $d(x)$. Where necessary, we also make the assumption that $y \in M \Rightarrow d(y) \in M$ where $M = \{y_i \mid i = 0, \dots, n\}$ is our chosen discretisation of the state space.

As in previous chapters, we assume that the degradation process is a gamma process, so that the transition density of the process is given by

$$f_\tau(y \mid x) = \frac{\beta^{\alpha\tau} (y - x)^{\alpha\tau-1} e^{-\beta(y-x)}}{\Gamma(\alpha\tau)} \quad y > x$$

as we have seen this density function is singular along the line $y = x$, so that we apply the results for a diagonally singular kernel for a monotone process. The properties of this process as a degradation model are discussed in chapter 2.

Unlike the case in the previous two chapters, the cost function in the presence of maintenance is now very flexible. The function $C(x, d(x))$ represents the cost of a maintenance action that results in the degradation of the system being reduced from level x to level $d(x)$, including inspection costs. Clearly, it should be chosen to reflect the costs incurred by the system to which the model is being applied. For our example, we shall consider the cost function given by

$$C(x, d(x)) = C_0 + C_1(x - d(x)) + C_2(x - d(x))^2$$

Clearly, this is not ideal for most systems. In particular it is assumed that the cost of reduction in maintenance by a given amount is independent of the initial level of degradation. This is not the case for many systems. We emphasise however that this cost function is completely general and may take any form.

In all the following examples we assume that the discount rate $\delta = 0.01$. In addition, as in previous chapters we assume that the average rate of degradation is 1 per unit time so that $\alpha/\beta = 1$. We vary the parameter α to reflect different levels of variability of the degradation process. Similarly we assume that the failure threshold $c = 1$, in all the examples. These assumptions do not affect the results below, since other values can be achieved by either a scale change or time change or a combination of the two.

5.5.1 Optimal Inspection of a deterministically maintained system

In this section we shall consider the results of section 5.3.1 and 5.3.3, and give examples of optimal inspection policies for systems having a fixed and known deterministic maintenance policy.

There are many possible maintenance strategies depending upon the type of system we are considering. In the examples of this section, we assume that the maintenance function d takes the following form:

$$d(x) = \begin{cases} x & \text{for } x < r_0 \\ x - r_0 & \text{for } r_0 < x < r_1 \\ 0 & \text{for } r_1 < x < c \end{cases}$$

The effect of this policy is simple. We define two control limits: if the first is breached a partial repair takes place, while if the second is breached the system is repaired to a good as new condition. This policy is suggested simply as an example, to illustrate the effect of the maintenance policy on the inspection schedule.

We begin by looking at the case of periodic inspections. We arbitrarily fix the value of c to be $c=1$ and assume the cost function is given by

$$C(x, d(x)) = 1 + 5(x - d(x)) + 15(x - d(x))^2$$

with cost of failure is given by $C_F = 20$. Tables 5.1 and 5.2 below show the effect of variations in the maintenance policy on the optimal periodic inspection interval, in the cases of a high and low degradation process variability. In table 5.1 we assume that $\alpha = \beta = 7$, and in table 5.2 it is assumed that $\alpha = \beta = 25$. These result in process variance of 0.14 and 0.04 respectively.

		r_1		
		0.3	0.6	0.9
r_0	0.2	4 1632.13	2 1639.03	2 1548.38
	0.5		8 1905.82	4 1823.08
	0.8			∞ ~1950

Table 5.1 – Table showing effect of maintenance schedule with $\alpha = \beta = 7$, on discounted total cost and optimal inspection interval ($\times 10$)

These tables are computed using the numerical methods previously described. We assume that the state space $[0,1]$ is divided into 10 steps, of length 0.1 and the

inspection intervals considered are in multiples of 0.1 time units. It is clear that this is a crude approximation to the true integral equation, but even this rough approximation is sufficient to show the form of the optimal policy and the effect of various parameters.

		r_1		
		0.3	0.6	0.9
r_0	0.2	4 1457.46	2 1444.43	2 1396.49
	0.5		7 1856.23	3 1714.25
	0.8			∞ ~2050

Table 5.2 - Table showing effect of maintenance schedule with $\alpha = \beta = 25$, on discounted total cost and optimal inspection interval ($\times 10$)

It can be seen from these tables that model gives results which are intuitively reasonable. In both cases a similar pattern is observed, with any differences accounted for by the increased process variability. In the first case, with high variability, costs are on the whole higher, with inspection intervals being shorter. The opposite case is found when both values of r_1 and r_2 become large. In that case, the process with lower variability is more expensive than that with high variability.

The main feature of both tables is that the maintenance costs increase as the value of r_0 becomes larger. This behaviour is caused by the cost function chosen. The fact we have a convex cost function means that it is more optimal to have small r_0 because this results in maintenance actions with lower cost. We would expect to find this behaviour with any convex cost function. Other than this the results are largely as we would expect, and exhibit many of the features described in chapter 3. The behaviour here is, of course, more difficult to predict, as the maintenance complicates matters. It is of course very difficult to generalise these results, since they are heavily dependent on the form of the cost function involved. In the case we have considered with a quadratic

cost function, it is clear that systems with high values of r_0 will cost more to run, as the relative cost of maintenance will be more. This effect is clearly visible in the results. In both of the above examples, the cases in which $r_0 = 0.2$ and $r_1 = 0.9$ give the least cost. Again, this is due to the form of the cost function, which implies small reductions in degradation cost relatively less than large amount of reduction in degradation.

In the same cases as above, we can also obtain optimal non-periodic inspection schedules. The relationship between the two forms of policy is similar to that found in chapter 3, which is of course a special case. Again we assume that r_0 and r_1 may take values as in the above tables and consider the cases $\alpha = \beta = 7$ and $\alpha = \beta = 25$. The results are shown in table 5.3 overleaf.

As in the above case, the results for different values of process variability are largely similar. The only difference of note comes in the case when $r_0 = 0.8$ and $r_1 = 0.9$. In this case, with high process variability the optimal policy is to never inspect the system, and simply replace on failure. In the lower variability case however, this is only the case in certain states, the other states having a defined finite inspection policy. Looking at the corresponding periodic policy shows that the costs are almost identical. The approximation we have used in this example is very rough, and in this case, it would be necessary to use a finer subdivision of the state space in order to determine which policy should be used. We would expect to find a policy of no inspections would be optimal.

Within each group, the policies seem, at first sight, very strange. In the case of the replacement model of chapter 3, it was found that the inspection interval was a decreasing function of the system state, as we would intuitively expect. In this case, however this does not appear to happen. If we take into account that the state will be changed by the application of the maintenance action, the policies are very reasonable, in terms of what we would intuitively expect.

		Inspection interval in each state (state before maintenance)										DTC	Per. Policy	Per. DTC
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9			
$\alpha=7$	$r_0=0.2$ $r_1=0.3$	0.4	0.3	0.2	0.4	0.4	0.4	0.4	0.4	0.4	0.4	1601	0.4	1632
	$r_0=0.2$ $r_1=0.6$	0.4	0.3	0.2	0.3	0.2	0.1	0.4	0.4	0.4	0.4	1574	0.2	1639
	$r_0=0.2$ $r_1=0.9$	0.4	0.3	0.2	0.3	0.2	0.1	0.1	0.1	0.1	0.4	1432	0.2	1548
	$r_0=0.5$ $r_1=0.6$	0.7	0.6	0.5	0.4	0.3	0.2	0.7	0.7	0.7	0.7	1861	0.8	1905
	$r_0=0.5$ $r_1=0.9$	0.6	0.5	0.4	0.4	0.3	0.2	0.5	0.4	0.4	0.6	1757	0.4	1823
	$r_0=0.8$ $r_1=0.9$	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	~1950	∞	~1950
$\alpha=25$	$r_0=0.2$ $r_1=0.3$	0.4	0.3	0.2	0.4	0.4	0.4	0.4	0.4	0.4	0.4	1431	0.4	1457
	$r_0=0.2$ $r_1=0.6$	0.4	0.3	0.2	0.3	0.2	0.1	0.4	0.4	0.4	0.4	1358	0.2	1444
	$r_0=0.2$ $r_1=0.9$	0.5	0.4	0.2	0.4	0.2	0.1	0.1	0.1	0.1	0.5	1296	0.2	1396
	$r_0=0.5$ $r_1=0.6$	0.6	0.5	0.4	0.3	0.2	0.1	0.6	0.6	0.6	0.6	1745	0.7	1856
	$r_0=0.5$ $r_1=0.9$	0.6	0.5	0.4	0.3	0.3	0.2	0.5	0.4	0.3	0.6	1578	0.3	1714
	$r_0=0.8$ $r_1=0.9$	∞	∞	0.8	0.7	0.6	0.5	0.3	0.2	0.1	∞	~2049	∞	~2049

Table 5.3 – Optimal Non-periodic Inspection schedule for system subject to maintenance. Table shows optimal inspection policy and corresponding discounted total cost, alongside optimal periodic policy and cost.

Looking at all of the cases (with the exception of the final one), the following pattern seems to apply. We see that for $0 \leq x \leq r_0$, the policy is indeed a decreasing function, and is in line with what we might expect. In the case when $r_0 \leq x \leq r_1$, the policy is ‘restarted’, since the state of the system is reduced by r_0 if the system is found to be in state x . In the final case when $x > r_1$, the system state is reduced to zero and so, the policy is identical to that which would apply if we had in fact observed the state of the system to be zero.

With regard to the cost of the non-periodic policies, it is found in most cases that the cost of the non-periodic policy is less than the corresponding optimal periodic policy. The only exception being when a policy of no-inspections is optimal.

In summary, the model provides inspection policies for given maintenance policies that are both intuitively sensible and informative. In this case, it is more difficult to be precise about the properties of the model, as each case is highly dependent on the particular cost function used. As always, it would be very interesting to see the application of the model to a real system, where physical comparisons would be available to test the model.

5.5.2 Optimal Periodic Inspection and Maintenance Policies

In this section, we give an example of the computation of jointly optimal inspection and maintenance policies, given by section 5.4.3. We assume that the inspection is periodic and use a simple search algorithm to compute the least cost policy. We do not consider optimal random maintenance in this example. The methods used in the example may be easily applied to that case. The functions in that case are more difficult, and thus it takes much longer to compute the optimal policies. For similar reasons, we do not consider the joint optimisation of maintenance and non-periodic inspections. This case can be solved using the policy improvement algorithm given in section 5.4.3 above.

We make similar assumptions to those given in the section 5.5.1 above. In this case however, the maintenance policy is freely determined by the optimisation process. We do not, therefore, have to specify a maintenance function, and so the values of r_0 and r_1

are redundant. We now consider the effect of changes in the cost function parameters, but again assume a fixed form

$$C(x, d(x)) = 1 + C_1(x - d(x)) + C_2(x - d(x))^2$$

and assume that C_F may take different values. Once again, we consider both high and low process variability. Apply the methods described above, gives the results shown in Table 5.4 overleaf.

The table shows the maintenance action in each state, alongside the optimal periodic inspection policy, and corresponding discounted total cost. As before, we consider inspection intervals in multiples of 0.1 time units, and subdivide the state space $[0,1]$ into 10 intervals of length 0.1. We note that the ‘maintenance action’ in each state gives the ‘level to which degradation should be reduced to’. Therefore, a value of 0 implies a complete repair (or replacement) should be carried out, whereas a value of x in state x implies no action should be taken. Once again, the results are heavily dependent on the cost function, and it is difficult to generalise comments about the properties of the model results.

As in previous cases, the cost function chosen means that a small reduction in degradation is relatively cheaper than a large reduction. The effect of this can be seen in the results, which show that the optimal maintenance strategy is to reduce the degradation in each state by a small amount, rather than to have a complete repair. Correspondingly, the inspection intervals in each case are relatively small, so that, in general, the optimal policy seems to be to carry out small maintenance actions quite often, rather than have longer inspection intervals. Deciding which of these strategies is most appropriate is very important for many systems.

For given levels of α and C_F , it is clear that the maintenance functions behave as expected. For lower repair costs, it is optimal to apply a complete repair in most states, only when the level of degradation gets close to 1, is it more cost effective to apply partial maintenance. When costs of repair become high relative to costs of failure, it is still optimal to carry out low levels of maintenance, but again this depends largely on the cost function.

		Optimal Maintenance Action in each state											τ^*	DTC
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
$\alpha=7$ $C_F=20$	$C_1=5$ $C_2=5$	0	0	0	0	0	0	0	0.1	0.1	0.1	0.2	0.4	1061
	$C_1=5$ $C_2=15$	0	0	0	0	0.1	0.2	0.2	0.3	0.4	0.4	0.5	0.3	1376
	$C_1=5$ $C_2=30$	0	0	0.1	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.9	0.2	1671
$\alpha=7$ $C_F=40$	$C_1=5$ $C_2=5$	0	0	0	0	0	0	0.1	0.1	0.1	0.1	0.2	0.3	1130
	$C_1=5$ $C_2=15$	0	0	0	0.1	0.1	0.2	0.2	0.3	0.3	0.4	0.4	0.2	1432
	$C_1=5$ $C_2=30$	0	0	0	0.1	0.2	0.3	0.3	0.4	0.4	0.5	0.6	0.2	1778
$\alpha=25$ $C_F=20$	$C_1=5$ $C_2=5$	0	0	0	0	0	0	0	0.1	0.1	0.1	0.2	0.4	960
	$C_1=5$ $C_2=15$	0	0	0	0	0.1	0.2	0.2	0.3	0.4	0.4	0.5	0.3	1299
	$C_1=5$ $C_2=30$	0	0	0	0.1	0.1	0.2	0.3	0.4	0.4	0.5	0.9	0.2	1651
$\alpha=25$ $C_F=40$	$C_1=5$ $C_2=5$	0	0	0	0	0	0.1	0.1	0.1	0.2	0.2	0.2	0.4	965
	$C_1=5$ $C_2=15$	0	0	0	0	0.1	0.2	0.2	0.3	0.3	0.4	0.4	0.3	1302
	$C_1=5$ $C_2=30$	0	0	0	0.1	0.1	0.2	0.3	0.4	0.4	0.5	0.5	0.2	1651

Table 5.4 – Jointly Optimal Maintenance and Periodic Inspection schedule. Table shows optimal policy and corresponding discounted total cost

In the case of a linear cost function, the results are sensitive to the relative cost of repair and failure. In general, if repair cost per unit reduction in degradation is less than corresponding failure cost, employ a complete repair at each inspection. Otherwise, do nothing and replace the system on failure. Clearly, this is an extreme case, and unrealistic in most situations.

For high and low process variability, the optimal policies are very similar, the costs being slightly higher in the high variability case. Likewise, the cases of low and high cost of failure are very similar. This is particularly the case for low process variability, in which, the policies employed and costs are almost identical. The reason is clear: with low variability, chances of failure are very low, so the actual cost of failure has little effect on the optimal policy. For high process variability, the optimal policies are almost identical, but the case of high failure cost results in a higher discounted total cost, as we would expect.

As before, the maintenance policies found are intuitively reasonable, and fit the system model well. It is clear that the model, as in previous cases, provides useful and informative information about the nature of the system, which can be effectively used by decision-makers. While we have not considered variations in all the parameters, the model responds positively to those we have considered, and is not overly sensitive to changes in any particular parameter. As before, the crucial assumption lies in the cost function. Any change to the form of the cost function may result in completely different policies.

5.6 Conclusions

In this chapter, we have considered optimal maintenance and inspection policies for systems whose degradation is modelled by a Lévy process. We assume that the system may have a completely general maintenance and inspection policy, and that the system failure is modelled by a threshold-crossing model.

Using the methods of previous chapters, we may derive equations for the optimal discounted total cost in each case, allowing us to compute optimal maintenance

policies. In particular we consider the cases of periodic and non-periodic inspections, coupled with both deterministic and random maintenance.

The example of the gamma process degradation model is again used, and the model gives realistic and informative maintenance policies in this setting.

There are many possible extensions to this model that could be considered. Of these, the most obvious would be to consider the case of imperfect inspection, or inspection of covariates. This case is more difficult, since a separate model must be specified for the effect of maintenance on the observed degradation or covariate process. Clearly, if the true level of degradation cannot be observed with certainty, it is extremely difficult to model the effect of maintenance.

Another extension would be to allow failure to occur in any state, rather than employ the restrictive assumptions of threshold failure models. This could easily be accommodated for a system with constant killing rate, but would become more difficult if a degradation dependent hazard rate were to be introduced.

It would be very useful to apply the methods of this chapter to a real system, to see how the policy compares to that used in reality. This is of course the ultimate test of any model. In particular it would be interesting to consider whether intrinsic reliability of a system may be substituted by maintenance. For example, can a system which has low reliability, in conjunction with an appropriate maintenance policy, be more cost effective than a more advanced, higher reliability system.

Clearly, maintenance and inspection decisions are difficult. There is so much information available, and it is difficult to combine this in a coherent way, to achieve an optimal result. As we have seen in the examples, the maintenance policies themselves are not always obvious. We believe this model provides a basis in which relatively complex systems may be analysed to determine which form of maintenance is better. Clearly, there is an increasing need for models that can incorporate all the features of such complex systems, and while this model does not accomplish this, it provides a useful starting point for further models in this direction.

Chapter 6

Summary and Conclusions

6.1 Summary and Conclusions

In this thesis we have been largely concerned with the optimal inspection and maintenance of systems whose degradation may be directly or indirectly observed, and whose failure is a direct result of such degradation. To obtain optimal policies we make simplifying assumptions, most important of which are:

1. System degradation is modelled by a Lévy process
2. The system fails when the degradation of the system reaches a critical level

Subject to these assumptions, we have seen that this model provides a flexible way of looking at many problems in maintenance optimisation. Clearly, these assumptions are not appropriate for all systems, but some degree of simplification is necessary for progress to be made. We have focused on three important cases.

Firstly, we considered the case of optimal perfect inspection, for a given replacement policy. This corresponds to the standard case in the literature in which we seek an inspection policy that gives the inspection interval in terms of the observed level of degradation. Clearly, subject to the above assumptions, this type of policy is applicable to many types of system.

In some cases, however, the degradation of the system, while still being the cause of failure, cannot be observed. In this case it may be appropriate to use a related covariate process as the basis for an inspection policy. The case in which the covariate process is modelled by a time homogenous Markov process is given in chapter 4.

Thirdly, we look at the important case of maintenance optimisation. In many cases it is not appropriate to use an inspection/replacement policy, and instead a maintenance action may be undertaken, to reduce the level of degradation of the system. In chapter 5 we considered a general maintenance model, which applies to both deterministic and random maintenance.

The gamma process model is used throughout. In chapter 3, optimal perfect inspection policies are found, and these are extended to optimal imperfect inspection policies in chapter 4, when the degradation is observed subject to a Gaussian measurement error. The results of these examples are described in the appropriate chapter. Mainly, the results are as would be expected. In summary, for periodic inspection policies

1. The optimal inspection interval decreases, as the cost of failure increases over the cost of replacement.
2. Increasing degradation process variability results in increased inspection intervals for relatively low cost systems, but results in decreased inspection intervals for relatively higher cost systems.
3. The replacement limit chosen by the decision-maker is extremely important. Generally, an optimal value may be found minimising the overall cost of inspection. Extreme values of the replacement limit result in higher overall costs, showing that a structured inspection/replacement policy is better than an age replacement policy.

In the case of non-periodic inspection policies, similar results are found. In most cases, the optimal general policy is not periodic, and in cases of high cost systems, considerable reduction in costs can be obtained by using a non-periodic policy.

Extending these results to imperfect inspection, we find the same general pattern. The main effect here is that of observation error. In this case, it is found that the optimal inspection interval tends to be greater than in the case of perfect inspection, depending on the relative costs of replacement and failure. This result seems somewhat paradoxical but can be explained by looking at the effect of the error on the system

(Section 4.7.4). This result emphasises the importance of appropriate maintenance and inspection models.

In chapter 5 we consider the case of optimal maintenance. As before the results are largely as one might expect, but are more difficult to generalise since the cost structure is more complex. Based on a quadratic cost function, jointly optimal periodic inspection and maintenance policies were obtained. The main feature of this model was that the costs of maintaining a system with large degradation process variability were less than the costs for a system with small degradation process variability. Again, this is somewhat paradoxical.

In all of the models considered we have encountered the problem of computational complexity. All of the models require solution of integral equations, which in itself is computationally demanding. Additionally, in the case of imperfect inspection and general maintenance, the added degree of randomness results in an integral equation which is specified in terms of functions, which themselves are given by multiple integrals. This means that, for relatively simply defined systems, solutions take some time to obtain. Clearly, as computers become faster, this will become less of a problem. However, this problem must be borne in mind when considering extension of this model to more complex cases.

Bearing this in mind however, the methodology used in this thesis can be applied as a general model. The Lévy process structure implies that the future degradation of a system is independent of previous levels of degradation, thus allowing the costs of maintenance to be expressed recursively. Clearly, the picture is not so clear if any of our basic assumptions is relaxed. It is this possibility which we now consider.

6.2 Possibilities for Further Research

To make progress in the previous chapters we have made many simplifying assumptions. It is clear that some of these assumptions may be relaxed, providing more general models. We have commented on some of these possibilities in the chapter conclusions, and now focus on more general points.

Firstly, this thesis has focused entirely on the case of a Lévy degradation process. A useful extension to the model would be to generalise the form of the stochastic process used to model degradation. It is clear from the derivations of cost functions given in chapters 3, 4 and 5 that the results may be extended to time-homogeneous Markov processes (Karlin and Taylor, 1981) with little effort. Whether or not the results here can be simply extended to more general Markov processes remains to be seen.

Secondly, as we commented on briefly in chapter 4, there is some scope for Bayesian methods to be applied in the case of an unobservable degradation process. When considering imperfect inspection we have assumed the degradation process has known parameters. In addition, we have considered inspection policies based only on the current observed covariate value. This could be extended to the case in which the underlying degradation process has unknown parameter values, and the inspection policy is based on the entire history of the covariate process. However, even the simple case described in chapter 4 is extremely computationally demanding. So we would expect that the general case involving complete history dependence would suffer these problems to an even greater extent.

Thirdly, a clear extension of the model described here is to the case of multi-component systems, or systems with multivariate degradation processes. In previous chapters, we have assumed that the degradation of the system is modelled by a univariate stochastic process. In many practical applications this is not a realistic assumption, since failure of a system depends on more than one factor. From a theoretical point of view there should be little difficulty in extending the results and methods given here to this case, but we would expect that the resulting equations would be extremely computationally demanding to solve.

Finally, it may be possible to consider the case in which system failure is not modelled by threshold failure, but instead is modelled as a degradation dependent hazard rate. As we have commented, this assumption is more appropriate for certain types of system, and this may provide a useful extension to our results.

More generally, there is still a need for models of condition based maintenance of modern systems. The methods used in this thesis are not new, and the results required to carry out this work have been available for some years.

Appendix A1

Approximate Solution of Integral equations

A1.1 Approximate Solution of Fredholm Integral Equations

We follow Press et al (1992). The equation we wish to solve is of the form

$$v(x) = c(x) + u(x)v(0) + \int_0^r v(y)K(y, x) dy \quad (\text{A1.1})$$

This is identical to the equation considered by Press et al, except for the atom at zero. In this section we deal with the case of a non-singular Kernel, section A1.3 gives the corresponding result for the case of a diagonally singular Kernel, which appears in the case of gamma process degradation.

Again, we define a mesh $y_i = ih$, $i = 0, 1, \dots, N$ where $h = r/N$, and apply quadrature rule based on this mesh to the integral, giving, for a particular y_j ,

$$v(y_j) = c(y_j) + u(y_j)v(y_0) + \sum_{i=0}^N v(y_i)K(y_i, y_j)w_{i,j} \quad (\text{A1.2})$$

Rearranging this gives

$$-c(y_j) = v(y_0)u(y_j) + \sum_{i=0}^N v(y_i)K(y_i, y_j)w_{i,j} - v(y_j)$$

which can be written in Matrix Form as

$$-\mathbf{c} = (\mathbf{K}^T + \mathbf{U} - \mathbf{I})\mathbf{v} \Rightarrow \mathbf{v} = (\mathbf{I} - \mathbf{U} - \mathbf{K}^T)^{-1}\mathbf{c}$$

where $\mathbf{I}$ is the $(n+1) \times (n+1)$ Identity Matrix and

$$\mathbf{v} = \begin{bmatrix} v(y_0) \\ v(y_1) \\ \vdots \\ v(y_N) \end{bmatrix} \quad \mathbf{c} = \begin{bmatrix} c(y_0) \\ c(y_1) \\ \vdots \\ c(y_N) \end{bmatrix}$$

$$\mathbf{K}^T = \begin{bmatrix} K(y_0, y_0)w_{0,0} & K(y_0, y_1)w_{0,1} & \cdots & K(y_0, y_N)w_{0,N} \\ K(y_1, y_0)w_{1,0} & K(y_1, y_1)w_{1,1} & \cdots & K(y_1, y_N)w_{1,N} \\ \vdots & \vdots & \ddots & \vdots \\ K(y_N, y_0)w_{N,0} & K(y_N, y_1)w_{N,1} & \cdots & K(y_N, y_N)w_{N,N} \end{bmatrix}$$

$$\mathbf{U} = \begin{bmatrix} u(y_0) & 0 & \cdots & 0 \\ u(y_1) & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ u(y_N) & 0 & \cdots & 0 \end{bmatrix}$$

The solution is thus reduced to solving a system of $(n+1)$ linear equations.

Having obtained the cost vector $\mathbf{v}$, we can use the original approximation as an interpolation formula, thus giving

$$v(x) = c(x) + u(x)v(0) + \sum_{i=0}^N v(y_i)K(y_i, x)w_i$$

which gives $v(x)$ as a function of the elements of the cost vector $\mathbf{v}$.

A1.2 Fredholm equations with diagonally singular kernel

As in section A1.2, we have a Fredholm equation of the form

$$v(x) = c(x) + u(x)v(0) + \int_0^r v(y)K(y|x)dy \quad (\text{A1.3})$$

In this case however, we assume that the Kernel is singular when $y = x$. We follow Press et al (1992) and modify the equation to remove the problem.

To this end, we rewrite the above equation as

$$\begin{aligned} v(x) &= c(x) + u(x)v(0) + \int_0^r [v(y) - v(x)]K(y|x)dy + \int_0^r v(x)K(y|x)dy \\ &= c(x) + u(x)v(0) + \int_0^r [v(y) - v(x)]K(y|x)dy + v(x)q(x) \end{aligned} \quad (\text{A1.4})$$

where $q(x) := \int_0^r K(y|x)dy$, which we assume exists. In this case the integral becomes zero for $x = y$, and so for purposes of computing this integral we assume $K(y, x) = 0$

Again, we define a mesh $y_i = ih$, $i = 0, 1, \dots, N$ where $h = r/N$, and apply a quadrature rule to the integral, giving, for a particular y_j ,

$$v(y_j) = c(y_j) + v(y_0)u(y_j) + v(y_j)q(y_j) + \sum_{i=0}^N [v(y_i) - v(y_j)]K(y_i, y_j)w_{i,j}$$

which, after some rearrangement becomes

$$-c(y_j) = u(y_j)v(y_0) + \sum_{i=0}^N v(y_i)K(y_i, y_j)w_{i,j} - v(y_j)[1 - q(y_j)] - v(y_j)\left\{\sum_{i=0}^N K(y_i, y_j)w_{i,j}\right\}$$

Which can be written in matrix form as

$$-\mathbf{c} = (\mathbf{U} + \mathbf{K}^T - \mathbf{Q} - \mathbf{W})\mathbf{v}$$

Where the matrices are defined by

$$\mathbf{v} = \begin{bmatrix} v(y_0) \\ v(y_1) \\ \vdots \\ v(y_N) \end{bmatrix} \quad \mathbf{c} = \begin{bmatrix} c(y_0) \\ c(y_1) \\ \vdots \\ c(y_N) \end{bmatrix}$$

$$\mathbf{K}^T = \begin{bmatrix} 0 & K(y_0 | y_1)w_{0,1} & \cdots & K(y_0 | y_N)w_{0,N} \\ K(y_1 | y_0)w_{1,0} & 0 & \cdots & K(y_1 | y_N)w_{1,N} \\ \vdots & \vdots & \ddots & \vdots \\ K(y_N | y_0)w_{N,0} & K(y_N | y_1)w_{N,1} & \cdots & 0 \end{bmatrix}$$

$$\mathbf{U} = \begin{bmatrix} u(y_0) & 0 & \cdots & 0 \\ u(y_1) & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ u(y_N) & 0 & \cdots & 0 \end{bmatrix} \quad \mathbf{Q} = \begin{bmatrix} 1-q(y_0) & 0 & \cdots & 0 \\ 0 & 1-q(y_1) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1=q(y_N) \end{bmatrix}$$

and $\mathbf{W}$ is defined by $\mathbf{W} = \text{diag}(\mathbf{w})$, and $\mathbf{w}$ is a vector defined by

$$\mathbf{w} = \mathbf{K}^T \mathbf{e}$$

where $\mathbf{e}$ is an $1 \times (n+1)$ vector of 1's. So that if $\mathbf{w}^T = (w_0, w_1, \dots, w_n)$, the matrix $\mathbf{W}$ is given by

$$\mathbf{W} = \begin{bmatrix} w_0 & 0 & \cdots & 0 \\ 0 & w_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & w_n \end{bmatrix}$$

Assuming the matrix $\mathbf{M} = \mathbf{Q} + \mathbf{W} - \mathbf{U} - \mathbf{K}$ is non-singular, (which in all cases we have considered it is), the cost vector $\mathbf{v}$ is found by

$$\mathbf{v} = (\mathbf{Q} + \mathbf{W} - \mathbf{U} - \mathbf{K})^{-1} \mathbf{c} \quad (\text{A1.5})$$

In this case however, interpolation is not as easy, since the numerical solution requires the value of the function v at the point y_j . Another problem is that numerical matrix inversion is often a difficult problem, and some combinations of parameters may lead to an ill-conditioned matrix, which does not provide a stable solution. We have not, however, encountered such problems in the calculations made for this thesis.

Appendix A2

Results in the case of Wiener Process Degradation

A2.1 Introduction

As we have remarked in this thesis an important practical example of a Lévy degradation process is the Wiener process. The calculations for computing the optimal inspection and maintenance policies are generally no more complex than those for the gamma process examples given in chapters 3, 4 and 5. The most difficult part of these calculations has been found to be computation of the joint density of the degradation process (or observed degradation process) and the maximum variable of the degradation process. In this appendix we consider these calculations as they apply to the computation of these densities in chapters 3 and 4.

Only basic numerical computations based on these results have been carried out. Generally it is found that the added complexity of a non-monotonic degradation process results in the time taken to find the optimum policy being greatly increased. As in chapter 4 this is largely due to the number of multiple integrals which have to be numerically computed during calculation.

A2.2 Joint Density of X_τ and M_τ

It is shown by Rogers and Williams (1994), that the joint density of a standard Brownian motion B_t and its maximum variable S_t is of the form

$$f_{S_t, B_t}(a, x) = \frac{2(2a - x)}{\sqrt{2\pi t^3}} \exp\left\{-\frac{(2a - x)^2}{2t}\right\} \quad a > 0, x < a$$

with $B_0 = 0$. Now let $W_t = \sigma B_t$, so that

$$\begin{aligned} P_{S_t^W, W_t}(a, x) da dx &= P_{S_t, B_t}\left(\frac{a}{\sigma}, \frac{x}{\sigma}\right) da dx \\ &= P_{S_t, B_t}(a', x') da' dx' \\ &= \frac{2(2a' - x')}{\sqrt{2\pi t^3}} \exp\left\{-\frac{(2a' - x')^2}{2t}\right\} \frac{da dx}{\sigma^2} \end{aligned}$$

Now define $X_t = W_t + \mu t$, so that, upon dividing by σ , we get

$$\frac{X_t}{\sigma} = \frac{W_t}{\sigma} + \frac{\mu t}{\sigma} = B_t + \frac{\mu}{\sigma} t$$

Now applying the Cameron-Martin-Girsanov Theorem (Rogers and Williams, 1993), we can change the measure to one incorporating a drift term at the required rate. The Radon Nikodym derivative is thus given by

$$\frac{d\mathbf{P}^{\mu/\sigma}}{d\mathbf{P}^0} = \exp\left\{\frac{\mu}{\sigma} x' - \frac{1}{2} \frac{\mu^2}{\sigma^2} t\right\}$$

So applying the change of measure gives

$$\begin{aligned} P^{\mu/\sigma}(S_t^W = a, W_t = x) da dx &= P(S_t^X = a, X_t = x) da dx \\ &= \frac{2(2a' - x')}{\sqrt{2\pi t^3}} \exp\left\{-\frac{(2a' - x')^2}{2t}\right\} \frac{da dx}{\sigma^2} \times \exp\left\{\frac{\mu}{\sigma} x' - \frac{1}{2} \frac{\mu^2}{\sigma^2} t\right\} \\ &= \frac{2(2a - x)}{\sqrt{2\pi \sigma^6 t^3}} \exp\left\{-\frac{(x - \mu t)^2}{2\sigma^2 t}\right\} \exp\left\{-\frac{2a(a - x)}{\sigma^2 t}\right\} da dx \end{aligned}$$

giving the density as required, so that in the notation of chapter 3,

$$f_\tau(m, y) = \frac{2(2m - y)}{\sqrt{2\pi \sigma^6 \tau^3}} \exp\left\{-\frac{(y - \mu \tau)^2}{2\sigma^2 \tau}\right\} \exp\left\{-\frac{2m(m - y)}{\sigma^2 \tau}\right\}$$

This however is conditional on $X_0 = 0$. It is clear that

$$f_\tau(m, y | X_0 = x) = f_\tau(m - x, y - x | X_0 = 0)$$

$$\Rightarrow f_\tau(m, y | x) = \frac{2(2m - y - x)}{\sqrt{2\pi\sigma^6\tau^3}} \exp\left\{-\frac{(y - x - \mu\tau)^2}{2\sigma^2\tau}\right\} \exp\left\{-\frac{2(m - x)(m - y)}{\sigma^2\tau}\right\} \quad (\text{A2.1})$$

for $-\infty < y < \infty, m \geq x$ and $m \geq y$.

A2.3 Distribution function of X_τ and M_τ

We require $F_{X_\tau, M_\tau}(y, m | x) := P(X_\tau \leq y, M_\tau \leq m | X_0 = x)$ for $m \geq x, m \geq y$. Then we get

$$F_{X_\tau, M_\tau}(y_0, m_0 | x) = \int_{-\infty}^{y_0} \int_x^{m_0} \frac{2(2m - y - x)}{\sqrt{2\pi\sigma^6\tau^3}} \exp\left\{-\frac{(y - x - \mu\tau)^2}{2\sigma^2\tau}\right\} \exp\left\{-\frac{2(m - x)(m - y)}{\sigma^2\tau}\right\} dm dy$$

which we can simplify as follows:

$$\begin{aligned} F_{X_\tau, M_\tau}(y_0, m_0 | x) &= \int_{-\infty}^{y_0} \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y - x - \mu\tau)^2}{2\sigma^2\tau}\right\} \int_x^{m_0} \frac{2(2m - y - x)}{\sigma^2\tau} \exp\left\{-\frac{2(m - x)(m - y)}{\sigma^2\tau}\right\} dm dy \\ &= \int_{-\infty}^{y_0} \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y - x - \mu\tau)^2}{2\sigma^2\tau}\right\} \left[1 - \exp\left\{-\frac{2(m_0 - x)(m_0 - y)}{\sigma^2\tau}\right\}\right] dy \end{aligned}$$

where we have made the obvious substitution $u = \frac{2(m - x)(m - y)}{\sigma^2\tau}$ in the inner integral.

Expanding this integral gives

$$\begin{aligned} F_{X_\tau, M_\tau}(y_0, m_0 | x) &= \int_{-\infty}^{y_0} \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y - x - \mu\tau)^2}{2\sigma^2\tau}\right\} dy \\ &\quad - \int_{-\infty}^{y_0} \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y - x - \mu\tau)^2 + 4(m_0 - x)(m_0 - y)}{2\sigma^2\tau}\right\} dy. \end{aligned}$$

Expanding the exponent of the second integral gives

$$(y - x - \mu\tau)^2 + 4(m_0 - x)(m_0 - y) = [y - (2m_0 - x + \mu\tau)]^2 - 4\mu\tau(m_0 - x)$$

So that

$$\begin{aligned}
F_{X_\tau, M_\tau}(y_0, m_0 | x) &= \int_{-\infty}^{y_0} \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y-x-\mu\tau)^2}{2\sigma^2\tau}\right\} dy \\
&\quad - \exp\left\{\frac{2\mu(m_0-x)}{\sigma^2}\right\} \int_{-\infty}^{y_0} \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y-(2m_0-x+\mu\tau))^2}{2\sigma^2\tau}\right\} dy \\
&= \Phi\left(\frac{y_0-x-\mu\tau}{\sigma\sqrt{\tau}}\right) - \exp\left\{\frac{2\mu(m_0-x)}{\sigma^2}\right\} \Phi\left(\frac{y_0+x-2m_0-\mu\tau}{\sigma\sqrt{\tau}}\right)
\end{aligned}$$

For $y \geq m$, it is clear that

$$P(X_\tau \leq y, M_\tau \leq m | X_0 = x) = P(X_\tau \leq m, M_\tau \leq m | X_0 = x) = P(M_\tau \leq m | X_0 = x)$$

Hence the distribution function of M_τ alone is given by

$$F_{M_\tau}(m | x) = \Phi\left(\frac{m-x-\mu\tau}{\sigma\sqrt{\tau}}\right) - \exp\left\{\frac{2\mu(m-x)}{\sigma^2}\right\} \Phi\left(\frac{-m+x-\mu\tau}{\sigma\sqrt{\tau}}\right)$$

We also require the form of the functions

$$K_\tau(y | x) = \int_x^c f_\tau(y, m | x) dm$$

which represents the probability density that the process reaches state y by the next inspection, but does not fail before that time. Using the above derivation it is clear that

$$\begin{aligned}
K_\tau(y | x) &= \int_x^c \frac{2(2m-y-x)}{\sqrt{2\pi\sigma^6\tau^3}} \exp\left\{-\frac{(y-x-\mu\tau)^2}{2\sigma^2\tau}\right\} \exp\left\{-\frac{2(m-x)(m-y)}{\sigma^2\tau}\right\} dm \\
&= \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y-x-\mu\tau)^2}{2\sigma^2\tau}\right\} \int_x^c \frac{2(2m-y-x)}{\sigma^2\tau} \exp\left\{-\frac{2(m-x)(m-y)}{\sigma^2\tau}\right\} dm \\
&= \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left\{-\frac{(y-x-\mu\tau)^2}{2\sigma^2\tau}\right\} \times \left(1 - \exp\left\{-\frac{2(c-x)(c-y)}{\sigma^2\tau}\right\}\right)
\end{aligned}$$

which is the product of a normal density and an exponential term. This formula is also used by Whitmore, Lawless and Crowder (1998).

A2.4 Probability density function of Y_τ and M_τ

We now consider the joint density function of Y_τ , the observed degradation level, and M_τ , the true maximum variable of the process. This is required in chapter 4 for the example of an imperfectly observed degradation process. As is stated there,

computation in the general case of a covariate process depends largely on the assumed nature of the relationship between the degradation and covariate processes.

We compute now the density conditionally on the true initial level of degradation, which can then be averaged out later. Thus we let

$$f_\tau(u, m | x) = p_{Y_\tau, M_\tau | X_0}(u, m | x)$$

with p denoting probability density. We begin by conditioning on the true degradation level X_τ so that (with abuse of notation)

$$\begin{aligned} f_\tau(u, m | x) &= \int_{-\infty}^m p(Y_\tau = u, M_\tau = m | X_\tau = z, X_0 = x) p(X_\tau = z | X_0 = x) dz \\ &= \int_{-\infty}^m p(Y_\tau = u | X_\tau = z) p(M_\tau = m, X_\tau = z | X_0 = x) dz \end{aligned}$$

The first of these densities is simple to compute,

$$p(Y_\tau = u | X_\tau = z) = p(X_\tau + \varepsilon_\tau = u | X_\tau = z) = p(\varepsilon_\tau = u - z)$$

So that

$$p(Y_\tau = u | X_\tau = z) = \frac{1}{\nu \sqrt{2\pi}} \exp\left(-\frac{1}{2\nu^2} (u - z)^2\right)$$

The second of the densities in this integral is given by equation (A2.1) in section A2.2 above. Combining these equations gives

$$f_\tau(u, m | x) = \int_{-\infty}^m \frac{2(2m - z - x)}{2\pi\nu\sigma^3\tau^{3/2}} \exp\left\{-\frac{1}{2}\left(\frac{(u - z)^2}{\nu^2} + \frac{(z - x - \mu\tau)^2}{2\sigma^2\tau} + \frac{2(m - x)(m - z)}{\sigma^2\tau}\right)\right\} dz$$

The exponent of this expression may be factorised after much tedious algebra as

$$(\sigma^2\tau + \nu^2)\left(z - \frac{u\sigma^2\tau + (2m + \mu\tau - x)\nu^2}{\sigma^2\tau + \nu^2}\right)^2 + \frac{\nu^2\sigma^2\tau}{\sigma^2\tau + \nu^2}(2m + \mu\tau - x - u)^2 - 4\mu\tau\nu^2(m - x)$$

So we obtain

$$\begin{aligned} f_\tau(u, m | x) &= \exp\left\{-\frac{1}{2(\sigma^2\tau + \nu^2)}(u + x - 2m - \mu\tau)^2\right\} \exp\left\{\frac{2\mu(m - x)}{\sigma^2}\right\} \\ &\quad \frac{2}{2\pi\nu\sigma^3\tau^{3/2}} \left[(2m - x) \int_{-\infty}^m \exp\left\{-\frac{1}{2s_\tau^2}(z - m_\tau)^2\right\} dz - \int_{-\infty}^m z \exp\left\{-\frac{1}{2s_\tau^2}(z - m_\tau)^2\right\} dz \right] \end{aligned}$$

Where

$$m_\tau = \frac{u\sigma^2\tau + (2m + \mu\tau - x)\nu^2}{\sigma^2\tau + \nu^2}, \quad s_\tau^2 = \frac{\nu^2\sigma^2\tau}{\sigma^2\tau + \nu^2}$$

The first integral can clearly be expressed in terms of the normal cumulative distribution function, thus

$$(2m - x) \int_{-\infty}^m \exp\left\{-\frac{1}{2s_\tau^2}(z - m_\tau)^2\right\} dz = (2m - x)s_\tau\sqrt{2\pi}\Phi\left(\frac{m - m_\tau}{s_\tau}\right)$$

The second integral is a truncated first moment of a normal distribution. It is clear that making a simple substitution we have (see proposition immediately below)

$$2 \int_{-\infty}^m z \exp\left\{-\frac{1}{2s_\tau^2}(z - m_\tau)^2\right\} dz = 2s_\tau\sqrt{2\pi}\left[m_\tau\Phi\left(\frac{m - m_\tau}{s_\tau}\right) - s_\tau\phi\left(\frac{m - m_\tau}{s_\tau}\right)\right]$$

So the expression becomes

$$f_\tau(u, m | x) = \exp\left\{-\frac{1}{2(\sigma^2\tau + \nu^2)}(u + x - 2m - \mu\tau)^2\right\} \exp\left\{\frac{2\mu(m - x)}{\sigma^2}\right\} \times \\ \frac{2s_\tau}{\sqrt{2\pi\nu\sigma^3\tau^{3/2}}}\left[(2m - 2m_\tau - x)\Phi\left(\frac{m - m_\tau}{s_\tau}\right) + 2s_\tau\phi\left(\frac{m - m_\tau}{s_\tau}\right)\right]$$

Rearranging we find

$$f_\tau(u, m | x) = \frac{1}{\sqrt{2\pi(\sigma^2\tau + \nu^2)}} \exp\left\{-\frac{1}{2(\sigma^2\tau + \nu^2)}(u + x - 2m - \mu\tau)^2\right\} \exp\left\{\frac{2\mu(m - x)}{\sigma^2}\right\} \times \\ \frac{2}{\sigma^2\tau}\left[(2m - m_\tau - x)\Phi\left(\frac{m - m_\tau}{s_\tau}\right) + s_\tau\phi\left(\frac{m - m_\tau}{s_\tau}\right)\right]$$

The following result is used in the above derivation:

The first truncated moment of the normal distribution is given by

$$I(m) = \int_{-\infty}^m x \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2\sigma^2}(x - \mu)^2\right\} dx = \mu\Phi\left(\frac{m - \mu}{\sigma}\right) - \sigma\phi\left(\frac{m - \mu}{\sigma}\right)$$

This can be shown using a simple substitution. Let $u = \frac{x - \mu}{\sigma}$, then

$$\begin{aligned}
I(m) &= \int_{-\infty}^{\frac{m-\mu}{\sigma}} \frac{\mu + \sigma u}{\sqrt{2\pi}} \exp\left\{-\frac{1}{2}u^2\right\} du \\
&= \mu \int_{-\infty}^{\frac{m-\mu}{\sigma}} \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{1}{2}u^2\right\} du + \sigma \int_{-\infty}^{\frac{m-\mu}{\sigma}} \frac{u}{\sqrt{2\pi}} \exp\left\{-\frac{1}{2}u^2\right\} du
\end{aligned}$$

The first integral is clearly a normal cumulative distribution function, and the second may be easily integrated applying the substitution $v = u^2$. Then, after some simplification we get,

$$\int_{-\infty}^{\frac{m}{\sigma}} \frac{x}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2\sigma^2}(x - \mu)^2\right\} dx = \mu \Phi\left(\frac{m - \mu}{\sigma}\right) - \sigma \phi\left(\frac{m - \mu}{\sigma}\right)$$

Appendix A3

Derivation of Dynamic Programming Equation (3.60)

A3.1 Derivation of Dynamic Programming equation

We follow Ross (1970) is showing that the dynamic programming equation (3.60) given in section 3.5.1 results in an optimal solution.

We consider a general inspection process, in which X_τ represents the level of degradation and let $c(x, y, \tau)$ represent the cost incurred at time τ , if inspection interval τ is chosen in state x and the process jumps to state y at the next inspection. We assume that c is a bounded function for all $x, y \in \mathbf{R}$ and $\tau > 0$. Let C denote the critical failure set.

Let π be a policy which chooses inspection interval τ in state x , so that $\pi(x) = \tau$. Then define the cost function to be the expected total discounted cost, namely

$$v_\pi(x) = E_\pi \left(\sum_{n=1}^{\infty} e^{-\delta(t_1 + \dots + t_n)} c(X_{t_{n-1}}, X_{t_n}, \tau_{n-1}) \mid X_0 = x \right)$$

where t_i represent inter-event times (an event being an inspection or a failure) and τ_i represents the inspection interval chosen at the i^{th} epoch. Expanding the summation we have

$$\begin{aligned} v_\pi(x) &= E_\pi \left(e^{-\delta t_1} c(x, X_{t_1}, \tau) + \sum_{n=2}^{\infty} e^{-\delta(t_1 + \dots + t_n)} c(X_{t_{n-1}}, X_{t_n}, \tau_{n-1}) \mid X_0 = x \right) \\ &= \int_{-\infty}^{\infty} \int_0^{\infty} E_\pi^{X_0=x, X_{t_1}=y, H_x^c=h} \left(e^{-\delta t_1} c(x, X_{t_1}, \tau) + \sum_{n=2}^{\infty} e^{-\delta(t_1 + \dots + t_n)} c(X_{t_{n-1}}, X_{t_n}, \tau_{n-1}) \right) f_\tau(h, y \mid x) dh dy \\ &= \bar{c}(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_\pi(y) f_\tau(h, y \mid x) dh dy \end{aligned}$$

where $f_\tau(h, y \mid x)$ is the density of H_x^c, X_τ given $X_0 = x$.

Since, if $v_\delta(x)$ is the optimal policy given by $v_\delta(x) = \inf_{\pi} v_\pi(x)$. Then

$$v_\pi(y) \geq v_\delta(y)$$

So that

$$\begin{aligned} v_\pi(x) &\geq \bar{c}(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_\delta(y) f_\tau(h, y \mid x) dh dy \\ &\geq \inf_{\tau} \left\{ \bar{c}(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_\delta(y) f_\tau(h, y \mid x) dh dy \right\} \end{aligned}$$

Since π is an arbitrary policy, this implies that

$$v_\delta(x) \geq \inf_{\tau} \left\{ \bar{c}(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_\delta(y) f_\tau(h, y \mid x) dh dy \right\}$$

Now, let τ_0 be such that (which exists since both c and the Kernel of the integral are bounded)

$$\begin{aligned} \bar{c}(x, \tau_0) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau_0 \wedge h)} v_\delta(y) f_{\tau_0}(h, y \mid x) dh dy \\ = \inf_{\tau} \left\{ \bar{c}(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_\delta(y) f_\tau(h, y \mid x) dh dy \right\} \end{aligned}$$

Let π be the policy which chooses τ_0 at time zero; and if the next state is y , then views the process as originating in state y , and follows a policy π_y , which is such that

$$v_{\pi_y}(y) \leq v_\delta(y) + \varepsilon, \quad y \in R$$

where ε is arbitrary, hence

$$\begin{aligned} v_\pi(x) &= \bar{c}(x, \tau_0) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau_0 \wedge h)} v_{\pi_y}(y) f_{\tau_0}(h, y | x) dh dy \\ &\leq \bar{c}(x, \tau_0) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau_0 \wedge h)} [v_\delta(y) + \varepsilon] f_{\tau_0}(h, y | x) dh dy \\ &= \bar{c}(x, \tau_0) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau_0 \wedge h)} v_\delta(y) f_{\tau_0}(h, y | x) dh dy \\ &\quad + \varepsilon \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau_0 \wedge h)} f_{\tau_0}(h, y | x) dh dy \\ &= \bar{c}(x, \tau_0) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau_0 \wedge h)} v_\delta(y) f_{\tau_0}(h, y | x) dh dy + \varepsilon K(\tau_0) \end{aligned}$$

noting that $K(\tau_0)$ is greater than or equal to zero.

Since $v_\delta(x) \leq v_\pi(x)$ we must have that

$$v_\delta(x) \leq \bar{c}(x, \tau_0) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau_0 \wedge h)} v_\delta(y) f_{\tau_0}(h, y | x) dh dy + \varepsilon K(\tau_0)$$

Hence, we obtain

$$v_\delta(x) \leq \inf_{\tau > 0} \left\{ \bar{c}(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_\delta(y) f_\tau(h, y | x) dh dy \right\} + \varepsilon K(\tau_0)$$

Since ε is arbitrary, we may make it very small, and in the limit we have

$$v_\delta(x) = \inf_{\tau > 0} \left\{ \bar{c}(x, \tau) + \int_{-\infty}^{\infty} \int_0^{\infty} e^{-\delta(\tau \wedge h)} v_\delta(y) f_\tau(h, y | x) dh dy \right\}$$

Which is the dynamic programming equation for the optimal non-periodic policy.

In computation of the optimal solution, we make use of the relationship between the hitting time of a point and the maximum of the process, so that under certain circumstances we may use the distribution of the maximum variable rather than the hitting time.

References

- 1 Abdel-Hameed, M. A. (1975), A Gamma Wear Process, *IEEE Trans. Reliab.*, Vol. R-24, pp 152-154
- 2 Abdel-Hameed, M. A. (1977), Optimal Replacement Policies for Devices subject to a Gamma Wear Process, *In The Theory and Applications of Reliability* (I.N. Shimi and C.P. Tsokos eds.), pp 397-412, Academic Press, New York.
- 3 Anderson, R. F. and Friedman, A. (1977) Optimal inspections in a stochastic control problem with costly observations, *Math. Of Oper. Res.*, Vol. 2, pp 155-190
- 4 Anderson, R. F. and Friedman, A. (1978) Optimal inspections in a stochastic control problem with costly observations II, *Math. Of Oper. Res.*, Vol. 3, pp 67-81
- 5 Ansell, J. I., Archibald, T. W. and Thomas, L. C. (1998), Life and Decay, *IMA J. Math. Appl. Bus. Industry*, Vol. 9, pp 393-401
- 6 Ansell, J. I. and Phillips, M.J. (1989), Practical Problems in the Statistical Analysis of Reliability Data (with Discussion), *J. Roy. Stat. Soc. Series C*, Vol. 38, pp 205-247.
- 7 Antelman, G. R. and Savage, I. R. (1965), Surveillance Problems : Wiener Processes, *Nav. Res. Logist. Quart.*, Vol. 12, pp 35-55
- 8 Ascher, H. and Feingold, H. (1984), *Repairable Systems Reliability*, Marcel Dekker, New York.
- 9 Asmussen, S. and Kella, O. (1998) A Multi-dimensional Martingale for Markov Additive Processes and its Applications, *To Appear*.
- 10 Barlow, R. E. and Prochan, F., (1965), *Mathematical Theory of Reliability*, John Wiley and Sons, New York
- 11 Barndorff-Nielsen, O., Blaesid, P. and Halgreen, C. (1978), First Hitting time Modles for the Generalized Inverse Gaussian Distribution, *Stoch. Proc. and their Appl.*, Vol. 7, pp 49-54
- 12 Bernardo, J.M. and Smith, A.F.M. (1992), *Bayesian Theory*, John Wiley and Sons, New York.

- 13 Breiman, L. (1968) *Probability*, Addison-Wesley, Reading, Mass., 1968
- 14 Chhikara, R.S. and Folks, J. L. (1989), *The Inverse Gaussian Distribution. Theory, Methodology and Applications*, Marcel Dekker, New York.
- 15 Çinlar, E. (1972a), Markov Additive Processes I, *Z. Wahrscheinlichkeitstheorie verw. Geb.*, Vol. 24, pp 85-93
- 16 Çinlar, E. (1972b), Markov Additive Processes II, *Z. Wahrscheinlichkeitstheorie verw. Geb.*, Vol. 24, pp 94-121
- 17 Çinlar, E. (1977), Shock and Wear Models and Markov Additive Processes, In *The Theory and Applications of Reliability* (I.N. Shimi and C.P. Tsokos eds.), pp 193-214, Academic Press, New York.
- 18 Çinlar, E. (1980) On a generalization of Gamma Processes, *J. Appl. Prob.*, Vol. 17, pp 467-480
- 19 Cox, D. R. (1972) Regression models and life tables, *J. R. Statist. Soc. B*, Vol 34, pp 187-220
- 20 Cox, D.R. and Isham, V. (1980), *Point Processes*, Chapman and Hall, London.
- 21 Cox, D. R. and Miller, H. D. (1965), *Theory of Stochastic Processes*, Chapman and Hall, London.
- 22 Cox, D.R. and Oakes, D. (1984), *Analysis of Survival Data*, Chapman and Hall, London.
- 23 Dagg, R. A. and Newby, M. J. (1998), Optimal Overhaul intervals with imperfect inspection and repair, *IMA J. Math. Appl. Bus. Industry.*, Vol. 9, pp 381-391
- 24 Dagpunar, J.S. (1997), Renewal Type equations for a general repair process, *Qual. and Reliab.Engin. Int.*, Vol. 13, pp 235-245
- 25 Delves, L. M. and Mohamed, J. L. (1985), *Computational Methods for Integral Equations*, Cambridge University Press, Cambridge, U.K.
- 26 Desmond, A. (1985) Stochastic models of failure in random environments, *Canad. J. Statistics*, Vol. 13, pp 171-183
- 27 Dykstra, R.L. and Laud. P. (1981), A Bayesian non-parametric approach to Reliability, *Annals of Statistics*, Vol. 9, pp 356-367

- 28 Erdelyi, A., Magnus, W., Oberhettinger, F. and Tricomi, F. G. (1954a) *Tables of Integral Transforms (2 Vols.)*, Mc-Graw Hill Book Company, New York.
- 29 Erdelyi, A., Magnus, W., Oberhettinger, F. and Tricomi, F. G. (1954b) *Higher Transcendental Functions (3 Vols.)*, Mc-Graw Hill Book Company, New York.
- 30 Fusaro, R., Nielsen, J. P. and Scheike, T. H. (1993), Marker dependent Hazard estimation: an application to AIDS, *Statistics in Medicine*, Vol. 12, pp 843-865
- 31 Geddes, K. O., Glasser, R. A., Moore, R. A. and Scott, T. C., (1990), Evaluation of classes of definite integrals involving elementary functions via differentiation of special functions, *Appl. Alg. in Engineering, Commun. And Computing*, Vol 1., pp 149-165
- 32 Gong, L. and Tang, K. (1997), Monitoring machine operations using on-line sensors, *Eur. J. Oper. Res.*, Vol. 96, pp 479-492
- 33 Hontelez, J.A.M., Burger, H.H. and Wijnmalen, D.J.D. (1996), Optimum condition-based maintenance policies for deteriorating systems with partial information, *Reliab. Engin. and Syst. Safety*, Vol. 31, pp 267-274
- 34 Hougaard, P. (1986), Survival models for heterogeneous populations derived from stable distributions, *Biometrika*, Vol. 73, pp 387-396.
- 35 Hoyland and Rausand (1994), *System Reliability Theory*, John Wiley and Sons, Chichester.
- 36 Hull, J. C. (1996) *Futures, Options and Other Derivatives*, Prentice Hall, New York.
- 37 IEC (International Electrotechnical Commission) (1991), *International Vocabulary*, Chapter 191, Dependability and Quality of Service, IEC 50 (191)
- 38 Jardine, A.K.S. (1973), *Maintenance, Replacement and Reliability*, Pitman, London.
- 39 Jewell, N. P. and Kalbfleisch, J. D. (1996), Marker Processes in Survival Analysis, *Lifetime Data Analysis*, Vol. 2, pp 15-29
- 40 Johnson, R. A. and Wichern, D. A. (1992) *Applied Multivariate Statistical Analysis*, Prentice-Hall International, London.

- 41 Karlin, S. and Taylor, H. M. (1975), *A first course in Stochastic Processes*, 2nd Edition, Academic Press, London.
- 42 Karlin, S. and Taylor, H. M. (1981), *A second course in Stochastic Processes*, Academic Press, London.
- 43 Kaspi, H. (1984) Storage Processes with Markov Additive Input and Output, *Math. of Oper. Res.*, Vol. 9, pp 424-440.
- 44 Kebir, Y. (1991) On Hazard Rate Processes, *Nav. Res. Logist. Quart.*, Vol. 38. pp 865-876
- 45 Kijima, M., Morimura, H. and Suzuki, Y. (1988), Periodical Replacement problem without assuming minimal repair. *Eur. J. Operat. Res.*, Vol. 37, pp 194-203
- 46 Lee, M-L. T. and Whitmore, G. A. (1993) Stochastic processes directed by randomised time, *J. Appl. Prob.*, Vol. 30, pp 302-314
- 47 Lemoine, A. J. and Wenocur, M. L. (1985), On Failure Modelling, *Nav. Res. Logist. Quart.* Vol. 32, pp 497-508
- 48 Lefebvre, M. and Gaspo, J. (1996a), Optimal Control of Wear Processes, *IEEE Trans. Automat. Contr.*, Vol. 41, pp 112-115
- 49 Lefebvre, M. and Gaspo, J. (1996b), Controlled Wear Processes: A risk-sensitive formulation, *Engineering Optimisation*, Vol. 26, pp 187-194
- 50 Lim, T.-J. (1998), A stochastic regime switching model for the failure process of a repairable system, *Reliability Engineering and System Safety*, Vol. 59, pp 225-238
- 51 Luke, Y. L. (1969) *The Special Functions and Their Approximations (2 Vols.)*, Academic Press, New York.
- 52 Myers, L. E. (1981), Survival functions induced by stochastic covariate processes, *J. Appl. Prob.*, Vol. 18, pp 523-529
- 53 Newby, M.J. (1998), Analysing Crack Growth, *Proc. ImechE. Part G – Aero. Eng.*, Vol. 212, pp 157-166
- 54 Newby, M.J. (1991), Estimating of Paris-Erdogan law parameters and the influence of environmental factors on crack growth, *Int. J. Fatigue*, Vol. 13, pp 187-198

- 55 Park, K. S. (1988a), Optimal Continuous-Wear Limit replacement under Periodic Inspections, *IEEE Trans. Reliability*, Vol. 37, pp 97-102
- 56 Park, K. S. (1988b), Optimal Wear-limit Replacement with Wear dependent Failures, *IEEE Trans. Reliability*, Vol. 37, pp 293-294
- 57 Pham, H. and Wang, H. (1996), Imperfect Maintenance, *Eur. J. Oper. Res.*, Vol. 94, pp 425-438
- 58 Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. (1992), *Numerical Recipes in C, 2nd Edition*, Cambridge University Press, Cambridge, UK.
- 59 Protter, P. (1990), *Stochastic Integration and Differential Equations*, Springer-Verlag, Berlin.
- 60 Puterman, M. L. (1994), *Markov Decision Processes : Discrete Stochastic Dynamic Programming*, John Wiley and Sons, New York.
- 61 Rishel, R. (1991), Controlled Wear Process: Modelling Optimal Control, *IEEE Trans. Automat. Contr.*, Vol. 36, pp 1100-1102
- 62 Rogers, L. C. G. and Williams, D. (1987), *Diffusions, Markov Processes and Martingales, Volume 2 : Itô Calculus*, John Wiley and Sons, Chichester.
- 63 Rogers, L. C. G. and Williams, D. (1994), *Diffusions, Markov Processes and Martingales, Volume 1 : Foundations (2nd Edition)*, John Wiley and Sons, Chichester.
- 64 Ross, S.M. (1970) *Applied Probability Models with Optimization Applications*, Holden-Day Inc., San Francisco.
- 65 Shi, M., Taylor, J. M. G. and Munoz, A (1996), Models for Residual Time to AIDS, *Lifetime Data Analysis*, Vol. 2, pp 31-49
- 66 Singpurwalla, N. D. (1995), Survival in Dynamic Environments, *Statistical Science*, Vol. 10, pp 86-103
- 67 Sobczyk, K. (1987) Stochastic Models for Fatigue Damage of Materials, *Adv. Appl. Prob.*, Vol. 19, pp 652-673
- 68 Stadje, W. and Zuckerman, D. (1991) Optimal Maintenance Strategies for Repairable Systems with General Degree of Repair, *J. Appl. Prob.*, Vol. 28. pp 384-396

- 69 Stadje, W. and Zuckerman, D. (1992) Optimal repair policies with general degree of repair in two maintenance models, *Operations Research Letters*, Vol. 11, pp 77-80
- 70 Tijms, H. C. and Van der Duyn Schouten, F. A. (1984) A Markov decision algorithm for optimal inspections and revisions in a maintenance system with partial information, *Eur. J. Oper. Res.*, Vol. 21, pp 245-253
- 71 Van Noortwijk, J. M. (1996), *Optimal Maintenance Decisions for Hydraulic Structures under Isotropic Deterioration*, Ph.D. Thesis, Technical University of Delft, Netherlands.
- 72 Villemeur, A. (1992), *Reliability, Availability, Maintainability, and Safety Assessment*, John Wiley and Sons, Chichester. (Two Volumes)
- 73 Wenocur, M. L. (1988), A Reliability Model Based on the Gamma Process and its Analytic Theory, *Adv. Appl. Prob.* Vol. 21, pp 899-918
- 74 Whitmore G. A. (1995), Estimating Degradation By a Wiener Diffusion Process Subject to Measurement Error, *Lifetime Data Analysis*, Vol. 1, pp 307-319
- 75 Whitmore, G. A., Crowder, M. J. and Lawless, J. F. (1998), Failure Inference from a Marker Process Based on a Bivariate Wiener Model, *Lifetime Data Analysis*, Vol. 4, pp 229-251
- 76 Whitmore, G. A. and Schenkelberg, F. (1997), Modelling Accelerated degradation data using Wiener diffusion with a time scale transformation, *Lifetime Data Analysis*, Vol. 3, pp 27-45
- 77 Whitmore, G. A. and Seshadri, V. (1987), A Heuristic derivation of the inverse Gaussian distribution, *The American Statistician*, Vol. 41, pp280-281
- 78 Yashin, A. I. And Manton, K. G. (1997) Effects of Unobserved and partially observed covariate processes on system failure: A review of models and estimation strategies, *Statistical Science*, Vol. 12, pp 20-34
- 79 Yeh, R. H. (1996), Optimal inspection and replacement policies for multi-state deteriorating systems, *Eur. J. Oper. Res.*, Vol. 96, pp.248-259

- 80 Zuckerman, D. (1986) Optimal Maintenance Policy for Stochastically Failing Equipment: A diffusion Approximation, *Naval Res. Logist. Quart.* Vol. 33, pp 469-477