City Research Online

City, University of London Institutional Repository

Citation: Weston, S.P. (2005). Applications of robust optimal control to decision making in the presence of uncertainty. (Unpublished Doctoral thesis, City University London)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.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/8465/

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. 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.

City Research Online: http://openaccess.city.ac.uk/ publications@city.ac.uk

Applications of robust optimal control to decision making in the presence of uncertainty

by

Stephen Peter Weston

B.A., University of Liverpool, 1978
M.A., University of Kent (1983)

A thesis submitted in partial fulfilment of the requirements of the degree of Doctor of Philosophy

Cass Business School,
City University,
London.

Copyright of all materials, ideas
and data remain with the author.

19 November 2005

Supervisor
Professor Mark Salmon

Examination Panel
Professor Stephen Taylor (External)
Dr. Russell Gerrard (Internal)

Applications of robust optimal control to decision making in the presence of uncertainty

by

Stephen Peter Weston

Submitted to the Cass Business School
on 19 November 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis is concerned with robustness of decision making in financial economics. Feedback control models developed in engineering are applied to three separate though linked problems in order to examine the role and impact of robustness in the creation and application of decision rules. Three problems are examined using robust optimal control techniques to evaluate the impact of robustness and stability in financial economic models. The first problem examines the use of linear models of robust optimal control in the pricing of castastrophe based derivatives and finds its relative performance to be superior to the popular jump diffusion and stochastic volatility models in the pricing of these emerging instruments. The novelty of the approach arises from the examination of the impact of robustness and stability of the pricing solution. The second problem involves robustness and stability of hedging. An alternative method of creating hedging rules is developed. The method is based on robust control Lyapunov functions that are simple, robust and stable in operation, yet in practice are not so conservative that they eliminate all trading gains. The third problem involves the development of robust control policies for managing risk, using non-linear robust optimal control techniques to provide clear evidence of superior performance of robust models when compared with existing VAR and EVT approaches to risk management. The novelty in the approach arises from the development of a simple and powerful risk management metric.

Contents

1 Introduction and literature review 10

1.1 Motivation 10

1.2 Definitions of key concepts 13

1.2.1 Utility based decision making under uncertainty 13

1.2.2 Robustness in decision making 15

1.3 Thesis outline 17

1.4 Literature review 17

1.4.1 Utility based decisions models under risk and uncertainty 18

1.4.2 Optimisation based decision making in the presence of uncertainty 49

1.4.3 Robust statistical methods - robust estimation 50

1.4.4 Robust statistical methods - confidence interval problems 56

1.4.5 Classical optimisation 57

1.4.6 Robustness, stability and optimal control 59

1.4.7 Linear-Quadratic Gaussian optimal control 62

1.4.8 Robust optimal control 65

1.4.9 Model uncertainty, robust control and feedback 67

1.4.10 Feedback amplification and robustness 69

1.4.11 Generalisation and the use of transfer functions 70

1.4.12 State-space theory 72

1.4.13 $H\infty$ loop shaping 74

1.4.14 Solving the $H\infty$ state-space robust optimal control problem 80

1.4.15 Robust optimal control and decision making 87

1.4.16 Modelling software and implementation 93

1.5 Scope for research - thesis organisation 95

2 Robust optimal control and the pricing of catastrophe derivatives 98

2.1 Introduction and motivation for research 98

2.1.1 Instruments and markets for catastrophe risk 100

2.2 Motivation for research 101

2.3 Review of the literature 104

2.4 Robust optimal control and catastrophe derivative pricing 118

2.4.1 Uncertainty and feedback 120

2.4.2 Robust optimal control 121

2.4.3 Robust optimal control pricing of options 122

2.4.4 Robust optimal control and the pricing of catastrophe bonds and options 130

2.5 Empirical analysis of pricing methodologies 145

2.5.1 Models and data sources 145

2.5.2 CAT bond valuation results 147

2.5.3 CAT option valuation results 167

2.6 Conclusions 177

3 Robust control and optimal hedging 180

3.1 Introduction and motivation for research 180

3.2 Review of the literature 182

3.2.1 Statistically based hedging approaches 183

3.2.2 Optimal control based hedging approaches 192

3.3 Robust hedging with non-linear optimal control 197

3.3.1 The non-linear robust optimal control approach 197

3.3.2 A pointwise min-norm control hedging rule 200

3.3.3 An example of a simple non-linear robust hedging rule 207

3.4 Empirical comparison of optimal hedging rules 209

3.4.1 Setting up the problem 210

3.4.2 Simulation results - first pass 211

3.4.3 Simulation results - second pass 217

3.5 Conclusions 218

4 Robust optimal control and risk management 220

4.1 Introduction and motivation for research 220

4.2 Review of the literature 221
4.2.1 Value at Risk 223
4.2.2 Extreme value theory 228
4.2.3 Minimax risk management 231
4.3 Robust optimal control and risk management 238
4.3.1 Stabilisation, robustness and control 241
4.3.2 A piecewise min-norm robust optimal control law for portfolio risk management . 243
4.4 Empirical comparison of risk methodologies 243
4.4.1 Setting up the problem 244
4.4.2 Simulation results 246
4.5 Conclusions 251

5 Summary, Conclusions and Suggestions for Further Research 253

5.1 Summary of the thesis 253
5.2 Conclusions from the thesis 255
5.3 Limitations of the research 255
5.4 Suggestions for further research 256

6 Bibliography 257

7 Appendices 273

List of Tables

Chapter 1

Number	Description
1.1	State-contingent markets

Chapter 2

Number	Description
2.1	Monte Carlo calibration cases
2.2	CAT bond parameters
2.3	Time Steps v Execution Time v Accuracy (non-linear Lyapunov)

Chapter 3

Number	Description
3.1	Hedging strategies evaluated
3.2	Simulation results for alternative hedging strategies
3.3	Historical simulation: EUR low volatility environment
3.4	Historical simulation: EUR high volatility environment
3.5	Historical simulation: GBP low volatility environment
3.6	Historical simulation: GBP high volatility environment
3.7	Historical simulation: USD low volatility environment
3.8	Historical simulation: USD high volatility environment
3.9	Simulation results for flattened PWMN hedging rule

Chapter 4

Number	Description
4.1	Portfolio risk management measures
4.2	Simulated volatility environments

List of Figures

Chapter 1

Number	Description
1.1	Non-linear expected utility and fanning out
1.2	Optimum for individual decision maker
1.3	Risk aversion in the Yaari model
1.4	Wealth dynamics and risk aversion
1.5	State preference and optimal insurance
1.6	Optimal insurance with varying risk aversion
1.7	Simple feedback control system
1.8	Optimal path in state-time space
1.9	Model uncertainty
1.10	Schematic view of a feedback amplifier process
1.11	Using feedback in a simple linear system
1.12	System state feedback
1.13	$H\infty$ loop shaping in a simple feedback system
1.14	Simplified block diagram illustrating impact of perturbations
1.15	Disturbances and model uncertainty in a closed loop system

Chapter 2

Number	Description
2.1	Schematic representation of feedback control process
2.2	Discretisation using Markov decision process
2.3	Short term CAT bond valuation using Cox-PIDE
2.4	Short term CAT bond valuation using numerical HJI
2.5	Short-term CAT bond model valuation: numerical HJI v Cox-PIDE
2.6	Short term CAT bond valuation using linear Lyapunov
2.7	Short-term CAT bond model valuation: Linear Lyapunov v Cox-PIDE
2.8	Short term CAT bond valuation using non-linear Lyapunov
2.9	Short term CAT bond valuation: Non-Linear Lyapunov v Cox-PIDE
2.10	Long term CAT bond valuation using Cox-PIDE
2.11	Long term CAT bond valuation using numerical HJI
2.12	Long term CAT bond valuation using linear Lyapunov
2.13	Long term CAT bond valuation using non-linear Lyapunov
2.14	Long term CAT bond valuation: Numerical HJI v Cox-PIDE
2.15	Long term CAT bond valuation: Linear Lyapunov v Cox-PIDE
2.16	Long term CAT bond valuation: Non-Linear Lyapunov v Cox-PIDE
2.17	PCS call option spread value
2.18	PCS CAT option valuation: Cox-TGA v numerical HJI
2.19	PCS CAT option valuation: Cox-TGA v Linear Lyapunov
2.20	PCS CAT option valuation: Cox-TGA v Non-Linear Lyapunov
2.21	PCS CAT option valuation: Cox-TGA v Non-Linear Lyapunov
2.22	Unit price comparison - pre-loss period

Chapter 3

Number	Description
3.1	Mean-variance hedging
3.2	Comparison of non-linear robust hedging strategies
3.3	Comparison of output using non-linear robust hedging strategies
3.4	Comparison of robust and non-robust hedging strategies

Chapter 4

Number	Description
4.1	Relative robustness of PWMN control laws
4.2	MC simulation - low volatility environment
4.3	MC simulation - flat volatility environment
4.4	MC simulation - high volatility environment
4.5	MC simulation - skew volatility environment
4.6	MC simulation - 10% jump in volatility
4.7	MC simulation - 30% jump in volatility
4.8	MC simulation - 50% jump in volatility
4.9	MC simulation - 100% jump in volatility
4.10	MC simulation - 500% jump in volatility
4.11	Historical simulations of portfolio level risk

Acknowledgements

I have been extremely fortunate to have the invaluable assistance and support of many people in the course of researching and writing this thesis. Professors Forest Capie and Geoffry Wood were kind enough to tolerate my false starts and wise enough to pass me to Professor Mark Salmon where the real work took place. Mark's kindness, enthusiasm and tolerance for the endless missed deadlines and excuses when I only too frequently found it almost impossible to balance the demands of research, developing financial software through my own company, consultancy and a young family, have been nothing short of saintly. So Mark, thank you from the bottom of my heart for everything.

At the very core of my efforts and my emotional bedrock, has been my long suffering wife Dods. It is she more than anyone else that has suffered the endless evenings alone, missed weekends and holidays, which together were required to complete this thesis during my spare time. Without you my darling this thesis would simply not have been possible. I offer you all my love and thanks for everything you have selflessly given and surrendered. Finally, my children Thomas and Hannah have not always seen as much of me as I think they would like, but have invariably been cheerful, curious and understanding. So thank you for your love and support my sweethearts. This thesis is decidated to the three most wonderful and special people that anyone could ever have the expectation to be loved by and to love.

Carrying out an undertaking as onerous as a thesis on a part-time basis makes one fully aware of the length of the journey from sperm to worm. Chairman Mao is once supposed to have said that it is better to travel in hope than to arrive. I agreed with him at the beginning of this thesis when my curiosity and enthusiasm were both high, but not any longer. It is both satisfying and relieving to complete the task and to arrive at one's final destination ! The process has been incredibly hard work but has provided immense satisfaction.

Thank you one and all for everything. I am now and always will be, totally in your debt. This thesis is dedicated to Dods, Thomas and Hannah with all my love.

Stephen Weston

London

September 2004

Chapter 1

Introduction and literature review

1.1 Motivation

The objective of this thesis is to investigate the impact and role of robustness on decision making in the presence of uncertainty. This objective immediately begs the obvious question of how to define the fundamental concepts of robustness, decision making and uncertainty. The following four quotes sum up the ideas that originally stimulated the curiosity of this author to research the area of decision making under uncertainty

> **I'd rather be vaguely right than precisely wrong.**
>
> Attributed to J. M. Keynes in Forbes magazine 25 January 1999 issue.

> **Far better an approximate answer to the right question, which is often vague, than an exact answer to the wrong question, which can always be made precise.**
>
> J. Tukey, 1968.

> **How well do we need to know the answer in order to make a decision ?**
>
> G.J. Macrae and G. Wang, MIT, 2003.

The most fundamental of the three concepts mentioned above is decision making, yet it is difficult to arrive at an acceptably clear definition of decision making without first defining the concepts of risk and uncertainty. The very notion of uncertainty is in and of itself, fraught with disagreement, despite which, there exists a substantial body of research and opinion behind the notion that it is vital to make such a fine distinction between the concepts of risk and uncertainty. Many economists, psychologists and behavioural scientists remain unconvinced of the validity of such a fine distinction. However, having

introduced such a distinction, it is important to explore its validity and form a view of how it does or does not impact upon the work in this thesis.

Distinguishing between the concepts of risk and uncertainty can be fairly directly ascribed to Knight's treatise on Risk, Uncertainty and Profit first published in 1921. It was Knight's research that for the first time made a clear case for the economic importance of and distinction between, risk and uncertainty. Knight explicitly linked profits, entrepreneurship and free enterprise to what he claimed were the separable concepts of risk and uncertainty. Subsequent contributions by authors such as Hicks in 1931, Keynes in 1936 and 1937, followed by Kalecki in 1937, Stigler in 1939, Tintner in 1941, Hart in 1942 and Lange in 1944 also marked key points in the search within economic theory to ascribe a pivotal role to risk or uncertainty in the explanation of how the behaviour of central quantities such profits, investment and the demand for liquid assets (to name but a few), are influenced by uncertainty and risk. A critical barrier to this understanding was making precise the distinction between risk and uncertainty.

In chapter 7 of his 1921 work, Knight interprets risk as referring to situations where the decision maker can assign mathematical probabilities to the randomness he faces. Whereas Knight refers to uncertainty as being those situations when randomness cannot be expressed in terms of specific mathematical probabilities. Keynes wrote about Knight's distinction on several occasions (e.g. page 148 of the General Theory), but the following extract from Keynes's writings probably gets as close as is possible to an unambiguous distinction between the two ideas

> *By uncertain knowledge, let me explain, I do not merely mean to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty...The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence...About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.*
>
> J.M. Keynes, Economic Journal, 1937.

However, for many economists, the distinction between risk and uncertainty remains artificial. Those against the distinction, for example, argue that in Knightian uncertainty the problem is that the decision maker simply *does not* assign probabilities, not that the decision maker actually *cannot* assign probabilities. In other words, the problem is one of knowledge of probabilities rather than the existence of probabilities in the first place, i.e. the problem is epistemology versus ontology. Yet other economists argue that there are actually no probabilities to be known anyway, because probabilities are only beliefs in the first place. In other words, one could easily and defensibly view probabilities as just subjectively, as-

signed expressions of belief which would imply that they have no necessary connection to the randomness of the real world, which may or may not be random in and of itself.

Despite these objections, many economists view as crucial the distinction between risk and uncertainty. Shackle (1949, 1955, 1961 and 1979) and Davidson (1982 and 1991/1995) argue that Knightian uncertainty is the only relevant form of randomness in economics, especially when considered in conjunction with the fact that information is not and can not ever be always and everywhere known instantly, completely and accurately. Looked at from the Knightian perspective, risk is only possible in extremely contrived situations where circumstances are measurable and repeatable such as in casinos. The key point is that in the real world, decision making situations are frequently unique, with only a small set of the possible alternatives either known or measurable. In such situations, it is clearly impossible to make mathematical probability assignments, so that decision rules should be considered differently from the conventional approach of expected utility.

It is interesting to note at this point that the greater part of research concerning uncertainty has arguably been too narrowly focused on uncertainty with regard to information. Yet when viewed from a rational expectations perspective, research should also encompass the uncertainty surrounding the choice of model by agents. Is the model always relevant and correct, or is it merely one of a family of similar models ? Or yet again, is a model only the relevant model over a particular range of outcomes ? In short, model uncertainty exists and needs to be analysed and integrated into any explanation of decision making under uncertainty.

What *is* clear, is that the Knightian definitions are useful in making a distinction between those theories of decision making which do make assignments of probabilities to alternatives and those that do not. In the former group are obviously the expected utility theories that use von Neumann and Morgenstern type objective probabilities, whilst in the latter group are the state preference approach of Arrow and Debreu (1954). Both of these approaches to uncertainty will be examined in extensive detail in the literature review that forms the main body of this introductory chapter.

As this thesis seeks to bring together ideas which may not be totally familiar to many economists, a clear statement of concepts, approach, tools and structure is useful in order to better understand the research that has been undertaken. This initial chapter therefore provides four things

- Clear definitions of the fundamental concepts relevant to this thesis.
- A review of the existing research and literature on decision making under uncertainty.
- A review of the existing research and literature on approaches to modelling robustness.
- An explanation of the theory behind the tools and approaches used to carry out the research reported

in chapters 2, 3 and 4.

1.2 Definitions of key concepts

1.2.1 Utility based decision making under uncertainty

Economists define utility as the real or desired ability of a good or service to satisfy a human want, so that marginal utility is the change in utility due to a one unit change in the quantity of a good or service consumed. In decision theory, utility is a measure of the desirability of consequences of courses of action that applies to decision making in the presence of risk or uncertainty. The concept of utility applies to both single and multi-attribute consequences. The fundamental assumption is that the decision maker always chooses the alternative for which the expected value of the utility is maximised. If this assumption holds, then utility theory can be used to predict or prescribe the choice that the decision maker will or should make, among the available alternatives. Utility therefore has to be assigned to each of the possible (and mutually exclusive) consequences of every alternative.

A utility function is thus the rule by which this assignment is achieved and depends on the preferences of the individual decision maker. In utility theory, utility measures u of the consequences are assumed to reflect a decision maker's preferences in the following sense

1. The numerical order of utilities for consequences preserves the decision maker's preference order among the consequences.

2. The numerical order of expected utilities of alternatives (referred to, in utility theory, as gambles or lotteries) preserves the decision maker's preference order among these alternatives (lotteries).

Utility theory provides a basis for the assignment of utilities to consequences by formulating necessary and sufficient conditions to satisfy 1 and 2. A utility function is defined mathematically as a function that maps from the set of consequences into the real numbers that provides for satisfaction of 1 and 2. Whilst there exist various methods for constructing utility functions, the best-known is based on the indifference judgments of the decision maker about specially constructed alternatives(lotteries). One of the key features of utility theory is that it provides the ability distinguish between differing attitudes to risk, making it possible to delineate risk loving, risk neutrality and risk aversion. By convention, we use the term Bernoulli Utility Function to refer to a decision-maker's utility over wealth - since of course it was Bernoulli who originally proposed the idea that people's internal, subjective value for an amount of money was not necessarily equal to the physical value of that money. The term von Neumann-Morgenstern Utility Function is used to refer to a decision-maker's utility over lotteries, or gambles.

But what is a decision ? A decision is an irrevocable allocation of resources. Abstracting from both psychology and economics and viewing decision making at the most basic level, all individuals are faced with an inescapable need to make choices which are usually based on some rule or measure. For example, the choice of whether to have work or leisure, is usually based on a series of rules along the lines of needing a certain amount of income to pay bills, yet still have enjoyment. This deceptively simple choice is perhaps the most fundamental that consumers make, namely, what is the point beyond which further increases in income do not compensate for lost leisure opportunities. The need to choose arises for no other reason than the fact that wants almost always exceed available resources, making choice inevitable.

Decisions are therefore necessary as a way of selecting particular outcomes, independent of whether the choices are couched in terms of courses of action, or options, or moves, or even payoffs. As Gilboa and Schmeidler (2001) point out, there are two pre-eminent paradigms for the formal encapsulation of human reasoning that have also been applied to decision making. Namely, probabilistic reasoning based on the Bayesian approach, which when coupled with the ubiquitous concept of utility maximisation, represents the dominant paradigm for formal models of decision making under uncertainty. The second is based around the idea of rule-based systems[1]. The rule-based approach forms the underpinning of dynamic programming which was developed by Bellman (1957) to deal with multi-stage problems involving the need to produce the best possible final outcome when faced with sequences of decisions over time and to which considerable analysis will be devoted later in this chapter.

Making decisions in the presence of either risk or uncertainty demands the construction of a framework, usually in functional form, which allows the comparison of preferences such that it is possible to evaluate the level of satisfaction enjoyed by the decision maker bearing the risk or uncertainty. Once a functional metric has been derived, the classical theory of choice tells us that decisions can be made based on the simple criteria of maximisation of the level of satisfaction, or utility, $U(X)$, of the decision maker based on the range of outcomes X. Although Bernoulli first introduced the concept of expected utility as far back as 1738, with the possible slight exception of Ramsey's work in 1926, the formal incorporation of the concepts of risk and uncertainty into economic theory for the purposes of explaining choice and decision making was only really achieved in 1944 with the publication of von Neumann and Morgenstern's seminal book "Theory of Games and Economic Behaviour".

In 1951, Arrow provided the first real survey of the then contemporaneous thinking on decision making under risk and uncertainty. The survey highlighted the point that the principal difficulty facing researchers was the task of deriving a precise definition of what it means for uncertainty or risk to affect economic

[1]In fact, though not of interest here, Gilboa and Schmeidler offer a third and additional alternative in the form of case-based decision making, where decisions are made based on reasoning by analogy.

decision making. Much energy was devoted to the task of defining the exact process by which decision making agents evaluate alternatives whose payoffs are random, with particular emphasis on the way in which increasing or decreasing uncertainty consequently leads to changes in behaviour. However, much of this early work remained unsystematic and unfocused in its approach with heuristics rather than clear rules emerging from the research.

In retrospect, it is easy to see that the element required to complete the argument was the formalisation, in an axiomatic fashion, of choice in the context of risk or uncertainty. Hicks (1931) and Tintner (1941) had already hinted at the idea that agents form preferences over distributions of possible outcomes, but had not arrived at the issue of the separability of attitude to risk and uncertainty from pure preferences over outcomes, which was the core of the problem. Several approaches to deriving the form of the utility function have been suggested and an analysis of the most important alternative types and forms of utility function and the associated loss functions in the decision making process, forms a central element in the literature review that forms the bulk of this chapter.

1.2.2 Robustness in decision making

Having made a preliminary attempt at defining uncertainty and decision making, the third key concept of robustness can be tackled. The concept central to an understanding of robustness is that economic agents making a decision do so based on *some* (possibly) approximating model. Uncertainty inevitably leads a decision maker to doubt the model, because the current model is viewed as an approximation in the sense that the model belongs to an imprecisely specified set of models that are in the neighbourhood of the approximating model. The concern of the decision maker about model, parameter and data misspecification and hence about uncertainty, leads to a preference for decision rules that work well over the set of nearby models. The issues surrounding model uncertainty and model risk are dealt with in detail in the literature review later in this chapter, but for now it is sufficient to state that if a decision rule works well in a neighbourhood of potential models then it can generally be said to be robust.

The application of the concept of robustness in economics is not new. In the world of rational expectations, decision makers are assumed to act as if they know the relevant economic model, which they are assumed to perceive in the form of a transition law or transfer function that links state variables, controls and a description of the stochastic shocks. Rational expectations associates a distinct decision rule with each specification of a shock process. It is easy to see the tightness of the implied linkage, when it is recalled that under rational expectations decision rules are a function of the serial correlations of shocks. This, in effect, is the world of optimal control theory, where it is assumed that decision makers know the form of the transition law (or transfer function) such that a distinct decision rule is associated

with each specification of the shock process. When robustness enters the picture, it effectively loosens the temporal linkage between shocks and decision rules.

Robust control theory treats models as approximations. Robust optimal control seeks to find one optimal rule that performs the best for an entire set of models. The alternative models can be imprecisely stated in terms of shocks, but feedback from the shocks is assumed to occur and thereby affect states, so representing mis-specified dynamics. To be robust, a decision maker prefers rules that are optimal across the widest possible range of specifications. Notice that the concept of feedback has suddenly appeared. This raises two questions, what is feedback and why is it important for robust decision making ? Feedback is related directly to control theory to the point that it the latter could be regarded as the theory of feedback systems. Modelling feedback explicitly is required because of the disturbances and uncertainty surrounding the model upon which decisions are based.

Perhaps more fundamentally than the above, however, is the question of why study robustness in the first place. What extra insights can robustness add to the study of decision making under uncertainty ? There are three main answers to this question. First, as already mentioned above, doubts about the decision model. If agents doubt their hedging model so will policy makers trying to develop decision rules. This is the essence of the idea that gave rise to the chapter on robust hedging rules. Second, rational expectations models - in common with many other models - give rise to prices that imply too high a price for risk. This manifests itself in such issues as the so called "equity premium puzzle". The third reason is model misspecification and, more recently, specific work on model risk attempting to quantify the costs of using the wrong model.

The final issue relating to robustness is measurement. How should robustness be measured so that the relative and absolute performance of decision rules can be compared, given that the existence of uncertainty implies that models are likely to be incorrect and consequently any measures are also likely to be suspect ? Sargent (2001) suggests two frequency-domain preference specifications that will produce robust decision rules, namely, the familiar one of entropy from classical statistics and the less familiar one of $H\infty$ which derives from robust optimal control theory in the world of engineering. Used in the context of decision theory, the entropy objective function summarises doubts about model specification in a single parameter. In contrast, by using a broader norm based measure, the $H\infty$ specification is designed to measure the degree of robustness to varying levels of uncertainty. The $H\infty$ norm can be used to shape the frequency response function for the decision maker in order to meet certain robustness objectives.

1.3 Thesis outline

Much of the remainder of this chapter is therefore occupied by a two-part literature review which seeks to provide first a review and critique, followed by a synthesis that explores the breadth and depth of existing research on robustness and decision making in the presence of uncertainty and thereby serves to demonstrate the broad scope for new and original work. The final part of the chapter uses the literature survey as the basis for identifying the much narrower area of research interest and explaining the tools and techniques that form the basis for the core of the research work contained in this thesis. The three specific problems of robust decision making in the face of uncertainty that are considered in turn in chapters 2, 3 and 4. Chapter 2 develops a robust optimal control based model for pricing catastrophic derivatives, whilst chapter 3 investigates using the same approach for hedging both vanilla and catastrophic derivatives. Chapter 4 widens the scope of chapter 3 to address the wider issue of robustness in various aspects risk management.

1.4 Literature review

Armed with some broad ideas about the key concepts surrounding uncertainty, decision making and robustness, underpinned by the integrating notion of feedback, the aim of this section is to review and synthesize the existing literature on robustness of decision making under uncertainty, so that the final section of this chapter can lead to an identification of and justification for, the area of research that forms this thesis.

As research into decision making under uncertainty has tended to follow one of two broad lines of enquiry, namely, choice based utility type models and state-space type general equilibrium models, this broad distinction will therefore be used as a method to classify and analyse existing approaches. This literature review is therefore structured as follows:

- *Utility based decision models under risk and uncertainty*: expected utility, subjective expected utility, state preferences, characterisation of risk aversion.
- *Robust statistical based decision models under uncertainty*: robust estimation, confidence interval problems, classical optimisation.
- *Robust optimisation based decision making in the presence of uncertainty*: robustness, stability and optimal control, linear-quadratic-Gaussian models, state-space based robust optimal control.

The principal message to emerge from this review is that the decision rules produced by conventional utility based decision making models are frequently neither robust nor stable. In contrast, the robust

optimal control approach can be used to compute robust decision rules that provide robustness and stability both in theory and practice.

1.4.1 Utility based decisions models under risk and uncertainty

Expected utility

The expected utility hypothesis that underpins much of classical micro-economics derives from the solution by the Swiss mathematician Daniel Bernoulli in 1738 of the St. Petersburg Paradox which had been posed in 1713 by his cousin Nicholas Bernoulli[2]. The apparent Paradox challenges the idea that individuals value random outcomes based on a comparison of returns expected from each outcome. The Paradox posed a question based on the flipping of a fair coin until a head appears. If the first head appears on the n-th toss, then the payoff is n units. How much should be paid by a player to play such a game ? The paradox arises because whilst the expected return, expressed in terms of wealth w, is infinite

$$E(w) = \sum_{i=1}^{\infty} 2^n \left(\frac{1}{2}\right)^n = 2\left(\frac{1}{2}\right) + 2^2\left(\frac{1}{4}\right) + 2^3\left(\frac{1}{8}\right) = 1 + 1 + 1... = \infty \tag{1.1}$$

it is simultaneously obvious that no rational player would pay an infinite amount to play such a game. Bernoulli's solution to this paradox was based on two fundamental ideas. First, that utility derived from wealth increases at a decreasing rate - the principle of diminishing marginal utility, as captured by $u'(w) > 0$ and $u''(w) < 0$. Second, that the valuation of an outcome is not its expected return, but its expected utility, such that the value of the above game to a player possessing zero initial wealth is

$$E(w) = \sum_{i=1}^{\infty} u2^n \left(\frac{1}{2}\right)^n = u(2)\left(\frac{1}{2}\right) + u(2^2)\left(\frac{1}{4}\right) + u(2^3)\left(\frac{1}{8}\right) + ... < \infty \tag{1.2}$$

which Bernoulli argued to be finite due to diminishing marginal utility. Using a logarithmic utility function of the form $u(x) = \alpha \log x$, Bernoulli demonstrated that a player would only be willing to pay a finite amount of money to play, despite the expected return being infinite. His logic lead to the adoption of the following expression as the standard way of valuing a risky venture

$$E(u|p, X) = \sum_{x \in X} p(x)\, u(x) \tag{1.3}$$

[2]It is also common to note that Gabriel Cramer, another Swiss mathematician, also provided effectively the same solution ten years before Bernoulli.

where X is the set of possible outcomes, $p(x)$ is the probability of a particular outcome $x \in X$ and $u : X \rightarrow R$ is a utility function over the set of outcomes[3].

Despite being of central importance during the Marginalist Revolution of 1871-74 and being key to the works of Jevons (1871) and Walras, diminishing marginal utility was not given centre stage as an explanation of decision making until von Neumann and Morgenstern's work in 1944 axiomatised the concept and expressed it in terms of preferences over alternative random outcomes (which they termed lotteries). Their work assumes that probabilities are objective or exogenous and incapable of influence by players. But, the problem was for the player to decide or choose between lotteries to find the best. von Neumann and Morgenstern's major contribution was to show that if a player has preferences defined over lotteries then there must exist a utility function $U : \Delta(X) \rightarrow R$ which assigns a utility to every lottery $p \in \Delta(X)$ that represents those given preferences. Their insight was thus to confine themselves to preferences over distributions and then to deduce the implied preferences over the underlying outcomes.

It is not necessary to reproduce von Neumann and Morgenstern's proof of their utility function. It is, however, necessary and worthwhile to state their axioms of preference over simple lotteries as a basis for understanding the construction of their utility function. Thus, if $\geq_h$is a binary relation over $\Delta(X)$, such that $\geq_h \subset \Delta(X) \times \Delta(X)$, then it is possible to write $(p,q) \in \geq_h$, or $p \geq_h q$ to show that lottery p is preferred to lottery q and that $p \sim_h q$ means p is equivalent to q, whilst if p is not preferred or equivalent to q, then this is expressed as q being strictly preferred to p, so that von Neumann and Morgenstern's axioms for these preferences can be stated as

Axiom 1 $\geq_h$*is complete, i.e. either* $p \geq_h q$, *or* $q \geq_h p$ *for all* $p,q \in \Delta(X)$.

Axiom 2 $\geq_h$*is transitive, i.e. if* $p \geq_h q$ *and* $q \geq_h r$, *then* $p \geq_h r$ *for all* $p,q \in \Delta(X)$.

Axiom 3 *if* $p,q,r \in \Delta(X)$, *such that* $p \geq_h q \geq_h r$, *then there is an* $\alpha,\beta \in (0,1)$ *such that* $\alpha p + (1-\alpha) r >_h q$ *and* $q >_h \beta p + (1-\beta) r$[4].

Axiom 4 *for all* $p,q,r \in \Delta(X)$ *and any* $\alpha \in [0,1]$, *then* $p \geq_h q$ *if and only if* $\alpha p + (1-\alpha) r \geq_h \alpha q + (1-\alpha) r$[5].

[3] Technically speaking, utility must also be bounded in order for the concept of diminishing marginal utility to provide an answer to the St. Petersburg paradox. If there is no bound then it is possible to find a series of payoffs $x_1, x_2, x_3 ...$ capable of yielding infinite expected value. Karl Menger only pointed out this issue in 1934 almost 200 years after Bernoulli's original work !

[4] *This is often referred to as the Archimedean axiom as it functions akin to a continuity axiom over preferences.*

[5] *This is the so called independence axiom and asserts that the preference between p and q is unaffected if they are both combined in the same way with a third lottery r. To achieve this requires recourse to a two stage lottery type of argument, which is slightly artificial, but is useful to illustrate that preferences between two stage lotteries ought to depend solely on preferences between the alternative lotteries.*

von Neumann and Morgenstern's utility function, $U : \Delta(X) \rightarrow R$ therefore represents preferences over lotteries, such that the preference between p and q is unaffected if they are combined in some arbitrary way with a third lottery, r. They show that the U has a representation

$$U(p) = \sum_{x \in Supp(p)} p(x)\, u(x) \tag{1.4}$$

where $u : X \rightarrow R$ is an elementary utility function over the underlying outcomes X. Given that $\Delta(X)$ is a convex set of probability distributions on X, then if $p \in \Delta(X)$, p must have a finite support (shown as $Supp(p) \subset X$[6]. This provides sufficient background to be able to state von Neumann and Morgenstern's expected utility theorem

If $\Delta(X)$ is the set of all simple probability distributions on X and $\geq_h$is a binary relation on $\Delta(X)$, then $\geq_h$satisfies the axioms 1-4 above, if and only if there is a function $u : X \rightarrow R$ such that for every $p, q \in \Delta(X)$

$$p \geq_h q \text{ if and only if } \sum_{x \in Supp(p)} p(x)\, u(x) \geq \sum_{x \in Supp(q)} q(x)\, u(x) \tag{1.6}$$

Such conventional, non-stochastic utility functions are usually assumed to be ordinal in that they preserve the order of the indexes of preferences. From the above formulation, it is easy to fall into the trap of thinking that utility is cardinal and can be used as a measure of preferences. However, the key point is that even though the elementary utility function of equation 1.3 above is cardinal in outcomes, the function is not cardinal over lotteries. This is simply because the utility function is based on lotteries and not on outcomes.

This line of reasoning suggests that the Independence Axiom is the foundation of von Neumann and Morgenstern's work. This is only partially true and was demonstrated by Allais in 1953 in his famous paradox over two pairs of lotteries. The apparent paradox is best understood by means of the following simple example. Imagine a quartet of distributions $(p1, p2, q1, q2)$. These points are outcomes over the following lotteries: $x_1 = 0$, $x_2 = 100$ and $x_3 = 500$, such that the first pair of lotteries p_1 and p_2 have the following payouts:

p_1: pays \$100 with certainty.

p_2: pays \$0 with 1% chance, \$100 with 89% chance, \$500 with 10% chance.

When confronted with this choice, decision makers usually opt for p_1, so that $p_1 \geq_h p_2$, which if

[6]Convexity also implies that for any $p, q \in \Delta(X)$, $\alpha p + (1-\alpha) q \in \Delta(X)$ for any $\alpha \in (0,1)$ and that if p and q are no more than simple probability distributions, then

$$(\alpha p + (1-\alpha) q)(x) = \alpha p(x) + (1-\alpha) q(x) \tag{1.5}$$

for any $x \in X$.

depicted using indifference curves would imply that there existed an indifference curve U_p such that p_1 were above and p_2 below. Imagine now a second pair of lotteries that have the following payouts:

q_1: pays \$0 with 89% chance and \$100 with 11% chance.

q_2: pays \$0 with 90% chance and \$500 with 10% chance.

Now, standard expected utility theory would predict that given that indifference curves are parallel to each other, then q_1 should be preferred to q_2, which would imply the existence of an indifference curve U_q with q_1 above and q_2 below, i.e. $q_1 \geq_h q_2$. This much is guaranteed by the independence axiom, which would imply $p_1 \geq_h p_2 \Rightarrow q_1 \geq_h q_2$.However, Allais (1953) showed that when presented with these two pairs of lotteries, decision makers first opt for p_1 over p_2 in the first lottery, but then choose q_2 over q_1 in the second lottery, thereby contradicting the prediction of the independence axiom.

Allais's own answer to this apparent paradox was that indifference curves are parallel, but not linear, so that they are capable of "fanning out". Hagen (1972 and 1979) further exploits the fanning out idea and shows that it is indeed the case, thereby proving that the decomposition predicted by the expected utility hypothesis is incorrect. This is because the utility of a particular lottery p, is not $U(p) = E(u;p) = \sum_{x \in X} p(x)\, u(x)$, but is instead $U(p) = f[E(u;p), var(u;p)]$. In other words, the utility of the lottery is a function of both its expected utility as well as the variance of the utility. Interestingly, the fanning out idea encompasses Kahneman and Tversky's (1979) "common consequence" concept. In order to understand the common consequence effect, it is necessary to appeal to the independence axiom, which claims that if $p >_h q$, then for any $\beta \in [0,1]$ and $r \in \Delta(X)$, then $\beta p + (1-\beta) r >_h \beta q + (1-\beta) r$. In other words, the possibility of an extra alternative lottery r should not have any impact on the preferences between the original lotteries p and q.

Kahneman and Tversky's common consequence argument is that the addition of r affects preferences between p and q, such that p and q now become second best alternatives if r fails to occur. In effect, they are claiming that if the prize in r is substantial, then a greater degree of risk-aversion is likely, which modifies preferences between p and q, so that less risky choices are made. The idea is that if r is a substantial prize and fails to occur, then disappointment occurs instead in direct proportion to the size of r. Intuitively, Kahneman and Tversky's common consequence argument is that getting a \$50 consolation prize in a multi-million dollar lottery is considerably less exciting than merely finding \$50 on the street. Consequently, in order to compensate for the potential disappointment, there is less willingness to take on risks as an alternative - as doing so would only worsen the burden. In contrast, if r is not that good, then one might be more willing to take on risks.

The response to this line of argument against expected utility came from Marschak (1951) and Savage (1954) in the form of an argument claiming that expected utility is only meant to be normative,

as it tries to describe what rational behaviour *should* occur in the face of uncertainty, not what behaviour *actually* occurs. Allais (19539) and Hagen (1972) worked directly with different combinations of the elementary utility functions on outcomes, rather than modifying the underlying axioms of the von Neumann-Morgenstern utility function on lotteries and as a result, they did not make it clear that their "fanning out" hypothesis was a result of normal rational behavior. In the late 1970's and early 1980's new axioms were introduced that lead to the production of the Allais hypothesis as a result. Although the literature on alternative expected utility expanded enormously in the 1980s and 1990s, much of the flavour and most of the key points can be encapsulated by providing a mainly heuristic review of two alternative concepts: Weighted Expected Utility as suggested by Chew-MacCrimmon (1979) and Non-Linear Expected Utility put forward by Machina (1982). John Quiggin's (1982, 1993) Rank-Dependent Expected Utility is also interesting, but adds relatively little to the tenor or direction of the debate.

One of the first axiomatic treatments of the theory of choice under uncertainty which incorporates the fanning out hypothesis of Allais was the "weighted expected utility" introduced by Chew and MacCrimmon (1979) and further developed by Chew (1983) and Fishburn (1983). The final result of Chew-MacCrimmon analysis is the representation of preferences over lotteries

$$U(p) = \sum u(x_i)p_i / \sum v(x_i)p_i \tag{1.7}$$

where, u and v are two different elementary utility functions. Still other functional forms of weighted expected utility have been suggested by Kamarkar (1978), Kahneman and Tversky (1979). For a simple three-outcome case, this becomes

$$U(p) = \frac{[p1u(x_1) + (1 - p_1 - p_3)u(x_2) + p_3u(x_3)]}{[p_1v(x_1) + (1 - p_1 - p_3)v(x_2) + p_3v(x_3)]} \tag{1.8}$$

so, for any indifference curve, setting $U(p) = U^*$ it is easy to see that

$$u(x_2) - U^*v(x_2) = p_1[U^*(v(x_1) - v(x_2)) - (u(x_1) - u(x_2))] + \tag{1.9}$$

$$p_3[U^*(v(x_3) - v(x_2)) - (u(x_3) - u(x_2))] \tag{1.10}$$

The first thing to note is that the indifference curves remain linear as no probability terms enter the expression for the slope $\frac{dp_3}{dp_1}$. The second point is that all the indifference curves intersect at the same point prior to the origin in the lower left quadrant (at theoretically negative probabilities) because they do not depend on levels of U^*.

Contemporaneous work by Machina (1982) on non-linear expected utility was a further important

development of the concept of fanning out of indifference curves. Machina maintained the preference ordering axioms and the independence axiom, but did not appeal to the independence axiom and succeeded in demonstrating that fanning out was possible with the added benefit of non-linearity of the underlying indifference curves. The upward sloping property of the underlying indifference curves is achieved by using the idea of stochastic dominance. Continuing with the simple example above, remember that if outcomes are ranked $x_3 >_h x_2 >_h x_1$, then using any probability distribution, p, it is easy to see that a change in probability distribution either upwards or to the right implies an increase in the probabilities by which outcomes are weighted, as is shown clearly in figure ??, which provides an example of indifference curves implied by the Machina approach. As can be seen from the figure, his illustrates the principle of stochastic dominance, whereby distributions such as q or r stochastically dominate p, with the result that it is possible to produce a non-linear, but upward sloping indifference curve U, passing through point p[7].

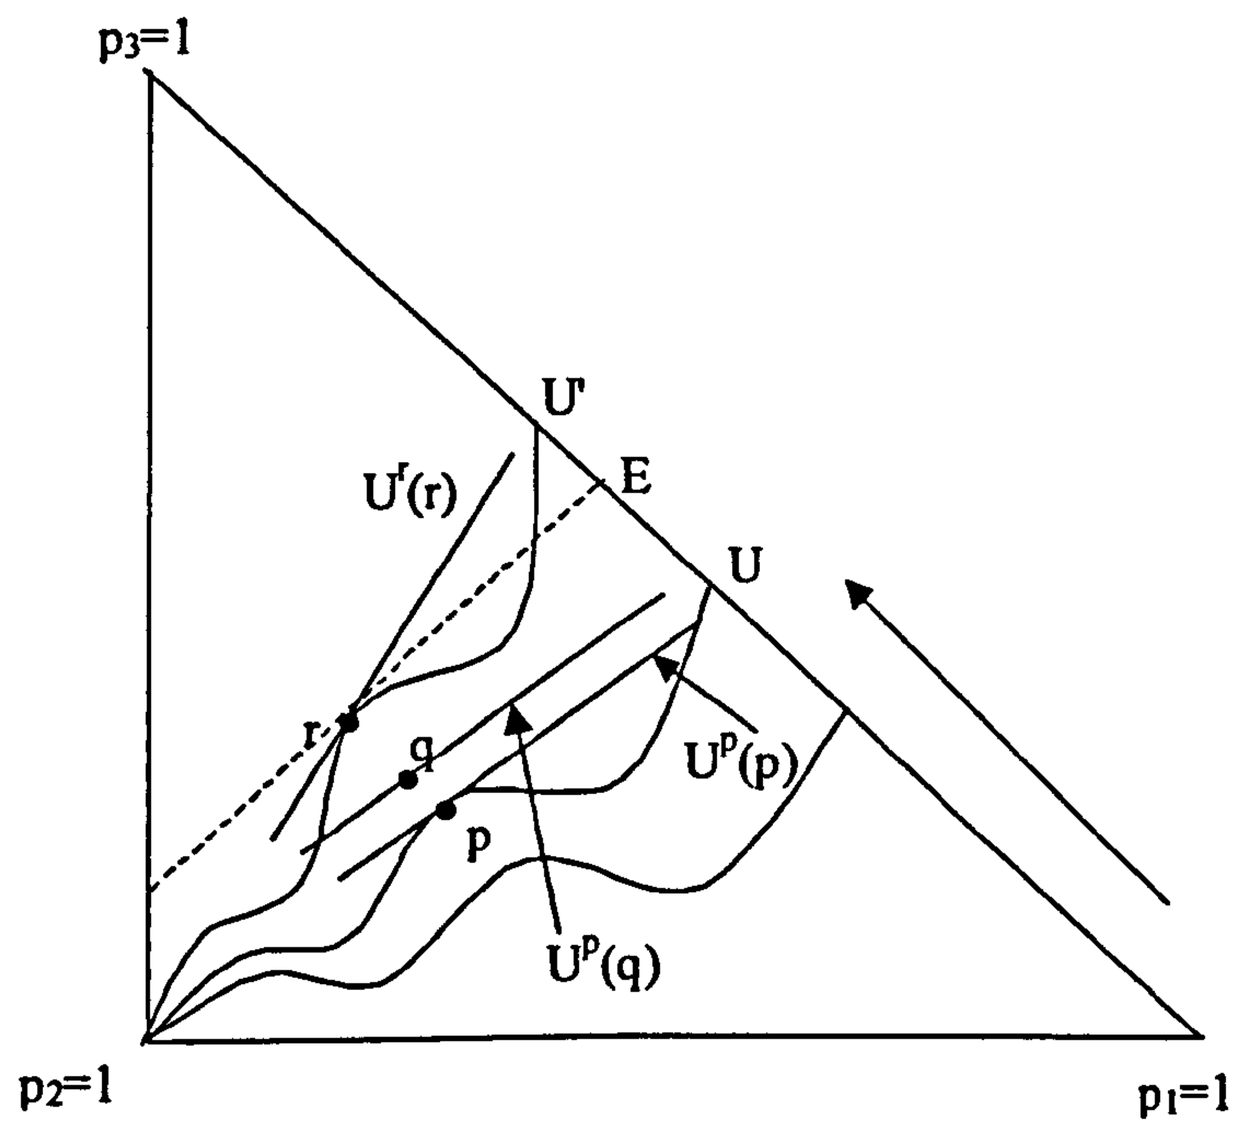

Figure1.1: Non-linear expected utility and fanning out

To obtain fanning out, Machina utilised the concept of "local expected utility" under which, at any given distribution p, there exists a utility function U_p, that possesses all of the von Neumann and Morgenstern properties of an actual utility function around p. The important point to note is that U_p is not a utility function over *all* distributions, but is specific or "local" to p, enabling it to be treated as a "local utility index". Machina therefore claimed that the behaviour around p of a decision maker acting with a general utility function U^P can be encapsulated by a local utility index $U_p(p)$. These local utility

[7]Stochastic dominance has two orders. First order occurs when there exist two random variables with cumulative distributions of say$F$ and G, such that if $F \geq_1 G \Rightarrow F(x) \leq G(x)$ for all $x \in [a,b]$. Whilst second order stochastic dominance helps rank distributions relative to riskiness in terms of the spread of the probability mass of the cumulative density functions and is expressed in analogous fashion as $F \geq_2 G$, if $T(x) = \int_a^x [G(t) - F(t)]\,dt \geq 0$ for all $x \in [a,b]$.

indexes are then used to prove two points. First, that every local utility function is increasing in x, which guarantees small stochastically-dominant shifts in the region of p (from p to q in figure ??) which is known to imply by linear expected utility that $U_p(p) < U_p(q)$, which also implies that $U(p) < U(q)$, so that the original indifference curves (which are non-linear) must rise from bottom right to top left.

The second point relates to fanning out and can obtained from Machina's second hypothesis. It is that local utility functions are concave in x for each distribution p, which means that movement towards stochastically-dominant distributions, causes the degree of local risk-aversion to increase, such that $-\left[\frac{U_{p''}(x)}{U_{p'}(x)}\right] < -\left[\frac{U_{r''}(x)}{U_{r'}(x)}\right]$ where r stochastically-dominates p - which is shown in figure ??. Remembering that "risk-aversion" implies that the indifference curve is steeper than any arbitrary contour value line (such as the dashed line E in figure ??). The local utility index around r, $U_r(r)$ is steeper than the local utility index around p, $U_p(p)$, making r more "risk-averse" than p. Machina therefore concludes that the Allais Paradox and all common consequence and common ratio problems can be explained by his simple set of axioms. However, one important point needs to be born in mind about work such as Machina's, namely, that it was designed to test for violations of the von Neumann and Morgenstern axioms, rather than to explicitly test the expected utility hypothesis.

The common consequence and common ratio effects as described in the Allais Paradoxes are seen as violations of the independence axiom, but a further common empirical finding is so called "preference reversal", which is a violation of the transitivity axiom. Evidence for such reversals was first uncovered in 1971 in the field of psychology, by Sarah Lichtenstein and Paul Slovic. The preference reversal phenomena is captured in the famous "P-bet, \$-bet" problem, which is framed as follows. Imagine two lotteries which can be described as follows:

P-bet: yielding outcomes (X, x) with probabilities $(p, (1-p))$.

\$-bet: yielding outcomes (Y, y) with probabilities $(q, (1-q))$.

X and Y are assumed to be large money amounts, whilst x and y are assumed to be very small, possibly negative, money amounts. The key point is that $p > q$ (so that the P-bet has higher probability of a large outcome) and that $Y > X$ (so that the \$-bet has the highest large outcome). The labeling of the bets therefore reflects the fact that those faced with the P-bet face a relatively higher probability of a relatively low gain, while in the \$-bet, a relatively smaller probability of a relatively high gain exists. For example:

P-bet: \$30 with 90% probability, and zero otherwise.

\$-bet: \$100 with 30% probability and zero otherwise.

The expected gain of the \$-bet is higher than that of P-bet and Lichtenstein and Slovic (1971, 1973 and others such as. Grether and Plott, 1983) have produced evidence that there is a tendency to choose

the P-bet over the \$-bet, yet there is also a willingness to sell the right to a P-bet for less than the right to engage in a \$-bet. In terms of risk-aversion, this means that when directly asked, the P-bet would be chosen, but that there was a willingness to accept a lower certainty-equivalent amount of money for a P-bet than for a \$-bet. In terms of the above example, this means that minimum selling prices would be \$25 for the P-bet and approximately \$27 for the \$-bet.

There have been numerous claims that this is tantamount to a violation of the transitivity axiom. Such claims are based on the argument that one is indifferent between the certainty-equivalent amount (which can be regarded as the "minimum selling price") of the bet or taking the bet. So, in utility terms $U(P\text{-bet}) = U(\$25)$and $U(\$\text{-bet}) = U(\$27)$. Using the simple principle of monotonicity, it can then be argued that more riskless money is better than less riskless money, so $U(\$27) > U(\$25)$, so forcing the conclusion that $U(\$\text{-bet}) > U(P\text{-bet})$. However, direct questioning usually reveals a preference for P-bet over \$-bet, which implies $U(P\text{-bet}) > U(\$\text{-bet})$, thereby demonstrating the intransitivity, which of itself is not necessarily a desirable result. Out of all the preference axioms it is generally agreed that transitivity represents the core of rationality, but it is interesting to note that modern general equilibrium theorists have apparently been able to eliminate it without creating difficulties in proving the existence of Walrasian equilibrium. (see Mas-Colell, 1974 for example). So the appropriate question would appear to be whether preference reversals observed in an experimental context do in fact reveal "intransitivity"?

The clearest doubt about this is contained in the work by Karni and Safra (1986, 1987) and Holt (1986), where it is indicated that experiment design may be affecting results, forcing the conclusion that it may not be intransitivity that is being observed, but an inability to identify certainty-equivalent amounts. The explanation being that decision makers often overstate minimum selling prices for lotteries for which there is less interest. The argument being that it is the independence and not transitivity that is being violated.

However, an alternative strand of thinking has tried to formulate an alternative reasoning for the existence of preference reversals in the form of non-transitive expected utility theory. The idea being to provide a rational underpinning for the apparently irrational aspect of preference results. The foremost among these explanations is regret theory suggested by Graham Loomes and Robert Sugden (1982, 1985), David Bell (1982) and Peter Fishburn (1982). The core of their argument is that choosing a bet not initially owned is a fundamentally different proposition to selling a bet that was not initially possessed. In terms of the above example, selling the \$-bet only to find that the buyer then wins the high-yield outcome, will engender more disappointment for having sold a subsequent winning bet than not having chosen the bet in the first place.

Based on this idea, Loomes and Sugden suggested a regret/rejoice function for pairwise lotteries that

have both the outcomes of the chosen and the foregone lottery. For example, suppose p and q are two lotteries, such that if p is chosen and q foregone and the outcome of p turns out to be x and the outcome of q is y, then the difference in the (elementary) utilities between the two outcomes can be taken as a measure of rejoice or regret

$$r(x,y) = u(x) - u(y) \tag{1.11}$$

which is negative in the case of regret and positive in the case of rejoice. The basic idea is therefore that faced with alternative lotteries the decision maker will not try to maximize expected utility, but instead to minimize expected regret (or maximize expected rejoicing). A simple example helps to illustrate the point. Imagine the following two lotteries with the following properties over the same set of outcomes $x = (x_1, ..., x_n)$

	Lottery p	**Lottery q**
Probabilities	$(p_1, .., p_n)$	$(q_1, .., q_n)$
Finite set of outcomes	$U(p) = \sum_i p_i u(x_i)$	$U(q) = \sum_j q_j u(x_j)$

(1.12)

The expected rejoice/regret function is therefore expressed as follows

$$E(r(p,q)) = \sum_i p_i u(x_i) - \sum_j q_j u(x_j) \tag{1.13}$$

$$= \sum_i \sum_j p_i q_j [u(x_i) - u(x_j)] \tag{1.14}$$

$$= \sum_i \sum_j p_i q_j r(x_i, x_j) \tag{1.15}$$

The obvious advantage of the regret/rejoice model is that the indifference curves over lotteries derived from it can be intransitive in so far as they yield up preference reversals. It is also intuitively appealing that minimizing expected regret is a valid criteria for rational choice. Importantly for this thesis, regret theory seems to be able to replicate fanning out while alternative expected utility theory cannot account for preference reversals, so it has been argued (see Sugden, 1986 for example) that regret theory is inherently more robust as a basis for decision making, but the Karni-Safra critique implies that such claims are far from conclusive.

Subjective expected utility

According to von Neumann-Morgenstern, probabilities were objective and in that they were simply assumed to exist. Three versions of this basic objectivist position exist. The classical view stated by de

Laplace (1795) argues that the probability of an event in a random trial is the number of equally likely outcomes that lead to that event divided by the total number of equally likely outcomes. This is based on two ideas, namely, the principle of cogent reason whereby physical symmetry is taken to imply equal probability and the principle of insufficient reason whereby equal probability should be assigned if it is impossible to know which outcome is more likely. The main problem with this approach is the meaning of symmetry and the possibly non-additive and the frequently non-intuitive results arising from applying the principle of insufficient reason. Subjective expected probability has therefore been challenged by many ideas, the most prominent being the relative frequentist concept put forward by von Mises in 1928 and applied by Reichenbach in 1949. The relative frequentist approach suggests that the probability of any given event in a particular trial is expressed in the relative frequency of occurrence of that event in an infinite sequence of similar trials and can be seen as an extension of Bernoulli's law of large numbers which claims that occurrence of a particular event a set number of times (k) in n identical and independent trials, then if the number of trials is arbitrarily large, such that k/n should be arbitrarily close to the objective probability of that event. The relative frequentists therefore claimed the high ground by defining probability as the limiting outcome of an experiment, independent of the existence of the notion of an objective probability. The idea of infinite repetition is clearly an idealization, but the idea caused problems for the objectivists. For how can it be possible to discuss the probability of the winners of events such as elections when they are deemed to be intrinsically unique ? As a result of this conundrum, some relative frequentists have opted to accept the inherent limitations of probabilistic reasoning to controllable situations, thereby conveniently confining unique random events to be not applicable. Yet others were unhappy with such a compromise on the scope of the applicability of probability reasoning. In attempt to reconcile these issues alternative directions have led to the work of Karl Popper (1959), resorting the propensity concept of objective probabilities as the way of explaining the tendency of Nature to yield a particular event on a single trial, independent of it being associated with any long-run frequency measure. These propensities are assumed to objectively exist, even if only in conceptually.

Statisticians and philosophers object to this view of probability, on the grounds that randomness is not an objectively measurable entity but rather a knowledge issue, making probabilities epistemological rather than ontological as an issue. Adherents of such a view, see a coin toss as not necessarily characterized by randomness, on the grounds that if factors such as shape and weight of the coin, the strength of the tosser, the atmospheric conditions of the room in which the coin is tossed, the distance of the coin-tosser's hand from the ground, etc., were all known in advance, then it would be possible to predict with certainty whether a head or tail will result from the toss. Such information is usually missing, so it is convenient to assume it to be a random event and assign probabilities to heads or tails. This implies that probabilities

are a measure of a lack of knowledge about conditions which might affect coin tossing, so that they represent prior beliefs about the coin tossing experiment. Knight expressed it neatly as follows

> **"...if the real probability reasoning is followed out to its conclusion, it seems that there is 'really' no probability at all, but certainty, if knowledge is complete".**
>
> Knight (1921, p. 219)

The knowledge based view of probability originates in the work of Bayes (1763) and de Laplace (1795) and can be divided into two broad groups, namely, the logical relationists and the subjectivists. The position of the former group was best set out in Keynes's Treatise on Probability (1921) and, later on, Rudolf Carnap (1950). Keynes (1921) argued that there is less subjectivity in epistemic probabilities than is commonly assumed as there is a sort of "objective" (although not necessarily measurable) relationship between knowledge and the probabilities that are deduced from such knowledge. For Keynes and logical relationists, knowledge is disembodied and not personal:

> **"In the sense important to logic, probability is not subjective. A proposition is not probable because we think it so. When once the facts are given which determine our knowledge, what is probable or improbable in those circumstances has been fixed objectively, and is independent of our opinion."**
>
> Keynes, 1921, p.4.

Ramsey (1926) fundamentally disagreed with Keynes, arguing instead that probability is related to the knowledge possessed by a particular individual alone. According to Ramsey, personal belief governs probabilities and not disembodied knowledge, thereby making probability subjective. Subjectivism in this form had been around since economists such as Irving Fisher (1906 and 1930) had first given it expression. But the main difficulty with the subjectivism is that it appeared impossible to derive precise mathematical expressions for probabilities from purely personal beliefs. If it is accepted that assigned probabilities are subjective, which implies that randomness itself is a subjective phenomenon, how can it then be possible to construct a consistent and predictive theory of choice under uncertainty? von Neumann and Morgenstern (1944) succeeded in this with objective probabilities, so the task was at least made manageable. But using subjective probability, that was far closer in meaning to Knightian uncertainty, made the task seem impossible.

Ramsey's seminal contribution in his 1926 paper (which was not published until after his death in 1931) was to put forward a method of driving out a consistent theory of choice under uncertainty that was capable of isolating beliefs from preferences, yet still maintaining subjective probabilities. Ramsey's

was the first attempt at an axiomatization of choice under uncertainty - more than a decade before von Neumann-Morgenstern's work. It is also worth noting that Bruno de Finetti (1931, 1937) had also (independently) suggested a very similar derivation of the subjective probability concept. To understand how subjective probability works, it is only necessary to consider problems like horse racing, where most wagering bets face more or less the same lack of knowledge about the horses, the track and the jockeys. Yet, while sharing the same knowledge (or lack of knowledge), there is a variety of different bets on the winning horse. The basic insight of the Ramsey-de Finetti work is that through observing bets, it is possible to infer the underlying personal beliefs about the outcome of the race. Ramsey and de Finetti therefore argued that subjective probabilities can be inferred from observation of people's actions as a sort of revealed preference or revealed belief theory of decision making.

Ramsay and de Finetti faced stiff competition from two other subjective probability theorists, Koopman (1940) and Good (1950, 1962), who both believed the Ramsey-de Finetti approach too constricting, because taken to the limit, the approach implies that a belief only qualifies as a belief if it is expressed through actual choice behavior. According to Koopman (1940) in contrast, the intuitive approach argues for probabilities having been derived directly from intuition and existing prior to objective experience. Subjective probabilities do not always necessarily have to be revealed through choice and even if they do, they usually occur in the form of intervals of upper and lower probabilities rather than precise point values, so that they are only partially ordered - an idea dating back to Keynes (1921, 1937), but most prominent in Shackle's work (1949, 1955, 1961).[8]

But probably more importantly, the intuitionists assert that not all choices reveal probabilities. That this may be so, can be seen from the fact that if the Ramsey-de Finetti analysis is taken to its logical extreme, choice behavior may reveal probability assignments that the owner had no idea about possessing ! Common examples of this abound, such as betting on horses in a race simply because the name of the horse appeals and not always because of a belief of winning - sentimentality can play a part. Adopting a Ramsey-de Finetti approach would imply that such a choice behavior would reveal a subjective probability assignment, even though the agent had actually made no such assignment or had no idea that he made one. It is therefore possible to assert that the implicit assumption underlying the Ramsey-de Finetti view is the existence of state-independent utility, which is dealt with later in this review (but see for example, Karni, 1996).

One final point to note is that one aspect of Keynes's (1921) work has resurfaced via the so-called Harsanyi Doctrine, or the common prior assumption (e.g. Harsanyi, 1968). This states that decision makers possessing the same knowledge, ought to have the same subjective probability assignments. Though

[8]It can be argued, that the Arrow-Debreu state-preference principle is an intuitionist view.

interesting, this notion is not stated either explicitly or implicitly in the subjective probability theory of either the Ramsey-de Finetti or intuitionists. Harsanyi grew mostly out of information theory and underpins parts of rational expectations theory - both of which enjoy tenuous links with uncertainty theory. Clearly, information theory cannot be joined to subjective probability too closely, because its rational is, to define an objective and deterministic linkage between information or knowledge and the choices made by decision makers, which makes it critical to filter out the idiosyncrasies permitted in subjective probability theory.

The axiomatisation and development into the full theory of subjective expected utility of the Ramsey-de Finetti view was achieved by Savage in 1954. Anscombe and Aumann's (1963) simpler axiomatization, however, has the advantage that it synthesised both objective and subjective probabilities into a single unified theory. The downside, however, is that the synthesised theory lost some of its generality during the synthesis process, such that their version of Savage's axiomatization of subjective expected utility theory is generally seen as an intermediate theory because it requires lotteries with objective probabilities. Anscombe and Aumann assume that an action f is no longer just a mapping from states S to outcomes X, but rather $f : S \rightarrow \Delta(X)$, where $\Delta(X)$ is the set of simple probability distributions on the set X, so that a consequence is no longer a particular value of x, but a distribution $p \in \Delta(X)$, so that the set of consequences $\Delta(X)$ are lotteries with objective probabilities, such that a result the components of the Anscombe and Aumann (1963) theory are the following:

S is the set of states

$\Delta(X)$ is the set of consequences (assumed to be objective lotteries on outcomes)

$f : S \rightarrow \Delta(X)$ is an action (e.g. a horse in a race, or a lottery combination)

$F = \{f \mid f : S \rightarrow \Delta(X)\}$ is the set of actions

$\geq_h \subset F \times F$ are preferences on actions

such that $\geq_h$are a binary relation on actions F that fulfill the von Neumann and Morgenstern axioms already described, where F is a mixture set such that for any $f, g \in F$ and for any $\alpha \in [0, 1]$, another element $\alpha f + (1-\alpha)g \in F$ can be associated, defined pointwise as $(\alpha f + (1-\alpha)g)(s) = \alpha f(s) + (1-\alpha)g(s)$ $\forall s \in S$. This is effectively a combination of subjective and objective probabilities (which behaves like a compound lottery but where the probabilities are unknown), with the payoff from the first lottery being a ticket for the second lottery. This means that the expected utility function is *state dependent* because the second lottery can only be played after the state $s \in S$ has occurred: $U_s : \Delta(X) \rightarrow R$, so that $U_s(f_s)$ is of the following form

$$U_s(f_s) = \sum_{x \in X} f_s(x_i) u_s(x_i) \tag{1.16}$$

This can be thought of as the expected utility of state $s \in S$ given that a particular action $f : S \rightarrow \Delta(x)$

has been chosen, so that assuming S is finite, gives

$$U(f) = \sum_{s \in S} U_s(f_s) \tag{1.17}$$

The important point to note is that the term on the right is not being multiplied by the probability that state s actually happens. This is because the probabilities are unknown. The Anscombe-Aumann approach is therefore a state-dependent expected utility representation of the utility of individual act, f.

Two key questions immediately follow from state-dependent utility. First, is the approach generalisable to preferences over actions; and second, is it possible to demonstrate state-independence ? Answering these questions requires the addition of three further axioms:

Axiom 5 (Null-states) *a state $s \in S$ is a null state if $(f_1, ..., f_{s-1}, p, f_{s+1}, ...f_n) \sim_h (f_1, ..., f_{s-1}, q, f_{s+1}, ...f_n)$ for all $p, q \in \Delta(X)$.*

Axiom 6 (Non-degeneracy) *there exists an $f, g \in F$ such that $f >_h g$, so that $>_h$is non-empty.*

Axiom 7 (State Independence) *let $s \in S$ be a non-null state and $p, q \in \Delta(X)$. Therefore: $(f_1, ..., f_{s-1}, p, f_{s+1}, ...f_n) >_h (f_1, ..., f_{s-1}, q, f_{s+1}, ...f_n)$, so that for every non-null state $t \in S$, the following can be stated:$(f_1, ..., f_{t-1}, p, f_{t+1}, ...f_n) >_h (f_1, ..., f_{t-1}, q, f_{t+1}, ...f_n)$*

Taken together, these three axioms make it possible to state that if $p >_h q$ at non-null state $s \in S$, then $p \geq_h q$ at any non-null state $t \in S$, thereby guaranteeing that the preference ranking between lotteries p and q is state independent. These axioms are also sufficient to be able to state Anscombe and Aumann's (1963) state independent expected utility representation:

Axiom 8 (Anscombe-Aumann) *suppose $S = [s_1, ...s_n]$ and that $\Delta(X)$ is a set of simple probability distributions, with $\geq_h$being a preference relation over the set $F = \{f \mid f : S \to \Delta(X)\}$. Then $\geq_h$satisfies the axioms above if and only if there is a unique probability measure π on S and a non-constant function $u : X \to R$ such that for every $f, g \in H$, it can be stated that*

$$f \geq_h g \text{ if and only if } \sum_{s \in S\pi} (s) \sum_{x \in X} f_s(x) u(x) \geq \sum_{s \in S\pi} (s) \sum_{x \in X} g_s(x) u(x) \tag{1.18}$$

The Anscombe-Aumann state-independent utility function is important because it allows the expression of preferences over actions via expected utility decomposition based on the utility of one action being greater than another. It is also important, because from it developed the work on state preferences and state contingent markets, to which we now turn.

State preferences

The state-preference approach to decision making in the presence of uncertainty was introduced by Arrow (1953) and further expanded by Debreu (1959), but was popularised by the work of Hirshleifer (1965, 1966) on the theory of investment and by Radner (1968, 1972) and others in finance and general equilibrium. The core of the idea is that it is possible to reduce choices under uncertainty to a more conventional indifference curve choice under certainty problem by an appropriate change in the commodity structure. It is this single feature that distinguishes the state-preference approach from the more microeconomic work on choice under uncertainty of von Neumann and Morgenstern (1944), because preferences are not directly formed over either simple or compound lotteries. Preferences are instead expressed with respect to state-contingent commodity bundles. In its reliance on states and choices of actions which are effectively functions from states to outcomes, it is therefore much closer in spirit to Savage (1954), but differs from Savage as it does not rely on the assignment of subjective probabilities (but such a derivation can be supported).

The fundamental premise of the state-preference approach to uncertainty is that commodities are differentiated by both their physical properties, location in space and time, but also by their state. To see that state is important, consider the simple example of an airline seat, whereby a seat on a flight from London to Paris on a Wednesday morning is a very different price from a seat on a London to Paris flight on a Friday evening. So, to use an abstract algebraic representation, if S is the set of mutually-exclusive states, then every commodity can be indexed by the state of nature in which it is consumed, enabling the construction of a set of state-contingent markets. Assume also that the number of states is n, giving nS commodities and thus nS prices, such that the commodity space X is a subset of R^{nS}. If x is the amount of commodity i delivered in particular state s and p is the price of commodity i in state s, then table 1 shows how state contingent markets for every good can be expressed in terms of a simple matrix of possibilities.

Table 1: State-contingent markets

← Goods →

		1	2	...	i	...	n	
↑	1	x_{11}	x_{12}	...	x_{i1}	...	x_{n1}	← x_1
States	2	x_{21}	x_{22}	...	x_{2i}	...	x_{2n}	← x_2
↓	...	...	...	...	...	...	...	
	s	x_{s1}	x_{s2}	...	x_{si}	...	x_{sn}	← x_s
	...	...	...	...	...	...	...	...
	S	x_{S1}	x_{S2}	...	x_{Si}	...	x_{Sn}	← x_s
$x_i : S \to R$		x_1	x_2	...	x_i	...	x_n	

[9]

The strength of this approach is that it is no longer necessary to rely solely on money outcomes. Choices can now be couched in terms of bundles of goods in any given state. For example, reading down a column gives the amounts of a good in alternative states, whilst reading across a row gives the amounts of each good in each state. In order to calculate the monetary value of bundles a set of state-contingent prices p_i, that will populate a matrix in analogous fashion to table 1. Decisions makers are therefore assumed to maximise their utility functions given a range of bundles of goods, which cost $p_{si}x_{si}$, available in alternative states which span the good space, $X \subseteq R^{nS}$. If preferences possess regular Arrow-Debreu properties over outcomes, X, then it is possible to define a quasi-concave utility function $U : X \to R$ representing preferences. Such preferences are not over lotteries in the sense of von Neumann-Morgenstern, but rather using state-contingent preferences, so that the idea of randomness is conveniently side-stepped, which interestingly, given that U is quasi-concave, means that the Arrow-Pratt concept of risk-aversion does not sit comfortably in this model.

State-preference theory can also be linked into Savage's (1954) subjective expected utility and the theories of risk-aversion, by simply assuming the existence of a state-independent utility function $u : C \to R$, that is. a real-valued mapping from the good space $C \subseteq R^n$, so that preferences over state-contingent goods can be summarized by the expected utility function

$$U(x) = \sum_{s \in S} \pi_s u(x_s) = \sum_{s \in S} \pi_s u(x_{s1}, x_{s2}, ..., x_{sn}) \tag{1.19}$$

such that the utility of a bundle of goods, x, is the sum-product of elementary (state-independent) utilities derived from state-contingent bundles of goods, $u(x_s)$, using subjective probabilities, $p_1, .., p_S$ as weights. Similarly to the Savage framework, a given p_s is the subjective likelihood of the emergence a particular

[9] It is assumed that references work row first column second. So that x_{12} means the amount of good 2 available in state 1.

state $s \in S$, so that a vector of subjective probabilities, $\pi = [\pi_1, ...\pi_S]$ where $\sum_{s \in S} \pi_s = 1$, summarise beliefs about the likelihood of occurrence of alternative sates. In analogous fashion, it is possible to see the link with the concept of risk-aversion by seeing that the relative quasi-concavity of the utility function U in state contingent goods represents the degree of risk-aversion.

Notwithstanding $U : X \rightarrow R$ captures beliefs about states and attitudes towards risk, it is not a necessary part of state-preference construction and the entire approach could proceed without using this assumption. It is also neither necessary nor desirable that preferences in this scenario be reconciled with the Savage (1954) axioms. To see that this is the case, consider the example that state-dependent utility may be required as a way of capturing the idea of random preferences. Then even if it were possible to extract subjective probabilities $\pi_1, .., \pi_S$, it could only be achieved through a decomposition along the lines of $U(x) = \sum_{s \in S} \pi_s u_s(x_s)$ where the important subscript s on the elementary utility function implies that utilities are themselves state dependent. In this case, the same good in a particular state is simply valued more by the consumer than the same good in another state independently of the probabilities of the states occurring.

However, assume for a moment that $U : X \rightarrow R$ has an expected utility construction with state-independent utility, then the individual optimum is defined by the following optimisation problem

$$\max U = \sum_{s \in S} \pi_s u(x_s) \tag{1.20}$$

$$s.t. \tag{1.21}$$

$$\sum_{s \in S} p_s x_s \leq \sum_{s \in S} p_s e_s \tag{1.22}$$

where $e_s = [e_{1s}, ..., e_{ns}]$ is a vector summarising outcomes in each state, which, using λ to represent the Lagrange multiplier, yields the Lagrangian

$$L = \sum_{s \in S} \pi_s u(x_s) + \lambda \left[\sum_{s \in S} p_s e_s - \sum_{s \in S} p_s x_s \right] \tag{1.23}$$

Differentiation with respect to every state-contingent good (assuming an interior solution) yields the following set of first order conditions

$$\frac{dL}{dx_{is}} = \pi_s u'(x_{is}) - \lambda p_{is}, \forall i = 1, ..., n \text{ and } s \in S \tag{1.24}$$

The budget constraint is fulfilled, so that $\sum_{s \in S} p_s e_s = \sum_{s \in S} p_s x_s$, so that as there is a single multiplier across the first order conditions, it implies that at the individual optimum, for any particular good i is given by

the condition that its expected marginal utility will be equated across states

$$\frac{\pi_1 u'(x_{i1})}{p_{i1}} = \frac{\pi_1 u'(x_{i1})}{p_{i1}} = ... = \frac{\pi_S u'(x_{iS})}{p_{iS}} \tag{1.25}$$

Arrow (1953) called this condition the fundamental theorem of risk-bearing. It is easy to see that if there was no expected utility decomposition, then the numerator of each equation would simply be $\frac{\partial U}{\partial x_{is}}$, which is the partial derivative of the original utility function with respect to state-contingent good x_{is}[10].

Figure 1.2 shows the individual optimum for a single commodity (which could be either money or consumption) and two states, $S = (1,2)$. Thus, a commodity bundle, in this case, is a pair of state-contingent goods, $x = (x_1, x_2)$ where x_1 is the amount of the good delivered in state 1 and x_2 the amount of the same commodity delivered in state 2. Assume that there exists a utility function $U : X \rightarrow R$ representing preferences over X and that U assumes the expected utility decomposition, such that $U(x) = \pi_1 u(x_1) + \pi_2 u(x_2)$, where π_1 and π_2 are the subjective probabilities of state 1 and 2 occurring (such that $\pi_1 + \pi_2 = 1$).Given that U is quasi-concave in X, then the upper contour set in Figure 1.2 is a series of convex indifference curves, with slope given by

$$\frac{dx_2}{dx_1}\,|_U = -\left(\frac{\pi_1}{\pi_2}\right) \cdot \left[\frac{u'(x_1)}{u'(x_2)}\right] \tag{1.26}$$

where the right hand side is the negative of the marginal rate of substitution of consumption in the two alternative states. So, when $x_1 = x_2$ (as at point c in Figure 1.2), $u(x_1) = u(x_2)$, so that the slope of the indifference curve is reduced to -(π_1/π_2), such that along the 45° (or certainty) line, the slope of each and every indifference curve is equal to -(π_1/π_2). The important point to note is that unless the decision maker assigns equal subjective probability assessments to both states, π_1/π_2 will generally not be equal to 1. This of course only holds true if it is possible to assume that decision makers have an underlying state-independent utility function. If this were not so and utility were state-dependent instead, then the slope of the indifference curve would be $dx_2/dx_1|u = -\pi_1 u_{1'}(x_1)/\pi_2 u_{2'}(x2)$, such that the reason that the slope of the indifference curve on the 45° might be different from ,1 could still be due to different beliefs about the probability of states occurring. But this may also be that true that decision makers have different assessments about the utility value of the same consumption in different states. In terms of the earlier airline seat example, it may simply be that a seat on a Friday is simply more valuable to the consumer than a seat on a Wednesday. Therefore, it may be that both beliefs and state-dependent preferences can together explain why the slope of the indifference curve is not equal to 1 on the 45° line.

[10] Alternatively, if subjective probabilities could be extracted, then it would be possible to retain state-dependent utility functions and so express the fundamental theorem of risk bearing througn $\frac{\pi_1 u_1'(x_{i1})}{p_{i1}} = \frac{\pi_2 u_2'(x_{i2})}{p_{i2}}$, etc.

But, by assuming that preferences are state-independent, as was required in order to obtain expected utility decomposition, then the only explanation that can rationally remain is differing beliefs.

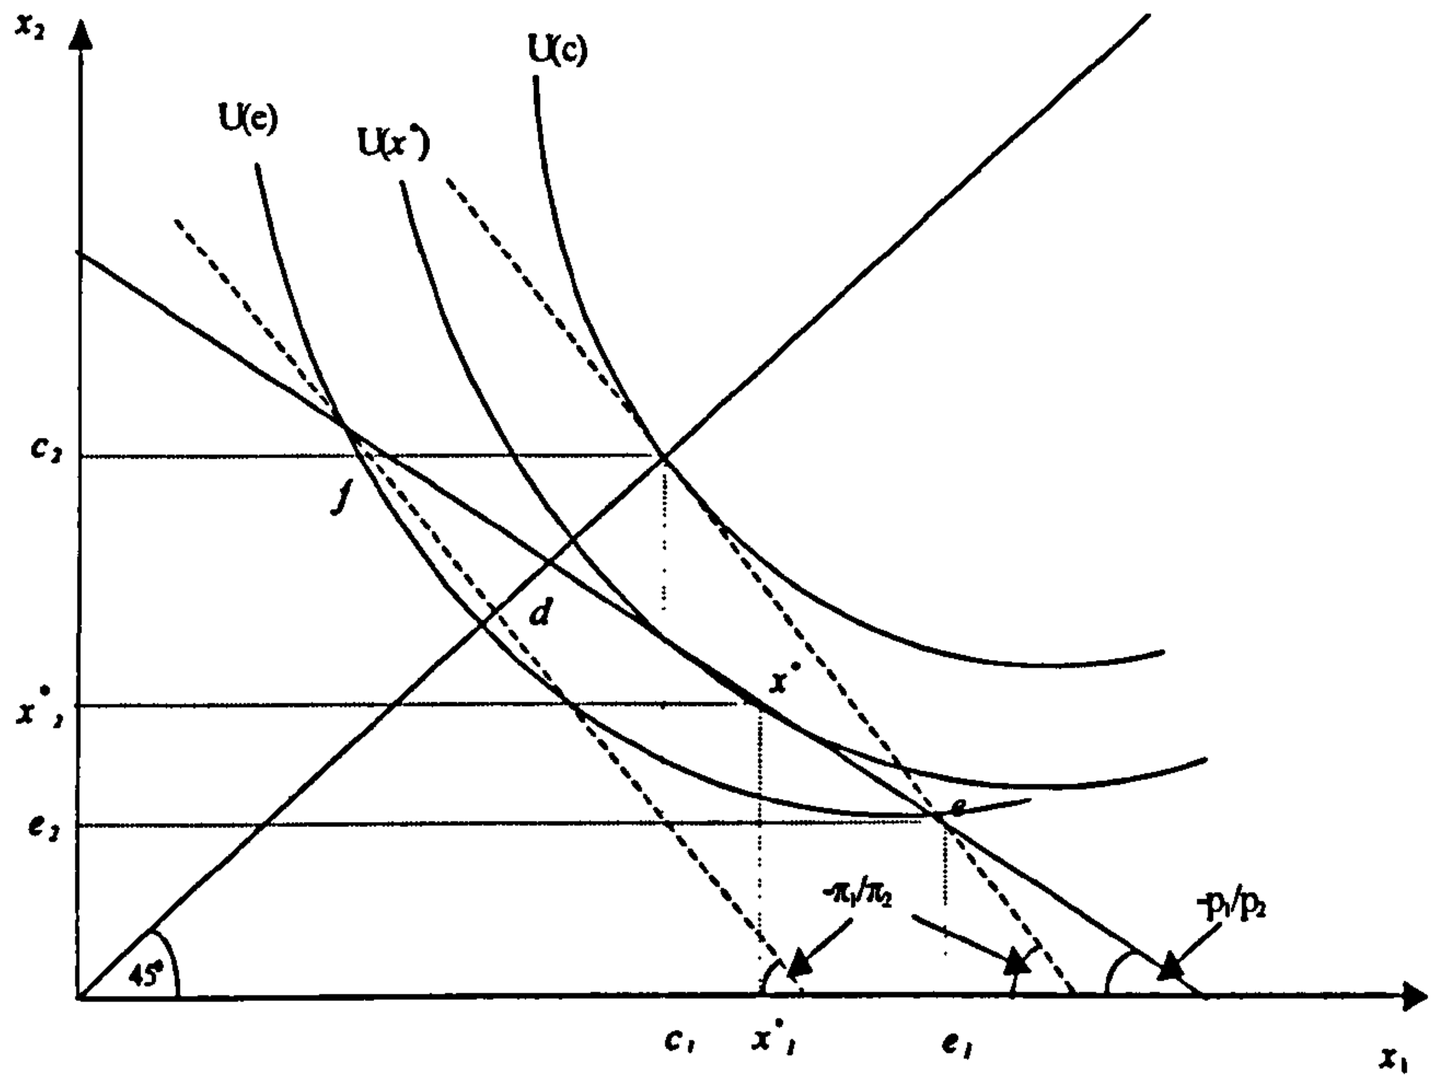

Figure 1.2: Optimum for individual decision maker

Given that the decision maker is assumed to possess a bundle of state-contingent goods, as shown by point $e = (e_1, e_2)$, shown in Figure 1.2 as point e. As $e_1 > e_2$, the endowment of the decision maker gives more of the bundle in state 1 than in state 2, such that if the decision maker consumes their endowment, utility level $U(e)$ will be achieved. If state-contingent markets exist, such that the decision maker can trade state-contingent goods, then let p_1 be the price of the good in state 1 and p_2 the price of the good in state 2 (which are assumed to be both known and given). The endowment and prices taken together define a budget constraint that passes through point e, as shown by the dotted line in Figure 1.2 with slope $-\pi_1/\pi_2$. The decision maker therefore faces the following optimization problem

$$maxU = \pi_1 u(x_1) + \pi_2 u(x_2) \tag{1.27}$$

$$s.t. \tag{1.28}$$

$$p_1 x_1 + p_2 x_2 \leq \pi_1 x_1 + \pi_2 x_2 \tag{1.29}$$

which yields the familiar Lagrangian and first order conditions for an interior solution

$$\frac{dL}{dx_1} = \pi_1 u'(x_1) + \lambda p_1 = 0 \tag{1.30}$$

$$\frac{dL}{dx_2} = \pi_2 u'(x_2) + \lambda p_2 = 0 \tag{1.31}$$

$$p_1 x_1 + p_2 x_2 = \pi_1 x_1 + \pi_2 x_2. \tag{1.32}$$

where λ is the Lagrangian multiplier. When the first two partial differentials are combined it can be seen that

$$-\frac{\pi_1}{\pi_2} \cdot \frac{u'(x_1)}{u'(x_2)} = -\frac{p_1}{p_2} \tag{1.33}$$

so that the decision maker selects the optimal bundle of state-contingent goods where the highest indifference curve is at a tangent to the budget constraint. The bundle of goods $x^* = (x_1^*, x_2^*)$ yielding utility $U(x^*)$ in Figure 1.2 is the individual optimum in this simple two-state example. The fundamental theorem of risk-bearing holds because the first order conditions imply that $\pi_1 u'(x_1)/\pi_1 = \pi_2 u'(x_2)/\pi_2$ which means expected marginal utility per monetary unit is the same for all states. As x^* does not fall on the certainty line $-x_1^* > x_2^*$, different amounts will be received in different states. The decision maker could have bought a bundle of goods containing no risk, e.g. d on the certainty line, but an indifference curve passing through d would produce utility lower than $u(x^*)$.

Figure 1.2 shows that this happens because the relative prices of the state-contingent goods do not match up with the subjective assessment by the decision maker of the likelihood of occurrence of both states, which can be seen from the fact that $\pi_1/\pi_2 > p_1/p_2$. This implies that market prices of goods in different states are not fair, which can be easily confirmed by supposing that probability assessments are $\pi_1 = 0.75$ and $\pi_2 = 0.25$ while the market prices are $p_1 = 0.5$ and $p_2 = 0.5$, so that $\pi_1/\pi_2 = 3 > 1 = p_1/p_2$. The decision maker therefore believes the probability of state 1 is much greater than that of state 2, yet the market generates the same prices for goods in both states. This in turn occurs because the decision maker is endowed mostly with state 1 good, so that a lower price results from selling it than if the market shared his probability assessments. Consequently, the decision maker will not sell most of good 1 and will move to a position, x^*, where he still has a random outcome. If prices coincided with the probability beliefs of the decision maker, such that the budget constraint had slope $-\pi_1/\pi_2$, then the budget constraint would be the dotted line passing through endowment e in Figure 1.2. In such a case, the consumer optimum is easily seen to be $c = (c_1, c_2)$ and the decision maker would move onto the 45° certainty line, thereby achieving a much higher level of utility $U(c)$ than previously.

The above analysis does not however mean that a decision maker would *always* prefer probabilistically fair prices and there will generally be both gainers and losers in most market situations, as can be seen from

the following simple example. Imagine that the decision maker had an endowment at point f in Figure 1.2, facing unfair prices (p_1, p_2) but keeping the same beliefs (π_1, π_2). In such a situation the decision maker would shift to the optimum $x^* = (x_1^*, x_2^*)$, resulting in a lot of state 2 good to sell. However, the market values state 2 good more than the decision maker thinks it likely that state 2 will occur, so that the decision maker is getting a good buy at the "unfair" prices. If prices were "fair" in this situation, then the budget constraint would be a line passing through f with slope $-\pi_1/\pi_2$ in Figure 1.2. In such a case, the decision maker would move to the certainty allocation at point d, so that it is possible to see that the utility achieved in this case will be lower than $U(x^*)$. Therefore, a decision maker with an endowment at f would lose if the market prices were made fair; because there would be a preference that the prices be kept at their "unfair" rate of $-(p_1/p_2)$.

Although a simple example, Figure 1.2 can be used to make two further points. The first, is that if a decision maker starts from a position of certainty (i.e. on the 45° line) and is offered fair prices, he will remain on the 45° line. This can be clearly seen by considering the decision maker starting at allocation c,with fair prices at $-\pi_1/\pi_2$ so that the budget constraint is the dashed line. $U(c)$ is therefore the highest utility the decision maker can achieve. This implies that starting from a position of certainty, a decision maker will not accept or seek "fair bets", but will undertake unfair bets if the odds are perceived to be advantageous (as in the case of an agent starting at d in Figure 1.2 and then being offered unfair prices which took him through to, say, x^*). The second point is that if a decision maker begins in an uncertain situation (such as e or f in Figure 1.2), then movement to a position of certainty occurs if the price is fair and a risky asset is purchased as a means of making the move! Note that the purchase of the risky asset is only undertaken to offset (or hedge from a trading perspective) the riskiness of the original endowment in order to provide certainty. One caveat is that the decision maker may not necessarily move to certainty position if the odds are unfair, but might optimally choose a risky situation, usually where the probabilities are perceived as favorable, such as a move from e or f to x^*.

Characterisation of risk aversion

It is clear that the shape of an indifference curve can reflect the subjective assessments of the probabilities of different states, but an equally interesting question is how an indifference curve can reflect "risk-aversion". Arrow (1965) and Pratt (1964) characterize risk-aversion by the concavity of the utility function over nominal income. The indifference curves in the simple two-state case of Figure 1.2 are a contour mapping of a quasi-concave utility function over a single commodity with two possible payoffs, which does not provide a clear way of seeing risk-aversion. However, if concavity is imposed on the utility function, thus forcing risk-aversion, such an assumption would also translate into the convex indifference curves in

the space of state-contingent goods as shown in Figure 1.2.

The theory of risk aversion using the state-preference approach was introduced by Yaari (1969). Figure 1.3 provides an example Yaari's approach, for a simple case of two agents, U and V, with two types of indifference maps. Given that the indifference curves U are more "convex" than the indifference curves V,does the extra convexity imply that U is more risk averse than V ? There are two answers to this and the first uses the risk-premium framework similar to Arrow-Pratt, whereas the second involves a new definition of risk-aversion. Using Figure 1.3, let $E = \pi_1 x_1 + \pi_2 x_2$ denote a particular level of (subjectively) expected return, so that the line labelled E with slope $-\pi_1/\pi_2$ represents all the combinations of good in state 1 and good in state 2 that yield identical expected return E. Analogously, let $F = \pi_1 x_{1'} + \pi_2 x_{2'}$ be another level of expected return, such that the line F also has slope $-\pi_1/\pi_2$ represents all the combinations of goods in both states that yield the same expected return F. Given that the line E is always below the line F, then every allocation of goods on E provides a lower expected return than any allocation of goods on F.

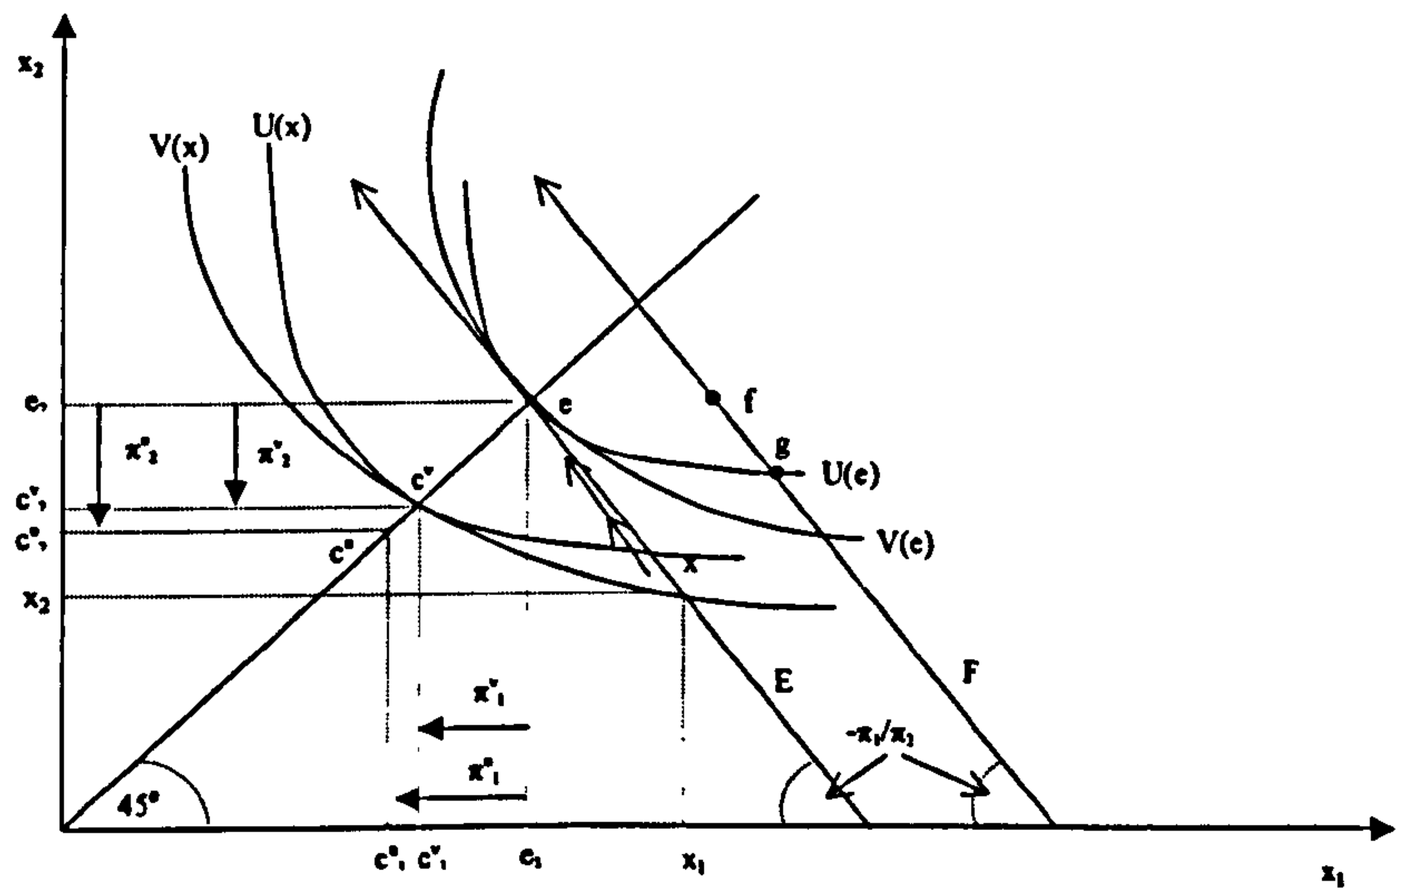

Figure 1.3: Risk Aversion in the Yaari model

Now examine the initial risky allocation $x = (x_1, x_2)$ in Figure 1.3, at which, there are utility levels $U(x)$ and $V(x)$ for the two decision makers. As the objective is to obtain risk premia, certainty-equivalent allocations can easily be found by moving along the indifference curves to the 45° certainty line. Allocation $c^u = (c_1^u, c_2^u)$ on the 45° line gives the same utility $U(x)$, as allocation x, so that c^u measures the decision maker U's "certainty-equivalent" allocation. $c^v = (c_1^v, c_2^v)$ similarly gives the same utility to the second decision maker $V(x)$, as was obtained at x, so that c^v is V's "certainty-equivalent" allocation. Calculation of the risk premia requires calculation of the amount of each commodity each decision maker would be willing to forgo to move from allocation x to the certainty-equivalent allocations c^u and c^v. Both decision makers give up a lot of good 1, but may demand only a small and variable amounts of good 2 by way of

compensation. It is clear that as allocations e and x lie on the same curve E, they share the same expected return, so one possibility would be to traverse the E line from x to e and then calculate the risk-premium as the amounts of goods required to move from e to c^u and c^v respectively. The "risk-premium" decision makers U and V would each have to pay would therefore be bundles of goods given by

$$\pi^u = (\pi_1^u, \pi_2^u) = (e_1 - c_1^u, e_2 - c_2^u) \tag{1.34}$$

$$\text{and} \tag{1.35}$$

$$\pi^v = (\pi_1^v, \pi_2^v) = (e_1 - c_1^v, e_2 - c_2^v) \tag{1.36}$$

Arrow-Pratt would argue that if any of the components in the risk-premia π^u or π^v are positive (or at least none are negative), then the decision maker is risk-averse. It is easy to see that both U and V are risk-averse as both π^u and π^v are positive in their elements. Similarly, a risk-neutral decision maker would exhibit a zero risk-premium in both elements, in which case, there would have to exist a linear indifference curve that passes through both points e and x, so that e would be the certainty-equivalent allocation (the converse argument holds for the risk-loving decision maker). This can also be used as a measure as it is easy to see that the more risk-averse decision maker U pays a higher premium (in both elements) than V, making U more risk-averse than V. Yaari's (1969) criteria for risk-aversion is therefore defined as follows

> U is more risk-averse than V if, beginning from the same allocation, the set of risky allocations acceptable to U is a subset of the set of risky allocations acceptable to V.

Figure 1.3 illustrates the implications of this via points f and g on the F line, because as $E < F$, then f and g give a higher expected return than e. Both U and V would reach a higher indifference curve (and therefore higher expected utility), than if given f, so that both U and V would accept the risky allocation f instead of e. But the more interesting point is g. At g, decision maker V would clearly obtain a higher level of utility at point g than e, so allocation g would be accepted. It is obvious that U would have lower expected utility at g than at e, so that U would not accept the risky allocation of g. Therefore, starting from point e, there must be at least one risky allocation that V would accept but U would not. From which, it can immediately be seen from the Figure 1.3 that the set of risky allocations acceptable to U forms a strict subset of the set of risky allocations acceptable to V - thereby establishing via the Yaari definition of risk-aversion, that U is more risk-averse than V.

It is therefore clear that given the existence of a state-independent utility function, the three ideas can be linked together in the following loose fashion. Namely, that U is more risk-averse than V if

1. The indifference curves of U are "more convex" than V.

2. That the risk-premium bundle paid by U is greater than that paid by V.

3. That V will accept risky allocations that U will not accept.

More interesting than this trivial characterisation of Yaari's (1969) work, is the characterisation of the degree of risk aversion with respect to changes in wealth, which can be traced out as a wealth-expansion path as shown in Figure 1.4 by the curve OE. Recalling that wealth is represented by initial endowment, then assuming prices are actuarially fair, such that market prices and subjective beliefs coincide, then at any particular level of endowment $e = (e_1, e_2)$, we can define total wealth as $W = \pi_1 e_1 + \pi_2 e_2$ which serves as our budget constraint. This is shown in Figure 1.4 by line W with slope $-\pi_1/\pi_2$.

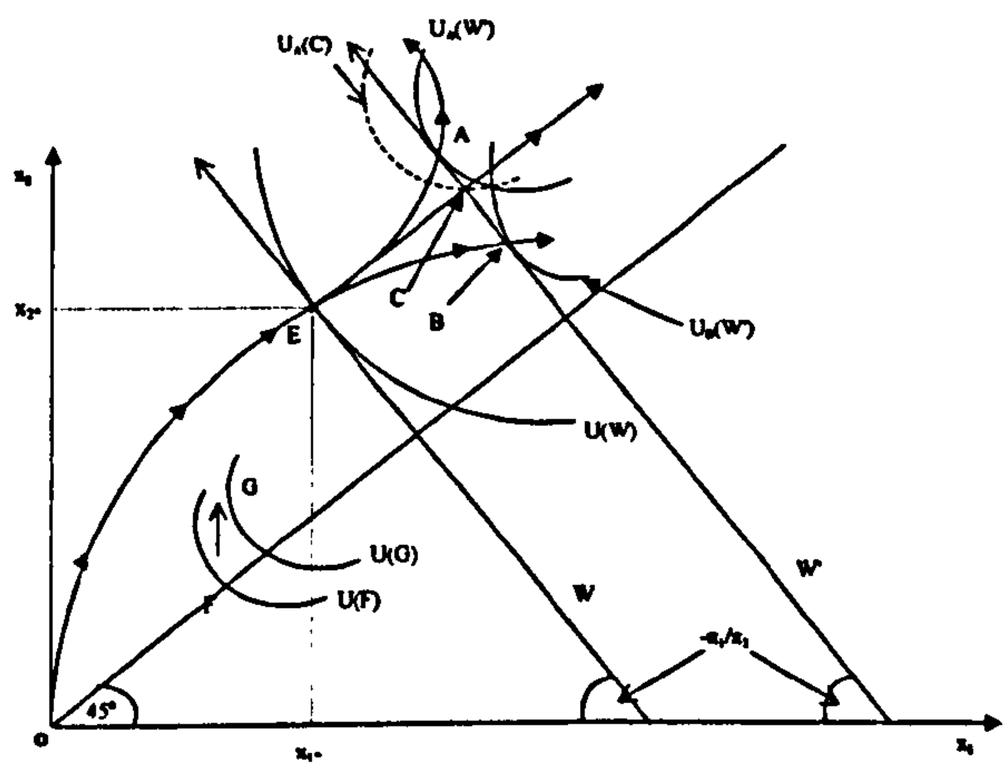

Figure 1.4: Wealth Dynamics and Risk Aversion

If initial wealth increases to $e' = (e_{1'}, e_{2'})$, with prices remaining constant, then wealth will rise to $W' = \pi_1 e_{1'} + \pi_2 e_{2'}$, so giving a higher budget line W' with slope $-\pi_1/\pi_2$. In familiar fashion, the optimum for the decision maker occurs at the point of tangency between the highest indifference curve and budget constraint, so that when wealth is W, then $E = (x_1^*, x_2^*)$ is the optimum utility level $U(w)$. The line OE traces the wealth expansion path. OE is constructed by tracing the individual decision maker optima at varying levels of wealth. The concept of decreasing/increasing rates of risk-aversion with respect to wealth can be seen in this framework, by considering what happens to wealth after point E.

There are clearly many possible paths for wealth beyond E, as shown be points A, B and C. Figure 1.4 shows points A and B as possible specific individual optima for a wealth level W', which yield indifference curves $U_A(w')$ and $U_B(w')$ respectively. An important point to observe is that the middle path EC is parallel to the 45° line, implying that the path EA is moving away from the 45° certainty line (increasingly risky) and the path EB is moving towards the 45° certainty line (increasing certainty). It is therefore tempting to say that if the wealth expansion path follows the EA track, it is decreasing absolute risk-aversion whereas if wealth follows EB, then increasing absolute risk-aversion is indicated, whilst if wealth follows EC track constant absolute risk aversion is implied.

This is easy to see, because as W is parallel to W', then simple geometry implies that the slopes of the indifference curves at the optimum when wealth is W' are the same as at E, so that if A is the new optimum at W', then the slope of $U_A(w')$ at A must be the same as $U(w)$ at E. Whereas if B were to be the new optimum on W', then the slope of $U_B(w)$ at B is the same as $U(w)$ at E. Recalling that the marginal rate of substitution (given by the inverse of the slope of the indifference curve) at point E is given by $\pi_1 u'(x_1^*)/\pi_2 u'(x_2^*)$, then differentiating with respect to the logs of this term, gives the following expression

$$dln(MRS) = \left[\frac{u''(x_1^*)}{u'(x_1^*)}\right] dx_1 - \left[\frac{u''(x_2^*)}{u'(x_2^*)}\right] dx_2 \tag{1.37}$$

which, given that the EC line is parallel to the 45° line, means that $dx_1 = dx_2$ as we move along EC. This implies that along EC

$$dln(MRS) = \left[\frac{u''(x_1^*)}{u'(x_1^*)} - \frac{u''(x_2^*)}{u'(x_2^*)}\right] dx \tag{1.38}$$

Now assume that there is an increase in wealth that moves the decision maker along EA to point A, then Figure 1.4 shows that if $U_A(w')$ is the optimum indifference curve, then the lower, dashed-line indifference curve $U_A(C)$ shows the indifference curve at point C if the line EA marks the true wealth expansion path. As the marginal rate of substitution at point A is clearly the same as that at E, then what is the implied rate at C? Logically, it must be less due to declining marginal rates of substitution

$$dln(MRS) = \left[\frac{u''(x_1^*)}{u'(x_1^*)} - \frac{u''(x_2^*)}{u'(x_2^*)}\right] dx < 0 \tag{1.39}$$

so that the marginal rate of substitution must have declined in the move from E to C, which implies that

$$\frac{u''(x_1^*)}{u'(x_1^*)} > -\frac{u''(x_2^*)}{u'(x_2^*)} \tag{1.40}$$

Given that $x_2^* > x_1^*$, then if consumption of state 2 good exceeds consumption of state 1 good, the Arrow-Pratt measure of absolute risk aversion of state 2 good is less than the absolute risk-aversion in state 2. However, remember that state 2 good is the same money good as state 1 good, so $x_2^* > x_1^*$ implies that shifting between x_1 to x_2 implies increasing wealth. Therefore, the result given by equation 1.40 implies that as wealth increases, the rate of absolute risk-aversion decreases. The dynamics of the wealth-expansion path EA therefore imply that there is decreasing absolute risk aversion (DARA) in the sense of Arrow-Pratt. It is also easy to see that if EB and not EA was the true description of wealth-expansion then the decision maker would be exhibiting increasing absolute risk-aversion (IARA). Finally, if EC was the true wealth expansion path, then the decision maker would be exhibiting constant absolute risk-aversion (CARA).

The more interesting question is how this notion of risk-aversion can be connected with the older one of increasing convexity indifference curves? The simple answer can be seen by supposing that the decision maker is at a point on the certainty line, such as point F with utility $U(F)$ in Figure 4, which would imply a marginal rate of substitution of

$$MRS_F = \frac{\pi_1 u'(x_1^*)}{\pi_2 u'(x_2^*)} = \frac{\pi_1}{\pi_2} \tag{1.41}$$

To map out the dynamics, keep x_1 constant and let x_2 rise, so that $dx_2 > 0$ and $dx_1 = 0$. This gives a vertical shift from F to G in Figure 1.4 and the convexity of the indifference curve implies that the MRS must rise, so that MRS_G must be greater than MRS_F, which in turn gives

$$dln(MRS) = -\left[\frac{u''(x_2^*)}{u'(x_2^*)}\right] dx_2 > 0 \tag{1.42}$$

so that the rise in dx_2 generates an increase in the marginal rate of substitution. For $dx_2 > 0$ to be correct, it must be that $-u''(x_2^*)/u'(x_2^*) > 0$, giving positive risk-aversion. If, however, the marginal rate of substitution did not change with the increase in x_2 (as would be the case if the indifference curves were linear) then

$$dln(MRS) = -\left[\frac{u''(x_2^*)}{u'(x_2^*)}\right] dx_2 = 0 \tag{1.43}$$

such that as $dx_2 > 0$, then $-u''(x_2^*)/u'(x_2^*) = 0$, giving risk-neutrality. The final point to note is that if there is a decision maker v whose marginal rate of substitution rises even more from the same change in dx_2, then it is possible to see that

$$\left[\frac{v''(x_2^*)}{v'(x_2^*)}\right] dx_2 > \left[\frac{u''(x_2^*)}{u'(x_2^*)}\right] dx_2 > 0 \tag{1.44}$$

Which clearly means that v displays a greater degree of risk-aversion.

A state preference example of why robustness is important

This chapter has so far analysed in detail the standard approaches to the way in which decisions are made in the presence of uncertainty. The clear question that arises from the foregoing analysis is how well does the utility maximisation paradigm explain decision making in the presence of uncertainty in the real world? To try to answer this question, recourse can conveniently be made to insurance, as an obvious application of the state-preference approach. Insurance is an explicit state-contingent contract paying an indemnity to the insured, if and only if, a particular event is deemed to have occurred - a situation analogous to conventional expected utility analysis. The state-preference approach to the problem of optimal insurance

was originally developed by Arrow (1963, 1965), Eisner and Strotz (1963), Borch (1968). The simplest formulation is posed as a two-state model with a fixed premium per monetary unit of insurance coverage, γ with the set of states being defined as $S = \{A, N\}$, where A is an event state whilst in N no event occurs. If income is $w = \{w_A, w_N\}$, then wA represents wealth if an event occurs, whilst wN measures wealth in the case of no event occurring. Given that $w_A < w_N$, then the loss incurred if a trigger event occurs will be $w_N - w_A > 0$, with the consequence, as shown in Figure 1.5, that state-dependent wealth, w, lies below the 45° certainty line. Assuming the existence of a state-independent utility function that is defined over payoffs, gives the following representation of the utility of the decision maker at the endowment point

$$U(w) = \pi_s u(w_A) + (1 - \pi_s)u(w_N) \tag{1.45}$$

which is shown in Figure 1.5 as $U(w)$, where π_s is the subjective probability that an insured event will happen, so that $(1 - \pi_s)$ is the probability that the insured event will not occur.

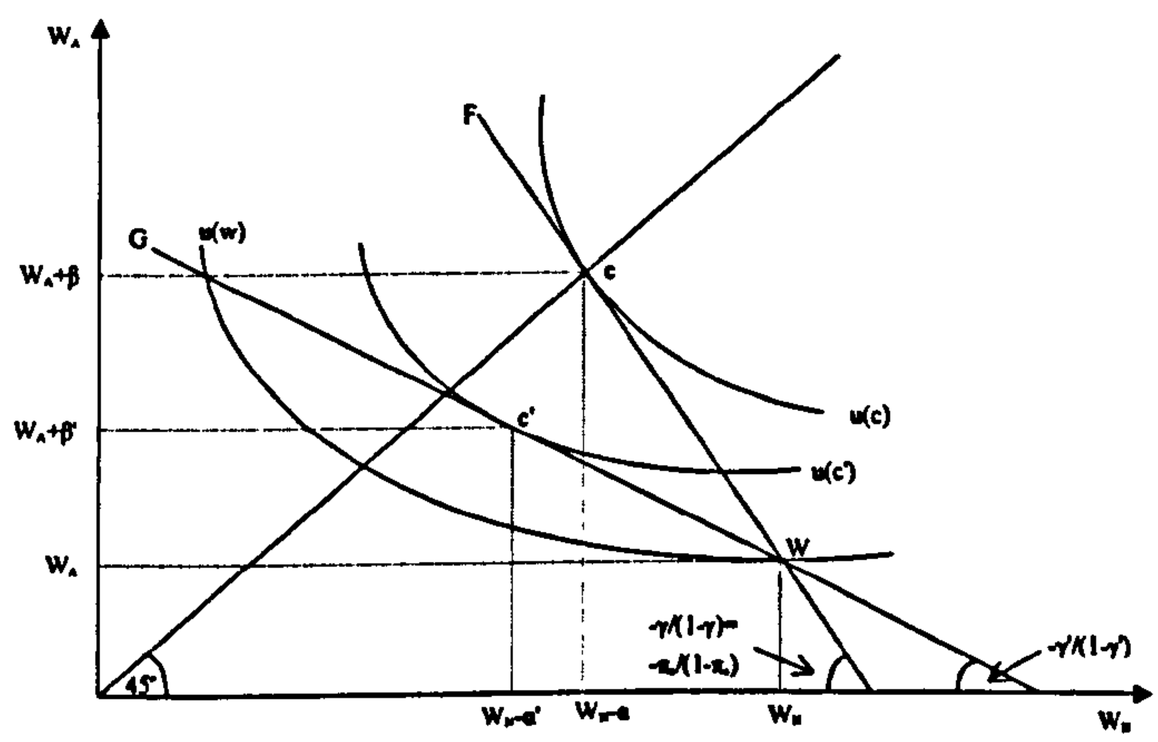

Figure 1.5: State preference and optimal insurance

An insurance contract can be summarised by the payoff function: $c = (\beta, \alpha)$ where α is the premium payment if no insured event occurs and β is the net indemnity if an insured event occurs, which means that if a decision maker purchases insurance expected utility is

$$U(w, c) = \pi_s u(w_A + \beta) + (1 - \pi_s)u(w_N - \alpha) \tag{1.46}$$

Neither α nor β are constant, but depend on C, the total amount of insurance cover chosen by the decision maker. Letting the total premium paid be proportional to the insurance cover, gives $\alpha = \gamma C$ where $\gamma \in [0, 1]$ is the premium per unit of insurance cover. The net amount paid (usually referred to as the net indemnity) if an insured event occurs is $\beta = C - \gamma C$, so that it is easy to see that the expected profits of the insurance company (assuming there is only one customer, or type of customer) are

$(1-\pi_s)\alpha - \pi_s\beta$.

To adapt the insurance concept into the subjective probability approach, begin with the simple assumption of zero profits, which gives $(1-\pi_s)\alpha - \pi_s\beta = 0$, which when rearranged gives $\frac{\pi_s}{(1-\pi_s)} = \frac{\alpha}{\beta}$, so the ratio of premium paid by the decision maker to his net indemnity is equal to the subjective probability of occurrence of an insured event. Replacing $\beta = (1-\gamma)C$ and $\alpha = \gamma C$, implies that $\frac{\pi_s}{(1-\pi_s)} = \frac{\gamma}{(1-\gamma)}$, which in turn implies $\pi_s = \gamma$, or simply put, the premium per monetary unit of insurance cover is equal to the subjective probability of an insured event. This "fair insurance" line, F, passing through w in Figure 1.5 therefore represents the series of insurance contracts where $\frac{\gamma}{(1-\gamma)} = \frac{\pi_s}{(1-\pi_s)}$ for different degrees of cover and acts as the budget constraint for the decision maker.

In terms of making an optimal decision, the decision maker's objective is to find the optimal amount of insurance cover C, given some unit monetary amount of cover γ, which implies the following optimization problem

$$maxU(w,c) = \pi_s u(w_A + C - \gamma C) + (1-\pi_s)u(w_N - \gamma C) \tag{1.47}$$

which upon application of the appropriate Lagrangian gives the first order condition

$$\frac{\partial U(w,c)}{\partial C} = \pi_s u'(w_A + C - \gamma C)(1-\gamma) + (1-\pi_s)u'(w_N - \gamma C)(-\gamma) = 0 \tag{1.48}$$

which after some simple algebra gives

$$\frac{u'(w_N - \gamma C)}{u'(w_A + C - \gamma C)} = \left[\frac{\pi_s}{(1-\pi_s)}\right]\left[\frac{(1-\gamma)}{\gamma}\right] \tag{1.49}$$

Then if insurance is "fair" such that $\gamma = \pi_s$, the above equation reduces to

$$\frac{u'(w_N - \gamma C)}{u'(w_A + C - \gamma C)} = 1 \tag{1.50}$$

which means that the marginal utility of an event state is equal to that of a non-event state. With state-independent utility, this implies that $w_A + C - \gamma C = w_N - \gamma C$, which of course implies that $C = wN - wA$, such that the decision maker takes full insurance cover so that his entire income loss from an insured event is recovered. The optimal insurance cover is illustrated in Figure 1.5 by the endowment $c = (w_N - \alpha, w_A + \beta)$,where the highest indifference curve $U(c)$ is tangent to the fair insurance line F on the 45° certainty line.

The full insurance cover result depends crucially on the assumption that $\gamma = \pi_s$ (so called "fair insurance"). However, if the insurer were to decide to make super-normal profits, such that $(1-\pi_s)\alpha - \pi_s\beta > 0$, then this would imply that $\gamma > \pi_s$, so that the premium per unit of insurance cover exceeds the

probability of an insured event. From the perspective of the decision maker, this is "unfair" insurance and is captured by the unfair insurance line G with slope $\frac{\gamma'}{(1-\gamma')}$ in Figure 1.5. Therefore, in the unfair insurance case , the first order condition for the decision maker implies that

$$\frac{u'(w_N - \gamma C)}{u'(w_A + C - \gamma C)} = \left[\frac{\pi_s}{(1-\pi_s)}\right]\left[\frac{(1-\gamma)}{\gamma}\right] < 0 \tag{1.51}$$

such that the marginal utility of a non-event state is less than the marginal utility of an event state. Given the assumption of quasi-concave utility, this implies that the utility of the decision maker in a non-event state exceeds utility in an event state, which implies that the decision maker must still be incurring some degree of loss in the case of an event such he cannot possibly be taking full insurance cover. In the case of a single decision maker, the optimum under "unfair" insurance is shown in Figure 1.5 at point c', the tangency of the unfair insurance line G with the highest indifference curve $U(c')$. Clearly, the decision maker does not have full insurance cover as he is off the 45° certainty line, so that the loss is not fully covered with $w_A + \beta' < w_N - \alpha'$.

To understand the implications of uncertainty in the case of unfair insurance, it is interesting to ask whether more risk-averse decision makers take more insurance rather than less risk-averse decision makers? Let u and v be two decision makers, where u is more risk averse than v, but both have the same subjective probabilities of an insured event and the same loss. Figure 1.6 shows the greater risk-aversion of u by the greater convexity of u's indifference curve relative to v's. With fair-insurance, such that $\pi_s = \gamma$ the decision maker u is on the F line, so that both u and v will have full coverage at c with $u(c)$ and $v(c)$. With unfair insurance, the decision makers are on the G line, so neither u or v takes full insurance. Figure 1.6 shows u's optimal contract will be at c^u and v's at c^v (giving utilities $u(c^u)$ and $u(c^v)$). Figure 1.6 also shows that $w_N - \alpha^v > w_N - \alpha^u$, which implies that u's insurance cover is greater than v's, thereby illustrating risk-averse decision makers take greater insurance cover.

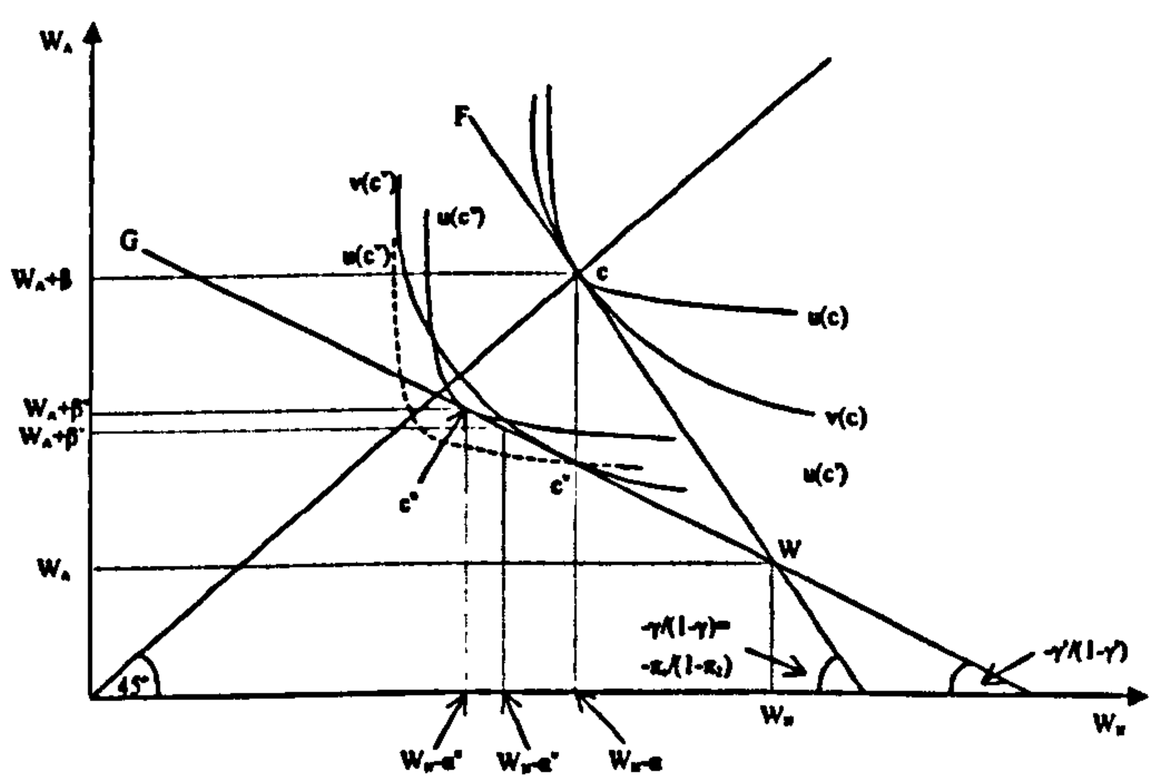

Figure 1.6: Optimal insurance with varying degrees of risk aversion

But how robust is this result? If u is more risk-averse than v, then by Arrow-Pratt $u = T(v)$, where T is a concave function. At v's optimum, the negative of the slope of v's indifference curve is equal to the premium ratio

$$\frac{-dw_A}{dw_{N|v}} = \frac{(1-\pi_s)v'(w_N - \gamma C^v)}{\pi_s v'(w_A + (1-\gamma)C^v)} = \frac{(1-\gamma')}{\gamma'} \tag{1.52}$$

where C^v is v's optimal coverage. Suppose instead u holds v's insurance cover, then the slope of the indifference curve of u at endowment c^v is

$$\frac{-dw_A}{dw_{N|u}} = \frac{(1-\pi_s)u'(w_N - \gamma C^v)}{\pi_s u'(w_A + (1-\gamma)C^v)} \tag{1.53}$$

or given that $u = T(v)$, so that $u' = T'(v)v'$

$$\frac{-dw_A}{dw_{N|u}} = \frac{(1-\pi_s)T'(v(w_N - \gamma C^v))v'(w_N - \gamma C^v)}{\pi_s T'(v(w_A + (1-\gamma)C^v))v'(w_A + (1-\gamma)C^v)} \tag{1.54}$$

which upon substitution gives

$$\frac{-dw_A}{dw_{N|u}} = \left[\frac{T'(v(w_N - \gamma C^v))}{T'(v(w_A + (1-\gamma)C^v))}\right]\left[\frac{-dw_A}{dw_{N|v}}\right] \tag{1.55}$$

so that the slope of u's indifference curve at c^v is a function of the slope of v's indifference curve at c^v. Clearly, as $v(w_N - \gamma C^v) > v(w_A + (1-\gamma)C^v)$ at the allocation c^v(as c^v lies below the 45° line), then as T is concave, it is known that

$$T'(v(w_N - \gamma C^v)) < T'(v(w_A + (1-\gamma)C^v)) \tag{1.56}$$

So that the ratio $\left[\frac{T'(v(w_N-\gamma C^v))}{T'(v(w_A+(1-\gamma)C^v))}\right] < 1$, therefore the above formula implies that at c^v

$$\frac{-dw_A}{dw_{N|u}} < \frac{-dw_A}{dw_{N|v}} \tag{1.57}$$

the marginal rate of substitution of u at c^v is less than the marginal rate of substitution of v at c^v, so that u's indifference curve is flatter than v's indifference curve. Figure 1.6 shows this by comparing the (dashed) indifference curve $u(c^v)$ and indifference curve $v(c^v)$,which plainly shows that u's optimal insurance contract must lie further up the G curve, nearer to the 45° line, at point cu. Therefore, $\alpha^u > \alpha^v$, which implies that $C^u > C^v$, so that the more risk-averse decision maker takes on greater insurance cover than the less risk-averse - as predicted.

The above analysis all seems to work quite predictably and cleanly. But how robust and stable is the analysis ? For example, what happens if instead of a well behaved Gaussian description of the probabilities

facing the two decision makers, their alternatives are governed by a distribution with heavier tails, so that extreme events are now much more likely to occur ? Do the same results on fair and unfair insurance still carry through ? What would happen for example if there were no market for the type of risk being insured in our simple example ? The insurance company would then presumably attempt to estimate unobservable parameters and calibrate some sort of model that gave reasonable prices for the insurance risk concerned. But what happens if their results are wrong ? These are the joint problems of model risk and model error.

Further, from the perspective of the decision maker, just how robust would the optimality results be to changes in the form of the utility function ? What would be the impact of changing the optimisation to optimise not with respect to a single model, but with respect to a family of similar models ? It is well known that the classical utility approach contains no formal account of how the actions of the decision makers are influenced by the occurrence of subsequent insured events - in other words, no account is taken of feedback either in terms of information or in terms of the structure of the model itself. For example, how would contagion type effects such as natural disasters or terrorism impact the analysis ? Is utility maximisation still the appropriate model ? How might habit persistence affect the robustness of the utility function ? Is the standard model robust to any financing costs that may occur if a claim takes a long time to settle ? Are decision makers sensitive to the early resolution of claim uncertainty ?

The above questions indicate two clear issues surrounding the classical utility approach. First, little or no account is taken of feedback in the standard utility maximisation model - actions or events in one period are not allowed to affect those in the next. The classical utility model is a static model that does not incorporate the dynamics of the underlying relationships. Second, uncertainty, where it is modelled, is generally only allowed to affect the parameters of the model - usually via the influence of probabilities. However, what if there is a possibility that the model itself may be incorrect in one or more significant respects ? Researchers such as Derman (1996) and Jacquier and Jarrow (1996/2000) were among the first to explicitly recognise model risk and model error. Their analysis identified three principal sources of model risk. First, is model identification and estimation errors. In the context of the insurance example, what if the premium calculation model is wrong, or what if the utility maximisation model is wrong ? Either way, even if the model is selected after meticulous analysis, both parties have no way of knowing ex-ante whether their model is correct or whether their parameter estimations are accurate. The second source of model risk is market imperfections. Problems for the insurance model might be that markets are not complete, that claims arrive discontinuously and that transactions costs exist. The final source of model risk is of course the improper use of a model. In terms of the insurance example, it may be that using a single number for premium may be incorrect as there may (for example) be size and market

segmentation effects to consider.

There are therefore two questions. First, just how robust are results from the classical utility maximisation model in such circumstances ? Second, what are the alternatives to the classical utility maximisation approach ? The aim of the next section is review the approaches to dealing with robustness from both a statistical as well as a mathematical perspective. To do this the section begins with a short discussion of the alternative approaches to robust optimal control, before proceeding to analyse the research applying these methods to decision making in the presence of uncertainty.

1.4.2 Optimisation based decision making in the presence of uncertainty

The previous section reviewed the research on the utility maximisation approach to decision making in the presence of uncertainty. The utility based approaches assume varying forms of optimal (usually maximising) behaviour on the part of the decision maker. However, beyond the use of simple probabilities, the utility maximisation approach pays little attention to how well the decision rules stand up to the existence and impact of uncertainty beyond the simple parametric uncertainty incorporated through the use of probabilities. The objective of this section, therefore, is to review alternatives to the utility based approaches to decision making in order to see how the development of rules in other disciplines copes with the wider impact of uncertainty. In other words, the question is "do other approaches to developing decision rules offer greater degrees of robustness in their modelling of the decision making process ?"

Statisticians and engineers have been the main interested research parties in the area of robustness. Both groups have taken an optimising approach, but their paths have diverged significantly in many respects. On the one hand, statisticians have mainly focused on ways to make the parameter estimates and testing methods more robust using optimisation rules such as the least squares minimisation rule. Whilst on the other hand, engineers have been predominantly concerned with modeling multi-stage decision making processes to produce more reliable engineering solutions to such problems as space shuttle re-entry or cruise control, by explicitly incorporating the effects of feedback as a way of helping to minimise the effects of uncertainty on the operation of their physical systems. Engineers initially tried to achieve this aim by developing approaches to decision making that began by concentrating on optimality, but subsequently moved on to focus on the broader and more complex issues of robustness and stability of their models.

This section therefore reviews each approach, beginning first with classical optimisation, followed by how robustness has been dealt with in statistics, then moving on to the different ways in which robustness has been approached in physical models via the far more explicit use of feedback loops. To achieve this objective, the following approach is followed. First, clear definitions are provided of the key concepts

necessary to understand and analyse the ideas of stability and robustness. There follows a detailed review and analysis of robust statistical methods. This is in turn followed by an introduction to and review of the literature and research on robust optimal control theory. The section introduces optimal control theory, the role of robustness and stability and uses a simple example to illustrate the critical principles in the context of models and decisions. There then follows a review of how robust optimality has hitherto been applied to dynamic models of multi-stage decision making problems in finance. In the final section of this chapter, the potential breadth and depth of opportunities for research in the area are examined and the theoretical basis and framework for the thesis is established.

1.4.3 Robust statistical methods - robust estimation

Initial interest in the robustness of decision making arose in engineering, where there was considerable focus on the problem of developing control systems that were both stable and robust in operation. The starting point for much of this work was classical regression which was used to estimate parameters of models that needed to continue to perform acceptably, even under extreme conditions. Robust operation was set up as an optimisation problem through seeking to find parameter estimates that were the best unbiased estimators with the minimum variance. Regression-based decision rules, often in the form of hypothesis tests, have often also been applied as part or the whole of a control system. This approach therefore lead engineers to concentrate on two main lines of research in regression analysis in an attempt to achieve their objective of finding control rules that were stable and robust in operation, namely, estimation and hypothesis testing.

With regard to estimation, one of the most widely used techniques in this area of optimization is ordinary least squares regression (OLS), which at its simplest uses a linear programming approach baed on a quadratic norm and consists of obtaining a functional model that relates the value of a "target" dependent variable, Y, with the values of the independent variables $X_1, X_2, ..., X_n$. Traditional approaches to the problem assume a particular form of the parametric function and use all data to obtain the values of the functional parameters that are optimal according to the least squares fitting criterion. This global parametric approach has been and continues to be used widely, giving reliable results when the assumed model is a close fit to the underlying data. Many financial series which are used for decision analysis are time series and regression analysis has an enormous plethora of techniques available to deal with a wide variety of problems encountered when modelling, controlling and predicting economic systems. In the study of both univariate and multivariate processes (such as are frequently encountered in utility theory), regression techniques provide a framework for describing not only the properties of the individual series, but also the possible cross-relationships among the series. As in many areas, the econometrician seeks

to model series jointly in order to understand the dynamic relationships among series over time and to improve the accuracy of forecasts.

The standard OLS approach involves computing the $\widehat{\beta}$ parameter estimates from all of the data in a single pass. In some cases it can be more convenient, or even vital to perform the estimation recursively by taking observations one at a time, recalculating the $\widehat{\beta}$'s and the associated covariance matrix each time. This approach has been used with some success in decision making problems such as enabling a time-varying model to track behaviour which is too complicated to be described adequately by a constant parameter model; or in the situation where a time-varying linear model can be used to track a highly non-linear system (see for example the work on adaptive filtering by Bolzern, Colaneri and Nicolao, 1999, which is discussed later). Arguably one of the most common problems with the standard univariate regression model is the assumption of normality of the residuals. In practice, residuals are rarely normally distributed due to the almost inevitable presence of outliers. Non-normality of residuals presents two problems in the context of the analysis of decision making problems, namely, in the derivation of estimators and hypothesis testing. With respect to the former, non-normality in and of itself may not necessarily be a problem, since least squares estimators remain unbiased and consistent, and asymptotically χ^2 and hypothesis tests are still available. However, least squares estimators are no longer asymptotically efficient, such that hypothesis tests critically lack robustness as the finite sample distribution can be altered dramatically by relatively small changes in the distribution of the errors and because a single observation can cause $\widehat{\beta}$ to take on any value - a situation all too familiar in time series data analysis.

The classical regression approach begins with a simple linear model of the form

$$Y = X\beta + u \tag{1.58}$$

and finds the $\hat{\beta}$ that minimises a simple quadratic norm of the sum of the squared errors

$$\hat{\beta} = (X'X)^{-1} X'u \tag{1.59}$$

If, at the limit, the error variance tends to ∞, then the least squares estimators are no longer minimum variance and conventional hypothesis tests become meaningless, making the analysis of decision rules extremely difficult. It was the combination of non-normality and possibly infinite variance in residuals that led to the development of robust estimators, which place less weight on outliers relative to classical least squares. The majority of the research in the univariate sphere has been concerned with the problem of deriving either estimates or test statistics that are insensitive to outliers and the distributions generally considered as alternatives to the normal have almost always been symmetrical with heavier tails than the

normal (e.g. Huber, 1964, 1973, 1977 and 1981).

A variety of univariate robust estimators have been suggested, of which the three major ones will be analysed here. The simplest and most popular is quantile regression, where the estimator is based on minimizing a criteria function of the form:

$$\sum_t \omega_\theta(u_t) \tag{1.60}$$

where, for $0 < \theta < 1$, either $\omega_\theta(u_t) = \theta|u_t|$, if $u_t \geq 0$,or, $\omega_\theta(u_t) = (1-\theta)|u_t|$, if $u_t < 0$. Clearly, as $\omega_\theta(u_t)$ is the weighted sum of the absolute values of the residuals, outliers receive less weight than under ordinary least squares. When $\theta = 0.5$, the least absolute error estimator is produced. The second approach extends the basic quantile regression idea, such that in the case of large positive or negative outliers, extreme values of θ can be used to heavily penalize abnormally varying observations. Using values of θ between 0 and 1 yields regression quantile estimators $\widehat{\beta}(\theta)$. The effects of the abnormal outliers will inevitably be most concentrated in the extreme values of θ, which has led to the polynomial combination of quantile estimators to produce estimators such as the trimean quantile estimator:

$$\widehat{\beta}_{trimean} = 0.25\widehat{\beta}(0.25) + 0.5\widehat{\beta}(0.5) + 0.25\widehat{\beta}(0.75) \tag{1.61}$$

and the trimmed regression quantile estimator:

$$\widehat{\beta}_{trimmed} = \frac{1}{(1-2\phi)} \int_{\phi}^{1-\phi} \widehat{\beta}(\theta)\, d\theta, \tag{1.62}$$

where $0 < \phi < 0.5$. The trimmed quantile estimator is calculated by first deriving $\widehat{\beta}(\phi)$ and $\widehat{\beta}(1-\phi)$, excluding all observations lying on or below the ϕth quantile and all those above the $(1-\phi)$th quantile, then applying ordinary least squares to the remaining observations. $\widehat{\beta}_{trimmed}$ has the attractive robustness property that its breakdown point is approximately 0.5. The breakdown point of a regression estimator is the largest fraction of data which may be replaced by arbitrarily large values without making the Euclidean norm $\left\|\widehat{\beta}\right\|$(defined as $\left\|\widehat{\beta}\right\|^2 = \sum_{i=1}^{p} \widehat{\beta}_i^2$, where p is the number of smallest squared residuals) of the resulting estimate tend to ∞. Notwithstanding the fact that the resulting estimator is an asymptotically normal estimator of β and has a high breakdown point, the fundamental weakness of the approach is that the choice of the quantiles is arbitrary, so that robustness and stability are only achieved asymptotically and can only gained for arbitrarily selected sub-sets of a data-set.

Robust estimation in the multivariate situation remains relatively underdeveloped, due principally to the problem that the variety of outliers (both in type and impact) in multivariate data can be enormous.

In a univariate situation it is frequently easy to detect atypical observations by casual inspection. In the multivariate setting, in contrast, observations can often only be considered atypical when the value of each variable is considered in relation to all other variables. Krzanowski (1996) points out the even more problematic issue of the simultaneous nature of outlier effects on the location, scale and orientation of multivariate data, frequently making it impossible to adjust for the marginal impacts of outliers in isolation without using an iterative procedure such as Huber (1964 and 1973) or Campbell (1980 and 1982).

Huber (1964) introduced the class of M-estimators which are based on the simple idea that instead of using the standard estimates of location and dispersion, iteratively calculated weights are used to penalise outliers or atypical observations. In order to take account of the joint problems of differential variances and correlations between variates, the Mahalanobis distance measure is used to provide a measure against which to judge the atypicality of the distance of an observation from the sample mean. The key problems with the Huber method are twofold. First, is the inconvenience of solving two non-linear equations for which there is no guarantee of the existence, uniqueness, or stability/robustness of the solution. Second, the robustness properties of the resulting estimates are not satisfactory, with the inaccuracies of variance estimates being as much as 3% (asymptotic bias) according to Campbell (1980 and 1982), thereby rendering them effectively useless for the analysis of many areas of economic decision making.

Krasker and Welsch (1982) generalize the M estimator concept by penalizing observations with abnormal residuals. Assuming that breakdown point (defined as the smallest proportion of arbitrary observations that an estimator can resist without becoming unbounded) is a reasonable measure of robustness; then the generalized M-estimators have a positive breakdown point. However, this decreases to zero as the number of predictors rises, which in turn means the estimators exhibit poor robustness and stability.

In an attempt to deal with this problem, Rousseeuw and Yohai (1984) introduced S-estimators which are consistent and asymptotically normal when the distribution of errors is symmetric around zero. The idea is to define T_n, as the set of parameters β that produce residuals with the smallest dispersion. By way of illustration, suppose that $(y_i, x_i), i = 1, ..., n$ satisfy the model:

$$y_i = x_i^{'}\theta + \sigma\epsilon_i, \tag{1.63}$$

where $i = 1, ..., n$ and $r_i(\beta) = y_i - x_i^{'}\beta$. Then the M-scale, $s_n(\beta)$,of the residuals $(r_1(\beta), ..., r_n(\beta))$ is defined as the solution to the equation:

$$\frac{1}{n}\sum_{i=1}^{n}\rho\left(\frac{r_i(\beta)}{s_n(\beta)}\right) = E_G(\rho(\epsilon)) \tag{1.64}$$

The corresponding S-estimator is defined by the property of minimizing $s_n(\beta)$:

$$T_n = \underset{\beta \in \Re^p}{\arg\min}\, s_n(\beta) \tag{1.65}$$

S-estimators are consistent for the true parameter θ and asymptotically normal when the distribution of the errors is symmetric around zero. Their drawback is that they are unable to simultaneously achieve high efficiency and high breakdown, resulting in a trade-off between efficiency and robustness. In 1987 Yohai suggested MM-estimators for regression in order to simultaneously achieve high efficiency and robustness. MM-estimators have a high breakdown point, are consistent and asymptotically normal, but rely on strong regularity conditions to achieve a high degree of efficiency - including a symmetric error distribution. Yohai and Zamar (1988) introduced the class of τ estimates, which assuming a linear model of the form: $y_i = x_i'\theta + \sigma\epsilon_i$ as described above, using two functions $\rho_1, \rho_2 : \Re \longrightarrow \Re$, such that for each $\beta \in \Re^p$ define the τ estimate of the scale of the residuals $(r_1(\beta), ..., r_n(\beta))$ by:

$$\tau_n^2(\beta) = s_n^2(\beta)\frac{1}{n}\sum_{i=1}^{n}\rho_2\left(\frac{r_i(\beta)}{s_n(\beta)}\right) \tag{1.66}$$

where $s_n(\beta)$ is the M-scale of the residuals calculated using the function ρ_1. The τ-estimates for the regression are then defined as: $\widehat{\theta} = \arg\min_\beta \tau_n(\beta)$. In particular, $\tau_n(\beta)$ is an efficient and robust estimator of σ. However, as in the case of the MM-estimator, strong regularity conditions (which once again include symmetry of the residuals) are required in order to achieve strong consistency and asymptotic normality.

Clearly, there are significant problems with robust estimation techniques in the univariate world. Unsurprisingly, the problems are, as already alluded to above, exacerbated in the multi-variate world. As far as stability is concerned, whilst it is true that ridge regression techniques can deal with the instabilities associated with high within-group correlations in a multivariate setting, Campbell (1980) finds it more revealing to apply these techniques to canonical variate estimation and achieves good stability. Krzanowski (1996) explores the use of canonical variate coefficient estimation in the presence of non-ideal data with the aim of addressing the issue of robustness. He examines the issues that arise when there is instability in the coefficients. The first problem arises, as in the univariate case, when there exist a number of atypical observations or outliers. The second arises when the within-groups covariance matrix, W, has a small eigenvalue and the between-groups sum-of-squares in the direction of the corresponding eigenvector is also small. Campbell (1982) uses robust M-estimation techniques to down-weight the outliers to improve robustness and his iterative procedure involves the simultaneous estimation of both the weights and the robust estimators.

A further and more recent direction in robust estimation techniques is the work on adaptive filtering by

authors such as Sayed and Kailath (1994) and Bolzern, Colaneri and Nicolao (1999), BCN for short. This work considers from the point of view of control theory, the general problem of adaptive filtering of a scalar signal variable which has been corrupted by noise. BCN model the signal as a linear regression depending on a drifting parameter and consider the mean square and worst case performances of the normalised least mean-squares, Kalman and central $H\infty$ filters. These three criteria are evaluated in terms of mean-square or H_2 performance, namely by filtering the error variance assuming white noise disturbances, but BCN note that knowledge of the spectral characterization of the disturbances required for such a technique is rarely if ever available. They use this lack of knowledge as a motivation for considering worst-case performance of adaptive filters in the face of arbitrary disturbances. BCN particularly consider the $H\infty$ approach which takes as a robustness index the maximum attenuation level from disturbances to the estimation error with respect to the set of admissible disturbances. Their approach is to regard the adaptive filtering problem as a special case state-space estimation problem and use $H\infty$ techniques to design robust algorithms. For a fixed regressor sequence, necessary and sufficient conditions exist to verify whether a given filter achieves a required attenuation level. However, for higher frequency processes where regressors are not known in advance the problem reduces to finding a method of ensuring the achievement of a required attenuation level across all possible regressor sequences.

BCN examine three alternative adaptive filtering approaches to dealing with both problems of process and measurement disturbances, namely, normalized least mean squares (NLMS), Kalman and central $H\infty$ filters. In the context of what is to follow in the section on pure $H\infty$ optimal control it is worthwhile to analyse briefly (without digressing into the technicalities of the vector algebra) the robustness properties of each of these filters. In the case of NLMS, when the parameter vector is constant, it can be shown (see for example Hassibi, Sayed and Kailath, 1996) that NLMS is $H\infty$ optimal in that it coincides with the central $H\infty$ filter that guarantees the minimum attenuation level. Unfortunately, this happens to occur at the point of the worst H_2 performance, whereas the central $H\infty$ filter is risk sensitive optimal and is also maximum entropy. However, the key objection to NLMS is that it cannot guarantee a finite attenuation level for all possible regressor sequences.

As far as the Kalman filter is concerned BCN show that, assuming unit variance for the measurement noise and slow parameter drift), then its performance can be arbitrarily close to zero (i.e. it is H_2 optimal). In addition, its $H\infty$ performance is also better than NLMS, but still not as good as a full $H\infty$ filter. But as BCN point out, a natural way to achieve robustness is to use the $H\infty$ filter which guarantees any desired attenuation level. Unfortunately, such an approach by itself results in poor H_2 performance. BCN show that the best way forward is to use a synthesis of H_2 and $H\infty$ methods, such that the target is to minimize H_2 performance subject to an $H\infty$ constraint. Interestingly, whilst this filter approach still

fails to produce an acceptable solution in the time-varying case (due to feed-forward effects on the $H\infty$ constraint), it appears to offer a promising direction of investigation when using an $H\infty$ optimal control approach to evaluate the behaviour of optimal decision rules. This is because $H\infty$ optimal control enables both the desired attenuation level and speed of adjustment to be varied either solely or jointly and so provides the ability to estimate the costs of a mixed $H_2/H\infty$ strategy in a very flexible fashion. $H\infty$ effectively provides the estimator for the worst case distribution.

1.4.4 Robust statistical methods - confidence interval problems

Approaches to the problem of obtaining robust and stable inference in linear regression models have proceeded in much the same direction as those of estimation. This is not entirely unsurprising given the fact that the issues of outliers and heavy tails obviously affect both estimators as well as the construction of robust confidence intervals. The research in robust hypothesis testing can be split into two main areas, namely, estimation of the asymptotic variance of the residuals and deriving a robust confidence interval.

Estimating the asymptotic variance of the residuals is a fundamental requirement for the development of any robust inference procedure. It is non-trivial to derive a formula for the asymptotic variance in the presence of asymmetry, as it is difficult to disentangle the simultaneities in the estimation of the scale parameter from the variance, because the distribution of the scale parameter depends on the estimate of the auxiliary position estimator. Whilst S-estimators can be used to break this recursive loop, significant finite sample problems remain for sample sizes less than 30, due to the fact that large locational values in the denominators of several of the estimators result in extremely large values for the asymptotic variance.

Work by Gross (1976) examined the robustness of confidence intervals for heavy-tailed symmetric distributions using jackknifed estimates and produced disappointing results which agreed with an earlier Monte Carlo study by Andrews et al (1972). Carol (1979) examined several robust estimates and techniques to estimate their variances under asymmetric distributions and found that the empirical approximation to the asymptotic variance to provide a severe under-estimate. Rocke and Downs's (1981) comparison of the empirical asymptotic variance, the jackknife and the bootstrap under a variety of different distributions both symmetric and asymmetric, found relatively poor robustness performance of the bootstrap variance estimator even under symmetric distributions with only slightly fat tails. Gosh et al (1984) examined the fat tail problem under a variety of symmetric and non-symmetric distributions and Aalse found that the bootstrap variance estimator performed poorly.

Shao (1990 and 1992) changed tack somewhat and studied the estimation of the variance of Frechet-differentiable statistical functionals when the generating distribution has fat tails. His proposal was for a modified bootstrap variance estimator which truncates the values of the re-sampled estimates to down-

weight the tails. Despite the fact that the estimator is asymptotically consistent, the truncation limits remain arbitrary and are not linked to the data in any rigorous mathematical way, which is critical in the case of a finite sample application.

Leaving aside the critical issue of estimating the asymptotic variance, it is possible to define a robust confidence interval for testing hypotheses as follows: for a fixed significance level, α, and a robust estimate T_n, find the percentile q_n such that:

$$1 - \alpha = P\left(|T_n - \mu_0| < q_n\right) \tag{1.67}$$

where μ_0 is the parameter of interest. Next, assume that the sequence T_n converges almost surely to a limit T_∞. Then the problem is that T_∞ need not necessarily be equal to μ_0 due to an asymptotic bias in T_n and that whilst it is possible to estimate the asymptotic variance, there is no estimator available for the asymptotic bias, though it is possible to bound it reasonably for most practical purposes such that relatively good robustness properties can be obtained within the bounded region.

How then is it best to sum-up the robustness and stability problems associated with regression based approaches to modelling decision making problems ? The first point to make is that, as Huber (1981) notes, robustness in econometric terms is all about doing well near a parametric model. Most commonly, robustness in econometric terms is defined as simultaneously controlling bias due to outliers and achieving high efficiency in the event that the underlying data is normally distributed. Second, whilst many econometric techniques exhibit highly satisfactory stability characteristics, most if not all, fail to simultaneously achieve both stability and robustness under anything other than highly restrictive conditions. Thirdly, all of the above estimators are either obliged to perform an exhaustive search or assume a known value for the amount of noise present in the data set (frequently referred to as the contamination rate), or equivalently an estimated scale value or inlier bound. When faced with more noise than assumed, all of the above estimators lack robustness and when the amount of noise is less than the assumed level, they lack efficiency, such that the parameter estimates suffer in terms of accuracy because not all the good data points are taken into consideration. Finally, all of the methods are also constrained to estimating a single component in a data set and are therefore not able to explicitly incorporate the dynamics of a decision making process into a meaningful model.

1.4.5 Classical optimisation

During the review of utility based decision making little was said about the nature of the objective function when actually making decisions, apart from the argument that the decision maker is assumed to pick the alternative that maximises utility. So it is arguable whether the concept of "best" or "optimal"

decisions naturally emerges as the fundamental approach for formulating decision problems. However, notwithstanding this, when the principle of optimality is applied, a single quantity that summarises the performance or value of a decision (such as the associated expected utility) is isolated and optimised (either minimised or maximised as appropriate) by a proper process of selection from among valid and available alternatives. The resulting decision is taken as being the optimal solution to the decision making problem. This approach has the virtues of being simple, elegant, precise and tractable.

At a purely mathematical level, the simplest form of this approach is easily and most often described in terms of the example of a single period production problem that is encountered when deciding how to use raw materials. Assume that it is decided to produce x_j units of a product, $j = 1, 2, .., n$, where the selling price of one unit of product j is σ_j. Assume also that the cost of producing one unit of product j is $\sum_{i=1}^{m} \rho_i a_{ij}$, where ρ_i is the unit value of the raw material and a_{ij} is the amount of raw material i required to produce a single unit of product j. Per unit net revenue is therefore given by

$$c_j = \sigma_j - \sum_{i=1}^{m} \rho_i a_{ij} \tag{1.68}$$

so that the net revenue corresponding to the production of x_j units of product j is simply $c_j x_j$, giving total net revenue of

$$\sum_{j=1}^{n} c_j x_j \tag{1.69}$$

The decision maker wishes to maximise this quantity. However, there are constraints on the production levels that can be selected. The two main types of constraint being that production quantity x_j must be non-negative ($x_j \geq 0, j = 1, 2, .., n$) and that it is not possible to produce more than can be supported by existing stocks of raw materials, implying the following constraint

$$\sum_{j=1}^{n} a_{ij} x_j \leq b_i \tag{1.70}$$

where b_i is the stock of raw material i. In a classical linear programming approach to maximising net revenue, the objective would be to find the optimal amount of each of the production quantities x_j to be produced that together maximise net revenue by solving the following maximisation problem

$$\text{maximise:} \quad \sum_{j=1}^{n} c_j x_j \tag{1.71}$$

$$\text{subject to:} \quad \sum_{j=1}^{n} a_{ij} x_j \leq b_i \text{ and } x_j \geq 0 \tag{1.72}$$

$$\text{where } i = 1, 2, ..., n \text{ and } j = 1, 2, ..., m \tag{1.73}$$

Depending on the precise form, this type of problem is usually solved by using one of a number of algorithms to find the values of the optimal c_j's. If the process can be completely, accurately and reliably described in terms of such a simple model, then all will be well and decisions made will indeed be optimal over any range of feasible sets of raw materials and products. This simple linear approach underpins much of classical mathematical economics and econometrics.

1.4.6 Robustness, stability and optimal control

A moment of reflection on the simple linear model soon reveals that the model pays no attention to modelling the dynamics of the underlying system. The first and obvious extension to the simple linear model is therefore to explicitly incorporate into the model one or more aspects of the dynamics of the underlying system. Introducing time leads to the state-space approach whereby different states represent the behaviour of a model over time and in different states of nature, the transition between states then being most often driven by difference (in the case of discrete time models) or differential (in the case of continuous time models) equations to capture and explain how the model evolves, or transitions, between states.

The problems begin to arise with classical optimisation as typified by the simple linear programming approach, however, once uncertainty is introduced into the picture. What happens for instance in the simple production model if there is an alternative model of the production process that when optimised produces a radically different set of optimal c_j's ? Which set of c_j's is the correct set and are there any unmodelled quantities that will affect the calculation of the optimal set ? Furthermore, how does the producer know when to stop ? What are the stopping rules ? For how long is the optimum valid ? What happens to the optimum if some of the raw materials run out or degrade over time, how can such extreme behaviour be modelled into the system ? Finally, how will the model respond to unforeseen events such as a fire at the factory or reject batches of raw materials ? Such a simple model cannot hope to handle uncertainties surrounding parameter values as well as the actual model itself.

In the early 1900's, classical statistics was based on a probabilistic approach that was frequently underpinned by a Gaussian distribution. The obvious advantage of such an approach was in only having to deal with means and variances and not having to confront the real distribution - frequently an unpleasant task. However, there were difficulties with Gaussian method, most notably a lack of robustness. As already seen, the ideas behind risk aversion and utility maximisation are based either implicitly or explicitly on the classical Gaussian probability approach. Gauss derived his famous distribution by appealing to what is now commonly known as the maximum likelihood principle. This is a classical example of a fallacy of composition and principally came about due to the omission of any explicit attempt to model feedback.

It was only in the early 1960's that Tukey used a simple example of adding numbers from two Gaussian distributions with unequal means and variances, to illustrate that the Gaussian distribution and related procedures were totally non-robust, thereby illustrating the need to model feedback in order to ensure robustness in the statistical method.

Writing in 1953, G.E.P. Box was the first to first to use the term "robust" in a quantitative sense. In its simplest statistical sense, robustness is generally taken to mean "insensitive to small departures from the idealised assumptions for which an estimator is optimised". As Huber (1981) points out, the word "small" can have two alternative but equally important interpretations, on the one hand it can refer to fractionally small departures for all data points, or else fractionally large departures for a small number of data points. The latter interpretation lead to the concept of outlier points generally more relevant for statistical procedures, whereas the former interpretation lead to the development of models of physical systems that produced more reliable predictions of behaviour in the neighbourhood of some base or idealised model.

As already discussed, the issue of robustness has been studied and applied by two broadly different groups, namely, statisticians and engineers. The approach taken by the statisticians was described in the previous section. In contrast, this section deals with the ways in which engineers have concentrated on adapting and refining optimisation techniques by the explicit incorporation of feedback loops into decision making as a means of capturing and reducing the effects of uncertainty on models of physical processes. By developing mathematical approaches and algorithms that explicitly incorporate feedback control, engineers have devised approaches that deal with the effects of uncertainty in such way that, within a broad range of circumstances, they are able to deal explicitly with the joint problems of robustness and stability.

Fundamental to understanding how such models work, is a clear appreciation of the critical role and overriding importance of incorporating feedback into the modelling of a process. That feedback is important is underlined by the existence and importance of examples of feedback in virtually every discipline from biology to economics, through to engineering and psychology. At its simplest, a feedback system is one where there is a process (the cause and effect relation) whose operation depends on one or more variables (the inputs) that cause changes in some other variables. If an input variable can be manipulated then it is referred to as a control input, otherwise it is considered to be a disturbance (or noise) input. The process variables that are monitored are referred to as the system outputs. The role of the modeler is to develop a process called the feedback controller whose role is to gather information about process behaviour by observing the outputs and then generating new control inputs in order to make the system behave in the desired fashion. Decisions taken by the controller process are therefore

critical. If they are incorrect they can lead to instability and catastrophe instead of improvement and stability. This is the fundamental reason that feedback controller design, which is the determining of the rules for automatic decisions taken by the feedback controller, is the critical issue for modelling a decision making process. Explicit modelling of feedback is the means that is used to reduce the uncertainty in a model of a system.

In a physical model, a feedback control system generally consists of four sub-systems:

- a process to be controlled
- sensors
- actuators
- controller

These sub-systems can be thought of as being connected as shown in the idealised system view in figure 1.7. The process is the actual physical model that cannot be modified.

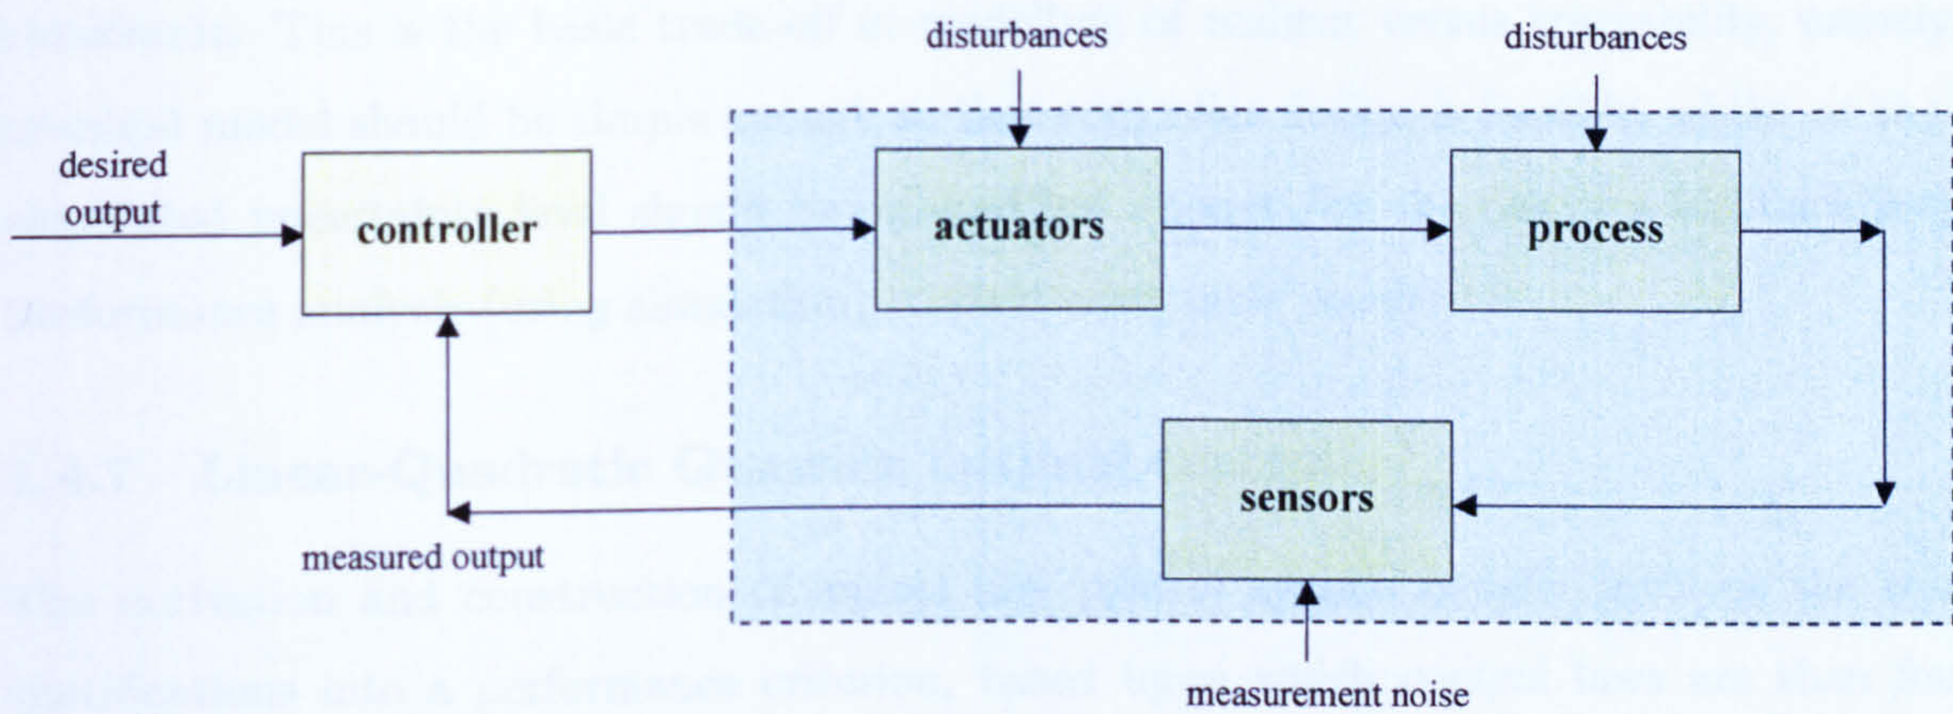

Figure 1.7: Simple feedback control system

In a physical system, actuators and sensors are chosen by engineers based on both physical and economic constraints such as the range of signals to be measured and/or generated, as well as the accuracy versus cost of the measurement devices. The controller must then be designed for the given system to achieve the required objectives. In engineering systems the controller is frequently a computer, whereas in economic systems the controller is usually the decision making agent. Designing an efficient and robust controller demands a good understanding of the cause effect relationship between the input and output variables. Engineering systems are frequently well described by physical laws of motion or nature, so that accurate mathematical models contain relatively low levels of uncertainty compared with economic models where input-output relationships can often depend on far more qualitative relationships.

The modelling of physical processes usually begins with the development of a mathematical model which describes the dynamic behaviour of the underlying process. This procedure should take account of modelling uncertainties, including parametric and intrinsic model uncertainties. The first step in the procedure is to derive a simple, elegant and accurate mathematical model of the underlying process. Once this is completed and performance objectives specified, then a feedback controller can be synthesised. This process generally involves simulation testing to establish how well the controller performs compared with the stated objectives. If performance is unsatisfactory, then the process model and controller must be evaluated and modified and the cycle of testing re-run. This iterative process continues until satisfactory results are obtained. Once complete, the result should be a nominal model of the process and an uncertainty description that represents the required confidence level for the system under consideration.

It is generally the case that the uncertainty magnitude can only be decreased and the confidence level consequently increased, by making the model more complex. This is usually achieved by increasing the number of variables and/or equations, as well as the linearity or non-linearity of the equations being used. Increasing complexity to achieve a better description of reality usually leads to complicated models which may be difficult to solve, or at the limit may fail to produce meaningful results within a usable timeframe. This is the basic trade-off in modelling of realism versus tractability, namely, that a useful nominal model should be simple enough so that controller design is feasible, whilst at the same time the associated uncertainty level should be reduced low enough, by the use of a feedback loop, to allow the performance analysis (using simulation) to yield acceptable results.

1.4.7 Linear-Quadratic Gaussian optimal control

The derivation and construction of almost any control system usually involves the transformation of specifications into a performance criterion, based upon which control laws are then found which will minimise the criterion. At its simplest, optimal control begins with an initial state x_0 of the system and a control history $u(t), t \in [0,T]$ of the process, so that the evolution of the system can be described by the state equation (frequently assumed to be a first-order differential equation) of the form

$$\dot{x}(t) = f(x(t), u(t), t), \; x(0) = x_0 \tag{1.74}$$

where the vector of state variables is $x(t) \in E^n$, the vector of control variables is $u(t) \in E^m$ and $f : E^n \times E^m \times E^1 \rightarrow E^n$, where the function f is assumed to be continuously differentiable. The path $x(t), t \in [0,T]$ is called the state trajectory, while $u(t), t \in [0,T]$ is called a control trajectory. Constraints are imposed on the control variables and an admissible control is defined (in a continuous time model) as

a control trajectory that is piecewise continuous and possesses the additional property

$$u(t) \in \Omega(t) \subset E^m, t \in [0,T] \tag{1.75}$$

with the set $\Omega(t)$ being determined by physical or economic constraints on the values of the control variables at time t. The objective function is frequently a quantitative measure of the time based performance of the system, so that an optimal control is defined to be an admissible control that maximises the objective function. In economics, a typical objective function might be a cost function of the form

$$J = \int_0^T F(x(t), u(t), t)\, dt + S[x(T), T] \tag{1.76}$$

where the functions $F : E^n \times E^m \times E^1 \rightarrow E^1$ and $S : E^n \times E^1 \rightarrow E^1$ are assumed to be continuously differentiable. In economic terms, $F(x, u, T)$ could be the instantaneous rate of profit and $S[x, T]$ could be the depreciated value of a company's assets, with x representing the system state and T the terminal time. The optimal control problem in this simple example would therefore be to find an admissible control u^* that maximises profit less the depreciated value of the assets contained in the objective function subject to the constraints

$$\left\{ \begin{array}{l} \max\limits_{u(t)\in\Omega(t)} \quad \left\{ J = \int_0^T F(x(t), u(t), t)\, dt + S[x(T), T] \right\} \\ \text{subject to:} \\ \qquad\qquad \dot{x} = f(x, u, t),\ x(0) = x_0 \end{array} \right\} \tag{1.77}$$

The controller u^* is an optimal controller and x^* (determined using the state equation with $u = u^*$), is known as an optimal trajectory. Being able to achieve optimality requires a clear definition of the concept. Bellman (1957) was the first to provide a formal statement of the principle of optimality:

> "An optimal policy has the property that, whatever the initial state and initial decision are, the remaining decision must constitute an optimal policy with regard to the outcome resulting from the first decision".
>
> Richard Bellman, "Dynamic Programming", Princeton Press, 1957.

To see how this important principle translates into practice, suppose that there is a value function $V(x, t) : E^n \times E^1 \rightarrow E^1$ that exits for all x and t in the relevant ranges, whose value is the maximum

value of the objective function for the control problem stated above, started at time t in state x

$$V(x,t) = \max_{u(t)\in\Omega(t)} \int_0^T F(x(t), u(t), t)\, dt + S[x(T), T]$$
$$\text{where for } s \geq t \tag{1.78}$$
$$\frac{dx}{ds} = f(x(s), u(s), s),\ x(t) = x$$

Using Bellman's principal of optimality, figure ?? provides a schematic representation of the optimal path $x^*(t)$ in state-time space using two nearby points: (x,t) and $(x+\delta x, t+\delta t)$, where δt is a small change in time and $x+\delta x = t+\delta t$. Using the principal of optimality, the change in the objective function is comprised of two parts. First, the incremental change in J from t to $t+\delta t$, which is given by the integral of $F(x,u,t)$ from t to $t+\delta t$. Second, the value function $V(x+\delta x, t+\delta t)$. The controller $u(\tau)$ should therefore be chosen so as to lie in $\Omega(\tau)$, $\tau \in [t, t+\delta t]$ and also to maximise the sum of these two terms.

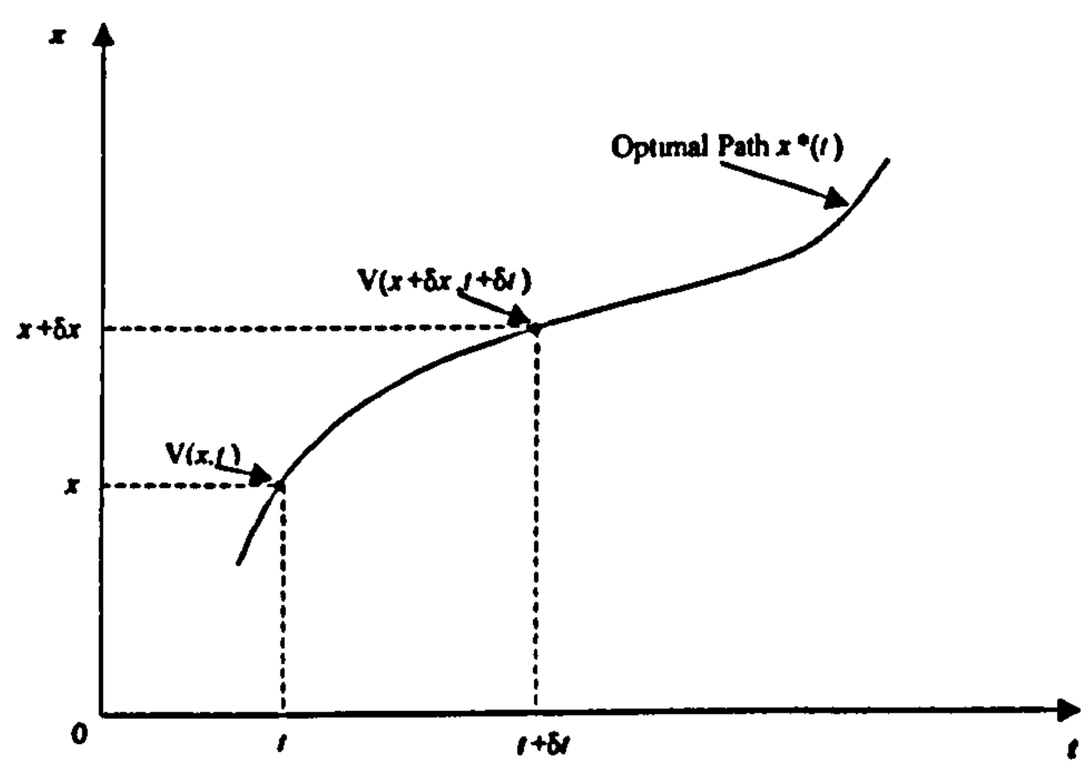

Figure 1.8: Optimal path in state-time space

Mathematically this gives rise to

$$V(x,t) = \max_{\substack{u(\tau)\in\Omega(t)\\ \tau\in[t,t+\delta t]}} \int_t^{t+\delta t} F[x(\tau), u(\tau), \tau]\, d\tau + V[x(t+\delta t), t+\delta t] \tag{1.79}$$

If V is assumed to be continuously differentiable, then employing a Taylor series expansion of V around δt and introducing an adjoint row vector $\lambda(t) \in E^n$ (where $\lambda(t)$ can be interpreted as the per unit change in the objective function for a small change in $x^*(t)$ at time t), then allows the stating of the Hamiltonian function

$$H(x,u,\lambda,t) = F(x,u,t) + \lambda f(x,u,t) \tag{1.80}$$

which finally leads to the well known Hamilton-Jacobi-Bellman (or HJB) equation

$$0 = \max_{u \in \Omega(t)} \left[H\left(x, u, V_x, t\right) + V_t \right] \tag{1.81}$$

which is the condition that must be solved to guarantee that the maximum required is satisfied. Without derivation, it is therefore possible to state the conditions for u^* to be the optimal controller as

$$\begin{gathered} \dot{x} = f\left(x^*, u^*, t\right), x^*\left(0\right) = x_0 \\ \dot{\lambda} = -H_x\left[x^*, u^*, \lambda, t\right], \lambda\left(T\right) = S_x\left[x^*\left(T\right), T\right] \\ H\left[x^*(t), u^*(t), \lambda(t), t\right] \geq H\left[x^*(t), u, \lambda(t), t\right] \end{gathered} \tag{1.82}$$

for all $u \in \Omega\left(t\right)$, $t \in [0, T]$ and where $u^*\left(t\right)$ must provide a global maximum of the Hamiltonian.

The key point to note with this standard statement of the HJB equation and the accompanying solution conditions, is that the maximum is calculated with respect to the 2-norm of the adjoint, namely with a norm of the form

$$\|z\|_2 := \sqrt{\sum_{i=1}^{n} |x_i|^2} \tag{1.83}$$

It is the use of this particular norm measure that almost single-handedly prevents simple optimal control from attaining robustness. This is because the 2-norm does not result in the optimisation being carried out with respect to the worst possible extremities that could be encountered - the search is over a restricted sub-space. To rectify this limitation involves computing the optimal controller with respect to the infinity norm

$$\|z\|_\infty := \max_{1 \leq i \leq n} |x_i| \tag{1.84}$$

such that optimality would then be guaranteed even in the face of the worst possible extremities of disturbance to the model because the search for the optimum had be carried out with respect to the entire space of possible values, in other words the optimisation is with respect to the extremum. It the idea of computing the controller with respect to the ∞–norm that is at the heart of robust optimal control and will therefore be examined in detail in the next section.

1.4.8 Robust optimal control

Having reviewed the optimal control problem, what are the principal similarities and differences between *non-robust* optimal control and *robust* optimal control ? Optimal control is a branch of mathematics that was originally developed to find optimal ways to control a dynamic system. Optimality will guarantee that for a given model, a certain objective function has been achieved subject to any stipulated constraints.

Optimality does not, however, guarantee either robustness. This because whilst a control solution may be optimal with respect to a given model, there is no guarantee that this is true for other models, even for those alternative models in close proximity to the initial model. To achieve guaranteed robustness, a control solution must satisfy the imposed constraints for an entire family of models in the neighbourhood of the initial model and the controller must be robust to disturbances. If a controller can be designed such that the system to be controlled remains stable when its parameters vary within certain expected limits, then the system is said to be robustly stable. If in addition, the controller can satisfy performance specifications such as steady state tracking of the target variables, disturbance rejection and required speed of response, then the controller is said to possess robust performance. The problem of designing controllers that satisfy both robust stability and robust performance requirements is called robust control.

Both robust and non-robust optimal control are focused on the design of control systems to achieve a target performance metric such as minimising a quadratic cost function. In the case of non-robust optimal control, the nominal model is assumed to be known with certainty, whereas in the case of robust optimal control the objective is specifically aimed at deriving a controller that is capable of guaranteeing performance and stability by the use of a dynamic feedback loop to reduce the level of uncertainty around the nominal model. This is illustrated schematically in figure 1.8.

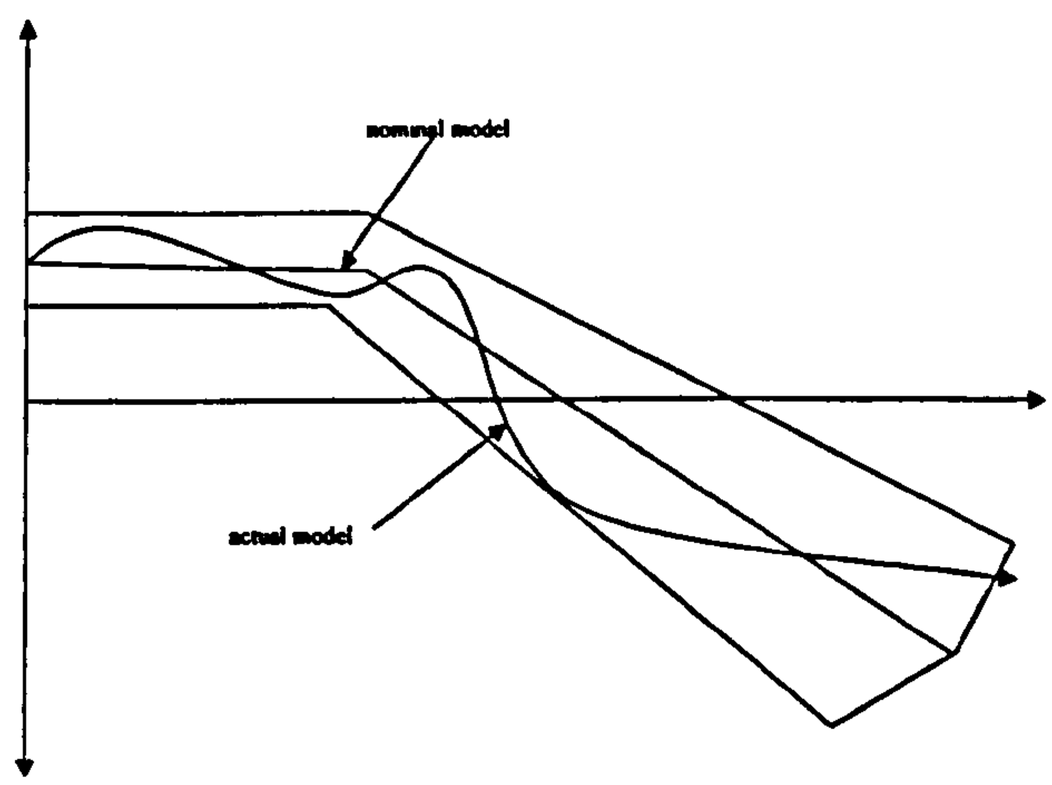

Figure 1.9: Model uncertainty

Nominal stability is achieved when the derived controller provides stable performance around the nominal model, whereas robust stability is only achieved when the controller provides stability for every model in the neighbourbood of the nominal model. Analogously, if the required performance objectives are satisfied for the nominal model, then the controller is said to possess nominal performance. Whilst if the performance objectives are satisfied for every model in the defined neighbourhood, the controller is said to possess robust performance.

1.4.9 Model uncertainty, robust control and feedback

The use of feedback mechanisms to control systems began during the 1920's. The then fledgling telephone industry had encountered an early problem with automatic control (in both theory and practice) when trying to construct feedback amplifiers whose properties remained constant despite component and supply variations. The problem was finally solved by Black in 1934 by an invention which had a tremendous impact and inspired much theoretical work. A novel approach to system stability was developed by Nyquist in 1932, the fundamental limitations of which were explored by Bode in 1940 (who also developed methods for designing feedback amplifiers, see Bode 1945). A systematic approach to designing controllers that are robust to feedback-gain variations was also developed by Bode, who described systems using transfer functions or frequency response functions, making it natural and consistent to state uncertainty in terms of deviations of the frequency responses. A number of measures such as amplitude and phase margins and maximum sensitivities were also introduced to describe robustness.

In contrast, the state-space theory that appeared during the 1960s was a fundamental paradigm shift, because it marked the beginning of systems being described using differential equations. Control design problems were formulated as optimization problems following work by Bellman in 1957 and 1958. Control of linear systems with Gaussian disturbances and quadratic criteria, the so called LQG problem, appeared particularly attractive because it admitted analytical solutions, such as those suggested by Bellman (1957) and Kalman (1960). Design computations were also improved because of the ability to build on advances in numerical linear algebra and efficient computer algorithms. The controller obtained from LQG theory also had a very interesting structure, because it was a combination of a Kalman filter and a state feedback.

State-space theory became the predominant approach. Safonov and Athans (1977) showed that the amplitude margin is infinite for an LQG problem where all state variables are measured. Unfortunately, Doyle (1978) showed that this does not hold for output feedback. There were numerous attempts to recover the robustness of state feedback using special design techniques called loop transfer recovery. The fundamental issue, however, is that it is not straightforward to capture model uncertainty in a state variable setting. Work by Horowitz and Shaked (1975) provided a focus for criticism of state-space theory which came to bear in the paper by Zames in 1981 which represented a paradigm shift that finally brought robustness to the forefront of the debate leading to the development of $H\infty$[11] theory. The idea behind $H\infty$ theory was to develop systematic design methods that were guaranteed to give stable closed loop

[11] Basar and Bernhard (1991) provide the following succinct defintion of $H\infty$:

The notation $H\infty$ stands for the Hardy space of all complex valued functions of a complex variable, which are analytic and bounded in the open right-half complex plane.

systems for systems with model uncertainty. Zames's original work was based on frequency responses and interpolation theory which led to compensators of high order. The seminal paper by Doyle, Glover and Khargonekar (1989) demonstrated how the problem could be solved using state-space methods. In an interesting and fertile departure, Basar and Bernard (1991) used game theory as an approach to $H\infty$ by formulating the problem as a game in which the objective is to find a controller in the presence of an adversary that changes the process.

Major advances in robust design were made by MacFarlane and Glover (1992) where they set the $H\infty$ control problem as a loop shaping problem, which provided effective design methods and it also reestablished the links with classical control. This line of research has been continued by Vinnicombe (1999) who obtained definite results relating modeling errors and robust control. To do this he also had to invent a novel metric for systems called the v-gap and his work brings $H\infty$ even closer to previously known classical results. This section presents the essence of the state-space method in the simple setting of single-input single-output system, before extending the approach to non-linearities.

At this point it is worthwhile noting that model uncertainty was perhaps the key motivation for introducing feedback, because classical control theory already had very effective ways of dealing with uncertainty both qualitatively and quantitatively. Process uncertainty was described very easily as a variation in the process transfer function with the qualification that disturbances do not change the number of right half plane poles or roots of the system. The theory gave important concepts and tools such as the transfer function, Nyquist's stability theory, the Nyquist curve, Bode diagrams, Bode's integrals and Bode's ideal loop transfer function. Robustness measures such as amplitude and phase margins and the maximum sensitivities were also introduced. Bode's ideal loop transfer function is probably the first design method that addressed robustness explicitly. Horowitz quantitative feedback theory is a continuation of this idea.

Before moving on to consider how feedback works in detail, it is necessary to pause to consider a number of definitions that are central to understanding model uncertainty. The first point to note is that there are four main sources of uncertainty

1. *Incomplete knowledge of a process*: This might arise due to the inability to accurately capture the parameters of the underlying process.

2. *Model simplification*: Most models involve a greater or lesser degree of simplification which is usually imposed in order to achieve tractability, so that even though a system might be known in detail, the model may have needed to have been simplified to facilitate a solution.

3. *Incomplete model structure*: Many models are linearised in the interests of tractability, which often

involves ignoring known non-linearities.

4. *Time variation in parameters*: Many systems can have different lists of parameters at different points in time.

One further distinction is also worthy of discussion, namely, that between disturbances and deviations. The former are usually external signals that are independent of the system inputs. The effect of disturbances can be aggregated as exogenous input d, at the output point of the process and added to the model input. In contrast, deviations are input (u) generated differences between the dynamics of the system and its model that lead to discrepancies. Model uncertainties can therefore be represented by an error with input z and output w.

1.4.10 Feedback amplification and robustness

One of the main reasons for using feedback is to reduce the effects of uncertainty which can appear in various alternative forms as either disturbances or other inadequacies in models of systems, be they of physical or financial processes. As the name implies, feedback works by repeatedly modifying a process in a predictable way, thereby reducing the scope and effect of uncertainty. For the effects of feedback to be maximised, the object is to find a feedback rule that produces the desired affect on the target system as quickly as possible. This means that feedback rules are chosen so as to modify the behaviour of the target system such that their impact or gain is maximised. However, using a feedback rule that has an amplified effect has the potential to also cause unstable behaviour in the underlying system for which the control rule is being developed.

Using feedback to control systems is central to the entire concept of developing processes to control other processes - a discipline known as control theory. A simple example of using a controller process to achieve a required objective is therefore useful and relevant . Figure 1.10 shows a schematic overview for a simple feedback amplifier.

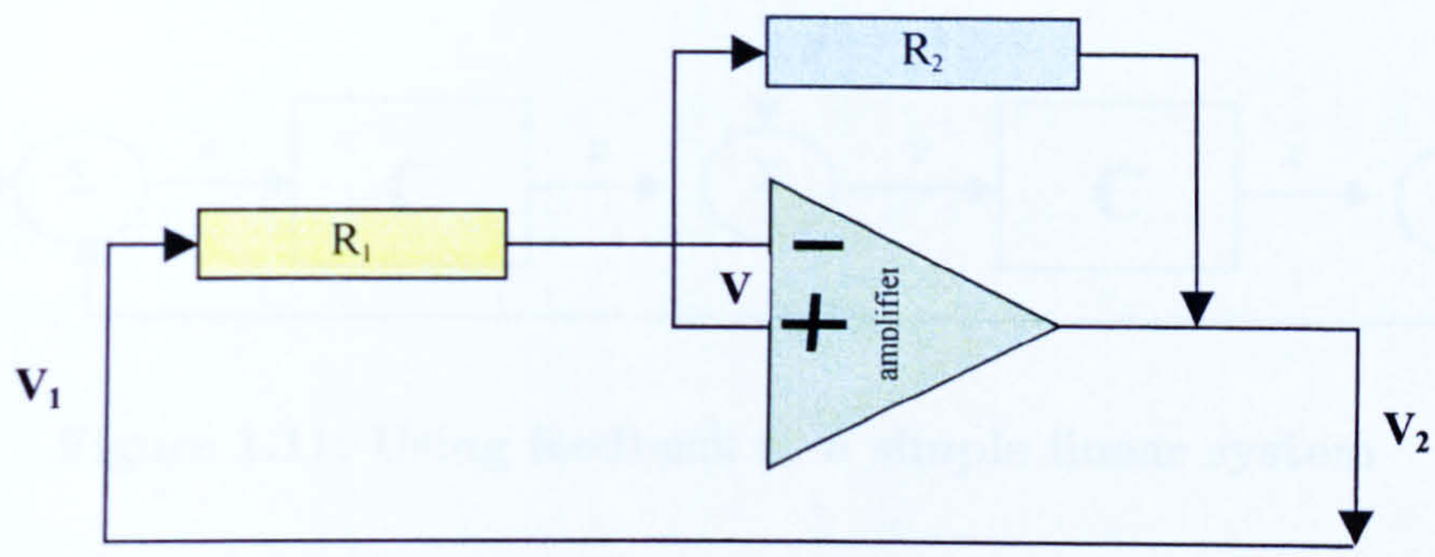

Figure 1.10: Schematic view of a feedback amplifier process

If the raw gain of the amplifier in Figure 1.10 is A, then the input-output relation of the feedback loop is given by

$$\frac{V_2}{V_1} = -\frac{R_2}{R_1}\frac{1}{1+\frac{1}{A}\left(1+\frac{R_2}{R_1}\right)} \tag{1.85}$$

where the gain as expressed by $\frac{V_2}{V_1}$ is given by the ratio $\frac{R_2}{R_1}$.The crucial point is that if the raw amplifier gain A, is large, then the gain is almost independent of the value of A. For example, if $\frac{R_2}{R_1} = 100$ and that $A = 10,000$, then a 10% change in A gives only a 0.1% variation in gain. Feedback therefore has the highly desirable property of being able to massively reduce the effects of uncertainty. However, not unsurprisingly, there is also potentially the huge risk of instability associated with this property. The trick is to use feedback to design an amplifier that is robust to variations in the gain of the process. Black (1934) achieved this objective in his work whilst at the same time producing a closed loop system that was highly linear.

1.4.11 Generalisation and the use of transfer functions

Using feedback to control behaviour is applicable to a wide range of systems. Figure 1.11 shows a basic feedback loop that consists of a process and a controller. The purpose of the system is to make the process variable x follow the set point r in spite the disturbances d and n acting on the system. The aim is that the properties of the closed loop system should also be insensitive to variations in the process. Two types of disturbances are important, namely, the load disturbance d which drives the system away from its desired state and the measurement noise n which corrupts the information about the system obtained from the system sensors. Figure 1.11 has three inputs r, d and n, in addition to four interesting signals x, y, e and u.

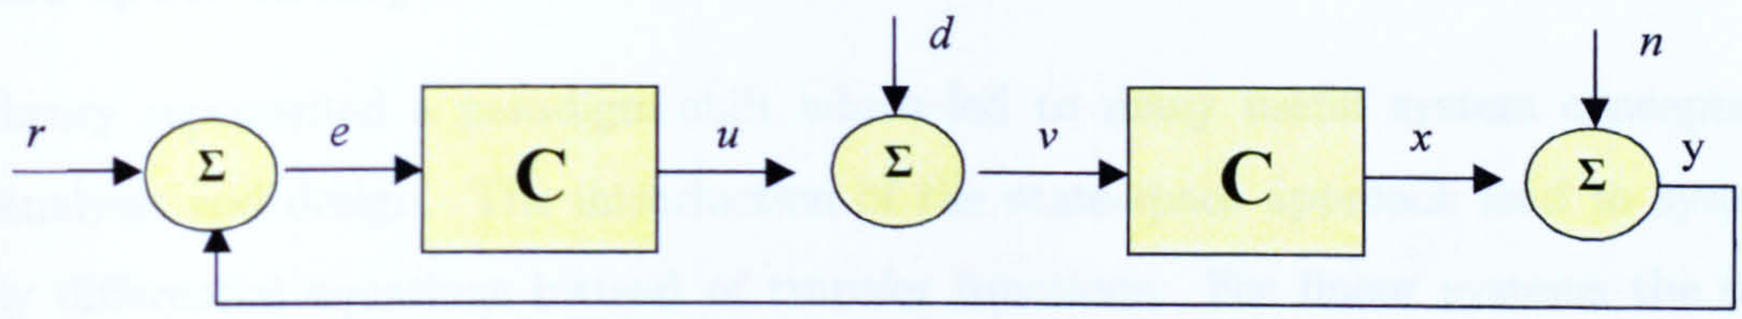

Figure 1.11: Using feedback in a simple linear system

This gives 12 relationships of interest. Assuming that the process and the controller are linear time-invariant systems that are characterized by their transfer functions P and C respectively, then the relations between the signals are given by the transfer functions

$$\begin{array}{lll} G_{xr} = \frac{PC}{1+PC} & G_{xl} = \frac{P}{1+PC} & G_{xn} = -G_{xr} \\ G_{yr} = G_{xr} & G_{yl} = G_{xl} & G_{yn} = \frac{1}{1+PC} \\ G_{er} = 1 - G_{xr} = G_{yn} & G_{el} = -G_{xl} & G_{en} = -G_{yn} \\ G_{ur} = \frac{C}{1+PC} & G_{ul} = -G_{xr} & G_{un} = -G_{ur} \end{array} \tag{1.86}$$

where G_{ij} denotes the transfer function from signal j to signal i. Clearly, there are only four independent transfer functions

$$\begin{aligned} G_{xr} &= \frac{PC}{1+PC} = T \\ G_{xl} &= \frac{P}{1+PC} \\ G_{yn} &= \frac{1}{1+PC} = S \\ G_{ur} &= \frac{C}{1+PC} \end{aligned} \tag{1.87}$$

where S is known as the sensitivity function and T is known as the complementary sensitivity function. Both S and T depend only on the loop transfer function $L = PC$ and the sensitivity functions are related by

$$S + T = 1 \tag{1.88}$$

According to Nyquist's stability criterion, the closed loop system is stable if

$$\frac{1}{2\pi} \Delta \arg_{\Gamma} (1 + L(s)) = -\mathcal{P}_{rhp}(L) \tag{1.89}$$

where $\Delta \arg$ is the argument variation when s traverses a contour Γ that encloses the right half plane and $\mathcal{P}_{rhp}(L)$ is the number of poles (or roots) of L in the right half plane.

1.4.12 State-space theory

State-space theory represented a paradigm shift which led to many useful system concepts and new methods for analysis and design. The introduction of the state-space approach lead to systems being represented by differential equations instead of transfer functions. For linear systems the state-space representation of the standard model is

$$\frac{dx}{dt} = Ax + Bu + v \tag{1.90}$$

$$y = Cx + e \tag{1.91}$$

where u is the input, y the output and x is the state. Uncertainty is represented by the disturbances v and e and by variations in the elements of the matrices A, B and C. The disturbances e and v are typically described as stochastic processes, so that the control problem can be formulated as to minimize the criterion

$$J = E \lim_{T \to \infty} \frac{1}{T} \int_0^T \left(x^T Q_1 x + u^T Q_2 u \right) dt \tag{1.92}$$

Since the equations are linear with stochastic disturbances and the criterion is quadratic, the problem became known as the linear quadratic Gaussian control problem (LQG). The solution to the control problem is given by

$$u = L(x_m - \hat{x}) + u_{ff} \tag{1.93}$$

$$\frac{dx}{dt} = A\hat{x} + Bu + K(y - C\hat{x}) \tag{1.94}$$

This control law has the highly intuitive interpretation as being feedback from the error $x_m - \hat{x}$, which also just happens to be the difference between the ideal states x_m and the estimated states $\hat{x}$, with the estimated states being provided by the Kalman filter. Controllability and observability are key conditions for solving this formulation of the problem. There are many other design methods based on the state-space formulation which give controllers with the above structure, such as pole placement. However, they differ from the LQG method in the sense that other techniques are used to obtain the matrices K and L.

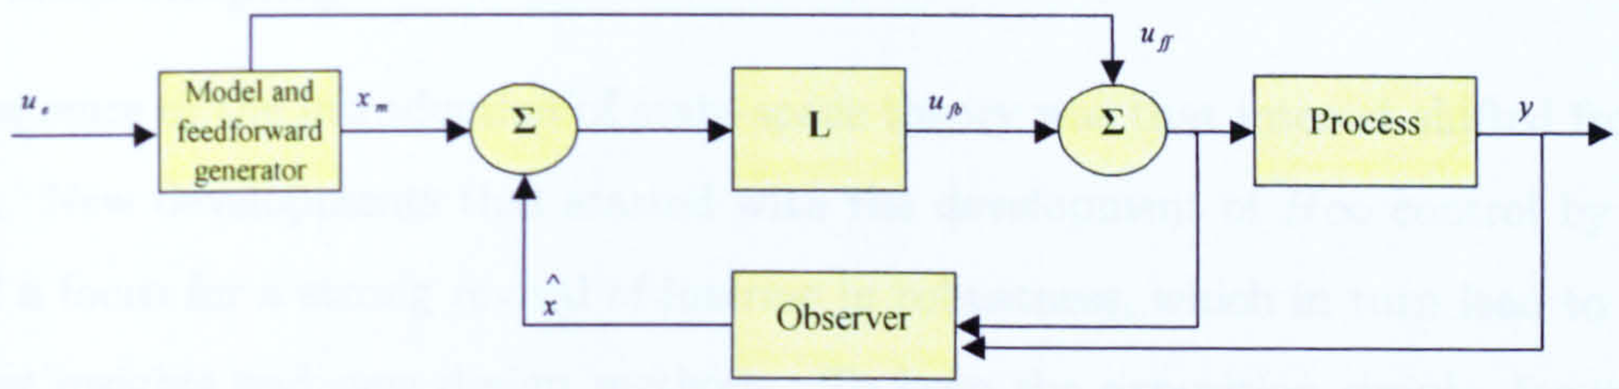

Figure 1.12: System state feeback

Figure 1.12 is known in the control systems literature as a "block diagram" and illustrates the controller obtained from LQG theory. The example system has two degrees of freedom, which is a very attractive structure. The observer or the Kalman filter delivers an estimate of the state based on a model of the system and the input and output signals of the system. It is worth pointing out that the state may also have components that represent the disturbances. There is a feedback from the deviations of the estimated state from its desired value x_m. In this simple model it is convenient to describe model uncertainties as variations in the elements of the matrices A, B and C. To do so, however, means using a very restricted class of perturbations which in turn means possibly neglected dynamics or small time delays. Such uncertainties are easier to describe in the frequency domain. LQG theory was also criticised heavily by classic control theorists because it failed to take robustness into account. However, very strong robustness properties can be established when all states are measured - which of course may present problems in the case of real physical or economic systems.

Unfortunately, the desirable robustness properties of systems with state feedback do not hold for systems with output feedback. For systems with output feedback attempts were made to recover the robustness of full state feedback by making very fast observers. This approach led to a design technique called loop transfer recovery. The only formal requirements on the system to be controlled in state-space theory is that the system is observable and controllable. There is no consideration of right half-plane poles and zeros or time delays. Because of this it becomes necessary to investigate the robustness of the design and to make appropriate modifications to achieve good robustness. This shows that it is not sufficient to check controllability and observability. Trying to design controllers which violate these limitations by making a closed-loop system that adjusts too fast results in a closed-loop system that has very poor stability margins even if the closed-loop poles are quite well damped. Feedback is not effective for disturbances having high frequencies, because disturbances will be amplified by the feedback, underlying the fact that it is important to be aware of the limitations when designing control systems.

1.4.13 $H\infty$ loop shaping

A central consequence of the introduction of state-space theory was that interest shifted from robustness to optimization. New developments that started with the development of $H\infty$ control by Zames in the 1980's provided a focus for a strong revival of interest in robustness, which in turn lead to developments that yielded new insights and new design methods. To keep the exposition simple, focus is limited to only systems having one input and one output, but techniques as well as results can be generalized to systems with many inputs and many outputs. If a system structure with two degrees of freedom is used, the problems of setpoint response can be dealt with separately and it is therefore possible to focus on robustness and attenuation of disturbances. Figure 13 shows a simplified representation of the system as having a system or plant model, P, a controller, C, two inputs, the measurement noise n and the load disturbance d.

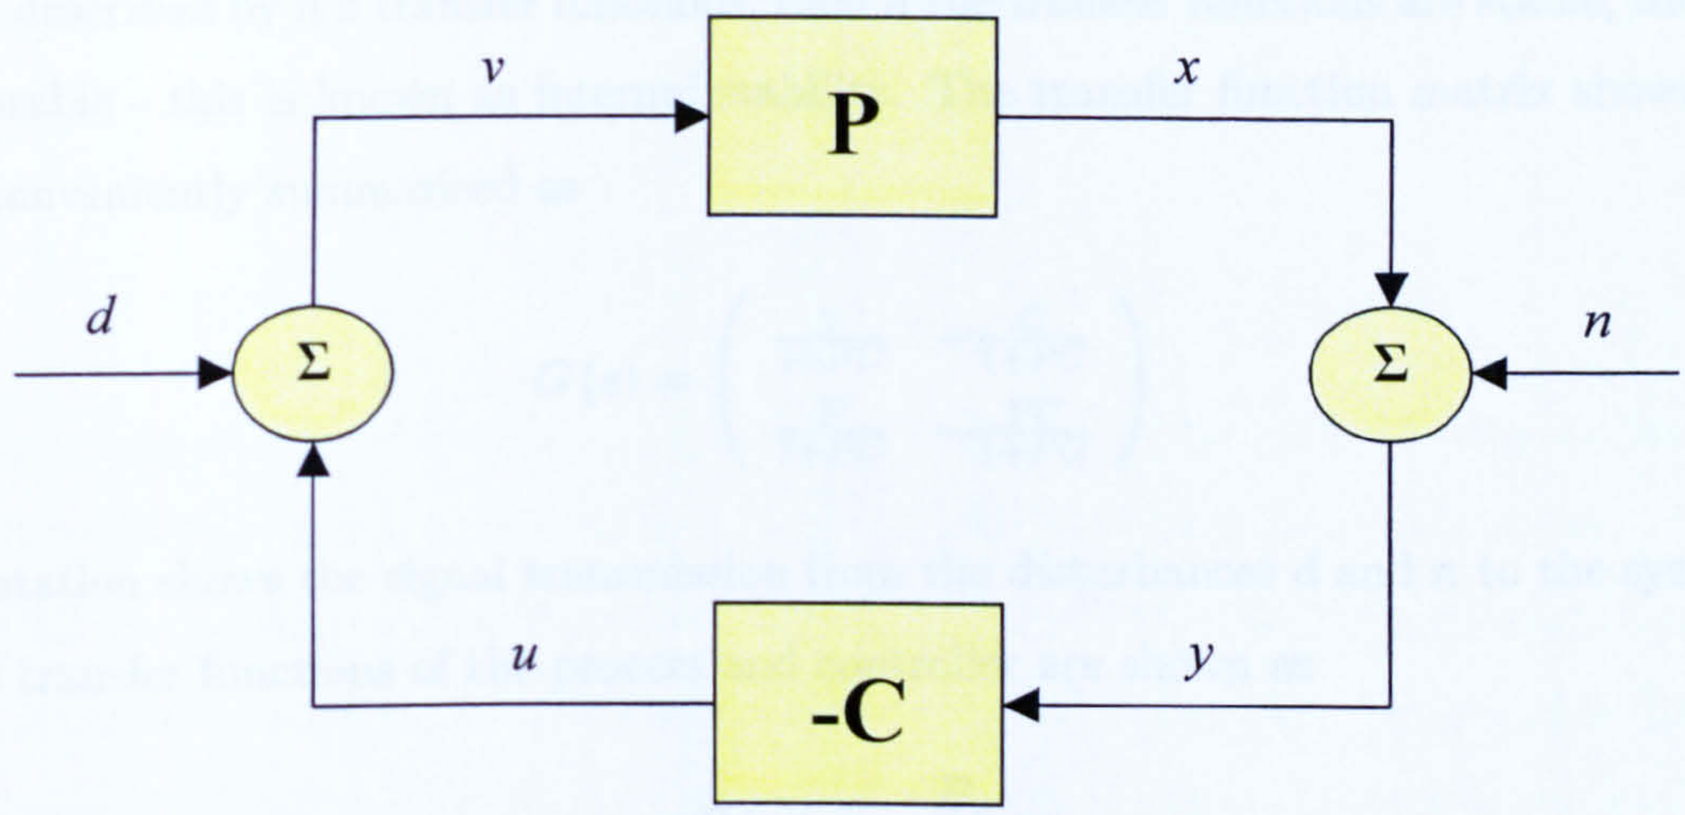

Figure 1.13: H∞ loop shaping in a simple feedback system

The problem is therefore to design a controller with the following two properties

- Insensitivity to changes in the properties of the process.
- Ability to reduce the effects of the load disturbance d. Does not inject too much measurement noise into the system.

Stability is a fundamental robustness requirement, but because stability is based solely on the loop transfer function, cancellations of poles and zeros may occur in either or both the process and the controller. This will not be problematic if the cancelled factors are stable, but results will be totally misleading if the canceled factors are unstable because internal signals in the system will diverge, rendering decision

analysis useless. To see how this may occur, consider the simple system shown in figure 13, based on

$$C(s) = \frac{s-1}{s} \tag{1.95}$$

$$P(s) = \frac{1}{s-1} \tag{1.96}$$

so that the loop transfer function is given by $L = \frac{1}{s}$, which makes the system appear stable. However, the transfer function from the disturbance, d, to output y is

$$G_{yd} = \frac{s}{(s+1)(s-1)} \tag{1.97}$$

which, upon further examination of the dynamics of the system shows that an input disturbance makes the system unstable. One way of dealing with this problem is to point out that if the closed loop system is completely described by it's transfer functions, then if the transfer functions are stable, then the system will also be stable - this is known as internal stability. The transfer function matrix shown in equation 1.86 can be conveniently summarised as

$$G(s) = \begin{pmatrix} \frac{1}{1+PC} & -\frac{C}{1+PC} \\ \frac{P}{1+PC} & -\frac{PC}{1+PC} \end{pmatrix} \tag{1.98}$$

This representation shows the signal transmission from the disturbances d and n to the system signals v and x. If the transfer functions of the process and controller are shown as

$$C(s) = \frac{B_c}{A_c} \tag{1.99}$$

$$P(s) = \frac{B_p}{A_p} \tag{1.100}$$

then 1.98 can be represented as

$$G(s) = \begin{pmatrix} \frac{A_cA_p}{A_cA_p+B_cB_p} & -\frac{B_cA_p}{A_cA_p+B_cB_p} \\ \frac{A_cB_p}{A_cA_p+B_cB_p} & -\frac{B_cB_p}{A_cA_p+B_cB_p} \end{pmatrix} \tag{1.101}$$

with the stability criterion

$$C_{pol} = A_cA_p + B_cB_p \tag{1.102}$$

which will have all of its roots in the left half plane, thereby guaranteeing internal stability.

Having considered stability, one of the remaining key problems for robustness analysis is how to compare and decide between systems when both behave similarly under a given feedback rule. For

example, it is entirely possible for two systems to exhibit similar behaviour under open loop, but different behaviour under closed loop feedback. It is also possible two systems to have different open loop behaviour, but similar closed loop behaviour. If either of these scenarios occurs, then it is clearly not possible to compare two systems by simply analysing their respective responses to input signals.

One approach to identification is to compare outputs when inputs are restricted to the class of inputs that give bounded outputs. This approach was first suggested by Zames and Sakkary (1980) and Sakkary (1985) who introduced a measure called the gap metric. However, an alternative approach called the graph metric was suggested by Vidyasagar (1985). To see how the graph metric works, assume that the process under consideration can be described by a rational transfer function

$$P(s) = \frac{B(s)}{A(s)} \tag{1.103}$$

where $A(s)$ and $B(s)$ are polynomials. Now introduce the stable polynomial $C(s)$, with a different degree to $A(s)$ and $B(s)$, such that the transfer function can be re-written

$$P(s) = \frac{B(s)/C(s)}{A(s)/C(s)} = \frac{D(s)}{N(s)} \tag{1.104}$$

then Vidyasagar's approach compares two systems by comparing the rational transfer functions D and N. The problem with the graph metric is that it can be difficult to compute in practice.

A convenient way to deal with this problem is to notice that $C(s)$ can be computed in a variety ways, one of which is known as coprime factorisation, which works as follows. The two rational functions D and N are called coprime if two conditions are met. First, that there exists rational functions X and Y such that

$$XD + YN = 1 \tag{1.105}$$

and secondly that $D(s)$ and $N(s)$ do not have any common factors. Taken together, these conditions make it possible to select $D(s)$ and $N(s)$ such that

$$DD^* + NN^* = 1 \tag{1.106}$$

where $D^*(s) = D(-s)$. Factorising equation 1.104 such that P, N and D satisfy equation 1.106, is known as a normalised comprime factorisation of P. Vinnicombe (1999) used the notion of coprime factorisation to develop a metric suitable for comparing two closely performing feedback systems. The metric works

as follows. Assume that there are two alternative systems with the following coprime factorisations

$$P_1 = \frac{D_1}{N_1} \tag{1.107}$$

$$P_2 = \frac{D_2}{N_2} \tag{1.108}$$

then using Nyquist's closed loop stability criterion of equation 1.89 to compare the two systems, it must be required that

$$\frac{1}{2\pi}\Delta \arg_\Gamma \left(N_1 N_2^* + D_1 D_2^*\right) = 0 \tag{1.109}$$

(where Γ is the Nyquist contour), so that Vinnicombe's v-gap metric is computed as

$$d = \frac{|P_1 - P_2|}{\sqrt{\left(1 + |P_1|^2\right)\left(1 + |P_2|^2\right)}} \tag{1.110}$$

which is effectively a Euclidean shortest distance between two points.

Normalised coprime factorisation can also be used to describe perturbations in a more flexible fashion than simply adopting additivity, which is essential to deal with feedback systems. Consider a system summarised by

$$P + \Delta P = \frac{N + \Delta N}{D + \Delta D} = ND^{-1} \tag{1.111}$$

where N and D are normalised coprime factorisations of P and the perturbations ΔN and ΔD are stable proper transfer functions. The interesting question is how large the perturbations can be before violating the stability condition. To see the answer to this question, figure 14 below,

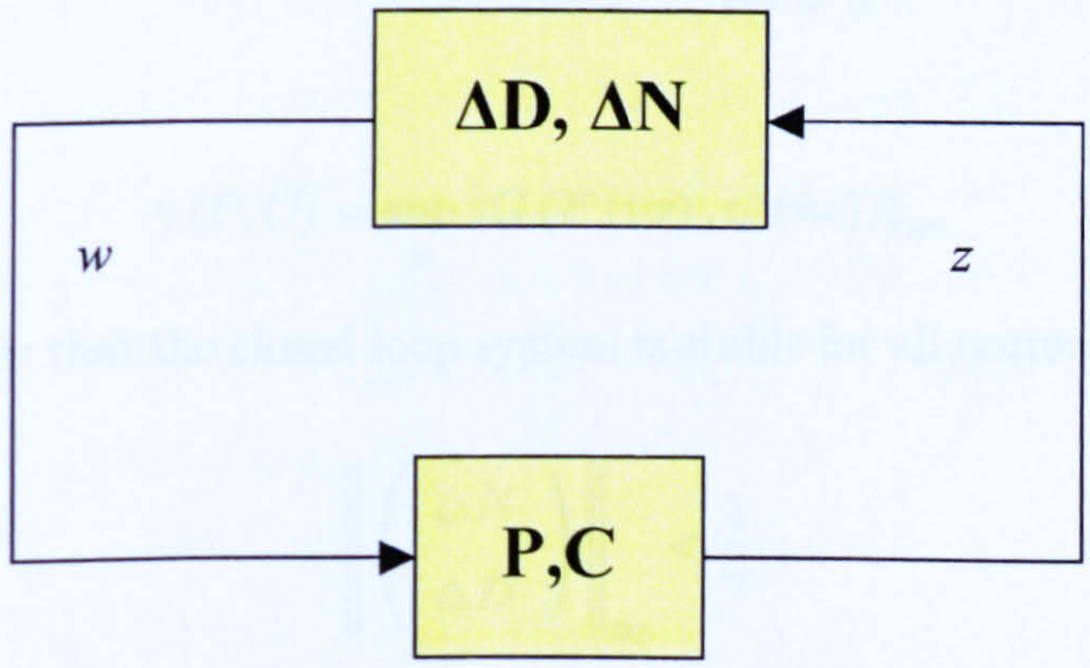

Figure 1.14: Simplied block diagram illustrating impact of perturbations

shows an extended version of the simple model of figure 13, which implies

$$z = \frac{D^{-1}}{1+PC}w_1 - \frac{D^{-1}C}{1+PC}w_2 = D^{-1}\left(\frac{1}{1+PC} - \frac{C}{1+PC}\right)\begin{pmatrix} w_1 \\ w_2 \end{pmatrix} \tag{1.112}$$

$$\text{and} \begin{pmatrix} w_1 \\ w_2 \end{pmatrix} = \begin{pmatrix} \Delta D \\ \Delta N \end{pmatrix} z \tag{1.113}$$

so that, invoking the small gain theorem[12] enables the conclusion that the perturbed system will be stable if the loop gain is less than 1, so that

$$\left\|\begin{pmatrix} \Delta N \\ \Delta D \end{pmatrix}\right\|_\infty \left\|D^{-1}\left(\frac{1}{1+PC} - \frac{C}{1+PC}\right)\right\|_\infty < 1 \tag{1.114}$$

which can be simplified by remembering that N and D are normalised coprime factorisations, to give

$$\left\|D^{-1}\left(\frac{1}{1+PC} - \frac{C}{1+PC}\right)\right\|_\infty = \left\|\begin{pmatrix} D \\ N \end{pmatrix} D^{-1}\left(\frac{1}{1+PC} - \frac{C}{1+PC}\right)\right\|_\infty \tag{1.115}$$

$$\left\|\begin{pmatrix} I \\ P \end{pmatrix}\left(\frac{1}{1+PC} - \frac{C}{1+PC}\right)\right\|_\infty = \|G(P,C)\|_\infty \tag{1.116}$$

where $G(P,C)$ is the system matrix

$$G(P,C) = \begin{pmatrix} \frac{1}{1+PC} & -\frac{C}{1+PC} \\ \frac{C}{1+PC} & -\frac{PC}{1+PC} \end{pmatrix} \tag{1.117}$$

and introducing

$$\gamma(P,C) = \sup_w \|G(P(iw), C(iw))\|_\infty \tag{1.118}$$

it is straight forward to show that the closed loop system is stable for all normalised coprime perturbations ΔD and ΔN such that

$$\left\|\begin{pmatrix} \Delta N \\ \Delta D \end{pmatrix}\right\|_\infty < \frac{1}{\gamma} \tag{1.119}$$

The quantity γ defined in equation 1.118 is a very important element in robustness. It is known as the H_∞ norm[13] and is used to evaluate the performance of the controller process in the feedback loop. The

[12] The small gain theorem states that a system is internally stable if the ∞-norm of the loop transfer function is less than 1.

[13] $\|\cdot\|$ is a norm if:

H_∞ norm has a role of central importance in this thesis, so a full consideration of it is warranted. But before considering its use, it is important to be clear on its definition. The familiar H_2 norm of a vector space used in simple optimal control is calculated as

$$\|x\|_2 := \sqrt{\sum_{i=1}^{n} |x_i|^2} \tag{1.120}$$

which is simply the square root of the sum of the absolute values of the members of the vector space, or, looked at from another perspective:

$$\|x\|_2 := \left(\int_{-\infty}^{\infty} v(t)^2 \, dt \right)^{\frac{1}{2}} \tag{1.121}$$

In contrast, the $H\infty$ norm taken as the maximum value over the entire vector space

$$\|x\|_\infty := \sup_t |v(t)| \tag{1.122}$$

which is the maximum (best or worst) value occurring over the vector space. The ∞-norm is useful when checking the boundedness of a signal from a model[14]. This means that using the H_∞ norm as a measure of performance will provide solutions that are insensitive to the worst possible variations in either disturbances or in the underlying system. It is this property that provides the key to ensuring robustness.

Zames (1981) was the first to suggest using the H_∞ norm as opposed to the more familiar and simple H_2 norm to evaluate performance. Doyle and Stein (1981) showed that model uncertainties can be described as norm bounded deviations, which, when taken together with the H_∞ norm, provide an ideal means of defining and controlling the robustness of a controller process. The objective of the H_∞ approach is therefore to design control systems that are insensitive to model uncertainty. This aim is achieved by deriving a controller C that gives a stable closed-loop system, whilst simultaneously minimising the H_∞ norm of the transfer function $(G(P,C))$. The obvious corollary is that the H_∞ approach also allows the largest deviation of the normalised coprime deviations.

---

1. $\|x\| \geq 0$
2. $\|x\| = 0$ if and only if $x = 0$
3. $\|\alpha x\| = |\alpha| \, \|x\|$ for any scalar α
4. $\|x + y\| \leq \|x\| + \|y\|$

[14]For both norms it is possible to define the linear space of signals $v(t)$ that has a bounded value for either norm. These function spaces are called $\mathcal{L}_2$ and $\mathcal{L}_\infty$ or Lebesgue spaces.

As already indicated, the link between stability and robustness is critical, so it is also worth noting that the H_∞ norm can also be used to as the basis of a generalised stability condition, namely, that

$$b(P,C) = \begin{cases} \frac{1}{\gamma} & \text{if } (P,C) \text{ is stable} \\ 0 & \text{otherwise} \end{cases} \tag{1.123}$$

which produces values in the range 0 to 1, with a value of 0 indicating system instability, whilst a value close to 1 implies a good margin of stability. Reasonable values that indicate stability would normally be expected to lie in the range 0.25 - 0.33. It is also clearly the case that the controller that maximises the stability margin is

$$b_{opt} = \sup_C b(P,C) \tag{1.124}$$

1.4.14 Solving the $H\infty$ state-space robust optimal control problem

In order to distinguish between the transfer function based approach and the state-space based approach to solving the $H\infty$ optimal control problem, it is convenient and helpful to employ a modified set of names for the model variables. As already explained, the performance measure to be minimised is the $H\infty$ norm of the closed loop transfer function

$$J_\infty(K) = \|F(G,K)\|_\infty \tag{1.125}$$

where G is the model or system under consideration, K is the compensator or controller and $F(G,K)$ is the transfer function of the closed loop system: $z = F(G,K)\,v$

Before proceeding, it is important to recall the distinction between disturbances and deviations, which is clearly illustrated in figure 15, where r is the reference output (for measurement purposes), m is actual output, e is an error signal, d is an exogenous input disturbance, y is the final output (actual plus disturbance), u is the input, model uncertainties can be represented by an error model with input z and output w, . It is arguably the case that the robust optimal control problem can be most convenient solved in the time domain by using the state-space approach. To illustrate the solution approach, assume that the underlying model, G, has the following state-space representation

$$\dot{x}(t) = Ax(t) + B_1 v(t) + B_2 u(t), \qquad x(0) = 0 \tag{1.126}$$

$$z(t) = C_1 x(t) + D_{12} u(t) \tag{1.127}$$

$$y(t) = C_2 x(t) + D_{21} v(t) \tag{1.128}$$

Clearly, the only relevant inputs are v and u, assuming a zero initial state. Based on the following series

of assumptions concerning the parameter matrices

1. The pair (A, B_2) is stabilisable.

2. $D_{12}^T D_{12}$ is invertible.

3. $D_{12}^T C_1 = 0$

4. The pair (C_1, A) has no unobservable modes on the imaginary axis.

5. (C_2, A) is detectable.

6. $D_{12} D_{12}^T$ is invertible.

7. $D_{12} B_1^T = 0$

8. The pair (A, B_1) has no uncontrollable modes on the imaginary axis.

The direct minimisation of the cost $J_\infty(K)$ is a very hard problem to solve directly. The way to approach a solution is therefore, to adopt the easier route of constructing conditions which define whether there exists a stabilising controller that achieves the $H\infty$ norm bound

$$J_\infty(K) < \gamma \tag{1.129}$$

for a given $\gamma > 0$.

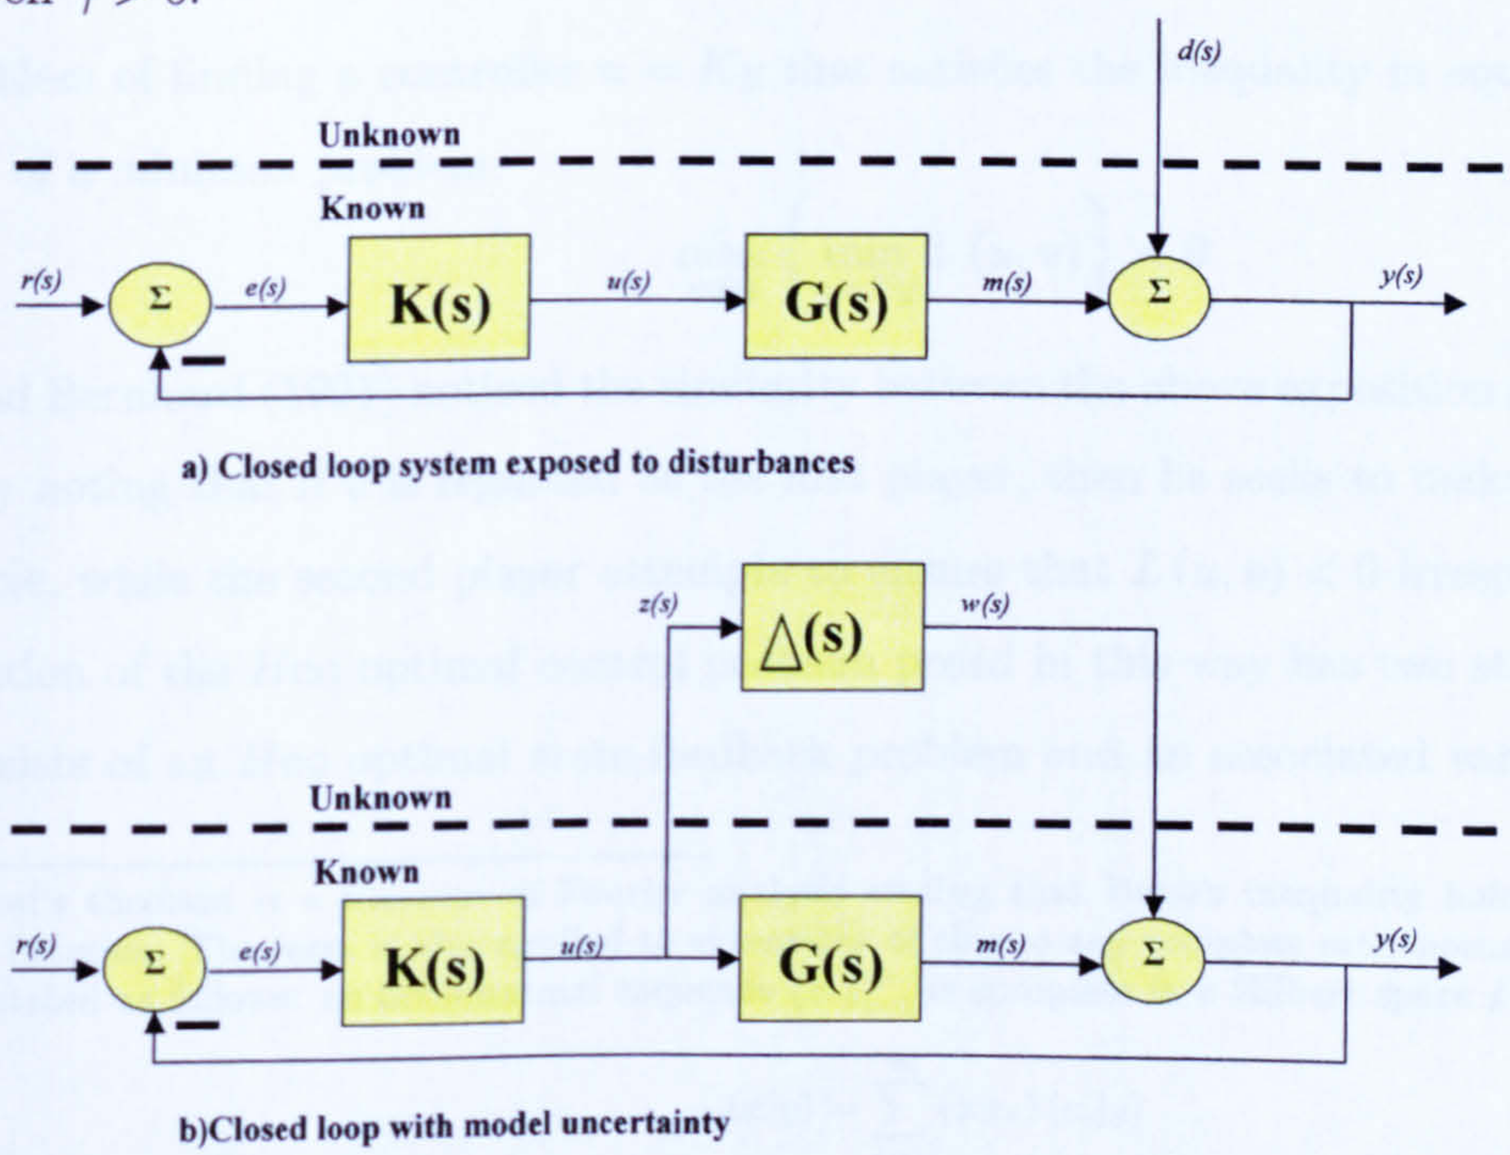

Figure 1.15: Disturbances and model uncertainty in a closed loop system

The conditions also provide a specific controller that achieves the bound of equation 1.129, which can then be used for various values of γ to iteratively determine the minimum of $J_\infty(K)$ to any required degree of accuracy. The construction of such controllers that satisfy the bound of equation 1.129 is generally referred to as the $H\infty$ optimal control problem. As already indicated, it is a hard computational problem, so in order to derive conditions for checking whether there exists a controller that achieves the bound it is useful to take advantage of the fact that the $H\infty$ performance measure can be characterised in terms of the worst case gain in therms of the $\mathcal{L}_2$ norm

$$J_\infty(K) = \sup\left\{\frac{\|z\|_2}{\|v\|_2} : v \neq 0\right\} \tag{1.130}$$

which is equivalent to

$$\frac{\|z\|_2}{\|v\|_2} < \gamma, \qquad \text{all } v \neq 0 \tag{1.131}$$

or

$$L(u,v) = \|z\|_2^2 - \gamma^2 \|v\|_2^2 < 0, \qquad \text{all } v \neq 0 \tag{1.132}$$

As the $H\infty$ optimal controller will be derived in the time domain, it is helpful to give an explicit time-domain expression of this latter inequality using Parseval's theorem[15]

$$L(u,v) = \int_0^\infty \left[z(t)^T z(t) - \gamma^2 v(t)^T v(t)\right] dt < 0, \qquad \text{all } v \neq 0 \tag{1.134}$$

The problem of finding a controller $u = Ky$ that satisfies the inequality in equation 1.134 can be stated in terms of a minimax problem

$$\max_{v \neq 0}\left\{\min_{u=Ky} L(u,v)\right\} < 0 \tag{1.135}$$

Basar and Bernhard (1991) noticed the similarity between the above exposition and the theory of dynamic games by noting that if v is regarded as the first player, then he seeks to make the cost $L(u,v)$ as large as possible, while the second player attempts to ensure that $L(u,v) < 0$ irrespective of the actions of v. The solution of the $H\infty$ optimal control problem posed in this way has two stages, namely, a first stage that consists of an $H\infty$ optimal state-feedback problem and an associated variable transformation, and

[15] Parseval's theorem is a theorem of Fourier analysis stating that Bessl's inequality holds as an equality for a square integrable function. The term is also applied to extensions of this to any complete orthonormal sequence in a Hilbert space. It can be stated as follows: an orthonormal sequence $\{e_i\}_{i=1}^\infty$ is complete in a Hilbert space H if and if for each x, y in H

$$(x|y) = \sum_{i=1}^\infty (x|e_i)(e_i|y) \tag{1.133}$$

a second stage that is an $H\infty$ optimal estimation problem. It is necessary to examine each step in turn.

The solution of the $H\infty$ optimal state-feedback problem gives the optimal state-feedback law $u(s) = K_\infty(s)\,x(s)$ that satisfies the $H\infty$ norm bound of equation 1.129, provided such a controller exists. Assuming that the controller has access to both present and historical state information, then there exists a state-feedback controller that satisfies equation 1.129, if and only if there exists a positive semi-definite solution to the algebraic Riccati equation

$$A^T X + XA - XB_2 \left(D_{12}^T D_{12}\right)^{-1} B_2^T X + \frac{XB_1 B_1^T X}{\gamma^2} + C_1^T C_1 = 0 \tag{1.136}$$

such that the matrix

$$A - B_2 \left(D_{12}^T D_{12}\right)^{-1} B_2^T X + \frac{XB_1 B_1^T X}{\gamma^2} \tag{1.137}$$

is stable (i.e. all negative eigenvalues have real parts). When the above conditions are satisfied, a controller that achieves the $H\infty$ norm bound of equation 1.129 is given by the static state-feedback controller

$$u(t) = K_\infty x(t) \tag{1.138}$$

where

$$K_\infty = -\left(D_{12}^T D_{12}\right)^{-1} B_2^T X \tag{1.139}$$

Assuming that the algebraic Riccati equation 1.136 has a positive semi-definite solution X such that the matrix in equation 1.137 is stable, then the right-hand side of equation 1.134 can be expanded as follows

$$\begin{aligned}
\int_0^\infty \left[z(t)^T z(t) - \gamma^2 v(t)^T v(t)\right] dt &= \int_0^\infty \left[\left[C_1 x(t) + D_{12} u(t)\right]^T \left[C_1 x(t) + D_{12} u(t)\right] - \gamma^2 v(t)^T v(t)\right] dt \qquad (1.140)\\
&= \int_0^\infty \begin{bmatrix} \left[u(t) - u^0(t)\right]^T D_{12}^T D_{12} \left[u(t) - u^0(t)\right] \\ -\gamma^2 \left[v(t) - v^0(t)\right]^T \left[v(t) - v^0(t)\right] \end{bmatrix} dt \qquad (1.141)
\end{aligned}$$

where

$$\begin{aligned}
u^0(t) &= -\left(D_{12}^T D_{12}\right)^{-1} B_2^T X x(t) && (1.142)\\
&= K_\infty x(t) && (1.143)\\
v^0(t) &= \frac{B_1^T X x(t)}{\gamma^2} && (1.144)
\end{aligned}$$

Equation 1.141 can be used to transform the $H\infty$ optimal output feedback problem into a $H\infty$ optimal

estimation problem by defining the following signals

$$\begin{aligned}
\tilde{z}(t) &= \left(D_{12}^T D_{12}\right)^{\frac{1}{2}}\left[u(t)-u^0(t)\right] && (1.145)\\
&= \left(D_{12}^T D_{12}\right)^{\frac{1}{2}} u(t) + \left(D_{12}^T D_{12}\right)^{-\frac{1}{2}} B_2^T X x(t) && (1.146)\\
\tilde{v}(t) &= v(t) - v^0(t) && (1.147)\\
&= v(t) - \frac{B_1^T X x(t)}{\gamma^2} && (1.148)
\end{aligned}$$

which gives

$$L(u,v) = \int_0^\infty \left[z(t)^T z(t) - \gamma^2 v(t)^T v(t)\right] dt = \int_0^\infty \left[\tilde{z}(t)^T \tilde{z}(t) - \gamma^2 \tilde{v}(t)^T \tilde{v}(t)\right] dt \qquad (1.149)$$

or

$$\tilde{L}(u,v) = \|\tilde{z}\|_2^2 - \gamma^2 \|\tilde{v}\|_2^2 < 0, \qquad \text{all } \tilde{v} \neq 0 \qquad (1.150)$$

which implies the transformed system

$$\begin{aligned}
\dot{x}(t) &= \left(A + \frac{B_1 B_1^T X}{\gamma^2}\right) x(t) + B_1 \tilde{v}(t) + B_2 u(t), \qquad x(0) = 0 && (1.151)\\
\tilde{z}(t) &= \left(D_{12}^T D_{12}\right)^{-\frac{1}{2}} B_2^T X x(t) + \left(D_{12}^T D_{12}\right)^{\frac{1}{2}} u(t) && (1.152)\\
y(t) &= \left(C_2 + \frac{D_{21} B_1^T X}{\gamma^2}\right) x(t) + D_{21} \tilde{v}(t) && (1.153)
\end{aligned}$$

The point is that in the output-feedback case, the base that can be achieved is to base the controller on an estimate of $x(t)$ or of the output $\tilde{z}(t)$, which naturally leads to the second stage, namely, the $H\infty$ optimal estimation problem, which can most easily be understood and a solution methodology derived by slightly simplifying the initial system of equations 1.126, 1.127 and 1.128 as follows

$$\begin{aligned}
\dot{x}(t) &= Ax(t) + B_1 v(t), \qquad x(0) = 0 && (1.154)\\
z(t) &= C_1 x(t) && (1.155)\\
y(t) &= C_2 x(t) + D_{21} v(t) && (1.156)
\end{aligned}$$

If F represents the stable causal estimators of the output z based on the measured output y such that $\hat{z}(s) = F(s)y(s)$, then in the $H\infty$ optimal estimation problem, define the $H\infty$ norm of the transfer function from the disturbance v to the estimation error $e = z - \hat{z}$, as

$$J_{e,\infty}(F) = \sup\left\{\frac{\|z-\hat{z}\|_2}{\|v\|_2} : v \neq 0\right\} \qquad (1.157)$$

Analogously with the state-feedback problem, consider the conditions for the existence of an estimator that achieves the $H\infty$ bound

$$J_{e,\infty}(F) < \gamma \tag{1.158}$$

which can also be equivalently stated as

$$L_e(v, \hat{z}) = \|z - \hat{z}\|_2^2 - \gamma^2 \|v\|_2^2 < 0, \qquad \text{all } v \neq 0 \tag{1.159}$$

The solution to which, though far too lengthy and complex to derive here, can be characterised in much the same way as for the optimal state-feedback problem by stating that there exists a stable estimator F that achieves the $H\infty$ norm bound of equation 1.158, if and only if there exists a symmetric positive definite or positive semi-definite solution Y to the algebraic Riccati equation

$$AY + YA^T - YC_2^T \left(D_{12}D_{12}^T\right)^{-1} C_2 Y + \frac{YC_1^T C_1 Y}{\gamma^2} + B_1 B_1^T = 0 \tag{1.160}$$

such that the matrix

$$A - YC_2^T \left(D_{12}D_{12}^T\right)^{-1} C_2 + \frac{YC_1^T C_1}{\gamma^2} \tag{1.161}$$

is stable, so that all its eigenvalues have negative real parts. When these conditions are satisfied, an estimator that achieves the bound $J_{e,\infty}(F) < \gamma$, is given by

$$\dot{\hat{x}}(t) = A\hat{x}(t) + L_\infty \left[y(t) - C_2\hat{x}(t)\right], \qquad \hat{x}(0) = 0 \tag{1.162}$$

$$\hat{z} = C_1\hat{x}(t) \tag{1.163}$$

where

$$L_\infty = YC_2^T \left(D_{21}D_{21}^T\right)^{-1} \tag{1.164}$$

The final stage is therefore to bring together the state-feedback result and the estimator result to provide a solution to $H\infty$ optimal controller problem for the system shown in figure ??. To begin with, recall that the $H\infty$ norm bound $J_\infty(K) < \gamma$ holds if and only if the inequality $\tilde{L}(u, v) = \|\tilde{z}\|_2^2 - \gamma^2 \|\tilde{v}\|_2^2 < 0$ holds for the associated system. Therefore, if assumptions 1-8 hold, then there exists a controller $u = Ky$ which achieves the $H\infty$ norm bound $J_\infty(K) < \gamma$ and it is possible to characterise this $H\infty$ optimal controller for the system described in equations 1.126, 1.127 and 1.128 if and only if the following conditions are satisfied

Theorem 9 *There exists a symmetric positive definite or semi-definite solution to X to the Riccati equation 1.136, such that the matrix in equation 1.137 is stable.*

Theorem 10 *There exists a symmetric positive definite or semi-definite solution Z to the algebraic Riccati equation associated with the given system*[16] *and the estimation performance bound*

$$\tilde{A}Z + Z\tilde{A}^T - Z\tilde{C}_2^T\left(\tilde{D}_{21}\tilde{D}_{21}^T\right)^{-1}\tilde{C}_2 Z + \frac{Z\tilde{C}_1^T\tilde{C}_1 Z}{\gamma^2} + \tilde{B}_1\tilde{B}_1^T = 0 \tag{1.165}$$

such that all of the eigenvalues of the matrix

$$\tilde{A} - Z\tilde{C}_2^T\left(\tilde{D}_{21}\tilde{D}_{21}^T\right)^{-1}\tilde{C}_2 + \frac{Z\tilde{C}_1^T\tilde{C}_1}{\gamma^2} \tag{1.166}$$

have negative real parts. When these conditions are satisfied, a controller that achieves the $J_\infty(K) < \gamma$ performance bound is given by

$$\dot{\hat{x}}(t) = \tilde{A}\hat{x}(t) + \tilde{B}_2 u(t) + L_Z\left[y(t) - \tilde{C}_2\hat{x}(t)\right] \tag{1.167}$$

$$u(t) = K_\infty\hat{x}(t) \tag{1.168}$$

where

$$K_\infty = -\left(\tilde{D}_{21}\tilde{D}_{21}^T\right)^{-1}B_2^T X \tag{1.169}$$

$$\textit{and} \tag{1.170}$$

$$L_Z = Z\tilde{C}_2^T\left(\tilde{D}_{21}\tilde{D}_{21}^T\right)^{-1} \tag{1.171}$$

The $H\infty$ optimal controller therefore consists of a $H\infty$ optimal estimator and a $H\infty$ optimal state-feedback on the state of the optimal estimator. As has been seen already, the estimator depends on the optimal state-feedback controller via the transformed system, such that the separation principle valid in the H_2 optimal control problem no longer holds. Notwithstanding this, it turns out that it is not necessary to solve the estimator Riccati equation associated with the transformed system (i.e. equation), but instead it is sufficient to solve the Riccati equation associated with the untransformed system due to the following relationship between the solutions X, Y and Z of equations 1.136, 1.160 and 1.165. If the the Riccati equation 1.136 has a symmetric positive semi-definite solution X, then the Riccati equation 1.165 has a symmetric positive semi-definite solution Z if and only if

a) there exists a symmetric positive semi-definite solution Y to equation 1.160; and

b) $\rho(XY) < \gamma^2$, where the term $\rho(XY)$ is the maximum eigenvalue of XY.

[16] *i.e. the system described by the equations 1.126, 1.127 and 1.128.*

When these two conditions hold, then the solution to Z of equation 1.165 is given by

$$Z = Y\left(I - \frac{XY}{\gamma^2}\right)^{-1} \tag{1.172}$$

The above analysis and construction only achieves a controller that attains the $H\infty$ bound for a required performance level $\gamma > 0$. To make the closed loop $H\infty$ norm as small as possible, it is necessary to iterate on γ until the required degree of accuracy has been attained. This minimum achievable $H\infty$ norm is usually referred to as γ_{inf} (infimum) as it is the greatest lower bound of all possible γ

$$\gamma_{\text{inf}} = \inf\{\|F(G,K)\|_\infty : u = Ky, \qquad K \text{ stabilising}\} \tag{1.173}$$

The final step to consider is the need to derive a controller that achieves a closed loop $H\infty$ norm which exceeds the minimum achievable norm by some required tolerance level $\delta > 0$, such that $\|F(G,K)\|_\infty \leq \gamma_{\text{inf}} + \delta$. This can be accomplished easily using a simple bisection strategy to iterate to the required level of accuracy, this iterative process is known as γ-iteration.

1.4.15 Robust optimal control and decision making

The second half of the literature review introduced and reviewed the concepts and theory associated with stability, robustness and optimal control. The next step is therefore to review how robustness has been applied to the problem of decision making in the presence of uncertainty. At first sight, there has been relatively little direct research involving the use of robustness in decision making. However, the area is not totally devoid of research findings. Research on robustness and decision making divides fairly neatly into two broad lines of enquiry. One direction has been the work on robust macro-economic policy rules aimed at deriving decision rules for macro-economic policy actions, most notably in the area of robust monetary policy decisions. The other direction has concentrated on applying the game theoretic route to solving the $H\infty$ robust optimal control problem to price options.

Robust, optimal macro-economic policy rules

The macro-economic research has mainly focused around work originally carried out by William Brainard in (1967) which was based on the notion of treating policy selection and therefore decision making, in an uncertain world as an optimal portfolio choice problem. Brainard's ideas arose after discussions with Arthur Okun in 1962, but only appeared in his seminal 1967 article, where Brainard was the first to deal with uncertainty in the construction of policy rules in a systematic fashion, by identifying the model, data and parameter uncertainties faced by those formulating and carrying out policy. His work on uncertainty

and the effectiveness of policy provided the first formal statement of the idea that uncertainty may reduce policy responsiveness because decision-makers back away from mechanistically applying optimal policy rules when faced with incomplete information.

Over the past 30 years, this simple notion has been used as a powerful heuristic, continuing to underpin the basic monetary policy process at many of the world's central banks. So much so in fact, that a paper by Martin and Salmon (1999) takes the notion even further, grading the degree of responsiveness from conservatism, to gradualism, through to caution. Brainard's original idea was that optimum monetary policy computed in a deterministic fashion (using a Tinbergen-Theil style approach) should, because of the combined effects of model and parameter uncertainties, be applied cautiously to attain chosen policy targets. This is so, Brainard argued, even if such action may result in an overall worse outcome. In short, the idea is that blind application of a deterministic optimal rule could in certain circumstances, actually exacerbate uncertainty over the outcome for the economy as a whole. As Blinder (1999) puts it, monetary policy makers should, "compute the direction and magnitude of their optimal policy move in the way prescribed by Tinbergen-Theil and then do less - a little stodginess at the central bank is entirely appropriate".

Brainard commences his exposition with a single-target/single-instrument approach using the following simple model:

$$y = aP + u \tag{1.174}$$

where y is the target variable, P is a policy instrument, u is a vector of exogenous variables and a determines the response of y to policy actions. In Brainard's world, the policy maker faces two principal kinds of uncertainty. First, uncertainty about the impact of u on y - which reflects the inability to perfectly forecast the value of exogenous variables, or the response to changes in their values. Second, uncertainty of the response of y to a policy change. Both types of uncertainty mean that the policy maker will be unable to move y to its target value y^*, but they have significantly different implications for policy action. Clearly, the first type is endogenous and is beyond the control of the policy maker. However, the second is the very reason for the principle of certainty equivalence, namely that the policy maker should act on the basis of expected values as if he were certain that they will actually occur.

The essential element of the Brainard style approach is one of a traditional optimisation (usually minimising) of some (generally quadratic) target loss function subject to a series of constraints. Typically, the loss function usually turns out to be some measure of excess inflation, with interest rates being the target variable. Brainard states his optimisation target cost function as follows:

$$U = (y - y^*)^2 \tag{1.175}$$

Brainard accepts the obvious practical objection to the above formulation. Namely, the symmetrical treatment of positive and negative deviations from target, as both are assumed to be equally important. To illustrate the linkage between policy action and the volatility of y, Brainard uses the simple assumption that a is a random variable dependent on some unobserved variables. a is also assumed to be correlated with u. Then, given the well known expression for the variance of a sum of two random variables and the assumption that a and u are independent, Brainard's model can be restated in terms of how much of the expected policy gap should be filled by policy action:

$$P^* = \frac{g}{(1+V^2)} \tag{1.176}$$

where

$$g = \frac{(y^* - \bar{u})}{\bar{a}} \tag{1.177}$$

and $V = \sigma_a/\bar{a}$. Only if the policy maker is certain (i.e. $V = 0$) will policy be aimed at closing the gap between u and y^*.

Brainard's work to estimate optimal control rules is based on simple non-stochastic minimisation of a linear-quadratic cost function subject to linear constraints, the so called linear-quadratic or LQ approach. Brainard assumes that the decision maker knows the model and associates a distinct optimal decision rule with each model specification. As such, the LQ approach makes three distinct assumptions. First, it takes no account of model misspecification or any uncertainties regarding parameter misspecification within the model. Second, it takes no account of disturbances or perturbations to the model and third it assumes that all required information is fully available at the time of decision making. Having used the LQ approach to generate an optimal control rule, there is no method for testing how well a control rule produced from one model performs in other similar but distinctly different models. In other words, there is no unambiguous method for establishing the robustness of the performance of the optimal control rule.

In a distinctly rational expectations vein, Sargent (2001) has written a substantial monograph about decision makers who doubt their model, drawing the distinction between these decision makers and those in a rational expectations world, where the model is assumed to be known with certainty. Sargent's focus is on the issue of the decision maker being concerned about the performance of his model in comparison to other nearby models and the extension of rational expectations models to include a fear of model misspecification. Sargent cites three reasons for extending rational expectations models to include a fear of misspecification. First, on purely behavioural grounds both observers and decision makers should be allowed to have uncertainty about models enter their analysis. Second, studying robustness of decisions may help to shed further light on the inability of current models to explain current market phenomena

such as the equity risk premium and the costs of business cycles[17]. Third, that a body research reaching back to Friedman (1953) and Brainard (1967) advocates framing policy rules in the light of doubts about model misspecification.

Sargent also points to a powerful feature of rational expectations models, namely,

> **"that they incorporate a mapping from probabilistic laws of motion for shocks to decision rules"**
>
> (Sargent, 2001, p.2)

He goes on to state that the mapping is implied by the assumption that decision makers use ordinary LQG control theory. LQG assumes that decision makers know the model in the form of a transition law linking state variables and controls, including a stochastic description of the shocks. LQG associates a distinct decision rule with each specification of the shock processes. Sargent also points to the fact that many rational expectations models flow from this association. He provides the specific example that the centre-piece of the Lucas Critique (1976) is the finding that under rational expectations, decision rules are functions of the serial correlation of shocks. Further, he also points to work by Salmon (1997) on how the econometric encapsulation of rational expectations achieves parameter identification by exploiting the structure of the function mapping shock serial correlation properties to decision rules.

Sargent (2001) applies robust control theory to illustrate the loosening of the mapping from shock temporal properties to decision rules, by showing how the former treats models as approximations, thereby trying to identify a single decision rule that works well over an entire set of models. Alternative models are loosely specified in terms of (possibly) serially correlated shifts in the means of the shock processes. Sargent argues that the shifts to the shocks can feedback arbitrarily on the history of the states, thereby representing mis-specified dynamics. As robust decision making implies a preference for rules that work across a range of model specifications, Sargent argues that the Lucas critique is therefore suspended for situations that are shifts within a set of models surrounding the approximating model.

Using the LQG approach, Sargent argues that the robustness specification provides an alternative means of sweeping a decision maker's doubts about the model unto an altered objective function. He uses invariant subspace methods for solving LQG problems to compute robust decision rules as solutions to zero sum two player games. Sargent uses two methods of altering the objective function for an ordinary LQG control problem to produce robust decision rules. The first is to use the known relationship between

[17]The equity risk premium acts as an indicator of portfolio return expectations and as such influences asset allocation policy. Sargent's premise is that a possible explanation for the persistent under measurement of the premium is that decision makers doubt their economic model and so may act rationally by not investing as much in equities as would be predicted by the standard captial asset pricing model, resulting in lower demand and price, therefore reinforcing the effect. The result is a lower equity risk premium.

the LQG problem and the two-player zero sum game where nature selects from a set of models in such as way as to induce the decision maker to opt for robust decision rules; whilst the second involves modifying the value function or indirect utility function in a zero-sum two player game based explicitly on the use of the $H\infty$ criteria.

Game theoretic analysis of options

At its simplest, game theory is the study of multi-person decision problems. The key development came in 1928 with the proof of the minimax theorem by von Neumann. Until the linkage between game theory and the LQG approach was noticed and exploited, game theory faced severe methodological limitations in handling uncertainty and timing decisions in dynamic models. Once the linkage with the state-space $H\infty$ solution to the robust optimal control problem was noticed and developed in the 1960's, interest increased, but the essential reference and wider distribution of the ideas is mainly attributable to the book by Basar and Olsder written in 1982. After the connection between the $H\infty$ problem and dynamic game theory was realised and more fully exploited, it became clear by a sort of ironic feedback loop, that some of the key results required to solve the state-space version of the $H\infty$ robust optimal control problem were already available in the game-theory literature. Most notably, the optimal state-feedback result of equation 1.136 had already been fully developed in classical LQ game theory.

The main work in the area is that of Ziegler (1999) who develops a game-theoretic approach to pricing options by framing the option pricing problem as a two person game between the option buyer and the option seller. The essence of Ziegler's approach is to separate the problem of the strategic interactions from the problem of valuing payoffs. Ziegler deals with the former using the problem of backward induction so that each sub-game can be replaced with its equilibrium payoff - an idea generally known as subgame perfection. Limiting attention for the moment to finite-horizon option valuation problems, it is then possible to say that in a finite game of perfect information, backward induction and subgame perfection are equivalent (see Fudenberg and Tirole, 1991).

For valuing the option payoff, consider a simple contingent claim whose underlying is some asset S, whose behaviour follows a standard geometric Brownian motion

$$dS = \mu S dt + \sigma S dZ \tag{1.178}$$

where μ is the drift and σ is the instantaneous standard deviation of the process and dZ is the increment of a standard Wiener process[18]. Then by the usual Black-Scholes style analysis, let S be the current value

[18] In two papers (1985 and 1987) Huang derives sufficient conditions under which equilibrium prices can be shown to follow a standard Wiener process.

of the underlying at time t, r the risk free rate of return, a the payoff to the holders of S per unit of time (i.e. a dividend) and let b be the payoff to holders of the contingent claim per unit of time. Then using subscripts to denote partial derivatives, Merton (1977) tells us that if $F(S,t)$ is the value of the contingent claim, it must satisfy the following differential equation

$$\frac{1}{2}\sigma^2 S^2 F_{SS} + (rS - a) F_S + F_t - rF + b = 0 \tag{1.179}$$

It is easy to see that Merton's equation contains parameters and is subject to certain definable boundary conditions. Ziegler's game-theoretic analysis of options is a three step procedure. In step 1 the game between the players is defined in terms of the action sets, the choice sequences and the resulting payoffs. Step 2 involves using standard arbitrage theory to value the future underlying payoffs of the players. The third step is to solve the optimal strategies of the players using backward induction or subgame perfection, working backwards from the final period.

The game theoretic approach thus in effect replaces the maximisation of expected utility that is usually encountered in classical game theory, with the maximisation of the value of an option. This provides the arbitrage-free value of the payoff of the option to the player, which acts as a proxy for expected utility. This has the advantage that the time value of money and the market price of risk are automatically taken into account in the analysis. However, as Ziegler points out, the greatest strength of the game-theoretic approach is in its ability to separate the valuation problem from the strategic interaction between the two theoretical players, which clearly makes the approach most useful in the valuation of non-European exercise options such as the problem of valuing the perpetual put option. The usefulness of this feature is that it makes it possible to solve complex decision making problems under uncertainty by applying no more than classical optimisation techniques to the value of an option. This capability is useful in turn because it can frequently be reduced to finding the first order conditions for either a maximum or minimum.

To understand how the method actually works, it is useful to consider the following example. Consider a two-person game between players I and II with the following structure. First, player I selects strategy A, followed by player II selecting strategy B. Taken together, the two strategies and the future value of S combine to determine the payoffs to I and II. If $G(A,B,S)$ and $H(A,B,S)$ are the current payoffs to players I and II respectively, then these values can be derived by solving a differential equation of a form similar to equation 1.179, subject to appropriate boundary conditions. The strategy for each player involves the selection of one of the parameters of equation 1.179, so as to maximise the value of their respective payoffs. In the final stage of the game, player II selects a strategy B so as to maximise the

value of his expected payoff $H(A,B,S)$ such that

$$\frac{\partial H(A,B,S)}{\partial B} = 0 \tag{1.180}$$

given that B is not a boundary condition. Clearly, this simple first order condition can be solved to produce an optimal strategy $\bar{B} = \bar{B}(A,S)$, which is highly likely to depend on the choice of strategy made by player I (i.e. A). The key point is that at the time player I makes his decision, he must attempt to anticipate player II's subsequent choice, which involves setting

$$\frac{\partial G(A,\bar{B},S)}{\partial A} = \frac{\partial G(A,\bar{B},S)}{\partial A} + \frac{\partial G(A,\bar{B},S)}{\partial \bar{B}}\frac{d\bar{B}}{dA} = 0 \tag{1.181}$$

which yields the optimal strategy $\bar{A} = \bar{A}(S)$, with the second term $\frac{\partial G(A,\bar{B},S)}{\partial \bar{B}}\frac{d\bar{B}}{dA}$ reflecting the indirect effect of player I's strategy choice on his expected payoff that results from the influence of his choice on player II's optimal strategy B. As Ziegler points out, it is this important incorporation of feedback that captures the essence of backward induction, as it explicitly incorporates the fact that player I must anticipate player II's action when making his choice.

There are three principal limitations to applying Ziegler's game-theoretic approach. First and most obviously, the resulting mathematical complexity is highly likely to preclude closed form solutions - a situation that would weigh heavily against its use in practical situations where time and tractability are frequently preferred to elegance and completeness. The second problem is that the game-theoretic approach will not work if the optimal strategies are stochastic, that is they depend on the value taken by S. The third problem arises in the case of convex payoff functions. When the payoff function is non-concave it may be difficult to state strategies unambiguously.

1.4.16 Modelling software and implementation

Though clearly not part of the literature review, this section describes the essential features of the software design approach to programming the analytical models which were developed to test the ideas contained in this thesis. As the body of software expanded it was decided to give the main components names to aid identification. The overall product name given to the components written for the option pricing, risk management and monetary policy software is Robusta.

It was decided not to use the popular Matlab toolboxes as the computational engine for the calculations required in this thesis. This is because although Matlab is an excellent tool it has a number of widely known and well documented limitations

1. *The use of the dense complex matrix as the main data structure for linear algebra computations:*

The need to utilise complex computations to solve computational systems design problems leads to very inefficient implementations of functions to solve problems such as Algebraic Riccati and Lyapunov equations which are central to many areas in robust optimal control and widely used in this thesis.

2. *The trade-off required to balance the Matlab matrix handling power with the ability to exploit intrinsic structural aspects of certain types of problem:* Attempting to exploit the structural features of computational problems frequently results in an offsetting increase in execution time due to the interpreted nature of the Matlab package. Using a language such as C or C++ means that the functions can be compiled, use more efficient data structures, make more efficient use of memory and execute faster.

3. *Lack of numerical robustness in Matlab:* when faced with problems which exhibit a significant degree of scaling in their structure or parameters, Matlab frequently produces spectacular failures such as general protection faults (resulting in the hapless user inadvertently addressing video RAM !). However, more seriously, the more dangerous situation arose on a number of occasions that Matlab produced seemingly correct results which when cross checked with appropriate boundary conditions or other packages such as Mathematica proved to be completely incorrect.

Given the generally non-portable nature of Matlab based applications, it was decided to utilise an alternative source of robust optimal control functionality. Accordingly, the Slicot free-source subroutine library was settled upon as providing the required functionality to an acceptably high standard of precision and rigorous methodology. Six separate C/C++ dynamic link libraries (DLLs for short, see Appendix 1 for full details) containing the functions required to carry out the computations required to test the robust optimal control approach developed in this thesis (see below for full details)[19]. These DLLs were then linked to a front-end graphical user interface written in Visual Basic and a back end Access database for storing source data and results data. The entire package together was christened Robusta.

One initial obstacle that had to be overcome was the fact the Slicot library is written in Fortran, whereas the computational libraries for Robusta are written in C++. There were therefore two options:

- Write mixed language C++/Fortran code to call the required Slicot library functions in their native Fortran form.

[19]For the more technically minded, the underlying algorithms were written in straight ANSI standard C for speed and then wrapped in a series of C++ classes for the purposes of providing a dual interface to both Excel and Visual Basic via a single compilation. The result is a single DLL that can used in either Excel or Visual Basic without the need for either XLLs or complex API calls in Visual Basic.

- Re-code the required Fortran functions from the Slicot library into C++.

Given the inherent difficulties associated with mixed language programming in Windows environments and the desire to use a single language for the computational libraries in Robusta, it was therefore decided to port the required parts of Slicot into C++. Unsurprisingly, this proved to be a non-trivial task, requiring several months of patient coding and testing. In order to achieve the recoding efficiently, the initial idea was to implement a two stage process. First, to run the required parts of the Slicot library through the shareware programme known as f2c which automatically takes Fortran code and performs the appropriate translations into straight C. However, initial efforts proved patchy due to a number of technical problems with the f2c parsing algorithms (e.g. the literality of translating Fortran GOTO statements results in poorly structured C code which then has to be restructured . It was therefore decided to re-write the C-code for the f2c parser, simultaneously improving on the command line interface of the original product by adding a more user-friendly front-end graphical user interface in the process. Once again, this proved to be a non-trivial task !

Once complete, the resulting C code then had to be wrapped into C++ - which proved to be another time consuming task ! However, once complete the resulting library proved to be both as fast as its Fortran predecessor in execution and almost as small in memory footprint. Crucially, results are just as accurate and execution does not cause general protection faults in Windows ! Appendix 1 provides a complete listing of the Slicot library functions that were re-coded into C++.

1.5 Scope for research - thesis organisation

The ideas presented and the research reviewed in this first chapter are a clear indication of the significant potential of applying robust optimal control to problems where robustness and stability of the optimal choice is required even in the presence of uncertainty. Apart from the work by Sargent, Ziegler and an unpublished paper by Bernhard, relatively little research has so far been carried out on the application of robust optimal control and the $H\infty$ approach in particular, to decision theoretic problems where meaningful insights could be gained from investigating the separation of intertemporal substitution and relative risk aversion. The question therefore arises of precisely which areas of research this thesis is attempting to contribute to in order to add to the existing body of knowledge in the area. As outlined at the beginning of this chapter, there are three main areas of contribution, namely, derivative pricing, dynamic hedging and portfolio risk management.

Chapter 2 moves beyond the work by Ziegler, to more recent work based on game theory by Bernhard (2000) that applies robust control techniques to option pricing. Bernhard's work is of particular interest

as it offers the possibility of being able to price options without recourse to any particular underlying probability distribution, thereby directly addressing a key weakness of the traditional Black-Scholes (1973) approach. Bernhard (2000) achieves this by using trajectory sets and bounded variation to derive a robust control theoretic framework for the option pricing problem of finding the price of a European exercise option paying:

$$\max(Spot - Strike, 0) \tag{1.182}$$

in the case of a call, or

$$\max(Strike - Spot, 0) \tag{1.183}$$

in the case of a put. This is an intriguing use of robust control theory and is potentially extremely valuable as it is distributionally independent. However, its greatest value as an approach may lie in the area of valuing options on extreme values or events, often referred to as catastrophe options. These are options on discrete or highly non-predictable events - a situation where the traditional probability driven Black Scholes approach is known to be extremely unreliable because of its reliance on a parameterized distribution and a known, fixed volatility (this is so, even in the more recent flavours which use approaches such as jump diffusion and stochastic volatility in an attempt to capture the inherently uncertain nature of the underlying process). At a practical level the Chicago Board of Trade offers cash-based index options on a number of instruments. The Property Claims Service (PCS) indices reflect estimated insured industry losses for catastrophes that occur over a specific loss period. The current market valuation practice is to use some form of GARCH model to estimate volatility and then use this within a Black-Scholes type of model. Such an approach may lead to a fundamental mispricing of options since the two most common distributional assumptions of either nomality or log-normality of the underlying price process are both known to be highly inaccurate in the case of extreme outcomes. The contribution of this thesis is in applying robust optimal control theory to provide an approach that will permit the calculation of robust and stable options prices in the case of options on extreme underlyings.

Two aspects of risk management appear to offer considerable profitable scope for the application of $H\infty$ optimal control techniques. The first is the more micro-economic aspect of hedging of individual options, whilst the second aspect is more macro-economic and is concerned with the control of overall risk across a financial institution. Chapter 3 examines the micro risk issues first.

Traditional option theory is constructed around the notion of risk neutrality, the idea being that for an option to be correctly priced arbitrage must be impossible and that full hedging is only optimal when

the current forward price is equal to the expected future spot price. This view is very much predicated on the notion of a quantifiable stochastic disturbances. However, Lien (2000) introduces the idea of Knightian uncertainty into the futures hedging problem, examining the situation where an agent has imprecise information about the underlying probability density function of futures prices. This imprecision generates ambiguity, to which Lien argues, agents are averse, resulting in hedging inertia with regard to the decision over full or partial hedging. He finds that when carrying out forward hedging using futures, a full one-to-one hedge ratio is more likely to occur under Knightian uncertainty in a region of prices around the current forward price. The size of the region is positively related to the degree of ambiguity over the probability density function. This kind of optimal hedging problem is exactly where the application of $H\infty$ optimal control techniques could be extremely useful. For example, little or no work has been carried out on the combined stability and robustness properties of dynamic hedging strategies.

Chapter 4 examines risk management at the macro-risk level. The VAR or "value at risk" approach is one of the most widely used methods of measuring and controlling risk among financial institutions. VAR summarises the expected maximum loss over a target horizon within a given confidence interval. Institutions using VAR typically look at a 95% confidence interval over periods ranging from one to ninety days. One obvious problem with VAR is the reliance on the first two moments of the selected distribution - namely the mean and standard deviation - in order to calculate the confidence interval for a situation where it is acknowledged the tails of the distribution are unlikely to be an accurate guide to the likely frequency and magnitude of large market movements (which is one of the key reasons for using VAR). The frequent and implicit use (due mostly to familiarity, though not necessity) of a quadratic loss function through the desire to pursue a minimax loss strategy is also likely to produce misleading conclusions if the underlying utility function is non-quadratic (especially in the presence of large discontinuous changes), which once again leads back to the variational utility concept where the separability of intertemporal substitution and risk aversion is critical to an understanding of how the risk decision is made. An organization cannot be indifferent to risks at different time horizons - bankruptcy tomorrow renders irrelevant any concept of discounting or future decisions. Both institutions using VAR and regulators assessing the risk profile of a firm usually do so with the intention of using VAR results to provide feedback into improving the future risk management of an institution. Applying robust optimal control techniques provides control rules that are both more robust and stable in operation than a wide range of the more popular hedging rules.

Chapter 5 draws together the main findings and conclusions in the thesis and offers ideas on directions for future research in the area of decision making in the presence of uncertainty.

Chapter 2

Robust optimal control and the pricing of catastrophe derivatives

It is better to understand a little than to misunderstand at lot.

Anatole France, nom de plume of Jaques Thibault, French novelist and critic, 1844 - 1924

2.1 Introduction and motivation for research

According to the 2003 report by the International Institute for Applied Systems Analysis, more than 700 major "natural" catastrophes occur every year. Since the 1950's alone, economic losses from disasters of either natural or man-made causes have increased 14-fold. Although the developed world generally suffers the greatest absolute economic losses, seen as a percentage of GNP, the losses in the developed world are 150 times greater in terms of human victims and 20 times greater in terms of economic losses. The major factors behind these escalating costs are economic and population growth with accompanying land-use changes and capital movements of capital to vulnerable regions. In other words, the main cause is global change in its broadest sense. At the same time as such profound and sustained global change is occurring, it is expected that long-term worsening of weather extremes due to climate change will also accelerate if human action is not taken.

Specific events such as the stock market crash of 1987, hurricane Andrew in 1992, the Northridge earthquake in 1994, or even the bombing of the World Trade Centre in 2001, have all served to underline the increasingly stretched capacity of global insurance markets. This has in turn sparked an increasingly active debate among actuaries, economists and politicians over whether financial markets with their vastly greater risk bearing capacity could, or indeed should, be used to hedge risk that has previously

been covered using other channels. A symptom of the debate and of the increasing linkage and possible convergence between the insurance and financial markets has been the steadily increasing stream of financial instruments that contain some form insurance product and vice versa. The lines between the industries have become increasingly blurred and this has stimulated the tailoring of products to the risk preferences, capacities and demands of both buyers and sellers of risk based products. Securitisation of risks in both areas has generated a concommitant development in the pricing of products containing both insurance and financial elements, the result of which has been a search for an appropriate class of processes to solve the problem of price determination for extreme risks. Arguably one of the most important aspects in this process is the way in which risk and uncertainty are handled in each discipline. Given that there is broad consensus between the two areas over the treatment of risk as a utility maximisation problem, the most interesting issue is therefore how to deal with uncertainty in the pricing of derivative instruments and in particular, derivative instruments whose underlying is some form of extreme or catastrophic event.

Hedging of catastrophic risks originated in the insurance industry, which makes it logical to begin the analysis of the problem from an insurance perspective. Catastrophe insurance is essentially a put option that locks in the value of some underlying asset. The purchaser pays a premium to the insurance company to buy a policy (which specifies the extent and limitations of cover). If and when some pre-determined event occurs that reduces the value of the protected asset, then the insurance company replaces the asset according to the details of the policy. According to the US Insurance Services Office, for example, a catastrophe is an event which causes at least $25 million in damage and must affect multiple parties. Unlike standard insurance policies for risks such as motor cars or whole life cover, catastrophes usually strike an entire area or group, frequently inflict massive damage and occur uncertainly, making prediction impossible. These factors have been the main reason why insurance companies were unwilling to assume catastrophe risk until relatively recently. Another factor has been the presence of high adverse selection, meaning that only those planning to use such policies buy them in the first place.

From an actuarial perspective, an insurance company writing catastrophe insurance can be analysed as if it were a portfolio of put options, each on an individual catastrophe, where the overall portfolio return takes on a Poisson distribution depending on the frequency and magnitude of the sum of the catastrophic events[1] - zero payout or 100% payout - with correlations very close to unity. Such risk can and will only be assumed if adequate reinsurance cover is available, which gave rise to the development of reinsurance, which allows primary insurers to hedge away all or part of their risk by selling it to another company. Reinsurance also allowed primary insurers to write catastrophe risk policies with premiums low enough

[1] In the sense that the Poisson distribution is used to model the number of events occurring within a given time interval. So, either an event or events do or do not occur, thereby affecting the payout. The arrival of individual catastrophes would be most appropriately described by a Bernoulli distribution due to the binary nature of catastrophic events.

to encourage sales, yet high enough to fund hedging via risk transfer. Reinsurance companies are willing to assume catastrophic risks because they are responsible for only a small part of the exposure (typically that beyond some limit as in the case of excess of loss reinsurance) with the least likelihood of occurrence as it is in the tail(s) of the perceived distribution. To complete the risk transfer process, the reinsurer sells back to the original insurer another put on the same underlying catastrophic risk, but with a lower strike price. The net of this process is to produce a risk profile that to an options modeler would be immediately recognisable as the payoff for a range forward. The important point is that the reinsurer is willing to engage in the process because the catastrophe risk is likely to have a very low correlation with the other risks in their portfolio, thereby allowing diversification whilst simultaneously earning premium income.

The emergence of reinsurance effectively made possible catastrophe insurance, but the process still has some problems. First and most obviously, is that the reinsurance market is not efficient as there are few competing firms who are able to vary their premia at will. Second, due to the relatively low levels of capitalisation of the competing firms, credit risk is a potentially serious problem. Third, due to the diversity of exposures, pricing is inconsistent and tradability of risks is poor due to lack of liquidity. It is not surprising therefore, that insurance companies have turned to the transparency, efficiency and capitalisation of the capital markets for solutions. For example, according to the 2002 Sigma survey, the entire primary and secondary (or reinsurance) industry had access to capital of $240 billion at the end of 2002. A catastrophe of say, $75 billion would therefore bankrupt approximately one third of all insurers. In contrast, the global capital markets have an annual turnover of approximately $26 trillion, with an average daily volatility of approximately 70 basis points, or $133 billion and so would be far more capable of providing a deeper and more liquid means of consistently pricing, hedging and transferring catastrophe risk.

2.1.1 Instruments and markets for catastrophe risk

1998 was a significant year for the catastrophe risk market, for three reasons. First, measurable, reliable indexes of catastrophe losses emerged which could act as reference points for reliable and accurate pricing of catastrophe linked products. Second, new products began trading on recognised financial exchanges, providing the beginning of transparency and liquidity required for an efficient catastrophe risk market. Finally, the over-the-counter market in the underlying catastrophic risk and associated derivative products began with the issuance of the first bonds whose payouts were linked directly to indexes of catastrophic events. The subject of catastrophe based derivatives is returned to in considerable detail in the subsequent section on the pricing of catastrophe derivatives using robust optimal control techniques.

Continuing the theme of robust decision making that underpins this thesis, this chapter seeks to address the relatively sparsely researched area of robustness in pricing certain types of catastrophe derivative instrument. The problem considered is the relatively narrow one of the robustness of the pricing of options on extreme and catastrophic events, such as natural disasters, credit default options or stock market crashes. Hitherto, the broadly predominant approach has been via expected utility theory, under which the price of an option is found to be the discounted value of the expectation of the underlying price process under some chosen risk-neutral measure. Research in this area has been carried out in two broad areas.

On the one hand, has been the work carried out in mathematical finance which has mainly concentrated on finding ways of extending the bench-mark (generally lognormal) Black-Scholes (1973) option pricing paradigm, which relies on the standard diffusion type processes and arbitrage-free arguments to price options under a risk-neutral measure. Within mathematical finance the standard BS approach has been extended to encompass the effects of well known shortcomings such as non-constant or non-smooth volatility, alternative measures (as well as their associated transforms and pricing processes) and distributional irregularities such as jumps and fat tails. This is hardly surprising when it is remembered that a substantial part of current option pricing theory continues to be based around the Black-Scholes (1973) and Merton (1974) (BSM) option pricing approach. The principal reason for this focus is practical, namely, that many practitioners like the computational simplicity and speed, as well as the strong intuitions underlying the BSM approach.

On the other hand, has been the work carried out in the area of actuarial mathematics where issues of pricing insurance contracts which cover eventualities where jumps or discontinuities are frequently encountered, such as ships sinking, or car accidents, life related events, or in extremis, catastrophes such as hurricanes or earthquakes. Actuarial mathematics has tended to focus on the use of risk processes of the compound Poisson type, where markets are by definition and practicality, rarely if ever complete. As a result, risk can rarely if ever, be hedged away and in many cases there will be an infinite number of equivalent martingale measures, so that pricing becomes directly linked to risk preferences. This has been a central reason for the use of techniques such as the Esscher transform, where the linkage between the financial and actuarial mathematics can be made via the theory of semi-martingales.

2.2 Motivation for research

What then is the motivation for research into the robustness of pricing catastrophic derivatives and other extreme events in options theory ? The motivation for the research in this chapter has three main sources. The first is the increasing number, frequency and intensity of catastrophic losses in both the

capital markets and insurance worlds that have occurred in the past 30 years and have brought to the attention of many researchers the importance of studying the impact and pricing of such events. Over the last 20 years in the financial markets, adverse events such as the stock market collapse of 1987 and the collapse of Long Term Capital Management (LTCM), have precipitated massive losses for institutions and private individuals alike. Over a similar period in the insurance markets there have been events such as hurricane Andrew and the Northridge earthquake. Whilst substantial research effort has been applied to these types of problems from the insurance area, little work has been carried out from the mathematical finance field. It is therefore especially interesting that the robustness or otherwise of the prices and values produced by current option modelling approaches have hitherto not been a major source of interest to researchers.

Second, is decision making under uncertainty, which is the underlying theme of this thesis. Problems in option pricing have witnessed some of the most interesting work in decision theory in the period since the publication of the seminal option pricing work by Fisher Black and Myron Scholes in 1973. Though still ubiquitous in its use, the Black-Scholes model suffers from a number of well known and increasingly well researched shortcomings, a small number of which are the central focus of this chapter. For example, from a theoretical perspective, uncertainty enters the decision making problem via the notion of utility and the concept of risk aversion, whilst possible alternatives would be equilibrium theory, or the state/belief dependent utility approach of Veronesi (2001). But either way, a multitude of possible martingale pricing measures can arise. The difficulty to be addressed is, that if the objective is to remain within the risk-neutral, arbitrage free approach, it is not possible to be certain which of the possible pricing measures should be selected.

The third motivation for this chapter is the apparent lack of robustness in much of option pricing theory. The substantial losses suffered by delta hedgers in the wake of the 1987 crash and the almost total failure of delta hedging in the case of LTCM, suggests that the underlying Black-Scholes model is not robust to shifts in one or more of the underlying parameters in the model of the price process, the volatility process or the arbitrage free nature of the underlying pricing theory. For example, work by Berkowitz (2001) tests the robustness of the pricing process to model selection. In particular, Berkowitz examines the impact of arbitrarily frequently recalibrating an implied volatility grid and using it to price standard Black-Scholes options. He finds that as the frequency of the recalibration increases arbitrarily, the prices of the options in turn become asymptotically more accurate. In other words, it would appear to be possible to get the "correct" option prices from the "wrong" model, just by continuously re-calibrating and re-hedging !

The final motivation for this chapter is the possibility of using feedback control modelling techniques

to form a bridge into option pricing theory by using control theory as a common thread between actuarial mathematics and mathematical finance. Robust optimal control is an area of mathematical control theory that has been in wide use in engineering for many years, but its use in finance is relatively sparse, with the substantial body of work by Sargent et al (e.g. 1999 and 2000) still proving the exception rather than the rule. At its simplest, robust optimal control involves finding a controller process which will make the results of the dynamic model of the system under consideration robust to uncertainties in data, parameters, model selection, model specification and types of perturbation. One of its key attractions is the facility to design models which do not require the specification of a particular probability distribution in order to generate a unique and robust solution to the problem of pricing an option in the presence of uncertainty. The particular attraction in the case of catastrophic derivatives is the impact of applying optimal control techniques to develop a class of pricing model that provides greater robustness in the presence of fundamental model and data uncertainty.

The work presented in this chapter therefore represents new ideas and methods of applying existing techniques in both linear and non-linear robust optimal control and state-space techniques to the emerging area of catastrophe based derivatives in order to produce a pricing paradigm more suited to the extreme nature of the underlying catastrophic events . The application and extension of existing work therefore arises in three main ways:

1. Treatment of catastrophe based derivatives as a problem in robust optimal control within the context of decision making.

2. Application of robust optimal control techniques to the pricing of catastrophe based derivatives.

3. Devlopment of valuation approaches to price a range of catastrophe based derivatives - most notably catastrophe options and catastrophe linked bonds.

The chapter therefore proceeds as follows: section 3.2 provides a review and critique of existing literature and approaches to the pricing of options on extreme events. The chapter continues with section 3.3 which provides derives a new approach to pricing catastrophe based derivatives using robust optimal control techniques. Section 3.4 provides a description and analysis of the data and modelling approach and preliminary results. The chapter concludes with section 3.4, which offers preliminary conclusions and suggestions for further work.

2.3 Review of the literature

Characterising extremes by example, Hurricane Hugo in 1989, Hurricane Andrew in 1992 and the Northridge earthquake in 1994 constitute the three most costly catastrophes in US history and highlight the vast exposure to catastrophic losses faced by the US insurance industry. Studies such as ISO 1996, indicate that the US insurance industry could face losses in the region \$75 - \$100 billion from a serious earthquake or hurricane. According to Sigma 2002, such losses would potentially find one third of all US insurers insolvent and leave losses of up to \$50 billion to be covered by remaining insurers and policy holders. The problem has two dimensions. First, the sheer scale of the losses which could be occasioned by the type of catastrophic events quoted above. For example, in 1992 hurricanes Andrew and Iniki struck the US and a record 63 property and casualty insurers became insolvent, whilst in the 10 years to the end of 1998, US property/casualty insurers suffered catastrophe losses of \$98 billion. This compares with only \$49 billion in similar losses in the 39 years to the end of 1988. A 1995 ISO study also showed that the worst hurricanes now have the potential to cause insured property losses in excess of \$75 billion, whilst earthquakes of around 8.5 on the Richter scale in an area such as New Madrid in the USA could cause insured property losses in excess of \$115 billion. The second dimension is that the population in those parts of the USA now exposed to hurricanes and/or earthquakes rose by 25% in the 20 years to the end of 1990. By way of comparison, the bombing of the World Trade Centre on 11 September 2001 resulted in losses of \$49 billion in a single highly contained event.

In order to be able to quantify such losses, actuarial and financial mathematics have generally started at different points. Actuarial driven research has generally focused on the issues surrounding the pricing of reinsurance contracts that will provide cover for primary insurers. Mathematical finance in contrast, when faced with problems such as the 1987 stock market crash, or the failure of Long Term Capital Management has traditionally looked to option pricing techniques in general and Black-Scholes in particular. This has lead the two groups in two broadly different directions. One the one hand, the use of jump diffusion type approaches have proved to be the mainstay of the actuaries, whilst continuous time stochastic volatility models have constituted one of the key areas of attention in the modelling of catastrophic events in mathematical finance.

From a practical perspective, numerical examples such as those above are useful. However, a rigorous mathematical framework is a fundamental prerequisite for a tractable theoretical and practical study of extreme events. So, how should such events be characterised ? In classical statistics, extreme events are usually regarded as outlier events - those that occur with extreme rarity, but can still be subjected to formal analysis within a precise mathematical framework. At a theoretical level, such events are generally characterised as having low expectation of occurrence and therefore of being in the extreme tails of a

distribution. They are difficult to predict, but will have a substantial and sustained impact on a given part of the financial sector of the economy, e.g. the insurance market or the capital/financial markets.

At the level of a formal model of such types of events, there are several important questions to consider when analysing candidate models for dealing with catastrophic events:

1. What features characterise extreme or catastrophic events and what are the features required of a model to capture such features realistically, usefully and tractably ?

2. How to deal with incomplete markets in hedging arguments.

3. What is the most appropriate model - what are the criteria for deciding ?

4. Is the model robust with respect to distributions, parameter and model uncertainty ?

5. Is data available for the chosen model parameters ?

6. Computational considerations - models must work reasonably quickly to be useful to practitioners.

These are the key issues to keep in mind throughout the review of the models presented in this section. However, the first point is that extreme or rare events are currently defined with respect to some form of discontinuity in the process for the underlying. At this point, it is important to differentiate discontinuity from volatility and this is where the broad approaches of the insurance and capital markets diverge, which makes an ideal point to commence a review and critique of the literature.

Characterising the differences in methodology in the pricing of catastrophe based products between the insurance and financial/capital markets can be approached in a number of ways, all of which yield varying insights. The approach adopted in this chapter can therefore be regarded as somewhat subjective in its perspective, but has the virtue that it provides a simple starting point to tackle an area that is among the most complex in the area of pricing derivative instruments. Given the possibility of viewing a catastrophe derivative as a put option (or even as a range forward) and given the ubiquitous nature of the Black-Scholes (BS) approach, it makes sense to begin by using the Black-Scholes framework as a starting point for a review and critique.

At its simplest, the BS option pricing formula is summarised by the familiar linear partial differential equation linking the price of a derivative to the price of an underlying instrument:

$$\frac{\partial f}{\partial t} + rS\frac{\partial f}{\partial S} + \frac{\sigma^2 S^2}{2}\frac{\partial^2 f}{\partial S^2} = rf \tag{2.1}$$

where S is the price of the underlying, f is the price of a derivative instrument, contingent upon the price of S; t is time, r is the risk free rate used for discounting and σ^2 is the variance of S. In its original

form, both σ^2 and r are assumed to be constant. The BS equation describes a single dimensional diffusion equation with well known boundary conditions, which when varied provide the solution space for a variety of alternative derivative instruments.

The BS model is underpinned by a range of assumptions, but one of the most fundamental is that S follows a specific type of Wiener process generally referred to as a geometric Brownian motion of the form

$$dS = \mu S dt + \sigma S dW, \quad \text{where } dW = \varepsilon(\sqrt{dt}) \text{ and } \varepsilon \sim N(0,1) \tag{2.2}$$

with constant mean (μ) and variance σ^2. N is assumed to be normally distributed with zero expected value and unit variance. Here is the heart of the critique of the BS approach, namely, that the assumption of constant variance does not reflect reality and that S moves in a smooth and continuous fashion. Broadly speaking, pure option theorists in financial mathematics have concentrated on developing models that attempt to expand the way in which volatility is treated within an option model, whilst actuaries have concentrated on developing models that modify or remove the assumption of the continuity of S, usually through the introduction of some form of jump feature in the behaviour of the underlying. But most fundamentally, neither group has shown much if any interest in the robustness or otherwise of the pricing and hedging results of their models. The following sections therefore analyse the problems with the relevant BS extensions, mainly from the perspective of catastrophic derivatives, concentrating substantially on those key areas which particularly illustrate the need for an alternative approach to the pricing problem. The analysis begins with the volatility type approaches favoured by mathematical finance and is then followed by the jump based approaches pursued with such single-minded vigour by the actuarial community.

Constant volatility models

In the 30 years since its introduction, the basic BS approach has been expanded in a number of ways in an attempt to adjust for the drawbacks that accompany an assumption of constant volatility. The adjustments can be grouped into three main types. First, models that treat volatility as purely deterministic. Second, are models that make volatility an independent stochastic variable. Third, are those models which do not attempt to correct volatility but instead opt to predict it over the remaining life of the option. Prediction is carried out using some form of autoregressive conditional heteroskedastic (ARCH) process, which restricts volatility to be constant over the remaining life of an option, while allowing conditional volatility (conditional on historic data) to vary as a function of previous prediction errors.

As far as deterministic volatility models are concerned, there have been three main strands of research. The first was the constant elasticity of variance (CEV) model, which treats the underlying and its volatility

are negatively correlated. This assumption is based on the idea that in the case of a leveraged firm, a fall in the stock price raises gearing and therefore riskiness, whilst for a non-leveraged firm, falling operating income lowers stock price performance and therefore also raises riskiness. Cox and Ross (1976) modelled this apparent inverse relationship between the underlying and its volatility using the assumption that changes in the underlying can be predicted by following relationship

$$dS = \mu S dt + \sigma S^{\alpha/2} dW, \quad \text{where } dW = \varepsilon(\sqrt{dt}) \text{ and } \varepsilon \sim N(0,1) \tag{2.3}$$

where μ, σ and α $(0 < \alpha < 2)$ are constants, with the latter measuring elasticity. At the point that $\alpha = 2$, the model becomes the regular BS model, whilst $\alpha = 0$ produces the absolute diffusion model and $\alpha = 1$ produces the so called square root model. Replacing the usual geometric Brownian motion under this approach results in a very similar result to the BS model, but with the unfortunate addition of an infinite sum in the solution to the differential equation. Approximations to avoid this problem exist for both the absolute diffusion (Cox and Ross (1976) and square root models (Cox (1975)), but neither model is widely used. This is probably due to two factors. First, the evidence that volatility is is fact negatively correlated with the underlying. Second, the lack of stability over time of the optimal choice of elasticity coefficient (Emmanual and MacBeth (1982)) and the apparent volatility of elasticity measures for particular stocks over some periods (Ang and Peterson (1984)). Interestingly, attempts to model time series of elasticity coefficients as either AR(1) or AR(2) processes indicate that they are uncorrelated with past values, which raises the question of whether assuming fixed elasticity provides any real improvement over the standard BS approach.

The second approach to varying the volatility assumption is the compound option model (COM), which like the CEV model, is based on the assumption that volatility is negatively correlated with the underlying. Based on the idea that the firm is leveraged with total value, V, but with constant debt, D, (comprised totally of zero coupon bonds), then the firm's value at expiration of its debt will be $\max(0, V - D)$ which implies that the value of the option on the underlying is an option on an option, or a compound option. Geske (1977) produced the first model for compound options, but his model is based on the assumption that the volatility of the value of the firm, σ_v, is constant and expresses the underlying as a solution to the BS equation with V replacing S and D replacing strike in the boundary condition for a call. His solution makes volatility a function of both V and t, such that the COM has a solution which is linear combinations of bivariate cumulative normals. The COM represents only a slight advance from the standard BS model given its assumption of constancy of σ_v, which is not much better that the constant volatility assumption of the standard BS model.

The third approach is the displaced diffusion model (DDM) which again prices equity options based on

assumptions about the structure of the underlying company (Rubenstein (1983)). The critical difference is that the DDM model produces positive correlation between the underlying and its volatility. Rubenstein assumes that a company is the sum of risky assets (exhibiting constant volatility) and riskless assets (used to pay off the debts of the company). Using Rubenstein's own notation, V is the value of the company, with α being the proportion accounted for by risky assets, where αV behaves according to a geometric Brownian motion, whilst the riskless assets compound at the risk-free rate, r. After some time period, t, the value of the company will have grown to

$$\left[\alpha e^{y} + (1-\alpha)(1+r)^{t}\right] V, \quad \text{where } y \sim N(0, \sigma_v \sqrt{t}) \tag{2.4}$$

If β is the company's gearing ratio, then the value of the company, S, at time t in the future is given by

$$S_t = \left[\alpha e^{y} + (1-\alpha)(1+r)^{t}\right](1+\beta) S - S\beta(1+r)^{t}, \quad \text{where } a = \alpha + \alpha\beta \tag{2.5}$$

If $a \leqslant 1$ then $\beta \leqslant (1-\alpha)/\alpha$ making debt riskless, so that the above equation becomes

$$S_t = \alpha(1+\beta) S e^{y} + (1-\alpha-\alpha\beta) S (1+r)^{t} \tag{2.6}$$

where the first term is the risky component of the company's value and the second is the risk free part. If $a > 1$, the COM model results, leading once again to negative correlation between the underlying and its volatility. The other main drawback of the DDM is that it is difficult to use in practice as it requires more variables than the COM and their estimation is every bit as fraught with practical difficulties. The final problem is that very few companies are actually capable of being decomposed into risky and riskless components in such a idealistic fashion.

Stochastic volatility models

The past 10 years have witnessed the production of a wide range and large number of models that treat volatility as a stochastic variable in its own right. Introducing a second source of stochastic variation produces two new problems - correlation and risk-neutrality. Numerous models assume correlation between the two stochastic variables to be zero - which considerably simplifies the solution of the resulting partial differential equation. There is currently no completely closed form solution in the case of non-zero correlation. However, there are a number of numerical approaches to deal with non-zero correlation.

The issue of risk-neutrality raises a trickier problem. The standard BS approach generates a solution hinged on the assumption that is it possible to form an instantaneously risk-free portfolio comprised of the

derivative and the underlying. This assumption is no longer valid in the presence of a second stochastic input because volatility is not in and of itself a tradable asset. The extra source of randomness can be eliminated by algebraic manipulation, but the problem of non-tradability of volatility remains a difficulty. One way around this problem has been to pose a model in terms of a volatility risk premium, λ^*. Setting $\lambda^* = 0$ results in a model where stochastic volatility is not priced such that the risk is non-systematic. Clearly, if there is no need to compensate for volatility then a risk-free portfolio can be constructed in the familiar way.

The stochastic volatility model of Hull and White (1987) produces a simple analytic expression for the price of an option under stochastic volatility (though their paper is actually about the stochastic process of the variance) using a property of diffusion state processes first introduced by Garman (1976). Garman's paper showed that the following must hold for any underlying, f, that depends on the state variable, θ_j

$$\frac{\partial f}{\partial t} + \frac{1}{2}\sum_{ij}\rho_{ij}\sigma_i\sigma_j\frac{\partial^2 f}{\partial\theta_i\partial\theta_j} - rf = \sum_i \theta_i\frac{\partial f}{\partial\theta_i}\left[\beta_i\left(\mu^* - r\right) - \mu_j\right] \tag{2.7}$$

where r is the riskfree rate, σ_j is the instantaneous standard deviation of θ_j, ρ_{ij} is the instantaneous correlation between θ_i and θ_j, μ_j is the drift rate of θ_j, β_j is the vector of regression betas for the regression of $\frac{\delta\theta_j}{\theta_j}$ on the market portfolio and μ^* is the vector of instantaneous expected returns on the market portfolio. When combined with a two equation form for both the underlying and the volatility

$$dS = \phi S dt + \sigma(t) S dW \tag{2.8}$$

$$dV = \mu V dt + \xi V dZ \tag{2.9}$$

where $V(t) = [\sigma(t)]^2$, gives

$$\frac{\partial f}{\partial t} + \frac{1}{2}\left[\sigma^2 S^2\frac{\partial^2 f}{\partial S^2} + 2\rho_{sv}\sigma^3\xi S\frac{\partial^2 f}{\partial S\partial V} + \xi^2 V^2\frac{\partial^2 f}{\partial V^2}\right] - rf \tag{2.10}$$

$$= -rS\frac{\partial f}{\partial S} + \left[\beta_v\left(\mu^* - r\right) - \mu\right]\sigma^2\frac{\partial f}{\partial V} \tag{2.11}$$

Making the assumption that the term $\beta_v(\mu^* - r) = 0$ is equivalent to $\lambda^* = 0$, which is a frequently used practical choice, but is less general than replacing μ by $\mu - \beta_v(\mu^* - r)$. The final expression for the price of a European exercise option on an underlying that pays no cash flows is then the discounted value

$$f\left(S_t, \sigma_t^2, t\right) = e^{(-r(T-t))}\int f\left(S_T, \sigma_T^2, T\right) p\left(S_T | S_t, \sigma_t^2\right) dS_T \tag{2.12}$$

where t is current time and T is expiration time; and $p\left(S_T | S_t, \sigma_t^2\right)$ is the distribution of S_T conditioned

upon the price and volatility of the underlying at time t. Assuming $\rho_{sv} = 0$ reduces the above expression to

$$f\left(S_t, \sigma_t^2\right) = \int C(v)\, g(v)\, dv \tag{2.13}$$

where

$$v = \frac{1}{T-t} \int_t^T \sigma_v^2 dv \tag{2.14}$$

measures the average variance over the life of the derivative instrument, $C(v)$ is the value of the BS call as a function of volatility and $g(v)$ is the risk-neutral density. Evaluating the integral $f\left(S_t, \sigma_t^2\right)$ using Hull and White's approach (which they argue converges rapidly for small values of $\xi^2(T-t)$), requires the power series expansion about $v = 0$ of $C(v)$. Hull and White's approach is simple mainly because of its unrealistic assumption of zero correlation between the underlying and its volatility.

Of all of the stochastic volatility models, Heston's 1993 paper has arguably received the most attention from practitioners in the finance market. This is probably because it does not use Garman's identity to derive the partial differential equation for the underlying. Heston focuses on the situation where λ^* is proportional to $v(t)$ and assumes the underlying and its volatility have the following stochastic processes

$$dS = \mu S dt + (v(t))^{\frac{1}{2}} S dZ_1 \tag{2.15}$$

$$d(v(t))^{\frac{1}{2}} = -\beta (v(t))^{\frac{1}{2}} dt + \delta dZ_2(t) \tag{2.16}$$

which when subjected to Ito's lemma enables Heston's so called "square root model" to provide the following expression for the stochastic process of the variance

$$dv(t) = \kappa [\theta - v(t)]\, dt + \sigma (v(t))^{\frac{1}{2}} dZ_2(t) \tag{2.17}$$

This expression shows clearly that Heston is using an identical stochastic process to that used by Hull and White to describe the evolution of the underlying. Unfortunately, although Heston's solution is semi-closed form, it contains characteristic functions that give rise to integrals that can only be computed by numerical methods. Heston also considers the case where the interest rate is allowed to vary stochastically by following a standard Cox-Ingersoll-Ross type of process - but this extra step suffers from the drawback that the variance of the underlying and the variance of the interest rate are determined by the same stochastic process, $v(t)$.

Wiggins (1987) offered an early attempt at building investor utility into a stochastic volatility model. He restricts his model to pricing options on indexes, thereby neatly justifying the assumption of $\lambda^* = 0$. Wiggins defines a risk-free portfolio that yields a risk free return of μ_p and a volatility of σ_p, which produces

a simple expression for expected loss as a per-unit function of risk: $\phi(\bullet) = \frac{(\mu_p - r)}{\sigma_p}$. This then allows him to show that the assumption that investors have log-utility functions leads to $\phi(\bullet) \equiv 0$, thereby giving a partial differential equation capable of being solved for any given ρ_{sv}using simple numerical techniques. Stein and Stein (1991) adopt a similar line to Heston, such that their result is couched in terms of the use of characteristic functions that once again produce a numerical integration, but at least avoids the Hull and White power series problem.

What is clear from the stochastic volatility models is that they all have two major drawbacks. First, market incompleteness arises due to the fact that the number of assets is insufficient to span the entire state space of contingencies. Or put another way, there is only one underlying asset, but two sources of Brownian motion. This makes impossible to perfectly replicate the aribtrage-free option price, rendering perfect hedging impossible. Hofmann (1992) was the first to suggest a way around the hedging problem, by suggesting that whilst there are no perfect hedges available in the stochastic volatility models, optimal portfolio type strategies could be used instead. For example, the target could be to minimise a hedging loss function (as suggested by Duffie and Richaradson) or search for dominating strategies (as suggested by Bensaid).

Second, there consequently exists an infinite number of option prices consistent with arbitrage-free conditions. One way of circumventing this problem, employed by both Heston and Wiggins is to employ the concept of a representative agent that can typically trade in both the underlying and the option. This is a neat standard device, that has the undesirable side effect of making the option price preference-dependent.

A further point on stochastic volatility models is that although suggesting that $\lambda^* = 0$ may be too simplistic in most cases, it may not be entirely unreasonable in the case of options on a market index. Heston's claim that the value of volatility risk should be proportional to the volatility itself has intuitive appeal, but to date is unsupported by any body of empirical evidence.

GARCH models

In contrast to the stochastic volatility models, one line of research into adapting the standard BS model that has attracted considerable attention is the forecasting of future volatility with the intention of using it to price derivatives more effectively. The simplest approach is to assume that markets are efficient so that implied volatility can then be assumed to form an accurate predictor of expected future volatility. This is the main idea behind the implied binomial/trinomial tree approach of both Rubenstein (1994) and Brown and Toft (1997). The main aim of both approaches to correct the standard BS model for non-constant interest rates and non-constant volatility. The second approach is to accept that markets

are not completely efficient and use historical volatility data to forecast future volatility. This was the approach that lead to the backward looking generalised autoregressive conditional heteroskedastic (ARCH) model of Engle (1982). ARCH is based on the basic autoregressive moving average (ARMA) model for data smoothing, which can only be used to model stationary processes due to their inability to describe the dynamics of the underlying price change process. Integrating a description of the dynamics of the underlying process, lead to the development of the autoregressive integrated moving average models (ARIMA).

The ARCH approach involves the econometric fitting of variable-volatility models to the underlying data. ARCH models are based on the assumption that markets are not efficient, such that expected future volatility is dependent on implied volatility, so ARCH models can in a sense be regarded as a means of performing smoothing estimation of unobserved volatilities which can then be used in a Black-Scholes type model to price derivatives more effectively. A wide variety of versions of the basic ARCH approach have been tried, but the evidence remains mixed over whether implied or historical volatility is the better predictor of future volatility over time, which suggests that there is a distinct possibility that the optimal volatility forecasting technique is time-dependent.

ARCH models extend the basic ARMA approach by making conditional variance a function of past errors over time. ARCH therefore implies constant variance, but *non-constant* conditional variance. The ARCH approach is based on the assumption that the conditional variance is a function of past values of the underlying

$$V\left(x_t|x_{t-1}, x_{t-2}, ..., x_{t-n}\right) = \alpha_0 + \sum_{j=1}^{n} \alpha_j \left(x_{t-j} - \mu\right)^2 \tag{2.18}$$

where $V\left(x_t|\bullet\right)$ is the conditional variance function and μ is the average of the time series. α_0 must be strictly positive to guarantee stationarity, which also restricts the remaining α series to be $\geqslant 0$ and sum to less than unity. This gives an unconditional variance

$$V(x_t) = \frac{\alpha_0}{(1 - S)} \tag{2.19}$$

which in most instances is generally augmented by changing notation using new variables such that $\varepsilon_t = x_t - \mu$ and $q_t = V\left(x_t|\bullet\right)$, resulting in $\varepsilon_t|x_{t-1,}x_{t-2}, ... \sim \mathrm{N}\left(0, q_t\right)$. Unfortunately, higher dimensionality causes problems in the standard ARCH approach as α can become negative resulting in unrealistic negative conditional variances. Tackling this drawback resulted in the development of the generalised ARCH approach (GARCH) which extends standard ARCH in much the same fashion as ARIMA extends standard

ARMA. GARCH adds an extra summation term to the standard ARCH model

$$q_t = \alpha_0 + \sum_{j=1}^{n} \alpha_j \varepsilon_{t-j}^2 + \sum_{j=1}^{m} \beta_j h_{t-j} \tag{2.20}$$

where $\beta_j \geqslant 0$ and the α_j ARCH restrictions still apply. This GARCH(m,n) model provides a process for ε_t that remains stationary with unconditional variance $V(x_t) = \frac{\alpha_0}{(1-S)}$ and the α's and β's sum to less than unity.

Practical use of GARCH tends to favour a (1,1) process such as

$$q_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1} = \frac{\alpha_0}{(1-\beta_1)} + \alpha_1 \sum_{j=1}^{m} \beta_1^{j-1} \varepsilon_{t-j}^2 \tag{2.21}$$

where the process for ε has constant unconditional variance

$$V = \frac{\alpha_0}{(1-\alpha_1-\beta_1)} \tag{2.22}$$

A further variant of GARCH is the GARCH mean or GARCH-M model (e.g. Choi and Wohar (1992)), which includes an extra equation containing either the conditional variance or its square root.

At the abstract level, one of the main problems with ARCH type approaches is that their predicted volatility models the squares of the innovations instead of the actual innovations. A further problem with ARCH type models is that they are unable to incorporate the observed negative relationship between underlying returns and volatility in the equity markets because they restrict the modelling of volatility to changes in the magnitudes of the innovations. Whilst it is true that EGARCH models attempt to remedy this problem by modelling the logarithm of h_t, the basic criticisms remain.

Jump based models

ARCH and stochastic volatility models are essentially designed to deal with issues such as volatility smiles which are frequently observed in the normal operation of the markets for options such as interest rate caps and floors and in equity options where the probability of occurrence can be assumed to be drawn from some stable, smooth, time related and non-extreme distribution.

In contrast, what distinguishes rare or extreme events is the way their size and probability of occurrence varies (or does not change in the case of the standard BS model) with the size of the observation interval. The key assumption is that as the observation interval, h, becomes smaller and eventually approaches 0, the size of the normal event shrinks correspondingly, but the probability of occurrence remains non-zero. In contrast, a rare or extreme event is by definition supposed to occur less frequently, so that as $h \to 0$

their probability of occurrence goes to zero but their size may not be reduced. Recalling from standard Black-Scholes analysis, the variance of the shock component $\sigma^2 \Delta W_t$ of asset prices is given by:

$$E\left[\sigma_t \Delta W_t\right]^2 = \sigma^2 h \tag{2.23}$$

during a small interval, with the size of the unpredictable changes being given by $\sigma_t \sqrt{h}$. The variance or standard deviation is the product of event size and probability of occurrence. A variance proportional to h can be derived either by probabilities that depend on h while their size is independent, or by probabilities that are independent of h while the size is dependent. The first case characterises rare or extreme events while the second corresponds to normal events. In the former case, if unusual jumps in an underlying occur, the usual continuous time Wiener process needs to be modified to cope with the effects of discontinuities. The most popular method of achieving this augmentation is to use the Poisson process, which is discontinuous. If N_t is the number of extreme shocks occurring in an asset market up to and including time t, then increments in N_t can have only one of two possible values. Namely, zero if no extreme event has occurred, or 1 if an event has occurred[2].

The result of this approach is that current modelling of extreme events involves characterising them as being comprised of two components, namely, one that is predictable given the available state of information and another that is unpredictable. Therefore, in small intervals of time of length, h, it is possible to posit the following

$$S_k - S_{k-1} = a\left(S_{k-1}, k\right) h + \sigma\left(S_{k-1}, k\right) \Delta W_k \tag{2.24}$$

where $k = 1, 2, ...n.$, so that as $h \rightarrow 0$, the continuous time version becomes valid for very small intervals

$$dS_t = a\left(S_t, t\right) dt + \sigma\left(S_t, t\right) dW_t \tag{2.25}$$

such that the current approach simply adopts the standard stochastic differential equation method to account for rare events. The only difference between this and the standard Wiener process being the non-continuity of the sample paths. The current approach has a separate model for the random or unpredictable events, so that it is possible to deal with random jumps in the underlying that occur only rarely, but with time varying probability. Dividing the error term into two parts: those normal events represented by the standard Wiener process, dW_k and those extreme events, dN_k, where at any instant

[2]Note that a variance proportional to h is common to Levy processes. The important distinction in the case under discussion is that between the smooth variation associated with a typical continuous Brownian motion process and the discrete, compound Poisson jump type processes. A discontinuous shift in the former will be extreme, whereas in the latter case it need not be so.

$k-1$, the following holds

$$N_k - N_{k-1} = \{\begin{Bmatrix} 1 \text{ with probability } \lambda h \\ 0 \text{ with probability } 1-\lambda h \end{Bmatrix} \tag{2.26}$$

where λ does not depend on the information available at time $k-1$, so that the jumps of size 1, ΔN_k, occur with a constant rate λ.

N_k has most popularly been modelled using a Poisson counting process, with two important modifications. First, the rate of occurrence of jumps in the Poisson process is highly likely to vary with time, but the vanilla Poisson process has a constant rate of occurrence. Second, the increments of N_t have a non-zero mean, which is in direct contrast to the usual stochastic differential equation approach of zero mean processes. Taking these two modifications together, suppose there is a modified variable

$$J_t = (N_t - \lambda t) \tag{2.27}$$

where the increments, ΔJ_t, have the dual properties of zero mean and unpredictability. Multiplication of J_t by a time dependent constant such as $\sigma_2(S_{k-1}, k)$, will then make the size of the jumps time dependent, making $\sigma_2(S_{k-1}, k)\Delta J_k$ an ideal variable to represent unexpected jumps in the underlying. Therefore, if the market for an underlying instrument is affected by rare events, an appropriate stochastic differential equation would be

$$S_k - S_{k-1} = a(S_{k-1}, k)h + \sigma_1(S_{k-1}, k)\Delta W_k + \sigma_2(S_{k-1}, k)\Delta J_k \quad \text{where } k = 1, 2, ..., n \tag{2.28}$$

which, as $h \longrightarrow 0$, becomes

$$dS_t = a(S_t, t)dt + \sigma_1(S_t, t)dW_t + \sigma_2(S_t, t)dJ_t \tag{2.29}$$

It is important to note that dW_t and dJ_t must be statistically independent with zero correlation at every point in time.

Many insurance companies have used reinsurance to provide a hedge against catastrophic loss. A good example of the problem is the excess of loss contract where the premium at time 0 , given a constant interest rate, is

$$E\sum_{i=1}^{N_t}\{Max(Z_i - b, 0)\} = E(N_t).E[Max(Z_i - b, 0)] = E(N_t).E[(Z_i - b)^+] \tag{2.30}$$

where Z_i is the amount of the claim, N_t is the number of claims up to time t, which are assumed to be independent and identically distributed with a distribution function $H(z)$, $(z > 0)$ and b is a retention

limit. However, for catastrophic events, the assumption of a Poisson process is inadequate as claims for catastrophic events by their very nature occur discontinuously, so that an alternative point process needs to be used to capture the incidence of claims arrival. A popular choice is the doubly stochastic Poisson, or Cox, process popularised by authors such as Bremaud (1981)

$$\Pr\left\{N_{t_2} - N_{t_1} = k | \lambda_s; t_1 \leqslant s \leqslant t_2\right\} = \frac{\exp\left(-\int_{t_1}^{t_2} \lambda_s ds\right)\left(-\int_{t_1}^{t_2} \lambda_s ds\right)^k}{k!} \tag{2.31}$$

A popular choice for measuring the impact of catastrophic events is the shot noise process of Cox and Isham (1980, 1986) which can be used as a parameter of the doubly stochastic Poisson process as a means of measuring the number catastrophic claims

$$\lambda_t = \lambda_0 e^{-\delta t} + \sum_{all\ i,\ s_i \leqslant t} y_i e^{-\delta(t-s_i)}, \quad \text{where } s_i < t < \infty,\ g(y) = \alpha e^{-\alpha y} \text{ and} \tag{2.32}$$

where s_i is the time of occurrence of catastrophe i, λ_0 is the initial value of λ, y_i is the jump size associated with catastrophe i, with a distribution function $G(y)$ where $y > 0$ and $E(y) < \infty$, δ is exponential decay (asymptotically converging towards zero, but never reaching zero).

A popular method for bringing tractability to the Cox and shot noise processes is the piecewise deterministic Markov process theory developed by Davis (1984). This method is used to calculate the mean of the number of claims and the mean of the claim intensity - both of which are central to the pricing of catastrophe related derivatives. Given the non-observability of the claim intensity, state estimation techniques can then be employed to derive the distribution of the claim intensity. Examples of this type of approach to insurance derivatives are Jang (1998, 2000), Dassios (1987) and Kluppelberg and Mikosch (1995).

Aase (1999) similarly uses a compound Poisson process to deal with catastrophe futures and derivatives on catastrophe futures. As these instruments were tradeable - which unfortunately ceased in 1995 - arbitrage arguments were available and a unique price process derivable by specifying the utility function of the representative agents under uncertainty within the given partial equilibrium framework. Aase uses a gamma distribution for loss sizes, coupled with a negative exponential utility function to generate closed form pricing formulae. Cummins and Geman (1994) and Geman and Yor (1997) derive pricing formulae for insurance derivatives by modelling the underlying index as a jump diffusion, also allowing for randomness between the occurrences of the catastrophic events. The approach to ensuring market completeness is quite imaginative in that it is assumed that there exists a series of layers of reinsurance

that serves to ensure market completeness, thereby enabling the derivation of semi-analytical call option pricing formulae using an Asian options approach.

Murmann (2001) foresakes the usual Cox process approach and opts instead for classifying the possible price processes solely on the basis of excluding arbitrage strategies. Under every fixed martingale measure, Murmann's approach enables the derivation of an inverse Fourier transform of the price process (in closed form) for a general class of derivatives without imposing any assumptions concerning the distribution of the jump sizes. Following Aase (1999), Murmann specifies the utility function of a representative agent, thereby allowing him to characterise the equivalent martingale measure and thus establish the unique price process for the catastrophic derivative. Murmann proceeds to generalise his technique to Levy processes (which can be thought of as random walks in continuous time) deriving the same results.

There are two fundamental problems with all of the jump type approaches. The first is that valuation based on arbitrage type arguments can only be justified when all underlying assets are completely and explicitly defined. This is clearly not the case for catastrophe based derivatives where the instrument is based on an index which is intrinsically non-tradeable, notwithstanding its observability or measurability. Even the Geman type approach of assuming layers of reinsurance only provides at best a partial answer to the problem. The second and related problem is that the stochastic jump sizes of the underlying index result in an incomplete market. Both issues together point to the question of whether unique price processes can be argued to exist solely on the basis of precluding the possibility of arbitrage.

Extreme volatility, explosions and the failure of the Martingale approach

The discussion so far has been couched in terms of the normal behaviour of option pricing models, but what happens in the presence of extreme behaviour in volatility ? There are two cases to consider: volatility explosions and volatility extremes. In most economically reasonable volatility models, actual volatility is modelled as a recurrent process. If the recurrent process is assumed to be stationary but unbounded, then the singular boundary at infinity becomes unreachable. However, in many volatility models, such as the GARCH diffusion model, the risk adjusted volatility process is different, because by assuming log-utility, the volatility process can explode, reaching infinity within a finite expected time. It is also the case that the volatility of the volatility (volofvol) process can easily explode with similar results. This latter occurrence can lead to the collapse of the martingale pricing assumption or the restricting of the pricing process to be only a local martingale. For example, Lewis examines a GARCH-diffusion model using the Feller boundary conditions with either risk neutrality or logarithmic utility. Lewis (2000) shows that it is possible for volatility to reach ∞ within a finite time (i.e. a volatility explosion according to the Feller criteria) depending on the type of boundary condition.

To deal with such failures, Lewis (2000) suggests the addition of an explosion correction term based on the work of Sin (1998). Lewis examines the impact on a GARCH-diffusion model, a CEV model and a square root type model, of adding an adjustment term (to either the volatility of the underlying, or to the volofvol) to take account of volatility explosion (which he interprets by taking volatility explosion probabilities as being stock price absorption probabilities). He finds in almost all of the cases that the addition of the explosion adjustment is sufficient to retrieve local martingale pricing, a finding that he confirms using Monte Carlo simulation.

The behaviour of option pricing models in the face of either the price of the underlying or its volatility tend to ∞ within a finite time. Assuming a zero interest rate and zero dividend growth, the BS model tells us that infinite spot price when combined with unit strike and normal levels of volatility simply gives unit option price, irrespective of the risk measure. But the extreme volatility behaviour is more interesting in the case of stochastic volatility models. If BS theory held (despite having stochastic volatility) then infinite volatility would return unit option price for many models such as GARCH-diffusion. However, Lewis supplies a counter example where the option price is less than unity in such circumstances, illustrating his arguments and approach with his three chosen models. In the case of the GARCH-diffusion and square-root type models he indeed shows that unit option price results from infinite volatility. In his modified square-root model Lewis finds that at large volatility values the drift term differs only by a half-integer power from his original model, despite the model being completely different, but, most importantly, the option value only reaches just under a third of the predicted BS value. Lewis checks his results by Monte Carlo simulation and finds that as the volatility is increased (with accompanying reduction in the time step) the observed errors fall as the volatility approaches infinity. He interprets this behaviour as the process becoming dominated by the deterministic limit as volatility approaches unity.

2.4 Robust optimal control and catastrophe derivative pricing

As can be readily appreciated from the previous section, all of the existing research into catastrophe derivative pricing takes as its starting point a model based on some type of distribution. In so doing the key question of whether the true probability law for the posited stochastic process selected is actually known is invariably overlooked. It is surprising that in the area of catastrophe derivatives pricing where it is particularly difficult if not impossible (by the very definition of the nature of catastrophes) to know their true probability law, that the issue of model misspecification in the face of uncertainty has not been considered. Research in decision making under uncertainty contains several strands of work to account for imprecise knowledge about the underlying probability law. A popular approach has been to model parameter uncertainty or estimation risk using a Bayesian prior which makes the strong assumption that

the prior belief can be modelled by a probability measure. However, this assumption is clearly inconsistent with evidence from experimental economics and psychology such as the Elsberg (1961) paradox where uncertainty exists about the states of the world and about the model itself. The key issue is that whilst there is some agreement concerning whether states of the world can be described by an objective probability law, Elsberg directly challenges the notion that model uncertainty can be described by a subjective probability prior. This is particularly interesting in the context of catastrophe derivatives where catastrophes occur in a totally arbitrary manner, because forming rational priors in such circumstances is virtually impossible due to the clearly unknown type, size and incidence of catastrophes. For example, how is it possible to form a rational prior about the likely size and impact of terrorist activity ? It is highly questionable whether for instance any rational prior could have handled the likelihood of the attack on the World Trade Centre.

In answer to such concerns two broad alternatives to the Bayesian approach have been developed. On the one hand the approach of Dow and Werlang (1992) examines the problem of portfolio choice under Knightian uncertainty using the axioms of expected utility theory. This is extended by the 1994 paper on discrete time by Epstein and Wang and by the 2000 paper on continuous time by Chen and Epstein to incorporate Knightian uncertainty via the use of multiple priors. Getting closer to the problem at hand, Epstein and Miao (2000) and Uppal and Wang (2002) extend the approach by examining both model uncertainty and varying level of ambiguity concerning the marginal probability laws that govern alternative states. Uppal and Wang offer a more appealing approach because they use a reference model as a means of differentiating between the alternative priors concerning the true model. The key point is that this feature leads to a problem formulation containing sufficient differentiability for deriving the HJB equation required to ensure robustness in the resulting pricing model.

On the other hand has been the work of Hansen and Sargent et al (1999, 2001) which introduces model misspecification and a preference for robustness into the Lucas rational expectations model. In the Hansen and Sargent et al approach the possibility of model misspecification is taken account of in decision making through the use of a parameter reflecting the overall level of ambiguity. Uppal and Wang extend this idea by accounting for differences in the degree of ambiguity about the marginal probability laws for the alternative states. Their approach is particularly attractive since it is sufficiently general to take account of different levels of ambiguity over the joint distribution of the underlying and of multiple subgroups.

Uppal and Wang's paper is primarily concerned with portfolio choice, so the detailed mechanics of their approach are not relevant to the problem of catastrophe derivatives. However, two points are relevant. First, the fundamental assumption underlying the intertemporal additive expected utility approach is

that the representative agent knows precisely the true probability law is clearly a difficult assumption to defend in the case of catastrophic events as already indicated. An alternative is of course to use recursive utility (as in Epstein and Zin, 1989). Unfortunately, Uppal and Wang show that once varying degrees of uncertainty ambiguity are allowed across assets, then representative agent preferences are no longer equivalent to recursive utility - so another approach is clearly required.

2.4.1 Uncertainty and feedback

As can be seen from the above analysis of the current approaches to the pricing of options on almost any type of underlying, concerns about the robustness or the stability of the pricing approaches and algorithms has not been a matter of significant concern to the overwhelming majority of researchers. This is particularly surprising in the face of both the increased number and magnitude of extreme or catastrophic events. This lack of concern may arguably be because such problems have so far presented much of financial and insurance mathematics with challenges that have been predominantly met by the use or abuse of a wide range of variations in martingale measure theory.

However, options on extreme or catastrophic events present two distinct though inevitably related problems, which on closer inspection, are not well treated by classical approaches when the results of applying those approaches are subjected to tests for robustness and stability. First, is the problem of pricing options on extreme or catastrophic events and the second is developing a model for such underlyings that is robust and stable in the face of a wide variety of disturbances. The previous sections of this chapter have already dealt with the current range of techniques that have so far been applied to the problem of pricing options on extreme or catastrophic events. The weaknesses of the Black-Scholes approach have also been examined from the traditional perspective of martingale measure theory. However, remembering that the Black-Scholes approach is effectively a bounded quadratic variation case, it is worth pointing out a further weakness. Namely, that it is not sufficient to assume that the quadratic variation be bounded in order to be able to calculate a value, indeed, the exact volatility must also be known before valuation is possible.

This section will instead concentrate on developing a robust and stable approach to pricing options on extreme and catastrophic events. Conceptually, the construction of the approach is as follows. First, the need for feedback and control is analysed. Second, basic state space approach inherent in the use of linear programming techniques is extended to robust optimal control. Third a robust optimal control pricing model is constructed for European options for non-catastrophic underlyings, which is then extended to catastrophic derivatives.

2.4.2 Robust optimal control

The reader of this thesis is advised to be familiar with control theory in general and with robust optimal control in particular, as covered in chapter one of this thesis. Robust control theory originated from the need to deal with systems that contain modelling uncertainty which occurs due to uncertain parameters and unmodelled dynamics. In addition, model simplification (due causes such as linearization and model reduction) and inadequate model identification (resulting from incomplete or imprecise data from identification experiments) can also cause problems in modelling the dynamics of a system. At its simplest, robust control theory is a method of dealing with the analysis and synthesis of control systems to satisfy various types and forms of stability and performance criteria given model uncertainty.

At a relatively high level of abstraction, there are several mathematical techniques that have been developed for robust control systems. The techniques correspond broadly to the type of uncertainty encountered in the modelling process. At the simplest level, is the choice between linear and non-linear models, followed by the question of parametric or dynamic models and finally the issue of structured or unstructured uncertainty. Attempting to characterise the various types of approach, could proceed as follows. The stability analysis of linear systems with uncertain parameters can be regarded as the issue of determining whether a family of polynomials has roots only in the left half-plane. Moving on to uncertain dynamics, there are operator-theoretic methods such as the small-gain theorem available. Whilst at the most general level there is the structured singular value (μ) theory.

Getting closer to the issue of option pricing, it will be recalled from the section on the use of linear programming, that option pricing can be regarded as an optimisation problem. The issue therefore centres around which technique to use and why. Thus, the performance of a control system can be extremely well characterised by suitably weighted signal norms such as keeping tracking errors with defined limits. Control design theories result from using different norms as signals. For example

- Bounding the $\mathcal{L}_2$ norm of an output for a fixed exogenous input leads to H_2 optimisation.
- Bounding the $\mathcal{L}_2$ induced norm from input to output leads to $H\infty$ optimisation.
- Bounding the $\mathcal{L}_\infty$ induced norm leads to $\mathcal{L}_1$ optimisation.

Of the three above approaches, it is $H\infty$ that is the focus of interest in the next section, so a brief introduction is relevant. $H\infty$ control has emerged as an effective control technique, as it brings together both performance and robustness requirements into an integrated index (the $H\infty$ norm). An optimal balance between system performance and robustness is then obtained by minimisation of this index. The controller developed can either have 2 degrees of freedom, consisting of a feed-forward prefilter and a

backward controller, or it can be a single degree of freedom form, possessing only a feedback controller. The $H\infty$ controller therefore handles system modelling errors as well as noise uncertainty, providing robust and stable controllers. It is sometimes referred to as the minimax controller as its objective is to minimise the maximum estimation error. Simon (2000) provides a highly intuitive and accessible introduction to the theory, construction and interpretation of $H\infty$ controllers.

2.4.3 Robust optimal control pricing of options

Wystup (2000) deals with the H_2 problem, by examining the treatment of the pricing of an option as a singular stochastic control problem, concentrating in particular on the pricing of exotic options such as barrier options. Whilst it is true that one of Wystup's central concerns is with the use of control theory in the pricing of barrier options which are indeed options on extreme values, such options are not the focus of this chapter, but as noted in the conclusions to this chapter, could form the focus for a possible direction for future research. Still within the H_2 paradigm, Dempster (1999) treats American exercise options as a linear programming problem, whilst Martini (2000) develops the pricing of American options as a degenerate stochastic control problem and finds that the option price is a unique bounded and continuous viscosity solution of a fully non-linear parabolic equation of the form

$$-\frac{\partial u^*}{\partial t}(t,x) = (Au^*)^+(t,x)\,,\ t<T,\ x>0,\ u^*(T,x) = \varphi(x)\,, \tag{2.33}$$

where A is the infinitesimal generator of the Black-Scholes model.

Following work by Rapaport (1998) and Bernhard, Crepey and Rapaport (2000) in contrast, suggest the use of a game theoretic approach to characterise the optimal control problem. Bernhard (2000) proposes a non-linear $H\infty$ robust control theoretic approach to the pricing of options. The particular attraction of Bernhard's (2000) approach is that it is based on describing the set of possible trajectories for the underlying's price process, but does not require any probabilistic rule for the price process - an issue which is returned to in the following section.

Having already examined the theoretical and practical aspects of pricing options using a linear programming approach, the control theoretic approach can be seen as an extension of the concept of viewing the option value function as a dynamic system that is influenced by the two exogenous inputs the price of the underlying and the trading strategy of the holder. From a control theory perspective there is a natural analogue which suggests that the goal of replicating the option payoff with a riskless portfolio can also be viewed as one of attempting to control the value of the replicating portfolio in the presence of disturbances. Extending this concept a little further, leads to the concept of attempting to protect a desired result (i.e. matching the value of the replicating portfolio to the value of the option) against

all possible disturbances contained in some bounded set of disturbances. What follows closely follows Bernhard (2000), by providing both a discrete and continuous time version of the approach. The two extensions that are new are the application of Bernhard (2000) to the pricing of catastrophe options and the extension to using viscosity solutions to solve the resulting Hamilton-Jacobi-Bellman equations.

Assume that there exists some underlying instrument, that possesses a time dependant and unpredictable market price $S(t)$ at time t. Further assume that there exists Ω, which is the set of possible time functions $S(\cdot)$. The economy within which the problem is set contains a riskless bond, whose unit value at expiry equals unity, but at any time $t \in [0,T]$ is $R(t)$, given its rate ρ, such that it produces either of the following discount factors in the discretised time interval where $N = \frac{T}{h}$ and $k = \frac{t}{h}$

$$R(t) = \exp(-\rho(k-N)h) \text{ in continuous time, or} \tag{2.34}$$

$$R(t) = (1+\rho)^{(k-N)h} \text{ in discrete time} \tag{2.35}$$

Assuming that the initial objective is to replicate a European option using a theoretical security whose value at time T is some convex function, M, of $S(T)$, given the following payoff

$$M(s) = \max\{s-K,0\} \text{ for a call option and} \tag{2.36}$$

$$M(s) = \max\{K-s,0\} \text{ for a put option} \tag{2.37}$$

then the problem is therefore to deal with a value function consisting of x units of the underlying instrument and y units of the riskless bond, such that the value of this portfolio at time t is given by

$$w(t) = x(t)S(t) + y(t)R(t) \tag{2.38}$$

is controlled by trading strategies of the form

$$x(t) = \varphi(t, S(t)) \tag{2.39}$$

In discrete time this simply involves buying and selling at time t the required amounts of the underlying instrument, which is then held until the next trading event. In continuous time, the idea of instantaneous re-hedging using the underlying instrument is less clear, but by assuming that the functions $S(\cdot)$ are limited to the set Ω of continuous functions, the result is at least unambiguous. As far as the riskless bond is concerned, trading activity is bounded by the requirement that the replicating portfolio be self-

financing, such that in continuous time

$$S(t)\,dx(t) + R(t)\,dy(t) = 0, \text{ or in discrete time} \tag{2.40}$$

$$S(t)(x(t) - x(t-1)) + R(t)(y(t) - y(t-1)) = 0 \tag{2.41}$$

such that, given that $dR(t) = \rho R(t)\,dt$ and that $y(t)\,R(t) = w(t) - x(t)\,S(t)$ the model becomes

$$dw(t) = x(t)\,dS(t) + \rho(w(t) - x(t)\,S(t)\,(dt)) \tag{2.42}$$

in continuous time, or with a time shift in discrete time

$$w(t+1) = S(t+1)\,x(t) + \rho(w(t) - x(t)\,S(t)) \tag{2.43}$$

Given the approach of the linear programming solution to the option pricing problem, it is easy to see that the equation describing the dynamics of $w(t)$ defines a system which has inputs of $S(\cdot)$ and $x(\cdot)$ and output of $w(t)$. The replicating portfolio must have a value at least as good as the option it is intended to hedge, which in this context translates to controllability in the face of all admissible disturbances. This constraint acts as an upper bound on the equilibrium price which leads to a value for the hedging portfolio greater than or equal to that of the option for all admissible disturbances $S(\cdot)$, given the existence of a trading strategy where $x(t)$ depends solely on past and present information. The problem is therefore to find the set of initial hedging portfolios capable of generating a set of admissible terminal states. To ensure that this set is capturable, it is usual to represent $S(\cdot)$ in continuous time as being the output of a first order controlled system of a form such as

$$\dot{S}(t) = S(t)(\mu + \sigma v(t)), \text{ where } |v(t)| \leqslant 1 \tag{2.44}$$

which, where $v(\cdot)$ is a measurable noise, whose only purpose is to define the set of admissible disturbances $S(\cdot)$ *not* a Brownian motion. This in turn implies restricting $S(\cdot)$ to contain only absolutely continuous, bounded variation positive functions; or, equivalently, in discrete time

$$S(t+1) = S(t)(1 + \mu + \sigma v(t)) \tag{2.45}$$

$$= S(t)(m + \sigma v(t)), \text{where } |v(t)| \leqslant 1 \tag{2.46}$$

where $v(\cdot)$ ranges over the continuous interval $[-1, +1]$ and *not* over the finite set $\{-1, +1\}$ (as in the

case of Cox Ross and Rubinstein's binomial lattice approach). Defining the following quantities

$$m - \sigma = a \tag{2.47}$$

$$m + \sigma = b \tag{2.48}$$

$$1 + \mu = m \tag{2.49}$$

$$\mu - \rho = m - r = \lambda \tag{2.50}$$

Following both Zhou, Doyle and Glover (1996) and Bernhard (2000), the classical approach is then, given the two equations for a dynamic system as expressed by S and w, with v as the disturbance and x as the control and

$$\Lambda_T = \left\{ \begin{pmatrix} S \\ w \end{pmatrix} \in (\mathbb{R}^+)^2 \,|w \geqslant M(S) \right\} \tag{2.51}$$

as the set of admissible terminal states. Ziegler (1999) develops the link between game theory and the pricing options and points out that option valuation can be viewed as a dynamic game, the solution to which involves optimal decision making behaviour, whilst Bernhard (2000) looks in the reverse direction by seeing the control problem from the perspective of game theory by noting the correspondence between the modelling of dynamic games and robust control, by pointing out that

$$\forall v(\cdot), \begin{pmatrix} S(T) \\ w(T) \end{pmatrix} \in \Lambda_T \tag{2.52}$$

can also be written as

$$\inf_{v(\cdot)} [w(t) - M(S(T))] \geqslant 0 \tag{2.53}$$

and that the existence of an admissible strategy φ that ensures that 2.53 is at least equivalent to

$$\max_{\varphi} \inf_{v(\cdot)} [w(t) - M(S(T))] \geqslant 0 \tag{2.54}$$

if the maximum exists - a property that will be used shortly to propose a solution methodology. But for the purposes of option valuation, the problem to be solved is to find the initial states that can be controlled to the set Λ_T in the dynamic system

$$\dot{S} = S(\mu + \sigma v(t)), \ |v(t)| \leqslant 1 \tag{2.55}$$

$$\dot{w} = \rho w + Sx(t)(\lambda + \sigma v(t)) \tag{2.56}$$

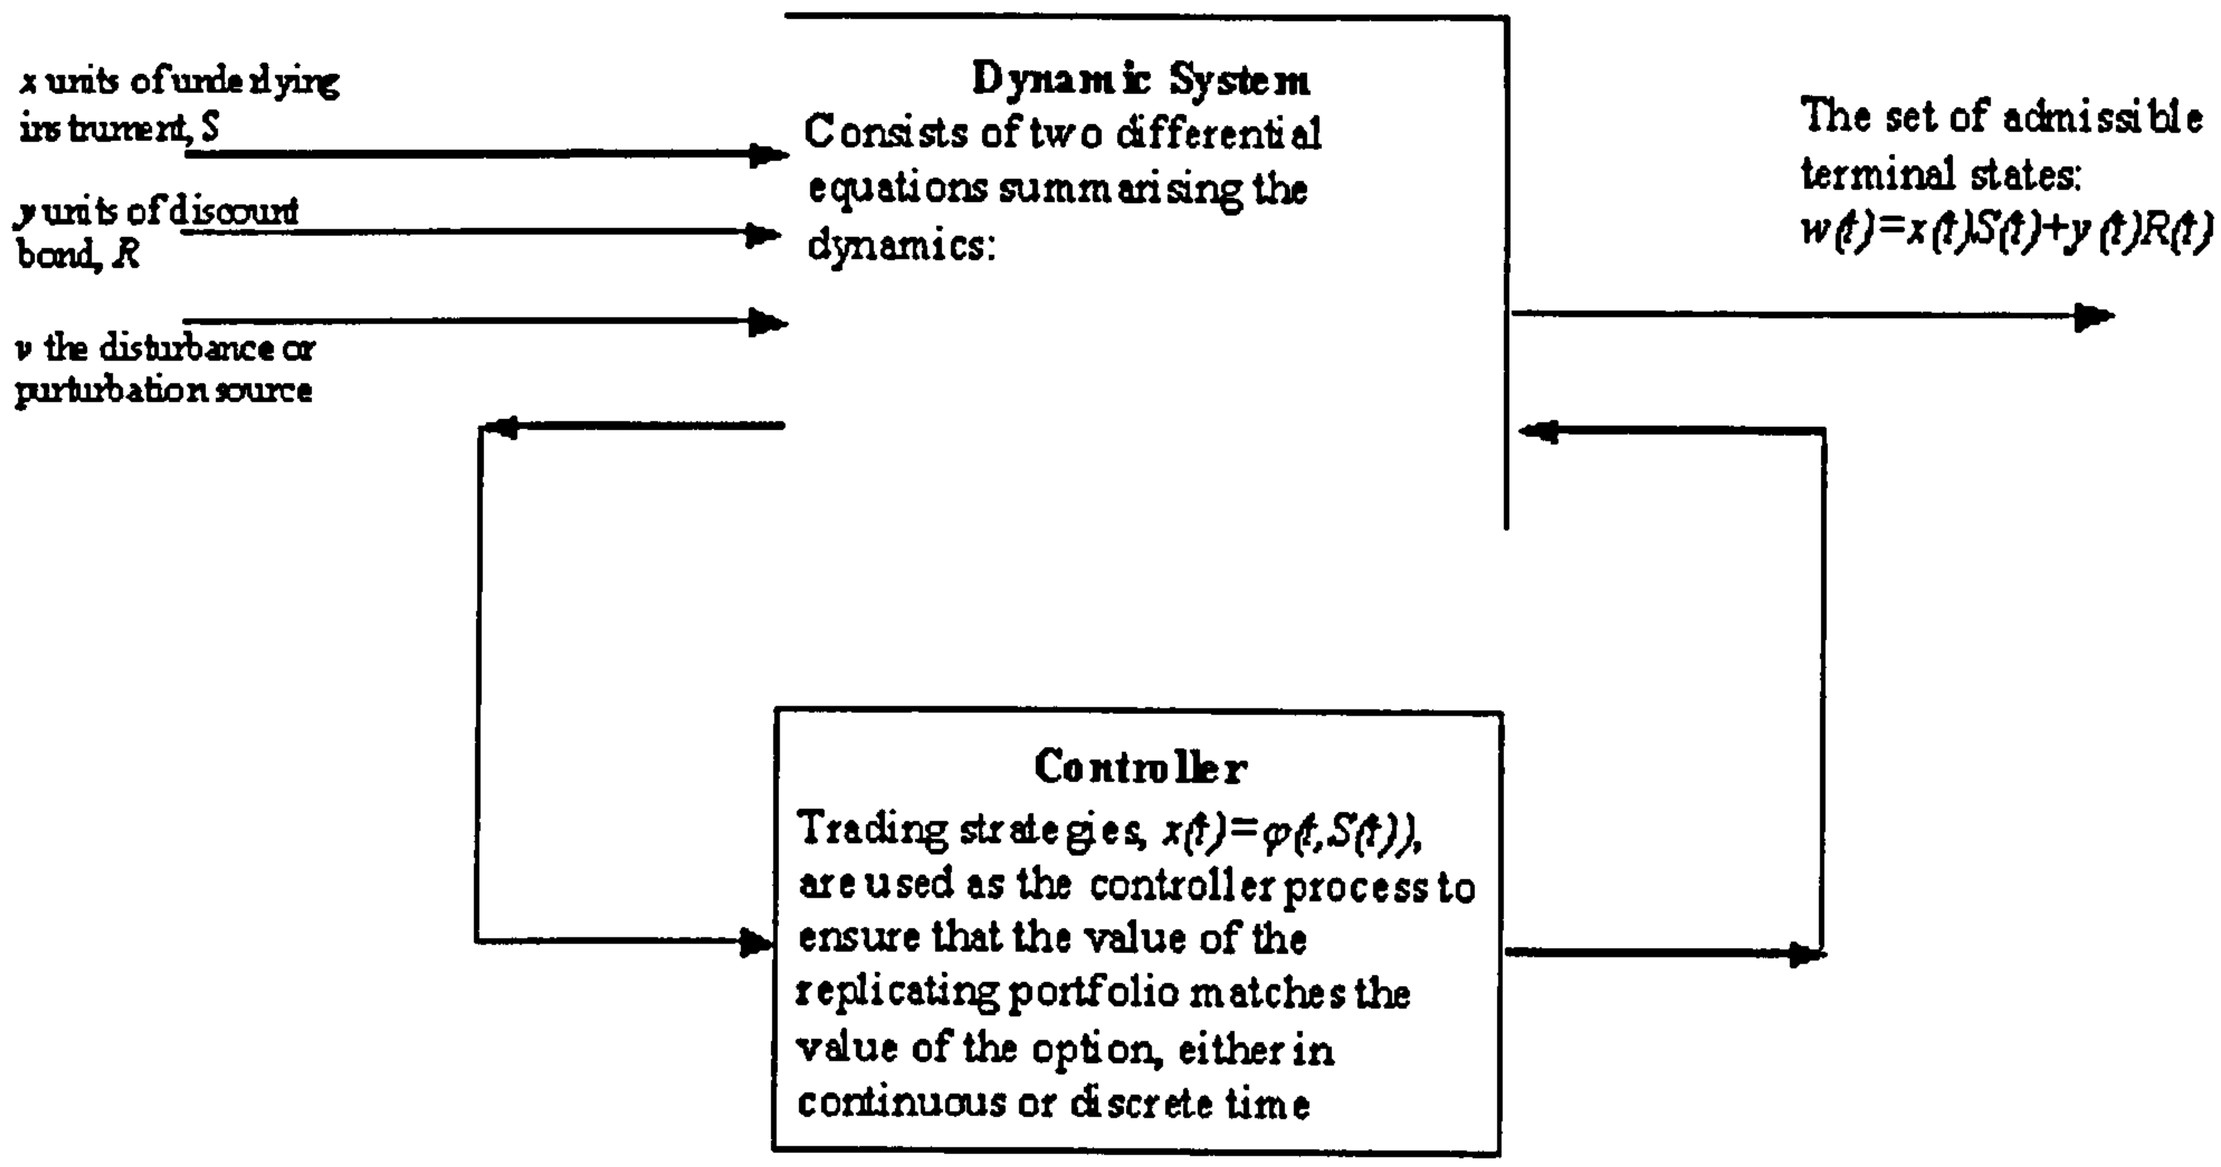

Figure 2-1: **Figure 2.1: Schematic representation of feedback control process**

given $x(t)$ as the control. Diagrammatically, this can be represented as follows

From the theory of dynamic games, see for example Basar and Bernhard (1995), it is known that the solution to this type of game can be found via the backwards construction of a barrier from the boundary of the capture set. Bernhard (2000) shows that a barrier is in fact the 2D manifold parameterised by t and S as $w = W(t,S)$ where $W(\cdot,\cdot)$ is given by

$$W(t,s) = \exp(\rho(t-T)) \tag{2.57}$$

This is because at a point (t,S,w) on this manifold, it is possible to compute a normal of the form

$$\eta(t,S,w) = \begin{pmatrix} \eta_t \\ \eta_S \\ \eta_w \end{pmatrix} = \begin{pmatrix} \rho\left[-RM\left(\frac{S}{R}\right) + M'\left(\frac{S}{R}\right)S\right] \\ -M'\left(\frac{S}{R}\right) \\ 1 \end{pmatrix} \tag{2.58}$$

using $R = \exp(\rho(t-T))$, so that it is then possible to form

$$H = \eta_t + \eta_S \dot{S} + \eta_w \dot{w} \tag{2.59}$$

$$= (\lambda + \sigma v) S \left(x - M'\left(\frac{S}{R}\right)\right) \tag{2.60}$$

which, given that $\lambda < \sigma$, means that v can be chosen such that $\inf_v H$ is non-positive. This in turn means that $\max_x \inf_v H$ occurs at $x = M'(\frac{S}{R})$ at a value of zero, such that the manifold $w = W(t, S)$ definitely constitutes a barrier, so that selecting a trading strategy $x(t) = M'\left(\frac{S(t)}{R(t)}\right)$ from any state on the manifold prevents the system from exceeding the barrier. Whilst the continuous time case indicates the core of the approach it does not provide a clear exposition of the solution methodology that can be applied to solving the optimal control problem arising from the application of such principles. To obtain such a clear view, it is essential to explore the discrete time case.

Continuing with the concept of the equivalence of single-player games and optimal control, it is well known (e.g. Bernhard (1995, 2000) and Ziegler (1995)) that the method of dynamic programming can be used to produce the optimal solution by solving a sequence of static optimisation problems stepping backwards through time. In two-player dynamic games the counterpart to this is a recursive equation that involves determining the saddle-point solutions of static games by stepping backwards through time. The equation is known as the Isaacs equation[3]. The discrete version of the continuous time dynamic system can be stated as follows

$$S(t+1) = (m + \sigma v(t))\, S(t) \tag{2.61}$$

$$w(t) = rw(t) + (\lambda + \sigma v(t))\, S(t)\, x(t) \tag{2.62}$$

changing variables to discounted versions, gives

$$u(t) = \frac{S(t)}{R(t)}, \qquad v(t) = \frac{w(t)}{R(t)} \tag{2.63}$$

and changing control variables to

$$\tau(t) = \frac{1}{r}(\lambda + \sigma v) \in \left[\tilde{a}, \tilde{b}\right] = \left[\frac{a}{r}, \frac{b}{r}\right], \qquad \psi(t) = x(t)\, u(t) \tag{2.64}$$

so that

$$u(t+1) = \tau(t)\, u(t) \tag{2.65}$$

$$v(t+1) = v(t) + (\tau(t) - 1)\, \psi(t) \tag{2.66}$$

[3]Isaacs equation was developed by Rufus Isaacs during the early 1950's - it is also known as the Hamilton-Jacobi-Isaacs equation and is a generalisation of the Hamilton-Jacobi-Bellman partial differential equation that provides a sufficient condition in optimal control.

where the set of admissible states remains

$$\Lambda_T = \left\{ \begin{pmatrix} u \\ v \end{pmatrix} \in \left(\mathbb{R}^+\right)^2 | v \geqslant M(u) \right\} \tag{2.67}$$

Therefore, let Λ_t be the set at time t, of sets capturable to Λ_T at the terminal time such that $(\bar{u}, \bar{v}) \in \Lambda_t$ iff there exists an admissible strategy

$$\psi\left(t'\right) \varphi\left(t', u\left(t'\right), v\left(t'\right)\right), \qquad t' \geqslant t \tag{2.68}$$

such that the system described by equations 2.65 and 2.66 initialised at $(\bar{u}, \bar{v})$ driven by that particular strategy is guaranteed to terminate in admissible state at time T for all admissible disturbances $\tau(\cdot)$. Following Ziegler (1999) and Bernhard (2000), note that if a state $(\bar{u}, \bar{v})$ belongs to Λ_t, then so do all $(\bar{u}, v)$ states for $v \geqslant \bar{v}$, so that if

$$V_t(u) = \min \left\{ v | \begin{pmatrix} u \\ v \end{pmatrix} \in \Lambda_t \right\} \tag{2.69}$$

where the function $V_t(\cdot)$ completely describes the set Λ_t and its epigraph (i.e. the set of points (u, v) such that $v \geqslant f(u)$). The classical dynamic programming solution to this problem is to solve it by stepping backwards one state at a time and solving each arising static optimisation problem. The sequence of sets Λ_t and the associated hedging strategy $\psi(t) = \varphi(t, S)$ are simultaneously defined by Isaacs' equation

$$\Pi_{\Lambda_t}(u, v) = \max_{\psi} \inf_{\tau} \Pi_{\Lambda_{t+1}}(\tau(u, v) + (\tau - 1)\psi) \tag{2.70}$$

where

$$\Pi_{\Lambda_{t+1}}\left(u', v'\right) = \left\{ \begin{array}{l} 1 \text{ if } v' \geqslant V_{t+1}\left(u'\right) \\ 0 \text{ if } v' \leqslant V_{t+1}\left(u'\right) \end{array} \right\} \tag{2.71}$$

and

$$V_t(u) = \min \{ v | \Pi_{\Lambda_t}(u, v) = 1 \} \tag{2.72}$$

This is a fairly simple computational procedure which provides the facility to calculate $V_0(S)$ for every admissible $S(0) = S$, which therefore forms the underpinning of a method for option valuation.

There are two further issues to consider before extending the robust optimal control approach to catastrophe options, namely, the convexity or otherwise of M and step size. Taking convexity of $M(\cdot)$ first, it is clear that if $M(\cdot)$ is convex, then the set Λ_T must also be convex and that for a given state

$(u(t-1), v(t-1))$ and a fixed strategy $\psi(T-1)$, that the set of possible $(u(T), v(T))$ is a line segment in (u, v) space. This fixed strategy can only be contained in the convex set Λ_T if and only its end points lie within Λ_T, so that it is necessary to determine those (u, v) combinations for which there is a fixed strategy ψ that satisfies the following conditions

$$v + (\tilde{a} - 1)\psi \leqslant M(\tilde{a}u) \tag{2.73}$$

$$v + \left(\tilde{b} - 1\right)\psi \leqslant M\left(\tilde{b}u\right) \tag{2.74}$$

Recalling that $\tilde{a} < 1 < \tilde{b}$, it is therefore clear that equation 2.73 and equation 2.74 are satisfied if

$$\frac{1}{\tilde{b}-1}\left[-v + M\left(\tilde{b}u\right)\right] \leqslant \frac{1}{\tilde{a}-1}\left[v - M(\tilde{a}u)\right] \tag{2.75}$$

and that there must exist a strategy ψ, if and only if equation 2.75 is satisfied, which in turn provides a description of the set Λ_{T-1}. Bernhard (2000) shows that due to the convexity of M, the value function V_{T-1} is convex, so that the identical calculation procedure can be used by stepping backwards in time by defining a sequence of functions $V_t(\cdot)$, which after returning to the original untransformed variables and setting $W(t, S) = R(t)V_t\left(\frac{S}{R_t}\right) = r^{t-T}V_t\left(r^{T-t}S\right)$, allows the above recursion to be re-written as

$$W(t-1, S) = \frac{1}{r}\left[\frac{b-r}{b-a}W(t, aS) + \frac{r-a}{b-a}W(t, bS)\right], \qquad W(T, S) = M(S) \tag{2.76}$$

As with equations 2.65 and 2.66 above, equation 2.76 provides the relationship to calculate a sequence of functions $W(t, S)$ for a given $M(S)$ and also yield an equilibrium price $W(0, S(0))$ for the underlying asset. It is readily apparent from the results presented in the next section that equation 2.76 gives a higher price than the vanilla Black-Scholes model. This arises for two reasons. First, and most obviously, as the optimal condition is being attained by being robust against all possible disturbances, then the hedging portfolio is bound to be of greater value as it is being used to replicate an option whose value is expected to be far more volatile and have a higher likelihood of attaining more extreme values. Second, as the trading of the replicating portfolio is being carried out less frequently than under the assumed continuous re-hedging associated with Black-Scholes, then the value will move only discontinuously, with the likelihood that larger divergences between the value of the hedging portfolio and the price of the option.

2.4.4 Robust optimal control and the pricing of catastrophe bonds and options

It will be recalled from the previous section, that a key feature of the robust optimal control approach is the fact that no notion of a probability distribution is required to be able to unambiguously derive the price for a derivative security. This concept is now taken and applied to the pricing of derivatives on extreme events such as catastrophe options. The approach is particularly attractive as it does not rely on any particular distribution to produce a price, which in turn ensures that the option price is robust to variations in distribution. The argument will be developed as follows: the first step is to apply the Hamilton-Jacobi-Isaacs (HJI) equation to the pricing of catastrophe options; the second step, is to examine the computational issues of convexity and the impact of discretisation on the parameterisation of the state space; the final step is to develop a practical pricing algorithm and apply it to real data in the cases of both single event catastrophes and catastrophe indexes, in the cases of both options and bonds.

Non-convexity of the payoff function

The first point to note is that it is not uncommon for the HJI equations not to have a continuously differentiable solution, particularly in cases where the value function is not continuously differentiable. In such cases, a solution concept that replaces the classical robust optimal control approach is the viscosity solution, first introduced by Crandall and Lions (1983). The concept is that under quite general assumptions on the structure of the first-order partial differential equation (without necessarily bounding the conditions to any specific form or class of optimisation problem), Crandall and Lions (1983) show that the viscosity solution is unique and that in the context of optimal control problems, the value function is necessarily the viscosity solution of the associated HJI equation. The final link to the problem at hand is provided by Souganidis (1985) who showed that a counterpart of this result exists in zero-sum two-player differentiable games, with the value function being the unique viscosity solution of an associated HJI equation. This relationship between viscosity solutions and value functions was further studied in Lions and Souganidis (1985), which established, using viscosity sub-solutions and viscosity super-solutions, the equivalence between the concept of viscosity solutions and various definitions of value for differentiable games. Gomes (2001) extends the work further by explicitly linking viscosity solutions and optimal control. The main reason for introducing these concepts is the need to deal with the issue of stability and robustness under small perturbations when establishing the notion of a robust and optimal control for options on extreme value derivatives.

When pricing options on non-extreme or non-catastrophic events using robust optimal control, crucial to the approach was the assumption of convexity of the payoff function $M(s)$ at terminal time T. However,

the payoff for a catastrophe option is

$$M(s) = \min\{\max\{0, s-K\}, L-K\} \tag{2.77}$$

which is non-convex, which in turn means that the requirement that the end points of the segment described by equations 2.65 and 2.66 must be in Λ_{t+1}, is no longer sufficient to ensure that (u_{t+1}, v_{t+1}) be contained in Λ_{t+1} for all possible values of $v(t)$. As such, no direction can be provided for a possible hedging strategy, which is of course a fundamental and non-acceptable weakness.

This chapter explores two possible strategies for finding solutions to this problem, namely, numerical approximation and second the use of Lyapunov equations.

Numerical implementation of the Isaacs equations

First, is the most obvious one of using a numerical implementation of the Isaacs' equations (see equations 2.70 and 2.71), such that for each time step t, the state space in u is traversed searching for admissible v values, which once found can be used to trigger a change in u. By following this procedure, both a valuation and hedging strategy can be defined simultaneously. As Bernhard (2000) points out (and as the results presented in this chapter confirm), this procedure produces more robust valuation and hedging results than the equivalent binary tree methodology - at least for the period and markets studied in this chapter. It is shown that the binary tree methodology under-prices the value of the replicating portfolio and provides no efficient computational route for option hedging parameters other than finite differencing of the binary tree input parameters. It is also demonstrated that the popular Cox or Poisson point processes that give rise to partial integro-differential equations suffer from similar limitations.

Parameterisation of the state space therefore proceeds as follows. The first step is feasible and efficient parameterisation of the state space. One of the major weaknesses of the regular discretisation approach to the pricing of options is the fact that the price may not improve asymptotically with either the addition of extra time steps in the time dimension or with reducing step size in the price dimension. Beyond some point the option price simply oscillates around a value but does not converge. One of the reasons for this is that the dispersal of coverage of the state space is not being improved, even by reducing step size and increasing frequency of occurrence. In the context of this thesis, this problem will be addressed by using a method of discretisation of the state space which simultaneously incorporates variable step sizes in both the time and price dimensions. Aside from the obvious computational burden imposed by regular discretisation, uniform discretisation of state spaces suffers from impractical computational requirements when the size of the discretisation step is small. One method of addressing this problem in the context of optimal control is suggested by Munos and Moore (2000) who use a refining process which starts with

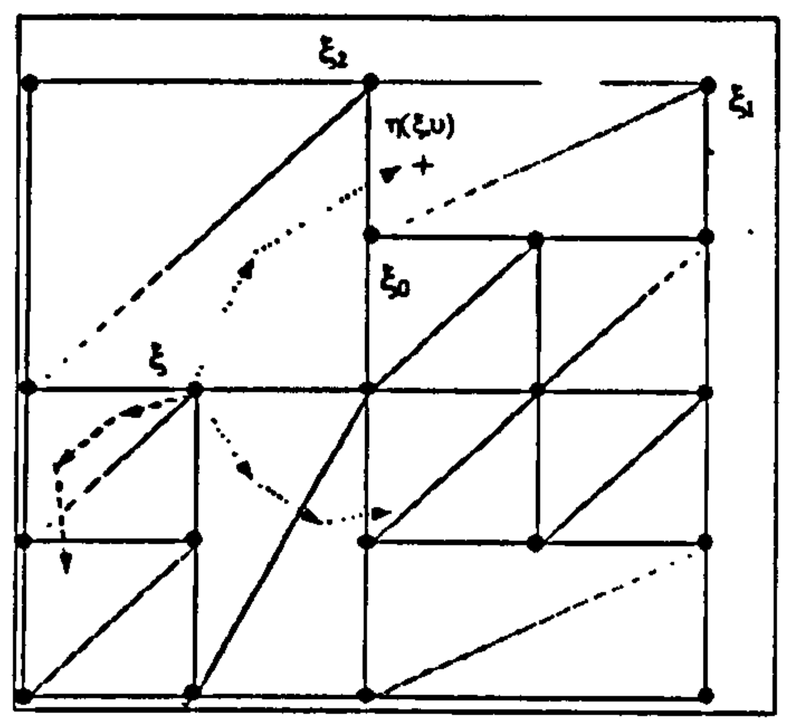

Figure 2-2: **Figure 2.2: Discretisation using Markov decision process**

an initial coarse grid which is successively refined by a splitting process until some required degree of accuracy is achieved.

Munos and Moore (2000) utilise the structure of a *kd* tree to produce a variable resolution discretisation of the state space, with the root of the tree covering the whole hyper-rectangle of the state space. For each "leaf" of the *kd* tree they use a Kuhn triangulation to linearly interpolate inside the rectangle. This defines a class of functions called barycentric interpolators that are piecewise linear, continuous inside each rectangle, but may be discontinuous at the boundary between two rectangles. The discretisation process is based on finite element method of Kushner and Dupuis (1992), which approximates for any discretisation of the state space, the continuous deterministic control process by a Markov Decision Process (MDP). It is important to note that the stochastic aspect of the MDP emanates from the discretisation process used and not from the continuous problem itself. Figure 2-2 provides a simplified schematic of the operation of the discretisation process.

To construct the *kd* tree, view the state space as the set Ξ of corners of the tree, such that for every corner ξ and control u, the corresponding trajectory $x(t)$ is approximated by integrating the state dynamics from the initial state ξ for a constant control u, during some time period, t, until it enters inside a new grid at some point (e.g. $\eta(\xi, u)$ in figure 2-2). The value interpolated at $\eta(\xi, u)$ is just a linear combination of vertices with positive coefficients that sum to unity, such that performing the interpolation is mathematically equivalent to probabilistically jumping to a vertex. Munos and Moore (2000) study a number of possible schemes for performing the interpolation, but conclude that it is necessary to combine both local and global splitting criteria to produce an efficient state space propagation mechanism. They utilise the dual notions of influence and variance. The states of highest influence are those where there is a change in the optimal control. The concept of state variance is slightly more complex, as these are the states with the highest uncertainty on the quality of the approximation of the value function and

so represent those states where the approximation of the value function could improve the most for any increase in accuracy associated with splitting. The variance is a measure of the extent of uncertainty in the value function attributable to the discretisation process and therefore provides an estimate of the quality of the approximation of the value function for a given discretisation. It was therefore decided that for catastrophic derivatives, the state space discretisation would be performed using a variance-influence global splitting approach using the Munos and Moore (2000) methodology.

Lyapunov equations

Of the two methods for dealing with the non-convexity of the payoff function, the Lyapunov equation is arguably the more elegant and flexible as it can be generalised to deal with more realistic non-linear problems and can be cast in a state-space setting making it ideal for dealing with securities that have prices that depend on different states. As F&K demonstrate, the existence of a robust control Lyapunov function implies robust stabilisability. As explained in chapter 1, it is possible to compute the value of a pointwise min-norm control law at any point by solving a static minimisation problem that is convex on the control space and is completely determined by the available data. The following sections apply the F&K approach and develop robust catastrophe bond and catastrophe option valuation models, that are both flexible, stable and robust.

Pricing catastrophe bonds using the traditional approach

Catastrophe bonds, or CAT bonds for short, are a way for insurers to access the greater risk bearing capacity of capital markets which are approximately 75 times the size of the insurance markets. CAT bonds are generally issued through some form of special purpose vehicle which acts as a trust thereby removing almost entirely associated credit risk. Though the CAT bond market is still in its relative infancy, Penalva-Zuasti (1997) found CAT bonds to be significantly more expensive than comparable competitive reinsurance contracts. An interesting question for this research, in addition to the issue of robustness, is therefore the source of this valuation disparity. Is it merely a consequence of lack of investor familiarity with CAT bonds ? Bantwal and Kunreuther (1999), approaching the question from the perspective of explaining the uncertainties inherent in existing models, suggest that ambiguity aversion, loss aversion and uncertainty avoidance may account for the reluctance of investors to trade in CAT bonds. They attempt to illustrate the attractiveness of CAT bonds by simulating their behaviour under a variety of Monte Carlo generated scenarios.

Before applying the Lyapunov approach to the valuation of catastrophe bonds, it is necessary to consider in some detail the construction and features of such bonds. Catastrophe bonds are a particular

version of a type of generic instrument most usually known as a threshold bond. Although there are numerous variations in the structure of threshold bonds, most conform to either one of two structures. The first broad type of structure is very similar to a defaultable bond, whereby issue occurs at some value, coupons may or may not then be paid at a stated coupon rate (with or without a spread) and at an agreed, pre-determined frequency. After issue and prior to maturity, premature termination can occur if either accumulated losses exceed an agreed threshold level, or if there is a single catastrophic event that can trigger premature termination, whereupon a certain or random recovery amount may or may not be payable.

The second type of structure is more akin to an insurance policy where a bond is issued at some given value - say par - for an agreed period. At maturity, the investor receives an uncertain amount. During the life of the bond catastrophic losses may occur. If the accumulated losses exceed a threshold level, losses on the principal of the bond occur and the investor will receive some (random) recovery amount. The losses are usually linked proportionally to the original principal of the bond. If no losses occur prior to bond maturity, or if accumulated losses fail to breach the threshold level during that time, then at maturity the investor will receive his original principal plus interest and any agreed spread. The interest payments may or may not be guaranteed.

Such potential variation can give rise to a plethora of subtly different structures. However, Baryshnikov et al show that the payment or not of either terminal or regular periodic coupons make no material difference to the valuation process such that in both broad cases the calculation of the no-arbitrage price can be reduced (under fairly simple and non-restrictive conditions) to the following simple threshold idealisation. Issue occurs at some level, say par, at time 0 and maturity at time T. There is an agreed accumulated loss process, L_s and a threshold loss level, D. At maturity, the bond either terminates after repaying a certain pre-determined amount C_τ- referred to as catastrophic termination. Else if time τ (when the threshold loss level is breached) exceeds T, termination occurs at T, nothing is paid - referred to as normal termination. The traditional approach to this pricing problem is to treat τ as the first instant of a Poisson point process, N_t, on the interval $[0,T]$ independent of all other variables, but in such a way that the intensity λ of the process is progressive and predictable so that for any time interval $[t_1,t_2] \subset [0,T]$ the process $N_t - N_{t_1}$ on $[t_1,t_2]$ is measurable with respect to the σ-algebra generated by $\int_{t_1}^{t_2} \lambda_s$. This assumption is usually justified on the grounds that the threshold time is usually contingent upon specified natural event(s) which is/are independent of either or both economic activity and/or financial data. The additional assumption required is that there is a progressive process of discounting

rates r, such that the value of a single unit paid at time $t > s$ is given by

$$\exp(-R(s,t)) = \exp\left(-\int_s^t r(\xi)\,d\xi\right) \tag{2.78}$$

so that the price of a threshold bond is given by

$$V = (N, r, C) \tag{2.79}$$

where V_t should satisfy

$$V_t = \mathbf{E}\left(C_\tau \exp(-R(t,\tau))\,|\mathcal{F}_t\right) \tag{2.80}$$

where $\mathcal{F}_t$ is an adapted filtration in the usual fashion and where τ is the stopping time at which the bond terminates[4]

$$\tau = \inf\{t : N_t = 1\} \tag{2.81}$$

which, given the assumptions about the structure of the point process and that λ_s is the stochastic intensity process, therefore gives

$$V_t = \mathbf{E}\left(\int_t^T \exp(-R(t,s))\,C_s\lambda_s ds|\mathcal{F}_t\right) \tag{2.82}$$

Looking first at the case of variable payment at maturity, define the process

$$Z_s = \mathbf{E}(Z|\mathcal{F}_t) \tag{2.83}$$

where the required condition is that Z_s is a predictable process. This can be validly and usefully interpreted as an assumption that the maturity payment is not linked directly to the occurrence or timing of the breaching of the threshold. So that the price of a bond paying Z at maturity ($\{\tau \geqslant T\}$) at time $t < \tau$ can be represented (using $N_T - N_t = \int_t^T dN_S$) by

$$V_t = \mathbf{E}\left(Z\exp(-R(t,T)) - \int_t^T Z_s \exp(-R(t,s))\,dN_s|\mathcal{F}_t\right) \tag{2.84}$$

Incorporating regular coupon payments, P_s, prior to maturity is simple. Assuming that coupons cease at

[4]The question of whether the arbitrage and real measures governing the underlying process should coincide is dealt with by Baryshnikov et al which refers to work by Froot & O'Connell (1996). The latter demonstrate that the rise in reinsurance prices that followed the surge in the data on the frequencies and aggregate losses due to natural catastrophes were due to the exhaustion of industry reserves and not to a change in the fundamental probability measures.

time τ, then if $\tau < T$

$$V_t = \mathbf{E}\left(\int_t^T \exp\left(-R\left(t,s\right)\right) P_s \left(1 - N_s\right) ds | \mathcal{F}_t\right) \tag{2.85}$$

which, upon integration by parts gives

$$V_t = \mathbf{E}\left(\left(1 - N_T\right) Q_T + \int_t^T dN_s \int_t^s \exp\left(-R\left(t,s\right)\right) P_u du | \mathcal{F}_t\right) \tag{2.86}$$

where $Q_s = \int_s^t \exp\left(-R\left(\xi,s\right)\right) P_u du$, which given that Q_s is predictable collapses to the standard threshold case.

The final stage required to arrive at a useful specific pricing situation is to specify the type of events to underlying the CAT bond. At the most abstract level the bond is described by the accumulated loss process L_s and its threshold loss D (where the threshold event occurs at time τ when the accumulated losses exceed the threshold level: $\tau = \inf\left\{t : L_t \geqslant D\right\}$, so that $N_t = \mathbf{1}\left\{L_t \geqslant D\right\}$. The usual approach is to assume that there exists a Poisson point process M_s which describes the flow of potentially catastrophic events and that this process is assumed to be a predictable and bounded process m_s. The losses produced by each event in the flow are assumed to be IID random values $\left\{X_i\right\}_{i=1,\ldots}$ with a distribution function $F_x = \mathbf{P}\left\{X_i < x\right\}$, so that

$$L_t = \sum_{t_i < t} X_i \tag{2.87}$$

So that, as Baryshnikov et al show, the non-arbitrage price of the CAT bond with threshold D, catastrophic flow M_s and the distribution function F for losses of the form

$$F\left(x\right) = 1 - \frac{1}{1 + x^e} \tag{2.88}$$

is given by

$$V_t = \mathbf{E}_t\left(\int_t^T \exp\left(-R\left(t,s\right)\right) C_s \left(1 - F\left(D - L_s\right)\right) 1 L_s < D m_s ds | \mathcal{F}_t\right) \tag{2.89}$$

which, after reduction to an ordinary differential equation can ultimately be solved in quadrature as soon as m and C are known using well known solutions for partial differential equations with integral terms.

One of the major problems with the developing a pricing model of the type above for catastrophe derivatives is that of hedging the risks created. Even in the case of zero coupon CAT bonds, risk is uncorrelated with underlying financial market variables such as interest rate levels or aggregate consumption. The result is that cash-flows for a CAT bond can not be hedged by a portfolio of traditional (riskless) bonds because it is impossible to replicate the payoff profile of the CAT bond using the usual portfolio approach. The pricing framework is therefore set within an incomplete markets approach, with Cox and

Pedersen (2000) providing arguably the seminal paper using this methodology, by valuing CAT bonds using equilibrium pricing based on a state-space type model. They use a two step procedure which involves first estimating the interest rate dynamics under varying degrees of risk for states of the world that do not involve a catastrophe, which from a state-space perspective is equivalent to estimating local state prices for states that are independent of a catastrophe. This is achieved by constructing a term structure of interest rates model. The second step is to estimate the probability distribution for varying degrees of severity of the catastrophic risk, so that valuation can be completed by discounting the cashflows.

Pricing catastrophe bonds using a robust optimal control approach

Many historical studies of CAT bond pricing (see for example Penalva and Zuasti, 1997) identify historically high spreads in the CAT bond market with respect to domestic bonds and attempt to explain the existence of an apparent CAT bond premium. However, Froot et al (1995) show that CAT bonds out-performed domestic bonds on a rate of return basis and are less volatile than either stocks or bonds. The question therefore arises of whether CAT bond spreads are too high to be explained by standard financial theory. Bantwal and Kunreuther (1999) attempt to use a number of ideas from behavioural finance to isolate and explain the persistence of the so-called "CAT bond premium". One of their key points is that models typically depend on the frequency and severity of CAT events - both of which are sources of uncertainty with respect to the surrounding measurability.

Lee and Yu (2002) extend the CAT bond pricing debate by incorporating default risk, stochastic interest rates and a more generic loss process, which taken together, produce lower CAT bond prices. Their approach is to use Monte Carlo simulation to demonstrate the robustness of the results provided by CAT bond pricing models to a variety of realistic disaster scenarios. In particular, Lee and Yu (2002) show that both moral hazard and basis risk drive down CAT bond prices substantially. As already discussed in chapter 1, a Monte Carlo approach that is based on some form of probability distribution, will have a significant possibility of generating conclusions which are likely to be neither robust nor stable. Which is where robust control Lyapunov analysis comes into its own in three main ways. First, through the ability to introduce closed-loop feedback into the modelling process allowing more complete, systematic and realistic capturing of the effects of uncertainty. Second, through the ability to capture non-linearities, such as convexity of the payoff function, as well as extreme behaviour inherent in the CAT bond process. Third and perhaps most importantly, by not placing any reliance on a particular form of probability distribution.

The model that is now developed in this section follows the spirit of Cox & Pedersen (2000) in so far as it combines primary financial market variables with catastrophic risk variables to produce a valuation

model first for CAT bonds and then for CAT options. However, the critical and defining difference and therefore the net new idea and work, is to extend the Cox & Pedersen approach by tackling the sources of uncertainty remaining within their model by incorporating stability and robustness into the valuation process in a coherent, elegant and efficient fashion. This is achieved by recasting their approach to the hedging problem using the concept of an interest rate model to span the otherwise incomplete market space as a state-space driven robust optimal control problem. This allows the valuation problem to be cast as a robust control Lyapunov problem that can be solved using F&K approach.

What alterations are required to translate the Cox & Pedersen approach to deal with CAT derivative pricing in a robust fashion ? Beginning with the simplest possible form of their model, assume that the CAT bond with a notional value of \$1 pays regular coupons of amount c and a final repayment of $\$1 + c$ at maturity time T, providing that a catastrophe does not occur. If a catastrophe occurs during a coupon period the bond makes a fractional payment of $f(1+c)$ then terminates. Under arbitrage conditions, financial economics asserts that in an arbitrage-free market there exists a probability measure, $\mathbb{Q}$, known as the risk-neutral measure such that the value at time 0 of an uncertain cash flow stream $\{c(k)\,|k = 1, 2, ..., T\}$ is given by the usual discounted expected value of the stream under the given probability measure

$$\mathsf{E}^{\mathbb{Q}}\left[\sum_{k=1}^{T} \frac{1}{[1+r(0)]\,[1+r(1)]\,...\,[1+r(k-1)]} c(k)\right] \tag{2.90}$$

where the process $r(k)$ is some single period interest rate process. If no catastrophic event occurs, then the above formula requires no further modification. However, if a catastrophic event occurs $c(k)$ can take on one of two possible values

$$c(k) = \left\{ \begin{array}{c} c1_{\{\tau>k\}} + f(1+c)\,1_{\{\tau=k\}} \\ k = 1, 2, ...T-1 \\ (c+1)\,1_{\{\tau>T\}} + f(1+c)\,1_{\{\tau=T\}} \\ k = T \end{array} \right\} \tag{2.91}$$

This formula will of course be simplified if only the coupons are at risk. Cox and Pedersen continue by assuming that a CAT bond can be traded in an arbitrage free market under a risk-neutral valuation measure $\mathbb{Q}$ with the timing of the occurrence of a catastrophe being independent of the interest rate environment under $\mathbb{Q}$. So, relating equations 2.90 and 2.91 together, gives the value at time 0 of the CAT

bond cash flow stream

$$c\sum_{k=1}^{T} P(k)\,\mathbb{Q}(\tau > k) + P(T)\,\mathbb{Q}(\tau > T) \tag{2.92}$$

$$+f(1+c)\sum_{k=1}^{T} P(k)\,\mathbb{Q}(\tau = k) \tag{2.93}$$

Note that the term $\mathbb{Q}(\tau > k)$ is the probability (under the risk-neutral measure) that the catastrophe does not occur in the first k periods. The cash flow,X, received by the bond holder given the occurrence of a catastrophe could be random, therefore requiring an adjustment to the model. If $Y(x)$ is the conditional severity distribution of the cash flow to the bond holder given a catastrophe, then using results from Tilley (1995 and 1997) then equation 2.92 becomes

$$c\sum_{k=1}^{T} P(k)\,\mathbb{Q}(\tau > k) + P(T)\,\mathbb{Q}(\tau > T) \tag{2.94}$$

$$+\sum_{k=1}^{T} P(k)\,\mathbb{Q}(\tau = k)\int_0^{\infty} x dY(x) \tag{2.95}$$

Note that generally the conditional severity distribution is embedded as part of the risk-neutral measure, $\mathbb{Q}$. Following Tilley, suppose that the catastrophe risk structure θ_0 represents the conditional probability under the risk neutral measure of no catastrophe occurring for a period. However, should a catastrophe occur, then there is assumed to be a single severity level that generates a payment of $f(1+c)$ at the end of the period in which the catastrophe occurs. If $\theta_1 = 1 - \theta_0$, then Cox and Pedersen simplifies to

$$c\sum_{k=1}^{T} P(k)\,(1-\theta_1)^k + P(T)\,(1-\theta_1)^T \tag{2.96}$$

$$+f(1+c)\sum_{k=1}^{T} P(k)\,\theta_1\,(1-\theta_1)^{k-1}$$

There are several important points here with respect to the treatment of uncertainty and robustness. First, in order to be able to apply Tilley's formula as in Cox and Pedersen it is necessary to know the conditional risk-neutral probability, θ_0, but θ_1 has not been linked to the probability of a catastrophe occurring, so that equation 2.96 is not closed. In order to close their model it is necessary to link equation 2.96 with observable quantities that can be used to estimate the parameters needed to apply the valuation model. Interpreting θ_1 as the empirical conditional probability of a catastrophe occurring, effectively means assuming away the uncertainty surrounding its measurement, such that assumptions can be made about the form of its conditional distribution and about the parameterisation of that distribution.

Second, as the model is incomplete, there does not exist a single value for the CAT bond. The best that can be achieved is to place upper and lower bounds on value. By embedding the catastrophe probabilities directly into the model by using them to span the state-space, a unique price can be generated. For example, if potential bond holders agree that q represents the probability of a catastrophe occurring and the the CAT bond price should be discounted using this risk adjusted expectation, then the discounted average (over both the catastrophe and non-catastrophe) cash flows (using an assumed fractional payment of f) is

$$(1+c)\left(fq+1.0\left(1-q\right)\right)\frac{1}{1+r} \tag{2.97}$$

Given that the probability distribution for the catastrophe and the assumption that values are discounted expected values over both risks, then unique prices result. The key problem with this approach is the necessity of either knowing or assuming the form and parameterisation of the conditional distribution of the catastrophes. Faced with the inherent uncertainty of the catastrophe modelling process, making such an assumption cannot be regarded as the ideal solution with respect to the likely robustness and/or stability of the final valuation result.

How then does this basic idea translate into to a robust optimal control Lyapunov approach ? To begin with, consider a simple representation of a system of catastrophic and financial variables, (complete with a number of simplifying assumptions, most of which will be successively relaxed) and assume there is some interest rate generating process, $F(x)$, where x is a vector of state variables describing the interest rate. Assume also that there exists a hedging strategy, $G(x)\,u$ that uses a non-catastrophe related zero coupon bond (whose value is determined by x) to hedge the CAT bond. The final element is a disturbance input, $H(x)\,w$, capable of capturing catastrophic shifts. It is assumed that u, the control or hedging strategy, is used to balance the hedging portfolio. These elements are linked together to form the following system

$$\dot{x} = F(x) + G(x)\,u + H(x)\,w \tag{2.98}$$

F, G and H are all assumed to be continuous functions. It is also assumed that the system is stabilisable and that the state is available for feedback - not an unreasonable assumption given that G is the hedging policy which will be dynamic and feedback into the model. The key assumption, however, is that a control Lyapunov function is known for this system. In other words, assume that a $\mathbf{C}^1$ (i.e. continuous in the first derivative), positive definite function of the form

$$V : \mathcal{X} \longrightarrow \mathbb{R}_+ \tag{2.99}$$

is known, such that

$$\inf_{u \in U} \nabla V(x) \cdot [F(x) + G(x)u] < -\alpha_V(x) \tag{2.100}$$

for all $x \neq 0$ and for some function α_V. The critical concept is to use the control Lyapunov function V as a robust control Lyapunov function for the uncertain system of equation 2.98. As F&K point out, this robust control Lyapunov function can be chosen independently of the uncertainty so that there is no knowledge of the structure of the disturbances, H. This means that for the CAT bond, it is possible to derive a robust control Lyapunov function without any knowledge of the structure of the catastrophic disturbance. This is a very strong feature of the model compared with the Cox and Pedersen approach where the assumption is made that there is a probability distribution for the catastrophic events and is one that ensures its robustness in the presence of uncertainty surrounding the likely arrival of catastrophic events.

From control theory it is known that to be a robust control Lyapunov function V must satisfy

$$\inf_{u \in U} \sup_{w \in B} \nabla V(x) \cdot [F(x) + G(x)u + H(x)w] < -\alpha_V(x) \tag{2.101}$$

for all $x \neq 0$.

Having provided an overview of the model, how do these optimal control theory concepts map into the Cox and Pedersen approach ? In order to be able to answer this question, it is first necessary to provide a precise definition of the variables in the model. As far as the financial market variables are concerned, Cox and Pedersen assume these to be modelled on the filtered probability space $\Omega^{(1)}, \wp^{(1)}, \mathrm{P}_1$, where $\Omega^{(1)}$ is taken to be finite such that it represents all paths that the financial variable can take over the time $k = 0, 1, ..., T$. However, the point of the robust approach is to move away from using a specific form of probability distribution to characterise the state space for the variables in the model. As Cox and Pedersen point out, their results also hold for infinite sample spaces, so the extension to a more general notion of a state-space seems intuitively acceptable. The key concept in making this transition for the purposes of robustness is the need to deal with the initial information state. It is known (e.g. Helton and James 1999) that careful choice of the initial state makes an enormous difference in the implementability of the controller or hedging process $G(x)$ and strongly affects the dynamic behaviour of the system.

Therefore, within the robust control Lyapunov approach, we will consider four finite dimensional Euclidean spaces: the state space (interest rate or financial variable such as the price of a discount bond) χ, the control or hedging space $\mathcal{U}$, the disturbance or catastrophe generating space $\mathcal{W}$ and the measurement space $\mathcal{Y}$. Given a continuous function $f : \chi \times \mathcal{U} \times \mathcal{W} \times \mathbb{R} \rightarrow \chi$, a differential equation can be

formed

$$\dot{x} = f(x, u, w, t) \tag{2.102}$$

where $x \in \chi$ is the state variable,$u \in \mathcal{U}$ is the control or hedging input, $w \in \mathcal{W}$ is the catastrophic disturbance input and $t \in \mathbb{R}$ is the time variable. Associated with the differential equation 2.102 are admissible measurements, admissible disturbances and admissible controls - with each being characterised by a set-valued constraint.

Taking the admissible measurements first, a measurement for equation 2.102 is a function $y : \chi \times \mathbb{R}$ such that $y(\cdot, t)$ is continuous for each fixed $t \in \mathbb{R}$ and $y(x, \cdot)$ is locally L_∞ for each fixed $x \in \chi$ (i.e. bounded on a neighbourhood of every point). Assuming a measurement constraint of the form $Y : \chi \times \mathbb{R} \rightsquigarrow \mathcal{Y}$, then a measurement $y(x, t)$ is deemed admissible when $y(x, t) \in Y(x, t)$ for all $(x, t) \in \chi \times \mathbb{R}$. The importance of this definition is that it allows for measurement uncertainty due to imperfections in the measurement process, perhaps because there may be several different measurement trajectories associated with a single state trajectory.

In equation 2.102, a disturbance is a function $w : \chi \times \mathcal{U} \times \mathbb{R} \rightarrow \mathcal{W}$, such that $w(\cdot, \cdot, t)$ is continuous for each fixed $t \in \mathbb{R}$ and $w(x, u, \cdot)$ is locally L_∞ for each fixed $(x, u) \in \chi \times \mathcal{U}$. Therefore, given a disturbance constraint $W : \chi \times \mathcal{U} \times \mathbb{R} \rightsquigarrow \mathcal{W}$, it is possible to state that a disturbance $w(x, u, t)$ is admissible when $w(x, u, t) \in W(x, u, t)$ for all $(x, u, t) \in \chi \times \mathcal{U} \times \mathbb{R}$. This is central to the modelling of the catastrophe space because admissible disturbances can include both exogenous disturbances such as catastrophes and feedback disturbances, such that they encompass a large class of memoryless model and input uncertainties and form part of the basis of the approach in yielding guaranteed stability framework for robust non-linear control.

In equation 2.102, a control is a function $u : \mathcal{Y} \times \mathbb{R} \rightarrow \mathcal{U}$ such that $u(\cdot, t)$ exhibits continuity for each fixed $t \in \mathbb{R}$ and $u(y, \cdot)$ is locally L_∞ for each fixed $y \in \mathcal{Y}$. Following the same approach, given a control constraint $U : \mathcal{Y} \times \mathbb{R} \rightsquigarrow \mathcal{U}$, it is possible to say that a control is admissible when $u(y, t) \in U(y, t)$ for $(y, t) \in \mathcal{Y} \times \mathbb{R}$ and that $u(y, t)$ is jointly continuous in (y, t). As F&K point out, it might be expected that a constant control constraint $U(y, t) \equiv U_0$ should be enough but for the purposes of our model there are valid and desirable reasons for allowing the constraint to depend on the measurement y. The most glaringly obvious example is that it might be desired not to hedge the the CAT bond using some possibly expensive strategy when the value of the CAT bond remains within an acceptably "normal" region.

The function f, taken with the set valued constraints U, W and Y, comprises a system

$$\sum = f(x, u(y(x, t))) \tag{2.103}$$

and a solution, $x(t)$, to this system solves the initial value problem

$$\dot{x} = f(x, u(y(x,t), t), w(x, u(y(x,t), t), t), t) \qquad x(t_0) = x_0 \tag{2.104}$$

given a measurement $y(x,t)$, a disturbance $w(x,u,t)$, a control $u(y,t)$ and an initial condition $(x_0, t_0) \in \chi \times \mathbb{R}$. Classical existence theorems from control theory guarantee that the right hand side of equation 2.104 is continuous in x and locally L_∞ in t, which means that solutions to $\sum$ always exist (locally in t) but need not necessarily be unique. It is also important to note that the above formulation can also include fixed order dynamics by re-defining the system $\sum$. For example, fixed order dynamics can be imposed by adding auxiliary variables to the state, control and measurement variables but the essential point that emerges from this problem statement is that solutions to $\sum$ are robustly, globally and asymptotically stable.

In order to derive a catastrophe derivative valuation framework based on this approach, it is necessary to take into account three particular issues. First, it must be remembered that for non-linear systems the feedback gain between inputs and outputs at each state depends on initial conditions. Second, the non-linearities inherent in the model, such as convexity of the payoff function, must be modelled as part of the initial conditions. Third, there must be an existing methodology for calculating the required quantities. Fortunately, robust control Lyapunov analysis satisfies all three demands and is the approach upon which the following analysis is constructed.

At its simplest, a control Lyapunov function for a system of the form $\dot{x} = f(x,u)$ is a $\mathbb{C}^1$ positive definite, radially bounded function $V(x)$ such that

$$x \neq 0 \tag{2.105}$$

$$\inf_{u \in U} \nabla V(x) \cdot f(x,u) < 0 \tag{2.106}$$

where U is a convex set of admissible values of the control variable, such that the derivative of the function can be made negative pointwise by the choice of control values. A function $V \in \mathcal{V}(\mathcal{X})$ is a robust control Lyapunov for a system Σ when there exist $c_v \in \mathbb{R}_+$ and $\alpha_v \in \mathcal{P}(\mathcal{X})$ such that

$$\inf_{u \in U(y,t)} \sup_{x \in Q(y,c,t)} \sup_{w \in W(u,t)} [L_f V(x,u,w,t) + \alpha_v(x,t)] < 0 \tag{2.107}$$

for all $y \in \mathcal{Y}$, all $t \in \mathbb{R}$ and all $c > c_v$; and where $L_f V$ is a Lyapunov derivative This formulation of the robust control Lyapunov function is important as it is generalisable in a number of directions such that it provides a significant degree of flexibility. Note that both control and disturbance inputs enter

the equation and that the definition copes with both measurement feedback and state feedback. This capability to deal with feedback is particularly valuable when devising valuation models as it means that more realistic hedging strategies can be represented. Notice also that the term c_v enables the modeler to cope with the three issues in stabilisability (in addition to asymptotic stabilisability) identified above.

Finding a function V that is a solution to Σ and is also robustly globally uniformly and asymptotically stable and which also converges to a residual and compact set $\Omega \in \mathcal{X}$ necessitates finding admissible controls known as pointwise min-norm control laws, which are so called because at each point x, their value is the unique element of $\mathcal{U}$ of a minimum norm that satisfies the control constraint $U(x)$ whilst also making the worst-case Lyapunov derivative at least as negative as $-\alpha_v(x)$. The good news from a computational perspective is that it is possible to compute the value of a pointwise min-norm control law at any point x by solving a convex, static minimisation programming problem that is completely determined by the data Σ, V and α_v. The further good news is that this static problem has a simple explicit solution in a wide variety of circumstances, a number of which are directly applicable to the CAT bond valuation problem. The only restriction is that the system must be jointly affine in u and w.

To see how this works in practice, take an example of a system $\dot{x} = f(x, u, w)$ for continuous functions f_0, f_1 and f_2

$$\dot{x} = f_0(x) + f_1(x)u + f_2(x)w \tag{2.108}$$

and suppose that V is a robust control Lyapunov function for this system such that $D : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$, then

$$D(x, u) := \max_{w \in W(x)} [L_f V(x, u, w) + \alpha_v(x)] \tag{2.109}$$

which, upon substituting, gives

$$\begin{aligned} D(x,u) &= \nabla V(x) \cdot f_0(x) + \nabla V(x) \cdot f_1(x) u && (2.110) \\ &+ \| \nabla V(x) \cdot f_2(x) \| + \alpha_v(x) && (2.111) \end{aligned}$$

Using the simplifications

$$\begin{aligned} \psi_0(x) &:= \nabla V(x) \cdot f_0(x) + \| \nabla V(x) \cdot f_2(x) \| + \alpha_v(x) && (2.112) \\ \psi_1(x) &= [\nabla V(x) \cdot f_1(x)]^T && (2.113) \end{aligned}$$

and defining $K : \mathcal{X} \rightsquigarrow \mathcal{U}$, gives

$$K(x) = \{u \in U : \psi_0(x) + \psi_1^T(x) u < 0\} \tag{2.114}$$

which finally gives the simplified expression for the pointwise min-norm control law

$$m(x) = \left\{ \begin{array}{ll} -\frac{\psi_0(x)\psi_1(x)}{\psi_1^T(x)\psi_1(x)} & \text{when } \psi_0(x) > 0 \\ 0 & \text{when } \psi_0(x) \leq 0 \end{array} \right\} \tag{2.115}$$

for all $x \in V^{-1}(c_v, \infty)$. This then is the expression that will enable the calculation of the robust control law that will provide both a hedging strategy in the face of both a catastrophe and disturbances to the underlying interest rate environment. The initial rigidity of joint affineness of the control and the disturbance can be relaxed through an integral back stepping procedure described by F&K, but it was found that such relaxation added relatively little to the robustness or stability of the results described in the next section.

2.5 Empirical analysis of pricing methodologies

This section provides a comparison and analysis of robust with non-robust derivatives pricing models. The "underlying" for the comparison is an index of catastrophic events. The relative performance of the robust and non-robust valuation models using this underlying is presented in terms of accuracy, stability and robustness of performance. The results were achieved by using four models to value both a catastrophe bond and a catastrophe option using the catastrophic loss index published by the Property Claims Service (PCS). In the case of both robust and non-robust models, the results were calibrated and satisfactorily cross-checked for reasonableness using a Monte Carlo simulation, for which the results are not reported.

2.5.1 Models and data sources

This section therefore takes the robust pricing framework developed in the preceeding sections of this chapter and applies it to the practical problem of producing robust valuations for both CAT bonds and CAT options where the underlying is an index of catastrophic events. The section presents and analyses two sets of empirical comparisons of a traditional non-robust model with three variants of the robust derivative pricing framework developed in the previous sections of this chapter. Section 2.5.2 analyses the comparative results of pricing CAT bonds using the standard Cox partial integro-differential equation approach (referred to as Cox-PIDE and which is standard in much of the insurance industry), with three versions of robust methods, namely, numerical HJI, linear robust control Lyapunov and non-linear robust control Lyapunov. Section 2.5.3 analyses the results of pricing CAT options using the standard translated gamma approximation and Monte Carlo simulation (to provide a cross check on the accuracy of the gamma model), with three versions of robust methods namely, numerical HJI, linear Lyapunov and non-linear

Lyapunov.

Calibrating the robust catastrophe bond and option models was carried out with respect to interest rate and catastrophe index data for the US market. The exact data used for all results reported in this chapter is detailed in Appendix 2. The following models were used in the comparative CAT bond computations:

1. Cox partial-integro differential equation (Cox-PIDE)

2. Numerical H∞ robust optimal control (i.e. numerical HJI)

3. Linear Lyapunov robust optimal control

4. Non-Linear Lyapunov robust optimal control

The following models were used in the comparative CAT option computations:

1. Cox-Compound Poisson using a translated gamma approximation (Cox-TGA).

2. Numerical H∞ robust optimal control (i.e. numerical HJI)

3. Linear Lyapunov robust optimal control

4. Non-Linear Lyapunov robust optimal control

The underlying data for both the CAT bond and the CAT option results was the catastrophe loss index produced by the PCS. After assigning a unique identifier to each catastrophe, the PCS provides daily data on the number and sizes of claims for each catastrophe for as long as the data continues to change. According to the PCS, the term "catastrophe" denotes a natural disaster that affects many insurers and where claims reach a certain threshold. Initially, the threshold was set at $1 million, but was subsequently increased first to $5 million and then to $25 million in 1997.

Although PCS catastrophe claim data is available from 1949 onwards, the data used in this thesis stops at end 2003. However, for the CAT bond research, only the daily data from January 1990 to December 2003 was used for the calculations reported in this thesis. The CAT option results are based on the daily data for the period that the PCS options were quoted on the COBT, namely, September 1995 to August 1999. This means that the PCS options data covers just over 500[5] exchange traded CAT options and as

[5]CBOT provide PCS CAT option data data on 474 call spread option contracts, 20 pure put option contracts and 36 pure call option contracts. During the period that PCS CAT options were traded on CBOT, the PCS index was updated in the following way. When a catastrophe has been identified, the PCS published a first estimate within 48-72 hours. The PCS then continued to refine the estimate of the losses. The PCS index is the accumulated loss estimates for all identified catastrophes within the defined reion for the duration of the occurrence period.

it focuses on US based claims with respect to US catastrophes, the data is denominated in US dollars. A more detailed description of the practical operation of the PCS index and the contract definition details required for the calculation of PCS option values is provided in Appendix 2.

Prior to using the CAT bond and CAT option models, the following two stage calibration process was carried out:

1. Fitting a loss distribution to the raw PCS claims data supplied. The heavy tailed Burr distribution was chosen for reasons explained in Appendix 2, where the details and results of the calibration process are also described.

2. Testing the above four models using Monte Carlo simulation (using 10,000 trials in each case for consistency) to ensure consistent valuation under boundary conditions.

2.5.2 CAT bond valuation results

CAT bonds[6] made their appearance in the mid-1990's, when a market in catastrophe insurance risk emerged in order to facilitate the direct transfer of reinsurance risk associated with natural catastrophes from corporations, insurers and re-insurers to capital market investors. The primary instrument developed to satisfy this need was the CAT bond. The distinguishing feature of CAT bonds is that the ultimate repayment of principal depends on the outcome of an insured, naturally occurring catastrophic event, such as earthquakes and hurricanes, both of which are beginning to have a dominating impact on the insurance industry.

This impact is partially due to the rapidly changing and heterogeneous distribution of high-value property in many vulnerable areas of the USA. A consequence of this has been an increased need for a primary and secondary market in catastrophe related insurance derivatives. The creation of CAT bonds, along with allied financial products such as catastrophe insurance options, was motivated in part by the need to cover the massive property insurance industry payouts of the early and mid-1990's. They also represent a "new asset class" in that they provide a mechanism for hedging against natural disasters - a risk which has been shown to be substantially uncorrelated with the capital market indices (Doherty, 1997). Subsequent to the development of the CAT bond, the class of disaster referenced has grown considerably. As yet, there is almost no secondary market for CAT bonds which hampers using arbitrage-free pricing

[6]Catastrophe (CAT) bonds are an example of a class of securities known as risk-linked securities. The class includes securities such as share quota transactions, life insurance securities, catastrophe options and other insurance-related financial instruments. The results reported in this section focus on just CAT bonds, which are privately placed securities that are sold to qualified institutional investors as defined under the Security and Exchange Commission Rule 144A. Generally speaking, a qualified institutional investor, as defined under Rule 144A, owns and invests on a discretionary basis at least $100 million in securities that are not affiliated with the investor.

models for the derivative and has also lead to the emergence and persistence of models that are frequently neither robust nor stable.

The basic CAT bond structure can be summarized as follows:

- The sponsor of the CAT bond establishes a special purpose vehicle (SPV) as an issuer of bonds and as a source of reinsurance protection.
- The issuer sells bonds to investors.
- The proceeds from the sale of the CAT bonds are invested in a collateral account.
- The sponsor pays a premium to the issuer; this and the investment of bond proceeds are a source of interest paid to investors.
- CAT bonds are usually zero coupon and therefore generate no cashflows prior to maturity.[7]

If the specified catastrophic risk is deemed to have been triggered, the funds are withdrawn from the collateral account and paid to the sponsor; at maturity, the remaining principal - or if there is no event, 100% of principal - is paid to investors. There are three types of triggers: indemnity, index and parametric. An indemnity trigger involves the actual losses of the bond-issuing insurer. For example, the event may be the insurer's losses from an earthquake in a certain area of a given country over the period of the bond. An index trigger involves, in the US for example, an index created from property claim service (PCS) loss estimates. A parametric trigger is based on, for example, the Richter scale readings of the magnitude of an earthquake at specified data stations, or the Saffir/Simpson classification scale in the case of hurricanes. The results reported in this section of the current chapter only address the issue of pricing CAT bonds that feature index triggers.

Property insurance claims of approximately USD 60 billion between 1990 and 1996 (Canter, Cole, and Sandor; 1996) caused great concern to the insurance industry and resulted in the insolvency of a number of firms. These bankruptcies were brought on in the wake of hurricanes Andrew (Florida and Louisiana affected, 1992), Opal (Florida and Alabama, 1995) and Fran (North Carolina, 1996), which caused combined damage totalling USD 19.7 billion (Canter, Cole, and Sandor; 1996). These, along with the Northridge earthquake (1994) and similar disasters, led to an interest in alternative means for underwriting insurance. In 1995, when the CAT bond market was born, the primary and secondary (or reinsurance) industries had access to approximately USD 240 billion in capital (Canter, Cole, and Sandor;

[7]Note that all CAT bonds have been set up as zero coupon. A zero coupon structure was used for two reasons. First, to provide comparability with other research results. Second, to simplify the modelling and make it easier to identify the impact of catastrophic events on dimunition in redemption value.

1996; Cummins and Danzon; 1997). Given the capital level constraints necessary for the reinsuring of property losses and the potential for single-event losses in excess of USD 100 billion, this was clearly insufficient.

The international capital markets provided a potential source of risk appetite for the reinsurance market. An estimated capitalisation of the international financial markets, at that time, of about USD 19 trillion underwent an average daily fluctuation of approximately 70 basis points or USD 133 billion (Sigma; 1996). The under-capitalisation of the reinsurance industry (and the consequent potential default risk) meant that there was a tendency for CAT reinsurance prices to be highly volatile. This was reflected in the traditional insurance market, with rates on line being significantly higher in the years following catastrophes and dropping off in the intervening years (Sigma; 1997; Froot and O'Connell; 1997). This heterogeneity in pricing has had a very strong damping effect, forcing many re-insurers to leave the market, which in turn has adverse consequences for the primary insurers. A number of reasons for this volatility have been advanced (Cummins and Danzon; 1997; Winter; 1994).

Some of the traditional assumptions of derivative security pricing are not correct when applied to these instruments due to the properties of the underlying contingent stochastic processes. There is evidence that certain catastrophic natural events have (partial) power-law distributions associated with their loss statistics (Barton and Nishenko; 1994), which if true, would overturns the traditional log-normal assumption of derivative pricing models. There are also well-known statistical difficulties associated with the moments of power-law distributions[8], thus rendering it impossible to employ traditional pooling methods and consequently the central limit theorem. Given that heavy-tailed or large deviation results assume, in general, that at least the first moment of the distribution exists, there will be difficulties with applying extreme value theory to this problem (Embrechts, Resnick, and Samorodnitsky; 1999). It would seem that these characteristics may render traditional actuarial or derivatives pricing approaches ineffective.

There are additional features to modelling the CAT bond price which are not to be found in models of ordinary corporate or government issue (although there is some similarity with pricing defaultable bonds). In particular, the trigger event underlying CAT bond pricing is dependent on both the frequency and severity of natural disasters. In the model described here, we attempt to reduce to a minimum any assumptions about the underlying distribution functions. This is in the interests of generality of application. The numerical examples will have to make some distributional assumptions and will reference some real data. Given the daily availability of PCS loss data, it is also appears to be reasonable to

[8]This has become a significant research topic in its own right - see for example, "Multifractal Power Law Distributions: Negative and Critical Dimensions and Other "Anomalies, Explained by a Simple Example", by Benoit B. Mandelbrot, in the Journal of Statistical Physics, Vol. 110, Nos. 3–6, March 2003.

assume that loss levels are instantaneously measurable and updatable. It is straightforward to adjust the underlying process to accommodate a development period.

There is a natural similarity between the pricing of catastrophe bonds and the pricing of defaultable bonds. Defaultable bonds, by definition, must contain within their pricing model a mechanism that accounts for the potential (partial or complete) loss of their principal value. Defaultable bonds yield higher returns, in part, because of this potential defaultability. Similarly, CAT bonds are offered at high yields because of the unpredictable nature of the catastrophe process. With this characteristic in mind, a number of pricing models for defaultable bonds have been advanced (e.g. Jarrow and Turnbull, 1995, Duffie and Singleton, 1999, Zhou and 1997). The trigger event for the default process has similar statistical characteristics to that of the equivalent catastrophic event pertaining to CAT bonds.[9]

With this in mind, the Cox model has been used as the benchmark for non-robust models of catastrophic processes. The underlying assumption is that there is a Poisson point process (of some intensity, in general varying over time) of potentially catastrophic events. However, these events may or may not result in economic losses. It is assumed that the economic losses associated with each of the potentially catastrophic events is independent and has a certain common probability distribution. This is justifiable for the Property Claim Loss indices used as the triggers for the CAT bonds. Within this model, the threshold time can be seen as a point of a Poisson point process with a stochastic intensity depending on the instantaneous index position.

Having defined the data and the models used in the previous section, this section now presents and analyses the results of the computations. The basic zero coupon CAT bond analysed in this section has the following structure. It is assumed to pay an amount, Z, at maturity, T, contingent upon a threshold time $\tau > T$. The no arbitrage, present value (discounted at a continuously compounded rate of R) of the zero coupon CAT bond associated with a threshold loss level, D, catastrophic flow, M, an aggregate loss process, L and a distribution of incurred losses, F, that pays Z at maturity, is given by

$$V_t^1 = E\left[Z \exp\left\{-R(t,T)\right\}(1 - N_t)\,|\mathcal{F}_t\right]$$

where $\tau = \inf(t : L_t \geqslant D)$ and $N_t = I(L_t \geqslant D)$[10]. It is assumed that the threshold event is the time at which the accumulated losses exceed the threshold level, D, i.e. $\tau = \inf(t : L_t \geqslant D)$. To simplify the

[9] In an allied application to mortgage insurance, the similarity between catastrophe and default in the log-normal context has been commented on (Kau and Keenan; 1996).

[10] Baryishnikov et al (1998) also show that this is a doubly stochastic Poisson process with intensity

$$\lambda_s = m_s\left\{1 - F(D - L_s)\right\} I(L_s < D)$$

computations, the zero coupon CAT bond valuations reported are assumed to redeem at par of \$100 at maturity if aggregate losses do not exceed the threshold level, D. However, unlike Burnecki, Kukla and Taylor (2001) it is assumed that if accumulated losses exceed the threshold, then the bond holder receives a recovery amount, B_t, calculated as

$$B = Z - N_t$$

where N_t is stated as percentage of Z. The following results are therefore reported for varying levels of accumulated catastrophic losses incurred prior to maturity, expressed as a percentage. To achieve this scaling, the PCS loss data was simply re-based to 100 at the beginning of the calculation period.

Figure 2.3 therefore illustrates the behaviour of the price of a series of such zero coupon CAT bonds (assumed to have been issued at a discount with accretion to par and with price expressed as a percentage of par) for given combinations of time to maturity and percentage loss. The valuations were produced using the Cox-PIDE model. The CAT bond valuations are at increasing monthly maturities from 1 month out to 12 months (e.g. 1m maturity, 2m maturity, 3m maturity etc), all with identical issue date of 01 August 1992. The results in figure 2.3 are for CAT bonds based on PCS loss data for the 12 months beginning 01 August 1992 for the National index, which was chosen specifically because it contained the largest (at that time) and most costly world insurance loss in the form of hurricane Andrew which occurred on 23 August 1992 and produced total insured losses of \$15.5bn in 1992 dollar terms (or \$20.8bn in 2004 dollar terms).

Figure 2.3 presents a number of interesting features. First, it is clear that increasing the threshold loss level increases the value of the CAT bond. This behaviour makes intuitive sense, since raising the threshold loss level means less likelihood of losses consuming the entire value of the bond. For comparative and sanity check purposes, it is encouraging to note that the profile shown in figure 2.3 is consistent with results produced by both Burnecki, Kukla and Taylor (2003) and Baryshnikov, Mayo and Taylor (2001) from similar studies of CAT bond pricing using PCS data.

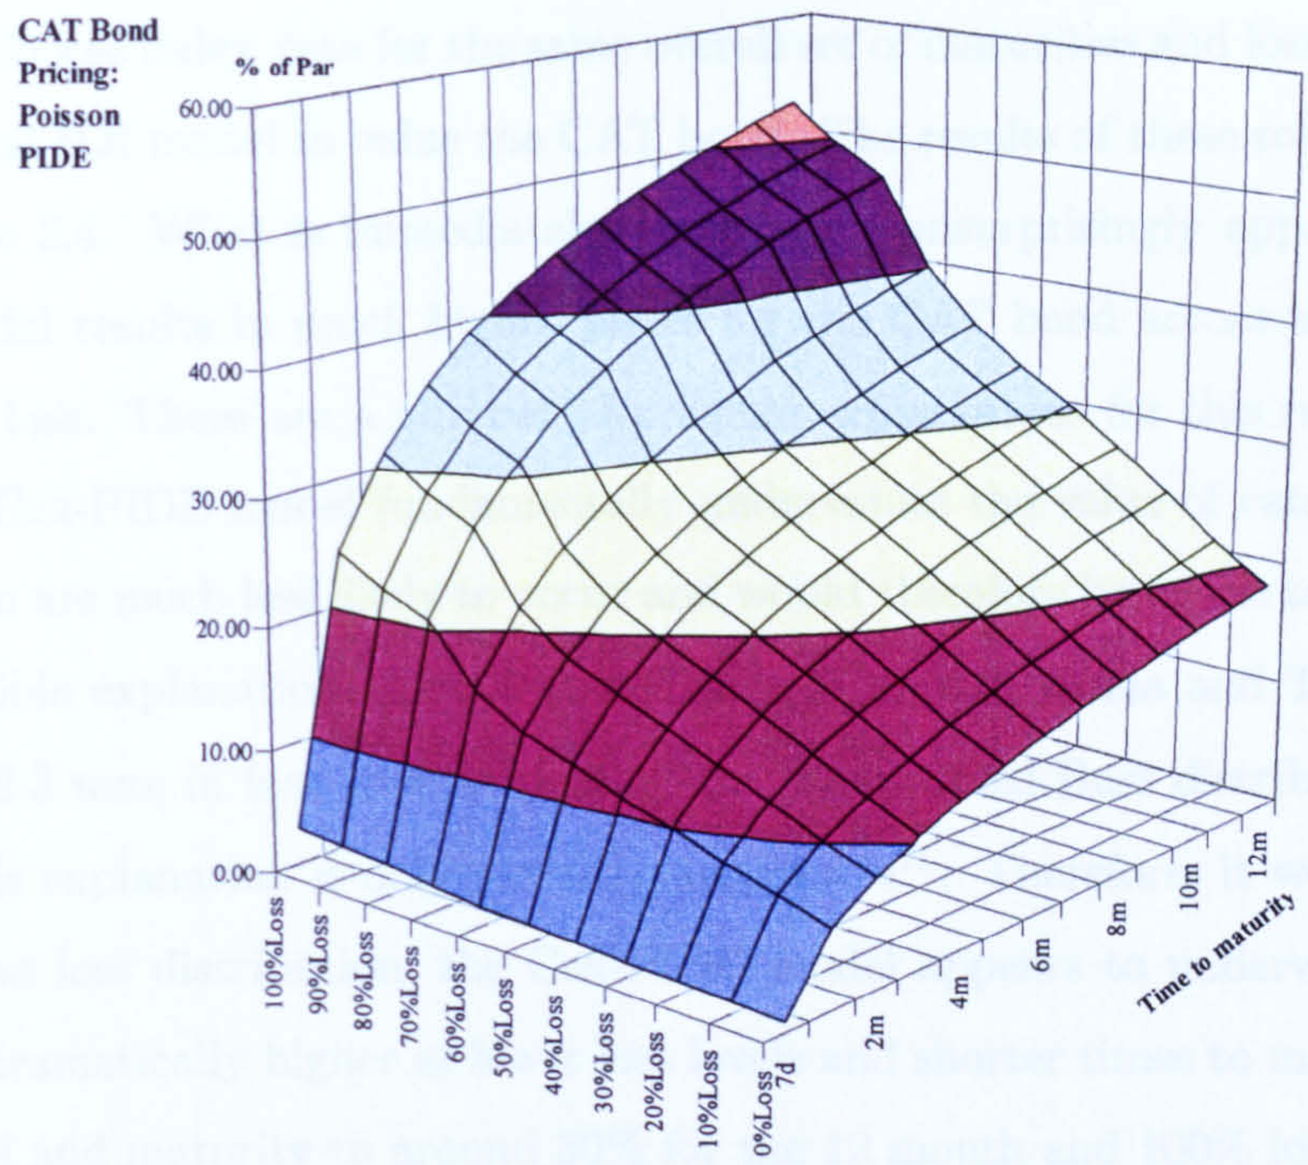

Figure 2.3: Short-term CAT bond valuation using Cox-PIDE

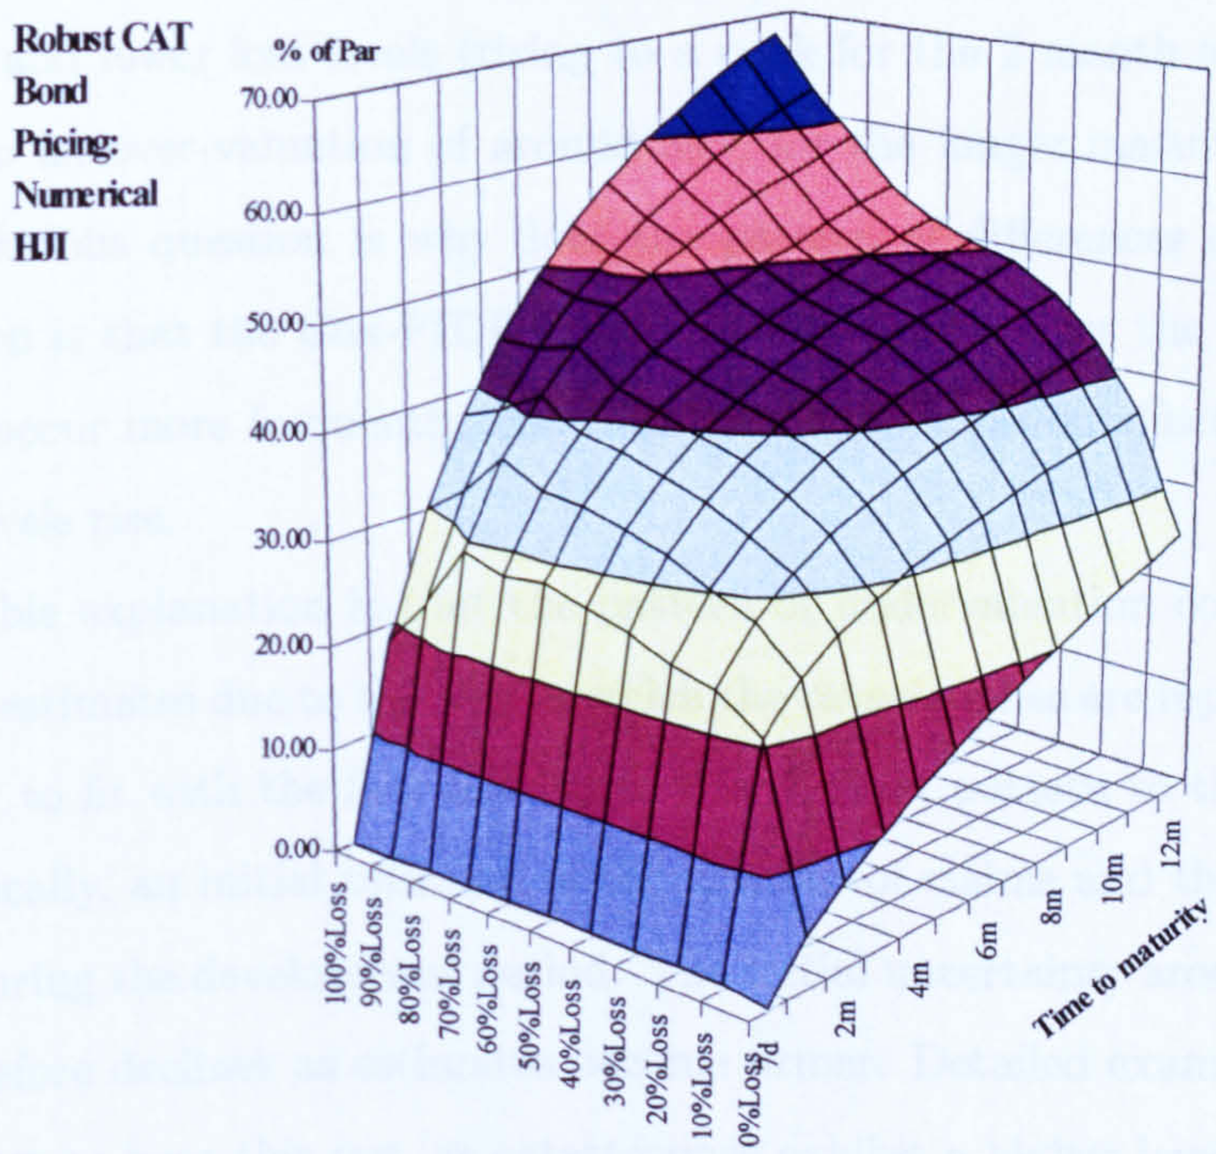

Figure 2.4: Short-term CAT bond valuation using numerical HJI

Now compare the results in figure 2.3 with the results of performing the same valuation exercise using

the same period PCS loss index data for the same overall set of maturities and loss levels, but using instead the robust numerical HJI model to value the CAT bond. The results of these robust model computations are shown in figure 2.4. What is immediately, though not unsurprisingly apparent, is that the robust numerical HJI model results in much higher prices for the CAT bond across almost all maturities and levels of threshold loss. There are a number of possible explanations for this result. First and trivially, could be that the Cox-PIDE model fundamentally undervalues the value of catastrophic events, that by their very definition are much less likely to occur and would therefore be in the tail of a distribution. This is clearly one possible explanation. However, following Burnecki, Kukla and Taylor (2003), the results reported in figure 2.3 were in fact generated using the heavy-tailed Burr distribution to fit the PCS loss distribution, so this explanation is only partially acceptable[11]. Therefore, it would appear that despite using a heavy tailed loss distribution, the Cox-PIDE model appears to undervalue the CAT bond; the undervaluation is dramatically higher at lower loss levels and shorter times to maturity and then declines with both loss level and maturity to around 30% for the 12 month and 100% loss level combination.

What is also interesting is the pattern of differences between the two approaches. Figure 2.5 shows the percentage undervaluation between Cox-PIDE and numerical HJI. The pattern of differences are worthy of note. First, it should be born in mind that the results in figures 2.3 and 2.4 are limited to bonds with maturities ranging from 1 month to 1 year. The undervaluation appears much greater for the shorter maturities and lower loss levels (rising to a peak for the 2 month and 30% threshold loss level), before tailing-off to an over-valuation of around 30% for the longer maturities and higher loss levels[12]. An immediately obvious question is why does the pattern of differences change so significantly? One possible explanation is that the Cox-PIDE fundamentally undervalues the impact of smaller, individual catastrophes that occur more frequently, but that the undervaluation effect is eroded with the effects of time and as loss levels rise.

A second possible explanation is that the pattern of undervaluation could also be a function of the volatility of claims estimates due to the way in which the catastrophes are reported and the index adjusted. This would appear to fit with the fact that there is a distinct pattern to the way in which catastrophes are reported. Typically, an initial estimate of the number of claims and their total value gets published and then refined during the development period. The initial uncertainty around the losses associated with a catastrophe therefore declines as estimates become firmer. Detailed examination of data for individual catastrophes appears to bear this out, as catastrophes exhibit a higher level of volatility in initial claims estimates (due to the initial lack of hard facts as losses take time to assess and claims then take further

[11] See Appendix 2 for details of the fitting procedure and some of the results of the calibration.

[12] Note that all of the results used to plot the figures contained in this chapter are provided in numerical form in Appendix 3).

time to complete and submit), which then declines as facts emerge, loss estimates crystalise and claim numbers and values stabilise.

A third possible explanation is that the mis-valuation could be attributable to seasonal effects. There is a well known and heavily documented[13] pattern of hurricane and tropical storm occurrence in the southern USA during the late summer and autumn. This weather pattern accounts for a significant proportion of the insured losses that form the PCS index. Even though the heavy-tailed Burr distribution was used to fit the PCS data, it may be the case the Burr distribution is simply not capable of adequately capturing the effects of the well known pattern to the occurrence of hurricanes and storms in the south eastern USA during this period. The spike in catastrophes is clearly evident in the steep slope in both figures 2.3 and 2.4, which both clearly show the dramatic impact of hurricane Andrew in August and September 1992. Figure 2.5 shows that the Cox-PIDE model significantly undervalues the impact on the CAT bond price of this catastrophe for short dated maturities when the number and size of claims is still at its most volatile, indicating a lack of robustness with respect to the occurrence of large catastrophic events.

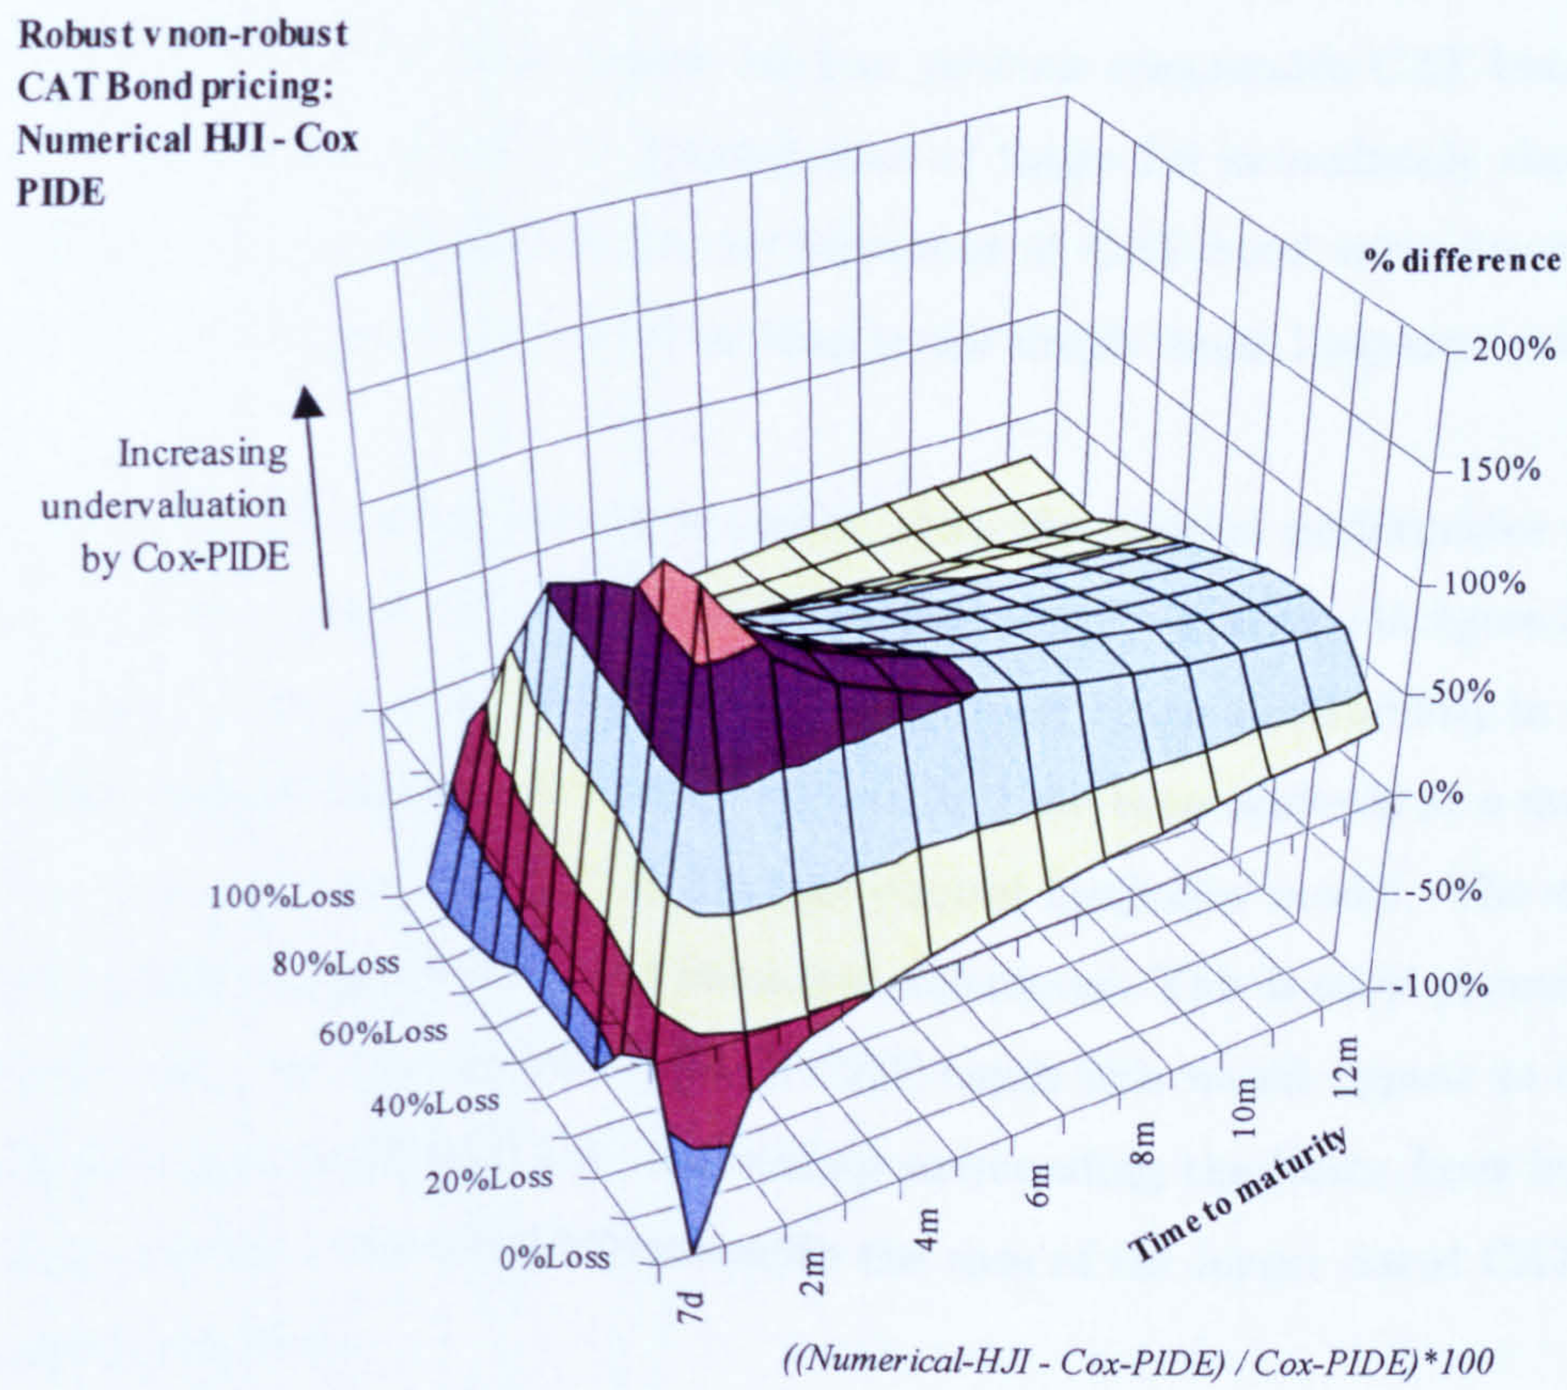

Figure 2.5: Short-term CAT bond model valuation: Numerical HJI - Cox-PIDE

[13] See for example the Insurance Information Institute website at: www2.iii.org/facts.

An immediately interesting question is whether or not the cost of robustness varies according to the robust model approach employed. Therefore, to investigate whether the pattern exhibited in the results of the numerical HJI model was mainly a function of a particular facet of the HJI model, a linear robust control Lyapunov model was used to value CAT bonds using the same range of maturities, the same underlying PCS loss data and the same range of loss levels as used for the Cox-PIDE and numerical HJI models. Figure 2.6 provides the results of these CAT bond valuations for the linear robust control Lyapunov valuation model developed earlier in this chapter. What is immediately clear is that all three models display relatively smooth monotonic valuation profiles throughout the ranges of loss levels and time to maturity. This is to some extent to be expected as all three models contain significant linearisations. The numerical HJI model used in this analysis is, in particular, an explicitly linear model - which, as already explained earlier in this chapter means that significant control effort (and therefore associated higher control cost, which is in turn reflected in a higher valuation) can be wasted attempting to combat inherent non-linearities. The final logical step is therefore to extend the analysis by including the non-linear version of the robust control Lyapunov model in its piece-wise min-norm form. Once again, the same maturities, loss levels and PCS data were used to produce comparable CAT bond valuations, the results of which are reported in figure 2.8. Examination of figure 2.8 immediately shows the benefit of explicitly incorporating the non-linearities, as the behaviour of CAT bond value around low loss levels and short time to maturity is now much smoother than in the simple linear Lyapunov case shown in figure 2.6.

Worthy of note is the interesting behaviour arises when the relative performance of the non-linear Lyapunov model is compared with the non-robust Cox-PIDE model - as shown in figure 2.9. Two features are worthy of comment. First, is that there is now a much lower range of variability in valuation around the Cox-PIDE model, suggesting that robustness appears to have been achieved at a much lower cost by explicitly incorporating non-linearities into the robust control Lyapunov model. The second interesting feature is the pronounced double peakiness in valuation differences. This is most pronounced around the 50% and 80% loss levels. For the very short dated CAT bonds this would appear to coincide with the re-estimation volatility associated with the uncertainty surrounding the losses from hurricane Andrew. The fact that this behaviour is far less pronounced in the case of the longer dated CAT bonds seems to lend support to such conjecture.

The contrast in performance between the three robust models is clearly seen when comparing the results in figure 2.5 (numerical HJI v Cox-PIDE), figure 2.7 (linear Lyapunov v Cox-PIDE) and figure 2.9 (non-linear Lyapunov v Cox-PIDE). The first point to note is that the numerical HJI model exhibits a

much smoother difference profile with respect to the Cox-PIDE model compared with the linear Lyapunov model, but at a higher cost than the non-linear Lyapunov model. Figure 2.7 shows an interesting pattern of differences which is most marked around the lower loss levels and shorter times to maturity. The numerical HJI model imposes a much greater cost penalty to robustness than the linear Lyapunov model, which is reflected in much higher valuations. This pattern is particularly pronounced in the short-term and low loss cases which in the context of the current analysis are precisely those CAT bonds most subject to the impact of hurricane Andrew. The pattern of differences then falls away, becoming far less significant in the case of increased time to maturity and higher loss levels.

The final step in this CAT bond research was to analyse the importance of time to maturity in determining the cost of robustness. Focus has so far been limited to short dated CAT bonds, but the critical question is whether examining bonds with a maximum maturity of only 12 months is likely to exacerbate or hide any valuation patterns that may be associated with achieving robustness. On the one hand, identifying the cost of robustness may be argued to be a simpler task by concentrating on short-dated CAT bonds. Unfortunately, on the other hand, little can be inferred about the dynamics of robustness over time if attention is restricted to such a short space of time. The next logical step is therefore to extend the maturity of the CAT bonds for all four models. Accordingly, all four CAT bond models were therefore re-run using the same loss levels, but using instead 10 years worth of PCS data beginning 01 January 1990 and ending 31 December 1999. The time to maturity of the longest CAT bond was extended to 10 years at 6 monthly intervals, i.e. 6m, 12,m, 18m,...,108m, 114m, 120m. In other words, CAT bonds with 20 different maturities ranging from 6 months to 10 years (but all with an identical start date of 01 January 1990) were valued using each of the four models. The results of each set of valuations is reported in figures 2.10 (Cox-PIDE), 2.11 (numerical HJI), 2.12 (linear Lyapunov) and 2.13 (non-linear Lyapunov), with comparisons to Cox-PIDE being presented in figures 2.14 (numerical HJI v Cox-PIDE), 2.15 (linear Lyapunov v Cox-PIDE) and 2.16 (non-linear Lyapunov v Cox-PIDE).

The results for these longer maturity CAT bonds provide a number of further insights into the robustness and stability of the three robust valuation models. The first and most obvious feature to emerge from the long-dated CAT bond valuations is the fundamentally different shapes of the valuation surfaces when compared with those generated for the short dated CAT bonds. The most interesting set of results is for the numerical HJI valued bonds shown in figure 2.11, which exhibit an extremely high implied cost of robustness as can be vividly seen in figure 2.14. Detailed examination of the results revealed that the principal reason for this behaviour was that for the longer dated bonds severe cost penalties were being incurred by the numerical HJI algorithm in order to ensure stable solutions. These cost penalties translated directly into higher valuations as the HJI model consumed increasing numbers of processing

cycles searching for a stable solution to satisfy the robustness and stability constraints.

The second feature of interest is the presence of a much clearer valuation differential between the numerical HJI model on the one hand and the two robust control Lyapunov models on the other hand. The numerical HJI algorithm used in the computations follows the standard power series approach of Al'brecht (1961) for solving infinite time optimal control problems[14]. Why does this differential occur and how should it be interpreted? To answer these questions, consider that in contrast to the numerical HI algorithm, the two Lyapunov based models both use the Freeman and Kokotovic approach of finding a meaningful cost function such that the given robust control Lyapunov function is the corresponding value function. This implicitly provides a solution to the equivalent linear HJI equation, thereby enabling the direct computation of the robust optimal control law. Therefore, providing that the cost function belongs to a meaningful class of cost functions, the resulting control law is robust and guaranteed to inherit all the required optimality properties. The robust control Lyapunov approach uses an inverse optimal robust stabilisation problem of finding a meaningful cost function, such that a given robust control Lyapunov function is the corresponding value function. This results in a solution to the equivalent linear HJI problem that is both stable *and* robust. In the case of the non-linear robust optimal control Lyapunov model, this further translates into solutions that take advantage of the non-linearities in the valuation problem that exist explicitly because the catastrophic events are driven by highly complex non-linear relationships. The outcome in the case of the non-linear robust control Lyapunov model is smoother and less expensive robustness - in other words, the cost of robustness is lower in the non-linear case as the solution takes advantage of the non-linearities rather than fighting against them, which translates directly into lower robustness costs.

The final issue to consider is whether the relative performance of the three robust models presented so far provide sufficient information to draw definitive conclusions about the cost of robust valuation in the face of massive catastrophic events such as hurricane Andrew? Looking first at the short-term valuation results, all three robust models can be seen to exhibit substantial differences compared to the Cox-PIDE model for the 1-6 month securities and up to around the 50% loss level. The pattern of differences then appears to be less pronounced for the 6-12 month securities and higher loss levels. This may at least in part be due to the volatility of the claims estimates referred to above being handled differently in the cost function within the models. the liner and non-linear Lyapunov approaches

As far as the long-term valuation results are concerned, the impact and tail effects of hurricane Andrew

[14] Al'brecht's (slightly modified) method solves the HJI partial differential equation in the neighbourhood of the origin using a power series method. This ultimately reduces the quadratic terms of the HJI pde to an easily solvable Riccati equation and a linear optimal feedback rule. This is then solved using function SB02PD ported from the Slicot library as explained in Appendix 1. This function solves the continuous algebraic Riccati equations using the matrix sign function method with condition and forward error bound estimates.

can be seen quite clearly between the four models. Once again, the numerical HJI model exhibits extreme cost penalties right across the loss level spectrum for the shorter maturities and in the middle of the loss threshold range, but these penalties fall off very rapidly as the initial impact effect of Andrew decays. The cost of robustness for the numerical HJI is therefore far higher even for the very longest bonds at all but the very lowest loss levels. Detailed examination of the diagnostics for the numerical HJI once again reveals the same explanation as in the shorter dated case. What is also worthy of note is that the pattern of extreme cost penalties appears to have quite a lengthy tail to its decay structure. The tail has three discernible phases, which can best be observed by looking at the longest dated bonds. The first phase covers the initial impact of Andrew and appears to last out to around 3 years. During this phase the cost penalties begin extremely high, but fall off very rapidly. The second phase is from around 3 to 5 years, during which time the cost penalties continue to fall but at a much slower pace. The final phase, from 5 to 10 years sees the cost penalties flattening out, but still remaining high.

In contrast, the two Lyapunov models no longer continue to attract extremely high cost penalties compared with the Cox-PIDE model as can be clearly seen when comparing figures 2.14, 2.15 and 2.16. What is even more interesting is that the two Lyapunov models actually exhibit valuations below the Cox-PIDE for some combinations of shorter maturities and higher loss levels. Closer examination of the diagnostics for the linear Lyapunov model revealed that incorporation and consequent influence of feedback yielded smoother solutions so that the cost penalties associated with achieving robustness were substantially reduced. In the case of the non-linear Lyapunov model the extra influence of the lower cost penalties associated with incorporating the non-linear dynamics further reduced the costs of robustness. This finding for the non-linear robust control Lyapunov function is a significant finding as it underlines that the costs of robustness to uncertainty may not be so high as to make robust strategies unaffordable. The point is that in times when catastrophic events do not occur - which by their very definition tends to be most of the time - the costs of robustness make it totally uneconomic as a valuation methodology. However, when catastrophic events are brought into the picture, the costs of robustness become far more acceptable compared the possible levels of loss, which may include bankruptcy or ruin at the limit.

Notwithstanding the above possible explanations for the valuation differences between the Cox-PIDE and the numerical HJI models, arguably the more interesting question is whether the robust model actually *overvalues* the benefits of robustness. One way of answering this question is by resorting to a comparison of the CAT bond valuations with the out-turn in the PCS index[15]. The answer to this question can be only partially inferred from the results presented in figures 2.3, 2.4 and 2.5. Hurricane Andrew occurred on

[15] An equally valuable cross-check would be to compare the results of the CAT bond models with the traded prices of puts and calls in the CBOT options prices. However, the option contracts did not trade for the entire period of interest of the short-dated CAT bonds. See Appendix 2 for details of the periods covered by the PCS options data available from CBOT.

23rd August 1992, so the first month in which estimates of likely loss were fully available was September 1992. The relative differences in value before and after the impact of hurricane Andrew can be seen in figure 2.5.

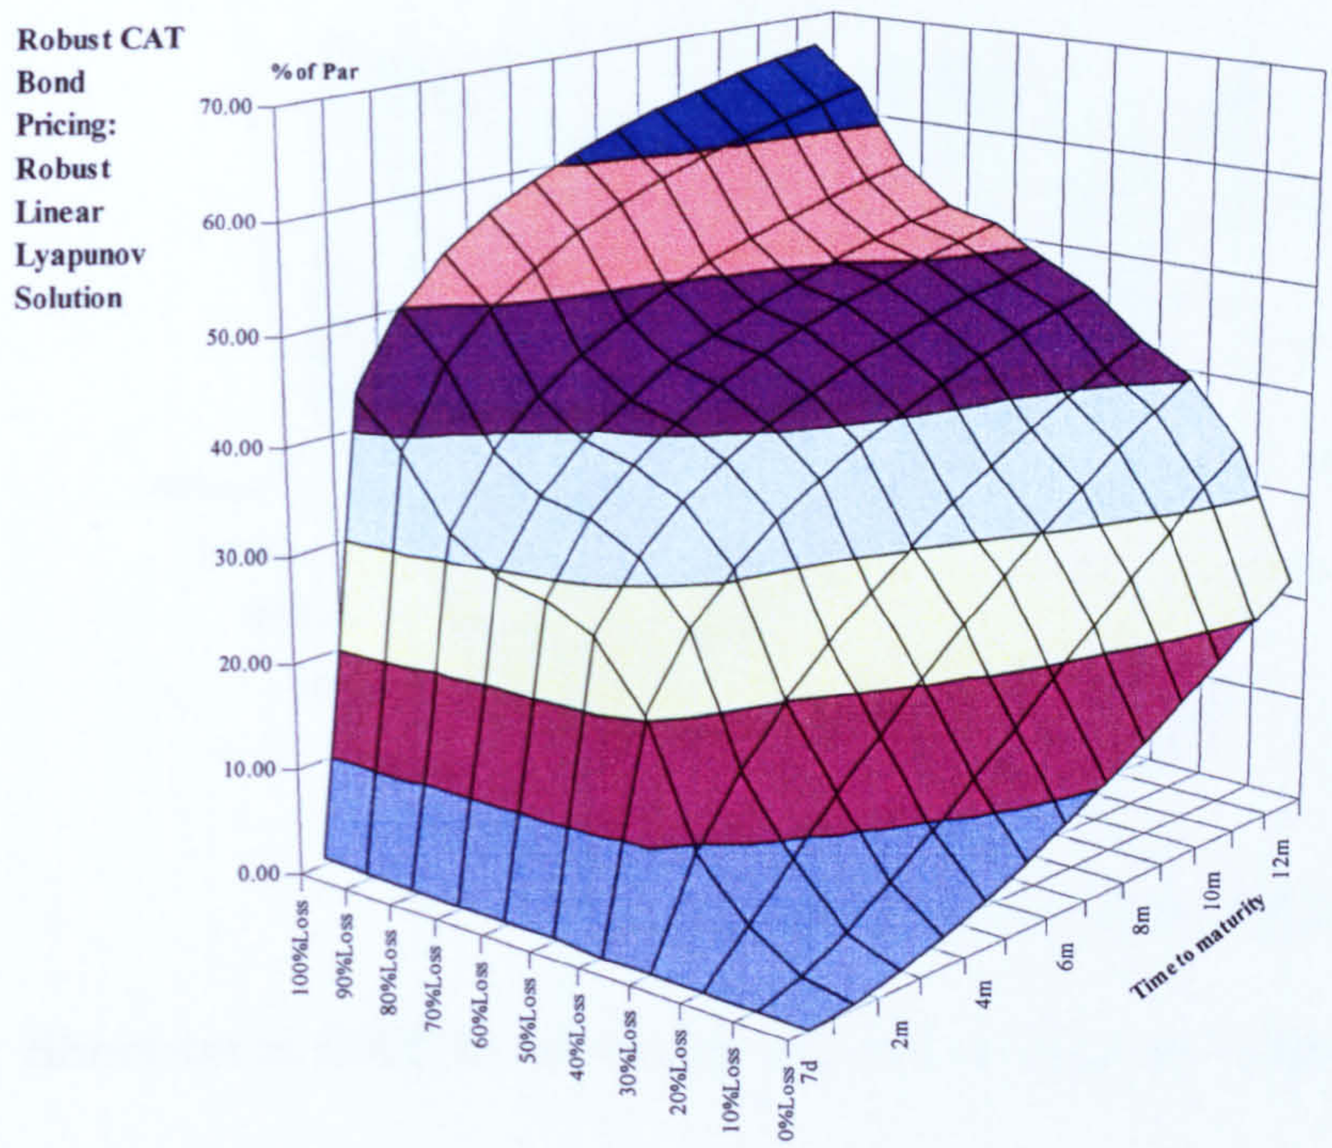

Figure 2.6: Short-term CAT bond valuation using Linear Lyapunov

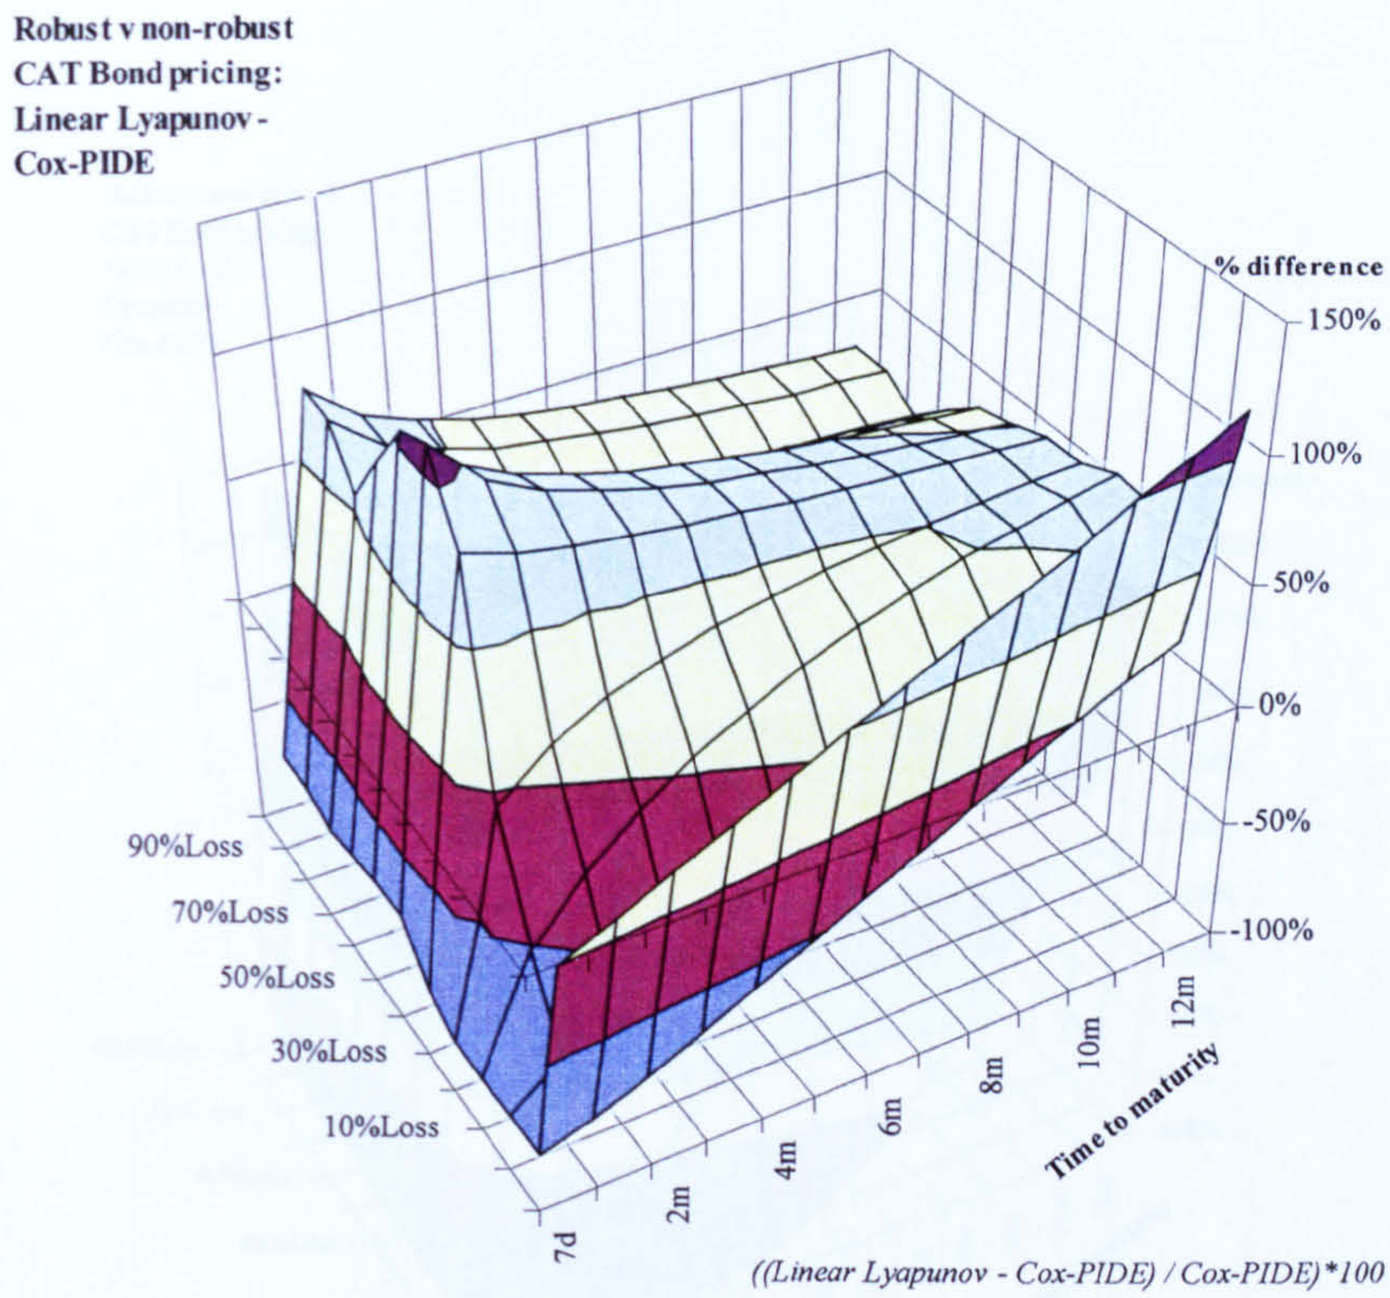

Figure 2.7: Short-term CAT bond model valuation: Linear Lyapunov - Cox-PIDE

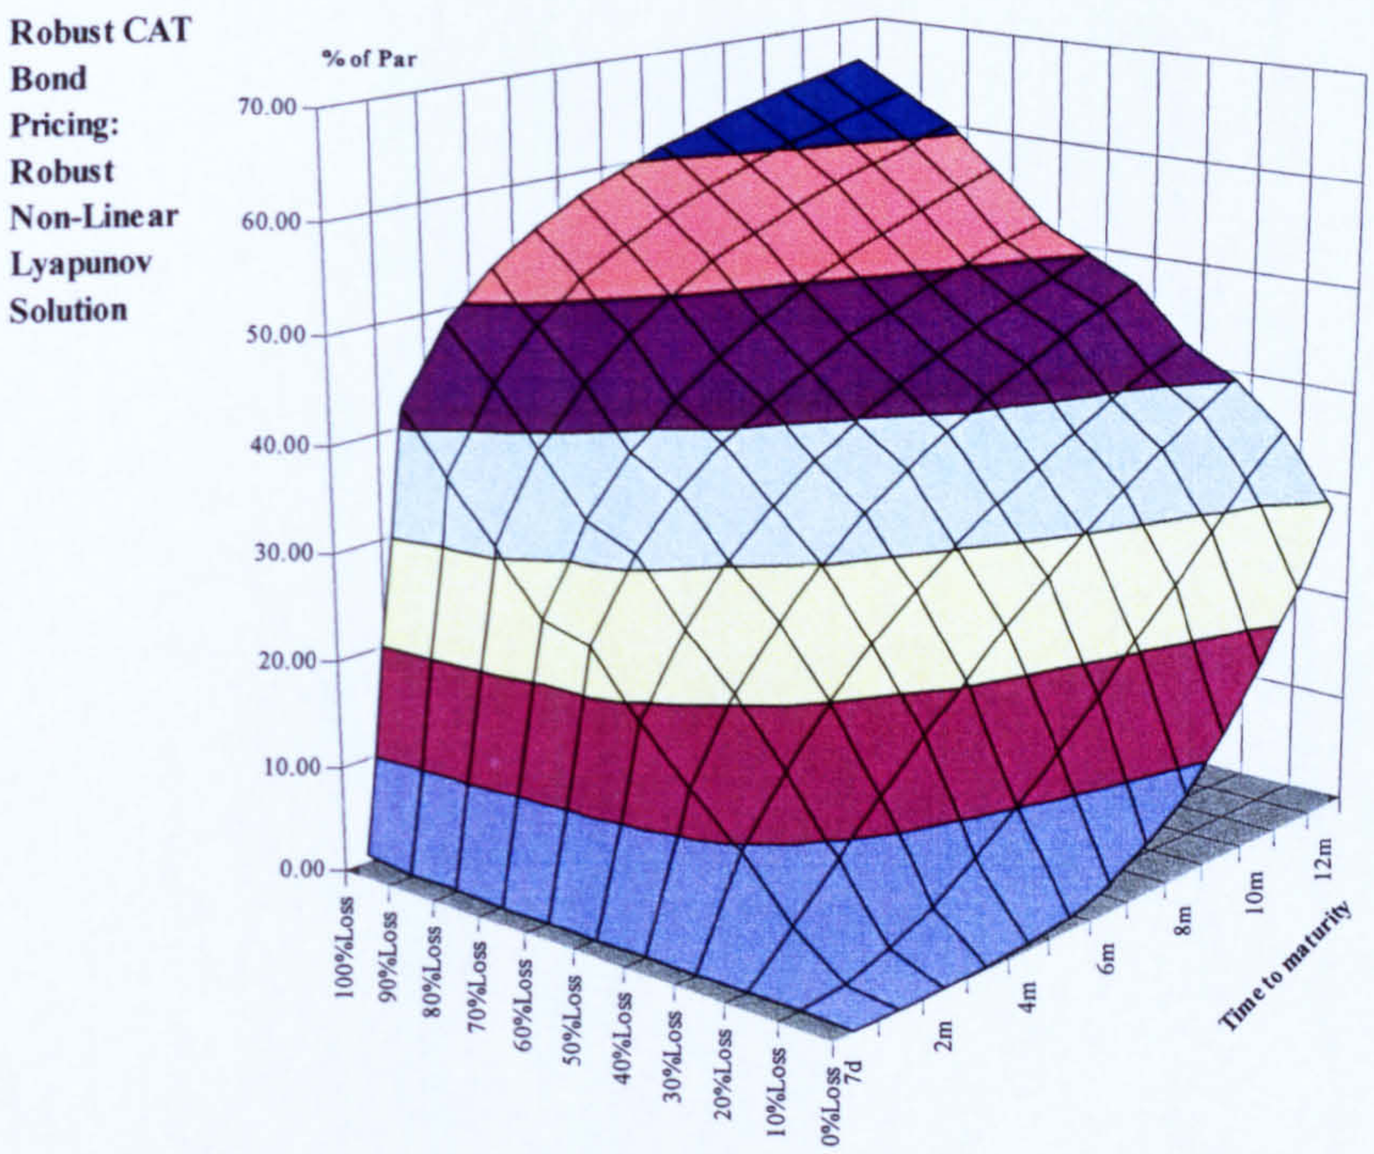

Figure 2.8: Short-term CAT bond valuation: Non-Linear Lyapunov

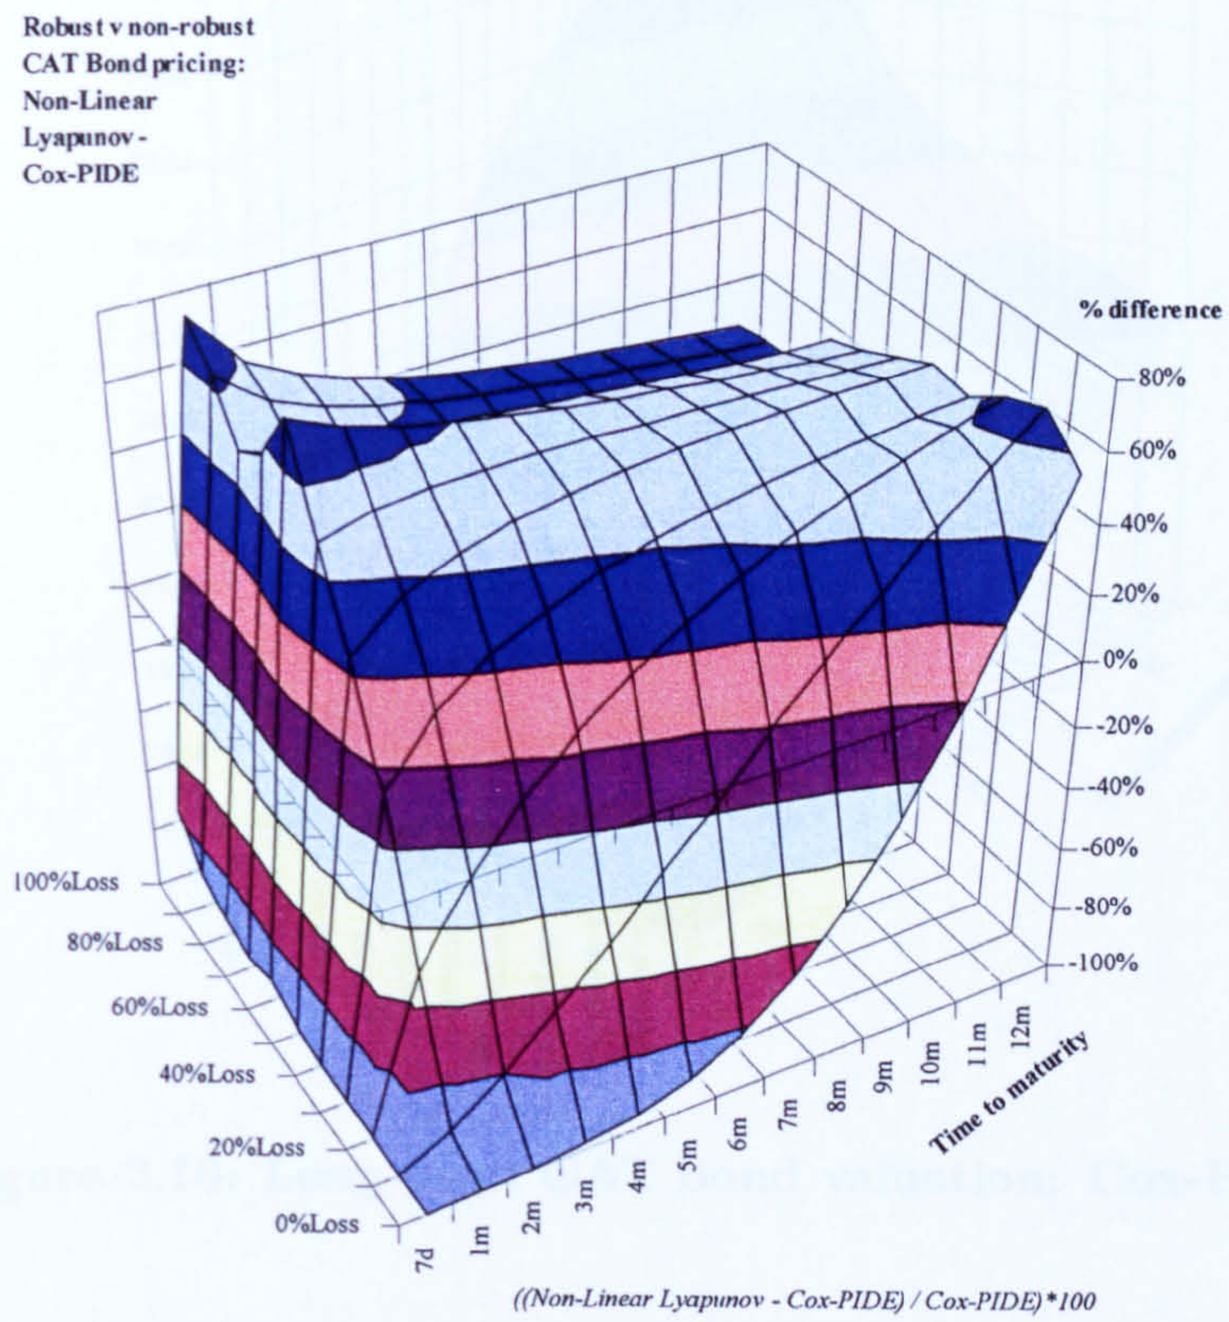

Figure 2.9: Short-term CAT bond model valuation: Non-Linear Lyapunov v Cox-PIDE

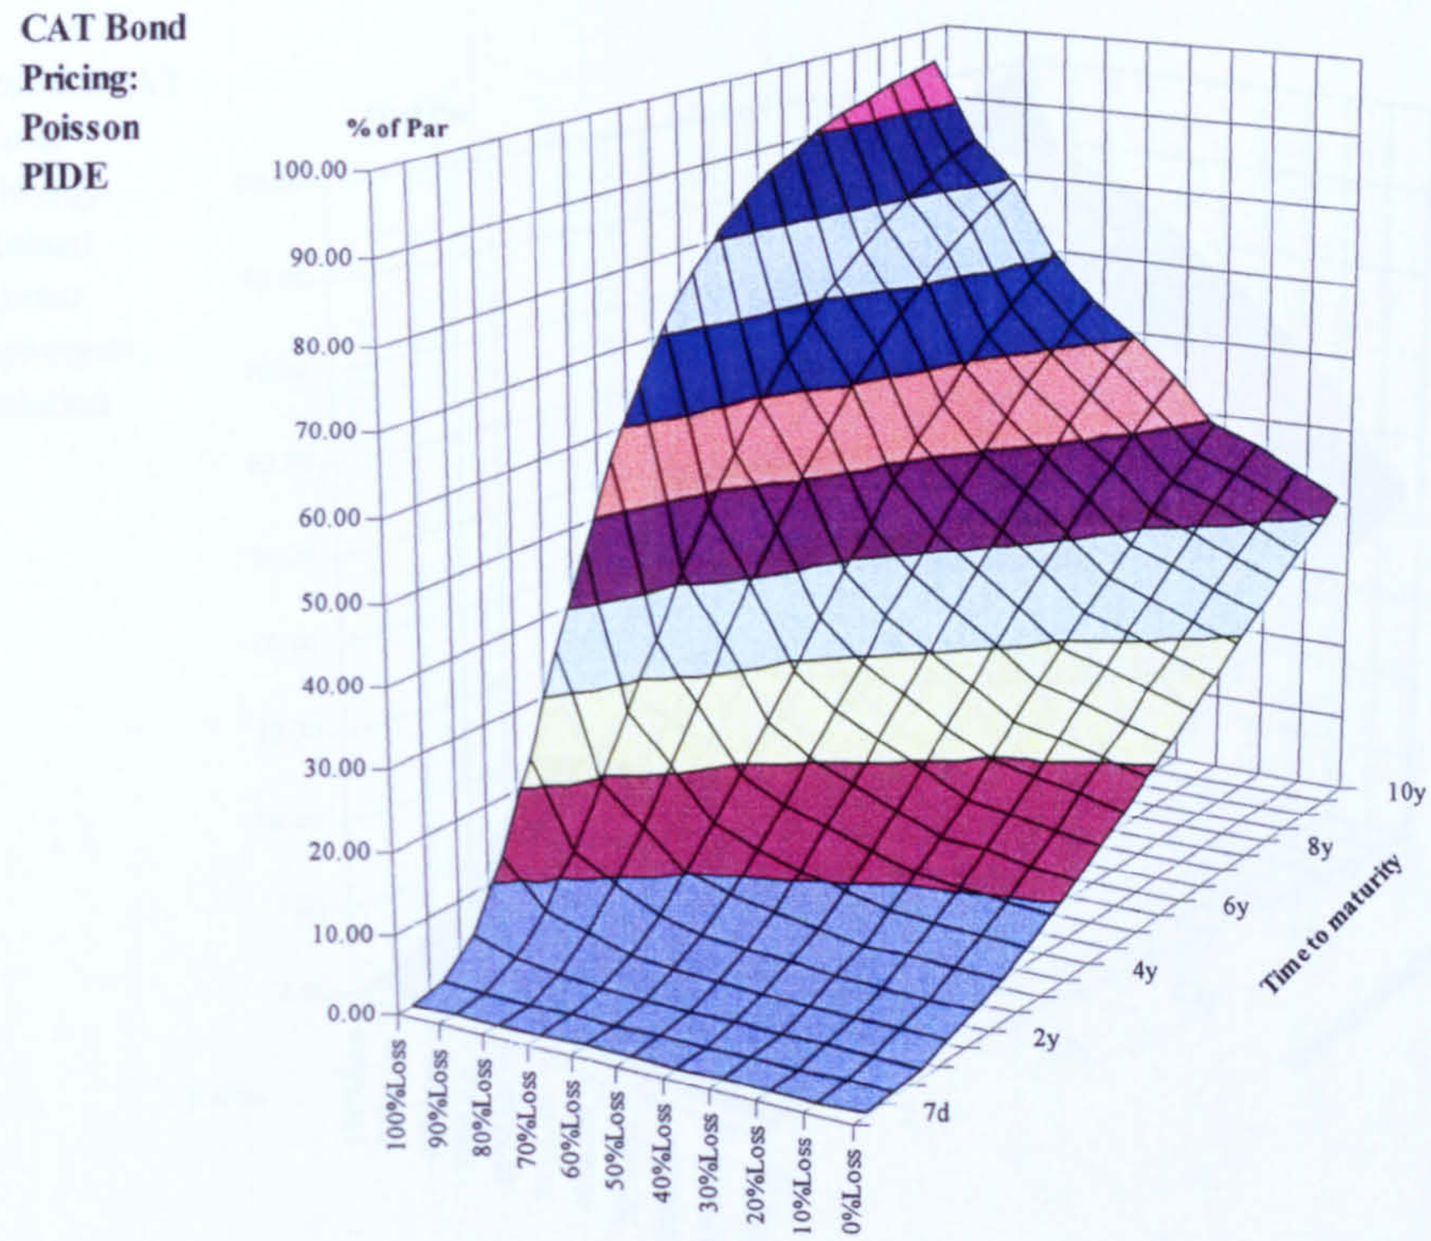

Figure 2.10: Long-term CAT bond valuation: Cox-PIDE

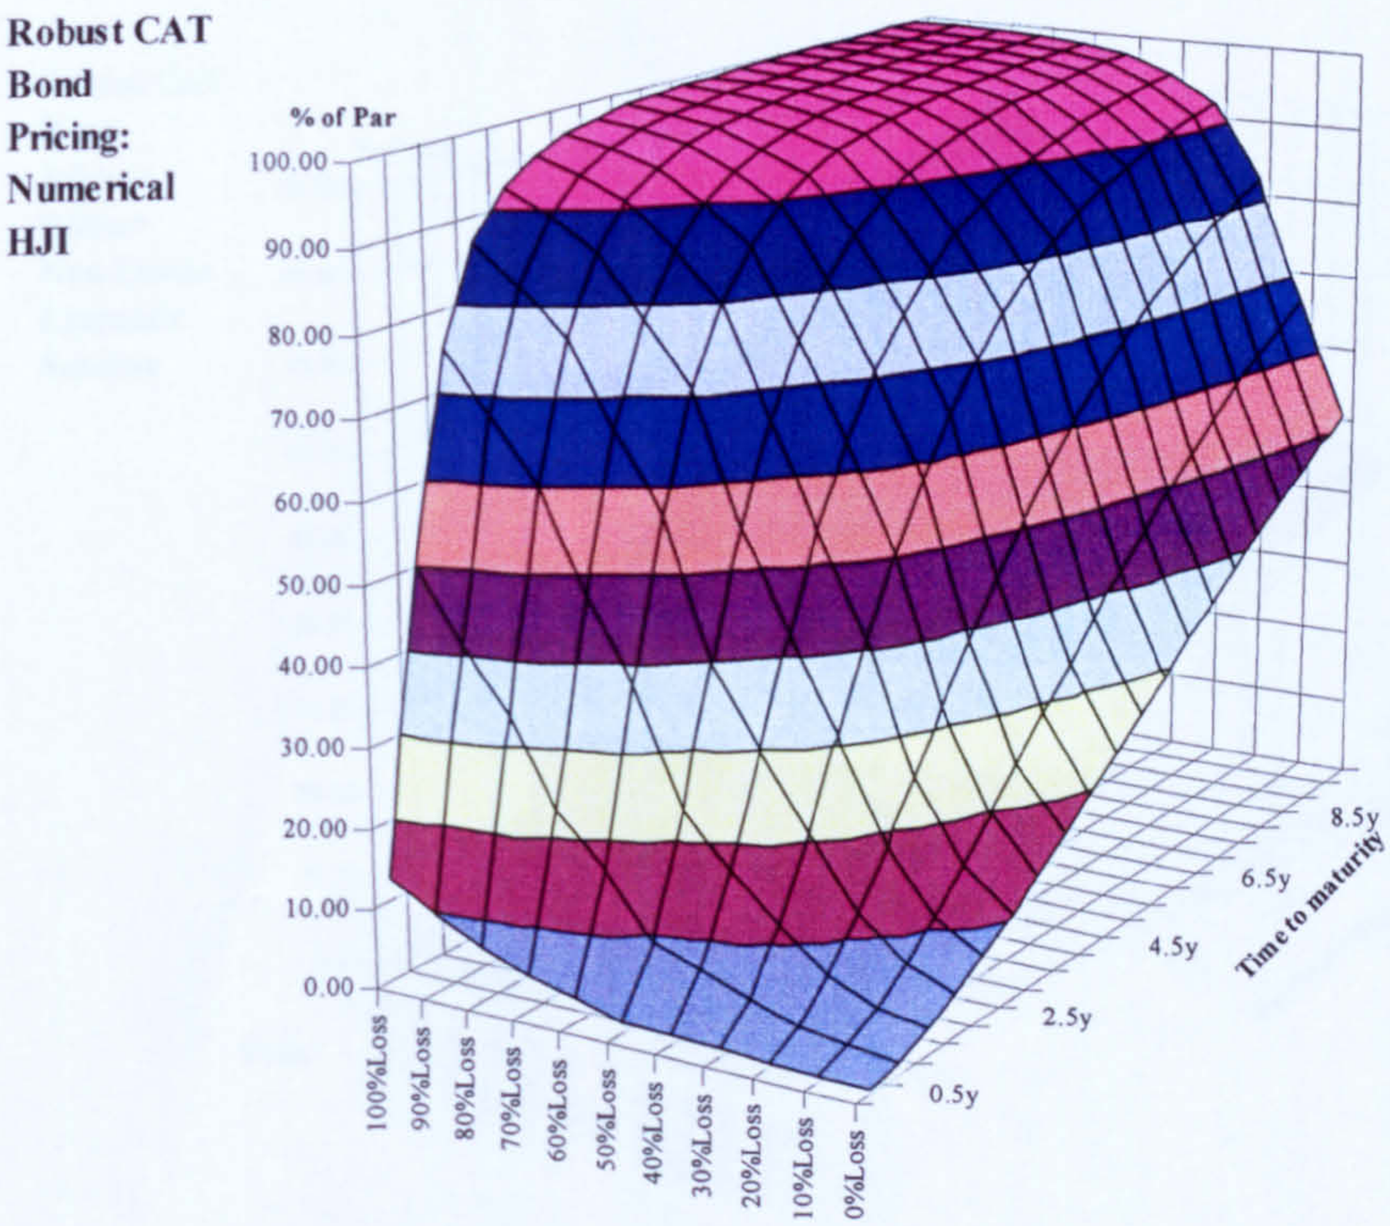

Figure 2.11: Long-term CAT bond valuation: Numerical HJI

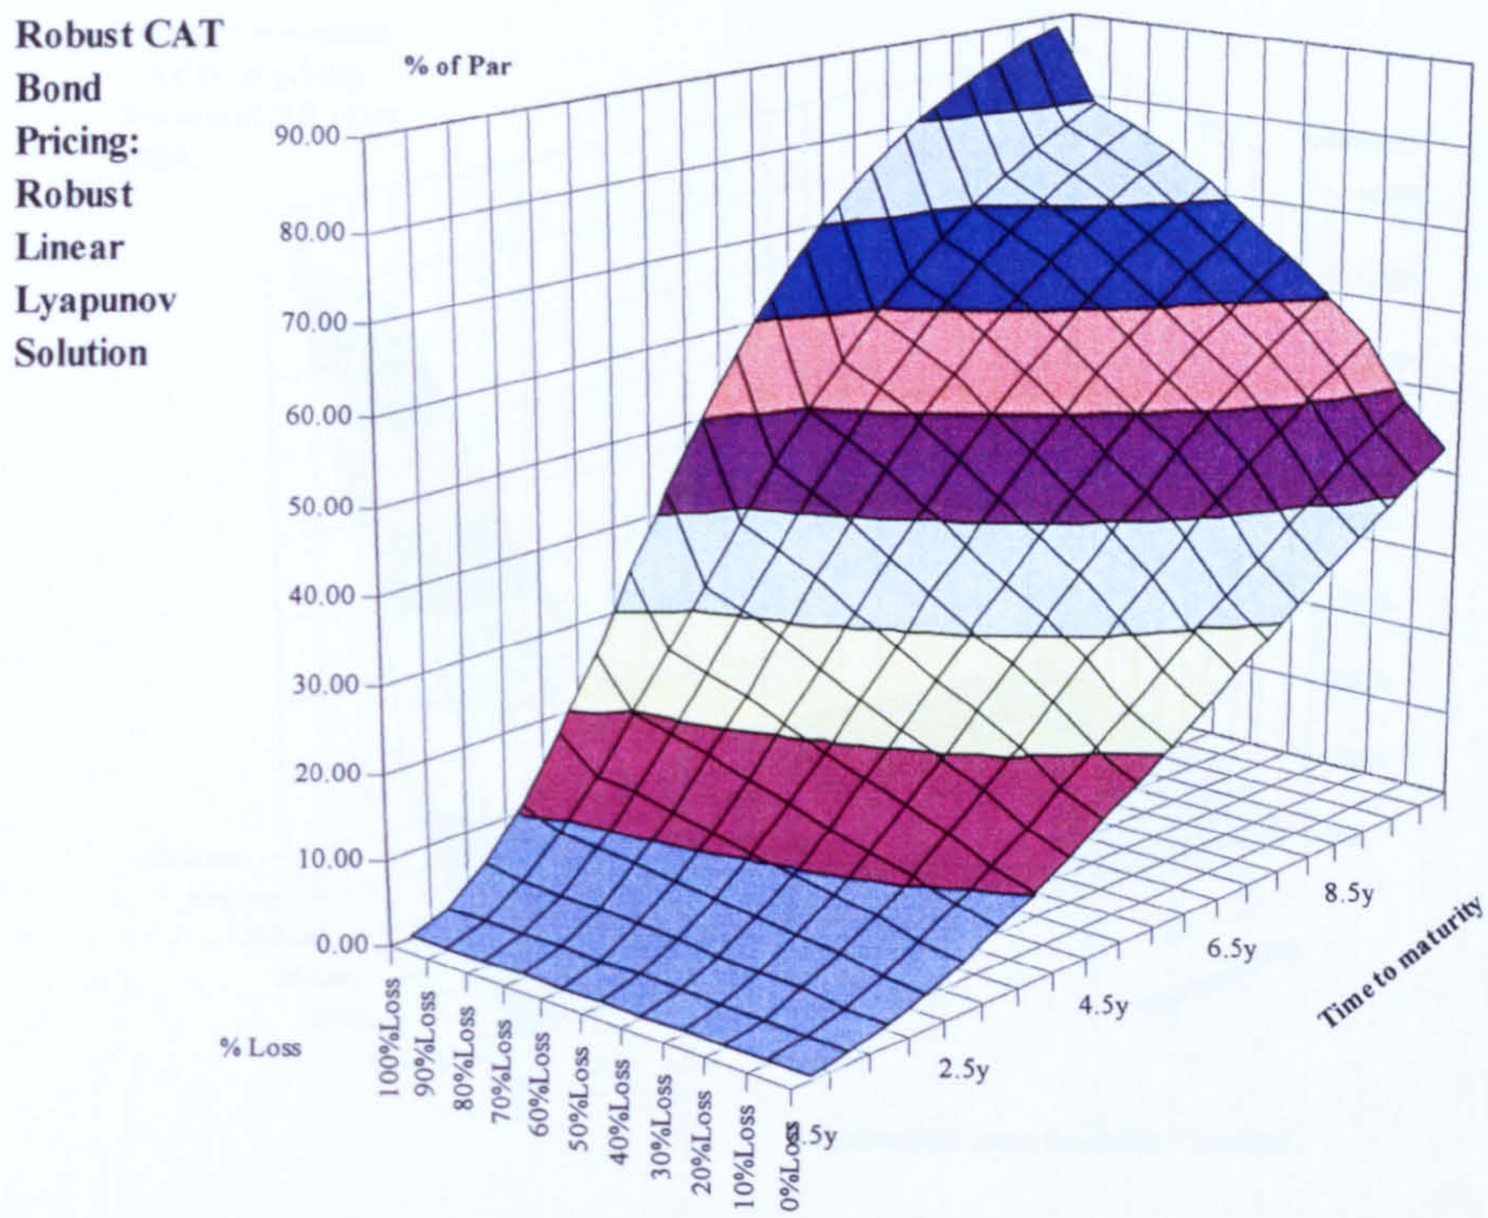

Figure 2.12: Long-term CAT bond valuation: Linear Lyapunov

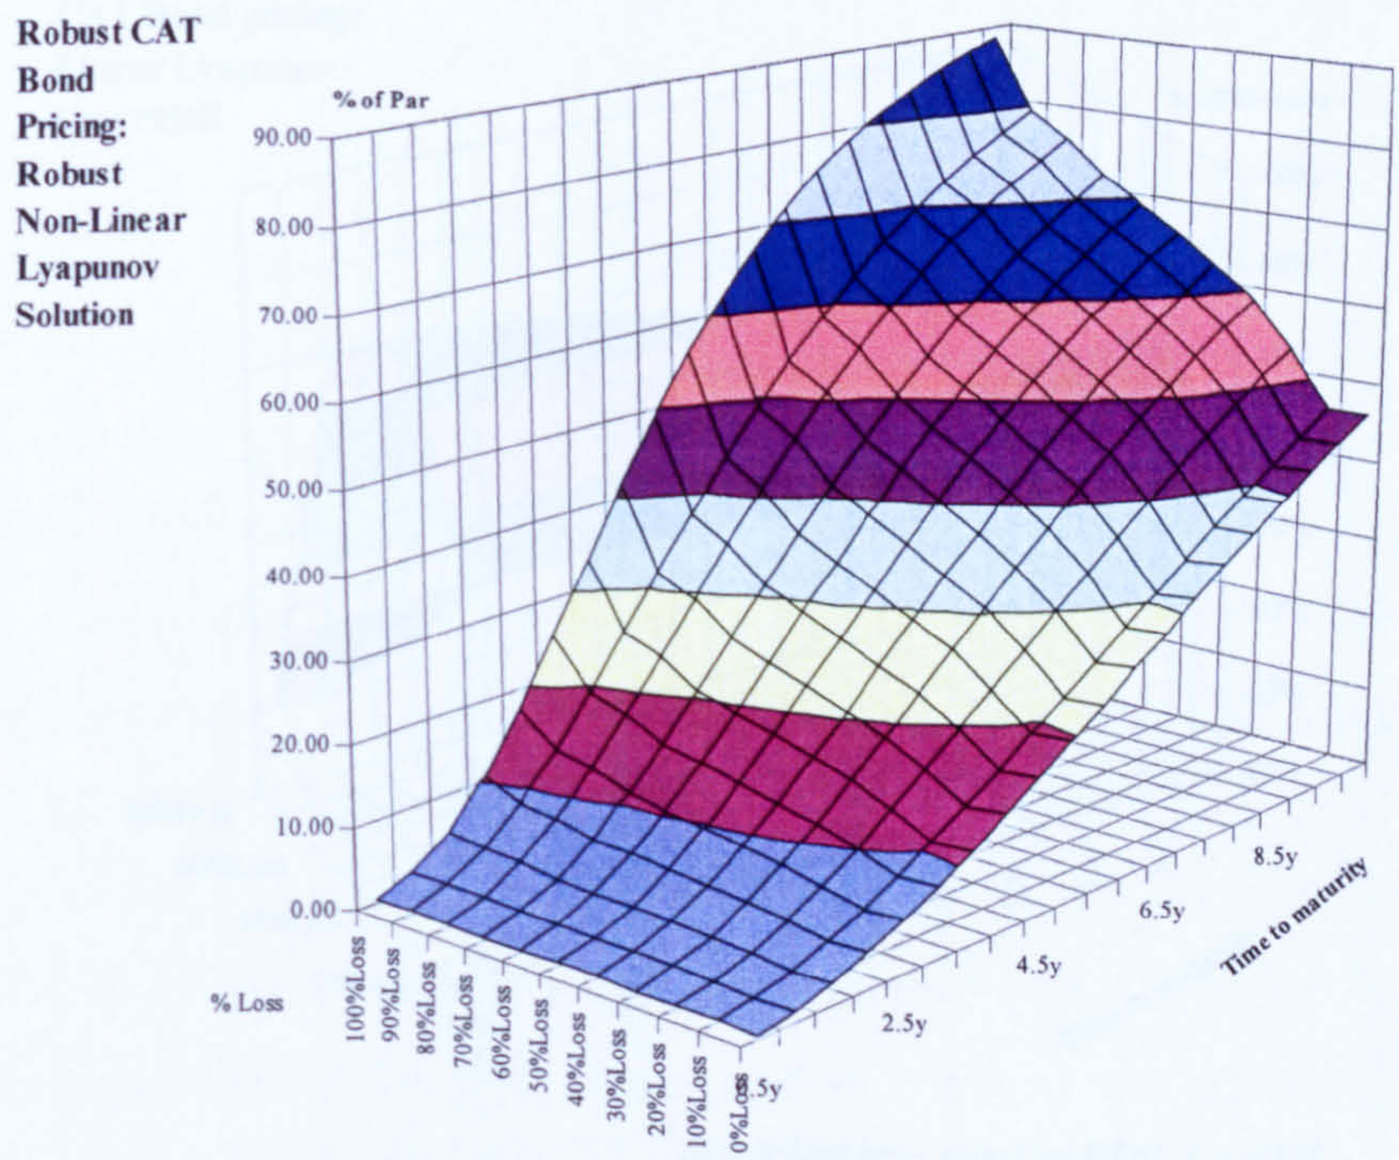

Figure 2.13: Long-term CAT bond valuation: Non-Linear Lyapunov

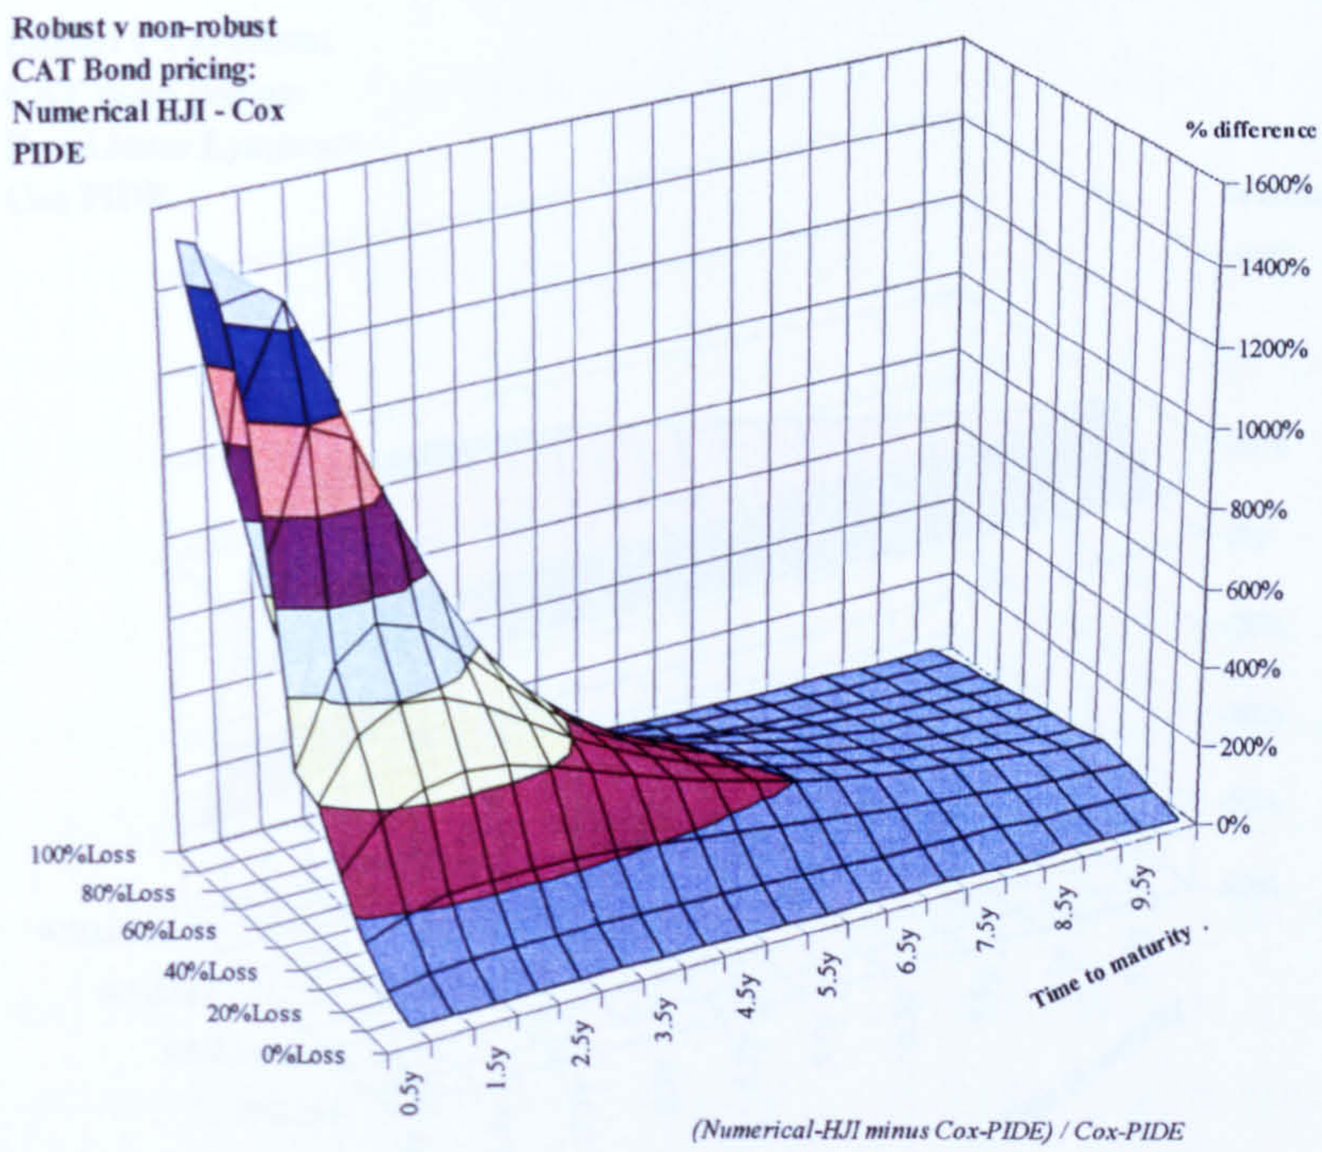

Figure 2.14: Long-term CAT bond valuation: Numerical HJI v Cox-PIDE

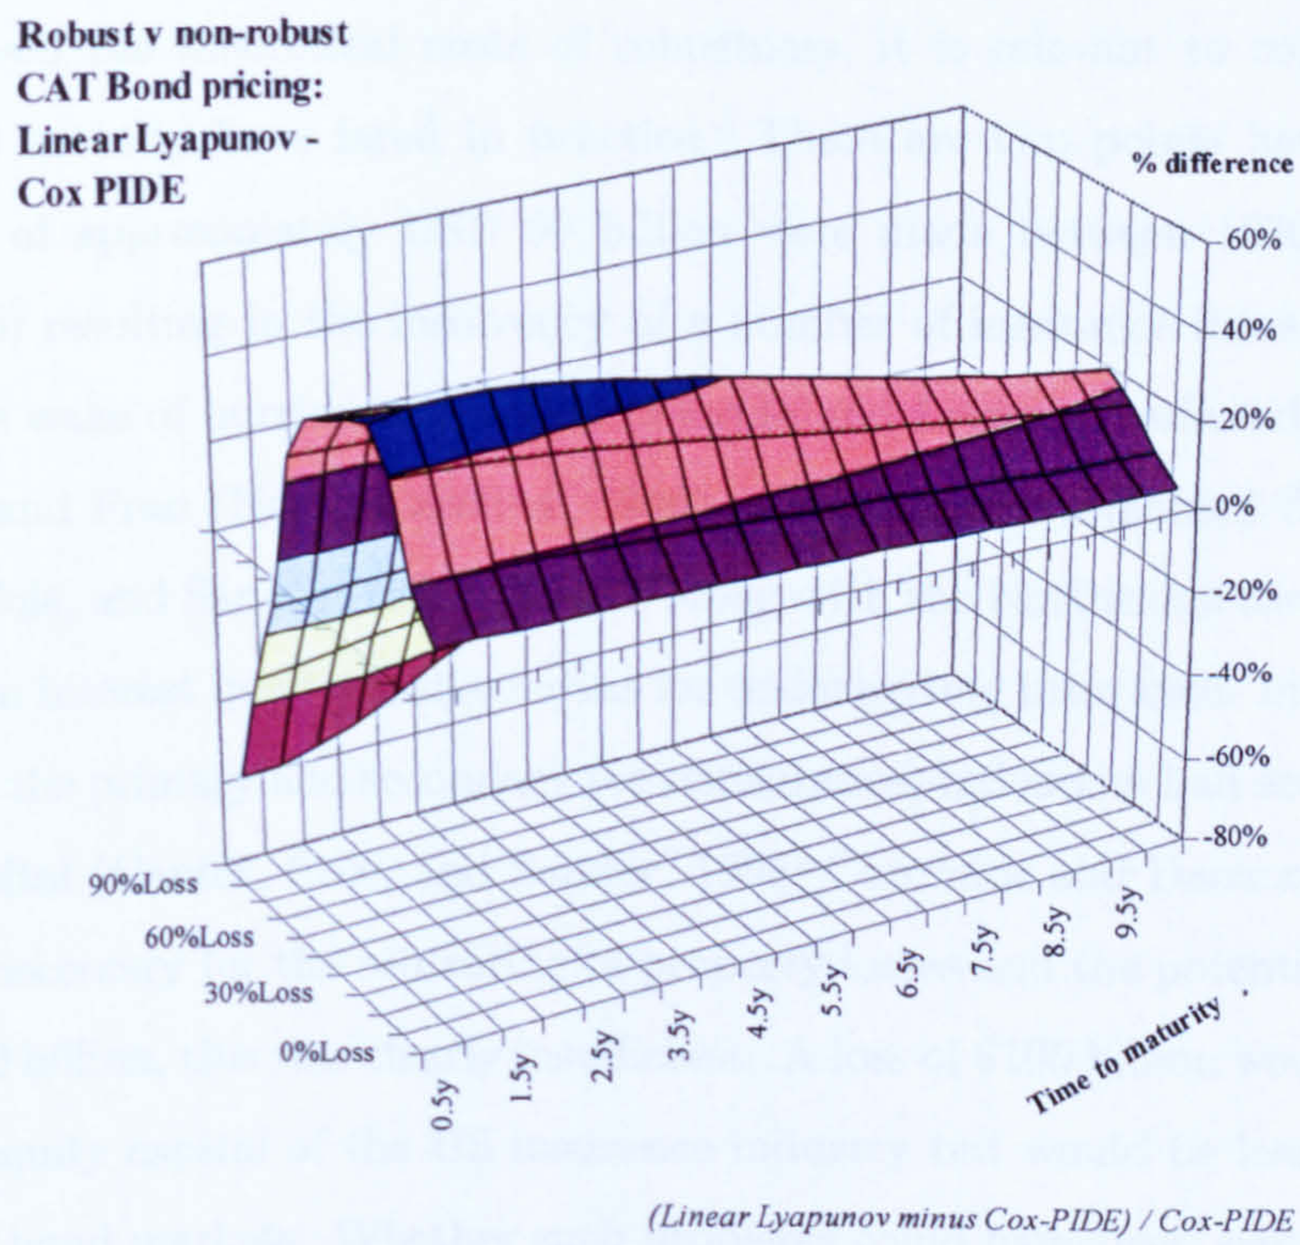

Figure 2.15: Long-term CAT bond valuation: Linear Lyapunov v Cox-PIDE

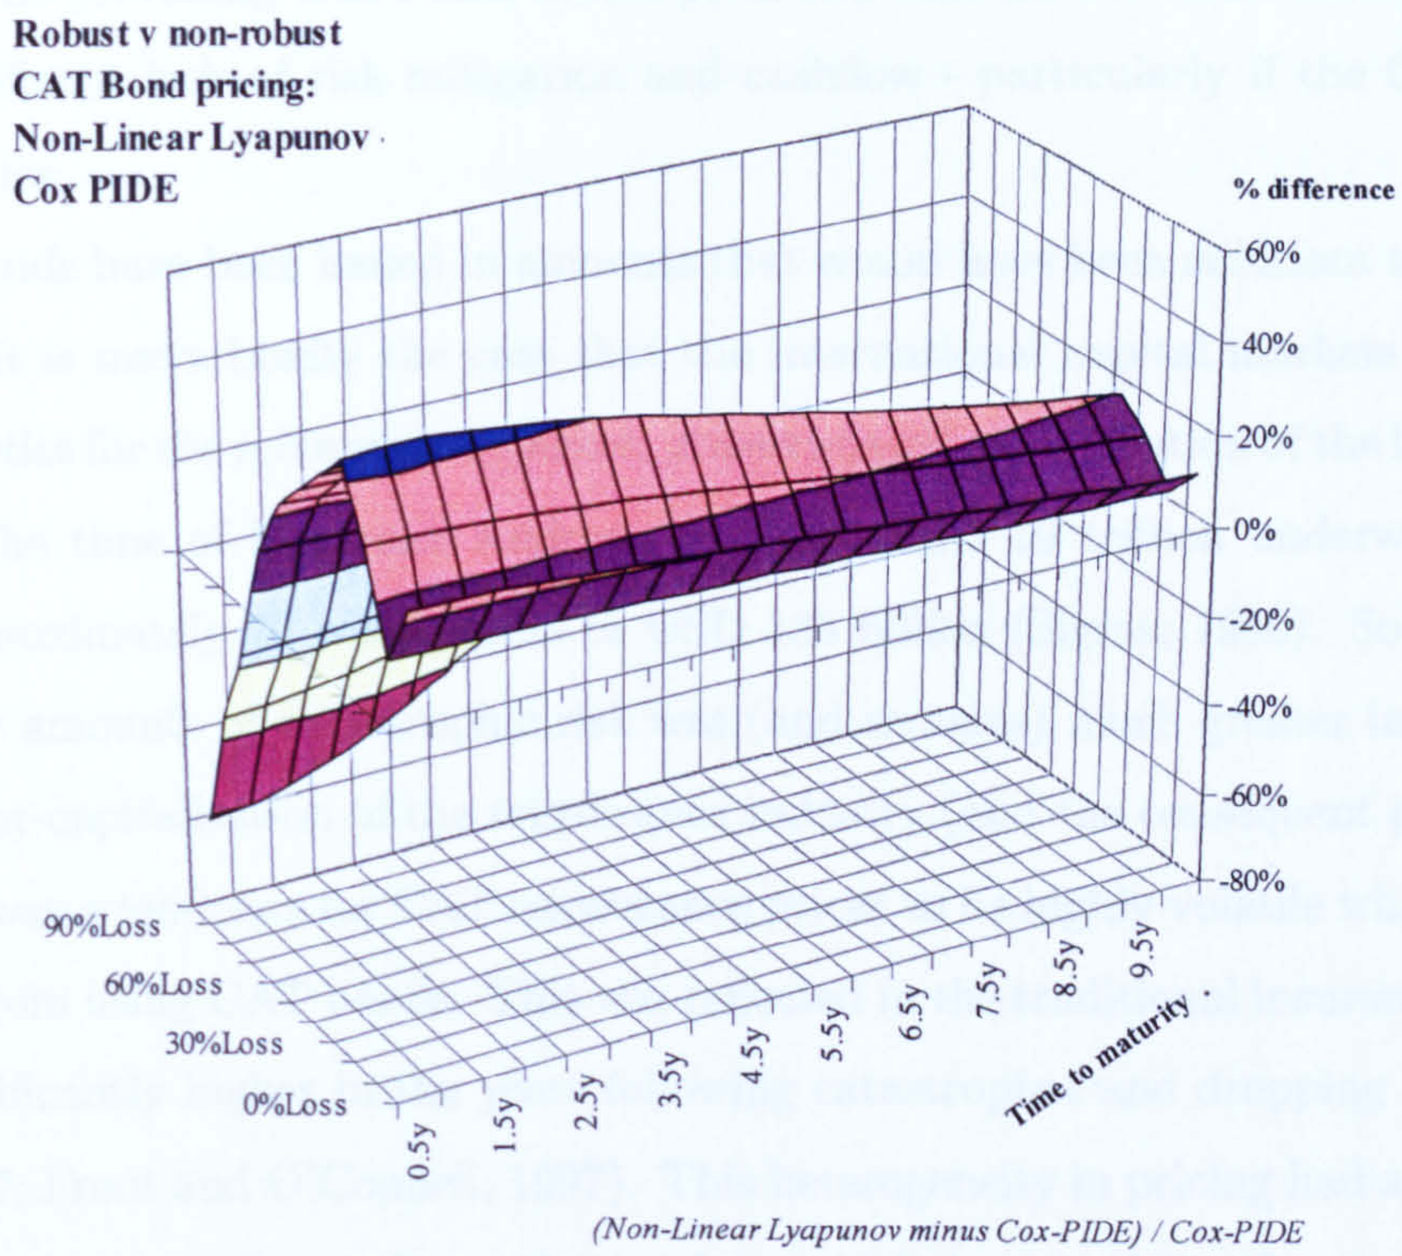

Figure 2.16: Long-term CAT bond valuation: Non-Linear Lyapunov v Cox-PIDE

Having analysed the theoretical costs of robustness, it is relevant to consider how such valuation techniques would actually have fared in practice. There are two points here. First, actual property insurance claims of approximately USD 60 billion were made between 1990 and 1996 (Canter, Cole, and Sandor; 1996) resulting in the insolvency of a number of insurance firms. These bankruptcies were brought on in the wake of hurricanes Andrew (Florida and Louisiana affected, 1992), Opal (Florida and Alabama, 1995) and Fran (North Carolina, 1996), which caused combined damage totalling USD 19.7 billion (Canter, Cole, and Sandor; 1996). These, along with the Northridge earthquake (1994) and similar disasters, led to an interest in alternative means for underwriting insurance. In 1995, when the CAT bond market was born, the primary and secondary (or reinsurance) industries had access to approximately USD 240 billion in capital (Canter, Cole, and Sandor; 1996; Cummins and Danzon; 1997). Given the capital level constraints necessary for the reinsuring of property losses and the potential for single-event losses in excess of USD 100 billion, this was clearly insufficient. A loss of $100 billion would consume approximately 30 - 40% of the equity capital of the US insurance industry but would be less than 0.5% of the value of the US stock and bond markets. Whether such problems could have been avoided by insurers attempting more sophisticated risk management through the issuance of CAT bonds that had been valued using robust valuation methods is difficult to assess directly. However, what can be concluded is that in terms

of total revenue, the prevailing Cox-PIDE model produced fundamental mis-valuations that would have resulted in a significant lack of risk mitigation and cashflow - particularly if the CAT bonds had been short-term securities.

Could CAT bonds have been issued in amounts that would have been sufficient to enable the required risk mitigation? It is undoubtedly the case that the international capital markets provided a potential source of risk appetite for the reinsurance market. An estimated capitalisation of the international financial markets around the time of hurricane Andrew, of about USD 19 trillion underwent an average daily fluctuation of approximately 70 basis points or USD 133 billion (Sigma; 1996). So, clearly the capacity to bear such large amounts of catastrophic risk was (and remains) much greater in the capital markets. However, the under-capitalisation of the reinsurance industry (and the consequent potential default risk) meant that there was a tendency for CAT reinsurance prices to be highly volatile which discouraged many potential issuers from using CAT bonds. This was reflected in the traditional insurance market, with rates on line being significantly higher in the years following catastrophes and dropping off in the intervening years (Sigma; 1997; Froot and O'Connell; 1997). This heterogeneity in pricing had a very strong damping effect, forcing many re-insurers to leave the market, which in turn has adverse consequences for the primary insurers. A number of reasons for this volatility have been advanced (see for example Cummins and Danzon; 1997 and Winter; 1994).

Some of the traditional assumptions of derivative security pricing are not correct when applied to these instruments due to the properties of the underlying contingent stochastic processes. There is evidence that certain catastrophic natural events have (partial) power-law distributions associated with their loss statistics (Barton and Nishenko; 1994), which if true, would overturns the traditional log-normal assumption of derivative pricing models and makes robustness hard if not impossible to achieve without using non-linear models. There are also well-known statistical difficulties associated with the moments of power-law distributions[16], thus rendering it impossible to employ traditional pooling methods and consequently the central limit theorem. Given that heavy-tailed or large deviation results assume, in general, that at least the first moment of the distribution exists, there will be difficulties with applying extreme value theory to this problem (Embrechts, Resnick, and Samorodnitsky; 1999). It would seem that these characteristics may render traditional actuarial or derivatives pricing approaches ineffective.

Although there is some similarity with the valuation of defaultable bonds, there are additional features to modelling the CAT bond price which are not to be found in models of ordinary corporate or government securities. The main feature is that the trigger event that underlies CAT bond pricing is dependent on

[16]This has become a significant research topic in its own right - see for example, "Multifractal Power Law Distributions: Negative and Critical Dimensions and Other "Anomalies, Explained by a Simple Example", by Benoit B. Mandelbrot, in the Journal of Statistical Physics, Vol. 110, Nos. 3-6, March 2003.

both the frequency and severity of natural disasters. The Cox-PIE model in this section is used to reduce to a minimum any assumptions about the underlying distribution functions in the interests of generality of application. However, given the daily availability of PCS loss data, it is also appears to be reasonable to assume that loss levels are instantaneously measurable and updatable, which makes it straightforward to adjust the underlying process to accommodate a development period and it is this feature that is explicitly included in the next section where results of CAT option valuation are reported that once again use PCS loss data.

There is a natural similarity between the pricing of catastrophe bonds and the pricing of defaultable bonds. Defaultable bonds, by definition, must contain within their pricing model a mechanism that accounts for the potential (partial or complete) loss of their principal value. Defaultable bonds yield higher returns, in part, because of this potential defaultability. Similarly, CAT bonds are offered at high yields because of the unpredictable nature of the catastrophe process. With this characteristic in mind, a number of pricing models for defaultable bonds have been advanced (e.g. Jarrow and Turnbull, 1995, Duffie and Singleton, 1999, Zhou and 1997). The trigger event for the default process has similar statistical characteristics to that of the equivalent catastrophic event pertaining to CAT bonds.[17]

2.5.3 CAT option valuation results

Before proceeding to present and analyse the results of the valuation of the CAT options based on PCS index underlyings, it is necessary to explain a number of points on the PCS data, option contracts and the models to be compared, beyond those already made earlier in this chapter. To begin with, the first publicly available, exchange-tradeable product explicitly designed to transfer catastrophe risk from the insurance markets to the financial markets was the catastrophe (CAT) future, which was first introduced to the Chicago Board of Trade in 1992. However, due to technicalities around its construction, the CAT future never became popular amongst potential users and in 1995 it was replaced by the PCS-index based CAT option. In all, CAT derivatives based on one or more PCS indices were traded on CBOT from 11 December 1992 through to 03 November 1997[18]. However, the results reported in this section are concerned solely with PCS CAT options, which only traded for the more limited subset of dates indicated in table 2.1 below.

Although CAT options were available as both puts and calls, they were most frequently traded as spreads on one of the PCS indices. This was because trading a spread had the effect of creating the

[17] In an allied application to mortgage insurance, the similarity between catastrophe and default in the log-normal context has been commented on (Kau and Keenan; 1996).

[18] Appendix 2 contains full product specifications for all of the option contracts traded, together with details of the underlying indices. See Appendix Table 2.1 for a summary listing of the CBOT contract codes and periodic coverage.

same overall hedging profile as an excess of loss catastrophe reinsurance contract which would be the usual means of providing cover for a given slice of the catastrophe liability structure. From a practical perspective, CAT options were available in 3 month and 12 month loss occurrence periods, with the choice of either a 6 month or 12 month development period after the loss occurrence period which provided the time to get a clearer picture of the cost of catastrophes that may have happened late in the occurrence period. During the period that the PCS CAT options traded, the PCS indices were published daily by ISO. However, in practice, the time and cost of re-estimating the claims levels, meant that the PCS issued claims updates daily, but re-calculated the actual financial loss indices only once or twice a month - depending on the level of claim activity and the severity of the catastrophes in hand[19]. There were four quarterly contracts, namely: March/June/September/December:

- March contracts cover losses occurring in the first quarter.
- June contracts cover losses in the second quarter.
- September contracts cover losses in the third quarter.
- December contracts covers losses in the fourth quarter.

There was also a single annual contract covering losses for the whole calendar year. Contracts also traded with either a 6m or 12m development period. From a practical perspective, liquidity was concentrated in those option contracts with a 3m loss period coupled with a 6m development period; so, on the grounds of liquidity[20] and representative valuation, this thesis looked solely at just these contracts. Table 2.1 summarises the option contracts and the periods over which they traded. In line with liquidity and market convention at the time, the valuations reported in this section were calculated based on using the puts and calls to create 0.5 basis point spreads (25 basis points either side of the strike) around the main loss index value strike levels indicated in Appendix 2.

PROPERTY CLAIM SERVICES (PCS) INSURANCE OPTIONS		
Area Covered	**Ticker**	**Period Covered**
NAT'L INS. LARGE CAP-6 MONTH CALL	DNC	09/29/1995-11/03/1997
NAT'L INS. LARGE CAP-6 MONTH PUT	DNP	09/29/1995-11/03/1997
NAT'L INS. SMALL CAP-6 MONTH CALL	QNC	09/29/1995-11/03/1997
NAT'L INS. SMALL CAP-6 MONTH PUT	QNP	09/29/1995-11/03/1997

Table 2.1: CBOT PCS CAT option contracts

[19]During the period which the PCS CAT options traded, the indices were updated daily based on the total number of claims, but multiplied by the average cost per claim. This estimation exercise was far less costly than a full calculation of the actual financial loss data and took only a fraction of the time. Full calculation occurred only once or twice per month - with the actual frequency depending on the volume of new claims and the size of any catastrophes that had occurred.

[20]Open interest in the quarterly contract was always substantially greater than that of the annual contract according to CBOT statistics.

Having reviewed the practicalities of the PCS CAT option contract, it is now necessary to fill-in a number of the contract specific details required to value a PCS option contract. The first point to note is that the strike prices of PCS CAT options were defined with respect to particular cumulative loss levels. Small cap contract strikes covered losses in the range from 0.1 ($10 million) to 200 points ($20 billion), whilst the large cap contracts covered losses ranging from 250 ($25 billion) to 500 points ($50 billion). Therefore, let T_1 and T_2 be the start and final maturity times of the option contract period respectively, such that $0 < T_1 < T_2$. If the loss occurrence period is $(0, T_1]$ and the development period is $(T_1, T_2]$, with catastrophes occurring at times $0 < \tau_1 < \tau_2 < ...$, then if each catastrophe is modelled using an individual index $\{L_t^i\}$ and the number of catastrophes in $(0, T_1]$ is denoted by N_t so that at time t, the PCS index is therefore given by

$$L_t = \sum_{i=1}^{N_{t \wedge T_1}} L_t^i$$

So, if $L(0, T_1; T_1, T_2; t)$ is the estimated industry wide dollar denominated loss amount at time t, then the corresponding PCS index value for the 500 point contract measured in tenths of an index point (where each one tenth of an index point is equivalent to a loss level of $10m) is given by

$$L^*(0, T_1; T_1, T_2; t) = \left[\frac{L(0, T_1; T_1, T_2; t)}{10,000,000} + 0,5\right] \frac{1}{10}$$

where each on tenth of an index point is worth $20. Note that the exercise price of a single option contract at strike k_1 is also measured in index points, as well as the option premium $C(s)$ at the the time of purchase s, then $NV^i(t)$ is the US dollar value of the PCS long call position at time t. To see how this worked in practice, consider the particular case of a large cap contract with a strike of k_1 points, so that the gross value of a long PCS call with loss period $(0, T_1]$ and final maturity T_2 is given by

$$\begin{aligned} C(0, T_1; T_1, T_2) &= 200. \min\{\max\{0, L^*(0, T_1; T_1, T_2; t) - k_1\}; 500 - k_1\} \\ &= 200. \min\left\{\max\left\{0; \left[\frac{L(0, T_1; T_1, T_2)}{10,000,000} + 0,5\right] \frac{1}{10} - k_1\right\}; 500 - k_1\right\} \end{aligned}$$

So that if $C(s)$ is the option premium paid at the time, s, of purchasing the option, $NV(t)$ gives the net value of a long position in the option at time t. As already mentioned, most trading of PCS CAT options took place in the form of call option spreads. So, if k_1 and k_2 are the lower and upper strike prices, respectively, then intuitively they can be thought of as representing the lower and upper limits of a traditional excess-of-loss catastrophe reinsurance contract. Selling a call-spread contract $k_1/k_2 (k_1 < k_2)$ is equivalent to selling a call option with strike price k_1 and automatically buying a call option with a

strike price k_2. By selling a call spread option contract, losses are capped by $k_2 - k_1$, with risk being reduced relative to a simple call. If $L(T)$ is the aggregate PCS index value, then the final payoff at the maturity of the option will be

$$\max\left(\min\left(L(T) - k_1, k_2 - k_1\right), 0\right)$$

Figure 2.17 illustrates the payoff profile of the PCS call option spread.

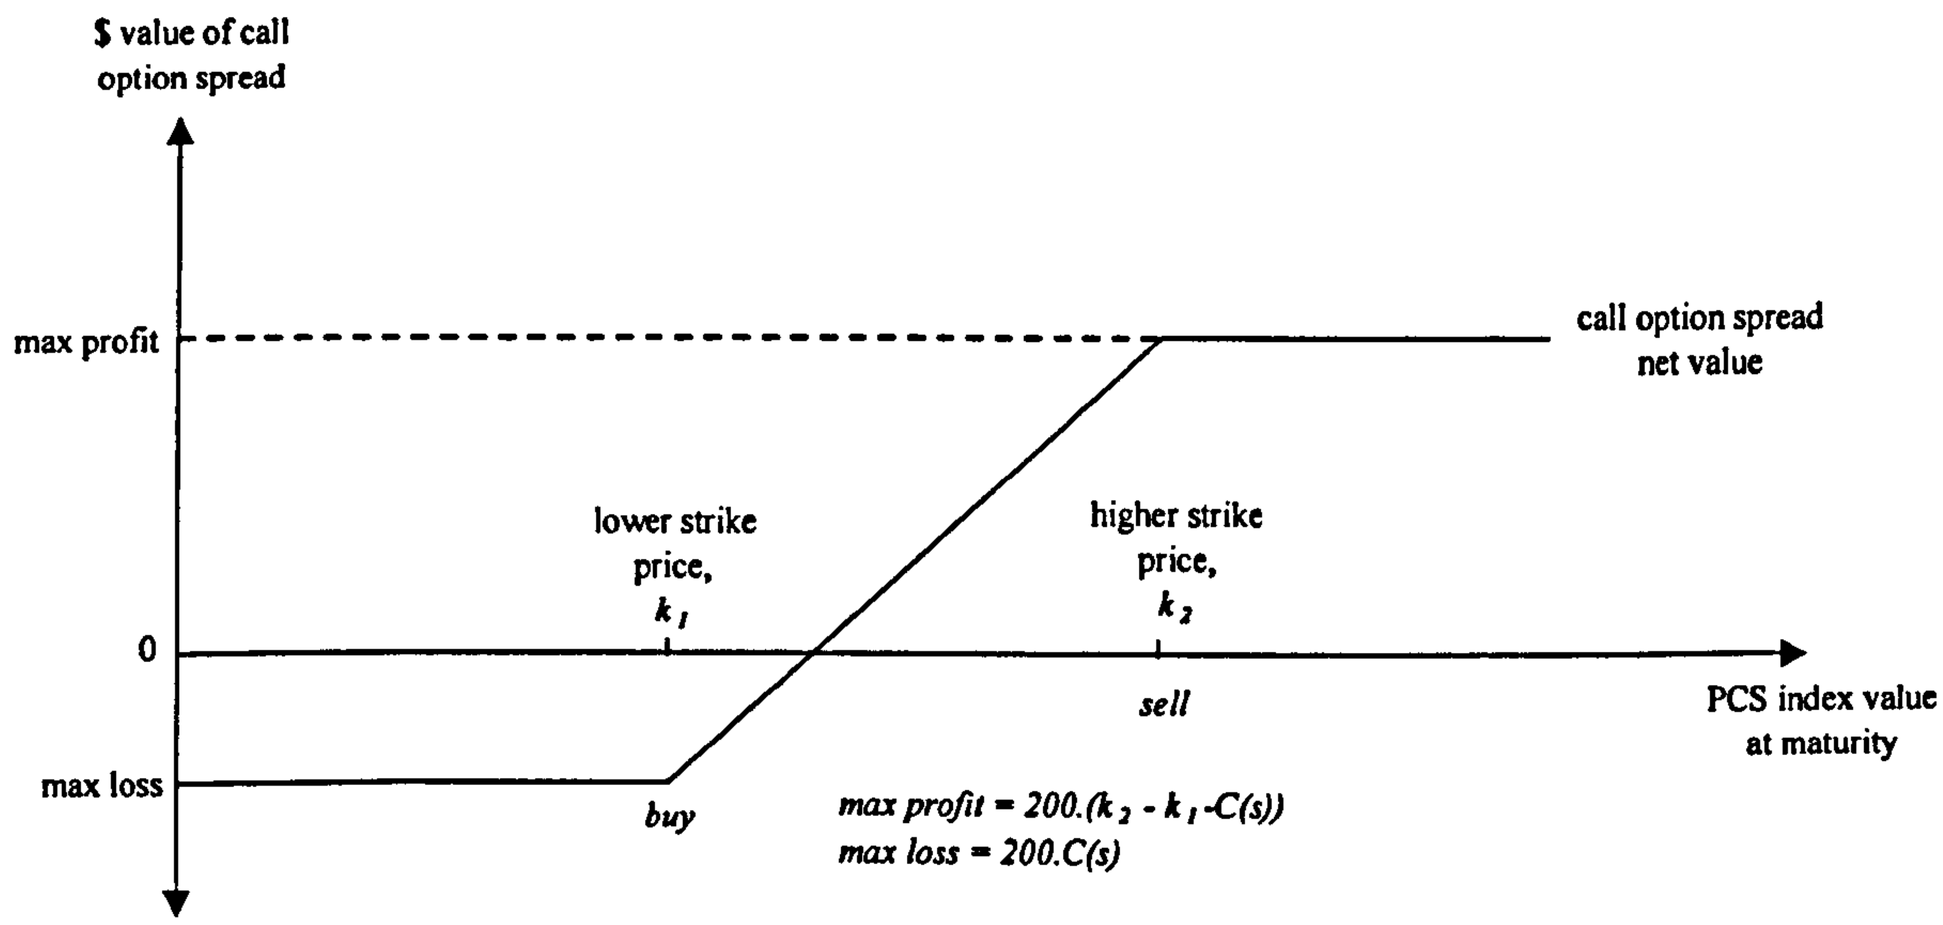

Figure 2.17: PCS call option spread value

Those details of the PCS option contracts required in order to interpret the computations have now been described. The next stage was therefore to use the same overall robust control valuation framework that was developed for pricing the CAT bonds, to derive values for a range of CAT options based on the PCS index using the following four (one non-robust and three robust) models:

1. Cox-Compound Poisson using a translated gamma approximation (Cox-TGA) - following Christensen (2000) and Schmidli (2003)

2. Numerical H∞ robust optimal control (i.e. numerical HJI)

3. Linear Lyapunov robust optimal control

4. Non-Linear Lyapunov robust optimal control

Figures 2.18 to 2.20 show the valuation differences between the Cox-TGA model and each of the robust models for the varying strike levels for PCS call spreads[21]. The strikes used were the main PCS

[21] All of the results presented in figures 2.18 - 2.20 are presented in numerical form in tables 1 - 4 in Appendix 4.

index points, with the call strike spread being set at 25 basis points of index value either side of the given PCS index strike. So, for example, in the case of the 200 point strike, the call option spread was 199.75 - 200.25. The spread options were all valued at point of first issue and the valuation results are undiscounted actual dollar amounts, having been multiplied by 200 (as an index point was worth $200). All of the theoretical price results reported below were calculated at the beginning of the loss period. The primary reason for this was that in the case of contracts traded prior to the commencement of the loss period, there are neither time decay nor event effects to consider, (i.e. no catastrophes have occurred that can affect the option price as the loss period has not started), which, all things being equal, should make the option valuations more stable and therefore easier to compare between models. The secondary reason is that as the volumes and open interest on the contracts were always very small, the traded prices tended to be unrepresentatively volatile and frequently substantially greater than the theoretical prices (even from the numerical HJI model) during the loss and development periods. Using values calculated prior to the commencement of the loss period, helps to minimise what otherwise might be a misleading effect (see Sun (2002) for confirmation of this behaviour).

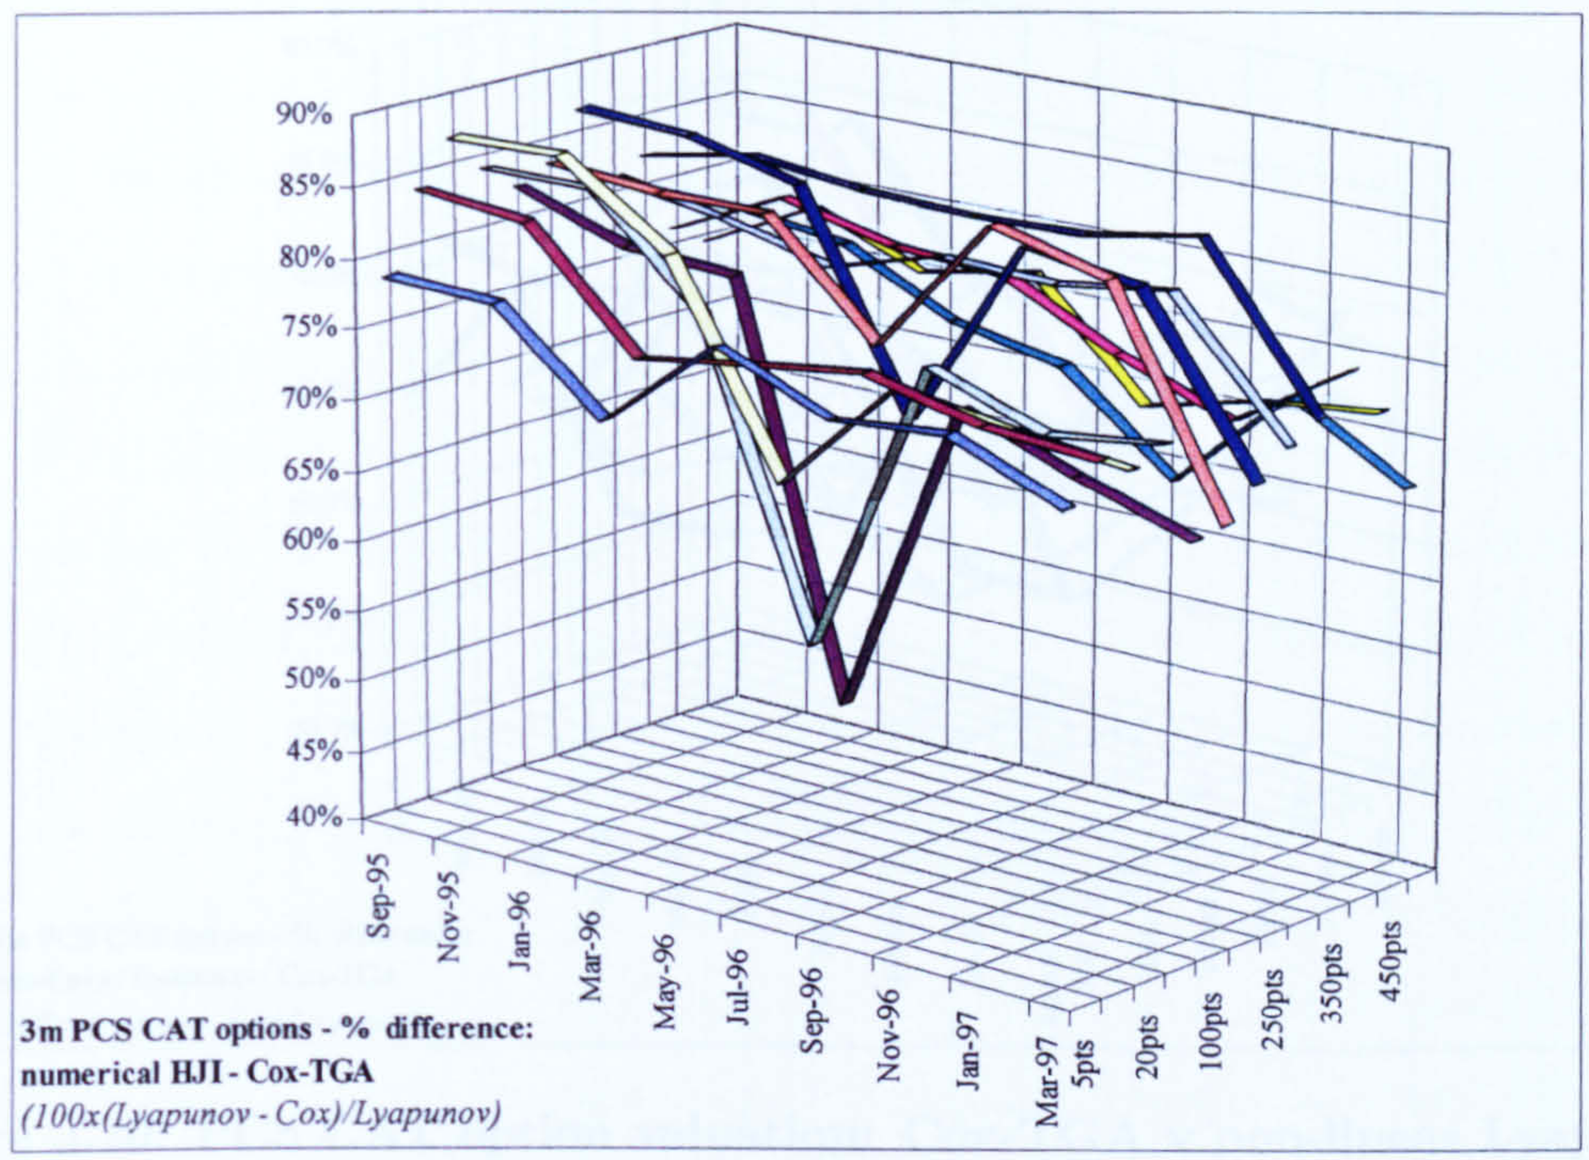

Figure 2.18: PCS CAT option valuation: Cox-TGA v numerical HJI

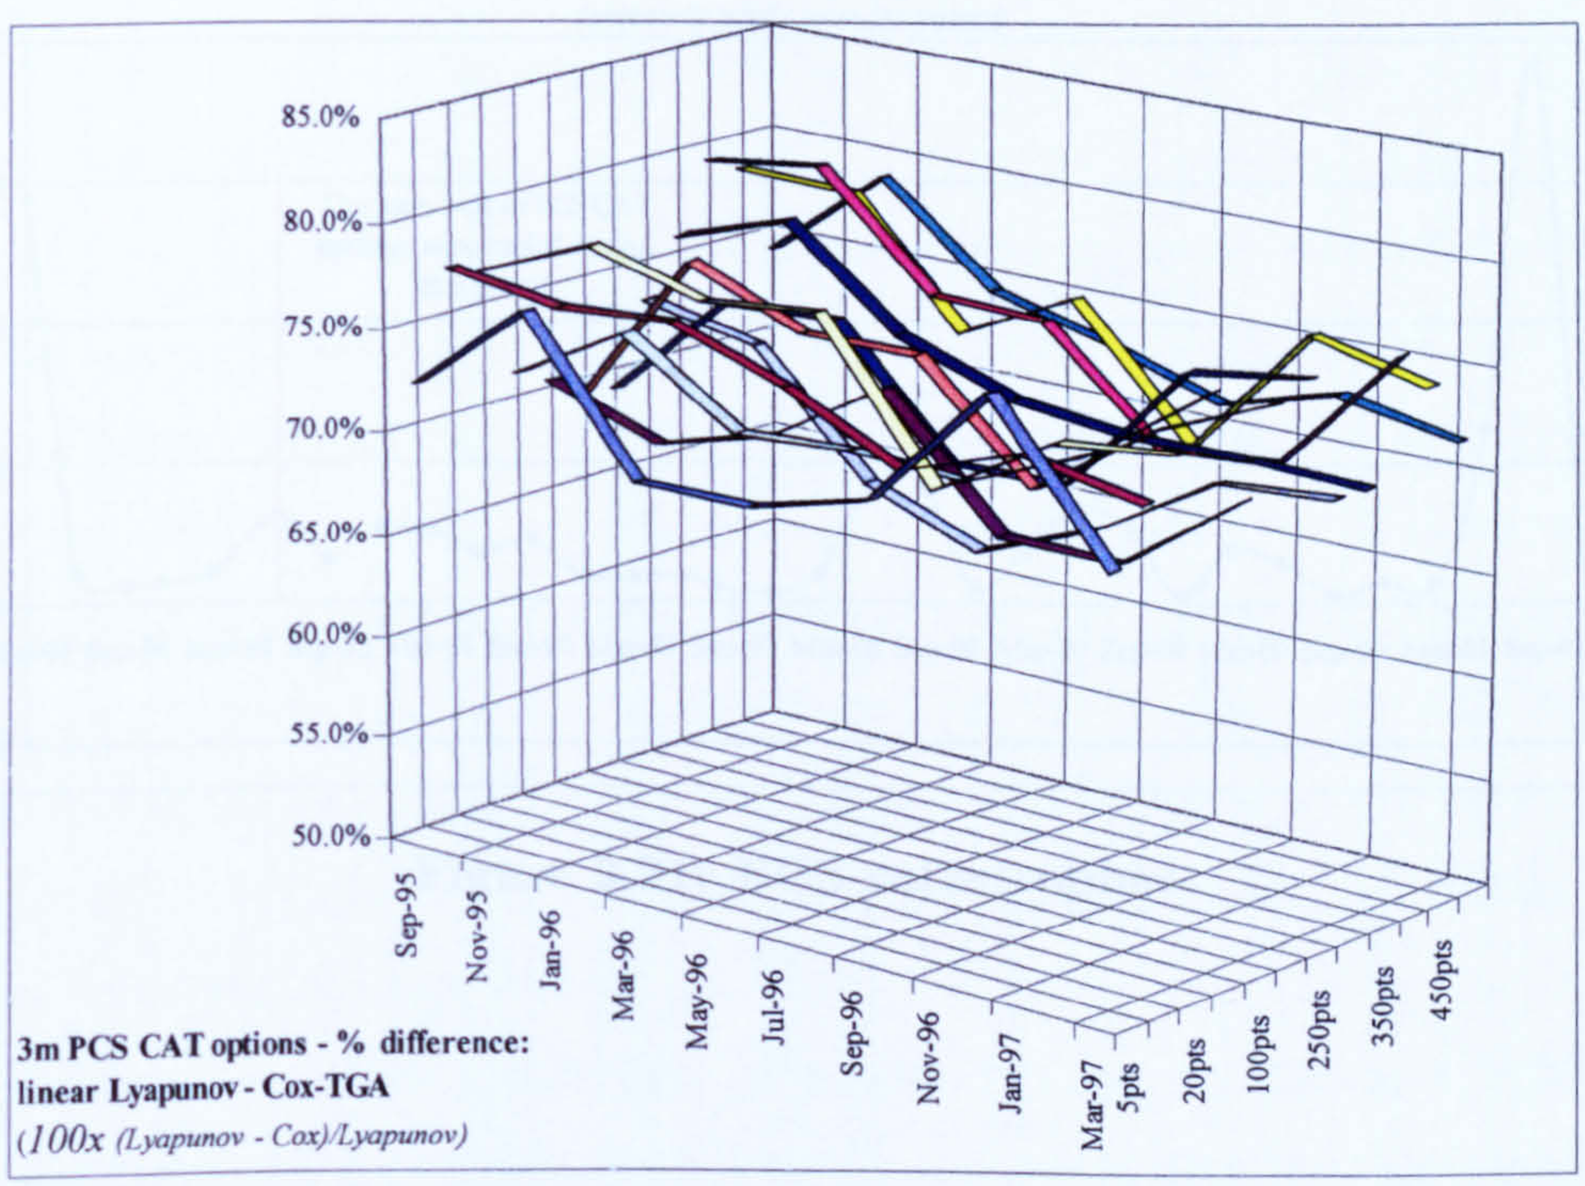

Figure 2.19: PCS CAT option valuation: Cox-TGA v linear Lyapunov

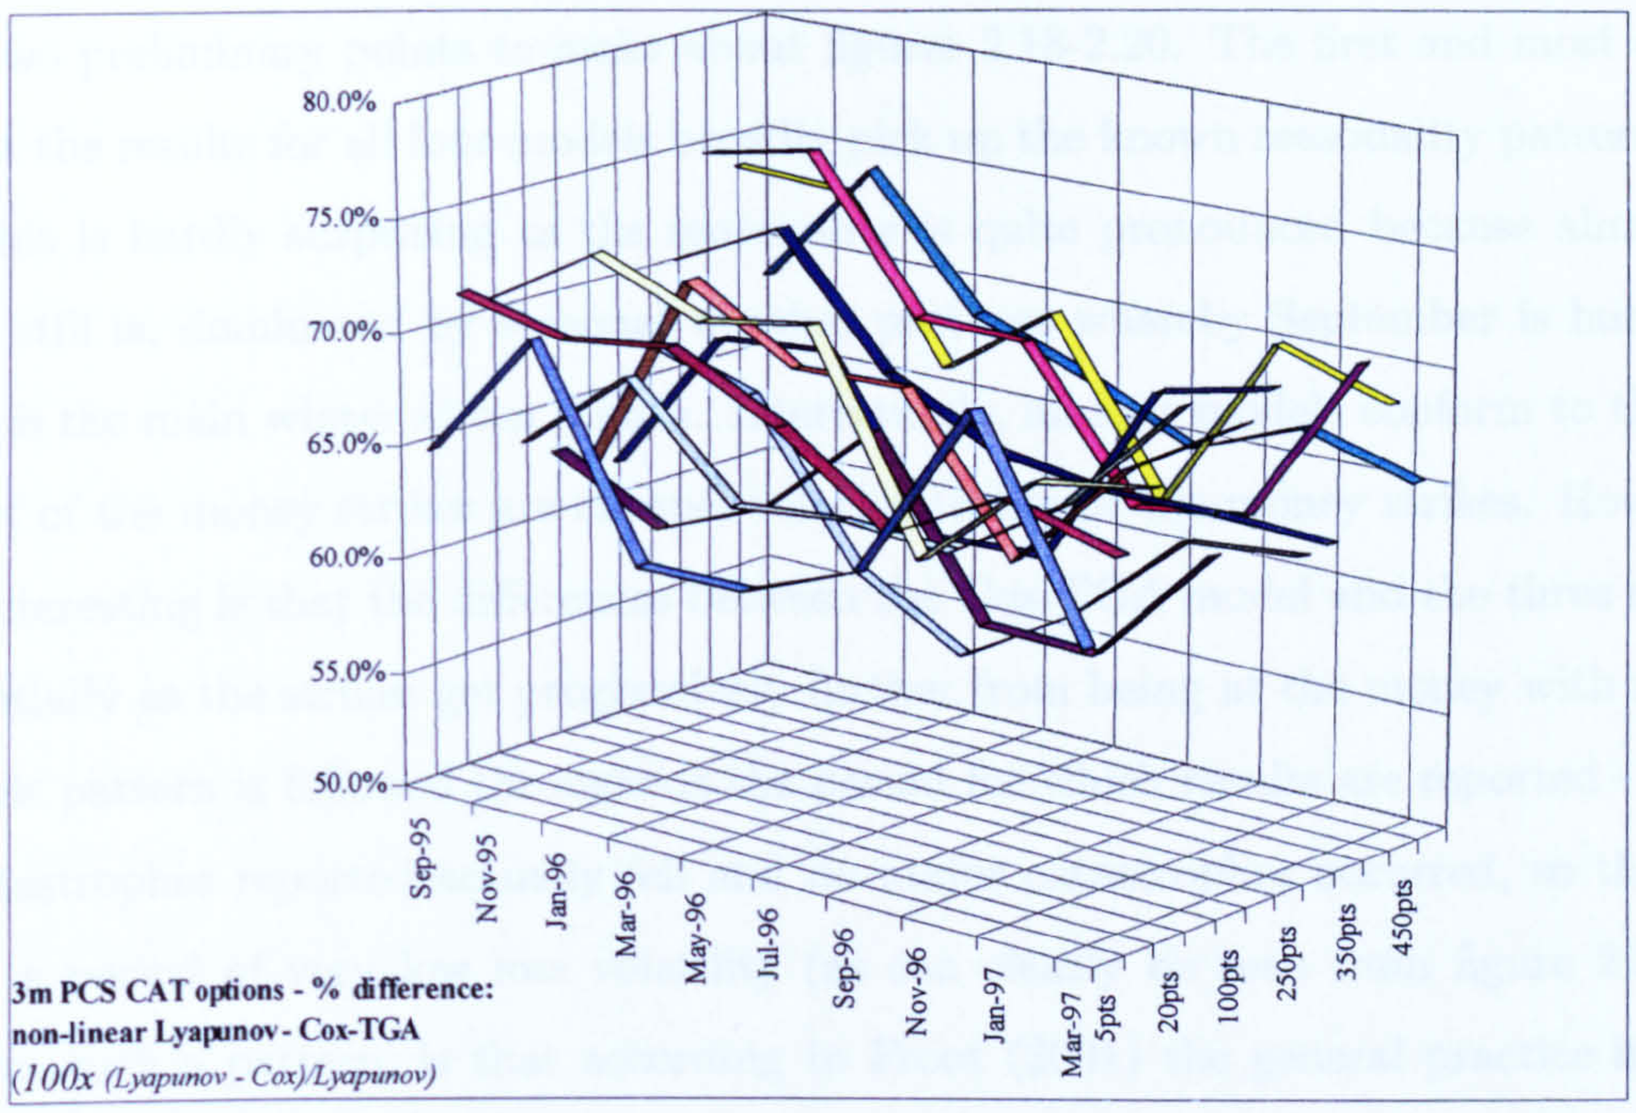

Figure 2.20: PCS CAT option valuation: Cox-TGA v non-linear Lyapunov

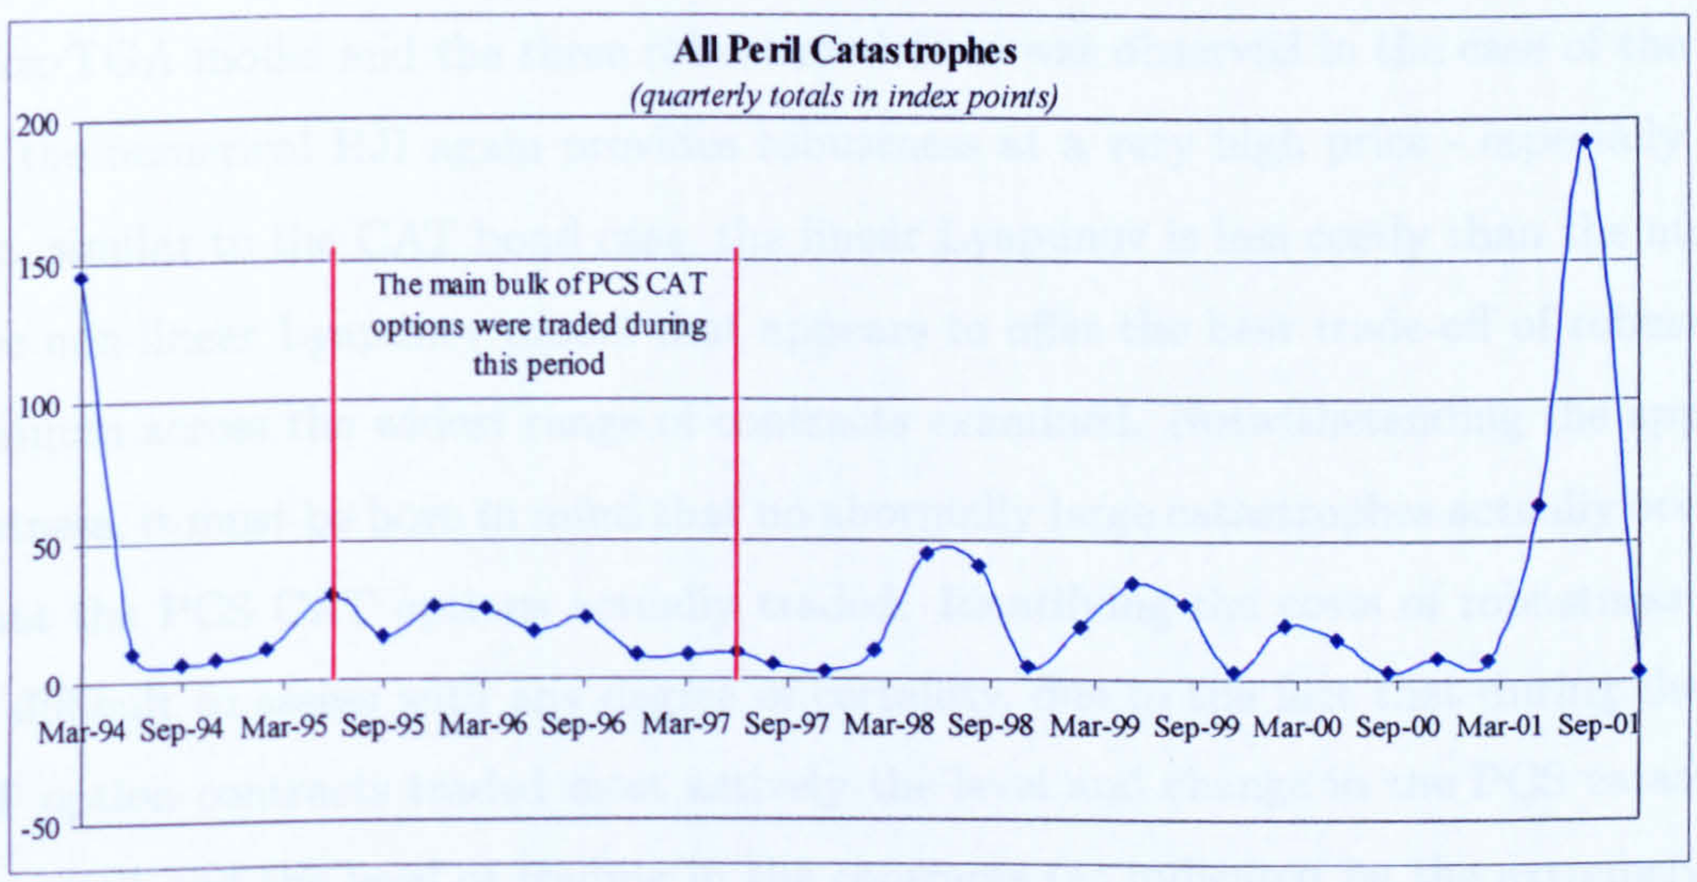

Figure 2.21: PCS catastrophes

There are two preliminary points to make about figures 2.18-2.20. The first and most obvious point to notice is that the results for all four models broadly pick up the known seasonality patterns in the PCS index data. This is hardly surprising as the seasonality is quite pronounced because almost the entire index was and still is, dominated by seasonal weather patterns whereby September is hurricane month and December is the main winter storm month. Interestingly, all four models conform to the usual prior that deeper out of the money strikes are cheaper relative to the at the money strikes. However, what is slightly more interesting is that the differences between the Cox-TGA model and the three robust models decline substantially as the strikes get progressively further from being at the money with respect to the index level. This pattern is followed throughout the period for which results are reported - during which the level of catastrophes reported actually fell and no major catastrophes occurred, so that the period was effectively a period of very low loss volatility (as can clearly be seen from figure 2.21). What is significant about such a pattern, is that according to Froot (2001) the general practice in insurance is that cover is rarely (if ever) taken for low probability but high cost events. This may help to explain why the differences between the Cox-TGA model and the robust models are proportionately much less in the case of the deep out of the money strikes.

The second point to make is that the same pattern of valuation differences can be observed between the non-robust Cox-TGA model and the three robust models as was observed in the case of the CAT bonds. Namely, that the numerical HJI again provides robustness at a very high price - especially at the lower strikes. Again, similar to the CAT bond case, the linear Lyapunov is less costly than the numerical HJI, but it was the non-linear Lyapunov model that appears to offer the best trade-off of robustness against increased premium across the widest range of contracts examined. Notwithstanding the apparently high price of robustness, it must be born in mind that no abormally large catastrophes actually occurred during the period that the PCS CAT options actually traded. Identifying the costs of robustness is made still further more difficult to assess with any degree of certainty, due to the fact that during the period that the PCS CAT option contracts traded most actively the level and change in the PCS catastrophe index was relatively small and the level of trading in the contracts (as indicated by the extremely low level of open interest in all PCS CAT option contracts) remained very low (less than 600 trades occurred across all 53 possible contracts over a 3 year period).

Arguably a more interesting question is how, given that robustness comes at a high price compared with the traditional Cox-TGA type model, do the theoretical prices from the four models compare with the actual traded prices of the quoted contracts? Froot (2001) and Sun (2002) have both examined this issue and found that the theoretical prices from the traditional models have managed to account for no more than 60% of the observed exchange-traded market prices, irrespective of whether the comparison is

made before (when price was at its least volatile and the ratio of theoretical to observed was around 43%), during (when price was often extremely volatile) or after the loss period (i.e. even in the development period). Froot's work was also of considerable interest as it also compared the price of PCS CAT option cover with actual contemporaneous reinsurance costs. Unfortunately, the same data was not available for this research (in particular the Guy Carpenter data set containing costs and levels of reinsurance activity was not available), so that detailed comparison between Froot's results and the results in this thesis is not possible. However, this research did have access to the CBOT data so was able to make the comparison between theoretical and actual quoted option prices.

Figure 2.22 shows the comparative results for actual CBOT traded prices versus theoretical, model based valuations for the 50pt strike call-spread[22]. The results were calculated for as near as possible to the beginning of the pre-loss period. What is immediately clear is that the Cox-TGA model does indeed undervalue the PCS CAT option most of the time. It is interesting to note that the Cox-TGA does in fact produce overvaluations in a number of instances. This is in direct disagreement with Froot's and Sun's work. The most obvious explanation for this result is that their results were for comparisons made of prices for the worst point during the loss period when differences between market and theoretical valuations may have been at their greatest, but also most volatile. Performing the comparison for the pre-loss period, as explained above, provides a less distorted view of the comparative real-world performance of the models.

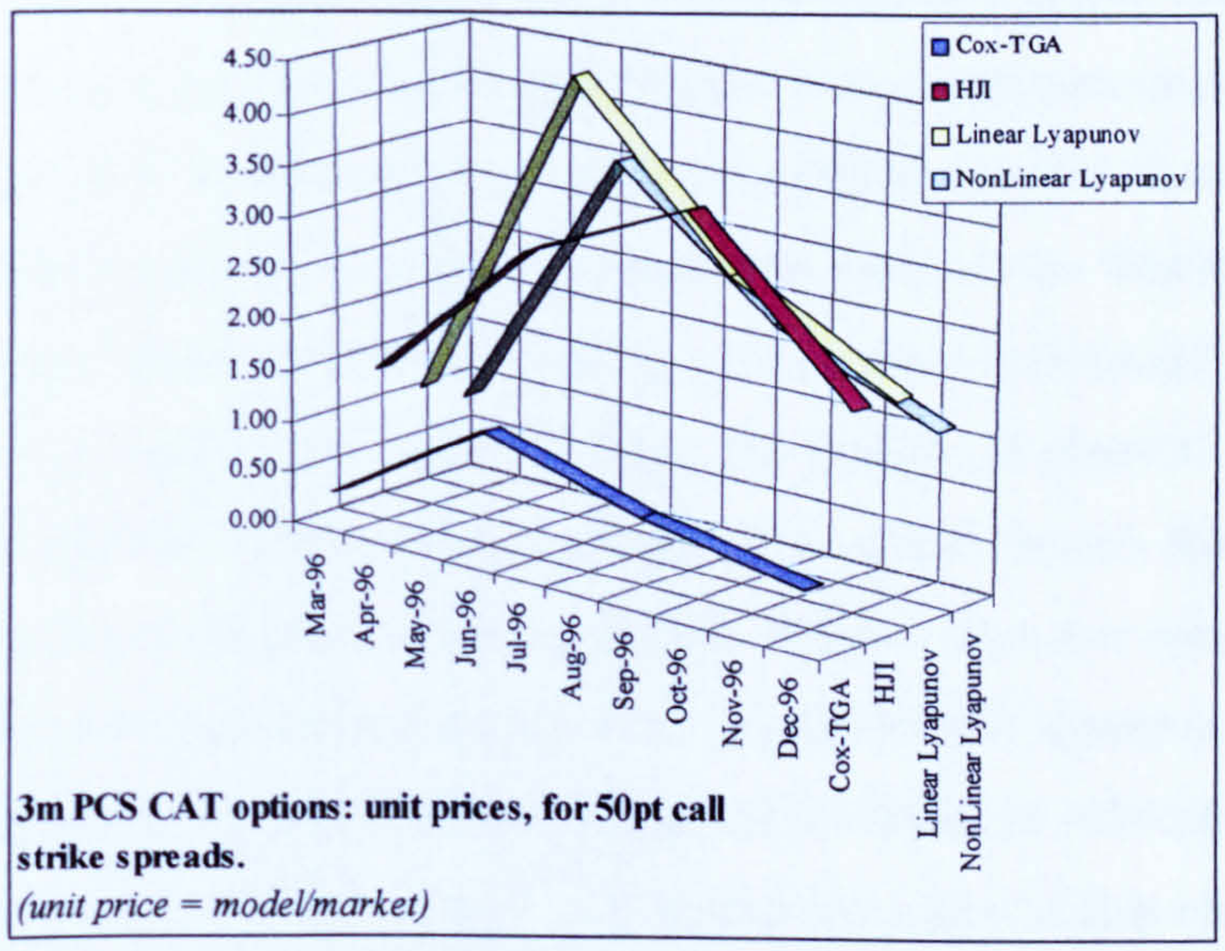

Figure 2.22: Unit price comparison - pre-loss period

[22]This comparison was selected as the 50pt strike was the most actively traded during the period under consideration.

However, what is also highly significant is that the robust models, without exception, produce unit prices greater than 1, implying over-valuation with respect to traded market prices. Once again, the numerical-HJI is the most expensive robust model, whilst the non-linear Lyapunov model valuations exhibit the least expensive robust alternative to the Cox-TGA model. Froot (2001) argued that supply and demand factors were the fundamental explanation for the inability of theoretical pricing to match open-market traded prices. Sun also investigated the use of alternative distributions in an attempt to address the issue of robustness, but as already explained earlier in this chapter, such an approach misses the key point that robustness requires more than simple insensitivity to distributional variations. In the light of the performance of the robust models, Froot's conclusions would therefore appear to be incomplete and the real explanation for underpricing would appear to be the non-robustness of the classical Cox-TGA type of approach. The results in this section therefore lend support to those of Sun (2002), by showing that, in line with the traditional reinsurance market, the actual observed market prices at which catastrophic-linked options traded were significantly higher than the theoretical prices and that the over-pricing pattern in the catastrophic-linked option market is strikingly similar to that found in the traditional reinsurance market. In addition, the results in this section go substantially beyond the work of Sun by indicating the extent to which even the higher market prices were substantially below robust prices.

It is worth bearing in mind that the cost of robustness appears excessive in all cases - the numerical-HJI in particular. However, it should also be remembered that the cost of ensuring robustness against all possible shocks is bound to be extremely high in the absence of a major catastrophic event (as the performance of the four models clearly shows during a period when there were no significant catastrophes), but could possibly be far more acceptable, if there is a sustained upward trend in catastrophic events such as hurricanes and earthquakes. How acts of terrorism such as the World Trade centre bombing in 2001 affect the valuation of catastrophe options is a much more interesting question. Valuing such binary events is far more complex than simply valuing the impact of physical phenomena. According to Hanson (2005), issues such as combinatorics, manipulation, moral hazard, hiding prices, and decision selection bias would all need to be factored into a model. Hanson does not consider any of these issues insurmountable and views terrorism futures as practical though morally questionable[23]. Notwithstanding this, given that such diverse factors have not formed part of the frame of reference of this research, it was concluded that the valuation of such binary acts fall beyond the scope of this research.

One final speculative theoretical point concerns Epstein and Wang's (1994) work on intertemporal asset pricing under Knightian uncertainty referred to in chapter 1. In their paper, Epstein and Wang show

[23] Without descending into the moral issues of the debate, the interested reader should pause for thought and read Stiglitz's (2003) Los Angeles Times article where he expresses his moral outrage and intellectual horror at the poor grasp of basic economics on the part of those suggeting that such a market could or should be developed.

that the model with uncertainty may lead to indeterminate equilibria with associated extreme volatility levels as the agents search in vain for a unique price to clear the market. When considered in conjunction with the relative scarcity of catastrophe data which means that those trading PCS CAT options might not have unique priors on the loss distributions or loss severity, this point may help to explain the highly volatile option values for the very low strikes. Such an effect would be most pronounced during the loss period, when uncertainty about the possible occurrence of catastrophes would be at its most influential, remaining high (though likely to be lower than during the loss period) in the development period as uncertainty continues to be resolved as claims on any catastrophes that have occurred during the loss period come to fruition.

Finally, given that the non-linear Lyapunov approach is robust and stable in operation, would it be practical to use in a real world trading environment? Table 2.2 provides representative timing information for the case of a 10 year CAT bond. Based on daily gridding of the state space, the model generated 3,653 steps, which yielded a valuation result in just under 2 seconds, which indicates that as far as speed of execution is concerned, the use of the pointwise min-norm control laws is certainly comparable with other widely used option valuation approaches using techniques such as trinomial trees[24].

Valuation	Execution Time Per Valuation	Time Steps
88.792168659	24.87687600	100,000
88.792168659	12.45350000	50,000
88.792168659	7.62625000	10,000
88.792168659	3.18196000	5,000
88.792168534	1.89474600	3,653
88.827452698	0.78785000	1,000
88.874878529	0.38767000	500
88.912989688	0.07876500	100
89.018273854	0.04746400	50
89.687757877	0.00123457	10

Table 2.2: Time Steps v Execution Time v Accuracy (non-linear Lyapunov)

2.6 Conclusions

Consistent with uncertainty as the central theme of this thesis, the work in this chapter has investigated the concept of dealing with uncertainty in the valuation of derivatives whose underlying are catastrophic events. The work has added a number of new ideas and some preliminary results to the subject of valuing catastrophe derivatives, focusing in particular on CAT bonds and CAT options. The key ideas developed are the application of robust control Lyapunov functions to deal with the difficulties associated with

[24] All timings were based on calculations performed on a 2Mhz Pentium-4 personal computer with 512Mb of RAM, running the Microsoft Windows XP operating system, code written in C++ and compiled using version 5 of Microsoft Visual Studio.

valuing assets with convex payoff functions whose underlying variables are discontinuous, with particular reference to the non-linearities associated with the behaviour of the underlying variables. The central focus was the harnessing of the desirable properties of the non-linear robust control Lyapunov model, which was compared with a number of the existing models using both simple comparisons and actual catastrophe data.

The new approach synthesizes three main ideas. The first is the Dempster and Hutton (1999) treatment of valuing an option using mathematical programming techniques. The second is the work by Cox and Pedersen (2002) who develop equilibrium pricing theory to deal with the problem of creating a unique arbitrage free valuation technique and deal in particular with the need to be able to generate a hedging portfolio. The final and novel feature of the theoretical work is to synthesise these elements together to produce a tractable and flexible state-space valuation framework that is arbitrage free, produces robust and stable results and is efficient in execution. Preliminary empirical results indicate that the non-linear robust optimal control Lyapunov model produces more stable results in the face of uncertainty through the explicit modelling of feedback and incorporation of non-linearities, which enables it to exhibit greater robustness to discontinuous behaviour than the current popular double Cox -PIDE model.

Two factors are specifically excluded from the work on CAT bonds. Namely, moral hazard and basis risk. In other words, no consideration has been given to the default riskiness or moral hazard on the grounds of scope. These factors, together with other issues such as stochastic interest rates, constitute potentially fertile and interesting areas for future research. They could be incorporated relatively easily by relatively simple modifications to the basic state space approach and via the type, timing and impact of perturbations applied within the non-linear Lyapunov model. As far as the Lyapunov model is concerned, possible extensions include removing the matching uncertainty restriction, introducing robust back-stepping and investigating more fully the problems associated with measurement disturbances. Other potentially fertile areas include investigating the impact of removing the assumption of arbitrarily fast time variation where uncertain non-linearity may not depend explicitly on time and the consequent use of robust integral action.

The results on the PCS CAT options were mainly concerned with the fairly narrow issue of the relative performance of the robust models and non-robust model. As the PCS CAT option contract only traded on CBOT for a short space of time, any analysis of the pricing information has to be treated with great care. However, when valuations from both the non-robust and the robust models were compared with the market prices of traded catastrophe options from CBOT, it was found the robust models were substantially better at capturing the behaviour of observed option prices. It was argued that this is due to the Cox-TGA model being ill-suited to valuing options on low probability, high-loss catastrophic events,

even when using a heavy tailed distribution such as the Burr distribution.

Chapter 3

Robust control and optimal hedging

It were not best that we should all think alike; it is difference of opinion that makes horse races

Pudd'nhead Wilson, by Mark Twain, 1894

3.1 Introduction and motivation for research

The previous chapter examined the problem of robustness in the pricing of options on a particular type of underlying, but viewed robustness in a somewhat disconnected fashion as the pricing problem was considered in isolation from the real world setting where any typical trading or strategic position in an underlying or derivative instrument is generally accompanied by some form of control rule that is intended to at least maintain the value of the position in the underlying instrument, or to protect some target terminal value. Indeed, a number of factors have been suggested as possible explanations for the application of such control rules[1], which are collectively most often referred to as a hedge.

In the options world, the concept of delta hedging is extremely familiar, referring to a strategy that consists of establishing an options position whose value varies in accordance with changes in the price of the underlying, so that a profit or loss on the underlying position is offset by a loss or profit on the option position. According to classical Black-Scholes theory, an option will only be fully delta hedged when the hedge is constantly adjusted. However, as this is not practical (due mainly to the practical issue of trading costs), discrete hedging is generally the norm, so that hedging error inevitably results. For

[1]See for example Stulz (1984), Smith and Stulz (1985), Stulz (1990), Breeden and Viswanathan (1990), Lessard (1990), DeMarzo and Duffie (1992), Froot, Scharfstein and Stein (1993), etc

example, Bossaerts and Hillion (1994) find that the so called "volatility smile" observed in equity markets, could be attributable to the impossibility of hedging in continuous time instead of model misspecification.

In a more general sense, the idea behind the act of hedging is that it affords the hedger the opportunity to reduce the uncertainties associated with any or all of price, quantity or timing. In the case of a simple static futures-based hedge, for example, the hedge at any point in time is a function of the current futures price and the parameters of the forecast quantity and price equations. So, if the current futures price is an unbiased estimator of the cash price at maturity and the interest rate is zero, then the optimal hedge is myopic - that is the same solution will hold for both static and dynamic hedging strategies. However, this result no longer holds when the expectation at time t of the cash price at maturity, T, differs from the futures price at time t (where $t \leqslant T$). In addition, an expected increase or decrease in the amount hedged is consistent with the current futures price being either an upwardly or downwardly biased estimator of the cash price at maturity, with the sign of the expected change in the hedge depending on the magnitude of bias relative to the degree of risk aversion.

Construction of a hedging policy generally requires some formalisation of the hedging objectives. Hedging is carried out by both companies and individuals alike, but the reasons are often extremely diverse, ranging from a desire to hedge the adverse impact of exchange rate fluctuations, through to making financial planning easier by reducing the uncertainty surrounding future cash-flows. For example, at the naive end of the spectrum is the strategy of simply taking an equal and opposite position in a hedging instrument to that held in the underlying instrument. The most obvious problem with such a strategy is the uncertainty around the correlation between the two instruments. Hedging policies are therefore often more in terms of optimisation of some target, with a common target being the construction of the optimal hedge ratio.

This leads to the second stage of the hedging process, namely, developing a metric to measure or assess the effectiveness of the hedging policy. Such metrics are obviously linked to the nature of the hedging policy. For example, in the case of using exchange traded futures to hedge foreign currency cash flows, the appropriate metric might be the tracking error of the hedge expressed in zero-deviation form. A great deal of research has been devoted to examining hedge effectiveness - mainly in the area of static hedging strategies[2], but such static hedging may significantly under-perform a more dynamic or evolving hedging strategy. Dynamic hedging is distinguished from the static or passive variety by the fact that it requires the initial position to be adjusted periodically in the light of newly available information. The frequency of the re-adjustment process often depending on mundane but vital factors such as transaction costs and desired activity levels. Hitherto, research on dynamic hedging has tended to concentrate on

[2]For example, see work by Benninga (1997), and Derman, Deniz and Kani (1994) and the references contained therein.

myopic decision rules which would be optimal in a relatively restricted set of circumstances.

Static hedging is open loop: decision makers form a plan at the beginning of the decision period and do not then respond to subsequent observations/information . This is a static decision problem as all the decisions for the entire period are calculated simultaneously at the beginning of the period and then all further information is ignored. A modification to this open loop strategy is one that recalculates a new open loop plan for the remainder of the period once each new observation arrives. In this way decision makers would carry out only the first period's decision of each open loop sequence. This strategy is known as open loop feedback and uses the available observations but ignores subsequent feedback because at each stage, the plan formulated by the decision maker is open loop and assumes that no further feedback is possible. In contrast, closed loop uses current information to calculate the current decision but also recognises that the environment will be observed in the future.

The objective of the research presented in this chapter is therefore two-fold. First, to develop robust optimal hedging rules using robust optimal control techniques based on closed loop feedback concepts explained in chapter 1, but extending the robust optimal techniques to incorporate suspected non-linearities. Second, to examine the use and impact of such robust optimal hedging rules. The advantages of robust dynamic hedging rules stem from the already analysed ability to reduce uncertainty by incorporating current information through feedback. Robust optimal hedging rules are compared with a range of other simple static and dynamic but non-robust optimal hedging rules for benchmarking purposes. This chapter therefore proceeds as follows. The next section provides a review of the existing literature. The following section examines and develops various robust optimal hedging rules and also investigates the relevance of a number of hedge efficiency evaluation metrics. This is followed by the results of empirical research carried out for this thesis that applies the robust optimal rules developed in the previous section and uses various metrics to assess their impact in alternative hedging scenarios. The final section offers conclusions and suggestions for further research.

3.2 Review of the literature

In financial markets, hedging errors arise from three sources. The first arises from hedging at discrete points in time. This source of error occurs when the underlying model is based on the assumption of continuously rebalancing a hedge, whereas hedging can only be practically carried out at discrete points in time. The second source of hedging error arises due to a lack of precise knowledge of volatility - misestimation; whilst the third source of hedging error is attributable to selling an instrument at the wrong premium.

In an attempt to estimate the size and impact of hedging errors, much of the work on hedging

strategies has been carried out on products traded on the exchange traded derivatives markets, where product specification and valuation is far more homogeneous than is the case in the over-the-counter markets. This is mainly because exchange traded instruments such as futures are standardised contracts and are commonly used as hedging instruments due to their liquidity, narrow bid-offer spreads and almost total absence of credit risk.

From a high-level perspective, there have been two broad approaches to robustness in hedging. The first has involved generating relatively wide range of hedging strategies based on first and second moments of the underlying price distribution, such as calculating optimal hedge ratios using mean-variance analysis through to delta hedging suggested by options theory, of the price of the underlying instrument and then applying statistical techniques of increasing sophistication to examine the efficiency of the hedge ratio of interest. Interestingly, little or no interest has been shown in the stability or robustness properties of the hedging strategies. In contrast, the second approach - which has so far seen relatively little in the way of research - involves the use of optimisation techniques for calculating optimal control rules which are then used as hedging strategies. The control rule based strategies have then been compared with a range of simpler hedging strategies in order to assess efficiency. Virtually no research interest has been shown in the robustness and stability of the hedging rules. The major contributions of each broad approach will the examined in turn.

3.2.1 Statistically based hedging approaches

Using Markowitz mean-variance portfolio theory, Ederington (1979) provided one of the first studies that formally defined hedge effectiveness as the reduction in the variance of the value of a position hedged with futures. He defined the stated objective of a hedge as being to minimise the risk of a given position, which is measured in terms of the variance of the returns. Ederington's hedge efficiency measure is

$$e = \frac{\sigma_{SF}^2}{\sigma_s^2 \sigma_f^2} = \rho^2 \tag{3.1}$$

where σ_s^2 and σ_f^2 represent the subjective variances and covariances of the possible price change between the spot and forward prices. ρ^2 is therefore the population coefficient of determination between the change in the spot price and change in the futures price. In contrast, Howard and D'Antonio (1984) define hedging effectiveness as the ratio of the excess return per unit of risk of the optimal portfolio of the spot commodity and the futures instrument to the excess return per unit of risk of the portfolio containing the spot position alone

$$HE = \left(\frac{\theta}{\left[\frac{(r_s - i)}{\sigma_s} \right]} \right) \tag{3.2}$$

where θ is the excess return per unit of risk, r_s is the expected one-period return for the spot position, i is the risk-free return and σ_s is the standard deviation of the one-period return for the spot position.

Smith and Stulz (1985) examined the hedging problem more widely and developed a positive theory of the hedging behaviour for value-maximising companies, treating hedging as simply one part of a company's financing decisions. Smith and Stulz examine taxes, contracting costs and the impact of the company's hedging policy on investment decisions as explanatory variables in their attempt to answer three key questions:

- Why some firms hedge and some do not;
- Why firms hedge some risks and not others; and
- Why some firms hedge their accounting risk exposure whilst others hedge their economic value.

Both financial and non-financial firms routinely implement hedging policies to mitigate their exposure to changes in asset prices. However, while these policies may perform satisfactorily in the limited sense of hedging the exposure under consideration, they might also increase the overall likelihood of financial distress due to the liquidity risks that they create.

Duffie and Richardson (1991) examine the problem arising when a hedger is faced with a commitment in one asset and the opportunity to continuously trade futures contracts on another asset whose returns are correlated with those of the committed asset. Optimal futures trading strategies are presented in closed form for several mean-variance and quadratic objectives. In contrast, Jorion (1991) analyses the performance of a minimum variance hedge that is rebalanced periodically based on the arrival of new information and finds that the Sharpe ratio of a dynamic hedging strategy is significantly better than that of a static hedge. However, his minimum variance criterion ignores the influence of expectations about future changes in the value of the underlying instrument.

Grossman and Vila (1992) solve for the optimal dynamic trading strategy of an investor who faces a leverage constraint in the form of a limitation on the ability to borrow for the purposes of investing in a risky asset. They assume constant relative risk aversion and that the risky asset follows a geometric Brownian motion. In the absence of the leverage constraint, they find that the optimal strategy involves investing a fixed proportion of wealth in the risky asset. They show that in the presence of a leverage constraint, the optimal investment also involves a strategy of investing a fixed proportion of wealth in the risky asset when the leverage constraint is not binding. However, the two proportions are different, which Grossman and Vila claim reflects the extent to which the investor alters strategy even when the leverage constraint is not binding because of the possibility that the leverage constraint may become binding in the future. In contrast to Grossman and Vila, Steil (1992) applies an expected utility analysis

to derive contingent claims for hedging foreign exchange transaction exposures over the complete range of probabilities, as well as the optimal forward and option hedge alternatives.

Ghose and Kroner (1994) apply the common persistence in GARCH models based on the work by Bollerslev, Engle and Nelson (1994) for example, to the performance of hedging strategies in financial markets. Dynamic hedging, which they define as continual rebalancing of the hedging portfolio based on time varying volatility of spot and futures returns, is compared with what they term "constant hedging" where changing volatility is not incorporated. While one would expect dynamic hedging to improve upon constant hedging in reducing the risk associated with a portfolio, Ghose and Kroner in fact find that the gain can sometimes be quite small and that the relative performance of dynamic hedging strategies tends to vary across markets. They conjecture that a potential reason for the difference in performance of these two hedging strategies across assets could be as follows. Namely, that dynamic and constant hedging will be approximately equivalent in the long-run if: 1. there is common persistence in conditional variances of spot and futures returns and 2. certain parameter restrictions in the generating processes of spot and futures returns hold. Ghose and Kroner derive conditions under which dynamic hedging will perform significantly better, which is important because dynamic hedging can have a significantly higher cost compared to that of constant hedging.

Lence and Hayes (1994) and Lence (1995) examine the minimum variance hedging problem, considering explicitly the importance of estimation risk in setting a hedge strategy. Estimation risk is defined as arising when the population moments of the joint probability density function used in a decision problem are unknown. As the hedger does not know the population parameters of the densities for the spot and futures prices, it is therefore necessary to use the "noisy" sample estimates in forming a hedging strategy. The conventional approach in this situation is to use the sample estimates as if they were population parameters. But it is relatively easy to think of occasions where such sample estimates are likely to be very imprecise. A simple example would be a hedger with a very noisy estimate of a hedge ratio who considers aggressive shorting of the related futures contract as a hedging strategy. The obvious question is whether such hedgers are merely exchanging commodity price risk for estimation uncertainty ? Lence and Hayes (1994) utilise Bayesian approaches to provide explicit methods for managing the overall risk to the hedger that includes price and estimation exposures, finding that incorporating the estimation risk leads to significant changes in the minimum variance hedge. They highlight the importance of using Bayesian techniques to properly incorporate sample and prior information in reducing risk to the hedger. Although Lence and Hayes (1994) and Lence (1995) both point out the use of numerical Bayesian procedures for calculating optimal hedges, they instead specify some aspects of the problem on an a priori basis in order to be able to concentrate on some of the issues surrounding how prior information might impact upon

hedging.

de Jong, de Roon and Veld (1995) note that existing research on the hedging effectiveness of currency futures assumes that futures positions are continuously adjusted which is clearly unrealistic in practice. They examine the effectiveness for futures positions which are not adjusted during the hedge period based on an out-of-sample approach using three models to determine hedge effectiveness: Ederington (1979) minimum variance model, Fishburn (1977) $\alpha - t$ model (a model in which the disutility of loss is minimised) and the Howard and D'Antonio (1984, 1987) model based on the Sharpe ratio. They find that the Ederington and Fishburn measures yield a higher effectiveness than the unadjusted model, whereas for the Sharpe ratio model they find that both naively and model based hedge positions lead to a lower hedging effectiveness than unhedged positions. Using a slightly different version of the variance-minimisation problem to that developed by Lence (1995), Foster and Whiteman (1997) provide a technique for estimating the hedge ratio that more fully considers estimation risk and their approach enables them to handle a variety of specifications for the time series model relating spot and futures prices.

Pennings and Meulenberg (1997) examine the hedging efficiency of futures contracts, concentrating on the extent to which hedgers are able to reduce cash price risk by using futures contracts. They concentrate on hedge effectiveness, which they define in terms of portfolio return. Their work is unusual, if not unique, in so far as it focuses on the hedging problem from the perspective of the futures exchange. A hedging efficiency measure is produced that measures the distance between the actual and the perfect hedge. The measure divides the distance into a systematic part (which can be managed) and a random part (which cannot be managed). Similarly to the coefficient of variation, Pennings and Meulenberg define their futures trading risk measure ($FRTM$) as the square root of the futures trading risk relative to the net price for the hedger if an ideal futures contract is used

$$FRTM = \frac{\sqrt{E\left(FTR_{t+1}\right)^2}}{PF_t^1 - C} \tag{3.3}$$

where FTR[3] is futures trading risk and

$$E(FTR) = B_{t+1} + DC_{t+1} \tag{3.8}$$

such that B_{t+1} is the basis of the futures contract and DC_{t+1} is the market depth cost when initiating the futures position, PF_t^1 is the price of the futures contract at the moment of hedge initialisation, C is the trading cost per futures contract, B_{t+1} is the basis of the futures contract and DC_{t+1} is the market depth cost when entering the futures position. An exactly analogous measure for cash price risk, $CPRM$ is defined as

$$CPRM = \frac{\sigma_{CP}}{E_t(CP)} = \frac{\sqrt{E\left(CP_t - \overline{CP}\right)^2}}{\overline{CP}} \tag{3.9}$$

where CP is the cash price and $\overline{CP}$ is the mean of the cash price. Combining $FRTM$ and $CPRM$ gives the Pennings and Meulenberg hedge efficiency measure

$$E = \frac{FRTM}{CPRM} \tag{3.10}$$

where $E \geqslant 0$. Upon substitution and simplification, this gives

$$E = \frac{\sqrt{E\left(FTR_{t+1}\right)^2}\overline{CP}}{\left(PF_t^1 - C\right)\sqrt{E\left(CP_t - \overline{CP}\right)^2}} = \frac{\left[\sqrt{\sigma_A^2 + \mu_A^2}\right]\overline{CP}}{\left(PF_t^1 - C\right)\sqrt{E\left(CP_t - \overline{CP}\right)^2}} \tag{3.11}$$

The intuition being that if E is less than unity, hedgers will reduce their risks because they exchange a larger cash price risk for a smaller futures trading risk. If futures trading risk rises compared to the cash price risk, then hedging efficiency falls. Similarly, if commission costs rise, hedging efficiency falls.

In his 1997 article on dynamic hedging in currency crisis, Krüger examines the hedging problem at a macro country level and presents results indicating that interest rate changes apparently have little effect on dynamic hedgers when volatility is high. This finding is in direct contrast to Garber and Spencer (1995, 1996) who argue that dynamic hedging may lead to perverse results when central banks try to

[3]The actual price realised upon settlement of a futures trade, ARP_{t+1} is given by

$$ARP_{t+1} = PF_t^1 - B_{t+1} - DC_{t+1} - C \tag{3.4}$$

where ARP_{t+1} is assumed to be a stochastic variable and so is comprised of an expected (or systematic) term and a random (or variance) term:

$$\mu_A = E(ARP_{t+1}) = PF_t^1 - C - E(B_{t+1} + DC_{t+1}) \tag{3.5}$$

$$\sigma_A^2 = E(ARP_{t+1} - \mu_A)^2 \tag{3.6}$$

so that

$$E(FRT^2) = \sigma_A^2 + \mu_A^2 \tag{3.7}$$

defend parities with higher interest rates. However, Krüger cites two reasons that this effect is likely to be small. First, in a period of currency crisis bid-ask spreads are likely to be widening, which would increase the cost of a dynamic hedging strategy, possibly making such an approach unusable for many market participants. Second, during a currency crisis, volatility (measured as conditional variance), is generally found to rise sharply. Krüger claims that this mitigates the effects of interest rate changes on the hedge ratio. He also argues that the option market functions as a kind of buffer in the following way. If a speculator or an 'ordinary' hedger sells, say, one million francs in the spot market or the forward market, he creates an excess supply of one million francs.3 But if he buys options with a notional value of one million francs and if the seller of the option uses a dynamic hedging strategy, less excess supply will be created - only delta times one million francs. Thus, the market makers in the option market actually keep away some of the pressure on central bank reserves. Therefore, on the whole, the notion that dynamic hedging impairs the use of the traditional 'interest rate weapon' does not seem warranted.

Frey and Stremme (1997) analyze the manner in which the demand generated by dynamic hedging strategies affects the equilibrium price of the underlying asset and derive an explicit expression for the transformation of market volatility under the impact of such strategies. They find that volatility increases and becomes time and price dependent, but the strength of these effects is found to depend both on the share of total demand that is due to hedging as well as (significantly) on the heterogeneity of the distribution of hedged payoffs. Wojakowski (1998) examines the problem of constructing an optimal hedge of the intertemporal long-term exchange rate using stochastic optimal control within a framework where the exchange rate follows a parity-reverting Gaussian process. The main finding is that firms should hedge more if the level of the exchange rate is above parity and less if it is below.

Dudenhausen and Schlogl (1999) examine the effect of model and parameter misspecification on the effectiveness of Gaussian hedging strategies for derivative financial instruments and show that Gaussian hedges in the "natural" hedging instruments are particularly robust. This is true for all models that imply Black/Scholes type formulas for option prices and hedging strategies. In this paper we focus on the hedging of fixed income derivatives and show how to apply these results both within the framework of Gaussian term structure models as well as the increasingly popular market models where the prices for caplets and swaptions are given by the corresponding Black formulas. By explicitly considering the behaviour of the hedging strategy under misspecification we also derive the El Karoui, Jeanblanc-Picque and Shreve (1995, 1998) and Avellaneda, Levy and Paras (1995) result that a superhedge is obtained in the Black/Scholes model if the mis-specified volatility dominates the true volatility. Furthermore, we show that the robustness and superhedging result do not hold if the natural hedging instruments are unavailable. In this case, we study criteria for the optimal choice from the instruments that are available.

Laurant and Pham (1999) consider the mean-variance hedging problem when asset prices follow Itô processes in an incomplete market framework. The hedging numeraire and the variance-optimal martingale measure appear to be a key tool for characterising the optimal hedging strategy. Their paper examines the hedging numeraire $\tilde{a}$ and the variance-optimal martingale measure $\tilde{P}$ using dynamic programming techniques and obtain explicit characterisations of $\tilde{a}$ and $\tilde{P}$ in terms of the value function of a stochastic control problem. When applied to a stochastic volatility problem, they derive an explicit form of the value function and then of the hedging numeraire and the variance optimal martingale measure, which provides explicit computational methods for calculating the optimal hedging strategies for the mean-variance hedging problem within a stochastic volatility model.

Dacorogna (2001) et al use high frequency data to study the problem of the hedge ratio and the neutral point in the case of currency hedging[4]. They pose the hedging problem as one of needing to compute the efficient frontier by optimising the return given the risk or optimising the risk given the return as shown in figure ??.

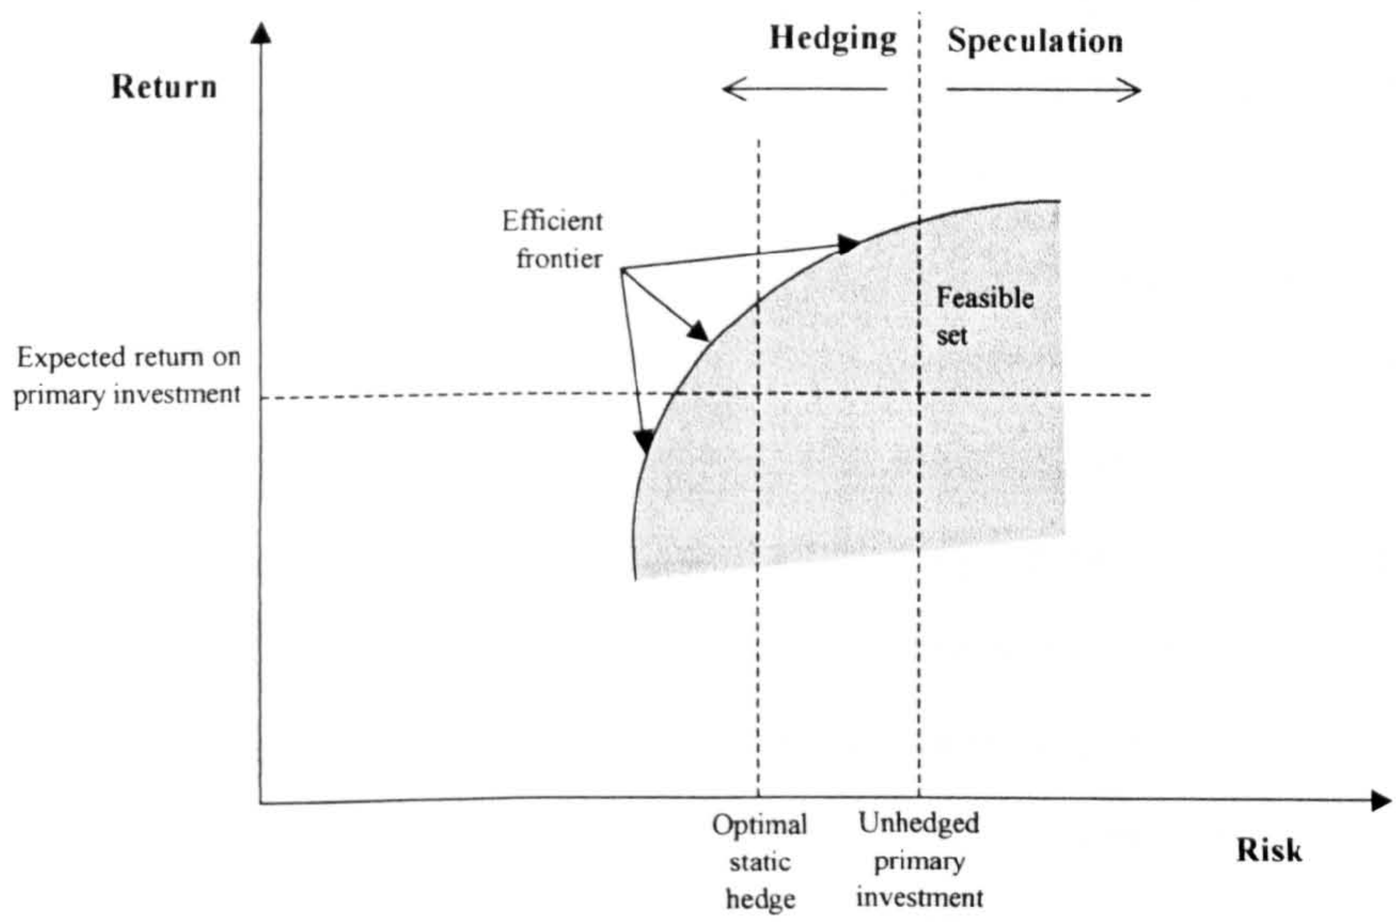

Figure 3.1: Mean-variance hedging

Rustem and Howe (2002) examine the problem of how minimax can provide a robust hedging strategy for written call options by formulating a minimax strategy that minimises the effect of a pre-defined worst case scenario, mainly in terms of bounds on the underlying source of uncertainty - i.e. the future price of the asset that underlies the option. They define robustness in two parts - first, as producing performance that is better than delta hedging for the set of options for which their model is designed and second that it does not perform worse than delta hedging. Their results suggest that minimax is particularly useful

[4]Currency hedging means the following: in the case of an investor that has purchased a foreign currency denominated asset, such as a dollar equity, holding a short position of size $-s_h$ in the foreign currency in order to minimise the volatility of the value of the total position due to foreign exchange rate fluctuations.

in the case of hedging the risk of writing an option when the price of the underlying stock is both highly volatile and crosses over the exercise price several times.

Agliardi and Andergassen (2002) study the destabilising effect of dynamic hedging strategies on the price of the underlying in the presence of sunk costs of transaction. Once sunk costs of transaction are taken into account, continuous portfolio rehedging is no longer an optimal strategy. Using a non-optimising (local in time) strategy for portfolio rebalancing, explicit dynamics for the price of the underlying are derived, focusing in particular on the excess volatility and feedback effects of these portfolio insurance strategies. Further, we show how these latter depend on the heterogeneity of the insured payoffs. Finally, conditions are derived under which it may still be reasonable, from a practical viewpoint, to implement Black - Scholes strategies.

Deep (2002) examines the case of hedging price risk using derivative contracts that are marked to market (such as futures contracts) and hence subject to margin calls. Deep shows that liquidity risk arising from margin calls on futures positions, can be a significant source of risk, possibly leading to financial distress and ruin, despite the fact that a firm remains "hedged". Deep argues that this type of risk should therefore be taken into account in the formulation of an optimal hedging policy. His paper derives a possible dynamic hedging strategy for a firm using futures contracts to hedge a spot market exposure. The risk emanating from the margin requirement on futures contracts is incorporated into the hedging decision by restricting the borrowing capacity of the firm. It is shown that this leads to a substantial reduction in the firm's optimal hedge, especially if the hedging horizon is long. The results provide some support for the low level of hedging observed empirically.

Deep's 2002 article on optimal dynamic hedging using futures under a borrowing constraint examines both financial and non-financial firms routinely implement hedging policies to mitigate their exposure to changes in asset prices. However, while these policies may perform satisfactorily in the limited sense of hedging the exposure under consideration, they might increase the overall likelihood of financial distress due to the liquidity risks that they create. This paper examines the case of hedging price risk using derivative contracts that are marked to market (such as futures contracts) and hence subject to margin calls. It is shown that liquidity risk, stemming from the need to meet margin calls on the futures position, can be a significant source of risk and can even lead to financial distress even though the firm remains "hedged". Such risks should therefore be taken into account in the formulation of an optimal hedging policy. This paper derives the dynamic hedging strategy of a firm that uses futures contracts to hedge a spot market exposure. The risk emanating from the margin requirement on futures contracts is incorporated into the hedging decision by restricting the borrowing capacity of the firm. It is shown that this leads to a substantial reduction in the firm's optimal hedge, especially if the hedging horizon is long. The

results provide theoretical support for the low level of hedging observed empirically.

De Losso and Bueno (2003) study the dynamic hedging problem using three different utility specifications: stochastic differential utility, terminal wealth utility and a new utility transformation which includes features from the two previous approaches. In all three cases, we assume Markovian prices. While stochastic differential utility (SDU) has an ambiguous effect on the pure hedging demand, it does decrease the pure speculative demand, because risk aversion increases. We also show that in this case the consumption decision is, in some sense, independent of the hedging decision. In the case of terminal wealth utility (TWU), we derive a general and compact hedging formula which nests as special cases all of the models studied in Duffie and Jackson (1990). In the case of the new utility transformation we find a compact formula for hedging which encompasses the terminal wealth utility framework as a special case; we then show that this specification does not affect the pure hedging demand. In addition, with CRRA- and CARA-type utilities the risk aversion increases and consequently the pure speculative demand decreases. If futures prices are martingales, then the transformation plays no role in determining the hedging allocation. Our results hold for a number of different price distributions. We also use semigroup techniques to derive the relevant Bellman equation for each case.

Moosa (2003) investigates the effect of the choice of the model used to estimate the hedge ratio on the effectiveness of futures and cross-currency hedging using data from the stock and foreign exchange markets. Four different models are used for this purpose to estimate the hedge ratio. The results show that model specification has little effect on the hedging effectiveness. It seems that what matters most is the correlation between the prices of the unhedged position and the hedging instrument. Results obtained by Moosa (2002) show that for an effective hedge, the correlation coefficient between the underlying and the hedge must be at least 0.50 to produce variance reduction of about 25 per cent. Lien (1996) argues that the estimation of the hedge ratio and the hedging effectiveness may change sharply when the possibility of cointegration between prices is ignored. In Lien and Luo (1994) it is shown that although GARCH may characterise the price behaviour, the cointegration relationship is the only truly indispensable component when comparing the ex post performance of various hedging strategies. Ghosh (1993) concluded that a smaller than optimal futures position is undertaken when the cointegration relationship is unduly ignored. Ghosh attributed the under-hedge results to model misspecification. Lien (1996) provides a theoretical analysis of this conjecture by assuming a cointegrating relationship of the form $f_t = p_{A,t} - p_{U,t}$, which is a simplified error correction model, implying that prices adjust in response to disequilibrium.

Valiani (2004) examines the portfolio decision problem for global investors as a joint choice problem over the financial assets and the relevant currencies. His paper investigates the currency risk hedging when volatilities and correlations of forward currency contracts and underlying asset returns are all time

dependent. Valiani uses a multivariate GARCH model with time-varying correlations to fit the dynamic structure of the conditional volatilities and correlations. The dynamic conditional risk-minimizing model is estimated for different hedge strategies and considers different international portfolios for the time period of January 1985 till December 2002. His empirical results show that the optimal dynamic hedge strategy using multivariate GARCH method can reasonably capture the currency fluctuations and significantly reduce the currency exposure risks and enhance the risk-adjusted performance of the international portfolios.

Ilhan and Sircar (2004) study optimal hedging of barrier options using a combination of a static position in vanilla options and dynamic trading of the underlying asset. The problem reduces to computing the Fenchel-Legendre transform of the utility-indifference price as a function of the number of vanilla options used to hedge. Using the well-known duality between exponential utility and relative entropy, we provide a new characterization of the indifference price in terms of the minimal entropy measure, and give conditions guaranteeing differentiability and strict convexity in the hedging quantity, and hence a unique solution to the hedging problem. We discuss computational approaches within the context of Markovian stochastic volatility models.

3.2.2 Optimal control based hedging approaches

In contrast to the statistical based approaches, Rustem and Howe (2002) use the minimax approach to examine optimal hedge construction in the presence of uncertainty when specifically faced with the worst case scenario. Worst-case risk optimality seeks to find the best possible outcome in the face of the worst possible outcomes. The objective function is usually expressed in terms of some form of cost or penalty function, so that stated more formally, worst-case hedging seeks to simultaneously determine the minimum of the cost function under the maximum or worst case scenario - hence the term minimax. Optimality is therefore expressed over all possible values of uncertainty. This is an important point as it distinguishes the minimax approach from the highly parametric, statistical approaches described in the previous section that are based on the assumption of some underlying distribution and which are only concerned with performance against some arbitrarily selected limit. In contrast, minimax weights all outcomes because of its enforced linearity, which means that it is forced to try to capture non-linear behaviour by over-compensating in the control rules it produces.

At a generic level, the maximum inner function is often couched in terms of a disutility or error function, such that the outer minimisation involves searching over the outcomes associated with the worst-case disutility scenario in order to find the best possible alternative. In mathematical terms, the

worst-case problem can be stated as

$$\min_{x \in \Re^n} \max_{y \in \mathcal{Y}} f(x, y) \tag{3.12}$$

where x is a vector of decision variables (represented by real numbers in n-dimensional Euclidean space, $\Re^n$) and y is a vector of uncertain variables defined over the feasible set $\mathcal{Y}$, with the solution being either discrete or continuous depending on whether $\mathcal{Y}$ is a discrete or continuous set. An equivalent, slightly more convenient representation is

$$\min_{x \in \Re^n} \Phi(x) \tag{3.13}$$

where

$$\Phi(x) = \max_{y \in \mathcal{Y}} f(x, y) \tag{3.14}$$

So that for the solution x^*

$$\Phi(x^*) = \max_{y \in \mathcal{Y}} f(x^*, y) \geqslant f(x^*, y), \quad \forall_y \in \mathcal{Y} \tag{3.15}$$

which states that the performance of the solution x^* is guaranteed to be non-inferior for any y. This is the specific feature that provides robustness from the minimax solution and ensures that performance will be better if the worst-case scenario is not realised.

From a financial risk management perspective, minimax has been employed in two main ways, namely, in discrete form as a robust strategy for discrete rival scenarios and in continuous form in problems such as option hedging. Analysing the former application first offers a slight advantage as it provides a framework for dealing with a discrete set of possible scenarios. Minimax thus arises from the ability to reduce the set of alternatives to single possibility, such that optimality is not determined by a single scenario, but simultaneously over all scenarios. Work by Rustem (1987, 1994) on policy optimisation examined the pooling of objective functions from rival models to generate an optimal policy based solely on a single model and then evaluates its impact if the second model proves to be the correct representation of the underlying system.

Larry Karp has authored three main papers which go beyond the minimax approach by applying control theory to the problem of optimal hedging. In the first paper, Karp (1985) examines a generalisation of the linear quadratic Gaussian control problem that provides a family of control rules which result in different combinations of moments of the quadratic payoff. He provides a recursive formula for calculating the second moment, using a dynamic optimal tariff to illustrate the method. Karp (1987) further develops his earlier work by formulating and solving a dynamic hedging problem with stochastic production. The optimal feedback rules that his model produces recognise that future hedges will be chosen optimally

based on the most current information. Karp (1988) examines a generalisation of the linear quadratic Gaussian problem using additive noise and the constant absolute risk aversion utility function, providing an explicit solution by using a limiting form of the discrete time linear exponential Gaussian control problem.

Karp studies the hedging problem facing a farmer that is assumed to exhibit constant absolute risk aversion wishing to create an optimal hedging strategy that will maximise the utility of future terminal wealth. Karp's (1988) work is a generalisation of a static problem used by Bray (1981), in which the discrete time dynamic problems is posed as a variation on the linear exponential Gaussian control problem first solved by Jacobson (1973). Karp's work differs from that of Bray in that it is the profit function which is exponentiated (linearly in the control - which is the hedge). The continuous time stochastic control model that arises from Karp's formulation of the optimal hedging problem produces a non-standard control problem, but as the discrete version remains standard and solvable, Karp therefore studies the limiting case where the interval between hedging opportunities goes to zero. Karp uses an expansion of the discrete time dynamic programming equation that assumes that the value function and its derivatives are continuous as ε (the time interval between hedging opportunities) approaches zero.

Karp's model works as follows. Suppose that the decision maker is a farmer growing a crop for which futures are available. The farmer's total production is therefore uncertain being subject to the usual forces of nature. There are assumed to be $n+1$ trading dates occurring as regular intervals of ε. Further assume that futures are first traded at time 0, then at $n\varepsilon = T$ the futures position is closed and the underlying asset is sold on the spot market. At each trading date the farmer decides the number of futures contracts to hold based on his current information about prices and the future harvest, where the time of harvest is assumed to not coincide with the maturity of the futures contract being used as the hedge. Karp models this imperfect time hedge by allowing the cash-futures basis to be a random variable. Given an instantaneous interest rate, r, then the discount factor for any period is $e^{-r\varepsilon}$. So that if p is the futures price, b the cash-futures basis and f the sales of futures contracts[5], h_T the harvest then the discounted stream of future revenue is given by

$$\pi_t = \sum_{i=0}^{j} \beta^{i+1} \left[p_{t+i\varepsilon} - p_{t+(i+1)\varepsilon}\right] f_{t+i\varepsilon} + \beta^{j+1} \left(p_T - b_T\right) h_T \tag{3.16}$$

with $j = n - 1 - \frac{t}{\varepsilon}$. If $h_{i\varepsilon}$ is the farmer's forecast at time $i\varepsilon$ of harvest at time T, then if h, p and b obey

[5] On page 624 of his 1988 International Econimic Review article, Karp defines f as sales of futures contracts, such that $f > 0$ implies that the farmer takes a short position. So in a sense f is actually the "stock" or net position as sales can take both positive and negative values in Karp's model.

the following stochastic difference equations

$$p_{(i+1)\varepsilon} = c(\varepsilon)\, p_{i\varepsilon} + \Delta_{1,i\varepsilon} \tag{3.17}$$

$$h_{(i+1)\varepsilon} = h_{i\varepsilon} + \Delta_{2,i\varepsilon} \tag{3.18}$$

$$b_{(i+1)\varepsilon} = b_{i\varepsilon} + \Delta_{3,i\varepsilon} \tag{3.19}$$

where $\Delta_{i\varepsilon} = (\Delta_{1,i\varepsilon}, \Delta_{2,i\varepsilon}, \Delta_{3,i\varepsilon})' \sim i.i.d \quad N(0, \Sigma\varepsilon)$ and $c(\varepsilon) = e^{a\varepsilon}$. The drift parameter a and the variance matrix Σ are time independent. Equation 3.17 allows the current futures price to be a biased estimator of the next period futures price, whereas equation 3.19 is an unbiased estimator of basis in the next period. Solving the Karp model yields the following expression for the optimal hedge

$$f_t = G_t y_t \tag{3.20}$$

$$\text{and} \tag{3.21}$$

$$y_t = (p_t, h_t, b_t)' \tag{3.22}$$

where

$$G_t = (g_{1,t}, g_{2,t}) \tag{3.23}$$

where

$$g_{1,t} = \left[\frac{a^2}{k} - k\left(\sigma^2 - \rho^2\right) - e^{-r\tau}\right]\left(\frac{1 - e^{-r\tau}}{r}\right) + \rho e^{-r\tau}(1 - 2a\tau) - \frac{a}{k} \tag{3.24}$$

$$g_{2,t} = e^{-r\tau} \tag{3.25}$$

[6].

The key criticism of the Karp approach, as in the case of chapter 2 on option pricing, is that the calculation of the optimal hedge is based on the implicit use of the H_2 norm. As already explained, this norm behaves stably and reliably, though not necessarily robustly, even in the presence of non-extreme market moves. However, as has already been indicated in chapter 2, use of the H_2 norm is neither necessary nor sufficient to guarantee robustness of the optimal hedge rules.

Neuberger and Hodges (2002) do not follow the control theory approach, but instead explore hedging strategies based on no-arbitrage bounds, that are model independent. In particular, they determine the bounds on the price of a general barrier option given the price of a set of European call options and

[6] Where the variables have the following meanings: $\tau = t - T$, a is a drift parameter, k is the coefficient of constant risk aversion, ρ is a covariance term from $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & \sigma^2 \end{pmatrix}$

identify the hedging strategy that enforces the bounds. The strategy puts a floor on the maximum loss that can be incurred by the writer of the barrier option. They show how the strategy can be made dynamic and the floor raised over time. The distribution of hedge errors under the strategy is compared with that under alternative strategies. The risk management of complex derivatives poses a particular challenge, because while there are often many closely related instruments that are liquid enough to be used for hedging, standard models are not correctly designed to provide either efficient or effective hedge strategies. Neuberger and Hodges show how to exploit the properties of rational bounds in order to design strategies that make use of the range of instruments available and are robust to model specification and model estimation error. The robust hedging strategy based on rational bounds puts a firm floor under possible losses. This is particularly attractive in the context of capital adequacy regulations that focus on the most unfavorable events. Traditional hedging strategies by contrast are less well equipped for the purpose since they are heavily model dependent and can generate heavy losses when model assumptions are violated.

According to Neuberger and Hodges, robust hedging strategies have another interesting feature, namely, that after the hedge is established, they require no trading in the derivatives market. This is attractive since in many derivative markets transaction costs are significant and liquidity is not assured. Traditional hedging strategies need to be rebalanced. They require most rebalancing after major market moves; this may be precisely the time model and market prices diverge most widely. By contrast, the robust hedger can choose to rebalance at such times if there is an opportunity to trade one bounding portfolio for a cheaper one, but has no need to do so. The cheapest super-replicating portfolio varies over time. By revising the portfolio periodically, the hedge performance can be greatly improved while still retaining a firm floor on the maximum size of loss. The choice of hedging strategy will necessarily depend on a multitude of factors: the instrument to be hedged, the available hedge instruments, the costs of transacting, the predictability of asset price dynamics, the preferences of the agent. Neuberger and Hodges present evidence to show that their robust hedge compares well with more conventional alternatives such as delta hedging and the Carr, Ellis and Gupta (1998) static option hedge[7].

It is interesting to note that Neuberger and Hodges' view robustness in the sense that downside risk is bounded whatever the path of prices. This view contrasts with previous work, such as Ahn, Muni and Swindle (1997) which adopted a weaker measure that aimed to maximise expected utility in the worst

[7]An interesting aside is contained in the paper by Hipp and Vogt (2003), who approach the hedging problem from the perspective of insurance, considering a risk process modelled as a compound Poisson process. They derive the optimal dynamic unlimited excess of loss reinsurance strategy that minimises the infinite time ruin probability and show the existence of a smooth solution of the corresponding Hamilton-Jacobi-Bellman equation. However, the focus of their work is aimed more squarely at the more macro-level risk management issues, so it will be discussed in chapter 4 on robust risk management.

case outcome from a restricted set of volatility scenarios. Neuberger and Hodges measure is stronger in so far as it can be viewed as a limiting case where a decision maker wishes to minimise loss in the face of the worst possible outcome without artificially restricting the range of outcomes as in the case of Avellanda, Levy and Paras (1995) who allow uncertain volatility but restrict it between fixed upper and lower bounds.

The Neuberger and Hodges approach still suffers from one of the classical robustness weaknesses, namely, that their procedure depends critically upon the choice of the price process and only guarantees the least upper bound on the price of the barrier option given a particular set of price processes that are consistent with the absence of arbitrage. If too small a set of price processes is selected, the upper bound on the price of the barrier option cannot be guaranteed to be an upper bound and the hedging strategy will not dominate the barrier option on all possible paths. If on the other hand too large a set of price processes is selected, the search is practically hard to implement.

3.3 Robust hedging with non-linear optimal control

This section has two objectives. First, to develop a hedging rule in the presence of uncertainty that is both robust and stable. Second, to provide a simple example to illustrate how such a rule works - as a precursor for the following section which contains a detailed empirical study to establish how robust rules perform in comparison with the type of rule discussed in the literature review.

The previous review of the literature dealt in detail with the issues surrounding the non-robustness of optimal hedging rules. This section develops a methodology for calculating robust optimal hedges based on robust optimal control theory. What are the required state variables in a model designed to provide robust and stable hedging decisions in the presence of uncertainty ? When pricing a standard option in a BS world of "well-behaved" underlying assets, it is relatively easy to see what variables should be included as state, and controller variables, as both Dempster, Rustem and Howe (2002), as well as Neuberger and Hodges (2002) have shown. In the more complex cases of options on extreme values such as catastrophes or defaults, there is the extra dimension of having to deal with outliers and their effects on pricing and hedging decisions, as well as the need for robustness in the face of worst-case scenarios - which will be different in each type of case (i.e. different worst-cases for "normal" options compared to catastrophe options).

3.3.1 The non-linear robust optimal control approach

At its simplest, robustness guarantees that the essential functions of a model are maintained under conditions so adverse that the model no longer reflects the behaviour encountered. However, when attempting

to guarantee robustness it is also critical that the resulting hedge rule produces a stable net position. Unstable behaviour could for example be produced in a volatile environment when a robust hedging rule might imply large and frequent readjustment of a hedge. Such behaviour may result in a lack of stability in the portfolio and cause a large loss of income due to excessive re-hedging costs. The objective of robustness must therefore be simultaneously stated with respect to stability - namely robust stabilisability.

As already explained in chapter one, a typical version of the robust optimal control approach involves the interconnection of a model G, a controller K and a source of disturbance or perturbation Δ to represent uncertainty. Once the underlying model has been devised, the robust control problem is to construct K such that closed-loop stability and performance are simultaneously guaranteed in the presence of every Δ belonging to a defined family $\mathcal{F}_\Delta$ of admissible controllers. This is the optimal control equivalent of the HRS model in that it is worst-case because the occurrence of all uncertainties in $\mathcal{F}_\Delta$ are deemed to be equally likely.

Much of robust control is linear in so far as it gives rise to results that have been calculated under the assumption that G is linear. Finite dimensionality and time invariance are also frequent assumptions. A diverse range of admissible controllers have been subjected to testing in this way, ranging from structured to unstructured uncertainties, linear and non-linear uncertainties, time invariant and time variant uncertainties, to name but a few. These different measures of the size of the uncertainty have produced different metrics of robust control such as H_∞ and $\mathcal{L}_1$ with all being based around some bound on a norm. But the important point is that when a model exhibits behaviour which indicates the existence of non-linearities, the classical linear approach responds by making the set $\mathcal{F}_\Delta$ large enough to deal with the suspected non-linearities because G is restricted to be linear. The key disadvantage of this course of action is that it disregards information about the non-linearities that may be available, with the result that the controllers that are generated are likely to be too restrictive - a result likely to be exacerbated the more significant are the non-linearities. The obvious way to get around this problem is to allow G to be non-linear and therefore search for robust non-linear controllers. One of the more tractable approaches to this problem is to use state-space Lyapunov stability which makes use of input-to-stability techniques to produce a robust non-linear optimal controller . The idea involves the construction of a robust control Lyapunov function which deals with both control and uncertainty inputs, as well as providing a generalisation of the output control Lyapunov function. The existence of a robust control Lyapunov function is both necessary and sufficient to ensure the solvability of the robust control problem.

Given the desire to avoid the impact of the worst case scenario, how should such a scenario be conceived of and should the objective be to minimise hedging error in the event that such a scenario occurs, as in HRS, or some other objective ? For example, if a hedging strategy fails, is ruin assumed

to occur instantaneously, or over a finite period of time ? If the latter, what should be the length of the "ruin period" ? Should ruin be deemed to have taken place based on the occurrence of one or more events ? How can this be captured within a model ? The simple and most common form of a dynamic system and linear objective function for robust optimal control is

$$\dot{x} = f(x, u, w) \tag{3.26}$$

$$J = \int_0^\infty L(x, u)dt \tag{3.27}$$

where J is the total cost, x is the state variable, u the controller and w the disturbance variable. This gives rise to the steady-state Hamilton-Jacobi-Isaacs partial differential equation

$$0 = \min\max\left[L(x, u) + \nabla V(x) \cdot f(x, u, w)\right] \tag{3.28}$$

where the value function $V(x)$ is the unknown value function. Solving such an equation in a non-linear form is only possible for the simplest of non-linear systems. However, for a careful choice of L in equation 3.28, it is possible to derive a positive definite solution $V(x)$ that will lead to a continuous state feedback control $u(x)$, that is optimal, stable and robust with respect to the disturbance w. Freeman and Kokotovic (1996) show that a known robust control Lyapunov function can be used to explicitly derive an optimal control law without needing to solve equation 3.28 by solving instead an inverse optimal robust stabilisation problem. In order for this to be possible, the objective or cost function must impose penalties (which are known to be non-restrictive in the particular case of a hedging problem) on both state and control functions. Optimality can therefore be viewed as a means of choosing a stabilising control law from among the universe of possible control laws containing the required properties. Using a recursive back-stepping method, Freeman and Kokotovic (FK) provide an extremely efficient procedure, so that if a meaningful cost function can be found such that a given robust control Lyapunov function is the corresponding value function, then the corresponding HJI equation will have been implicitly solved and the robust hedging rule can be calculated directly using a relatively simple formula. Their approach is thus to invert the usual robust control approach and instead use a back-stepping procedure to find a meaningful cost function such that a given robust control Lyapunov function is the corresponding value function. In the classic problem of minimising the hedging errors, this translates into a realistic and relatively simple objective function that has meaning and tractability across a wide range of hedging situations.

The key feature of FK's work is that if a robust control Lyapunov function is known, then it is

possible to construct a feedback law that is optimal with respect to a meaningful cost function, without having to solve the associated non-linear Hamilton-Jacobi-Isaacs cost equation. In fact as they point out, it is not even necessary to construct the cost function as the optimal feedback control law can be calculated directly from the robust control Lyapunov function without recourse to the Hamilton-Jacobi-Isaacs equation. FK provide a formula to generate a class of optimal control laws which involves only the robust control Lyapunov function, the system equations and required design parameters. They refer to their class of admissible laws as pointwise min-norm control laws. The next section provides a detailed derivation of such pointwise min-norm control rules for use as hedging strategies. The following section provides a simple example that compares the HJI robust control rule and a pointwise min-norm (PWMN) control rule. This is followed by detailed empirical work applying the robust PWMN control rule to a series of hedging scenarios and comparing its performance with that of a number of the more widely used hedging rules.

3.3.2 A pointwise min-norm control hedging rule

FK (1996) show that the existence of a robust control Lyapunov function is equivalent to robust stabilisability. Using the FK approach therefore means first constructing a robust control Lyapunov function, followed by deriving a robustly stabilising feedback controller to make the derivative of the Lyapunov function negative. Each of these two steps will now be considered in detail. This will finally lead in the following section to a formal statement of both elements for a robust hedging function.

Constructing a robust control Lyapunov function

The technique described in this section is that developed by FK (1996) which involves the construction of families of robust control Lyapunov functions and therefore the implicit derivation of stabilising control rules. The technique is known as robust backstepping and is based on the idea of using a known control Lyapunov function for a version of the system without uncertainties (known as the nominal system) as the robust control Lyapunov function for the uncertain system. The technique leads to the construction of a robust control Lyapunov function which is non-quadratic and has three desirable properties. Namely, high gain, softer control laws (resulting in less violent and frequent re-hedging) and faster computation. The non-quadratic robust control Lyapunov function in turn reduces the amount of effort required to subsequently make the Lyapunov derivative negative and thereby construct the stabilising feedback controller.

Begin by defining the following class of uncertain n^{th}-order, non-linear systems

$$\dot{x} = F(x,w) + G(x,w) + u \tag{3.29}$$

where F and G are continuous functions. Make the following assumptions, based on which, a robust control Lyapunov function will be constructed

- There is state feedback $(Y(x) = \{x\})$.
- There is a single unconstrained control input $(U(x) \equiv \mathcal{U} = \mathbb{R})$
- There is a disturbance constraint $W(x)$ that is independent of the control u.
- W is continuous with nonempty compact convex values.
- $W(x) \equiv B$, where B is constant.
- F is of the structural form

$$F(x,w) = \begin{bmatrix} \phi_{11}(x,w) & \phi_{12}(x,w) & 0 & \cdots & 0 \\ \phi_{21}(x,w) & \phi_{22}(x,w) & \phi_{23}(x,w) & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \phi_{n-1,1}(x,w) & \phi_{n-1,2}(x,w) & \phi_{n-1,3}(x,w) & \cdots & \phi_{n-1,n}(x,w) \\ \phi_{n1}(x,w) & \phi_{n2}(x,w) & \phi_{n3}(x,w) & \cdots & \phi_{nn}(x,w) \end{bmatrix} x + F(0,w) \tag{3.30}$$

  for continuous scalar functions ϕ_{ij}[8].

- G is of the structural form

$$G(x,w) = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \phi_{n,n+1}(x,w) \end{bmatrix} \tag{3.31}$$

  for continuous scalar functions ϕ_{ij}.

- Each function ϕ_{ij} depends only on w and the state components x_1 to x_i

$$\phi_{ij}(x,w) = \phi_{ij}(x_1, ..., x_i, w) \tag{3.32}$$

  for $1 \leqslant i \leqslant n$ and $1 \leqslant j \leqslant i+1$.

[8] The decomposition of F is not required to be unique.

- $\phi_{ij}(x_1, ..., x_i, w) \neq 0$, for all $x_1, ..., x_i \in \mathbb{R}$ for all $w \in B$ and for all $1 \leqslant i \leqslant n$ (this ensures that the above system is controllable for each fixed $w \in B$.

FK describe a system that satisfies the above conditions as being strict feedback form (also known as lower triangular form). This is important, because any system that can be transformed into strict feedback form will have a robust control Lyapunov function. A further advantage of this form is that growth restrictions are not required on the non-linearities. The construction of the robust control Lyapunov function now proceeds as follows.

The first step is to construct a diffeomorphism[9] on the state space $\mathcal{X}$ using smooth scalar functions $s_1(x)$, $s_2(x_1, x_2)$, ..., $s_{n-1}(x_1, ..., x_{n-1})$ using the method described below, such that each function s_i will depend only on the state components x_1 to x_i. When these functions have been selected, it is then possible to define a transformed state vector z

$$z_1 \ := x_1 \tag{3.33}$$

$$z_2 \ := x_2 - z_1 s_1(x_1) \tag{3.34}$$

$$z_3 \ := x_3 - z_2 s_2(x_1, x_2) \tag{3.35}$$

$$\vdots \tag{3.36}$$

$$z_n \ := x_n - z_{n-1} s_{n-1}(x_1, ..., x_{n-1}) \tag{3.37}$$

[9] A diffeomorphism is a differentiable mapping that has a differentiable inverse.

When expressed in matrix form, the diffeomorphism and its associated inverse are

$$z \;:\; = S(x)\,x = \begin{bmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ -s_1 & 1 & 0 & 0 & \cdots & 0 \\ s_1 s_2 & -s_2 & 1 & 0 & \cdots & 0 \\ -s_1 s_2 s_3 & s_2 s_3 & -s_3 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ \pm s_1 \cdots s_{n-1} & \mp s_2 \cdots s_{n-1} & \pm s_3 \cdots s_{n-1} & \cdots & -s_{n-1} & 1 \end{bmatrix} x \tag{3.38}$$

$$x \;=\; S^{-1}(x)\,z = \begin{bmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ s_1 & 1 & 0 & 0 & \cdots & 0 \\ 0 & s_2 & 1 & 0 & \cdots & 0 \\ 0 & 0 & s_3 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & s_{n-1} & 1 \end{bmatrix} z \tag{3.39}$$

where the signs in the last row of equation 3.38 depend on whether the dimension of the state space vector is even or odd. The functions s_i are constructed so that $V(x) := z^T z$ is a robust control Lyapunov function for the system. The procedure also facilitates the subsequent construction of a robustly stabilising state feedback control law.

The first task is to calculate $\dot{z}$ from equations 3.38 and 3.39 by taking the derivative of $z = S(x)\,x$

$$\dot{z} \;=\; \left[\frac{\partial S}{\partial x_1} x \frac{\partial S}{\partial x_2} x \cdots \frac{\partial S}{\partial x_n} x\right] \dot{x} + S(x)\,\dot{x} \tag{3.40}$$

$$:\; = T(x)\,\dot{x} \tag{3.41}$$

where $T(x)$ can be easily calculated from equation 3.38. Letting $\circledS_i$ represent any function depending only on the states x_1 to x_i and the functions s_1 to s_i and their partial derivatives, gives

$$\dot{z} = T(x)\,\dot{x} = \begin{bmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ \circledast_1 & 1 & 0 & 0 & \cdots & 0 \\ \circledast_2 & \circledast_2 & 1 & 0 & \cdots & 0 \\ \circledast_3 & \circledast_3 & \circledast_3 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \circledast_{n-1} & \circledast_{n-1} & \circledast_{n-1} & \cdots & \circledast_{n-1} & 0 \end{bmatrix} \dot{x} \tag{3.42}$$

The robustly stabilising state feedback control law is then of the form $u(x) = z_n s_n$, where $s_n(x)$ is a yet to be determined smooth function. So, given the choice for u, equation 3.39 can be used to re-write equation 3.29 as

$$\dot{x} = \begin{bmatrix} \phi_{11} & \phi_{12} & 0 & \cdots & 0 & 0 \\ \phi_{21} & \phi_{22} & \phi_{23} & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \phi_{n-1,1} & \phi_{n-1,2} & \phi_{n-1,3} & \cdots & \phi_{n-1,n} & 0 \\ \phi_{n1} & \phi_{n2} & \phi_{n3} & \cdots & \phi_{n,n} & \phi_{n,n+1} \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ s_1 & 1 & 0 & 0 & \cdots & 0 \\ 0 & s_2 & 1 & 0 & \cdots & 0 \\ 0 & 0 & s_3 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & s_{n-1} & 1 \\ 0 & 0 & 0 & \cdots & 0 & s_n \end{bmatrix} z + F(0,w) \tag{3.43}$$

Substituting equation 3.43 for $\dot{x}$ in equation 3.42 gives

$$\dot{z} = \begin{bmatrix} \phi_{11} + \phi_{12}s_1 & \phi_{12} & 0 & \cdots & 0 \\ \bigstar_1 & \bigstar_1 + \phi_{23}s_2 & \phi_{23} & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \bigstar_{n-2} & \bigstar_{n-2} & \cdots & \ddots & \phi_{n-1,n} \\ \bigstar_{n-1} & \bigstar_{n-1} & \bigstar_{n-1} & \cdots & \bigstar_{n-1} + \phi_{n,n+1}s_n \end{bmatrix} z + T(x)F(0,w) \tag{3.44}$$

$$= \begin{bmatrix} \phi_{11} & \phi_{12} & 0 & 0 & \cdots & 0 \\ \bigstar_1 & \bigstar_1 & \phi_{23} & 0 & \cdots & 0 \\ \bigstar_2 & \bigstar_2 & \bigstar_2 & \phi_{34} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ \bigstar_{n-2} & \bigstar_{n-2} & \bigstar_{n-2} & \cdots & \bigstar_{n-2} & \phi_{n-1,n} \\ \bigstar_{n-1} & \bigstar_{n-1} & \bigstar_{n-1} & \bigstar_{n-1} & \cdots & \bigstar_{n-1} \end{bmatrix} z + \tag{3.45}$$

$$\begin{bmatrix} \phi_{12}s_1 & 0 & 0 & \cdots & 0 \\ 0 & \phi_{23}s_2 & 0 & \cdots & 0 \\ 0 & 0 & \phi_{34}s_1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \phi_{n,n+1}s_n \end{bmatrix} z + T(x)F(0,w) \tag{3.46}$$

where $\bigstar_i$ indicates any function depending solely on w, the states x_1 to x_{i+1} and the functions s_1 to s_i and their partial derivatives[10]. Using the $A(x,w)$ and $D(x,w)$ to replace the first and second terms in

[10] A function of the type $\circledast$ is also of type $\bigstar$, but not necessarily vice versa. This is because a $\bigstar$ function is allowed to

equation ?? gives

$$\dot{z} = [A(x,w) + D(x,w)]\, z + T(x)\, F(0,w) \tag{3.47}$$

The next step is to calculate the derivative of $V := z^T z$ as follows

$$\dot{V} = z^T \left[A(x,w) + A^T(x,w) + 2D(x,w)\right] z + 2F^T(0,w)\, T^T(x)\, z \tag{3.48}$$

Applying Young's inequality[11] to the final term of equation 3.48 gives

$$\dot{V} \leqslant z^T [A(x,w) + A^T(x,w) + 2D(x,w) + T(x)\, T^T(x)] z + \|F(0,w)\|^2 \tag{3.49}$$

and from equation 3.42

$$T(x)\, T^T(x) = I_{n\times n} \begin{bmatrix} 0 & \circledast_1 & \circledast_2 & \cdots & \circledast_{n-1} \\ \circledast_1 & \circledast_1 & \circledast_2 & \cdots & \circledast_{n-1} \\ \circledast_2 & \circledast_2 & \circledast_2 & \cdots & \circledast_{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \circledast_{n-1} & \circledast_{n-1} & \circledast_{n-1} & \cdots & \circledast_{n-1} \end{bmatrix} \tag{3.50}$$

Therefore, combining equations 3.49 and 3.50 gives

$$\begin{aligned} \dot{V} \;\leqslant\; & z^T z + \|F(0,w)\|^2 + 2z^T D(x,w)\, z \\ & + z^T \begin{bmatrix} 2\phi_{11} & \bigstar_1 & \bigstar_2 & \cdots & \bigstar_{n-1} \\ \bigstar_1 & \bigstar_1 & \bigstar_2 & \cdots & \bigstar_{n-1} \\ \bigstar_2 & \bigstar_2 & \bigstar_2 & \cdots & \bigstar_{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \bigstar_{n-1} & \bigstar_{n-1} & \bigstar_{n-1} & \cdots & \bigstar_{n-1} \end{bmatrix} z \end{aligned} \tag{3.51}$$

So that for any value of the design parameter $c > 1$, there will exist choices for the s_i which satisfy

$$\max_{w\in B} \dot{V} \leqslant -(c-1)\, z^T z + \max_{w\in B} \|F(0,w)\|^2 \tag{3.52}$$

Using the definition of $D(x,w)$ leads to

$$\dot{V} \leqslant -(c-1)\, z^T z - z^T M(x,w)\, z + \|F(0,w)\|^2 \tag{3.53}$$

depend on both w and x_{i+1}.

[11] Young's inequality is: $2ab \leqslant a^2 + b^2$

where the symmetric matrix $M(x,w)$ is given by

$$M = \begin{bmatrix} -c-2\phi_{11}-2\phi_{12}s_1 & \bigstar_1 & \bigstar_2 & \cdots & \bigstar_{n-1} \\ \bigstar_1 & \bigstar_1 - 2\phi_{23}s_2 & \bigstar_2 & \cdots & \bigstar_{n-1} \\ \bigstar_2 & \bigstar_2 & \bigstar_2 - 2\phi_{34}s_3 & \cdots & \bigstar_{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \bigstar_{n-1} & \bigstar_{n-1} & \bigstar_{n-1} & \cdots & \bigstar_{n-1} - 2\phi_{n,n+1}s_n \end{bmatrix} \tag{3.54}$$

The functions s_i can be selected so that the matrix M is positive definite for all $x \in \mathcal{X}$ and all $w \in B$. This leads to the functions s_i being constructed as follows

$$M_1(x_1,w) := -c - 2\phi_{11}(x_1,w) - 2\phi_{12}(x_1,w)\, s_1(x_1) \tag{3.55}$$

Once s_1 has been found the functions $\bigstar_1$ in M can be derived[12] and the procedure continues progressively. Determining the s_i functions means that the feedback control $u(x) = z_n s_n(x)$ is known and is a robust control Lyapunov function for the system.

Constructing a stabilising feedback controller

Having constructed a robust control Lyapunov function, FK (1996) provide a highly computationally tractable formula for the construction of robust stabilising feedback controllers, which produces what they term point-wise min-norm (PWMN) control laws that are optimal and robustly stabilising with respect to the chosen cost functional-form. Using a simple example serves to illustrate FK's formula for generating PWMN control laws, given an underlying robust control Lyapunov function and avoids further lengthy exposition using turgid matrix algebra. Consider a simple function

$$\dot{x} = f_0(x) + f_1(x)\,u + f_2(x)\,w \tag{3.57}$$

[12] The second leading determinant of M is given by

$$M_2(x_1,x_2,w) := \begin{bmatrix} M_1(x_1,w) & \bigstar_1 \\ \bigstar_1 & \bigstar_1 - 2\phi_{23}(x_1,x_2,w)\, s_2(x_1,x_2) \end{bmatrix} \tag{3.56}$$

where f_0, f_1 and f_2 are continuous functions and where $U(x) \equiv \mathcal{U}$ and $W(x) \equiv B$. Then let V be a robust control Lyapunov function for this simple system, so that

$$D(x,w) = \nabla V(x) \cdot f_0(x) + \nabla V(x) \cdot f_1(x) u + \|\nabla V(x) \cdot f_0(x)\| + \alpha_V(x) \tag{3.58}$$

$$= \psi_0(x) + \psi_1^T(x) u \tag{3.59}$$

where

$$\psi_0(x) \; : \; = \nabla V(x) \cdot f_0(x) + \|\nabla V(x) \cdot f_2(x)\| + \alpha_V(x) \tag{3.60}$$

$$\psi_1(x) \; : \; = [\nabla V(x) \cdot f_1(x)]^T \tag{3.61}$$

which gives the PWMN control

$$m(x) = \left\{ \begin{array}{ll} -\frac{\psi_0(x)\psi_1(x)}{\psi_1^T(x)\psi_1(x)} & \text{when } \psi_0(x) > 0 \\ 0 & \text{when } \psi_0(x) \leqslant 0 \end{array} \right\} \tag{3.62}$$

for $x \in V^{-1}(c_V, \infty)$.

One of the key points to arise from this approach is that once the robust control Lyapunov function and stabilising feedback controller have been constructed, the resulting hedging strategy provides a clear floor under possible losses. This floor is guaranteed as the rule is robust with respect to uncertainties on modelling, data and measurement. There is the additional beneficial side effect that the computations are simple and efficient in operation.

3.3.3 An example of a simple non-linear robust hedging rule

To help understand how FK's procedure works in practice and before applying it to complex hedging problems, it is instructive to consider a simple example of hedging a single 3-month option on a Eurodollar interest rate deposit future. The model presented below provides an over-stylised and simplified view of the world, comprising a system equation and a cost function, J. The price of the underlying futures contract0, x, is assumed to follow a naively simple process where u is an unconstrained control input, such that robustness is required with respect to a disturbance known to take values in the interval $[-1, 1]$

- A simple dynamic system equation: $x' = -x^3 + u_i + wx$
- A quadratic cost equation: $J = \int_0^\infty \left[x^2 + u_i^2\right] dt$

It is clear by simple inspection that a robust control Lyapunov function for this system is $V(x) = x^2$. The derivative of this can be made negative with the simple control law

$$u_1 = x^3 - 2x$$

This control law is the obvious one suggested by feedback linearisation and does indeed succeed in producing a robust and stable solution. However, two alternative control laws, u_2 and u_3 are examined to highlight the benefits of using a pointwise min-norm robust optimal control hedge

- A robustly stabilising feedback hedging rule based on solving a Hamilton-Jacobi-Isaacs (HJI) equation: $u_2 = x^3 - x - x\sqrt{x^4 - 2x^2 + 2}$
- A pointwise min-norm hedging rule based on solving a robust control Lyapunov (point-wise min-norm - or PWMN for short) equation: $u_3 = \begin{pmatrix} x^3 - 2x & \text{when } x^2 < 2 \\ 0 & \text{when } x^2 \geq 2 \end{pmatrix}$

From simple inspection it is easily seen that u_1 wastes a potentially useful source of non-linear dampening, u_2 does not, neither does it ever generate positive feedback. However, the superior characteristics of u_2 result from having to solve a steady-state Hamilton-Jacobi-Isaacs equation, which as FK point out, is only possible for very simple classes of non-linear systems. FK's approach provides a simple formula for generating a class of optimal control laws that involve only the robust control Lyapunov function , the system equations and the required design parameters. They refer to their controllers as pointwise min-norm control laws. In the case of the above simple example, the FK pointwise min-norm control law u_3 is

$$u_3 = \begin{matrix} x^3 - 2x & \text{when } x^2 < 2 \\ 0 & \text{when } x^2 \geqslant 2 \end{matrix} \tag{3.63}$$

which when compared against u_1 and u_2 shows identical qualitative behaviour with the added benefit that they do not waste control effort by trying to cancel out the beneficial non-linearity in the system equation. The relative performance of the two control laws, using scaled data and expressed in distance in basis points from the at-the-money strike is compared in figure 3.2. It is clear by simple inspection that the two methods produce extremely similar control rules, though the computational effort required is markedly different. The PWMN rule is rapidly computable on a normal desktop personal computer using a standard package such as Microsoft Excel. In the case of more complex examples, a relatively trivial amount of programming in a language such as Visual Basic for Applications or C#, will serve to produce a crude but effective result whose output can be easily viewed in Excel.

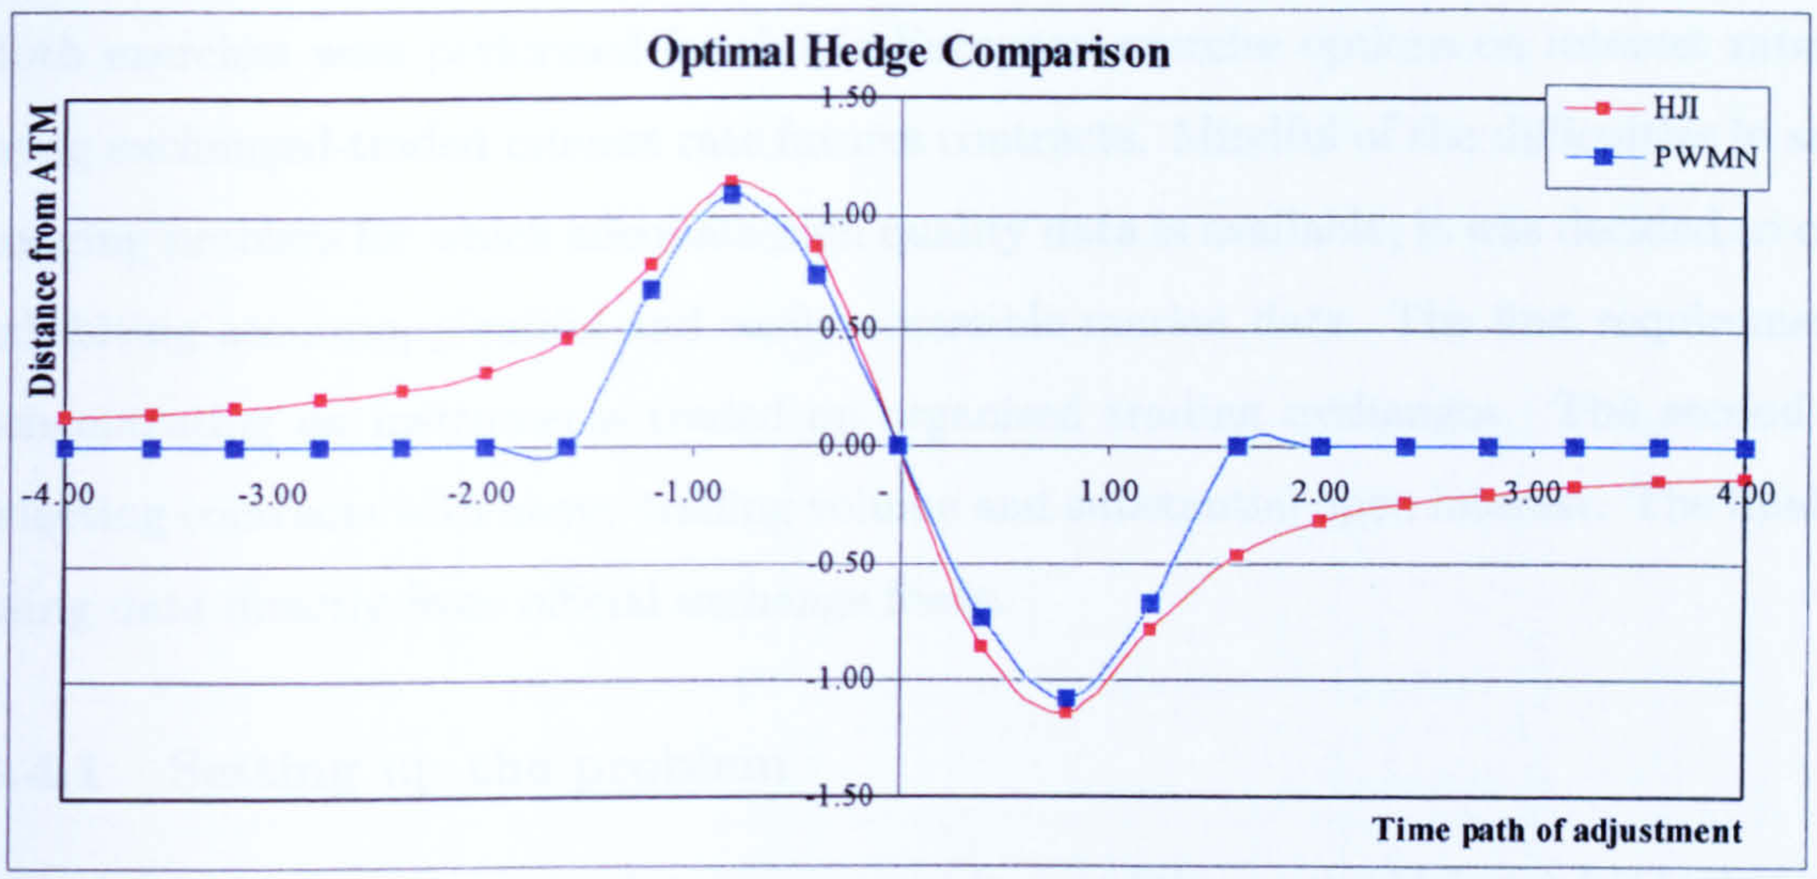

Figure 3.2: Comparison of on-linear robust hedging strategies

When either of the control laws above are applied to the hedging problem, it is clear from the following figure that the results are extremely similar in terms of the behaviour of the objective and system functions. The benefits of using the PWMN over the HJI are that the former is easier to formulate, easier to deal with computationally and quicker to solve, thereby making the approach more attractive from a practical perspective.

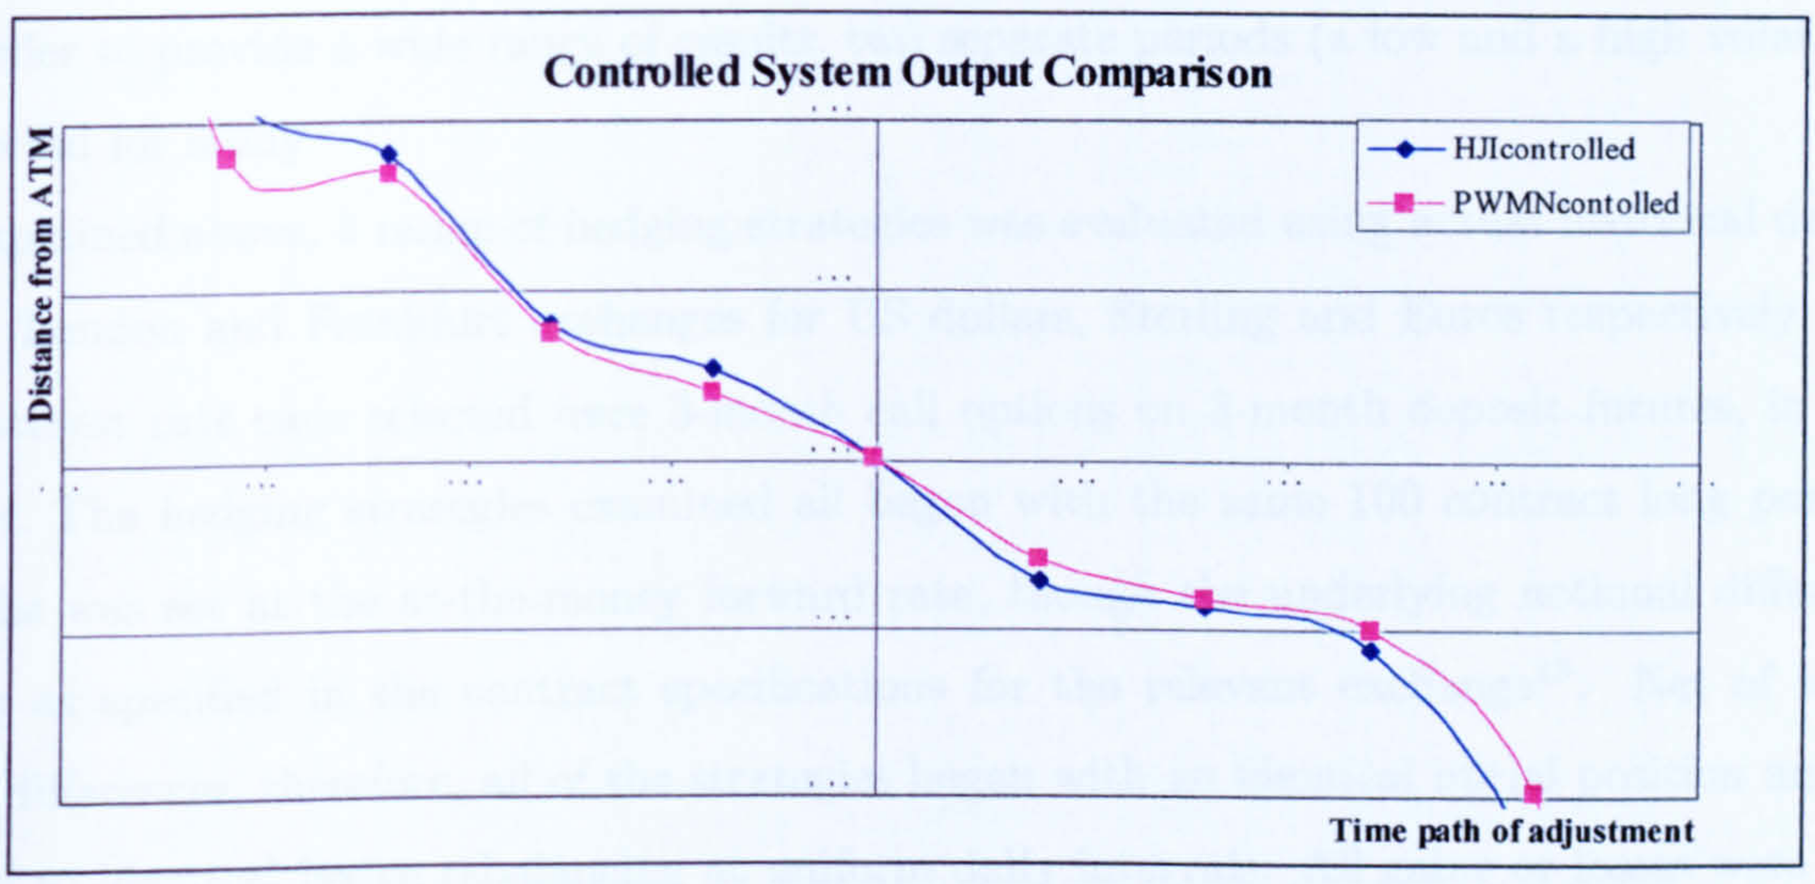

Figure 3.3: Comparison of controlled output using non-linear robust hedging strategies

3.4 Empirical comparison of optimal hedging rules

The above simple example clearly does no more than illustrate the PWMN approach. Testing on a more realistic and useful problem required a two stage process. The first stage involved running Monte Carlo simulations to provide a realistic and accurately calibrated alternative for comparison with the PWMN model. The second stage involved applying the PWMN control rules to real-world data. Both stages also included a number of the hedging metrics discussed in the literature review for comparative purposes.

Both exercises were performed for simple European exercise options on interest rates, that were hedged using exchanged-traded interest rate futures contracts. Mindful of the difficulties in selecting a real-world hedging problem for which adequate high quality data is available, it was decided to consider instruments exhibiting accurate, plentiful and easily accessible market data. The first requirement inevitably meant concentrating on instruments traded on organised trading exchanges. The second requirement meant selecting contracts with heavy trading volume and substantial open interest. The final requirement meant using data directly from official exchange feeds.

3.4.1 Setting up the problem

The three month interest rate futures contract is one of the most widely traded instruments in the financial markets, so based on its liquidity and high levels of open interest, it was decided to use it as the hedging instrument. In order to simplify interpretation of the results, an underlying option instrument with exactly similar frequency, day-count and interest basis was selected, namely, a European exercise interest rate cap whose underlying is the three month interest rate future. The major virtue of using exchange traded instruments is the plentiful supply of high quality detailed daily data. For comparative purposes and in order to provide a wide range of results, two separate periods (a low and a high volatility period) were selected for study.

As explained above, a range of hedging strategies was evaluated using actual historical data from the Chicago, London and Frankfurt exchanges for US dollars, Sterling and Euros respectively. The single period interest rate caps selected were 3-month call options on 3-month deposit futures, in USD, GBP and EUR. The hedging strategies examined all began with the same 100 contract long position where each strike was set at the at-the-money forward rate, though the underlying notional differed between contracts as specified in the contract specifications for the relevant exchange[13]. Net of idiosyncratic contract differences, therefore, all of the strategies began with an identical initial position and were then subjected to identical hedge rebalancing at uniform daily intervals. All gains or losses were assumed to be rewarded or penalised at the relevant futures margin-account interest rate. The interest rate costs, forward rates and discount factors were calculated using a standard yield curve bootstrap developed for this thesis. Volatilities required for valuation were a combination of those supplied by the relevant exchanges and those supplied by Credit Suisse First Boston. Hedging and funding flows were assumed to be calculated with respect to the bid or offer rates for the cash market for the relevant currency. All calculations were carried out with respect to end of day mark-to-market valuations of the relevant

[13] Detailed contract specifications for Chicago, Frankfurt and London are easily available via the internet through the portal sight www.numaweb.com.

variables. The final performance of a strategy was defined as the final cumulative value of the initial position, plus cash inflow, less cash outflow, normalised by the notional amount of the underlyings of the relevant currency. Table 3.1 below summarises the eleven strategies that were evaluated. Detailed explanations of each of the strategies along with a description of how the accompanying simulation data was generated, are all contained in Appendix 5. The results of the Monte Carlo calibrating simulations, together with the results of the historical simulations are presented in tables 3.2 - 3.8.

Table 3.1: Hedging strategies evaluated

Strategy	Objective Function	Conditions	Costs[14]
Delta	Delta neutrality	n.a.	Excluded
Minimax-95	Potential hedge error	95% level	Excluded
Minimax-99	Potential hedge error	99% level	Excluded
Heuristic-w	Potential hedge error	Weighted	Excluded
Minimax-95c	Potential hedge error	95% level	Included
Heuristic-95c	Potential hedge error	95% level	Included
ρ^2	Potential hedge error	n.a.	Included
$\theta/(r_s - i)/\sigma_s$	Potential hedge error	n.a.	Included
$r_h^{ce} - r_s^{ce}$	Potential hedge error	n.a.	Included
Linear $H\infty$	Potential hedge error	Uncertain	Excluded
Non-Linear $H\infty$	Potential hedge error	Uncertain	Excluded

Note that all strategies were evaluated for Euro, Sterling and US dollar futures and options. for both high and low volatility periods in each case (see Appendix 5 for details).

3.4.2 Simulation results - first pass

The results of running the hedging strategies contained in table 3.1 show the time-path of each of the strategies evaluated. The results are presented in tables 3.2 to 3.8 below[15]. The first set of results in table 3.2 are those generated by the Monte Carlo simulation using arbitrary flat volatility of 20% for benchmarking purposes. Although the results are not totally realistic is so far as they do not use historical volatilities, they are broadly representative of the flavour of the results in tables 3.3 - 3.8 inclusive. The first major point is that even very close to the at-the-money-forward strike (i.e. the 0.5% column in table

[14]Transaction costs can be modelled in a number of different ways. However, it was felt that adjusting the bid-ask spread was the most appropriate method in the interests of transparency and in line with most capital-market conventions.

[15]Note that the only the results for positive distances from the ATMF strikes have been reported. This is because the negatives were very similar, so that reporting them would have not added anything to the overall results. In addition, all results have been converted back into EUR in order to aid comparison.

3.2), the costs of robustness are relatively high for the robust strategies. Figure 3.4 re-expresses the data contained in table 3.2 in terms of the percentage of average profit lost for each strategy compared with the simple delta hedging strategy. What is immediately obvious from figure 3.4 is that the non-robust strategies perform remarkably similarly. However, the robust strategies are all hugely and disastrously more expensive than any of the non-robust strategies, with the linear H∞ being the most expensive compared to the standard delta hedging strategy.

Table 3.2: Monte Carlo simulation results for alternative hedging strategies
(average profit in '000s of EURs)

Distance from ATMF strike	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%
Delta	96.7098	92 0720	88 3192	84 8363	81.2908	78 1407	75.1868	71 8199	69.0689	66.0473
Minimax 95	95.3731	91 3861	86 6990	82.9219	79.5367	75.3133	71.8866	68 2621	65.0369	62.0451
Minimax 99	94.7941	88 6747	83 6654	78.5424	74.2267	69 4703	65.3871	61 6472	58 6867	54.5921
Heuristic w	95 4511	90.3891	85 8027	81 6071	77.8251	73.5963	70.3632	66 5998	63.4110	59 9469
Minimax 95c	86 9303	75 0096	65.1623	56 3755	48 7659	42 2997	36 5861	32 0085	28 1935	23 8390
Heuristic 95c	82 5940	68 2387	56 2045	46 3879	38 2998	32.0765	26.8618	21 4773	17 8312	14.7606
Ederington	96 2116	92 6845	88 8073	85.1297	82.3519	78.9539	75.5142	72 8251	69.9574	67.0624
Howard & D'Antonio	95.5990	90.6496	86.4468	82 0818	78.3820	74 5489	71.1139	67.3984	64.3321	61.5146
Pennings & M	94 4877	88 9052	83.3809	78.4518	74 0389	69.8370	65.5788	61 4651	57.7986	54.4239
Linear H-inf	63.7829	40 6543	25 9627	16 6455	11.2796	6.8010	4 6045	3 3926	1.7744	1.3305
Non-Linear H-inf	74.1826	54 5206	40.1854	30.1190	21.9686	16.2591	12.1230	8 8743	6.8716	5.3753

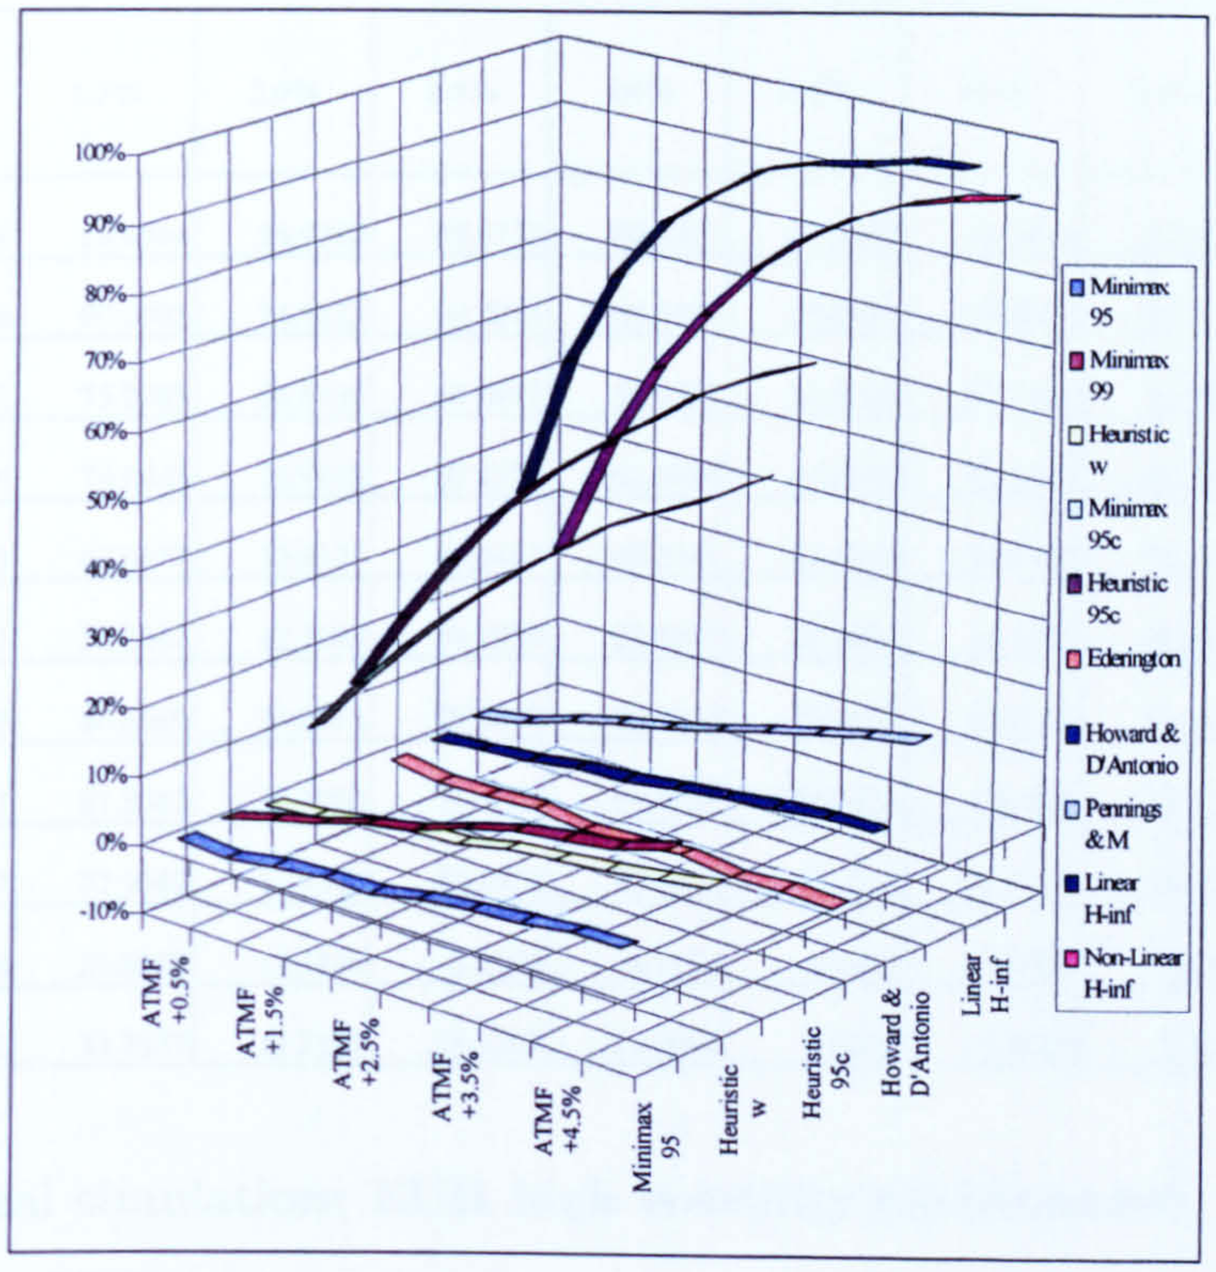

Figure 3.4: Comparison of robust and non-robust hedging strategies

What is interesting is that the pattern is broadly maintained throughout the historical simulation results shown in tables 3.3 - 3.8. In fact the pattern is even more dramatic as the distance from strike increases for the robust methods. Once again the situation appears clear that the cost of hedging robustly in a benign environment are hugely expensive and arguably not warranted under any kind of "normal trading conditions". The question then clearly follows as to whether such behaviour continues in the historical simulation results. The short answer is that it does as can be clearly seen by even a casual perusal of the information contained in tables 3.3 - 3.8. However, what is also clear throughout all of the historical simulation results reported in tables 3.3 - 3.8 is that with the exception of the high volatility environment for GBP, the non-linear robust strategy proves to be less punishingly costly than the linear robust strategy. This is probably due to the non-linear strategy "swimming with the tide" and therefore having to fight less hard against the losses. There is also some intuitive appeal in the argument that robustness represents severe overpayment for insurance against an environment so extreme, that in line with the rest of the insurance and reinsurance industry, such cover would never be viable at such prices.

Table 3.3: Historical simulation: EUR low volatility environment

(average profit in '000s of EURs)

Distance from ATMF strike	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%
Delta	97.1510	93.4856	89 6844	86.9288	83 0773	80.3523	77 4430	74 4893	71.9963	68 9881
Minimax 95	93 4362	86 4416	80.5405	74 8523	69.7025	64.4436	59 8644	55.4613	51 8588	48 0936
Minimax 99	91.2657	82 9412	75 2989	68 6086	62 0997	56.5392	51 9260	47.1468	42.8802	39.2861
Heuristic w	93 0112	85 4290	79 0449	72 9883	67 4522	62 3420	57.8777	53 3659	49.5395	45.5206
Minimax 95c	86 2313	73 3055	63.0377	53 9131	46 0497	40.2441	34 4793	29.0972	24.7727	21 7760
Heuristic 95c	80 6860	65 1087	52 8667	42.5063	34 1929	27 8357	22.3469	18 0297	15 2618	11.7792
Ederington	95 7777	90.5745	86 0669	81 8875	78 1042	74 1239	70.9932	67.0209	64.4290	61.1383
Howard & D'Antonio	93 6722	87 5191	81 8082	76 7290	71.5602	67.0195	62.7025	58 6298	55.1510	51 9492
Pennings & M	91 9617	84 5447	77 9040	71.7022	65.3677	60 1463	55.5762	51.0564	46 9567	42.7215
Linear H-inf	60 2164	36 3978	21 9387	13 4408	8 2323	4.7533	3.3965	2 0391	1.9229	0.9293
Non-Linear H-inf	69 6094	48 6731	33.7632	23.2260	16.1267	11.3093	7.8058	5 6020	3.7785	2.8415

Table 3.4: Historical simulation: EUR high volatility environment

(average profit in '000s of EURs)

Distance from ATMF strike	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%
Delta	93.9586	86 5284	81.0726	74 8013	69 6317	64 6813	60.3561	56.0750	52.0093	48.5947
Minimax 95	86 6890	75 6110	65 3448	56 4317	48 8568	42.7988	36 9249	31.8810	27.8804	23 9362
Minimax 99	83 4981	69 6409	58 4826	48 4537	40 9861	33 9226	28 3615	23.4195	19.5448	17.0828
Heuristic w	87 8084	76 4921	66 9457	57 8798	50.8631	44.2095	38 6425	33.5602	29.3465	25.7495
Minimax 95c	76 2463	58 0911	44.2459	34 4364	25 8069	19 6355	15.1843	11.5203	8.7589	7.0047
Heuristic 95c	71 2361	49 9791	35.5592	25 2473	17.5998	12 4670	8 9055	6.2697	4.5698	3.3589
Ederington	89 6057	80 6470	71 3707	64 1470	56 9938	50 9033	45 4713	40 6608	36 6480	32.5593
Howard & D'Antonio	87.2648	76 5322	66 3111	57 8572	50.3309	43.9573	38.3109	33.3847	29 0183	25 6432
Pennings & M	84 8910	72 1028	60 4432	51 1368	43 3517	36.8395	30 9307	26 1588	22 0093	19 0491
Linear H-inf	33 4014	11 1064	4 6253	1.2278	0 8799	0.3713	0.1050	0 0703	0.1447	0.3576
Non-Linear H-inf	51 0790	26 0525	13 3960	6.7315	3.5734	2.1841	0.8959	0.5383	0.2867	0.4097

Table 3.5: Historical simulation: GBP low volatility environment

(average profit in '000s of EURs)

Distance from ATMF strike	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%
Delta	95 3578	91 4259	87.1594	82.4908	79 0374	75.8036	71 6437	68 0526	64.7550	61.7232
Minimax 95	92 4879	84.7776	78 0027	71.7184	66.4328	60.8280	56 0395	51 9334	47.5674	43.7796
Minimax 99	89 8127	81.0284	72 9275	65 0343	58 4855	52.9881	47.5647	42 6329	38 0256	34.1343
Heuristic w	91.4573	84 0703	76 5091	69 9540	64 5175	58 5033	53 6135	48 9346	45 1590	41.0984
Minimax 95c	84 3950	71 6435	60 2861	50 6821	42 8523	35.8982	30.3465	25.4954	21.7690	18 1198
Heuristic 95c	79 9622	62 6481	49 5542	39.2598	31.2323	24.7908	19.7333	16.0062	12.2091	9.6983
Ederington	94 3483	89 0463	83 8136	78.7802	73.9431	70.1650	65.6378	61.9249	58 0434	54.6429
Howard & D'Antonio	92.7078	85 6920	79.2736	73.3240	67.9303	63.1937	58.5284	53.8378	49.7797	46.1380
Pennings & M	91.2064	82.5626	75.3338	68.1511	62 5313	56.3348	51.3513	46.4831	42.3585	38 6102
Linear H-inf	50 4710	26 2040	13 6043	7.0827	3 3377	1 6630	0.8312	0.4423	0.4287	0 1941
Non-Linear H-inf	64 0591	41 0107	26 6236	17.5189	10 9905	7.0459	5.1414	2.8315	1 9839	1.1520

Table 3.6: Historical simulation: GBP high volatility environment

(average profit in '000s of EURs)

Distance from ATMF strike	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%
Delta	92 4830	84 1093	77.1662	70 9561	64.7685	59.5038	54 4223	49 7978	45 6926	41.8773
Minimax 95	84 9506	72 0887	61.1943	52.3001	44 1906	37 6918	32.3800	26 9708	23 0173	19.4347
Minimax 99	81.3046	66 3978	53 5913	43 9916	35 7477	28 9739	23 3326	19 0805	15 7879	13.0546
Heuristic w	85.8882	73 4535	63 0692	54.1951	46 0857	39.5933	33.9813	29 8249	24.8104	21.5032
Minimax 95c	71.5224	51 1774	36 5959	25 9288	19 2559	13.2400	9.4244	6 7322	5.3051	3.5488
Heuristic 95c	65.2834	42 6161	27 8152	18 1572	12 4410	7 8794	5 1140	3 4119	2 2176	2.1688
Ederington	86 4370	74 8288	64.7972	56.1911	48 6510	42.0117	36.1795	31.4006	26.8149	23.1602
Howard & D'Antonio	79.9947	63 9139	51 1291	41.3193	32 6374	26 0938	21.1802	17 2225	14.1232	11.1229
Pennings & M	82 0253	67.3020	55.3858	45 4182	37.5865	30.4637	25.2949	20.9628	16 8199	13.9561
Linear H-inf	28 6909	8 4639	2 3741	0 6874	0 3502	0.7539	0.1938	0 6912	0 8576	0.4688
Non-Linear H-inf	46.2087	21 0865	10 4405	4 4573	2 6638	1.0378	0.5695	0.2175	0.3033	0.3008

Table 3.7: Historical simulation: USD low volatility environment

(average profit in '000s of EURs)

Distance from ATMF strike	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%
Delta	98 4209	97.2903	95 5889	94.5478	92 9881	90 8469	89.4360	87.9993	86.5636	85.2050
Minimax 95	95.7140	90 9585	86 8547	82.3886	78 6236	75.1788	71.7874	68 2304	64.6748	61.8245
Minimax 99	94 2123	88 3035	83 0611	78 5859	73.7193	68.8978	64 8242	61.1670	57.1256	53.7324
Heuristic w	93 3744	87 5632	81 1964	76 4290	70.7130	65.8772	61.7979	57 8421	53.4696	50 0365
Minimax 95c	88 9379	78 4121	69.3203	60 9710	53.8195	47.6739	42 0314	38 0478	32.9431	29 0134
Heuristic 95c	85 9613	73 2389	62.4630	53 6616	45 6638	39 8574	33.3607	28.5091	24 4401	21.4510
Ederington	96 4935	91.9583	87 6216	84.5133	80 1577	77.4937	74.1644	70.6388	67.4220	64.3254
Howard & D'Antonio	94 7941	89 5383	84 7659	80 2121	75 7908	72.3184	68 1325	64 6150	60.8768	57.5800
Pennings & M	93.7206	87.9716	82 2787	76.5581	71.9013	66.9926	62.6432	58 6589	54.9358	51.8318
Linear H-inf	70 4944	49 6369	34 8061	24 2138	17.0063	12.7301	8 4594	6 0460	4.1055	3.0783
Non-Linear H-inf	79 6353	63 9083	50 4183	40.4654	31.9591	25.5851	20.6878	16.1020	12.9093	10.2515

Table 3.8: Historical simulation: USD high volatility environment

(average profit in '000s of EURs)

Distance from ATMF strike	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%
Delta	95.0316	90 3065	85.8111	81 8668	77.3799	73.4917	69.7469	66 2939	63.2475	59.9261
Minimax 95	90 4037	81 4914	73 6560	66 6313	59 8530	54.3301	48.8608	44.7065	39.9499	36.0709
Minimax 99	87.2756	76 3038	66.9839	57 9295	50.6160	44 6714	38.7809	33.5649	29.5670	25 4151
Heuristic w	89.3931	79 9965	71.4085	64 0972	57.7771	50.9754	46.2559	40 8512	36.4452	32.6700
Minimax 95c	80 4672	65 0205	52 7129	42.2146	33.7707	27.2602	21.8104	17.6295	14 4995	11.5010
Heuristic 95c	77.3458	60 0894	46.5486	36.1963	27.6882	21.5287	16 7552	13 5550	10.1895	8 0030
Ederington	91 6134	83 9066	76 8943	70 9733	64 5389	59.1471	54.1017	49.5821	45 4842	41.7743
Howard & D'Antonio	89 8905	80.7982	72 8333	65.3748	58 4719	53.0182	47.2843	42.4799	38.7286	34.3396
Pennings & M	86.7125	74 7835	64.5038	56.2642	48.2576	41.7680	35.9517	31.5440	27.0455	23.7505
Linear H-inf	49 7554	24 2612	12 4914	5.7732	3.1825	1.9038	0 8509	0.5115	0.2274	0.3033
Non-Linear H-inf	73 4919	54 0621	40 1217	29.1323	21.5176	15.7075	12.3135	9 0405	6 8484	4.5913

3.4.3 Simulation results - second pass

The initial results for the PWMN reported in tables 3.3 - 3.8 were based on a robust control Lyapunov function that is quadratic in the set of non-linearly transformed z-coordinates. As can be seen from the results tables, this form produces hedging rules that whilst robust, are less conservative than those produced by a linear $H\infty$ approach. However, upon closer examination it became apparent that the control rules produced by this form of the Lyapunov function produced excessive responses to large shifts in the prices of the underlying options and consequently large changes in the hedge position and consequently much lower overall profitability. Upon investigation of the intermediate results of the simulations, this behaviour was found to be attributable to the fact that the PWMN hedge rule was exhibiting local gains that were growing extremely rapidly for options that were further out of the money and exhibited higher volatility.

Several approaches have been suggested and evaluated in the literature (see Marino and Tomei, 1993 for example), but FK suggest a method that involves modifying the robust control Lyapunov function to reduce the size of the local gains which has the effect of drastically reducing the control effort required to achieve robust stabilisation, with no reduction in computational performance. The method involves using a smooth, non-negative scalar function, ϱ, to penalise the *distance to the region around* the control manifold rather than the *distance to* the control manifold itself. This involved a relatively simple modification to the algorithm described above (with the additional trivial constraint that $\varrho_i(0) = 0$), which is both necessary and sufficient to guarantee that the choice of V will possess the required first order differential continuity, as well as being positive definite (which helps, though does not guarantee stability). Experimentation with the most suitable functional form for ϱ indicated, that as FK suggest, the following form works best

$$\varrho_i(x_i) = |x_i|^r \tag{3.64}$$

The value of r was investigated at some length in an attempt to determine an automated, algorithmic procedure for its generation. After some experimentation, it was found that higher order powers tended to increase the chattering effect when hedging deeper out of the money options. Higher hedge returns from the modified PWMN function worked best over all three currencies with $r = 2$. The new algorithm was therefore re-calibrated using the Monte Carlo simulation and the revised results are shown in table 3.9 below.

Table 3.9: Simulation results for flattened PWMN hedging rule

(average profit in '000s of EURs)

Distance from ATMF	Original Linear H-inf	Original Non-Linear H-inf	Flattened Non-Linear H-inf(2)
0.50%	63.7829	74.1826	88.9121
1.00%	40.6543	54.5206	85.0237
1.50%	25.9627	40.1854	60.7866
2.00%	16.6455	30.1190	44.6229
2.50%	11.2796	21.9686	36.3107
3.00%	6.8010	16.2591	30.1585
3.50%	4.6045	12.1230	25.7826
4.00%	3.3926	8.8743	22.1450
4.50%	1.7744	6.8716	20.0590
5.00%	1.3305	5.3753	17.9997

What is immediately clear from table 3.9 is that the results of using the flattened piecewise-min-norm, non-linear robust rule are far less punishing when compared with the unflattened non-linear H∞ hedging rule. However, when compared with the standard delta neutral approach, the results are still extremely expensive for what amounts to insurance that is not wanted and extremely unlikely ever to be exercised. And if that was the end of the matter, then the results would not be of great significance. However, that is not the end of the story for two reasons. First, although such hedging/insurance may appear to be expensive, when considered in conjunction with the observed rise in the number of natural catastrophes, the results may not prove to be quite so expensive. Although at the time of writing, natural catastrophes had not caused significant spill-over effects in the mainstream financial markets, it can only a matter of time given the increasing concentrations of population and assets in the United States for example, areas that are increasingly affected by natural catastrophes.

3.5 Conclusions

This chapter has categorised and reviewed alternative approaches to the problem of developing a dynamic hedging strategy that is robust and stable in a multi-period hedging problem. Robust optimal control techniques were applied and two forms of dynamic hedging strategy were developed - one linear, the other non-linear. The linear model used $H\infty$ optimal control techniques and produced hedging rules that were found to be robust but too conservative to be useful in practice. Piecewise min-norm robust optimal control rules were developed based on robust optimal control Lyapunov techniques in an attempt to exploit the suspected non-linearities of the the hedging problem. The first version of the PWMN rule was found to be less conservative than the $H\infty$ rule, but was observed to still be implying excessively punitive changes in hedging policy for options deep out of the money. The PWMN rule was modified by using a simple scalar function modification and found to produce encouraging modification to the previously conservative behaviour, whilst retaining desirable robustness and stability properties.

Two possible ideas occurred for further research in the area of softening PWMN control laws. The first would involve using a non-smooth robust control Lyapunov function, thereby implicitly allowing

non-differentiability of the Lyapunov function, which is not in and of itself a requirement for stability and therefore poses no significant constraint in the context of the current problem. The second idea is to use directional derivatives to eliminate the "chattering" in the hedge rules that can appear at much greater distances from the at the money forward strike and which undoubtedly had a significant impact on the apparent cost of such options.

Chapter 4

Robust optimal control and risk management

All models are wrong, but some are useful.

George Box

4.1 Introduction and motivation for research

This chapter continues the theme of robust decision making by examining the robustness of portfolio level risk management in the presence of uncertainty. The previous chapter examined the problem of developing and applying optimal, dynamic hedging strategies for individual exposures that are both robust and stable in operation. The objective of the current chapter is therefore to extend this framework to portfolios of instruments and investigate its performance relative to other established portfolio risk management approaches.

The principal motivation for the research contained in this chapter is threefold. First, is the fact that despite the considerable and ever increasing levels of risk being carried both on and off the balance sheets of many financial institutions, there has been relatively little interest in either the theoretical or practical robustness and stability of current measurement methodologies. The bankruptcies of Enron and WorldCom alone wiped out loans on bank balance sheets of some $34 billion[1]. Financial institutions continue to reduce collateralised low-return loans on their balance sheets in an ever accelerating search for higher return and inevitably riskier investments. Despite both the complexity and magnitude of the risks involved, the metrics for quantifying and controlling the risks remain relatively simple in their conception

[1]Source: "Who's carrying the can ?", The Economist 14th August 2003.

and operation. Second, current methods of measurement and management are based on the often implicit assumption that some arbitrary distribution can be acceptably parameterised to deal with the problem at hand. This is the fundamental premise behind such risk measures as value at risk and extreme value theory - both of which are examined in detail in the literature review in the following section In reality, such distributions have not proven to be robust with respect to the measurements of risk carried by financial institutions in particular. Third, current approaches take little or no account of the dynamics of the instruments or portfolios that are being managed, such that modelling is predominantly carried out using static optimisation techniques ignoring feedback. This chapter therefore attempts to deal with two of these distinct problems, namely, quantification of uncertainty at a portfolio level and managing the aggregate risk of a portfolio of individual derivative instruments in the presence of uncertainty.

This chapter therefore contributes to the body of existing research into managing portfolio level risk through the development of a new methodology for measuring and controlling portfolio risk based on the application of non-linear robust optimal control techniques. The research applies existing techniques to an existing problem in a novel fashion to enable the creation, execution and measurement of the performance of portfolio level risk management strategies. It is important to note that the focus of this chapter is *not* the robustness of optimal portfolio construction techniques (as studied by authors such as Lutgens (2004)), but rather the robustness of the rules and metrics used to measure, monitor and control portfolio level risk. The chapter therefore proceeds as follows. The next section provides a review of the current literature and research. This is followed in the succeeding section by the development of an alternative robust optimal control approach. An empirical comparison and evaluation of the current approaches with the new approach follows and a final section concludes.

4.2 Review of the literature

Actuarial and financial mathematicians have approached the problems encountered in the managing of uncertainty in a variety of ways, some of which share common features, assumptions and models, whilst the differences in the approaches arise from the diverse nature of the constraints under which each industry operates. Insurance companies are predominantly long term in their focus (e.g. matching the payouts on insurance policies with long liability tails, such as asbestosis for example where claims could have a 30 or 40 year tail), being overwhelmingly concerned with the probability of ruin - usually expressed as the total elimination of available capital. In contrast, the capital markets though still very much concerned with the probability of ruin, must contend with a very different maturity profile that is much shorter term (the average term in both the vanilla interest rate and credit default swap markets is around 5 years). The results have been the emergence of three main approaches to the problem of managing uncertainty, the

assumptions and predictions of which will be analysed in detail in this section. But before proceeding it is instructive to get a flavour for the nature of the constructs involved in each of the three methodologies.

Beginning with the areas of commonality of approach, it is not therefore surprising that the current focus when managing such financial risks is primarily concerned with analysing behaviour of possible losses in the tail quantiles of distributions that do not have heavy or fat tails, such as, for example, the value of x such that $P(X > x) = 0.05$ for large values of x. This is very much the idea behind the ubiquitous Value at Risk (usually referred to as VAR) approach, where extreme quantiles and probabilities are of special interest due to the view that the capability to accurately assess such quantiles of regular distributions provides the key to being able to manage the extremes associated with financial crises.

As explained below, traditional parametric methods that are usually based upon the estimation of entire densities are not the most appropriate tools for assessing extreme quantiles or probabilities of rare events. This is because parametric methods attempt to provide a good fit in the tails of the distribution, where by definition, few observations are likely to be found. Silverman (1986) also shows that the same poor estimation performance is true for non-parametric methods of estimating density such as kernel smoothing. A further complication is that it is also frequently necessary to estimate quantiles and probabilities both near and in many cases beyond the boundary. Attempting to estimate densities in such circumstances would appear to be at best a hopeless exercise due at least in part to the almost total lack of useful and relevant data.

Such is in fact not quite the case, according to the supporters of an approach known as extreme value theory (or EVT for short), which provides a methodology to estimate extreme quantiles and probabilities by fitting a model to the empirical survival function[2] of a dataset using only the extreme event data for the tail of the distribution. The result is attractive to many practitioners for two main reasons. First, fitting to the tail and only the tail of the distribution means that it is possible to concentrate solely on the area of specific interest without biasing the results by the use of the centre of the distribution. Second, it is possible to adapt any reasonable functional form and use it as a representation of the tail of a distribution.

The third approach considers the problem managing uncertainty as a minimax problem, on the grounds that decisions under uncertainty are incorrectly based on expected value optimisation which will never be robust as it ignores the possible impact on the system under consideration of realisation of the worst-case. Therefore, if the objective is to make decisions that are the best possible alternative and therefore robust in the face of the worst-case outcome. The appropriate course of action is then to adopt the minimax

[2] A survival function is one minus the cumulative density function, i.e. $1 - F(x)$, remembering of course that because $F(x)$ approaches unity as x grows, then the survival function approaches zero.

criteria as the decision metric. Looked at mathematically, the minimax criterion amounts to minimising a non-differentiable objective function that is defined by the maximum of an inner function - a procedure generally referred to as minimax. In the majority of applications this is stated in the form of a disutility or cost function. The problem of making decisions in the presence of uncertainty therefore reduces to finding the least worst outcome in the face of the worst-case event or scenario - behaviour that has been shown to satisfy certain robustness criteria, but not necessarily stability criteria. The remainder of this section is therefore devoted to a detailed analysis of each of the three methodologies in turn, beginning with VAR.

4.2.1 Value at Risk

According the Risk Metrics document by the US bank JP Morgan-Chase, Jorion (1997) and Deutsch (2001), VAR can be described as a method of assessing risk that uses standard statistical techniques to measure the worst expected loss over a given time interval, under normal market conditions for a given confidence level. In other words, how much is likely to be lost with a given probability over a given time-horizon. An alternative way of expressing this is that VAR is the lowest quantile of the potential losses that can accrue to a given portfolio within a pre-specified time-period. Expressed more succinctly and usably, the essence of the VAR approach can be stated as follows. The VAR of the value, V, of a financial instrument or portfolio is the upper bound for the loss which will not be exceeded with given confidence level over a stated time horizon

$$c = CP_{\delta V}\left(\delta V \geqslant VAR\left(c\right)\right) \tag{4.1}$$

where c is the confidence limit, V is the value of the financial instrument and $CP_{\delta V}$ is the cumulative probability function of the random variable δV over some time horizon δt. Or, more explicitly

$$c \neq 1 - CP_{\delta V}\left(\delta V \leqslant -VAR\left(c\right)\right) = 1 - \int_{-\infty}^{-VAR(c)} pdf_{\delta V}\left(x\right) dx \tag{4.2}$$

where $pdf_{\delta V}\left(x\right)$ is the probability density function of the random variable δV over some time horizon δt. Note that VAR is defined by the probability distribution of δV and not by the distribution of the associated risk factors. Given the simple case of a single risk factor, S, described by a standard geometric Brownian motion

$$d\ln S\left(t\right) = \mu dt + \sigma dW \quad \text{where } dW \sim N\left(0,1\right) \tag{4.3}$$

the process for S can be described by the stochastic variable $y := \ln S(t)$, such that y satisfies the standard geometric Brownian motion

$$dy(t) = \mu dt + \sigma dW \tag{4.4}$$

Then, choosing the function $f(y,t) = e^y$ and invoking Itô's lemma gives

$$df(S,t) = \left[\mu \frac{\partial f}{\partial S} + \frac{\partial f}{\partial t} + 0.5\sigma^2 \frac{\partial^2 f}{\partial S^2}\right] dt + \frac{\partial f}{\partial S}\sigma dW \tag{4.5}$$

which, since $f(y,t) = S(t)$ gives the following expression for the differential of S

$$dS(t) = S(t)\,\widetilde{\mu} dt + S(t)\,\sigma dW \tag{4.6}$$

where $\widetilde{\mu} = \mu + 0.5\sigma^2$. The next step is therefore to derive the process for a risk factor for a finite time period, δt, by solving equation 4.6. This can be achieved by defining a further stochastic variable $y = W(t)$ as the value of the Wiener process at time t, such that it satisfies $dy(t) = dW(t)$, then forming a function $S(y,t) := S_0 \exp(\mu t + \sigma y)$, where S_0 is any arbitrary initial value. Invoking Itô's lemma shows that the constructed process is a solution for equation 4.6, which upon making the substitution $t \to t + \delta t$ gives the required expression for a change in S over a finite time δt as

$$S(t+\delta t) = S(t)\exp(\mu dt + \sigma dW) \qquad \text{where } \delta W \sim N(0, \delta t) \tag{4.7}$$

where δt is assumed to be arbitrarily long and is usually known as the liquidation period.

If there is a portfolio consisting of a single position in N of the risk factor S, such that the factor N is the sensitivity of V with respect to S, then given that δV is a linear function of δS, the value change δV over the period δt is given directly from equation 4.7 as

$$\begin{aligned} \delta V &= NS(t+\delta t) - NS(t) && (4.8) \\ &= NS(t)\left[\exp(\mu dt + \sigma \delta W) - 1\right] && (4.9) \end{aligned}$$

which using the definition of VAR and recalling that the only stochastic variable is $\delta W \sim N(0, \delta t)$, gives upon further manipulation and simplification

$$\begin{aligned} CP_{\delta V}(\delta V \leqslant -VAR) &= CP_{\delta V}\left(\delta W \leqslant \frac{\ln\left(1 - \frac{VAR}{NS(t)}\right) - \mu \delta t}{\sigma}\right) && (4.10) \\ &= CP_{\delta V}\left(\delta W \leqslant a\sqrt{\delta t}\right) && (4.11) \end{aligned}$$

where

$$a := \frac{\ln\left(1 - \frac{VAR}{NS(t)}\right) - \mu\delta t}{\sigma\sqrt{\delta t}} \tag{4.12}$$

So that the probability that δW is less than or equal to any given value depends solely on the distribution of δW, such that $CP_{\delta V}$ can simply be replaced by $CP_{\delta W} = (X \leqslant a)$ which in turn becomes $CP_X = N(0,1)$

$$CP_{\delta V}(X \leqslant a) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{a} \exp\left(0.5x^2\right) dx \tag{4.13}$$

so that a therefore represents the $(1-c)\,th$ percentile of the standard normal distribution, $Q_{(1-c)}^{N(0,1)}$, such that the solution for VAR of a long position in N risk factors is given by

$$VAR(c) = NS(t)\left[1 - \exp\left(\mu\delta t + \left(Q_{(1-c)}^{N(0,1)}\right)\sigma\sqrt{\delta t}\right)\right] \tag{4.14}$$

[3] where the long and short positions will not be equal due to the presence of the drift μ and the fact assumption that changes are assumed to be lognormally and therefore asymmetrically distributed. Practical implementation of VAR is frequently based on the assumption of short liquidation periods with 10 days being common among many institutions (i.e. $\delta t = 10/365 = 0.0274$ years), which in turn means that $\exp(x) \approx 1 + x$ and that the drift can effectively be set to zero and ignored, which produces the following possible simplifications

$$\delta S(t) \approx \left\{ \begin{array}{ll} S(t)\left[\exp(\sigma\delta W) - 1\right] & \rightarrow \text{assuming } \delta t \approx 0 \\ S(t)\left[\mu\delta t + \sigma\delta W\right] & \rightarrow \text{a linear approximation for exp} \\ S(t)\sigma\delta W & \rightarrow \mu = 0 \text{ and linear approximation for exp} \end{array} \right\} \tag{4.16}$$

which simplifies the expressions for the VAR of a long position in N risk factors S to

$$VAR(c) \approx \left\{ \begin{array}{ll} NS(t)\left[1 - \exp\left(\sigma\sqrt{\delta t}Q_{(1-c)}^{N(0,1)}\right)\right] & \rightarrow \text{assuming } \mu \approx 0 \\ NS(t)\left[-\mu\delta t - \sigma\sqrt{\delta t}Q_{(1-c)}^{N(0,1)}\right] & \rightarrow \text{a linear approximation for exp} \\ -NS(t)\sigma\sqrt{\delta t}Q_{(1-c)}^{N(0,1)} & \rightarrow \mu = 0 \text{ and linear approximation for exp} \end{array} \right\} \tag{4.17}$$

Note that long and short VAR are only equal and offsetting in the third of the cases above and in the case where changes in the value of the underlying portfolio are approximately linear in the underlying

[3] Note the the VAR of a short position in $-N$ risk factors S is given by

$$VAR(c) = -NS(t)\left[1 - \exp\left(\mu\delta t - \left(Z_{(1-c)}^{N(0,1)}\right)\sigma\sqrt{\delta t}\right)\right] \tag{4.15}$$

risk factor (known as the Delta-Normal approximation).

In practice, there is often far more than a single risk factor to be considered, which implies the use of a covariance matrix for the coupled stochastic differential equations

$$d\ln S_i(t) = \mu_i dt + dZ_i \quad \text{for } i = 1, 2, ..., n \text{ factors} \tag{4.18}$$

where the dZ_i are correlation zero-drift Brownian motions with covariance matrix

$$cov\left[dZ_i, dZ_j\right] = d\Sigma_{ij}, \quad E\left[dZ_i\right] = 0 \tag{4.19}$$

with

$$d\Sigma_{ij} = \rho_{ij}\sigma_i\sqrt{\delta t}\sigma_j\sqrt{\delta t} = \rho_{ij}\sigma_i\sigma_j\delta t \tag{4.20}$$

such that

$$d\mathbf{Z} = \begin{pmatrix} dZ_1 \\ dZ_2 \\ \vdots \\ dZ_n \end{pmatrix} \sim N(0, d\Sigma) \tag{4.21}$$

As previously, the solution to the resulting stochastic differential equations are of the form

$$S_i(t+\delta t) = S_i(t)\exp\left(\mu_i\delta t + \delta Z_i\right) \qquad \text{with } \delta\mathbf{Z} \sim N(0, \delta\Sigma) \tag{4.22}$$

where the covariance matrix $\delta\Sigma$ is given by

$$\begin{pmatrix} \delta\Sigma_{11} & \delta\Sigma_{12} & \cdots & \cdots & \delta\Sigma_{1n} \\ \delta\Sigma_{21} & \ddots & & \ddots & \delta\Sigma_{2n} \\ \vdots & & \delta\Sigma_{ij} & & \vdots \\ \vdots & \ddots & & \ddots & \vdots \\ \delta\Sigma_{n1} & \delta\Sigma_{n2} & \cdots & \cdots & \delta\Sigma_{nn} \end{pmatrix} \quad \text{where } \delta\Sigma_{ij} = \rho_{ij}\sigma_i\sigma_j\delta t \tag{4.23}$$

VAR generally begins with a series of factors that have been identified as the principal contributors to risk. The problem is then to transform these initial (assumed to be) independently identically distributed values into correlated variables with covariance matrix $\delta\Sigma$. The usual approach to achieving this is to employ Cholesky decomposition, which proceeds as follows. Consider the matrix A which has the following property

$$AA^T = \delta\Sigma \tag{4.24}$$

(where A^T is the transpose of A), so that A transforms the initial uncorrelated variables into correlated variables with covariance $\delta\Sigma$. The matrix A is then derived via Cholesky decomposition by solving for A_{ji} for all j, beginning with $i = 1$ and $j = 1$

$$A_{ij} = \begin{pmatrix} 0 & \text{for } j < i \\ \sqrt{\delta\Sigma_{ii} - \sum_{k=1}^{i-1} A_{ik}^2} & \text{for } j = i \\ \frac{1}{A_{ii}}\left(\delta\Sigma_{ji} - \sum_{k=1}^{i-1} A_{ik}A_{jk}\right) & \text{for } j > i \end{pmatrix} \text{ where } \delta\Sigma_{ji} = \begin{pmatrix} \sigma_i^2\delta t & \text{for } j = i \\ \rho_{ij}\sigma_i\sigma_j\delta t & \text{for } j \neq i \end{pmatrix} \tag{4.25}$$

The results of the decomposition then provide sufficient information to construct correctly correlated Monte Carlo paths that can be used to simulate the behaviour of the VAR for the chosen portfolio. One of the most popular methods of doing this is to use the method known as the delta-normal method, which proceeds as follows. First calculate the sensitivities of the portfolio with respect to all required risk factors S_i, call these Δ_i. Next, multiply the covariance matrix $\delta\Sigma$ with the portfolio sensitivities Δ_i and the risk factor values S_i, to derive the portfolio variance

$$var\,(\delta V) = \delta t \sum_{i,j=1}^{n} \Delta_i \rho_{ij} \sigma_i \sigma_j \Delta_j \tag{4.26}$$

Now multiply the above portfolio variance with the liquidation period and the square of the percentile of the required confidence level for the standard normal distribution. Upon taking the square root, the VAR is given by

$$VAR_V\,(c) = Q_{1-c}\sqrt{\delta t}\sqrt{\sum_{i,j=1}^{n} \Delta_i \rho_{ij} \sigma_i \sigma_j \Delta_j} - \delta t \sum_i \Delta_i \mu_i \tag{4.27}$$

where the final term is the drifts of the risk factors.

In order to utilise VAR in practice, it is necessary to have a precise prediction of the probability of an extreme change in the value of a portfolio of instruments. Extreme movements are closely related to the tails of a distribution, using almost any reasonable definition of the tails of a generating process for a distribution. Indeed, beginning with the ground breaking work of Mandelbrot in 1963, almost all studies of financial data have indicated that such series exhibit fat tails. Defining a fat tail is not without problems, but a reasonable definition is what constitutes a thin tail is that a distribution can be said to have thin tails if the density reaches zero before a finite quantile[4].

[4]This should not be taken to imply that the normal is a fat-tailed distribution. On the contrary, it is simply that there is substantial evidence that many time series associated with financial and insurance claim data are fat-tailed, with a (much) higher probability of " outliers' compared with the normal distribution. However, standard tests, or variants of them, for the presence of unit roots assume a normal distribution for the innovations driving the series. Application of the former to the latter therefore involves an inconsistency.

Many studies of fat-tailed distributions have used distributions such as the log-normal, generalised error and mixtures of normal distributions. Unfortunately, all of these are thin-tailed according to the simple definition provided in the preceding paragraph since the tails of these distributions decay exponentially, notwithstanding the fact that they exhibit kurtosis in excess of the normal distribution. In many practical cases the distributions just mentioned fit the empirical data quite well up to reasonable quantiles, but the fit deteriorates rapidly in the higher quantiles or extremes.

4.2.2 Extreme value theory

Until relatively recently, VAR was considered to be the standard for those managing risk. In its simplest form, VAR is typically focused on the issue of the size of the maximum potential loss assuming either a normal or lognormal distribution. The key weakness is thus that VAR is based around an assumption of "business as usual" type losses - it is primarily concerned with the centre of the returns distribution. However, risk management tends to be concerned with low probability events in the tails of the returns distribution. The further into the tails of the distribution, the smaller is the probability of occurrence of an event, but the larger will be its consequences. Modelling the extremes of a standard normal or lognormal distribution is therefore relatively uninsightful, as there are few expected events in the extremes and the tails are small proportions of the whole distributions.

In contrast, Extreme Value Theory (EVT), as the name suggests, is totally concerned with the tails of the distribution, as its main aim is to provide asymptotic models for the tails of a distribution. EVT has been around for some time in the insurance industry. Its origins can be traced to the seminal theoretical work on block maxima by Fischer and Tippett (1928) and work on extremes of distributions by Gumbel (1958). More recent research by Balkema and de Haan (1974) and Pickands (1975) has focused on threshold based extreme value. In contrast, work by McNeil and Frey (2000) make the interesting extension to VAR of using EVT to generate the data for a VAR-type analysis.

Where VAR has proved the default method of choice for the finance industry, EVT has effectively occupied the same role in the insurance industry, but is increasingly being applied in the finance industry as the deficiencies of VAR become a major issue due to regulatory issues and the considerations of practical business management. The research on both theoretical aspects and practical applications of EVT is considerable, but the 1999 text by Embrechts et al probably remains the definitive reference.

At its simplest, EVT has two significant results that must be outlined before any discussion of the approach can take place. The first is that the asymptotic distribution of a series of extrema (under certain conditions) can be modelled such that the distribution of the standardised extrema of the series can be shown to converge to the Gumbel, Frecet or Weibull distributions, where the generalised extreme value

(GEV) distribution represents the general form. The second is that the limiting distribution for excesses over a given threshold is the generalised Pareto distribution (GPD). It is worth noting that some EVT techniques can be used to solve for very high quantiles - a fact which can prove very useful for predicting extreme loss situations.

The key principles of EVT can therefore be described as follows. Take a series of independent, identically distributed observations $X = (X_1, X_2, ..., X_n)$, with a distribution function, F, that maybe unknown, then $M_n = \max\{X_1, X_2, ..., X_n\}$ is the sample maxima. In insurance, sub-exponential or heavy-tailed distributions are the standard way of dealing with individual claim sizes, their defining property being

$$\lim_{x \longrightarrow \infty} \frac{P(X_1 + X_2 + ... + X_n > x)}{P(\max(X_1, X_2, ..., X_n) > x)} = 1 \tag{4.28}$$

for every $n \geqq 2$, so that the tails of the distribution of the sum and of the maximum of the first n claims are asymptotically of the same order, which implies that the largest claim has a significant effect on the total amount of all claims. The standard generalised extreme value distribution (GEV), $H_{\xi,\mu,\sigma(x)}$, is given by

$$H_{\xi,\mu\sigma(x)} = \left\{ \begin{array}{ll} \exp\left[-\left(1+\xi\frac{x-\mu}{\sigma}\right)^{-1/\xi}\right] & \text{if } \xi \neq 0, 1+\xi\frac{x-\mu}{\sigma} > 0 \\ \exp\left[-\exp(-\frac{x-\mu}{\sigma})\right] & \text{if } \xi = 0 \end{array} \right\} \tag{4.29}$$

where μ, σ and ξ are the location, scale and shape parameters respectively, with different values of ξ corresponding to different distributions[5]. The GEV is used to estimate values of X outside of the range of the existing data using either extreme events or exceedences of a specific level. The tail of the population is usually assumed to follow some form of the GEV.

Work by Pickands (1975) shows that the generalised Pareto distribution (GPD) is the limit distribution of excesses $Y := \max\{X - u, 0\}$ over sufficiently high threshold limits u and offers an acceptable approximation of the tail of F for some fixed ξ and β. Therefore, the distribution of Y is in effect the conditional distribution of X given $X > u$, with ξ as the shape parameter and β as the scale parameter

$$GPD_{\xi,\beta(y)} = \begin{array}{ll} 1 - \left(1+\xi\frac{y}{\beta}\right)^{-1/\xi} & \text{if } \xi \neq 0 \\ 1 - \exp\left(-\frac{y}{\beta}\right) & \text{if } \xi = 0 \end{array} \tag{4.30}$$

[5]The three individual distributions are:
Gumbel: $\xi = 0$
Frechet: $\xi = 1/\alpha$, $1/\alpha > 0$
Weibull: $\xi = -1/\alpha$, $1/\alpha < 0$

where

$$y = \begin{matrix} [0,\infty] & \text{if } \xi \geqslant 0 \\ \left[0, -\frac{\beta}{\xi}\right] & \text{if } \xi < 0 \end{matrix} \quad {}_{6} \tag{4.31}$$

It can then be shown that the mean excess function of the GPD is

$$e(u) = E(X - u | X > u) = \frac{\beta + \xi u}{1 - \xi} \tag{4.32}$$

where

$$\beta = \sigma + \xi (u - \mu) \tag{4.33}$$

such that the $\max_n Y_n$ follows a GEV distribution with parameters ξ, μ and σ.

Having defined the key theoretical pillars of EVT, the next step involves the construction of the series of extreme values. Two common approaches have dominated thinking and practice. The first is the block-extrema method, which involves dividing the data into a series of non-overlapping blocks of identical length and selecting the extreme value from each block. This approach effectively validates the assumption that the extreme observations are independent and identically distributed. This situation is common in the finance industry where a period of high volatility is frequently followed by a period of low volatility - a phenomenon usually referred to as volatility clustering. Increasing the block size mitigates the problem, but risks losing the extreme values within the block, thereby making block size selection highly subjective and therefore problematic.

The second approach is the so called peaks over threshold method and consists of selecting a threshold value over which extreme values are chosen. The problem is that choosing the threshold is obviously a case of trading variance for bias. To see this, consider the situation where, by increasing the number of obsevations for the series of extrema (which implies lowering the threshold value), it is inevitable that some observations from the centre of the distribution will be included, thereby increasing the precision of the tail estimate, but also simultaneously increasing estimation bias. Whilst if a high threshold is selected, bias falls but index volatility increases because of the fall in the number of observations. A further problem is that there may also be dependent observations, but Resnick and Starica (1996) suggest that a way around such a problem is to standardise the observations in order to fit the various parameters.

More recent work by Bystrom (2001) extends the EVT approach by making the further distinction

[6] There are three cases of ξ to consider:

1. $\xi < 0$: tail of F belongs to a collection of heavy-tailed distributions such as Pareto, log-Gamma and Cauchy, and behaves like a decaying power function $x^{-1/\xi}$.

2. $\xi = 0$: tail of F belongs to a collection of distributions with more or less "medium" tails such as normal and log-normal distributions.

3. $\xi > 0$: tails of F has a finite right end point, so that the distribution belongs to a collection of short tailed distributions such as the beta distribution.

between conditional and unconditional distributions, arguing that the latter are more suited to longer term investment decisions and the examination of very rare stress events, whereas the former are more applicable when examining day-to-day risks and short term risk management issues. Bystrom extends the earlier work by McNeil and Frey (2000) by using conditional threshold based EVT models to forecast VAR measures by first filtering the raw data using an AR(1)-GARCH model. Both studies find there to be little difference between the threshold and block maxima approaches when employed in their respective frameworks. Bystrom provides the more interesting perspective as he performs his estimations for both tranquil and volatile market periods.

4.2.3 Minimax risk management

In contrast to both VAR and EVT, neither of which are founded on the basis of optimality, Rustem and Howe (2002) consider the problem of optimal choice in the presence of uncertainty when specifically faced with the worst case scenario. Worst-case risk management seeks to find the best possible outcome in the face of the worst possible situation. The objective is usually expressed in terms of some form of a cost or penalty function, so that stated more formally, worst-case risk management seeks to simultaneously determine the minimum of the cost function under the maximum or worst case scenario - hence the term minimax. Optimality is therefore expressed over all possible values of uncertainty. This is an important point as it distinguishes the minimax approach from both VAR and EVT, both of which are only concerned with performance against some arbitrary limit - being 95% certain that losses will not be greater than some desired level in the case of VAR, for example.

At a generic level, the maximum inner function is often couched in terms of a disutility or error function, such that the outer minimisation involves searching over the outcomes associated with the worst-case disutility scenario in order to find the best possible alternative. Stated in mathematical terms, the worst-case problem can be stated as

$$\min_{x \in \Re^n} \max_{y \in \mathcal{Y}} f(x, y) \tag{4.34}$$

where x is a vector of decision variables (represented by real numbers in n-dimensional Euclidean space, $\Re^n$) and y is a vector of uncertain variables defined over the feasible set $\mathcal{Y}$, with the solution being either discrete or continuous depending on whether $\mathcal{Y}$ is a discrete or continuous set. An equivalent, slightly more convenient representation is

$$\min_{x \in \Re^n} \Phi(x) \tag{4.35}$$

where

$$\Phi(x) = \max_{y \in \mathcal{Y}} f(x, y) \tag{4.36}$$

So that for the solution x^*

$$\Phi(x^*) = \max_{y \in \mathcal{Y}} f(x^*, y) \geqslant f(x^*, y), \quad \forall_y \in \mathcal{Y} \tag{4.37}$$

which states that the performance of the solution x^* is guaranteed to be non-inferior for any y. This is the specific feature that provides robustness from the minimax solution and ensures that performance will be better if the worst-case scenario is not realised.

From a financial risk management perspective, minimax has been employed in two main ways, namely, in discrete form as a robust strategy for discrete rival scenarios and in continuous form in problems such as option hedging. Analysing the former application first offers a slight advantage as it provides a framework for dealing with a discrete set of possible scenarios. Minimax thus arises from the ability to reduce the set of alternatives to single possibility, such that optimality is not determined by a single scenario, but simultaneously over all scenarios. Work by Rustem (1987, 1994) on policy optimisation examined the pooling of objective functions from rival models to generate an optimal policy based solely on a single model and then evaluates its impact if the second model proves to be the correct representation of the underlying system.

Chow (1979) was the first to experiment with robustness and competing models. His work uses two models. The first model is used to generate an optimal policy, the performance of which is then analysed using the second model should the second model actually prove to be the correct representation of the system. Chow constructs a sort of payoff matrix for alternative strategies, so that the optimal strategy can be chosen based on the model that causes the least damage when used on the other model - the model is thus sequential. Rustem and Howe (2002) extend the approach of Chow by carrying out the calculation of the worst-case scenario simultaneously with the minimisation over x. The result is that policy choice is no longer constrained to be dependent on a single model as optimality can be based on more than a single model.

Rustem and Howe introduce the discrete minimax strategy by considering the pooling of rival objective functions using a vector of fixed pooling weights ϑ such that

$$\vartheta \in \mathbb{E}_{+}^{m^{sce}} \equiv \left\{ \vartheta \in \Re^{m^{sce}} \mid \vartheta \geqslant 0, \ \langle \mathit{1}, \vartheta \rangle = 1 \right\} \tag{4.38}$$

where $1 \in \Re^{m^{sce}}$ denotes the vector with every component equal to unity[7]. Then proceed to formulate the optimal decision problem as a constrained optimisation of the pooled objective functions subject to non-linear constraints

$$\min_{x} \{\langle \vartheta, f(x) \rangle \mid g(x) = 0,\ h(x) \leqslant 0\} \tag{4.39}$$

where all of the functions are twice differentiable in their relevant domains[8]. In equation 4.39 each element of f represents a separate objective function that corresponds to an alternative model or scenario, with the number of models or scenarios being less than the number of decision variables. In Rustem and Howe's model, robust pooling corresponds to the strategy that is supposed to be invariant to the scenario or model that actually turns out to represent the system and is given by the solution to the following minimax problem

$$\min \max \left\{\langle f(x), \vartheta \rangle \mid g(x) = 0,\ h(x) \leqslant 0,\ \vartheta \in \mathbb{E}_{+}^{m^{sce}}\right\} \tag{4.40}$$

which in turn ensures that the worst-case scenario is calculated at the same time as the minimisation over x and ensures robustness[9]. It is then therefore the case that irrespective of the model that turns out to be the correct representation of the system, the minimax strategy ensures that the value produced by the objective function will never deteriorate, because it will always be at least as good as the minimax value. If the minimax strategy turns out to be too cautious and/or costly, then the selection of the optimal policy could be based on equation 4.39 so that ϑ is selected in the near neighbourhood of the minimax value such that the policy is as robust as possible. The drawback of such a sub-optimal approach is that the selected policy has robustness that is limited to the alternative representations of the system. If a new alternative model arises then robustness can not be guaranteed.

The second application of minimax that bears on the problem of robust risk management is the work by Howe, Rustem and Selby (HRS, 1994 and 1996) which applies continuous minimax to provide a robust

[7] Where the *sce* superscript denotes scenario. $\vartheta \in \mathbb{E}_{+}^{m^{sce}}$, $x \in \Re^{n}$, $f : \Re^{n} \rightarrow \Re^{m^{sce}}$, $g : \Re^{n} \rightarrow \Re^{e}$, $h : \Re^{n} \rightarrow \Re^{i}$ and f, g and h are all twice differentiable functions.

[8] In the above equation, each element of f (f^{j}) represents a rival objective function corresponding to the jth rival model or scenario. In Rustem and Howe, the restriction $g(x) = 0$ and $h(x) \leqslant 0$ are the equality and inequality constraints imposed on the overall decision problem. Note that the number of scenarios or models is generally far fewer than the number of decision variables, i.e. $n \gg m^{sce}$.

[9] Rustem and Howe define robustness for minimax in terms of guaranteed performance, non-inferiority according to the following criteria in terms of equation 4.40:

1. There exists a minimax solution (x^{*}, ϑ^{*}) with associated mulltipliers $(\mu_{e}^{*}, \mu_{i}^{*}, \lambda^{*}, \eta^{*})$.
2. $f(x), g(x), h(x) \in \mathbf{C}^{1}$ at x^{*}.
3. $\vartheta^{*} = 0 \Rightarrow \lambda^{*} > 0$ and $\vartheta^{*} > 0 \Rightarrow \lambda^{*} = 0$.
4. Then for $i, j, \ell \in \{1, 2, ..., m^{sce}\}$:

- $f_{i}(x^{*}) = f_{j}(x^{*}), \forall i, j\,(i \neq j)$ iff $\vartheta_{i}^{*}, \vartheta_{j}^{*} \in (0,1)$
- $f_{i}(x^{*}) = f_{j}(x^{*}) > f_{\ell}(x^{*}), \forall i, j, \ell\,(\ell \neq i, j)$ iff $\vartheta_{\ell}^{*} = 0$ and $\vartheta_{i}^{*}, \vartheta_{j}^{*} \in (0,1)$
- $f_{i}(x^{*}) > f_{j}(x^{*}), \forall_{j}\,(j \neq i)$ iff $\vartheta_{i}^{*} = 1$
- $f_{i}(x^{*}) < f_{j}(x^{*}), \forall_{j}\,(j \neq i)$ iff $\vartheta_{i}^{*} = 0$

hedging strategy as a replacement for standard delta hedging for the seller of call options. The objective of HRS's continuous minimax strategy is stated in terms of minimisation of the effects of a pre-defined worst-case hedging error and expressed in terms of bounds on the price of the underlying asset for a risk averse writer of call options. The minimax strategy is stated with respect to a given time horizon and rebalancing of the hedge is specifically incorporated. HRS use Leland's 1985 model to explicitly incorporate hedging costs via an adjustment to volatility as follows

$$\hat{\sigma} = \sqrt{\sigma^2 \left[\frac{2}{\pi} \frac{K}{\sigma\sqrt{\Delta t}}\right]} \tag{4.41}$$

where K is the round-trip hedging transaction cost, expressed as a percentage of trading volume and defined as any complete transaction comprising a purchase followed by a sale (or vice versa). In the BS model, the call writer can construct a riskless portfolio by continuously rebalancing the hedge to remain delta neutral. Continuous rebalancing is supposed to keep the size of hedging errors very low, thereby keeping hedging cost very low. However, in practice only discrete rehedging is possible, so non-trivial hedging errors are likely to occur and it is minimisation of these errors which forms the basis of the minimax hedging strategy. HRS adopt the Merton (1973) no-arbitrage argument to construct the basic self-financing and non-stochastic (in terms of return) hedging portfolio, referred to as the "ideal portfolio", containing positions in the underlying and the riskless bond. The ideal portfolio is used as the benchmark in defining the objective function. HRS derive basic properties for the minimax hedging strategy on the basis of the ideal portfolio and conditional upon those results add the consequent costs.

The hedging error, $\mathcal{HE}$, over the period from t to $t+1$ for a portfolio comprised of a short position in a call option and a long position in the underlying (stock) is

$$\mathcal{HE} = N\left(B_t - B_{t+1}\right) + x_t\left(y_{t+1}^S - y_t^S\right) \tag{4.42}$$

where B_t is the Black-Scholes call price at time t. The minimax hedging strategy seeks to minimise the maximum potential hedging error over the hedge period, so that the objective function used is thus the potential hedging error. Following the spirit of Leland, HRS include the cost of funding the long position which may arise if the worst case does occur in the calculation of the maximum potential hedging error. If the worst case fails to materialise, the hedging error will be less than indicated by the minimax solution.

HRS present their minimax hedging framework within the context of stock options so that the minimising variable in which the hedge is expressed is the number of shares, x and the maximising variable is y_t^S the stock price at time t

$$\min \max f\left(x, y_{t+1}^S\right) \tag{4.43}$$

(where y_t^S is permitted to take any value within predefined bounds) subject to

$$y_{t+1}^{S,Lower} \leqslant y_{t+1}^S \leqslant y_{t+1}^{S,Upper} \tag{4.44}$$

and short sales are explicitly permitted. HRS examine two possible worst-case stock price scenarios. The first sets the upper and lower bounds as being plus or minus two standard deviations about the expected value of the stock price at $t+1$ - assuming a 95% confidence limit. Whereas the second scenario is set in terms of 1 to 3 standard deviations around the at the money forward stock price in order to capture the area of greatest elasticity in value of the option.

The minimax objective function used by HRS is

$$f\left(x_t, y_{t+1}^S\right) = \frac{1}{2}\left\langle U - U^d, Q\left(U - U^d\right)\right\rangle \tag{4.45}$$

where $U^d \in \Re^{k+1}$ is the vector of potential hedging errors which is set to zero to be be consistent with delta hedging. The remaining variables are defined as follows

$$x_t = \begin{bmatrix} x_{1,t} \\ \vdots \\ x_{k,t} \end{bmatrix} \quad \text{and} \quad y_{t+1}^S = \begin{bmatrix} y_{1,t+1}^S \\ \vdots \\ y_{k,t+1}^S \end{bmatrix} \tag{4.46}$$

$$U\left(x_t, y_{t+1}^S\right) = \begin{bmatrix} U_1\left(x_t, y_{t+1}^S\right) \\ \cdots \\ U_2\left(x_t\right) \end{bmatrix} \quad \text{and} \quad U^d = \begin{bmatrix} U_1^d \\ \cdots \\ U_2^d \end{bmatrix} = \begin{bmatrix} 0 \\ \cdots \\ 0 \end{bmatrix}^{10} \tag{4.47}$$

where

$$U_1\left(x_t, y_{t+1}^S\right) = \sum_{i=1}^{k} x_{i,t}\left(y_{t+1}^S - y_t^S\right) + \sum_{i=1}^{k} N_i\left(B_{i,t} - B_{i,t+1}(y_t^S)\right) \tag{4.48}$$

$$+\sum_{i=1}^{k}\left(-\left(x_{i,t} - x_{i,t-1}\right) y_{i,t}^S + C_{i,t-1}\left(1 + r\Delta t\right)\right) r\Delta t \tag{4.49}$$

where U_1 is the potential hedging error (including funding costs) between time t and $t+1$ and contains the potential shift in the underlying, the potential shift in the option position; and where

$$C_{i,t-1} = C_{i,t-2}\left(1 + r\Delta t\right) - \left(x_{i,t-1} - x_{i,t-2}\right) y_{i,t-1}^S - \hat{K}\left|\left(x_{i,t-1} - x_{i,t-2}\right) y_{i,t-1}^S\right| \tag{4.50}$$

[10] ℜWhere HRS define their variables as follows: $U_1 : \Re^k \times \Re^k \rightarrow \Re^1$, $U_2 : \Re^k \rightarrow \Re^k$, $U : \Re^k \times \Re^k \rightarrow \Re^{k+1}$, $x_t \in \Re^k$, $y_{t+1}^S \in \Re^{k+1}$ and Q is a $(k+1) \times (k+1)$ positive-definite weighting matrix.

and

$$U_2(x_t) = \begin{bmatrix} U_{1,2}(x_{1,t}) \\ \vdots \\ U_{k,2}(x_{k,t}) \end{bmatrix} \quad \text{where} \quad U_{i,2}(x_{i,t}) = \hat{K}(x_{i,t} - x_{i,t-1})\, y^S_{i,t} \tag{4.51}$$

and where $C_{i,t-1}$ is the cumulative net cashflow at time $t-1$, which when multiplied by $(1 + r\Delta t)$ includes interest costs[11]. Introducing transactions for the current period would mean non-differentiability in the objective function, so HRS use $U_{i,2}$ as a proxy for transaction costs for option i at time t instead. Q is a diagonal $(k+1)$ x $(k+1)$ positive definite weighting matrix, with a high q_1 implying a strong preference for minimising potential hedging error between t and $t+1$, whilst a high q_i $(i = 2, ..., k+1)$ showing a preference for minimising the penalty term. The optimisation then proceeds for each option $i = 1, ..., k$ $B_{i,t+1}\left(y^S_{i,t}\right)$ using the modified volatility estimate of equation 4.41. The objective is to find the mix of the risk free asset and underlying that minimises the deviation of the return on the hedge portfolio (inclusive of costs) from the return on the ideal portfolio (which will be zero according to Merton's no arbitrage arguments).

HRS implement the minimax approach in five different variants and compare the results with standard delta hedging using a structured simulation. The procedure generates paths that, in addition to the usual raw random values also include two other types of values. The first are termed cross over events which are selected from the simulated values such that the successive random values straddle the option strike; with the second type of event being called abrupt change events. HRS use the following expression to weight the number of units of the underlying recommended by the minimax strategy to reflect the estimation of the level of noise contained in the volatility and the expectation of the possibility of reversal of a cross-over event

$$k_1 = \frac{\left[w_1 \frac{|\Delta y^S_t|}{a^* sd} + w_2 \frac{|y^S_t - X|}{b^* sd}\right]}{(w_1 + w_2)} \tag{4.52}$$

where setting w_2 to a non-zero value happens when a cross-over event occurs during the period $t-1$ and t[12].

Of the minimax strategies compared, the most interesting is the version that applies the expression in equation 4.52 to weight the minimax hedge prescription by a factor between 0 and 1 in an attempt to represent the assessment of the information contained in the volatility of the underlying stock. The idea is to provide a simple mechanism to filter out noise, so that the hedger only responds to the pure signal

[11] Note that when setting up a minimax problem, it is possible to adopt any desired convenient target value. HRS use zero on the grounds that when delta hedging, the expected value of the hedging error is zero.

[12] The hedger can attach any weights to a and b to reflect the importance attached to sd (the standard deviation) of y^s_t. Note that Rustem and Howe require a non-zero value be applied to $w2$ when a cross-over occurs from $t-1$ to t, but there are no restrictions on the value of $w2$.

component of the volatility, with a higher level of volatility being assumed to be associated with greater noise. This clearly represents the beginnings of introducing feedback control and closed loop modelling of the relationship between volatility and hedge management, offering a degree of robustness for successive states. However, it does not take account of the dynamics involved in the state transition process and lacks precision in both theory and practice. This is not surprising, given the fact the HRS clearly state that their focus is static optimisation - in the case of American exercise bond options considered by HRS, the formulation is effectively a single-period minimax optimisation using values generated by the usual expectations type framework using a binomial or trinomial tree. On the basis of their empirical work, HRS conclude that the weighted minimax strategy produces robustly superior static performance compared with the standard delta hedging method, particularly in periods of high volatility and where the spot price exhibits frequent cross-over behaviour for at the money options. Whilst such a conclusion is certainly supportable based on their work, it is certainly not the whole story.

The work of HRS is based firmly in the realm of static optimisation, to which some drawbacks have already been alluded and reviewed. However, it is instructive to summarise the principal issues as a precursor to reformulating the objective function for uncertainty management within the area of robust optimal control. The following are the essential weaknesses of the minimax approach

- The fundamental premise is the static behaviour of the underlying system. Clearly, the behaviour of the underlying, be it a single stock or bond, or a portfolio of such instruments, is not well described by a static approach because the dynamics of the relationships or of the impact of time on relationships is not correctly handled in a purely static environment - even in its multi-period form minimax is a static model that is just too simplistic for a complex risk environment.

- The application of static minimax optimisation results in considerable unmodelled dynamics; the problem is that minimax takes no account of system dynamics and feedback or controller design, loop-shaping, closed v open loop, robust design/implementation/operation. The principal reason to use feedback is to reduce the effects of uncertainty which arise from modelling error or from an unmodelled disturbance or noise. Yet it is precisely such phenomena the minimax ignores.

- Does not deal with the occurrence or impact of alternative types of perturbation, so that the notion of robustness is therefore limited.

- Even with HRS's heuristic weighting adjustment mechanism, there is a lack of modelling of the feedback process between risk policies, subsequent information and actions.

- There is no single variable that provides an unambiguous measure to express the degree of robustness of a risk measure.

For ease of use, we shall refer to the Rustem and Howe approach as discrete-minimax (DM) or continuous minimax (CM).

4.3 Robust optimal control and risk management

Having examined the bases and weaknesses of VAR, EVT and DM/CM it is now appropriate to construct an alternative that attempts to address a number of the limitations of these approaches. Construction of such an alternative approach requires a fundamental consideration of the process by which decisions about managing exposures (both at the level of individual instruments and portfolios) are made in the presence of uncertainty. In order to be able to analyse this process and make the construction of a model more systematic, it is useful to distinguish three distinct classes of processes to which an alternative risk management methodology should conform.

First, are the processes or rules by which search is conducted - the search rules. If the search process is to be conducted step-wise (either discretely or continuously), whereby a piece of information is acquired or an adjustment is made (either dynamically in closed-loop mode as part of transition between states, or statically within a given state), what are the necessary and sufficient search rules and how should they be constructed ? Given the inherent need to make assumptions about the quality and quantity of the information that is available in each state, then taking account of the weaknesses of the three approaches considered above clearly involves specifying how feedback should be incorporated and modelled. This is of particular significance when considering the form of both the objective function and the types of constraints - as is made clear below.

Second, are the rules required to terminate the search for the most desirable choice among the possible alternatives - the stopping rules. If the rules are simple, as in the case of VAR, then no notion of optimisation is either necessary in the formulation of the underlying axioms of the theory or required in the actual execution of the search. However, the limitations of VAR have already been demonstrated, making it clear that simple heuristics do not constitute an acceptable framework for making decisions that will have robust and stable results in the face of complex uncertainties. Whilst EVT demonstrates the fundamental need to focus on extremes in the decision making process and DM/CM underline the requirement for optimisation, it remains clear that robustness and stability in both theory and practice need to be incorporated into the decision making process.

The third class of processes involved are those that govern the making of decisions once information has been acquired and a stopping rule has been applied - the decision rules. In the case of the static optimisation used in the minimax approach, the search is limited to the pre-defined universe of scenarios. Optimisation is carried out with respect to the given linear objective function, subject to the stated

constraints. The constraints and the objective function taken together determine the way in which the decision rules are constructed and applied. However, due to the absence of feeback in the minimax approach, the decisions produced by the optimisation need to be manually adjusted by a simple heuristic. This deprives the approach of its internal consistency and is likely to be non-optimal in practice.

It is therefore clear that explicit mathematical formulation is required for all relationships in the model - objective function, feedback, state variables, control variables and perturbations, in order for an alternative framework to constitute a viable practical alternative to existing approaches. The key aspect underlying construction of any approach to dealing with such uncertainty is how to construct a model that deals realistically but tractably with model, data and parameter uncertainties. At the simplest end of the scale, VAR adopts the conventional, essentially single period static approach to making decisions under uncertainty by basing its premises and predictions on expected value optimisation. The unfortunate and inevitable problem with such a methodology is that it neglects the worst-case effect of uncertainty in favour of expected values. This may very well be acceptable in some circumstances, particularly where the prime objective is generally accepted to be the estimation of possible losses based on disruption to business-as-usual distributions. However, decisions based on expected value static optimisation may require justification in the context of the worst-case scenario - a situation particularly prevalent in the case of managing risks associated with such derivatives as those based on catastrophes or credit defaults, where omitting to hedge against the possibility of catastrophic events could result in ruin.

Conversely, given a source of uncertainty, some worst-case realisations may be thought to be so improbable that even considering them as possible eventualities may be both impractical and unacceptable for many institutions. Notwithstanding this view, it remains the case that for instruments where extreme behaviour of the underlying is the very rationale for the creation of the instrument, it is essential to consider worst-case scenario based optimisation in order to attempt to ensure some degree of robustness in risk management decisions, even if such optimisation acts only as a benchmark. As was illustrated by Rustem and Howe in the case of DM, the minimax optimisation approach is non-inferior for any scenario and better for all other cases apart from the worst-case as it ensures guaranteed optimality for the worst-case. In the case of CM, optimal performance is guaranteed due to the continuum of scenarios. The critical weakness of minimax, however, is that even with a continuum of scenarios it still remains the case that optimality is only guaranteed for any given series of scenarios and no account is taken of the dynamics driving the system between states.

There are a number of possible ways of developing a new approach to the problem at hand. One possibility would be to simply state a new model, then justify it step by step, explaining the reasons for including each equation and their relationships. A second possibility would be to consider the key

requirements of an applicable model and build up a new model around those requirements. On balance, the latter seems a more justifiable and structured approach. So the questions are, given the subject matter of this thesis, what are the key features that are missing in the existing models and how can they be remedied by applying the robust optimal control approach ?

The first and most obvious is of course uncertainty - in model selection, in parameters, in measurement and in data. The second key feature is the need for feedback. Risk management in the majority of institutions managing financial risk is inherently a process that involves considerable feedback - hedging decisions are generally taken at the level of individual trades but are not viewed as single period static decisions, but monitored constantly, thereby giving rise to risk management decisions and activity which is generally discrete. Given the prevalence in the use of BS style delta hedging, it is clear that there is an awareness of the need to incorporate feedback into the risk management process to ensure greater robustness in decision making.

Two questions therefore remain to be answered. First, why apply robust optimal control to risk management ? Second, how should robust optimal control be applied to risk management in order to produce an understandable, practical and tractable model ? The answer to the first is relatively straight forward - the objective of optimal control theory is to find optimal ways to control a dynamic process. Managing financial risk - particularly on derivatives - certainly qualifies as a dynamic process, though in practice, management often tends to be static. One of the key attractions of optimal control is that it can be applied in both discrete and continuous forms. Optimal control is also more intuitively in tune with the risk management problem as it allows the explicit inclusion of control variables whose relationships to the state variables can be explicitly modelled to more accurately capture the nature of the feedback process.

In order to proceed to the development of robust risk management, it is necessary to make a short diversion on the technical tools needed to solve the problems encountered. The approach adopted in this chapter is to use nonlinear Lyapunov equations employing an approach developed by Freeman (1993) and Freeman and Kokotovic (1996). The previous chapter constructed the basic framework on which the robust optimal control for risk management will be built in the next section. The following section therefore extends the framework developed in chapter 3 by incorporating two modifications for softening PWMN control laws. The first involves using a non-smooth robust control Lyapunov function, thereby implicitly allowing non-differentiability of the Lyapunov function, which is not in and of itself a requirement for stability and therefore poses no significant constraint in the context of the current problem. Permitting non-smoothness in the robust control Lyapunov function allows a richer possible set of non-linearities to be investigated which will be useful in examining portfolio level risk and producing a portfolio risk

measure. The second modification involves the use of directional derivatives to eliminate the "chattering" in robust portfolio management rules that can appear at much greater distances for non-at-the-money instruments and therefore compromise the usefulness of a portfolio level risk measure. There then follows a more comprehensive portfolio level example of applying robust optimal control techniques. The final section of the chapter concludes.

4.3.1 Stabilisation, robustness and control

In chapter 3 it was shown that for non-linear systems, using the Freeman and Kokotovic (1996) non-linear robust optimal control approach involves a two stage process of first constructing a robust control Lyapunov function, then deriving a robustly stabilising feedback controller such that the derivative of the Lyapunov function is negative. Whilst the entire procedure will not be needlessly re-iterated in this chapter, some repetition is inevitable as the same framework is once again being brought to bear but in a rather different context and with the added development of a flattening the robust control Lyapunov function by the use of simple penalty conditions that penalise the distance to a region around the manifold rather that the distance to the manifold itself.

A simple example demonstrates how the "flattening" actually works. Consider the following simple second order system where x_i are the state variables, w is the disturbance and u the control variable

$$\dot{x}_1 = x_2 + |x_1|^{1+r} w \tag{4.53}$$

$$\dot{x}_2 = u \tag{4.54}$$

where w is a scalar disturbance with values in the interval $B = [-1, 1]$ and $r > 0$, a growth parameter that multiplies the effect of the disturbance. The following (quadratic) robust control Lyapunov function applies

$$V(x) = x_1^2 + p[x_2 - x_1 s(x_1)]^2 \tag{4.55}$$

where p is a design parameter and $s(x_1)$ is known to be a smooth function. The robust control Lyapunov function is quadratic of type $V(x) = z^T P z$, where

$$z_1 := x_1$$

$$\text{and} \tag{4.56}$$

$$z_2 := x_2 - x_1 s_1(x_1) \tag{4.57}$$

The worst possible derivative of V is

$$\begin{aligned}\max_{w\in B}\dot{V} &= \left|2x_1 - 2pz_2\left[s_1 + x_1 s_1'\right]\right| |x_1|^{1+r} \\ &\quad + \left[2x_1 - 2pz_2\left[s_1 + x_1 s_1'\right]\right] x_2 + 2pz_2 u\end{aligned} \tag{4.58}$$

where s_1'denotes the derivative of $s_1(x_1)$ with respect to x_1. FK show that every smooth control law $u(x_1, x_2)$ which makes the robust control Lyapunov function non-positive also satisfies

$$\left|\frac{\partial u}{\partial x_2}\right| \geqslant \frac{p}{4}|x_1|^{3r} \tag{4.59}$$

Which in simple terms means that the local gain of the control law in the x_2 direction grows by $|x_1|^{3r}$ as $|x_1| \to \infty$. The point is that the exponential term $3r$ which measures the hardening of the control law is unaffected by choices of the the function s_1 or the parameter in p above. Unfortunately, this applies to every control law that makes the Lyapunov derivative negative. The good news is that the local gain is not actually necessary for robust stabilisation and, furthermore, it is just a by-product of the quadratic form of the robust control Lyapunov function. FK demonstrate a simple technique for constructing an alternative type of robust control Lyapunov function that produces much slower growth by flattening the penalty term $|x_1|^{3r}$.

The procedure is simple. Recall from the current simple example that the second term of the robust control Lyapunov function penalises the distance to the manifold defined by $x_2 = x_1 s_1$. FK obtain an alternative, "flattened" penalty term by penalising the distance to a region around the manifold rather than the distance to the manifold itself

$$V(x) = x^2 + p\begin{cases}[x_2 - x_1 s_1(x_1) - \varrho_1(x_1)]^2 & \text{when} \quad z_2 \geqslant \varrho_1(x_1) \\ 0 & \text{when} \quad |z_2| \leqslant \varrho_1(x_1) \\ [x_2 - x_1 s_1(x_1) + \varrho_1(x_1)]^2 & \text{when} \quad z_2 \leqslant -\varrho_1(x_1)\end{cases} \tag{4.60}$$

The results of running a Monte Carlo simulation of the above flattened version of the robust control Lyapunov function for a simple portfolio of a European exercise option and a hedging future can be seen in figure 4.1 which shows the comparative effects of applying the flattening in the above simple example.

Figure 4.1: Relative robustness of PWMN control laws

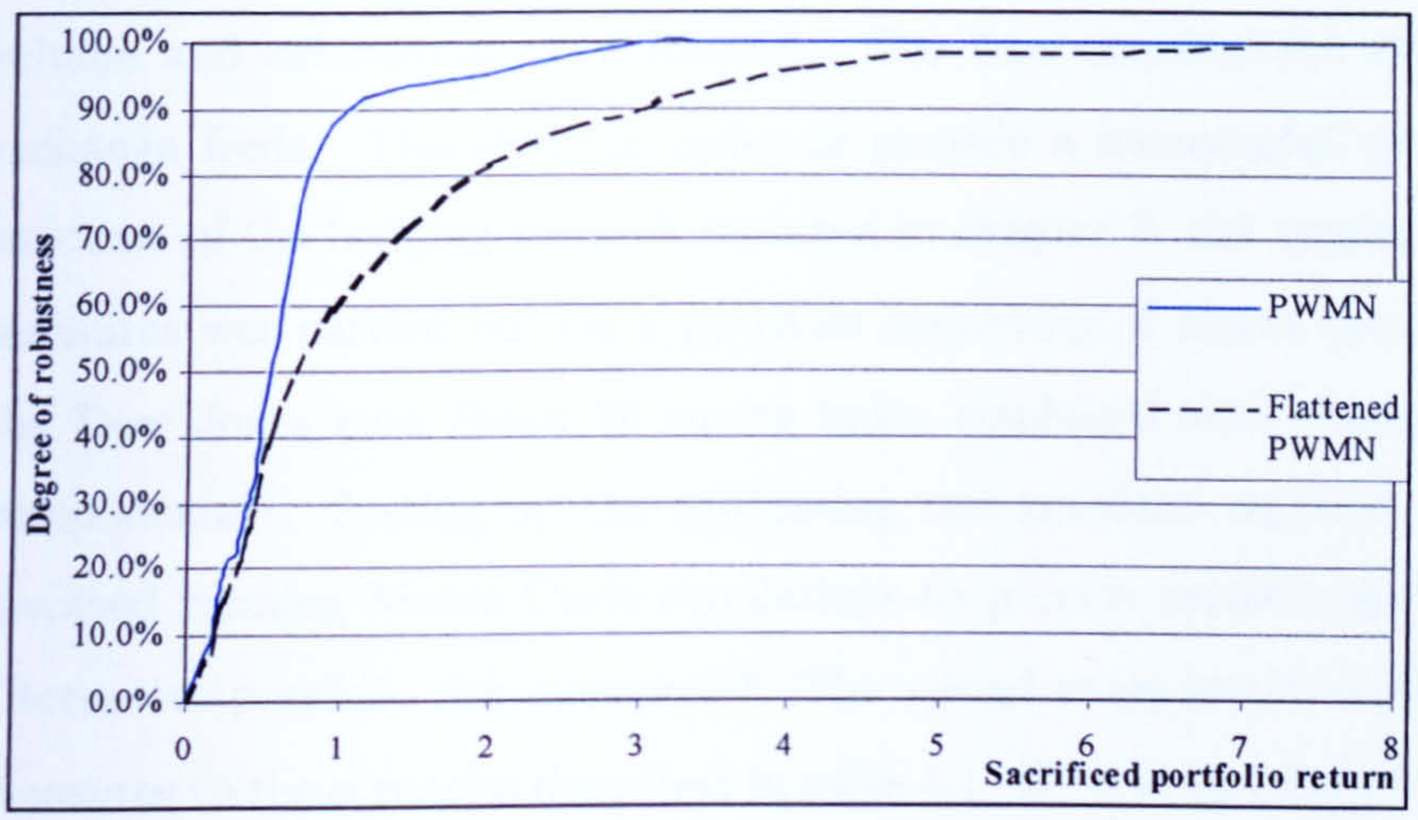

This simple example illustrates serves two purposes. First, to demonstrate the essence of the approach that is adopted in the next section. Second, to illustrate the simple point that it is possible to solve meaningul robust optimal control problems very simply and elegantly without the use of the HJI equation and its associated complex requirements.

4.3.2 A piecewise min-norm robust optimal control law for portfolio risk management

At the heart of the general PWMN framework is a simple cost function of the familiar form

$$J = \int_0^\infty \left[x^2 + u^2\right] dt \tag{4.61}$$

where x is the vector of state variables (equity prices for example) and u is the vector of control variables (the hedging instruments). The general risk management problem can be stated in terms of simply minimising J according to the required degree of robustness, then proceeding through the two part process of first constructing a robust control Lyapunov function for the system, then deriving a robustly stabilising feedback controller for the system such that the derivative of the Lyapunov function is negative. The solution approach then proceeds simply and elegantly given the target degree of robustness required from the control law.

4.4 Empirical comparison of risk methodologies

Mindful of the difficulties in selecting a real-world portfolio risk management problem for which adequate high quality data is available, it was decided to consider instruments exhibiting accurate, plentiful and easily accessible market data. The first requirement inevitably meant concentrating on instruments traded on organised trading exchanges. The second requirement meant selecting contracts with heavy trading

volume and substantial open interest. The final requirement meant using data directly from official exchange feeds. Therefore, in order to provide a meaningful contrast with the more micro-economic concerns of the hedging research reported in chapter 3, the empirical testing of alternative portfolio risk measures was carried out on a portfolio consisting of Eurex traded European exercise call options on the Dow Jones euro Stoxx 50 equity index combined with outright positions in the underlying index components[13]. Setting up the calibrating test portfolio required a two stage process. The first stage involved running Monte Carlo simulations to provide realistic and accurately calibrated results for the alternative portfolio risk measures[14]. The second stage involved applying the portfolio control rules and measures to the scenarios described in table 4.1. Both stages included all of the portfolio rules and metrics discussed in the previous section.

4.4.1 Setting up the problem

As explained above, a range of portfolio risk management strategies was evaluated using actual historical data from the Eurex exchange with results denominated in Euros. The hedging strategies examined all began with the same 100 contracts , though the underlying notional differed between contracts as specified in the contract specifications for the relevant exchange[15]. All strategies were subjected to identical hedge rebalancing at uniform daily intervals. All gains or losses were assumed to be rewarded or penalised at the relevant futures margin-account interest rate. The interest rate costs, forward rates and discount factors were calculated using a standard yield curve bootstrap developed for this thesis. Volatilities required for the portfolio risk evaluations were a combination of those supplied by Eurex, those supplied by Credit Suisse First Boston and those that needed to be calculated and substituted to ensure the consistency of the data. Hedging and funding flows were assumed to be calculated with respect to the bid or offer rates for the Euro cash market using the overnight index swap rate to fund all balances. All calculations were carried out with respect to end of day mark-to-market valuations of the relevant variables and the risk measures calculated after portfolio rebalancing. The final performance of a strategy was defined as the final cumulative value of the initial position, plus cash inflow, less cash outflow, normalised by the notional of the relevant currency, averaged over the 10,000 simulations performed for each scenario. Table 4.1 below describes the risk management measures that were evaluated, whilst table 4.2 provides a list of the simulated volatility environments.

[13] See the first part of Appendix 6 for full details of the data sources, time period covered and contract specifications. By way of comparison, the second part of Appendix 6 reports the results of the calibrating Monte Carlo simulation in tabular format.

[14] The results for which are reported in table 1 in Appendix 6.

[15] Detailed contract specifications for Eurex are easily available via the internet either through the portal web site: www.numaweb.com/ or via www.eurex.com.

Table 4.1: Portfolio risk management measures

Simulation Scenario	Description
Delta Neutrality	Simple delta hedge
10d Delta-Gamma VAR-95%	10d horizon VAR calculated using delta-gamma approximation - 95%
10d Delta-Gamma VAR-99%	10d horizon VAR calculated using delta-gamma approximation - 99%
Minimax	Minimisation of potential hedging error
EVT (GPD)	Peaks over threshold model
Rational Bounds	Assumed rebalancing
PWMN	No flattening
PWMN flattened	Flattened using distance to region around manifold [16]

Note: Transaction costs can be modelled in a number of different ways. However, it was felt that adjusting the bid-ask spread was the most appropriate method in the interests of transparency and in line with most capital-market conventions.

Table 4.2: Simulated volatility environments

Volatility Environment	Description
Low vol	Non-constant, but low & time declining volatility
Flat vol	Constant volatility
High vol	Non-constant, high but time declining volatility
Skew vol	Affine combination of normal (25%) and lognormal (75%) volatilities
Single Jump: 10%	10% jump shift on high volatility half way through life of option
Single Jump: 30%	30% jump shift on high volatility half way through life of option
Single Jump: 50%	50% jump shift on high volatility half way through life of option
Single Jump: 100%	100% jump shift on high volatility half way through life of option
Single Jump: 500%	500% jump shift on high volatility half way through life of option

[16] Many large institutions currently using VAR use a 10-day holding period and a 99% confidence limit and employ historical simulation based on the well known Taylor Series expansion of the P&L to generate the VAR estimates from the time series of each underlying risk factor (Equity Index Spot Price, Interest Rate Zero Rate, etc). In this thesis it was decided to calculate a change in the underlying (either proportional or absolute move) using the Taylor Series expansion to calculate a change in the portfolio P&L: $\Delta PL(t) = \delta m(t) + 0.5\Gamma m^2(t)$

Each of the risk management strategies was simulated for the scenarios detailed in table 4.2. All simulations were run 10,000 times to ensure sufficient convergence upon a stable and accurate result. However, a computational barrier arose as it became evident that a single 3.00Gb Pentium IV PC with 2Gb of RAM was insufficient for running the required number of simulations within a reasonable time. The author therefore adapted a share-ware job distribution algorithm to create a virtual compute grid out of 8 PC's of varying specifications, bringing the calculation time down to less than 5 minutes for a complete set of Monte Carlo runs for all risk management scenarios.

Each of the portfolio risk management strategies was evaluated using a single option on an underlying 3-month futures contract for the three currencies mentioned, over the entire trading life of the relevant contract. The interest rate costs, forward rates and discount factors were calculated using a standard yield curve bootstrap developed for this thesis. Volatilities required for valuation were a combination of those supplied by the relevant exchanges and those supplied by Credit Suisse First Boston. Hedging and funding flows were assumed to be calculated with respect to the bid or offer rates for the cash market for the relevant currency. All calculations were carried out with respect to end of day mark-to-market valuations of the relevant variables.

4.4.2 Simulation results

Two sets of results from the simulations are presented in figures 4.2 - 4.11. The first set of results, contained in figures 4.2 - 4.11, shows the results of the Monte Carlo simulations with one figure depicting the results of each volatility scenario. The final set of results contained in figure 4.11 presents the results of the historical simulation and compares all of the volatility environments in a single consolidated figure. Note that the numbers underlying figures 4.2 - 4.11 are also presented in tabular form in appendix 6.

The first and most obvious point to be evident from the simulation results is their broadly two cohort nature. In the first cohort are the non-robust strategies: VAR, minimax, EVT and rational bounds strategies. These strategies generally involved the incurring of the greatest levels, ranges and volatility of risk exposures. This excessive range is most notable in the case of the EVT results for the 100% and 500% jump volatility results. This is not entirely surprising as the simulation exercise uses the Generalised Pareto Distribution which is a heavy-tailed distribution which is therefore more likely to emphasise tail effects than the standard VAR based on the lognormal distribution.

In the second cohort were the PWMN and flattened PWMN risk management strategies. As anticipated, these were the risk management strategies provided the most interesting results from the perspective of robustness. The first point to notice is that this group of strategies do not perform that well in low volatility environments. This makes intuitive sense as robustness is likely to represent an overkill strategy

in a low volatility environment - an apparent case of paying for expensive insurance when faced with a low probability of occurrence event. Interestingly, in the flat volatility environment the picture reverses somewhat with the PWMN strategies becoming the most cost effective as they manage to produce the lowest level and range of risk of any of the measures. Not surprisingly, the PWMN strategies do best in the high volatility and large jump volatility environments when the guarantee of an absolute ceiling to risk makes robustness worth far more. This picture becomes even more pronounced in the highest jump shift environment, when the use of flattened control laws delivers lower levels and spreads of risk, making the trade-off of cost against robustness implied by the strategy attractive[17].

Figure 4.2: Monte Carlo simulations of portfolio level risk - low volatility environment

(risk expressed as a % of strike premium)

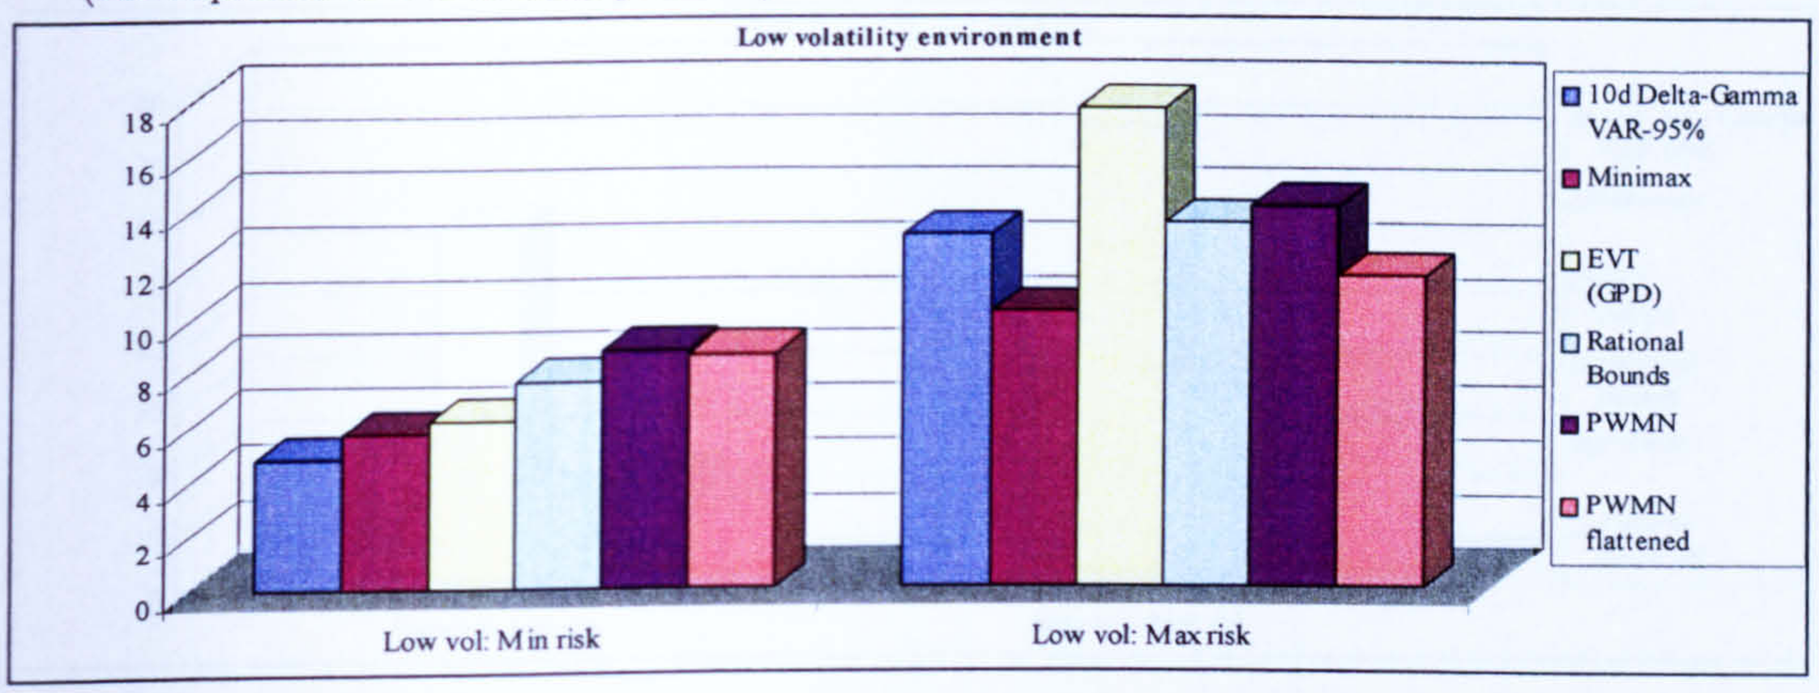

Figure 4.3: Monte Carlo simulations of portfolio level risk - flat volatility environment

(risk expressed as a % of strike premium)

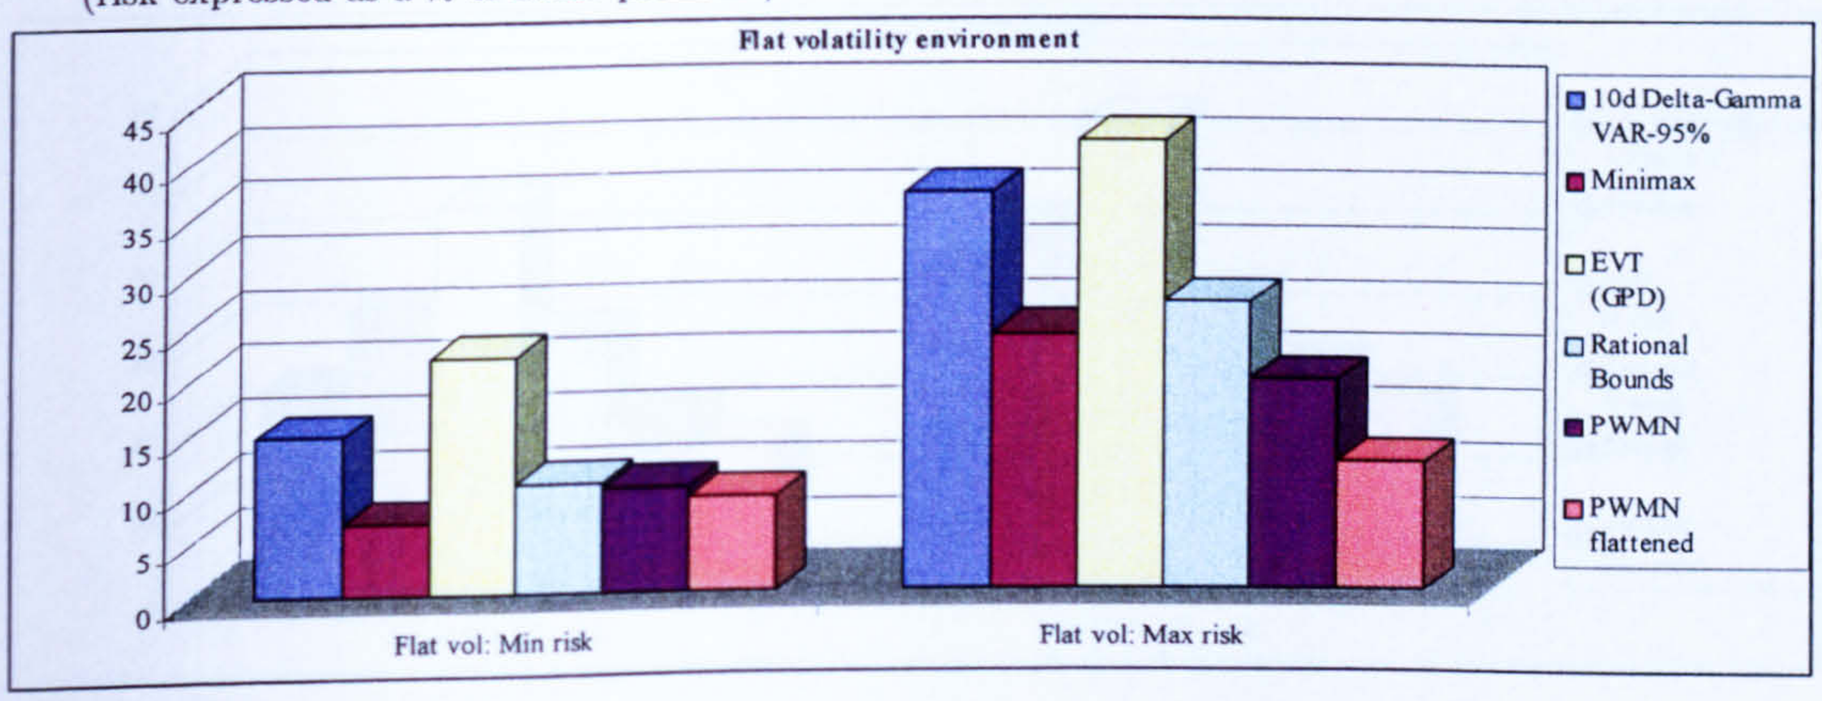

Figure 4.4: Monte Carlo simulations of portfolio level risk - high volatility environment

(risk expressed as a % of strike premium)

[17] It is also worth noting that results not reported here for a wider range of robustness criteria make relatively little difference to the overall picture.

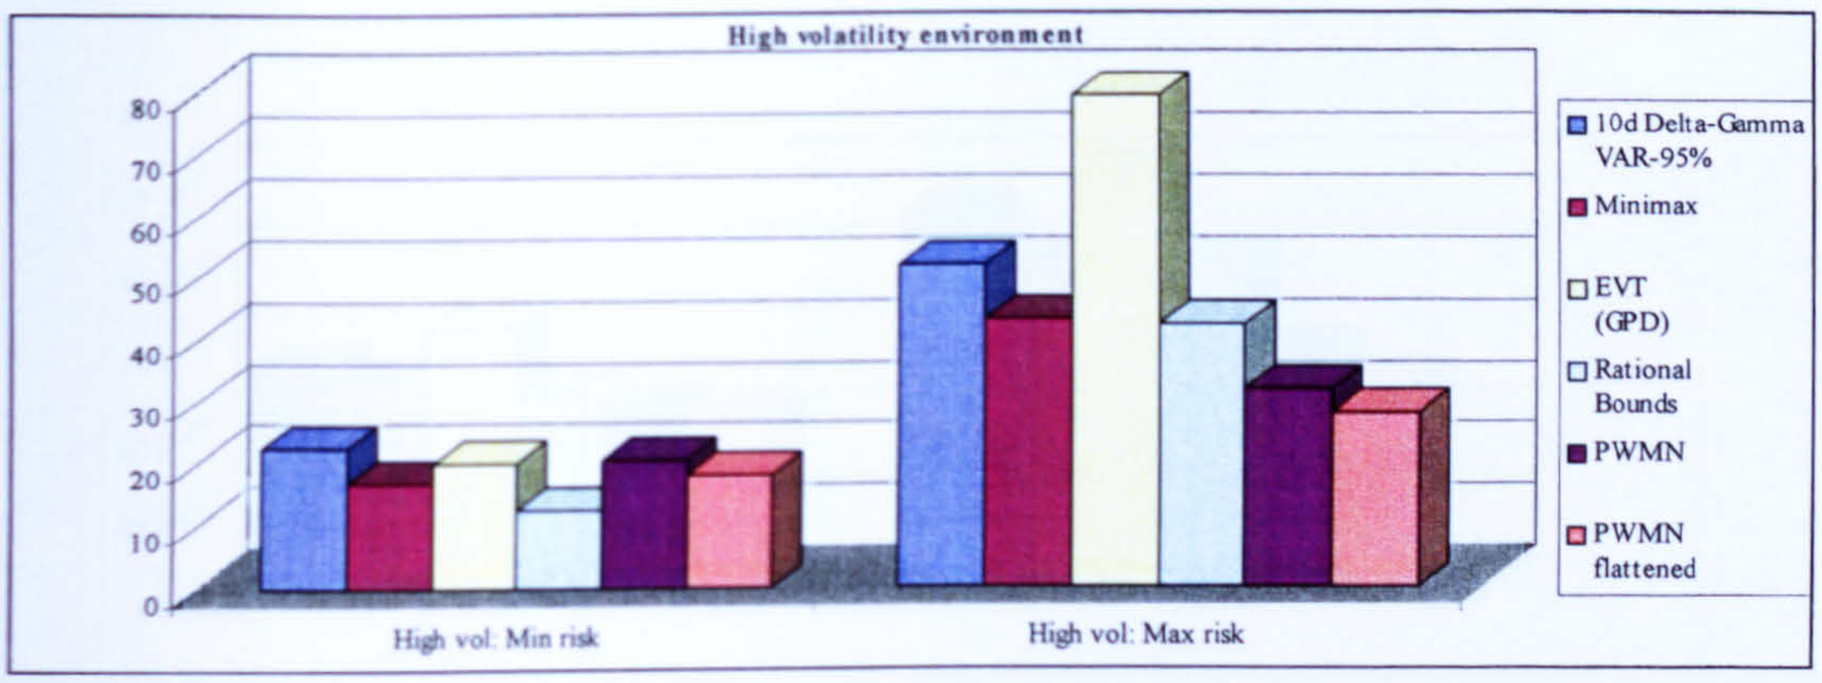

Figure 4.5: Monte Carlo simulations of portfolio level risk - skew volatility environment

(risk expressed as a % of strike premium)

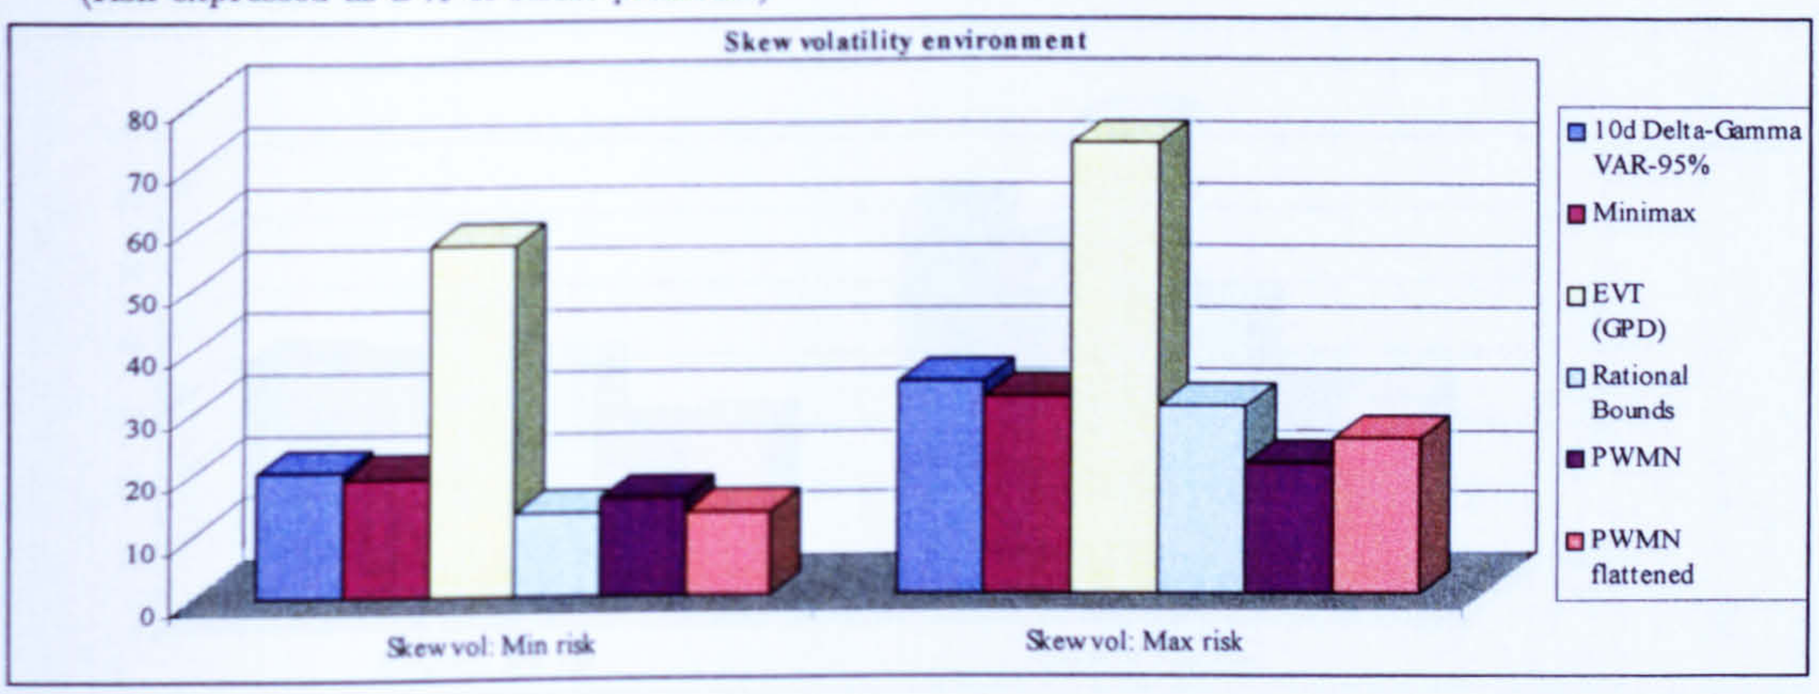

Figure 4.6: Monte Carlo simulations of portfolio level risk - 10% jump in volatility

(risk expressed as a % of strike premium)

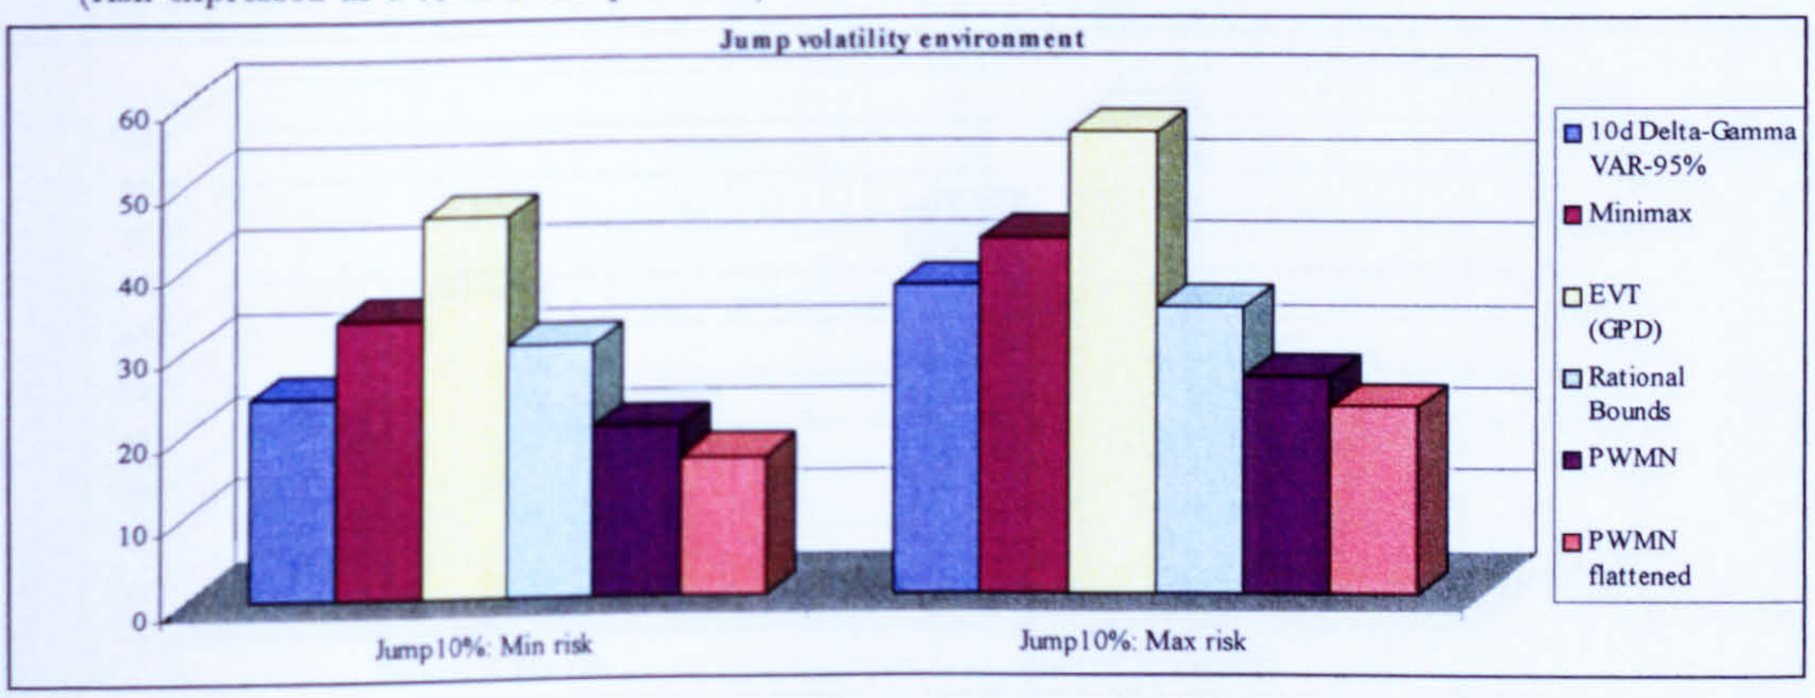

Figure 4.7: Monte Carlo simulations of portfolio level risk - 30% jump in volatility

(risk expressed as a % of strike premium)

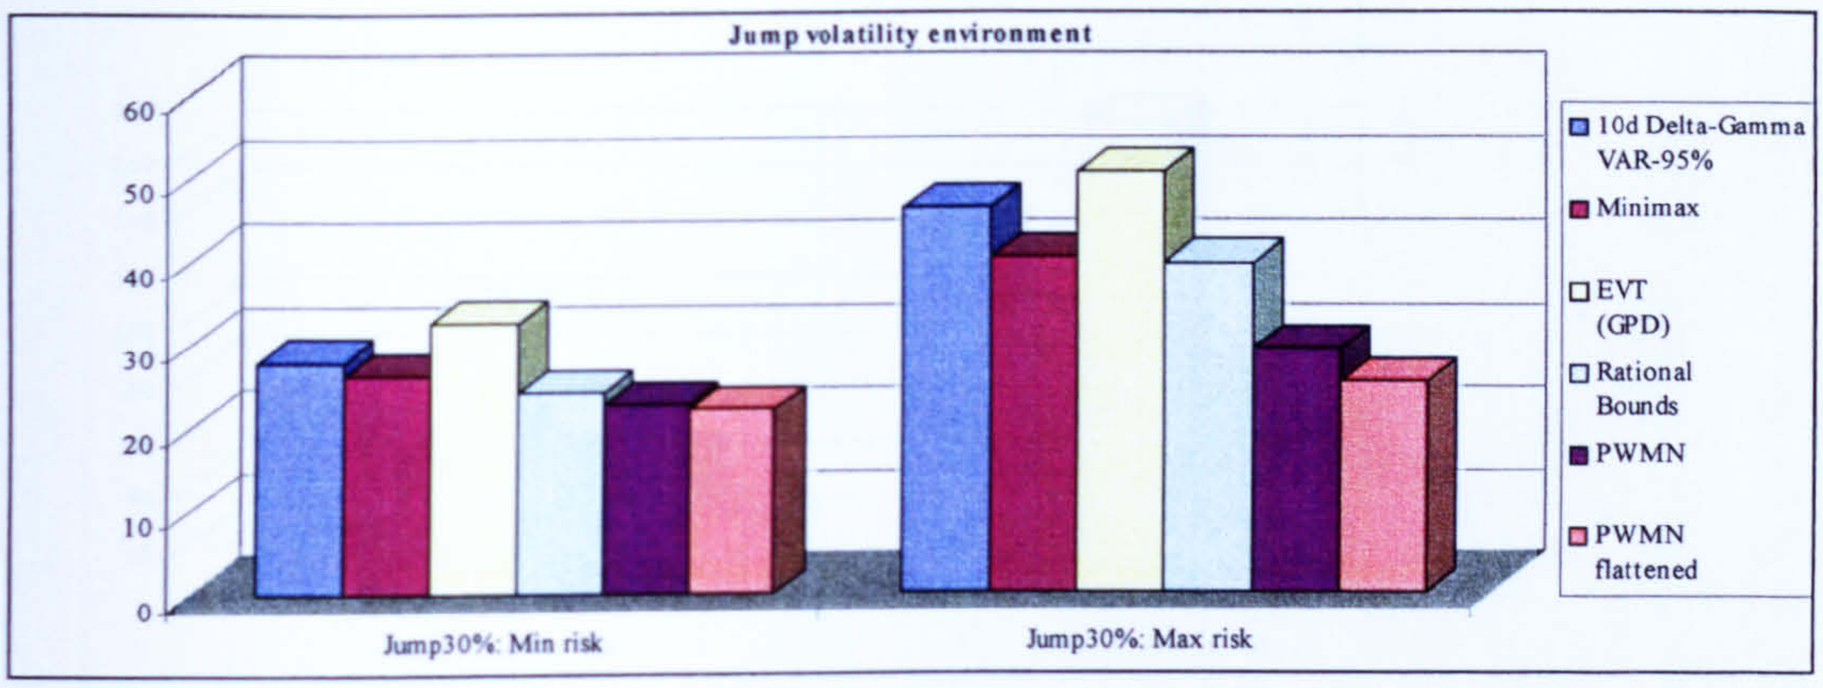

Figure 4.8: Monte Carlo simulations of portfolio level risk - 50% jump in volatility

(risk expressed as a % of strike premium)

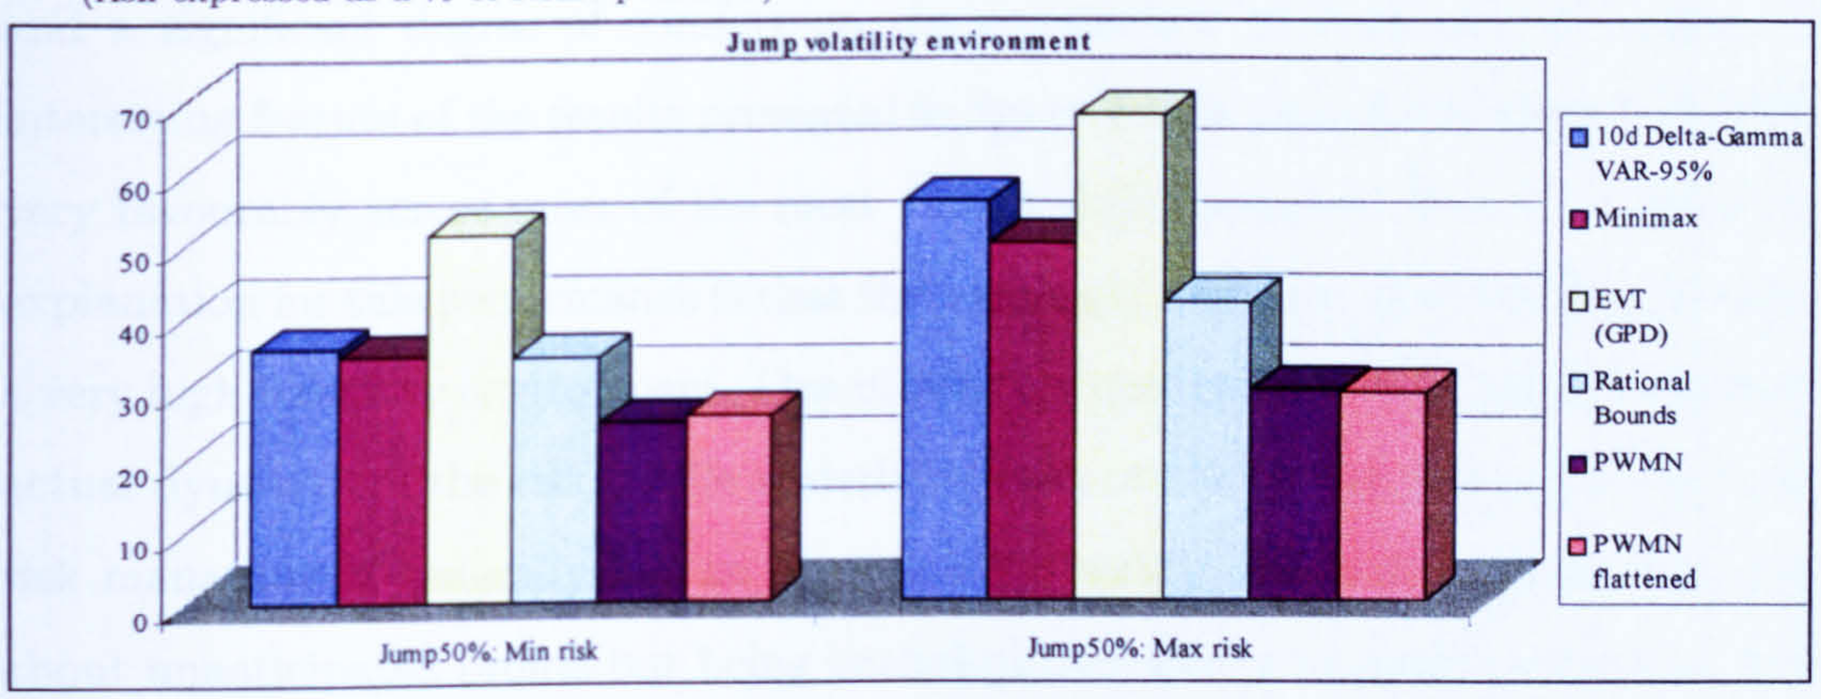

Figure 4.9: Monte Carlo simulations of portfolio level risk - 100% jump in volatility

(risk expressed as a % of strike premium)

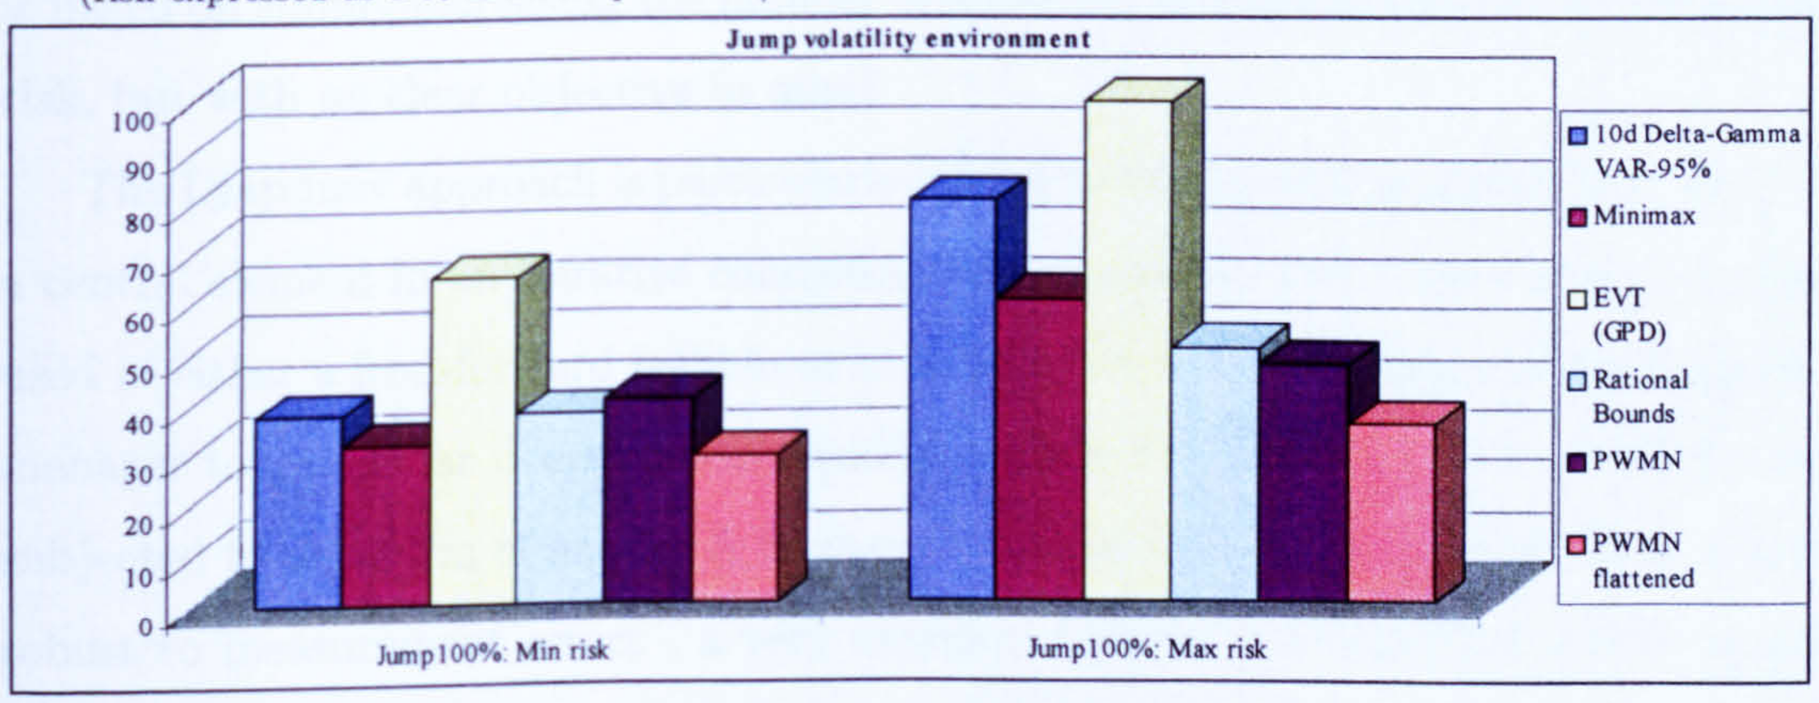

Figure 4.10: Monte Carlo simulations of portfolio level risk - 500% jump in volatility

(risk expressed as a % of strike premium)

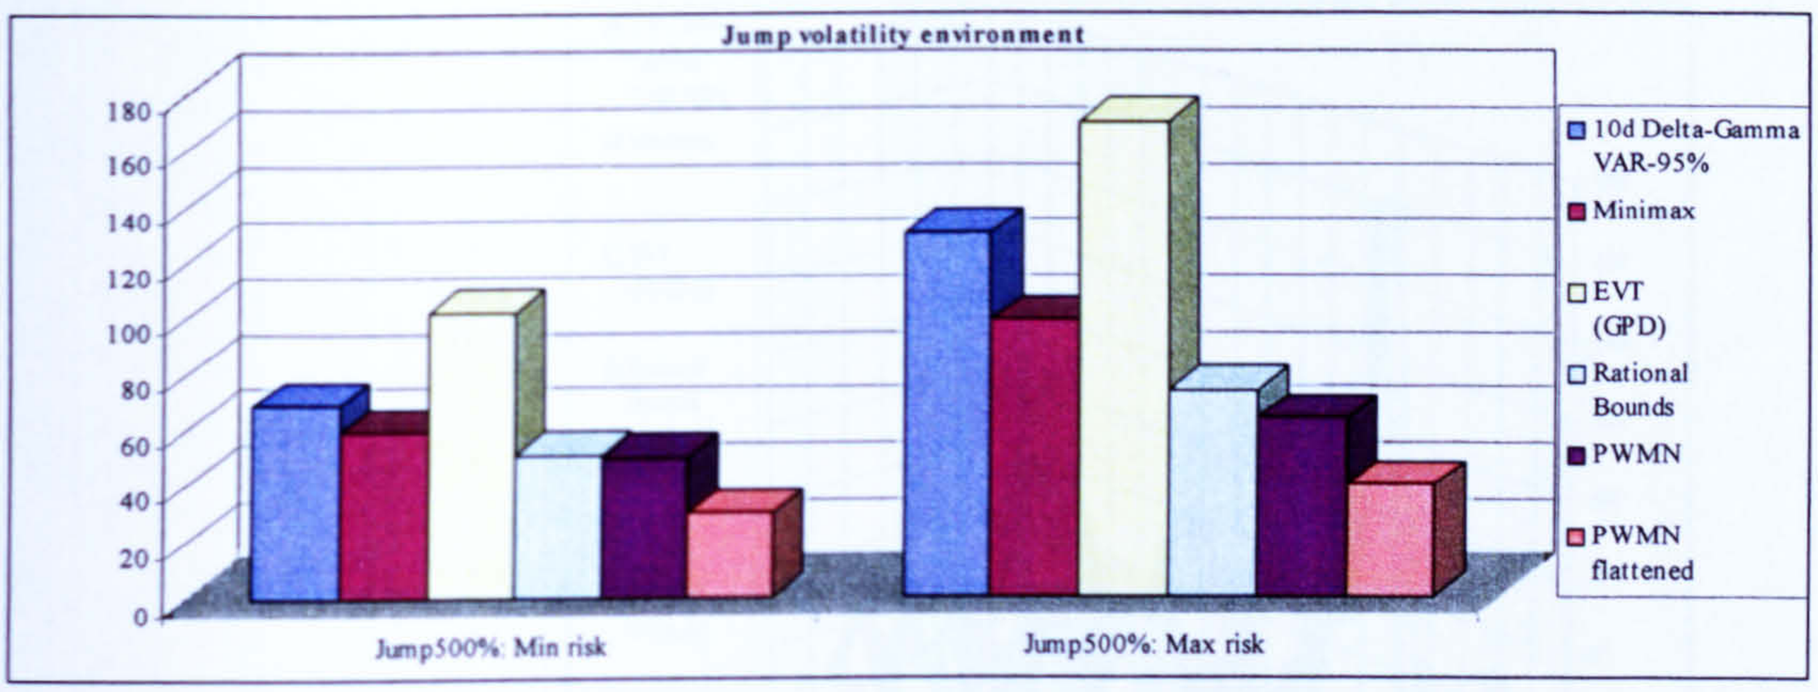

The broad picture to emerge from the Monte Carlo simulation is broadly maintained when the exercise is repeated using the same shocks applied to real historical data. This provides a useful cross check and a significant degree of comfort on the consistency of both sets of results. Potentially the most interesting feature of the results presented in figure 4.11 is that the flattened PWMN control law performs very favourably across most of the most violent shift scenarios. It would appear that the most rational explanation for this performance is that the extremely high cost of robustness becomes more worthwhile in a very high volatility environment. One interesting question not explicitly considered in this research is the actual dynamics of the risk profile underlying each of the robust and non-robust risk measures. Portfolio risk management generally has an asymmetrical attitude to risk management - being relatively benign about unanticipated profits but being pathologically averse to significant losses. This would suggest that potentially one of the most significant uses of this type of robustness analysis would be as a replacement for existing stress-test type approaches where the usual methodology of assessing maximum downside risk is based on simply increasing the number of standard deviations away from the expected level of portfolio risk, but with no clear objective in mind.

The Lyapunov approach is particularly suited to this type of analysis, as it was originally devised to be a central element in an iterative controller design process. The robust control Lyapunov function can be used in either a feed-forward fashion or in a recursive back-stepping approach which would allow the risk manager to gain a far deeper understanding of how the dynamics of a portfolio evolved over time when subjected to any form of shock. A further benefit of using this approach is that robustness means being robust to measurement errors - a very common problem for real-life portfolio managers where accurate profit and loss and risk information for alternative scenarios can frequently be far too costly in both computational and human resource terms to be practically feasible.

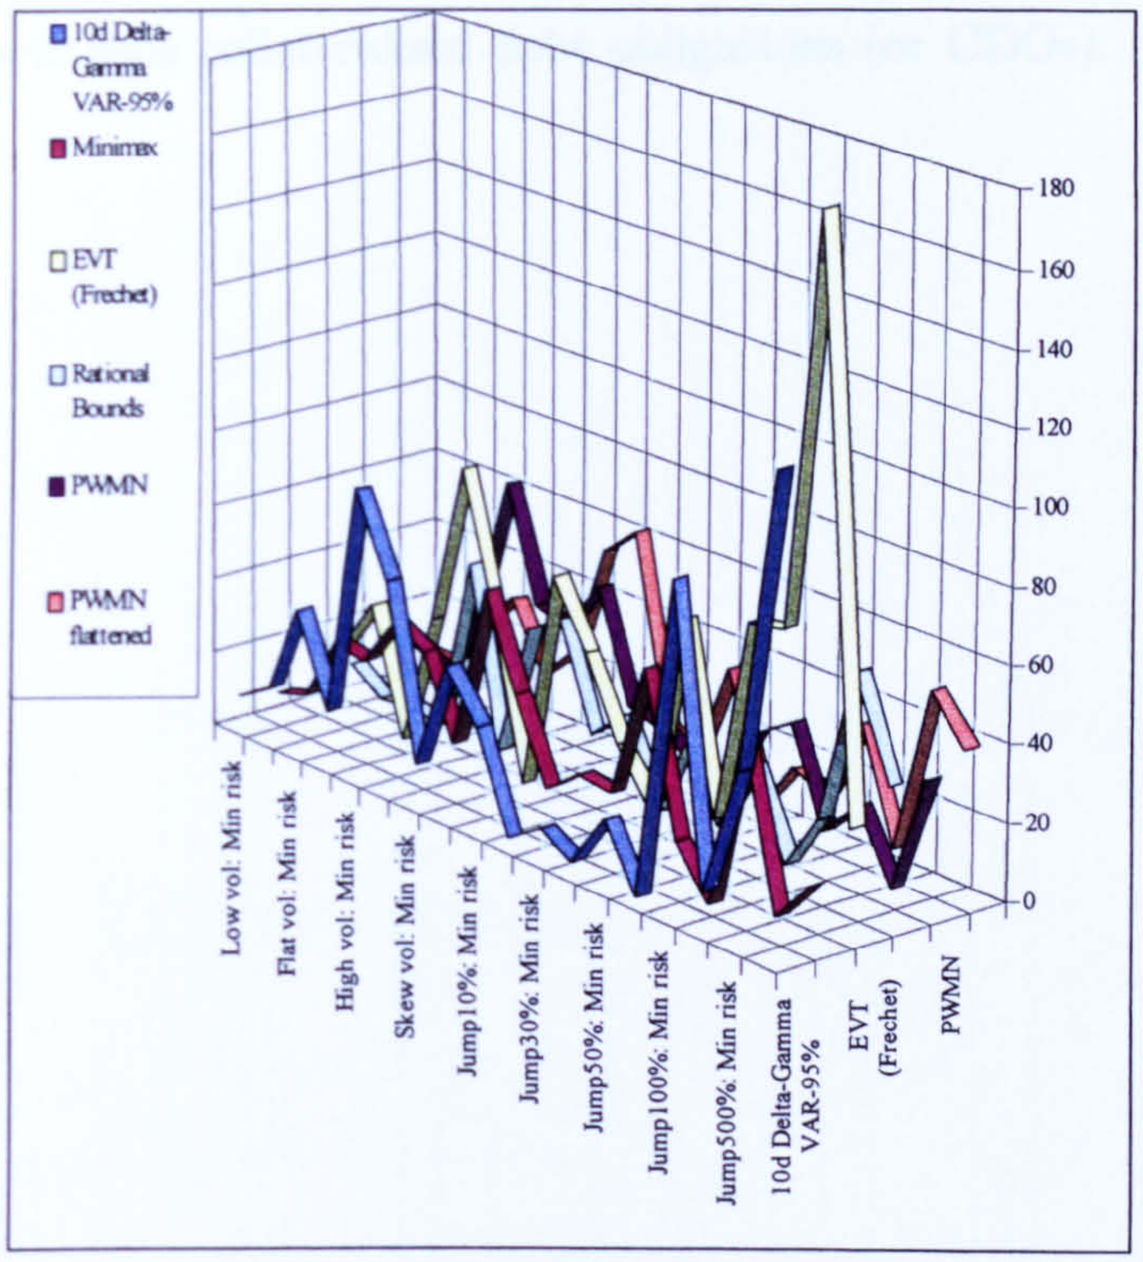

Figure 4.11: Historical simulations of portfolio level risk

4.5 Conclusions

This chapter has reviewed alternative non-robust risk management approaches. An alternative approach was developed using piecewise min-norm control laws based on non-linear robust control Lyapunov functions. The framework was compared empirically with a number of the more widely used portfolio risk measures. The research applies existing non-linear robust optimal control techniques developed in engineering to the problem of managing risk at a portfolio level to enable the creation, execution and measurement of the performance of portfolio level risk management strategies. The principal findings were that the PWMN control approach provides a framework for the risk manager to select the preferred trade-off between the required degree of robustness and the costs of portfolio management. An empirical comparison and evaluation of the current approaches based on the DJ euro Stoxx 50 index indicates that depending on the cost appetite and the volatility environment, ensuring robustness may not be as expensive as would at first be thought on an a priori basis.

Given the substantial growth in the market for credit derivatives, an obvious extension to the work in this chapter is to apply the same techniques and approach to evaluate the potential trade-off for the credit risk of a financial instrument portfolio, especially with respect to the magnitude and scope of default

risk on multi-credit products such collateralised debt obligations (or CDOs). This work is currently in progress.

Chapter 5

Summary, Conclusions and Suggestions for Further Research

Everybody sets out to do something and everybody does something, but no one does what he sets out to do.

George Moore, Irish novelist, art critic and essayist, 1852 - 1933

It is better to travel in hope than to arrive.

Chinese Proverb

The original idea that motivated this thesis was the robustness of decision making in the presence of uncertainty. Although this idea has not and did not change during the course of the research, the focus of its application did deviate in varying degrees (depending on the chapter and topic). However, the fundamental objective remained constant, namely, robustness of decision making in the presence of uncertainty, within financial economics. This is the final chapter and it provides a brief summary of the principal findings contained in the research, draws a number of conclusions based on the research and offers some tentative ideas for further research in the future.

5.1 Summary of the thesis

Chapter 1 provided the theoretic underpinnings of the framework which is the unifying theme underlying this thesis. The main themes of decision making under uncertainty were introduced and examined in detail. The chapter also introduced the concepts and mechanics associated with robust optimal control

theory. The central ideas behind the use of the $H\infty$ norm were explained and discussed, setting up the framework for use in later chapters.

Chapter 2 examined a number of aspects of robustness in the pricing of options where the underlying is an extreme event. Consistent with uncertainty as the central theme of this thesis, the work chapter 2 investigated the concept of dealing with uncertainty in the pricing of derivatives whose underlying are catastrophic events. A number of new ideas and some preliminary results were presented based on valuing catastrophe derivatives using a robust optimal control approach. The work focused in particular on catastrophe bonds and catastrophe options. Preliminary empirical results indicate that the new model produces more stable results in the face of uncertainty through the explicit modelling of feedback and exhibits greater robustness to discontinuous behaviour than the current popular double Cox PIDE model.

Chapter 3 categorised and reviewed alternative approaches to the problem of developing a dynamic hedging strategy that is robust and stable in a multi-period hedging problem. Robust optimal control techniques were applied and two forms of dynamic hedging strategy were developed - one based on linear $H\infty$ robust optimal control techniques, the other non-linear robust control Lyapunov functions. The linear model produced hedging rules that were found to satisfy robustness criteria, but proved to be too conservative (and therefore costly) to be useful in practice. Piecewise min-norm robust optimal control rules were developed based on non-linear robust optimal control Lyapunov techniques in an attempt to exploit the suspected non-linearities of the the hedging problem. The first version of the PWMN rule was found to be less conservative than the $H\infty$ rule, but was observed to still be implying excessively punitive changes in hedging policy for options deep out of the money. A simple scalar function modification to the PWMN rule was and found to produce encouraging improvement in the previously conservative behaviour, whilst retaining desirable robustness and stability properties.

Chapter 4 reviewed alternative non-robust risk management approaches. The piecewise min-norm control law framework developed in chapter 3 was extended to include flattened non-linear robust control Lyapunov functions to capture the trade-off between robustness and cost for a portfolio of index options, futures and simple equities. The PWMN approach was compared empirically with a number of the more widely used portfolio risk measures. The principal findings were that the PWMN control approach provides a transparent and easily applied framework for the risk manager to select the preferred trade-off between the required degree of robustness and the costs of portfolio management. An empirical comparison and evaluation of the current approaches based on the DJ euro Stoxx 50 index indicates that depending on the cost appetite, ensuring the required degree of robustness may not be as high as would at first be thought on an a priori basis.

5.2 Conclusions from the thesis

Chapters 2, 3 and 4 all point to the conclusion that incorporating feedback into the modelling of decision making clearly helps to dramatically reduce the effects of unmodelled uncertainty on the robustness of decision rules. The research also supports the (not unsurprising) conclusion that the cost of ensuring that decision rules possess the required degree of robustness is higher in the case of the linear $H\infty$ models than it is for the non-linear robust control Lyapunov models. This latter conclusion is highlighted most clearly by the results presented in table 3.9 in section 3.4, which illustrate the gains from adopting the non-linear, flattened robust control Lyapunov equation as the approach to produce robust decision rules.

The explicit modelling of feedback can never totally eliminate uncertainties. One of the interesting issues investigated during this research was the exact nature of the uncertainty that was being modelled. Results obtained but not reported include specific examination of different forms of uncertainty, structured or unstructured, for example. The form in which the uncertainty is modelled appeared to have little or no quantitative or qualitative impact upon either the robustness or stability of the control rules derived in chapters 2, 3 or 4.

5.3 Limitations of the research

During the course of the research it has become apparent that the finished thesis contains a number of limitations in its scope. In chapter 2 moral hazard and basis risk were specifically excluded from the work on CAT bonds, which means that no consideration has been given to the default riskiness or moral hazard. These factors, together with other issues such as stochastic interest rates, constitute potentially fertile and interesting areas for future research. They could be incorporated relatively easily by simple modifications to the basic state space approach and via the type, timing and impact of perturbations applied within the non-linear Lyapunov model.

In a non-robust setting default riskiness is frequently tackled using some form of copula function to proxy default events, the output of which is then mapped into actual default times. The approach then generates actual values using some form of Monte Carlo simulation to value the underlying payoff function (see Li's 2000 paper for Risk Metrics for what has become the standard application of this methodology).

Clearly, such approaches still suffer from a lack of robustness for all of the reasons already referred to in previous chapters. However, it is likely that such omissions could easily be incorporated into the non-linear Lyapunov model and preliminary work not reported in this thesis indicates that such changes do not affect the quantitative or qualitative nature of the results reported in chapters 2, 3 or 4.

5.4 Suggestions for further research

Extending the PWMN approach to credit derivatives would be a useful, interesting potentially fertile development. For example, applying PWMN to managing the portfolio risk of collateralised debt obligations (CDO's and their compound derivatives CDO's on CDO's or CDO^2's) would appear to be a potentially fertile area for future research. Given the substantial growth in the market for credit derivatives, an obvious extension to the work in this chapter is to apply the same techniques and approach to evaluate the potential trade-off for the credit risk of a financial instrument portfolio, especially with respect to the magnitude and scope of default risk as referred to in the previous section. This work is currently in progress.

Chapter 6

Bibliography

Bibliography

[1] Aase, K., (1999), "An equilibrium model of catastrophe insurance futures and spreads", Geneva Papers on Risk and Insurance Theory, Volume 24.

[2] Ahn, H. and Wilmott, P., (1999), "On trading American options", OCIAM, Oxford.

[3] Allais, M., (1953), "Le Comportement de l'Homme Rationnel devant le Risque: Critique des postulats et axiomes de l'¶Ecole Americaine", Econometrica.

[4] Andrews, D.F., Bickel, P.J., Hampel, F.R., Huber, P.J., Rogers, W.H. and Tukey, J.W., (1972), "Robust estimates of location: Survey and Advances", Princeton University Press.

[5] Aoki, M. (1987), "State Space Modelling of Time Series", Springer Verlag.

[6] Arrow, K.J., (1954), "Existence of an Equilibrium for a Competitive Economy", with Gerard Debreu, Econometrica

[7] Arrow, K.J. and Hahn, F.H., (1971), "General Competitive Analysis", Oliver & Boyd.

[8] Balas,G.J., Doyle, J.C., Glover, K., Packard, A., and Smith, R., (1995), "mu-Analysis and Synthesis Toolbox", The MathWorks Inc., Natick, Mass..

[9] Balkema, A. and de Haan, L., (1974), "Residual lifetime at great age", Annals of Probability.

[10] Barberis, N., Huang, M. and Santos, T., (1999), "Prospect theory and asset prices", Unpublished manuscript.

[11] Barro, R., (1986), "Reputation in a model of monetary policy with incomplete information", Journal of Monetary Economics.

[12] Basar, T. and Bernhard, P., (1995), "H∞ optimal control and related minimax design problems - a dynamic game approach", Birkhauser.

[13] Bates, S., (1996), "Testing option pricing models", in Handbook of Statistics, volume 14, edited by Maddala, G.S. AN Rao, C.R..

[14] Bellman, R. (1957), "Dynamic Programming", Princeton University Press.

[15] Benninga, S., (1997), "Financial Modeling", MIT Press, Cambridge, Mas-sachusetts.

[16] Bernhard, P., (2000), "A robust control approach to option pricing", Unpublished manuscript, University of Nice-Sophia Antipolis.

[17] Bernhard, P., Crepey, S. and Rapaport, A., (2000), "Comparison of two numerical schemes for barrier value of a simple pursuit-evasion game", Annals of Dynamic Games.

[18] Black, F. and Scholes, M., (1973), "The pricing of options and corporate liabilities", Journal of Political Economy.

[19] Bollerslev, Engle and Nelson, (1994), "ARCH MODELS", Chapter 49, In "Handbook Of Econometrics", Volume IV, North Holland, 1994.

[20] Bolzern, P., Colaneri, P. and Nicolao, G., (1999), "Robust performance in adaptive filtering", in Springer Verlag Lecture Notes in Control and Information Sciences, number 245 by Garulli, A., Tesi, A. and Vicino. A., Eds.

[21] Borici, A. and Luthi, H-J., (1998), "Pricing American put options by fast solutions of the linear complementarity problem",

[22] Brainard, W., (1967), "Uncertainty and the effectiveness of policy", American Economic Review.

[23] Bray, M., (1981), "Futures trading, rational expectations and the efficient markets hypothesis", Econometrica, p. 575-596.

[24] Brogan, W.L. (1991), "Modern Control Theory", Prentice Hall.

[25] Brown, G. and Toft, K., (1997), "Constucting binomial trees from multiple implied probability distributions", University of Austin-Texas, Working Paper.

[26] Buhlmann, H., Delbaen, F., Embrechts, P. and Shiryaev, A.N., (1995), "No-arbitrage, change of measure and conditional Esscher transforms", Working Paper, University of Zurich, Department of Mathematics.

[27] Campbell, N. A., (1980), "Robust procedures in multivariate analysis I: Robust covariance estimation", Applied Statistics.

[28] Campbell, N. A., (1982), "Robust procedures in multivariate analysis II: robust canonical variate analysis", Applied Statistics.

[29] Carroll, R.J., (1979), "On estimating variances of robust estimators when the errors are symmetric", Journal of the American Statistical Association.

[30] Chadha, J.S. and Schellekens, P., (1999), "Monetary policy loss functions: two cheers for the quadratic", Bank of England Working Paper Number 101.

[31] Chen, B.M., (1998), "H∞ control and its applications", Springer Verlag, Lecture Notes in Control and Information Sciences, Number 235.

[32] Chew, S.H. and MacCrimmon, K.R., (1979), "Alpha-nu Choice Theory: A generalization of expected utility theory", Working Paper No. 669, University of British Columbia, Vancouver.

[33] Chew, S.H., (1983), "A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox", Econometrica, Vol. 51, p.1065-92.

[34] Chew, S.H., (1989), "Axiomatic utility theories with the betweenness property", Annals of Operations Research.

[35] Chew, S.H., Karni, E. and Safra, Z., (1987), "Risk aversion in the theory of expected utility with rank dependent probabilities", Journal of Economic Theory.

[36] Choi, S. and Wohar, M., (1992) , "Implied volatility in options markets and conditional heteroskedasticity in stock markets", Financial Review.

[37] Christensen, C.V., (1999), "A new model for pricing catastrophe insurance derivatives", Centre for Analytical Finance, University of Aarhus, Aarhus School of Business, Working Paper Series, Number 28.

[38] Crandall, M.G. and Lions, P.L., (1983), "Viscosity solutions of Hamilton-Jacobi equations", Transactions of the American Mathematical Society, volume 277.

[39] Crouhy, M., Galai, D. and Mark, R., (1999), "Model risk", Journal of Financial Engineering, volume 7, number 3.

[40] Dacorogna, M.M., Gencay, R., Muller, U., Olsen, R.B. and Pictet, O.V., (2001), "An introduction to high frequency finance", Academic Press.

[41] Davidson, P., (1982), "Rational expectations: a fallacious foundation for studying crucial decision-making processes". Journal of Post Keynesian Economics.

[42] Davidson, P., (1995), "Uncertainty in Economics". In: Dow, S.,Hillard, J. (Eds.) "Keynes, knowledge and uncertainty", Aldershot.

[43] De Jong, A., De Roon, F. and Veld, C., (1995), "An empirical analysis of the hedging effectiveness of currency futures", Department of Business Administration, Tilburg University.

[44] De Losso, R. and Bueno, S., (2003), "Dynamic Hedging with Stochastic Differential Utility", Department of Economics, University of Chicago, September, 2003.

[45] Dempster, M.A.H. and Hutton, J.P., (1999), "Pricing American Options by Linear Programming", Judge Institute for Management Studies Working Paper.

[46] Dempster, M.A.H. and Richards, D.G., (1999), "Pricing exotic American Options fitting the volatility smile", Centre for Financial Research, Judge Institute of Management Studies, University of Cambridge.

[47] Derman, E., Deniz, E., Kani, I., (1994), "Forever Hedged", RISK, 1994, vol. 7, no. 9, p. 141–145.

[48] Derman, E., Deniz, E., Kani, I., (1994), "Static Options Replication", The Journal of Derivatives, 1995, vol. 2, no. 4, p. 78–95.

[49] Dow, J. and Werlang, S.R., (1992), "Excess volatility of stock prices and Knightian uncertainty", European Economic Review.

[50] Duffie, D., (1992), "Dynamic Asset Pricing Theory", Princeton University Press.

[51] Duffie, D. and Epstein, L.G., (1992), "Stochastic differential utility", Econometrica.

[52] Duffie, D., Geoffard, P.-Y. and Skiadas, C., (1994),"Efficient and equilibrium allocations with stochastic differentiable utility", Journal of Mathematical Economics.

[53] Duffie, D. and Richardson, H.R., (1991), "Mean-variance hedging in continuous time", The Annals of Applied Probability, 1991, volume 1.

[54] Dumas, B., Uppal, R. and Wang, T., (1997), "Efficient intertemporal allocations with recursive utility", University of British Columbia Working Paper, September.

[55] Ellsberg, D., (2001), "Risk, Ambiguity and Decision", Garland Publishing.

[56] Embrechts, P., Kluppelberg, C. and Mikosch, T., (1997), "Modelling extremal events", Springer Verlag.

[57] Epstein, L. (1987), "The global stability of efficient intertemporal allocations", Econometrica.

[58] Epstein, L. and Zin, H. (1989), "Substitution, risk aversion and the temporal behaviour of consumption and asset returns: a theoretical framework", Econometrica.

[59] Epstein, L. and Zin, H. (1990), "First order risk aversion and the equity premium puzzle", Journal of Monetary Economics.

[60] Epstein, L.(1992), "Behaviour under risk: Recent developments in theory and applications", in Advances in Economic Theory: Sixth World Congress, edited by J.-J. Laffont, Cambridge University Press.

[61] Epstein, L.(1999), "A definition of uncertainty aversion", Review of Economic Studies.

[62] Fishburn, P.C., (1983), "Transitive Measurable Utility", Journal of Economic Theory, Vol. 31, p.293-317.

[63] Fischhoff, B., (1991) " Value elicitation: Is there anything in there ?", American Psychology, volume 46.

[64] Fischer, R. and Tippett, L., (1928), "Limiting forms of the frequency distribution of the largest or smallest member of a sample", Proceedings of the Cambridge Philosophical Society.

[65] Freeman, R.A. and Kokotovic, P.V., (1996), "Robust nonlinear control design: state space and Lyapunov techniques", Birkhauser Press.

[66] Frey R. and Stremme A., (1997), "Market Volatility and Feedback Effects from Dynamic Hedging", Mathematical Finance October 1997, vol. 7, no. 4, pp. 351-374(24).

[67] Garber, P. M. and M. G. Spencer, (1995), "Foreign exchange hedging and the interest rate defense", IMF Staff Papers 42, 490-516.

[68] Garber, P. M. and M. G. Spencer, (1996), "Dynamic hedging and the interest rate defense", in: J.A. Frankel, Gianpaolo Galli and Alberto Giovannini, The Microstructure of Foreign Exchange Markets, Chicago and London: University of Chicago Press, 209-226.

[69] Geman, H. and Yor, M., (1997), "Stochastic time changes in catastrophe option pricing", Insurance Mathematics and Economics, volume 21.

[70] Gerber, H.U. and Shiu, E.S.W., (1996), "Actuarial bridges to dynamic hedging and option pricing", Insurance: Mathematics and Economics, volume 18.

[71] Geoffard, P.-Y., (1996) "Discounting and optimizing: capital accumulation as a variational minimax problem", Journal of Economic Theory.

[72] Giannoni, M.P., (1999), "Does model uncertainty justify caution ? Robust optimal monetary policy in a forward looking model", Unpublished Princeton University Working Paper.

[73] Gigerenzer, G. and Selten, R., (2001), "Bounded Rationality - The Adaptive Toolbox", Dahlem Workshop reported, published by MIT Press.

[74] Gilboa, I., (1987), "Expected utility theory with purely subjective non-additive probabilities", Journal of Mathematical Economics.

[75] Gilboa, I. and Schmeidler, D., (1989), "Maximin expected utility with a non-unique prior", Journal of Mathematical Economics.

[76] Gilboa, I. and Schmeidler, D., (1991), "Updating ambiguous beliefs", Northwestern University Working Paper.

[77] Gilboa, I. and Schmeidler, D., (1991), "A theory of case based decisions", Cambridge University Press.

[78] Glover, K. and Doyle, J.C., (1998), "State-space formulae for all stabilizing controllers that satisfy an Hinf norm bound and relations to risk sensitivity", Systems and Control Letters, vol. 11, pp. 167-172, 1988.

[79] Gomes, D.A., (2000), "Viscosity solutions of Hamilton-Jacobi equations and asymptotics for Hamiltonian systems", unpublished manuscript.

[80] Gomes, D.A., (2001), "Perturbation theory for viscosity solutions of Hamilton-Jacobi equations and stability of Aubrey-Mather sets", unpublished manuscript.

[81] Gosh, M., Parr, W.C., Singh, K. and Babu, G.J., (1984), "A note on bootstrapping the sample median", Annals of Statistics.

[82] Gross, A.M., (1976), "Confidence interval robustness with long-tailed symmetric distributions", Journal of the American Statistical Assocition.

[83] Greene, W.H., (1997), "Econometric Analysis", Prentice Hall.

[84] Grether, D. and Plott, C., (1979), "Economic Theory of Chocie and the Preference Reversal Phenomenon", American Economic Review, Vol. 69, p.623-38.

[85] Grossman, S.J. and Vila, J-L., (1992), "Optimal dynamic trading with leverage constraints", Journal of Financial and Quantitative Analysis, volume 27.

[86] Gumbel, E., (1958), "Statistics of Extremes", Princeton University Press.

[87] Hagen, O., (1972), "A New Axiomatization of Utility under Risk", Theorie A Metoda, Vol. 4, p.55-80.

[88] Hampel, F.R., Ronchetti, E.Z., Rousseeuw, P.J. and Stahel, W.A., "Robust Statistics. The approach based on influence functions", Wiley, New York.

[89] Hart, A.G., (1942), "Risk, Uncertainty and the Unprofitability of Compounding Probabilities", in Lange et al., editors, Mathematical Economics and Econometrics.

[90] Harvey, A.C. (1993), "Time Series Models", Prentice Hall, second edition.

[91] Hassibi, B., Sayed, A.H. and Kailath, T., (1996), "H∞ optimality of the least mean squares algorithm", IEEE Transactions in Signal Processing, AC-44.

[92] Heston, S., (1993), "A closed form solution for options with stochastic volatility with applications to bond and currency options", Review of Financial Studies.

[93] Hicks, J., (1931), "Theory of Uncertainty and Profit", Economica.

[94] Hipp, C. and Vogt, M., (2003), "Optimal dynamic excess of loss reinsurance", Astin Bulletin, volume 33, p. 193-207.

[95] Holt, C., (1986), "Preference Reversals and the Independence Axiom", American Economic Review, Vol. 76 (3), p.508-15.

[96] Huang, J. and Pang, J-S, (1998), "Option pricing and linear complementarity", Journal of Computational Finance, volume 2, issue 3.

[97] Huber, P. J., (1964), "Robust estimation of a location parameter", Annals of Mathematical Statistics.

[98] Huber, P. J., (1973), "Robust regression: asymptotics, conjectures and Monte Carlo", Annals of Statistics.

[99] Huber, P. J., (1977), "Robust statistical procedures", SIAM, Philadelphia.

[100] Huber, P. J., (1981), "Robust statistics", Wiley, New York.

[101] Ilhan, A. and Sircar, R., (2004), "Optimal Static-Dynamic Hedges for Barrier Options", Princeton University Occasional Paper.

[102] Jacobson, D.H., (1973), "Optimal stochastic linear systems with exponential performance criteria", IEEE Transactions on Automatic Control, p. 124-131.

[103] Jaillet, P., Lamberton, D. and Lapeyre, B., (1990), "Variational inequalities and the pricing of American options", Acta Applicandae Mathematicae.

[104] Jang, J.W., (2000), "Doubly stochastic Poisson processes and the pricing of catastrophe reinsurance contracts", Department of Actuarial Studies, University Of New South Wales.

[105] Jorion, P., (1997), "Value at Risk: The new benchmark for Controlling Deriatives Risk", McGraw Hill.

[106] Kahneman, D. and Tversky, A., (1979), "Prospect theory: an analysis of decisions under risk", Econometrica.

[107] Kalecki, M., (1937), "The Principle of Increasing Risk", Economica.

[108] Karatzas, I. and Shreve, S.E., (1998), "Methods of Mathematical Finance", Springer Verlag.

[109] Karmarkar, U.S., (1978), "Subjectively Weighted Utility: A descriptive extension of the expected utility model", Organizational Behavior and Human Performance, Vol. 21, p.61-72.

[110] Karni, E. and Safra, Z., (1986), "Preference Reversals and the Observability of Preferences by Experimental Methods", Economic Letters, Vol. 20 (1), p.15-18.

[111] Karni, E. and Safra, Z., (1987), "Preference Reversals and the Observability of Preferences by Experimental Methods", Econometrica, Vol. 51, p.675-85.

[112] Keynes, J.M., (1936), "General Theory of Employment, Interest and Money".

[113] Keynes, J.M., (1937), ""Alternative Theories of the Rate of Interest", Economic Journal.

[114] Knight, F., (1921), "Risk, uncertainty and profit", Houghton, Mifflin and Boston.

[115] Krasker, W.S. and Welsch, R.E., (1982), "Efficient bounded influence regression", Journal of the American Statistical Association.

[116] Koopmans, T.C., (1960), "Stationary ordinal utility and impatience", Econometrica.

[117] Krüger, M., (1997), "Dynamic hedging in currency crisis", University of Cologne Working Paper.

[118] Krzanowski, W. J., (1996), "Principles of multivariate analysis - A User's Perspective", Oxford Science Publications.

[119] Kushner, H.J. and Dupuis, S., (1992), "Numerical methods for stochastic control problems in continuous time", Applications of Mathematics, Springer-Verlag.

[120] Lange, O., (1944), "Price Flexibility and Employment", Cowles Commission Monograph.

[121] Laub, A.J. (1985), Numerical Linear Algebra Aspects of Control Design Computations, IEEE Transactions on Automatic Control, Vol. AC-30, No. 2, February, pp. 97-108.

[122] Laurant, J.P. and Pham, H., (1999), "Dynamic programming and mean-variance hedging", Finace and Stochastics, volume 3.

[123] Lewis, A.L., (1999), "Option valuation under stochastic volatility", Finance Press.

[124] Lichtenstein, S. and Slovic, P., (1971), "Reversal of Preferences Between Bids and Choices in Gambling Decisions", Journal of Experimental Psychology, Vol. 89, p.46-55.

[125] Lichtenstein, S. and Slovic, P., (1973), "Response-Induced Reversals of Preference in Gambling: An extended replication in Las Vegas", Journal of Experimental Psychology, Vol. 101, p.16-20.

[126] Lien, D., (2000), "Production and hedging under Knightian uncertainty", Journal of Futures Markets.

[127] Lions, P.L. and Souganidis, P.E., (1985), "Differential games, optimal control and directional derivatives of viscosity solutions of Bellman's and Isaacs' equations", Journal of Control and Optimisation, Society for Applied Mathematics.

[128] Lucas, R., (1976), "Econometric policy evaluation: a critique", Carnegie Rochester Conference Series on Public Policy 1: 19-46.

[129] Luenberger, D.G., (1969), "Optimisation by vector space methods", Wiley.

[130] Lutgens, F.J.W., (2004), "Robust portfolio optimisation", Ph.D thesis, Universiteit Maastricht.

[131] McNeil, A.J. and Frey, R., (2000), "Estimation of tail-related risk measues for heteroskedastic financial time series: an extreme value approach", Journal of Empirical Finance.

[132] Ma, C., Protter, P. and Yong, J., (1994), "Solving forward-backward stochastic differential equations explicitly - a four step scheme", Probability Theory and Related Fields.

[133] Ma, C., (1998), "Attitudes toward the timing and resolution of uncertainty", Journal of Economic Dynamics and Control.

[134] Ma, J. and Yong, J., (1999), "Forward-backward stochastic differential equations and their applications", Springer Verlag, 1702.

[135] Machina, M.J., (1982), "Expected Utility: Analysis without the Independence Axiom", Econometrica, Vol. 50, p.277-323.

[136] Major, J.A., (1997), "A synthetic history of the Guy Carpenter Catastrophe Index: Methodology, Data and Analysis", Guy Carpenter & Company Inc., New York.

[137] Marino, R. and Tomei, P., (), "Robust stabilisation of feedback linearisable time-varying uncertain nonlinear systems", Automatica, volume 38.

[138] Marschak, J., (1951), "Why "Should" Statisticians and Businessmen Maximize Moral Expectation?", Proceedings of the Second Berkeley Symposium.

[139] Martin, B. and Salmon C., (1999), "Should uncertain policy makers do less ?", Bank of England Working Paper, Number 99.

[140] Martinin, C., (2000), "American option prices as unique viscosity solutions to degenerated Hamilton-Jacobi-Bellman equations", INRIA Rapport de recherche, number 3934.

[141] Mas-Colell, A., (1974), "An Equilibrium Existence Theorem without Complete or Transitive Preferences", Journal of Mathematical Economics.

[142] Merton, R.C., (1973), "Theory of rational option pricing", Bell Journal of Economics and Management Science.

[143] Moosa, I.A., (2003), "The Sensitivity of the Optimal Hedge Ratio to Model Specification", Finance Letters, volume 1.

[144] Moscinski, J. and Ogonowski, Z. (1995), "Advanced Control with Matlab and Simulink", Ellis Horwood.

[145] Munos, R. and Moore, A., (2000), "Variable resolution discretisation for high-accuracy solutions of optimal control problems", Working Paper, Robotics Institute, Carnegie Mellon University.

[146] Musiela, M. and Rutkowski, M., (1997), "Martingale Methods in Financial Modelling", Springer Verlag.

[147] Myneni, R., (1992), "The Pricing of the American Option"The Annals of Applied Probability.

[148] Ogata, K. (1995), "Discrete Time Control Systems", Prentice Hall.

[149] Onatski, A. and Stock, J.H., (1998), "Robust monetary policy under model uncertainty in a small robust model of the US economy", Harvard University Department of Economics Working Paper.

[150] Overdahl, J.A. and Martin, P.G., (1994), "The exercise of equity options: theory and empirical tests", Journal of Derivatives.

[151] Petkov, P.Hr., Gu, D.W., and Konstantinov, M.M., (1998), "Fortran 77 routines for Hinf and H2 design of continuous-time linear control systems", Rep. 98-14, Department of Engineering, Leicester University, Leicester, U.K.

[152] Pickands, J., (1975), "Statistical inference using extreme order statistics", Annals of Statistics.

[153] Quiggin, J., (1982), "A Theory of Anticipated Utility", Journal of Economic Behavior and Organization, Vol. 3, p.323-43.

[154] Quiggin, J., (1993), "Generalized Expected Utility Theory: The Rank-Dependent Expected Utility model", Amsterdam: Kluwer-Nijhoff.

[155] Rabin, M., (1999),"Risk aversion and expected utility theory: a calibration theorem", Unpublished paper, Department of Economics, University of California, Berkley.

[156] Radford, A., (1957), "Central Intelligence and National Security", Time Magazine, February 25th, 1957 issue, page 27.

[157] Ramsay, F. (1929), "The Foundations of Mathematics, chapter on Philosophy.

[158] Ramsay, F.P., (1926), "Truth and Probability", in Braithwaite R.B. (ed), "The Foundations of Mathematics and other Logical Essays", London: Kegan, Paul, Trench, Trubner.

[159] Rapaport, A., (1998), "Characterisation of barriers of differential games", Journal of Optimisation, Theory and Applications, volume 97.

[160] Reinsel, G.C. (1997), "Elements of Multivariate Time Series Analysis", Springer.

[161] Rocke, D.M. and Downs, G.W., (1981), "Estimating the variances of robust estimates of location: influence curve, jackknife and bootstrap", Comupatational Statistics Bulletin.

[162] Rousseeuw, P. J., (1984), "Least median of squares regression", Journal of the American Statistical Association.

[163] Rousseeuw, P. J. and Leroy, A. M., (1987), "Robust regression and outlier detection", Wiley, New York.

[164] Rousseeuw, P. J. and Yohai, V.J., (1984), "Robust regression by means of S-estimators", in Robust and Non-Linear Time Series by Franke, Hardle and Martin, eds., Springer Verlag Lecture Notes in Statistics, No. 26.

[165] Rubenstein, M., (1994), "Implied binomial trees", Journal of Finance.

[166] Rudesbusch, G.D. and Svensson, L.E.O., (1998), "Policy rules for inflation targeting", in J. Taylor (Ed.), Policy Rules for Inflation Targeting (University of Chicago Press for the NBER).

[167] Rustem, B. and Howe, M., (2002), "Algorithms for worst-case design and applications to risk management", Princeton University Press.

[168] Sargent, T.J. (2001), "Robust control and model uncertainty", paper presented to American Economic Association, New Orleans, January.

[169] Savage, L.J., (1954), "Foundations of Statistics".

[170] Sayed, A.H. and Kailath, T., (1994), "A state-space approach to adaptive RLS filtering ", IEEE Signal Processing Magazine.

[171] Schmeidler, D., (1989), "Subjective probability and expected utility without additivity", Econometrica.

[172] Schroder, M. and Skiadas, C., (1999), "Optimal consumption and portfolio selection with stochastic differential utility", Unpublished working paper from Northwestern University.

[173] Schroder, M. and Skiadas, C., (2000), "An isomorphism between asset pricing models with and without linear habit formation", Unpublished working paper number 247, Michigan State University.

[174] Sefton, J.A., (1996), "Output gaps and inflation: an approach using a multivariable Beveridge-Nelson decomposition", National Institute of Economic and Social Research.

[175] Shackle, G.L.S., (1949), "Probability and Uncertainty", Metroeconomica.

[176] Shackle, G.L.S., (1955), "Uncertainty in Economics", Cambridge University Press.

[177] *Shackle, G.L.S., (1961), "Decision, Order and Time in Human Affairs", in Bowman, editor, Uncertainty and Business Decision.*

[178] *Shackle, G.L.S., (1979), "Imagination and the Nature of Choice".*

[179] Shreve, S., (1996), "Stochastic Calculus and Finance", Carnegie Mellon University.

[180] Shao, J., (1990), "Bootstrap estimation of the asymptotic variances of statistical functionals", Annals of the Institute of Statistics and Mathematics.

[181] Shao, J., (1992), "Bootstrap variance estimators with truncation", Statistics and Probability Letters.

[182] Silverman, B.W., (1986), "Density estimation for statistics and data analysis", Chapman and Hall, New York.

[183] Simon, D., (2000), "From here to infinity", Embedded Sytems Journal, July.

[184] Simon, H.A., (1957), "Models of Man", published by Wiley, New York.

[185] Sims, C.A., (1984), "Martingale-like behaviour of prices and interest rates", University of Minnesota, Mimeograph.

[186] Smith, C.W. and Stulz, R.M., (1985), "The determinants of firms' hedging policies", Journal of Financial and Quantitative Analysis, volume 20.

[187] Souganidis, A., (1985), "Max-min representations and product formulas for the viscosity solutions of Hamilton-Jacobi equations with applications to differential games", Non-Linear Annals of Transactions of the American Mathematical Society.

[188] Steil, B., (1992), "Currency options and the optimal hedging of contingent foreign exchange exposure", Discussions Papers in Economics, Number 72, Nuffield College Oxford.

[189] Stigler, J.G., (1939), "Production and Distribution in the Short Run", Journal of Political Economy.

[190] Stiglitz, J. E., (2003), "Terrorism: There's No Futures in It. Los Angeles Times, Thursday, July 31, 2003 "

[191] Taylor, J.B. (1993), "Discretion versus policy rules in practice", Carnegie Rochester Conference Series on Public Policy, Number 39.

[192] Taylor, J.B. (1998a), Introductory remarks on monetary policy rules, Monetary Policy Rules published by the National Bureau for Economic Research.

[193] Taylor, J.B. (1998b), An historical analysis of monetary policy rules, National Bureau of Economic Research Working Paper No. 6768.

[194] Taylor, J.B. (1999a), The robustness and efficiency of monetary policy rules as guidelines for interest rate setting by the European Central Bank, Unpublished Working Paper presented at a conference on Monetary Policy Rules at the Institute for International Economic Studies in June 1998.

[195] Taylor, J.B. (1999b), "Monetary policy rules, National Bureau of Economic Research", University of Chicago Press, ISBN 0-226-79124-6.

[196] Thaler, R.H. and Johnson, E.J., (1990), "Gambling with the house money and trying to break even: the effects of prior outcomes on risky choice", Management Science.

[197] Tintner, G., (1941), "The Theory of Choice under Subjective Risk and Uncertainty", Econometrica.

[198] Tintner, G., (1941), "The Pure Theory of Production under Technological Risk and Uncertainty", Econometrica.

[199] Tversky, A. and Kahneman, D., (1992), "Advances in prospect theory: cumulative representation of uncertainty", Journal of Risk and Uncertainty.

[200] Uppal, R. and Wang, T., (2003), "Model misspecification and underdiversification", Journal of Finance, volume 58, pp. 2465–2486.

[201] Valiani, S., (2004), "Application of Multivariate GARCH in Optimal Dynamic Hedging and International Asset Allocation", March 2004.

[202] Venter, G.G., (1991), "Premium calculation implication of reinsurance without arbitrage", ASTIN Bulletin 21.

[203] Vickers, J., (1986), "Signalling in a model of monetary policy with incomplete information", Oxford Economic Papers.

[204] von Neumann, J., (1944), "Theory of Games and Economic Behavior", with O. Morgenstern.

[205] Wojakowski, R., (1998), "Hedging level risk for a firm", Paper presented to the XIVth International Conference of the French Finance Association, 23-25 June 1997 in Grenoble, France.

[206] Woodford, M., (1996), "Control of the public debt: a requirement for price stability ?", NBER Working Paper Number 5684.

[207] Woodford, M., (1999), "Optimal monetary policy inertia", NBER Working Paper Number 7261.

[208] Wystup, U., (2000), "Valuation of exotic options under short selling constraints as a singular stochastic control problem", Ph.D thesis Carnegie Mellon University, Department of Mathematical Sciences.

[209] Yohai, V.J., (1987), "High breakdown point and high efficiency robust estimates for regression", Technical Report No. 66, Department of Statistics, University of Washington.

[210] Yohai, V.J. and Zamar, R.H., (1988), "High breakdown point estimates of regression by means of the minimization of an efficient scale", Journal of American Statistical Society.

[211] Zames, G., (1981), "Feedback and optimal sensitivity: model reference transformations, multiplicative seminorms and approximate inverses", IEEE Transactions on Automatic Control, Volume 26.

[212] Zhou, K., Doyle, J.C. and Glover, K., (1996), "Robust and optimal control", Prentice Hall.

[213] Ziegler, A., (1999), "A Game Theory Analysis of Options - Contributions to the Theory of Financial Intermediation in Continuous Time", Lecture Notes in Economics and Mathematical Systems, Number 468, Springer Verlag.

Chapter 7

Appendices

Appendix 1: Robusta analytics

1. Ported Slicot Library Functions

The following list contains the names of those functions ported from the Slicot Fortran library into the C++ Robusta equivalent function names

Synthesis Routines

SB01BD/[SSPole]: Pole assignment for a given matrix pair (A,B).

SB01DD/[SSEigenstructure]: Eigenstructure assignment for a controllable matrix pair (A,B) in orthogonal canonical form.

SB01MD/[SSFeedback]:State feedback matrix of a time-invariant single-input system.

SB02MD/[SSARESchur]: Solution of algebraic Riccati equations (Schur vectors method).

SB02MT/[SSConvert]: Conversion of problems with coupling terms to standard problems.

SB02ND/[SSOptimalFeedback]: Optimal state feedback matrix for an optimal control problem.

SB02OD/[SSRiccatiSchur] Solution of algebraic Riccati equations (generalized Schur method)

SB02PD/[SSRiccatiMatrix] Solution of continuous algebraic Riccati equations (matrix sign function method) with condition and forward error bound estimates

SB02QD/[SSRiccatiError] Condition and forward error for continuous Riccati equation solution

SB02RD/[SSRiccatiSchurV] Solution of algebraic Riccati equations (refined Schur vectors method) with condition and forward error bound estimates.

SB02SD/[SSRiccatiError] Condition and forward error for discrete Riccati equation solution

Lyapunov Equations

SB03MD/[LyapunovSep] Solution of Lyapunov equations and separation estimation

SB03OD/[LyapunovChol] Solution of stable Lyapunov equations (Cholesky factor)

SB03PD/[LyapunovDis] Solution of discrete Lyapunov equations and separation estimation

SB03QD/[LyapunovError] Condition and forward error for continuous Lyapunov equations

SB03RD/[LyapunovCont] Solution of continuous Lyapunov equations and separation estimation

SB03SD/[LyapunovDisErr] Condition and forward error for discrete Lyapunov equations

SB03TD/[LyapunovContErr] Solution of continuous Lyapunov equations, condition and forward error estimation.

SB03UD/[LyapunovDis] Solution of discrete Lyapunov equations, condition and forward error estimation.adbeat control state feedback matrix

Transfer Matrix Factorization

SB08CD/[TMFLeftCoprime] Left coprime factorization with inner denominator

SB08DD/[TMFRightCoprime] Right coprime factorization with inner denominator

SB08ED/[TMFLeftCoprimeSD] Left coprime factorization with prescribed stability degree

SB08FD/[TMFRightCoprimeSD] Right coprime factorization with prescribed stability degree

SB08GD/[TMFLeftCoprimeSS] State-space representation of a left coprime factorization

SB08HD/[TMFRightCoprimeSS] State-space representation of a right coprime factorization

SB08MD/[TMFSpectralCont] Spectral factorization of polynomials (continuous-time case)

SB08ND/[TMFSpectralDis] Spectral factorization of polynomials (discrete-time case)

Realization Methods

SB09MD/[RMSequence] Closeness of two multivariable sequences

Optimal Regulator Problems

SB10DD/[ORHinfinityDis] H-infinity (sub)optimal state controller for a discrete-time system

SB10ED/[ORH2Dis] H2 optimal state controller for a discrete-time system

SB10FD/[ORHinfinityCont] H-infinity (sub)optimal state controller for a continuous-time system

SB10HD/[ORH2Cont] H2 optimal state controller for a continuous-time system

SB10ID/[ORFBControllerCont] Positive feedback controller for a continuous-time system

SB10KD/[ORFBControllerDis] Positive feedback controller for a discrete-time system

Controller Reduction

SB16AD/[CRControllerSP] Stability/performance enforcing frequency-weighted controller reduction

SB16BD Coprime factorization based state feedback controller reduction

SB16CD Coprime factorization based frequency-weighted state feedback controller reduction

Generalized State-Space Synthesis - Generalized Lyapunov Equations

SG03AD Solution of generalized Lyapunov equations and separation estimation

SG03BD Solution of stable generalized Lyapunov equations (Cholesky factor)

State-Space Transformation Routines

TB01ID Balancing a system matrix for a given triplet

TB01KD Additive spectral decomposition of a state-space system

TB01LD Spectral separation of a state-space system

TB01MD Upper/lower controller Hessenberg form

TB01ND Upper/lower observer Hessenberg form

TB01PD Minimal, controllable or observable block Hessenberg realization

TB01TD Balancing state-space representation by permutations and scalings

TB01UD Controllable block Hessenberg realization for a state-space representation

TB01WD Reduction of the state dynamics matrix to real Schur form

TB01ZD Controllable realization for single-input systems

State-Space to Polynomial Matrix Conversion

TB03AD Left/right polynomial matrix representation of a state-space representation

State-Space to Rational Matrix Conversion

TB04AD Transfer matrix of a state-space representation

State-Space to Frequency Response

TB05AD Frequency response matrix of a state-space representation

TC - Polynomial Matrix

Polynomial Matrix Transformations

TC01OD Dual of a left/right polynomial matrix representation

Polynomial Matrix to State-Space Conversion

TC04AD State-space representation for left/right polynomial matrix representation

Polynomial Matrix to Frequency Response

TC05AD Transfer matrix of a left/right polynomial matrix representation

TD - Rational Matrix

Rational Matrix to Polynomial Matrix Conversion

TD03AD Left/right polynomial matrix representation for a proper transfer matrix

Rational Matrix to State-Space Conversion

TD04AD Minimal state-space representation for a proper transfer matrix

Rational Matrix to Frequency Response

TD05AD Evaluation of a transfer function for a specified frequency

TF - Time Response

TF01MD Output response of a linear discrete-time system

TF01ND Output response of a linear discrete-time system (Hessenberg matrix)

TF01OD Block Hankel expansion of a multivariable parameter sequence

TF01PD Block Toeplitz expansion of a multivariable parameter sequence

TF01QD Markov parameters of a system from transfer function matrix

TF01RD Markov parameters of a system from state-space representation

TG - Generalized State-space

Generalized State-space Transformations

TG01AD Balancing the matrices of the system pencil corresponding to a descriptor triple

TG01BD Orthogonal reduction of a descriptor system to the generalized Hessenberg form

TG01CD Orthogonal reduction of a descriptor system pair (A-sE,B) to the QR-coordinate form

TG01DD Orthogonal reduction of a descriptor system pair (C,A-sE) to the RQ-coordinate form

TG01ED Orthogonal reduction of a descriptor system to a SVD coordinate form

TG01FD Orthogonal reduction of a descriptor system to a SVD-like coordinate form

TG01HD Orthogonal reduction of a descriptor system to the controllability staircase form

TG01ID Orthogonal reduction of a descriptor system to the observability staircase form

TG01JD Irreducible descriptor representation

TG01WD Reduction of the descriptor dynamics matrix pair to generalized real Schur form

2. Robusta system overview

Robusta is a three tier system. The front-end graphical user interface is written in Visual Basic version 6. The middle tier analytics are written in C++ and implemented as dynamic link libraries available to other C++ and Visual Basic programs. The lower tier is an Access database that holds the market and

trade data required to run the simulations. Figure 1 shows the main menu screen the user sees upon logging in, whilst figure 2 shows the main pricing menu screen and figure 3 shows the

Figure 1: Robusta entry screen

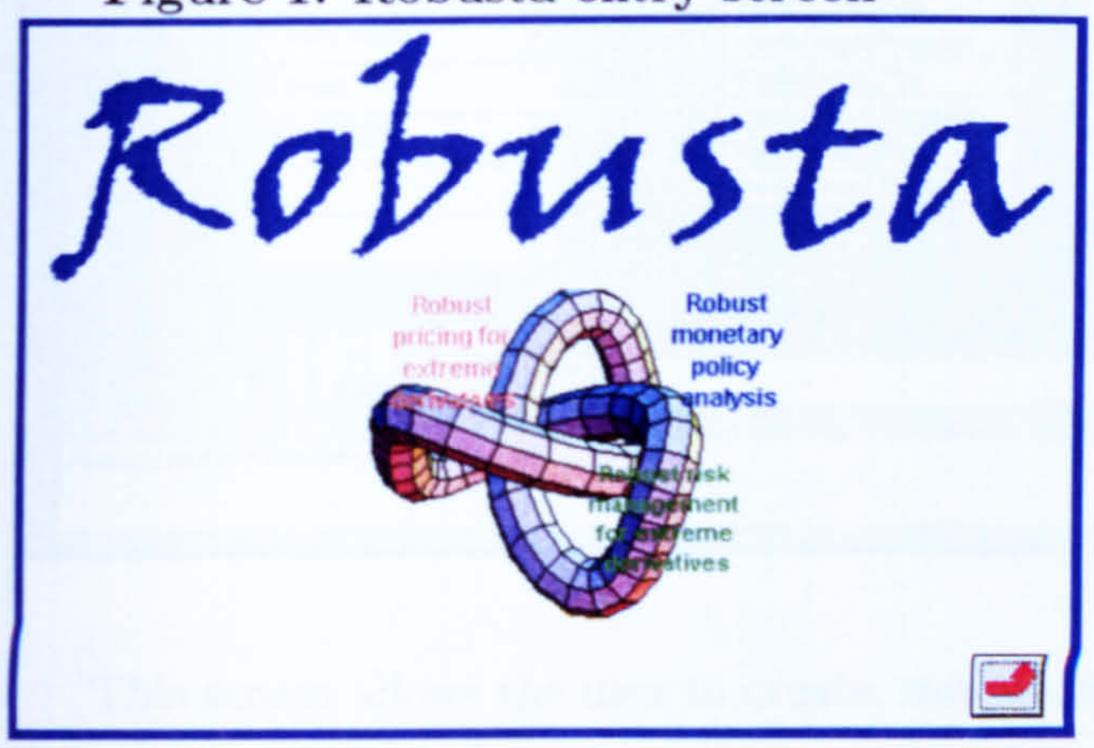

Figure 2: Robust pricing menu screen

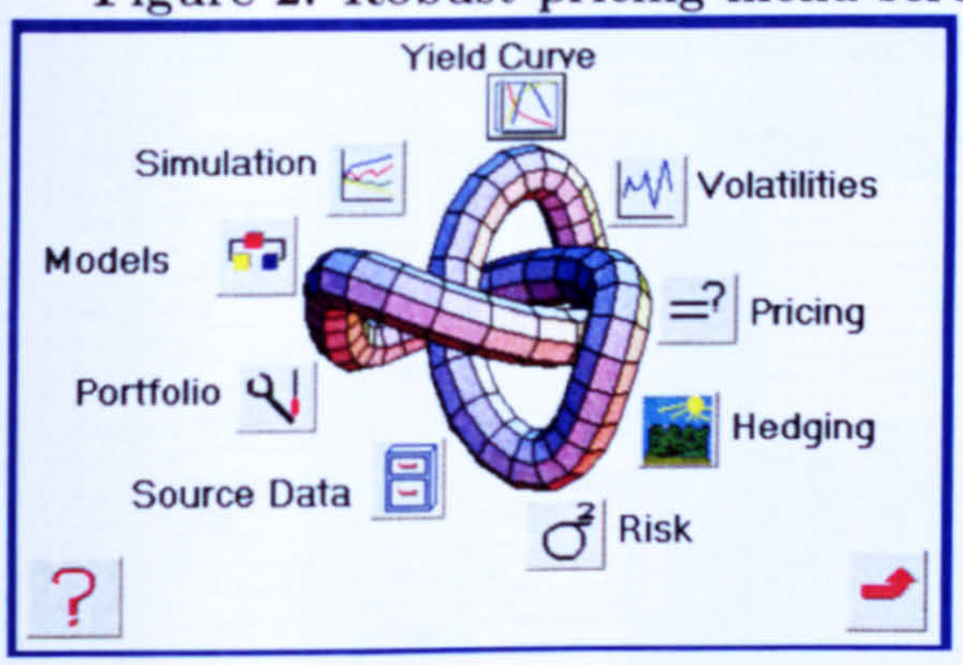

Figure 3: Robusta pricing screen

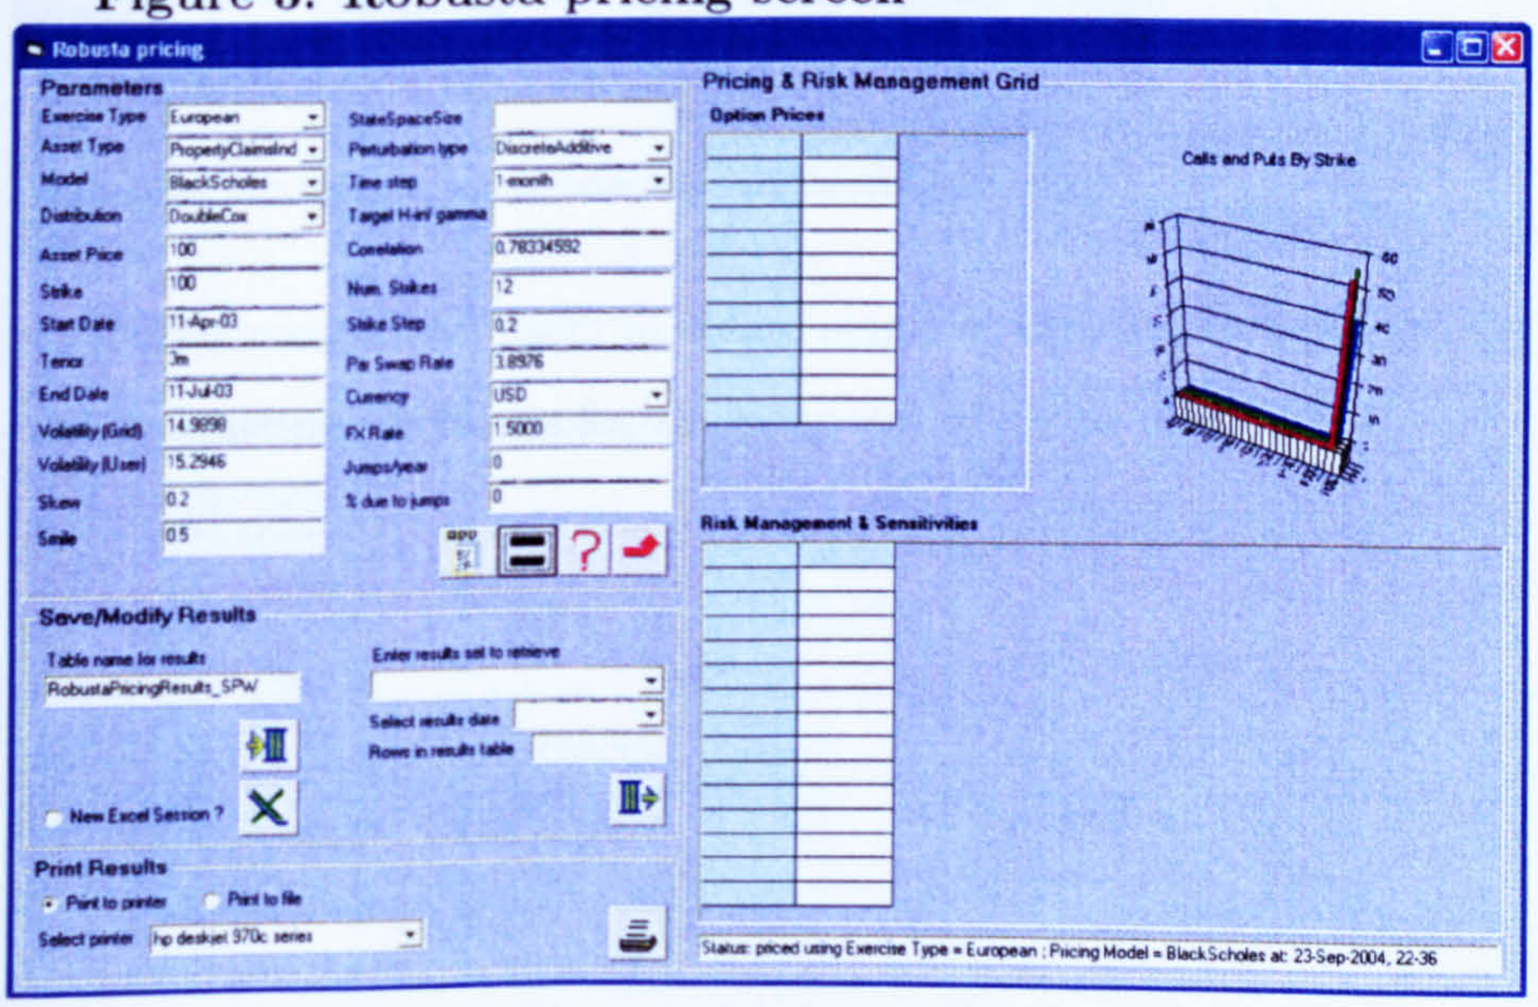

Figure 4: Robusta trade entry screen

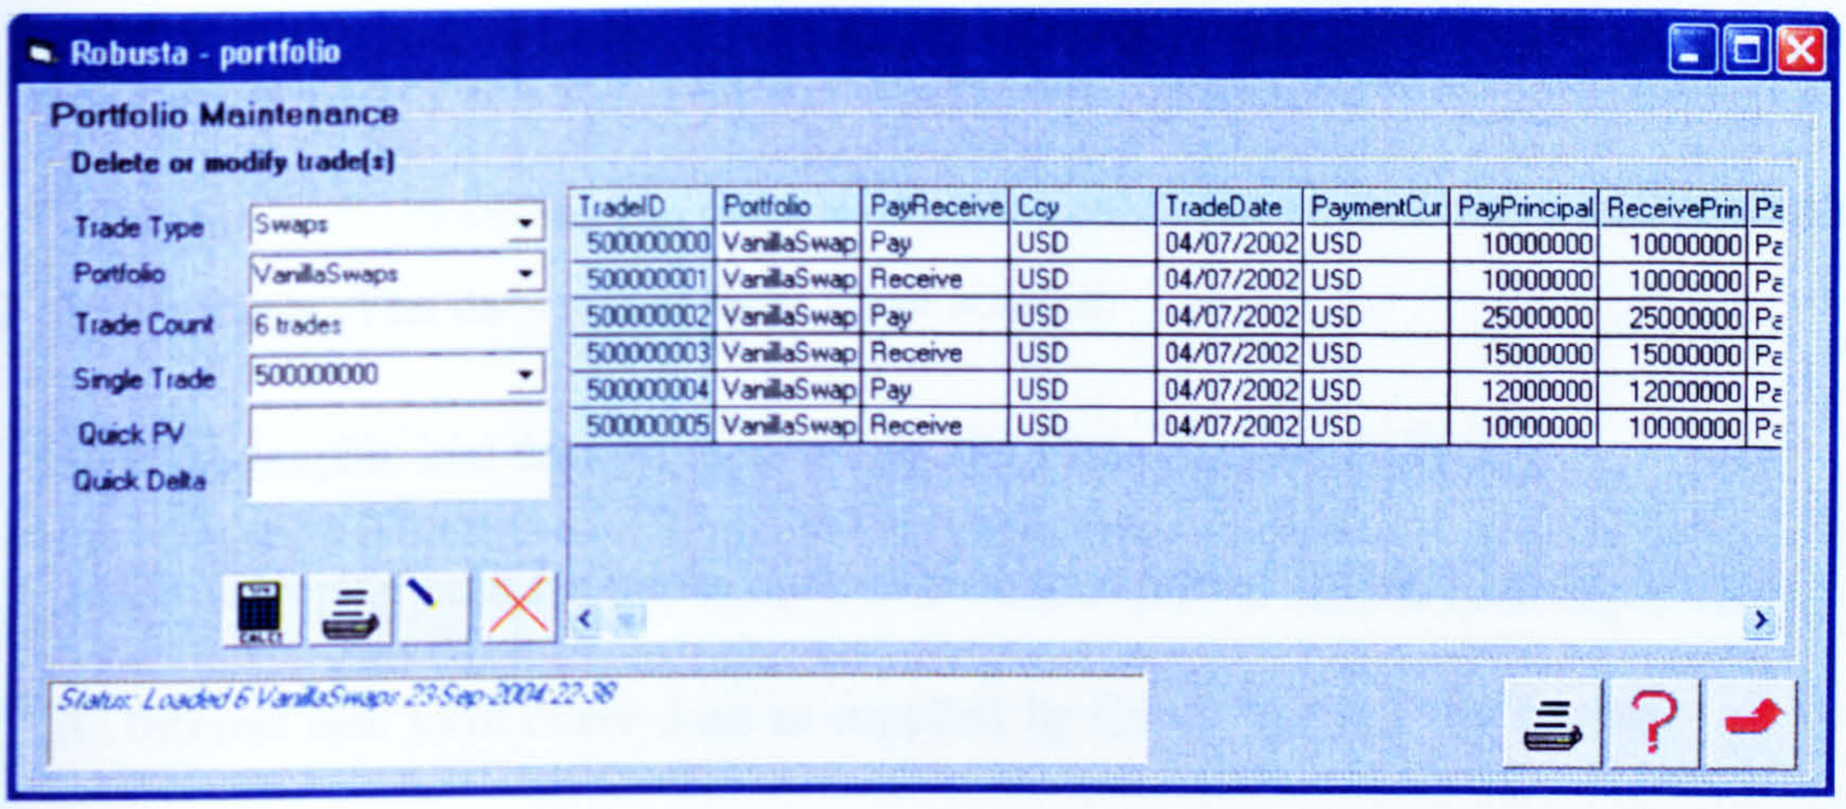

This screen allows the user to create, modify, add and delete trades to portfolios in order to simulate hedging and risk scenarios.

Figure 5: Robusta hedging screen

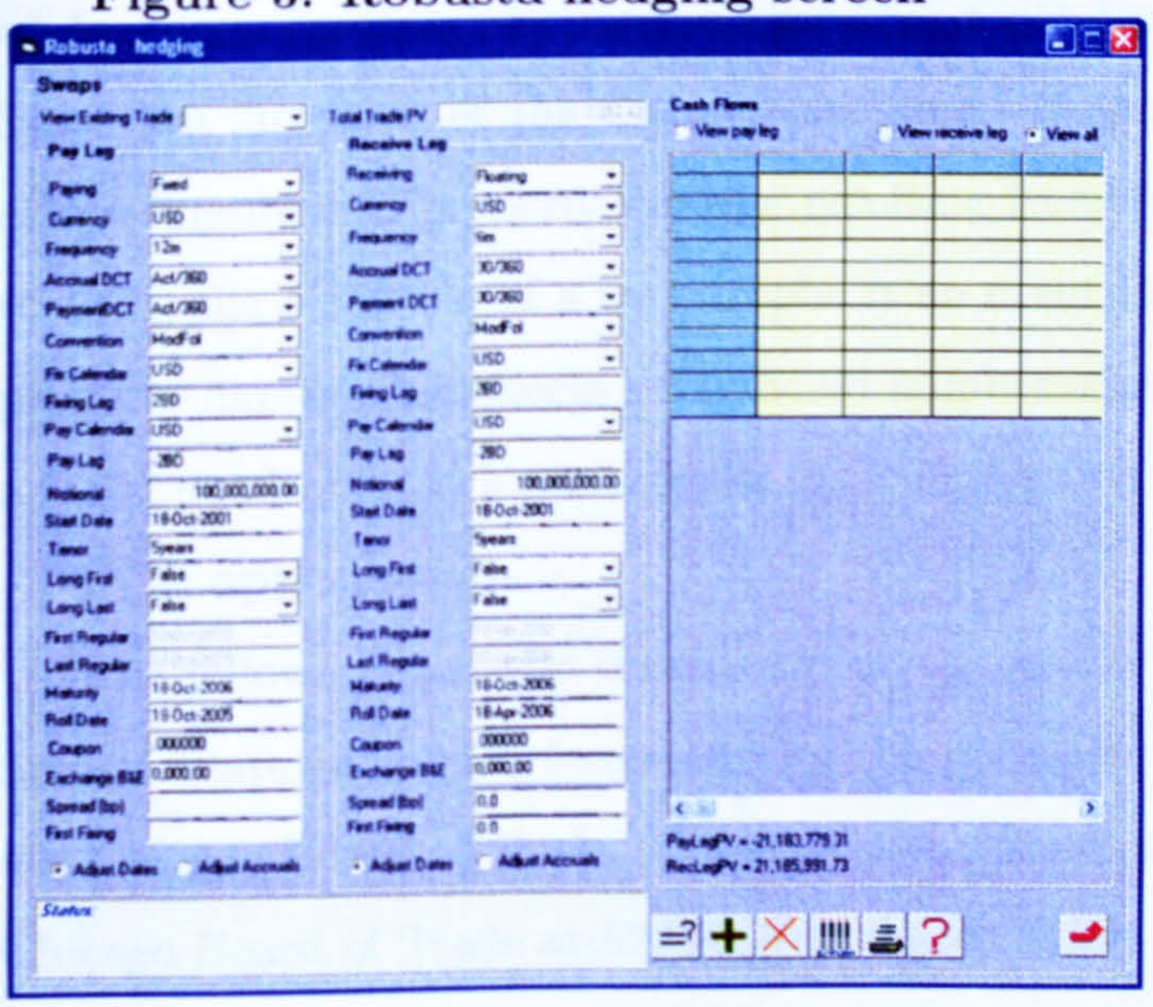

This screen allows users to create single or multiple hedging trades, price them and then add them to new or existing portfolios for the purposes of simulation.

Appendix 2: PCS data for CAT bond and CAT option valuations

This appendix describes the data used in the calculations performed and results reported in chapter 2 of this thesis. The data comes from three sources:

1. Catastrophic loss data as reported by the Property Claims Service.

2. Catastrophic Loss Insurance options data as reported by the Chicago Board Of Trade[1].

3. Interest rate yield curve data as supplied by Credit Suisse First Boston.

Each of these data sources will now be described in detail in the following sections:

Property Claim Services (PCS)

PCS is the insurance industry recognised authority on insured catastrophic events. PCS is a division of American Insurance Services Group Inc., a not-for-profit organisation serving the insurance industry. Since the inception of the Catastrophe Serial Number system in 1949, PCS has been responsible for estimating insured property damage resulting from catastrophes affecting the United States. According to PCS, the definition of a catastrophe is an event that causes losses in excess of $25 million of insured property damage and affects a significant number of policy holders and insurance companies. PCS assigns a serial number to each catastrophe for identification throughout the industry.

PCS options data

The following material draws extensively on the Appendix "PCS Catastrophe Insurance Options - Salient Features" which forms part of the document "A User's Guide to PCS Options", published by the Chicago Board of Trade. Copies of the document can be obtained from the European Office of the Chicago Board of Trade at 52-54 Gracechurch Street, London, EC3V 0EH. Further details are available via the following URL: **www.cbot.com.**

Features of the PCS and CBOT datasets used in this research

The loss claim data provided by the Property Claims Service of ISO cover the period from August 1949 to the end of December 1999. The data includes the dates, states and types of each catastrophic event, as well as the estimated total loss claims in then current dollars for the whole property/casualty insurance industry. PCS assigns a unique serial number to every event where the losses exceed a predetermined threshold value and which has a simultaneous impact on a large number of both policyholders and insurers. From inception to 1983 the threshold value was $1 million, but was increased to $5 million after 1983 and $25 million with effect from the beginning of 1997.

[1]The underlying for these options was the catastrophic loss index data as supplied by the Property Claims Service.

During the time that PCS options traded on CBOT, nine PCS indices were provided to the CBOT daily: a National index; five regional indices covering losses in Eastern, Northeastern, Southeastern, Midwestern, and Western exposures; and three state indices that covered losses in Florida, Texas, and California. Each PCS index value represents the then current PCS estimates for insured catastrophic losses, but divided by $100 million, then rounded to the nearest first decimal point. For instance, if estimated losses total $5,232,780,000, the index value would be 52.3. There is generally a three to five day lag between the initial PCS release of a first estimate of claims after a catastrophe. Re-estimates are released within 60 days if catastrophic losses exceed $250 million. PCS continues to update loss estimates until an accurate estimate is felt to have been achieved. The loss claim data was used directly as the underlying for the CAT bond research.

The Chicago Board of Trade began trading of CAT futures and options on the futures in 1992. These CAT futures and CAT options based on ISO's loss ratio index failed to gain a material level of market interest and were replaced by PCS options in September 1995. Unfortunately, after just over 3 years of trading activity, the PCS options also stopped trading in the summer of 1999 due to inadequate open interest. During this period, CAT bonds were found to be more popular than exchange-based products. However, the lack of secondary market activity for CAT bonds means that there is little information about the pricing patterns of such catastrophe-linked securities. As far as PCS options traded on the CBOT are concerned, a total of less than 600 transactions were completed over the three-and-half-year period during which the contracts traded. Although this is a small volume, the results of the open out-cry transaction mechanism do at least provide what were at the time realistic bid/ask market prices.

The following table shows the contracts that were actually used in the research reported in chapter 2:

Chicago Board of Trade

CATASTROPHE INSURANCE FUTURES/OPTIONS	Code	Period Coverage
NAT'L CATASTROPHE	UN	12/11/1992-11/15/1995
NAT'L CATASTROPHE CALL	UNC	12/11/1992-11/15/1995
NAT'L CATASTROPHE PUT	UNP	12/11/1992-11/15/1995

PROPERTY CLAIM SERVICES (PCS) INSURANCE OPTIONS	Code	Period Coverage
NAT'L ANNUAL INS LARGE CAP-6 MONTH CALL	DHC	10/04/1996-11/03/1997
NAT'L ANNUAL INS. LARGE CAP-6 MONTH PUT	DHP	10/04/1996-11/03/1997
NAT'L ANNUAL INS. SMALL CAP-6 MONTH CALL	QHC	10/04/1996-11/03/1997
NAT'L ANNUAL INS. SMALL CAP-6 MONTH PUT	QHP	10/04/1996-11/03/1997
NAT'L INS. LARGE CAP-6 MONTH CALL	DNC	09/29/1995-11/03/1997
NAT'L INS LARGE CAP-6 MONTH PUT	DNP	09/29/1995-11/03/1997
NAT'L INS. SMALL CAP-6 MONTH CALL	QNC	09/29/1995-11/03/1997
NAT'L INS. SMALL CAP-6 MONTH PUT	QNP	09/29/1995-11/03/1997

Appendix Table 2.1: PCS Derivative Codes and Operational Periods
(All data supplied in electronic form by Chicago Board of Trade)

PCS Methodology for Estimating Catastrophe Damage

PCS compiles its estimated of insured property damage using a combination of procedures, including

a general survey of insurers, its National Insurance Risk Profile and where appropriate its own on the ground survey. PCS estimates take into account both the expected dollar loss and the projected number of claims to be filed. A survey of companies, agents and adjusters is one part of the estimating process PCS also conducts confidential surveys of at least 70% of the market based on premium written market share. PCS then develops a composite of individual loss and claim estimates reported by these sources. Using both actual and projected claim figures, PCS then extrapolates to a total industry estimate by comparing this information to market share data.

PCS also relies on its National Insurance Risk Profile in preparing an insured property damage estimate. This PCS developed profile includes an inventory of buildings and insured vehicles in each of the over 3,100 counties of the USA and is based on census, taxation and other demographic data. The inventory is decomposed into residential, multiple family and commercial buildings. The inventory also includes PCS estimates of the number of vehicles protected by comprehensive insurance. Using this information, PCS can determine the approximate number of insurable risks within the specific geographic area affected by a catastrophe. This information is coupled with the results of insurer surveys to provide a specific outline of the scope of the damage.

Index Valuation

Each PCS loss index represents the sum of the then-current PCS estimates for insured catastrophic losses in the areas and the loss period covered divided by $100 million. each index is quoted in points and tenths of a point. So for example, if loss estimates for the PCS September Eastern index totalled $2 billion, that index would be valued at 20.0 points ($2,000,000,000/$100,000,000), with each index point being worth $200 in cash value. Table shows

PCS Loss Index Value	Index Loss Equivalent Value
0.1	$10 million
1.0	$100 million
20.0	$2 billion
50.0	$5 billion
100.0	$10 billion
200.0	$20 billion (small cap limit)
250.0	$25 billion
300.0	$30 billion
350.0	$35 billion
400.0	$40 billion
450.0	$45 billion
500.0	$50 billion (large cap limit)

Small/Large Cap Option contract specifications

Each PCS index will have both small cap and large cap option contracts listed for trading. Small cap contracts track aggregate estimated catastrophic losses from $0 to $20 billion, whilst large cap contracts track aggregate estimated catastrophic losses from $20 billion to $50 billion.

The loss period is the time during which a catastrophic event must occur in order for resulting losses to be included in a particular index. PCS options can have either quarterly and annual loss periods. Each loss period is tracked by small and large cap quarterly or annual contracts as appropriate, as shown in table .

Contract	Loss Period Covered
March	First quarter
June	Second quarter
September	Third quarter
December	Fourth quarter
Annual	Calendar year

The development period is the time after the loss period during which PCS estimates and re-estimates for catastrophes that occurred during the loss period continue to affect the PCS indices. PCS option users can choose either a 6 or 12 month development period. The PCS index value at the end of the chosen development period will be used for settlement purposes, even though PCS loss estimates may continue

to change.

Strike value are listed in integral multiples of 5 points. For small cap contracts, strike values range from 5 to 195. For large cap contracts, strike values range from 200 to 495. Premia are quoted in points and tenths of a point (the market tick is a tenth of a point), with each point being worth $200 (clearly each tenth of point is therefore worth $20), so that a premium of 5.2 equates to a value of $1,040. Exercise style is European, with in-the-money options being automatically exercised by 6.00 p.m. on the day of expiration. All options expire at 6.00 p.m. on the same day in which the settlement value of the underlying index in made publicly available, either 6 months or 12 months after the loss period (on the last business day of the period), depending on the development period of the option.

Small cap contracts settle in cash to the lesser of a) $200 x the settlement value of the index or, b) $40,000 (220 cap x $200). Large cap contracts settle in cash to the lesser of: a) $200 x the settlement value of the index or, b) $100,000 (500 cap x $200). Large cap contracts have a lower bound of 200, or $40,000 Last day of trading will be the last business day of the 6th calendar month following the loss period for options with a 6 month development period and the last business day of the 12th month following the loss period for options with a 12 month development period.

PCS call spread options are standardized, exchange-based contracts that track the PCS catastrophe loss indices. The options are European style in that they can not be exercised prior to the expiration date. However, they can be closed by selling or buying exactly the same option contract. Each contract is characterized by four factors: regional coverage, lower and upper strike prices, loss period, and development period. Other regions covered correspond to the nine regional PCS indices, so that nine regional contracts are available: Eastern contracts, Midwestern contracts, National contracts, Northeastern contracts, Southeastern contracts, Western contracts, California contracts, Florida contracts, and Texas contracts.

Strike prices are defined with respect to lower and upper strikes. So, if k_1 and k_2 are the lower and upper strike prices, respectively, they represent to the retention level and upper limit of a traditional excess-of floss catastrophe reinsurance contract. Selling a call-spread contract $k_1/k_2(k_1 < k_2)$ is equivalent to selling a call option with strike price k_1 and automatically buying a call option with a strike price k_2. By selling a call spread option contract, losses are capped by $k_2 - k_1$, with risk being reduced relative to a simple call. If $L(T)$ is the aggregate PCS index value, then the final payoff at the maturity of the option will be

$$\max\left(\min\left(L(T) - k_1, k_2 - k_1\right), 0\right)$$

If $[T, 0]$ is the loss period in months, then if $T = 12$, it is an annual contract, which covers the aggregate insured losses occurring within a whole calendar year; whereas if $T = 3$, it is a quarterly contract. There

are four quarterly contracts, such as March/June/September/December:

- March contracts cover losses occurring in the first quarter.
- June contracts cover losses in the second quarter.
- September contracts cover losses in the third quarter.
- December contracts covers losses in the fourth quarter.

Only annual contracts are available for the Western and California regions. Both annual and quarterly contracts are provided for the National index and there are two choices for development periods: six months or a year. During the development period, the PCS continues to update its nine regional index values for catastrophe events occurring during the loss period of each contract. The final payoff of a regional contract will be determined by the final, updated PCS index values on its expiration date.

A further benefit of the PCS call spread option contracts is that they had similar payoff functions to those of traditional excess-of-loss catastrophe reinsurance contracts, making the two reasonably good comparisons. Froot and O'Connell (1999) explain that limited capacity is the main factor that drives the reinsurance cycle problem. Exchange-based PCS options traded at the CBOT have no credit risks, so limited capacity would not be an issue. If the pricing cycle is not identified with PCS options, it would be a side support of their theory; if the cycle persists, then other reasons need to be explored. The following subsection briefly introduces some background knowledge of PCS options and the PCS index, and provides a summary of basic statistics on those transactions.

Interest rate data

Data for the construction of zero coupon yield curves were kindly provided by Credit Suisse First Boston for the period January 1970 to December 2003. The data provided and used was rates/prices for mid-market closing for cash market and swap market rates. Although interest rate futures were available for the entire period, their incorporation into the construction of a yield curve via a bootstrap process only became market practice in the mid-1980's. So for the sake of comparison and simplicity it was to be acceptable to exclude futures from the construction of the zero coupon yield curve. The impact on the forward rates produced by the bootstrap of omitting futures is minimal in the context of the reported research. The yield curve tenors used for the construction of the zero coupon yield curve were therefore:

Cash Curve:

the following maturity instruments:

O/N, 1W, 1M, 3M, 6M, 1Y

Swap Curve:

the following maturity instruments:

2Y, 3Y, 4Y, 5Y, 6Y, 7Y, 8Y, 9Y, 10Y, 11Y, 12Y, 13Y, 14Y, 15Y, 20Y, 25Y, 30Y

The Burr distribution

The cumulative density function for the Burr distribution, $Burr(\alpha, \lambda, \delta)$ with $\alpha > 0$, $\lambda > 0$, $\delta > 0$ is given by

$$cdf : F(x) = 1 - \left(\frac{\lambda}{\lambda + x^{\delta}}\right)\alpha$$

and the probability density function is given by

$$pdf : f(x) = \alpha\delta\lambda^{\delta-1}\left(\lambda + x^{\delta}\right)^{-\alpha-1}$$

Appendix 3: Catastrophe bond valuation results

Table 1: CAT Bond Pricing: Poisson PIDE

(short dated: maturity 12 months)

Cat-PIDE

Maturity		0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
7d		0.659697	0.787481	0.939768	1.094681	1.249111	1.403059	1.587164	1.892484	2.527527	2.767291	3.095048
1m		6.227245	7.353636	8.664816	9.965044	11.228771	12.457514	13.887839	16.170602	20.591231	22.155059	24.207401
2m		8.585207	10.092205	11.829336	13.534060	15.174116	16.753113	18.572009	21.433079	26.831864	28.698358	31.114567
3m		10.315650	12.086240	14.112720	16.086467	17.971536	19.773772	21.834605	25.043705	30.993198	33.018534	35.616782
4m		11.724384	13.699835	15.947786	18.124021	20.190407	22.155059	24.388625	27.839404	34.150462	36.273666	38.979000
5m		12.929334	15.073142	17.500737	19.838868	22.048121	24.138865	26.504335	30.135160	36.701952	38.890101	41.662996
6m		13.990725	16.277620	18.856131	21.328530	23.654720	25.847297	28.317689	32.088468	38.844252	41.077238	43.894065
7m		14.944188	17.355486	20.063822	22.650370	25.074776	27.351822	29.908013	33.790734	40.690173	42.954740	45.800240
8m		15.812848	18.334096	21.156051	23.841410	26.349871	28.698358	31.326127	35.300232	42.310919	44.597861	47.461648
9m		16.612708	19.232352	22.155059	24.927132	27.508565	29.918372	32.606766	36.656615	43.754509	46.057184	48.931914
10m		17.355372	20.063949	23.076923	25.925926	28.571429	31.034483	33.774834	37.888199	45.054945	47.368421	50.248756
11m		18.049570	20.839169	23.933702	26.851557	29.553828	32.063555	34.848858	39.015964	46.237231	48.557765	51.439737
12m		18.702060	21.565965	24.734701	27.714630	30.467580	33.018534	35.843026	40.055923	47.320316	49.645013	52.525626

Table 2: CAT Bond Pricing: Numerical HJI

(short dated: maturity 12 months)

Numerical HJI

Maturity		0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
7d		0.000475	0.613325	0.552955	0.325386	0.357017	0.376777	0.396924	0.218998	0.129831	0.009464	0.000717
1m		4.538871	20.649606	23.819885	24.610683	26.378369	27.441686	28.495086	28.072805	29.133104	23.036032	18.474267
2m		7.645458	24.584358	28.496233	30.106615	32.100966	33.290652	34.461937	34.775164	36.763999	32.715636	29.719156
3m		10.273936	27.108432	31.473698	33.625236	35.733505	36.984486	38.211290	39.019858	41.590127	39.239663	37.853081
4m		12.597721	28.994905	33.684705	36.239892	38.417212	39.704102	40.962477	42.140201	45.121047	44.130394	44.105660
5m		14.697949	30.510393	35.450925	38.326570	40.549529	41.859294	43.137108	44.605922	47.896010	48.009272	49.098953
6m		16.622333	31.780982	36.924342	40.064431	42.319022	43.643984	44.934172	46.641033	50.173786	51.198616	53.196724
7m		18.402490	32.876935	38.189535	41.553722	43.830832	45.166064	46.464150	48.370738	52.099435	53.888614	56.630204
8m		20.060977	33.841654	39.298664	42.856499	45.149851	46.492005	47.794983	49.872390	53.762685	56.201190	59.555006
9m		21.614726	34.703924	40.286273	44.013969	46.319054	47.665761	48.971527	51.197220	55.222977	58.219227	62.080400
10m		23.076923	35.483871	41.176471	45.054945	47.368421	48.717949	50.024988	52.380952	56.521739	60.001600	64.285714
11m		24.458125	36.196140	41.986773	46.000392	48.319722	49.670774	50.977965	53.449506	57.688999	61.591590	66.230115
12m		25.766974	36.851739	42.730325	46.866056	49.189274	50.540862	51.847362	54.422286	58.747213	63.021910	67.958735

Table 3: CAT Bond Pricing: Linear Lyapunov Solution

(short dated: maturity 12 months)

Linear Lyapunov

Maturity		0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
7d		0.000000	0.000000	0.000001	0.001032	0.080934	0.357017	0.201116	0.416272	0.475457	0.632940	0.750725
1m		0.417257	1.249111	2.682584	9.352678	20.391324	26.378369	26.381896	29.478761	32.328575	38.911581	43.065478
2m		1.343068	3.454130	6.355969	15.228432	26.591659	32.100966	32.865012	35.549188	38.664104	45.666611	49.951968
3m		2.640604	6.167313	10.312250	19.902108	30.731386	35.733505	37.008844	39.345816	42.573654	49.710366	53.998065
4m		4.235697	9.189388	14.319354	23.825357	33.875072	38.417212	40.074053	42.123001	45.408107	52.584877	56.840280
5m		6.071102	12.389934	18.258541	27.214569	36.417366	40.549529	42.507697	44.312998	47.628348	54.803756	59.015230
6m		8.098405	15.676086	22.064522	30.197596	38.553172	42.319022	44.524023	46.119318	49.449805	56.603042	60.766881
7m		10.275374	18.980109	25.702015	32.858638	40.394397	43.830832	46.243258	47.654727	50.991160	58.111047	62.226753
8m		12.564876	22.252603	29.153851	35.257023	42.011708	45.149851	47.739934	48.988457	52.324926	59.405307	63.473778
9m		14.934334	25.458022	32.414104	37.436543	43.452778	46.319054	49.063562	50.166148	53.498707	60.536240	64.558988
10m		17.355372	28.571429	35.483871	39.430648	44.751381	47.368421	50.248756	51.219512	54.545455	61.538462	65.517241
11m		19.803512	31.576018	38.368560	41.265573	45.932361	48.319722	51.320691	52.171489	55.488930	62.436746	66.373378
12m		22.257880	34.461197	41.076133	42.962326	47.014549	49.189274	52.298255	53.039218	56.346832	63.249427	67.145706

Table 4: CAT Bond Pricing: non-Linear Lyapunov Solution

(short dated: maturity 12 months)

Non-Linear Lyapunov

Maturity		0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
7d		0.000000	0.000000	0.000014	0.001032	0.005110	0.022602	0.100898	0.060106	0.140801	0.318226	0.711494
1m		0.002000	0.561841	4.331552	9.352678	13.913050	18.438574	24.207401	25.273087	30.838022	36.212343	41.745070
2m		0.031990	1.996220	8.846680	15.228432	20.794988	25.520854	31.114567	33.117638	38.671797	43.677932	48.600653
3m		0.161738	4.134260	13.164531	19.902108	25.855319	30.412226	35.616782	38.228436	43.575788	48.206167	52.652446
4m		0.509392	6.840651	17.221087	23.825357	29.898738	34.183332	38.979000	42.027652	47.139236	51.441420	55.509299
5m		1.234568	9.987511	21.005717	27.214569	33.271654	37.256026	41.662996	45.043596	49.925403	53.943975	57.701408
6m		2.526513	13.454949	24.526506	30.197596	36.162942	39.846720	43.894065	47.536184	52.202995	55.974620	59.470648
7m		4.581974	17.134294	27.798849	32.858638	38.689248	42.083039	45.800240	49.653764	54.121930	57.676319	60.947756
8m		7.571724	20.930723	30.840710	35.257023	40.928510	44.047201	47.461648	51.489374	55.774567	59.135967	62.211364
9m		11.599866	24.764581	33.670437	37.436543	42.935741	45.795532	48.931914	53.105330	57.221877	60.410318	63.312402
10m		16.666667	28.571429	36.305732	39.430648	44.751381	47.368421	50.248756	54.545455	58.506224	61.538462	64.285714
11m		22.649711	32.301138	38.763182	41.265573	46.406111	48.795873	51.439737	55.841753	59.658232	62.548447	65.156166
12m		29.314635	35.916368	41.058073	42.962326	47.923796	50.100810	52.525626	57.018314	60.700747	63.461069	65.942105

Table 5: CAT Bond Pricing: Poisson PIDE

(long-dated: maturity 10 years)

Cox-PIDE

Maturity	0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
0.5y	0.280476	0.267227	0.258148	0.251107	0.264929	0.279129	0.292051	0.325394	0.460364	0.585435	0.817947
1y	0.990099	0.990099	0.990099	0.990099	1.088032	1.185771	1.283317	1.477833	2.152642	2.818270	4.761905
1.5y	2.056936	2.114901	2.157280	2.191780	2.463469	2.734829	3.014709	3.530860	5.193568	6.863238	12.548064
2y	3.433304	3.597855	3.720013	3.820602	4.356710	4.892802	5.455020	6.447796	9.473607	12.496601	23.262469
2.5y	5.076939	5.395102	5.633776	5.831829	6.714070	7.595371	8.527198	10.120876	14.741029	19.263312	35.128939
3y	6.947967	7.461924	7.850202	8.174040	9.468693	10.757417	12.123470	14.396291	20.671604	26.626592	46.522117
3.5y	9.008365	9.753483	10.318904	10.791940	12.546166	14.282201	16.119472	19.098439	26.933559	34.094077	56.499306
4y	11.222112	12.225808	12.989601	13.629091	15.869681	18.069728	20.387112	24.051512	33.235553	41.290177	64.762892
4.5y	13.555520	14.837251	15.812852	16.630334	19.364440	22.023695	24.805000	29.095821	39.352028	47.974074	71.399547
5y	15.977558	17.548552	18.743462	19.743733	22.961042	26.056597	29.265700	34.097739	45.128564	54.022482	76.652750
5.5y	18.460095	20.324278	21.739512	22.921987	26.597780	30.092907	33.679772	38.953689	50.474245	59.398433	80.792957
6y	20.978128	23.132810	24.763993	26.123319	30.221906	34.070508	37.977036	43.589617	55.348113	64.120254	84.061919
6.5y	23.509316	25.946597	27.784910	29.311896	33.790024	37.940745	42.105811	47.957630	59.744778	68.237015	86.656855
7y	26.034915	28.742138	30.775298	32.457862	37.267823	41.667515	46.030891	52.031313	63.681920	71.811568	88.731804
7.5y	28.538645	31.499822	33.713022	35.537088	40.629361	45.225782	49.730892	55.800801	67.190615	74.909990	90.404637
8y	31.006996	34.203667	36.580428	38.530724	43.856085	48.599831	53.195474	59.268261	70.308474	77.595648	91.764835
8.5y	33.428914	36.840989	39.363903	41.424634	46.935736	51.781503	56.422728	62.444082	73.075097	79.926272	92.880259
9y	35.795544	39.402044	42.053385	44.208776	49.861260	54.768538	59.416899	65.343870	75.529269	81.952861	93.802535
9.5y	38.099993	41.879652	44.641874	46.876585	52.629776	57.563117	62.186493	67.986182	77.707366	83.719607	94.571157
10y	40.337083	44.268837	47.124944	49.424379	55.241656	60.170637	64.742784	70.390897	79.642584	85.264345	95.216537

Table 6: CAT Bond Pricing: Numerical HJI

(long-dated: maturity 10 years)

Numerical HJI

Maturity	0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
0.5y	0.414364	0.497512	0.648678	0.938433	1.308477	1.496807	2.949467	4.359904	7.061263	9.614215	13.191070
1y	1.477833	1.960784	2.912621	4.761905	6.976744	9.090909	16.666667	23.076923	33.333333	41.176471	50.000000
1.5y	3.078095	4.306220	6.820883	11.684876	17.127692	23.140110	37.583384	47.457087	60.085317	67.819596	75.066618
2y	5.130093	7.407407	12.114475	20.879915	29.788301	39.689592	56.825410	66.378257	76.692383	82.163943	86.808930
2.5y	7.553862	11.111111	18.381394	31.074589	42.566986	54.689956	70.709124	78.359890	85.785522	89.416996	92.348985
3y	10.272818	15.254237	25.169156	41.118646	53.898659	66.457575	79.849265	85.598865	90.831152	93.274640	95.195308
3.5y	13.214858	19.678715	32.071686	50.272469	63.219390	75.078020	85.765215	90.037426	93.774362	95.472584	96.787172
4y	16.313830	24.242424	38.776652	58.209370	70.588235	81.241111	89.649760	92.853272	95.585774	96.806708	97.743074
4.5y	19.510689	28.825623	45.076318	64.886490	76.313170	85.642647	92.266134	94.707655	96.755918	97.661120	98.351215
5y	22.754162	33.333333	50.855054	70.410414	80.741927	88.818365	94.078100	95.972566	97.543974	98.233298	98.756719
5.5y	26.000913	37.694704	56.067154	74.944663	84.178943	91.144445	95.367088	96.862946	98.093849	98.631008	99.037754
6y	29.215306	41.860465	60.714163	78.658975	86.865647	92.876835	96.306884	97.507236	98.489278	98.916235	99.238890
6.5y	32.368874	45.799458	64.826167	81.706578	88.985279	94.188726	97.007409	97.984838	98.781077	99.126298	99.386801
7y	35.439602	49.494949	68.448153	84.216850	90.674370	95.198094	97.539983	98.346430	99.001252	99.284567	99.498120
7.5y	38.411116	52.941176	71.630955	86.295284	92.034108	95.986289	97.952045	98.625312	99.170627	99.406184	99.583587
8y	41.271857	56.140351	74.425702	88.026346	93.139619	96.610232	98.275894	98.843952	99.303149	99.501256	99.650356
8.5y	44.014276	59.100204	76.880710	89.477078	94.046994	97.110365	98.534001	99.017868	99.408397	99.576711	99.703321
9y	46.634108	61.832061	79.039967	90.700538	94.798440	97.515865	98.742311	99.158011	99.493099	99.637403	99.745906
9.5y	49.129707	64.349376	80.942589	91.738743	95.425997	97.848090	98.912342	99.272257	99.562079	99.686808	99.780559
10y	51.501482	66.666667	82.622838	92.625060	95.954215	98.122895	99.052555	99.366369	99.618855	99.727457	99.809064

Table 7: CAT Bond Pricing: Linear Lyapunov Solution

(long-dated: maturity 10 years)

Linear Lyapunov

Maturity	0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
0.5y	0.304203	0.320502	0.361035	0.371028	0.384887	0.394973	0.403925	0.411809	0.408263	0.398440	0.327154
1y	1.088032	1.185771	1.380671	1.477833	1.574803	1.671583	1.768173	1.864573	1.912702	1.931941	1.951172
1.5y	2.275951	2.527191	2.994353	3.277642	3.543548	3.830690	4.126367	4.430704	4.630719	4.767193	5.402209
2y	3.814248	4.286581	5.131659	5.697505	6.213976	6.792983	7.397890	8.029154	8.488443	8.843475	10.766156
2.5y	5.653378	6.405011	7.713463	8.643509	9.476705	10.430026	11.431957	12.482876	13.288510	13.947372	17.731003
3y	7.745584	8.822642	10.655686	12.009885	13.204336	14.585897	16.038489	17.561128	18.763317	19.776401	25.718648
3.5y	10.044874	11.480235	13.873816	15.687691	17.264413	19.096690	21.016301	23.019016	24.626658	26.003593	34.077290
4y	12.507600	14.320895	17.286972	19.572141	21.530101	23.806386	26.175812	28.629056	30.615045	32.329732	42.246311
4.5y	15.093111	17.291582	20.821033	23.568038	25.887809	28.577783	31.354084	34.201084	36.513003	38.514688	49.839143
5y	17.764283	20.344239	24.410684	27.593120	30.241612	33.298339	36.422058	39.590603	42.162516	44.388010	56.647980
5.5y	20.487879	23.436474	28.000452	31.579563	34.514867	37.881736	41.285506	44.698243	47.460746	49.844600	62.605032
6y	23.234716	26.531859	31.544882	35.473992	38.649733	42.266396	45.881834	49.463884	52.351214	54.832620	67.733232
6.5y	25.979660	29.599883	35.008088	39.236480	42.605330	46.412243	50.174778	53.858525	56.812716	59.339265	72.104481
7y	28.701493	32.615677	38.362899	42.838933	46.355208	50.296731	54.148523	57.876022	60.848705	63.377746	75.810733
7.5y	31.382667	35.559548	41.589764	46.263230	49.884616	53.910855	57.802181	61.525839	64.478438	66.976896	78.946645
8y	34.009000	38.416434	44.675590	49.499327	53.187885	57.255552	61.145106	64.827266	67.730250	70.173613	81.600508
8.5y	36.569330	41.175309	47.612594	52.543500	56.266110	60.338681	64.193183	67.805098	70.636779	73.007744	83.850366
9y	39.055152	43.828581	50.397241	55.396802	59.125208	63.172627	66.966036	70.486591	73.231787	75.518869	85.763060
9.5y	41.460267	46.371534	53.029312	58.063762	61.774354	65.772465	69.485028	72.899422	75.548176	77.744428	87.394764
10y	43.780440	48.801806	55.511093	60.551327	64.224763	68.154623	71.771886	75.070408	77.616842	79.718767	88.792169

Table 8: CAT Bond Pricing: non-Linear Lyapunov Solution

(long-dated: maturity 10 years)

Non-Linear Lyapunov

Maturity		0% Loss	10% Loss	20% Loss	30% Loss	40% Loss	50% Loss	60% Loss	70% Loss	80% Loss	90% Loss	100% Loss
0.5y		0 342787	0 293872	0 335333	0 361191	0 363316	0 364879	0 361442	0 357817	0 385322	0 386352	0 317321
1y		1 224812	1 088032	1 283317	1 429276	1 487538	1 545732	1 584490	1 623217	1 806756	1 874203	1 893456
1.5y		2 558208	2 321481	2 786430	3 159424	3 350903	3 548149	3 706952	3 870280	4 380943	4 628720	5 247876
2y		4 278926	3 943453	4 782643	5 482042	5 885010	6 306278	6 669080	7 047137	8 048145	8 597287	10 475553
2.5y		6 327373	5 902766	7 202183	8 309456	8 991571	9 709881	10 350146	11 021417	12 633604	13 580267	17 288826
3y		8 646152	8 147322	9 970452	11 543565	12 554843	13 622449	14 592583	15 611177	17 894274	19 290264	25 138142
3.5y		11 180378	10 625199	13 011774	15 083324	16 453070	17 897722	19 224385	20 614576	23 564824	25 412872	33 392931
4y		13 878547	13 286044	16 252876	18 831169	20 568127	22 393570	24 078088	25 835285	29 395707	31 656782	41 501177
4.5y		16 693382	16 082358	19 625693	22 697845	24 792647	26 982078	29 004416	31 100509	35 177856	37 784882	49 073882
5y		19 582450	18 970503	23 069305	26 605505	29 034423	31 555507	33 880043	36 270762	40 753395	43 625877	55 894670
5.5y		22 508507	21 911373	26 531050	30 489182	33 218327	36 028588	38 610190	41 242510	46 015432	49 071448	61 885588
6y		25 439576	24 870760	29 966891	34 296978	37 286308	40 338040	43 127449	45 945818	50 901257	54 065501	67 060539
6.5y		28 348809	27 819437	33 341234	37 989303	41 196064	44 440335	47 388346	50 339242	55 382816	58 590884	71 484554
7y		31 214182	30 733019	36 626329	41 537562	44 918981	48 308564	51 368928	54 403746	59 457072	62 656917	75 244893
7.5y		34 018089	33 591673	39 801447	44 922547	48 437775	51 929086	55 060254	58 136801	63 137665	66 289229	78 433292
8y		36 746861	36 379714	42 851930	48 132780	51 744158	55 298350	58 464307	61 547286	66 448349	69 522222	81 136404
8 5y		39 390275	39 085152	45 768225	51 162933	54 836706	58 420130	61 590581	64 651393	69 418205	72 393899	83 431489
9y		41 941062	41 699214	48 544961	54 012417	57 719031	61 303245	64 453374	67 469536	72 078337	74 942530	85 385146
9.5y		44 394440	44 215871	51 180108	56 684173	60 398270	63 959764	67 069748	70 024128	74 459728	77 204669	87 053625
10y		46 747689	46 631402	53 674240	59 183673	62 883890	66 403640	69 458053	72 338047	76 591947	79 214089	88 483864

Appendix 4: Catastrophe option valuation results

Table 1: PCS option valuation results using Cox-TGA

(3m loss period contract, 6m development period)

Contract	5pts	10pts	20pts	50pts	100pts	200pts	250pts	300pts	350pts	400pts	450pts	500pts
Sep-95	12 759734	13 005635	15 034362	19 761823	24 481331	17.23834	12 642529	9.9622424	7.0133886	5.835347	5 6439259	4 9166117
Dec-95	12 919789	14 872636	15.058629	20 00971	28 813507	18 429487	12 801114	10.087206	7.1013629	5.9085443	5.7147221	4.9782846
Mar-96	11 140505	12 161225	13.375253	18 613303	20.983053	14.325566	11 677364	10.790374	7.5720809	6.5919005	6.2986626	5.2952994
Jun-96	11 979727	12.700317	13 105258	17.887413	18.250186	13 07257	11 850827	10 927223	7.6681137	6.3716974	5.3872312	5.3624569
Sep-96	15 113283	16.734279	19 495252	20.777879	27.401795	16 021821	13.292547	10.474452	7.3739828	7.135372	6.2692836	5.1693999
Dec-96	13 111073	15 045034	17.965363	19 860944	23.272455	15.288865	12 755021	10.039067	7 0901626	6 9070027	5.1501569	4.9826105
Mar-97	10 392248	11 873527	12.755789	13 94063	18 945616	14 19912	10.378618	8 5077062	7 9603098	6.3157497	5.9882543	5.5667952

Table 2: PCS option valuation results using numerical HJI

(3m loss period contract, 6m development period)

Contract	5pts	10pts	20pts	50pts	100pts	200pts	250pts	300pts	350pts	400pts	450pts	500pts
Sep-95	60 08926	82 662325	118 62135	127 62615	141.1708	106 15163	97.653002	53 988044	40 001383	25.029836	22.76392	21.387635
Dec-95	59 750765	89 988597	120.28177	128.75864	143.31736	108 97506	95 448412	51 003347	44 172202	30 263305	27.471019	20 631058
Mar-96	39 518473	49 614336	74.505884	82.321002	102.03369	85.394987	75.834649	49.652417	44.707996	30.632022	26.272246	18 863216
Jun-96	56 117282	54 588588	40 816412	40 647773	37.319478	53 425545	39 050083	56 059484	44.457012	29 053966	24 028383	18.041646
Sep-96	61 049925	75 253995	81 689423	89 989626	103.28697	109 62851	79.593711	52.749844	43.074608	27.489903	20.697909	14.270286
Dec-96	54 657936	62 339192	71 634867	76.340284	78 897485	91 263656	71 267213	53.615053	45.122566	23.816695	18.211623	17.476067
Mar-97	38 098051	47.973336	49.150188	55 883999	58.996413	44.775189	35.015866	30.634454	29.938189	26.425263	21.817499	16.71575

Table 3: PCS option valuation results using linear Lyapunov

(3m loss period contract, 6m development period)

Contract	5pts	10pts	20pts	50pts	100pts	200pts	250pts	300pts	350pts	400pts	450pts	500pts
Sep-95	46 265258	57 859636	65.260674	69.522682	83.585109	56.610196	41.393924	37.392915	29.346204	28.433185	26.428903	19.114815
Dec-95	56 021855	63.73714	72.372786	78 787363	92.123568	80.711297	50.540137	36.369698	32.18932	29.885323	26.605219	23 627548
Mar-96	37 107306	53 15488	59 833843	64 043447	70.691383	56.957811	46.77511	32.42811	28.563818	26.367426	23 091482	20.65376
Jun-96	39 973282	51 232273	60.296383	62.345903	69.689602	52.176712	38.569765	30.849863	26.897201	25.35128	22.003945	19.841711
Sep-96	53 025567	60.778423	67.461568	71.737844	85.615233	52.909087	43.482169	31 645241	25.36732	23.623744	20.559724	17 89929
Dec-96	58 878285	58 921771	70 633878	74 845979	72.562484	58.273571	53.182441	33.862103	24.404242	23.284989	21.428625	18.232174
Mar-97	35.507477	44 610899	51.442141	62 23785	67.863077	59.476452	44.888391	29 077281	27.241003	26 879348	23 628765	19.578109

Table 4: PCS option valuation results using non-linear Lyapunov

(3m loss period contract, 6m development period)

Contract	5pts	10pts	20pts	50pts	100pts	200pts	250pts	300pts	350pts	400pts	450pts	500pts
Sep-95	36 382815	45 514879	51 539098	54 942828	66.496052	45.175637	33.161476	30 091157	23.635504	23.090445	21.52243	15.587782
Dec-95	44 056954	50.13707	57.157307	62 264261	73.290538	64.407578	40.486997	29.267726	25.923244	24.270349	21 66738	19.268977
Mar-96	29 181405	41 811519	47.255379	50 610988	56 240405	45 452073	37.472459	26 097398	23 003831	21 413011	18 804543	16.842882
Jun-96	31 436092	40.298853	47.621835	49.271124	55 442804	41.635023	30.898518	24.826594	21 661482	20.587314	17.918696	16.180864
Sep-96	41 700484	47 810625	53 277286	56 692639	68 113766	42.221488	34 834369	25.466549	20.431069	19.184577	16.743106	14 597164
Dec-96	46.304396	46 34737	55.786558	59.149475	57.729908	46.501285	42.603268	27 249012	19 655378	18.90958	17.45163	14.868895
Mar-97	27.922349	35 091859	40 62797	49.184933	53 992511	47.464576	35 961177	23.399894	21.939566	21.828402	19.243104	15.9657

Appendix 5: Hedging simulation data

Explanation of hedging strategies

Table 1 (which is identical to table 3.1 in chapter 3) summarises the strategies that were evaluated as the basis of the empirical work reported in chapter 3 of this thesis. The object of this appendix is to provide precise meaning of each of the strategies - which for ease of exposition are referred to.by the reference number in the first column of table 1 below.

Table 1: Strategy summaries

Number	Strategy	Objective Function	Conditions	Costs[2]
1	Delta	Delta neutrality	n.a.	Excluded
2	Minimax-95	Potential hedge error	95% level	Excluded
3	Minimax-99	Potential hedge error	99% level	Excluded
4	Heuristic-w	Potential hedge error	Weighted	Excluded
5	Minimax-95c	Potential hedge error	95% level	Included
6	Heuristic-95c	Potential hedge error	95% level	Included
7	ρ^2	Potential hedge error	n.a.	Included
8	$\theta/(r_s - i)/\sigma_s$	Potential hedge error	n.a.	Included
9	$r_h^{ce} - r_s^{ce}$	Potential hedge error	n.a.	Included
10	Linear $H\infty$	Potential hedge error	Uncertain	Excluded
11	Non-Linear $H\infty$	Potential hedge error	Uncertain	Excluded

Strategy 1:

The objective of this strategy was to be precisely delta neutral. This was achieved by first calculating the delta of the position and then adding or subtracting exactly the right amount of futures contracts to produce an exactly zero delta position - both in total and by forward time tenor. This was repeated for each daily hedge re-balancing. Making this work proved to be an iterative procedure for which code had to be developed tested and then implemented.

Strategy 2:

This strategy follows very closely the Rustem, Howe and *Selby work* referred to in chapter 3. The objective of this strategy was to minimise the minimax hedging error. In this case a distinction is drawn between actual and potential hedging error. Actual error is defined to be inclusive of interest payments on borrowed money. It is calculated when actual B_t, y_t^S, B_{t+1} and y_{t+1}^S are used in the equation for U_1. Potential hedging error, including interest payments on borrowed money, is calculated when *actual* values

[2]Transaction costs can be modelled in a number of different ways. However, it was felt that adjusting the bid-ask spread was the most appropriate method in the interests of transparency and in line with most capital-market conventions.

of B_t and y_t^S and potential values of B_{t+1} and y_{t+1}^S are used in the equation for U_1. Note that potential y_{t+1}^S is taken from a predefined range that maximises the objective function. Potential B_{t+1} is the value of the call option based on the pricing model, given potential y_{t+1}^S , i.e. potential $B_{t+1} = B_{t+1}(y_{t+1}^S)$. The minimax hedging error at time t is defined as

$$\text{minimax hedging error} = U_1\left(x_t^*, y_{t+1}^{S^*}\right)$$

The minimax hedging error is the worst-case potential hedging error, including interest payments on borrowed money, given the solution of x_t^* and $y_{t+1}^{S^*}$. The limit is set with respect to the 95% confidence level.

Strategy 3:

Identical to strategy 2, except that the limit is set at 99% confidence level.

Strategy 4:

The objective of this strategy was to minimise potential hedging error based on weighting the number of shares to hold by a weighting factor in the range 0 to 1. The weight represented the hedger's view of the information contained in the price changes in the underlying stock. The weights were algorithmically generated based on the standard deviation of the underlying, based on the changes in volatility.

Strategy 5:

Identical to strategy 2, but including transactions costs.

Strategy 6:

Identical to strategy 4 but incorporating a 95% confidence limit.

Strategy 7:

The objective of this strategy was to assess the performance of the Ederington hedge efficiency measure, calculated as follows

$$e = \frac{\sigma_{SF}^2}{\sigma_s^2 \sigma_f^2}$$

where σ_s^2 and σ_f^2 are the subjective variances of the possible price changes of the spot and forward prices and σ_{SF}^2 is the variance of the spot-futures portfolio.

Strategy 8:

The objective of this strategy was to assess the performance of the Howard and D'Antonio hedge effectiveness ratio which is calculated as follows

$$\frac{\theta}{\left[\frac{r_s - i}{\sigma_s}\right]}$$

Strategy 9:

The objective of this strategy was to assess the performance of the Pennings and Meulenberg hedging efficiency measure

$$E = \frac{\sqrt{E\left(FTR_{t+1}\right)^2}\overline{CP}}{\left(PF_t^1 - C\right)\sqrt{E\left(CP_t - \overline{CP}\right)^2}} = \frac{\left[\sqrt{\sigma_A^2 + \mu_A^2}\right]\overline{CP}}{\left(PF_t^1 - C\right)\sqrt{E\left(CP_t - \overline{CP}\right)^2}}$$

where the terms are explained fully in text and footnotes in chapter 3.

Strategy 10:

The objective of this strategy was to assess the performance of the linear H∞ hedging rule, based on the following formulation

$$\dot{x} = f(x, u, w) \tag{7.1}$$

$$J = \int_0^\infty \left[x^2 + u^2\right] dt \tag{7.2}$$

where J is the (quadratic) total cost function, x is the state variable (the value of the portfolio), u the linear controller or hedge rule and w the disturbance variable.

Strategy 11:

The objective of this strategy was to assess the performance of the non-linear H∞ hedging rule, based on the following formulation

$$\dot{x} = f(x, u, w) \tag{7.3}$$

$$J = \int_0^\infty \left[x^2 + u^2\right] dt \tag{7.4}$$

where J is the (quadratic) total cost function, x is the state variable (the value of the portfolio), u the non-linear controller or hedge rule and w the disturbance variable.

Volatility regimes

Two volatility environments were used in the generation of the Monte Carlo simulation data set for the simulations reported in chapter three. In an attempt to make the results comparable with Howe Rustem and Selby, the low volatility environment was deemed to be a flat volatility of 20% whilst the high volatility environment was set at a flat rate of 60%.

Generating the simulation data

The simulations were performed using the simulation engine inside Robusta, which uses a widely available and tested version of a C++ implementation of the Mersenne Twister random number generator.

M⋊⋉≈C∂⋌⋖⋊~⊐⋗≅⋖∂≈⊐⋊⋉∂≈∂

For each of the eleven strategies, 100,000 paths were calculated for each combination of low and high volatilities versus 11 alternative strike levels (0.5% - 5.0% away from the at-the-money strike in 0.5% increments), which meant a total of 22 (i.e. 11 strikes times 2 volatilities = 22) paths for each strategy. Each strategy was evaluated over each of the possible paths, with the final result for the strategy being the simple arithmetic average of the all the path evalutions. The standard normal distribution was used to simulate the paths on the grounds that it provided an easy to interpret benchmark against which to compare the results of the historical simulation.

Historical simulation data

For each of the eleven strategies, 100,000 paths were calculated for each combination of low and high volatilities versus 11 alternative strike levels (0.5% - 5.0% away from the at-the-money strike in 0.5% increments), which meant a total of 22 (i.e. 11 strikes times 2 volatilities = 22) paths for each strategy. Each strategy was evaluated over each of the possible paths, with the final result for the strategy being the simple arithmetic average of the all the path evalutions. The standard normal distribution was used to simulate the paths on the grounds that it provided an easy to interpret benchmark against which to compare the results of the historical simulation.

Appendix 6: Portfolio risk simulation data sources and results

Data sources and market data calculations

The period selected for the portfolio risk research reported in chapter 4 was the three month period covered by the second quarter of 2004 (i.e. from Monday 01 March 2004 to Friday 02 July 2004). Data for interest rate, equity price and volatility data were supplied by Credit Suisse First Boston and cross checked with the exchange where possible. Contract specifications were downloaded from the Eurex web site. Daily data on both spot index values, as well as call option prices was purchased from Eurex. Contextual data on market depth (e.g. open interest to ensure deep and liquid pricing) was downloaded from the Dow Jones EuroStoxx50 index web page (**http://www.eurexchange.com/data/**). Where available, volatilities from Eurex were used. When thought to be unreliable, historical volatilities were calculated using the daily data using the following industry standard definition. Let a be a stochastic process, where its terms may represent prices, accumulated values, exchange rates, interest rates, etc. Then the volatility of the process at time $t-1$ is defined as the standard deviation of the time t return. Typically, log returns are used, so the definition becomes

$$\text{volatility} = std\left(\log\left(\frac{Q_t}{Q_{t-1}}\right)\right)$$

where *log* denotes a natural logarithm. However, simple returns are often used; which can often be true in the context of portfolio theory. Assuming that returns are conditionally homoskedastic, then the above definition is precise. However, if they are conditionally heteroskedastic, then the definition needs to be clarified. Does volatility at time $t-1$ represent the unconditional standard deviation of the time t log return? Or does it represent the standard deviation of the time t log return conditional on information available at time $t-1$? The answer is generally accepted to be the latter. To emphasize this, the definition is expressed as follows

$$\text{volatility} = std^{t-1}\left(\log\left(\frac{Q_t}{Q_{t-1}}\right)\right)$$

where the superscript $t-1$ on the standard deviation operator indicates that the standard deviation is conditional on information available at time $t-1$. Another issue in defining volatility is that of the unit of time on which the calculation and therefore the result is based. The standard deviation of a stock's price return over a day might be 0.01 and over a year, it might be 0.16. Accordingly, the volatility that was used in the portfolio risk computations was appropriate to the period for which the option was being valued.

Generating the simulation data for evaluating the risk measures

The simulations were performed using the simulation engine inside Robusta, which uses a widely

available and tested version of a C++ implementation of the Mersenne Twister random number generator. The method for generating the volatilities used in the valuations for each strategy, was therefore as follows:

1. For each given volatility environment, 10,000 option price paths and 10,000 index value paths were generated.

2. The simulations were run using: low, flat, skew, 10% jump, 30% jump, 50% jump, 100% jump, 500% jump and high volatility (90% flat) environments. Results were generated for a number of alternative distributions (normal, lognormal, Frechet and a Generalised Pareto) but only those for the extreme value Generalised Pareto distribution are reported on the grounds that the remainder of the results added little or nothing of any interest or significance to the reported results.

3. The risk strategies were then evaluated using the generated path information using the methodologies described in chapter 4, with the reported EVT measure being based on the Generalised Pareto distribution using a peaks over threshold method.

Description of the risk management strategies

Simulation Scenario	Description
Delta Neutrality	Simple delta hedge
10d Delta-Gamma VAR-95%	10d horizon VAR calculated using delta-gamma approximation - 95%
10d Delta-Gamma VAR-99%	10d horizon VAR calculated using delta-gamma approximation - 99%
Minimax	Minimisation of potential hedging error
EVT (GPD)	Peaks over threshold model
Rational Bounds	Assumed rebalancing
PWMN	No flattening
PWMN flattened	Flattened using distance to region around manifold

The above risk management strategies involved the following:

Delta Neutrality:

This strategy uses the same definition as used in chapter 3.

10d Delta-Gamma VAR-95%:

This strategy uses the industry standard delta-gamma method based on a 10 day window at the 95% confidence level (see for example Jorion, 1997).

10d Delta-Gamma VAR-99%:

This strategy uses the industry standard delta-gamma method based on a 10 day window at the 95% confidence level (see for example Jorion, 1997).

Minimax:

This strategy uses the industry standard delta-gamma method based on a 10 day window at the 99% confidence level (see for example Jorion, 1997). This strategy follows very closely the Rustem, Howe and Selby work referred to in chapter 3. The objective of this strategy was to minimise the minimax hedging error. In this case a distinction is drawn between actual and potential hedging error. Actual error is defined to be inclusive of interest payments on borrowed money. It is calculated when actual B_t, y_t^S, B_{t+1} and y_{t+1}^S are used in the equation for U_1. Potential hedging error, including interest payments on borrowed money, is calculated when *actual* values of B_t and y_t^S and potential values of B_{t+1} and y_{t+1}^S are used in the equation for U_1. Note that potential y_{t+1}^S is taken from a predefined range that maximises the objective function. Potential B_{t+1} is the value of the call option based on the pricing model, given potential y_{t+1}^S , i.e. potential $B_{t+1} = B_{t+1}(y_{t+1}^S)$. The minimax hedging error at time t is defined as

$$\text{minimax hedging error} = U_1\left(x_t^*, y_{t+1}^{S^*}\right)$$

The minimax hedging error is the worst-case potential hedging error, including interest payments on borrowed money, given the solution of x_t^* and $y_{t+1}^{S^*}$. The limit is set with respect to the 95% confidence level.

EVT using the Generalised Pareto distribution:

This strategy uses the standard peaks-over-threshold methodology (see Embrechts, Kluppelberg and Mikosch (1999).

Rational bounds:

This strategy uses the a simple heuristic of keeping portfolio risk within 95% of the previous day's risk as measured by the simple delta of the portfolio.

PWMN (piece wise min-norm):

This strategy uses a piece wise min-norm control rule to control portfolio risk. To see how the algorithm actually functions, the interested reader is referred to Kokotovic (1996) and the Slicot library documentation.

PWMN (piece wise min-norm) flattened:

This strategy uses a piece wise min-norm control rule modified such that local gains on the controller grow much more slowly, thereby reducing the control effort required to robustly maintain portfolio risk within the guaranteed bounds. This is achieved by using a penalty term that penalises the distance to

a region around the manifold rather than the distance to the manifold itself. The interested reader is referred to Kokotovic (1996) for explicit details of the algorithm.

Simulation results

Table 1: Summary Monte Carlo simulations of portfolio level risk

(expressed as a % based on ATMF strike premium)

Volatility Environment	10d Delta-Gamma VAR-95%	Minimax	EVT (GPD)	Rational Bounds	PWMN	PWMN flattened
Low vol: Min risk	4.7831	5.7183	6.1529	7.6136	8.7682	8.6358
Low vol: Max risk	12.9700	10.1373	17.7299	13.4837	14.0902	11.5733
Flat vol: Min risk	14.8462	6.6902	21.8704	10.1256	9.7047	8.7471
Flat vol: Max risk	36.5839	23.6490	41.5538	26.5349	19.4863	11.8452
High vol: Min risk	22.3887	16.8171	20.1649	12.4383	20.2752	18.1792
High vol: Max risk	51.9761	43.1728	79.1042	42.0361	31.9215	27.9094
Skew vol: Min risk	20.1382	19.1410	56.5888	13.4471	15.6188	13.4471
Skew vol: Max risk	34.0353	31.6734	72.8776	30.0539	20.9428	25.0403
Jump10%: Min risk	24.0309	33.5126	45.7869	30.0160	20.3220	16.5214
Jump10%: Max risk	37.0203	42.5999	55.6970	34.2276	25.8716	22.4910
Jump30%: Min risk	27.9093	26.3343	32.5102	24.2738	22.6814	22.0460
Jump30%: Max risk	46.1536	40.1775	50.3042	39.5581	29.1972	25.3424
Jump50%: Min risk	35.2842	34.1334	51.2126	33.9755	24.8864	25.8352
Jump50%: Max risk	55.5474	49.6453	67.6706	41.5597	29.2593	28.9724
Jump100%: Min risk	37.4579	31.4965	64.4147	37.4651	40.3429	28.9964
Jump100%: Max risk	79.2800	59.8575	98.9560	50.0443	46.7538	35.0948
Jump500%: Min risk	68.0541	58.4148	101.4690	50.3362	49.4917	29.6558
Jump500%: Max risk	129.7709	98.6785	169.4591	73.2804	63.0695	38.9674

Table 2: Historical simulations of portfolio level risk

(expressed as a % based on ATMF strike premium)

Volatility Environment	10d Delta-Gamma VAR-95%	Minimax	EVT (Frechet)	Rational Bounds	PWMN	PWMN flattened
Low vol: Min risk	8.901337502	6.6872754	9.5294096	8.8710034	13.816892	6.4846398
Low vol: Max risk	13.89454381	9.2552413	19.394049	2.4368806	13.519934	6.2860554
Flat vol: Min risk	38.42055898	26.770183	34.37765	3.8915639	9.798945	19.790382
Flat vol: Max risk	14.56586267	25.240415	1.2087348	1.9827425	30.175494	30.846321
High vol: Min risk	77.21869326	38.690593	37.273279	49.077158	68.087109	13.282882
High vol: Max risk	55.87174185	34.403723	80.666508	1.7514975	38.715533	22.794197
Skew vol: Min risk	10.65549708	13.451618	45.548332	38.436564	35.733705	53.623808
Skew vol: Max risk	40.34694127	57.242239	3.1506352	44.151369	49.943209	62.067165
Jump10%: Min risk	27.29342516	33.052179	61.309842	16.748973	21.541292	13.128707
Jump10%: Max risk	2.187893015	11.816301	44.666758	25.596927	18.674107	8.4501746
Jump30%: Min risk	7.863849895	18.940541	24.370843	3.2442942	13.195762	35.569141
Jump30%: Max risk	3.189848104	18.004995	11.497037	25.28611	9.9166706	2.5091388
Jump50%: Min risk	17.4403616	53.784907	63.516164	6.8961202	28.577785	15.939434
Jump50%: Max risk	1.434945302	12.80808	15.102957	26.254793	34.555644	12.133913
Jump100%: Min risk	85.7575873	0.050889	69.529149	4.0773441	9.6265993	35.235633
Jump100%: Max risk	10.28368041	45.895736	70.789234	19.460773	20.064724	5.9070918
Jump500%: Min risk	43.98615123	4.7296436	179.51098	60.895904	2.3132564	50.537462
Jump500%: Max risk	122.7914597	16.319405	27.392902	34.822791	33.907512	38.918415