Applications of Continuous Time Stochastic Processes in Sequential Clinical Research Design and Econometrics

Daniele Bregantini

Ph.D.

University of York

Economics

September 2014

Abstract

The principal subject of this thesis is hypothesis testing and related problems of estimation for stochastic processes. The thesis is concerned in particular with two areas: sequential hypothesis testing in a Bayesian setting and estimation of the parameters governing a continuous-time stochastic differential equation that drives data sampled at high-frequency. The former area is concerned with hypothesis testing for a newly developed healthcare technology and makes use of optimal stopping theory. The latter area sees the application of limit theorems for stochastic processes that allow to recover the true volatility process that can be estimated using the methods of moments estimator.

Contents

List of Tables

List of Figures

Preface

One of the subjects that fascinated me the most during my undergraduate studies was econometrics. I first encountered regression analysis in an introductory class in my second year where I became quickly interested in time series analysis. Many natural and economic phenomenons can be understood as events that evolve over time and being able to model them, through the use of time series econometrics, seemed to be of fundamental importance.

Graduate econometrics improved my understanding of the subject and I enjoyed discovering new techniques such as simulation methods. It was towards the end of my Graduate studies, while browsing books in the mathematics section of the library, that I found that the econometrics I was enjoying had a firm grounding in a well developed area of mathematics: stochastic processes. I then postponed entry to PhD and took the time to learn the mathematics involved in stochastic analysis. This was no easy task, understanding Brownian motion, new integration methods and Ito calculus proved to be tough. However, the application of these tools is very rewarding and the results contained in the thesis could not have been achieved without knowledge of these mathematical tools.

This thesis stems from my Graduate studies in Economics and Mathematics and is the result of my favourite research areas: probability, continuous time mathematics, econometrics and health economics. I became interested in the latter at the University of York, where, due to the large group of people working on Health Economics, it is only a matter of time before one get exposed to some of the issues pertaining to this fast-growing area of Economics.

I am extremely grateful to my supervisor Professor Jacco Thijssen for sharing his knowledge with me and for leading me into new fascinating areas within probability and statistics. Without

his support and outstanding supervision, this thesis would not have achieved such a diverse combination of topics and techniques. I also wish to thank Professor Peter Smith for helping transiting across different research areas and for supporting me in doing so.

I wish to thank my parents for their unconditional support over all these years. Without their help, this thesis would never have happened. I also wish to thank my partner Laetitia for supporting me in the difficult moments and for sharing with me all the exciting times of the PhD.

Acknowledgements

I wish to thank the members of the Thesis Advisory Panel, Professor Peter N. Smith and Martin Forster for their helpful comments. I wish to thank Paola Zerilli and Professor Peter N. Smith for introducing me to continuous-time financial econometrics and for their comments and guidance on chapter 7. This PhD thesis was supported by the Economic and Social Research Council ESRC Studentship.

Author's declaration

All the work contained in this thesis is original. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References. Chapter 5 is jointly authored with Professor Jacco Thjissen at the University or York.

Chapter 7 of this thesis has been presented at various conferences and workshops and benefited from comments of many participants: Brunel University Macro and Finance conference, the ESRC Student Conference in Leeds. Chapter 7 has been published in the Journal of Banking and Finance (2013), Vol. 37, Issue 2, pp. 4755-4764.

Chapter 1

Introduction

1.1 Introduction to Part II

1.1.1 Clinical Trials, Uncertainty and Health-care Technology Assessment

Decision makers face great challenges in coping with the pressure that health care systems face due to ageing populations, newly developed health-care technologies and changing population expectations about the use of resources in a health care system. In 1992, Australia became the first country to formalise and issue mandatory guidelines for health economic evaluations of pharmaceutical products as a requirement prior reimbursement. Other countries such as Canada, France, Finland and Portugal have developed similar guidelines. In the UK, the National Institute for Health and Care Excellence (NICE) appraises the clinical and economic benefits of new and existing health-care technologies and makes recommendations to the National Health Service (NHS) (Hjelmgren et al., 2001).

Health technology assessment (HTA) decisions are based on evidence of relative costs and effectiveness of alternative interventions. Decision makers, when evidence suggests that the incremental net benefit of the new intervention is positive, are faced with the decision of whether to adopt the new intervention over the existing one or, given the uncertainty surrounding the evidence, wait for more information.

When uncertainty about the net benefit of alternative treatments is present, there is a positive probability that the decision taken is wrong. Claxton (1999) argued that there are two conceptually separate but simultaneous decisions that must be made within a health care system: i) should a technology be adopted or reimbursed on the basis of existing evidence (and uncertainty surrounding outcomes and resources used) and ii) is further evidence required to support this adoption or reimbursement decision, and if existing evidence is deemed insufficient and further research is needed, what is the appropriate design for it ?

A source of information for the assessment of cost-effectiveness of newly developed health care technologies derives from clinical trials. The trials establishes the clinical efficacy and effectiveness of

medical therapies (Glick et al., 2007). A newly developed healthcare technology (e.g. a drug) needs to go through a number of stages of testing in human subjects before approval by the medical community. Phase I trials are concerned with aspects of clinical pharmacology and toxicity. A typical objective of a Phase I trial is to identify adequate dose levels that would avoid adverse side effects and involve sample sizes of 20 to 80 healthy volunteers. Phase II trials involve about 100 to 300 disease-affected patients and are concerned with effectiveness evaluation and safety aspects. Phase III trials take a further step in the evaluation of a newly developed drug and involve more than 1000 patients and can last more than 5 years. Patients suffering from the medical pathology are randomly assigned either to the new treatment or to standard treatment or, when possible, a placebo treatment. Phase IV trials involve additional testing and monitoring of the new treatment once it has been approved for general use. Phase IV trials are also referred to as post-marketing surveillance (Jennison and Turnbull, 2000). However, these categorisations are not strict and trials of different stages can overlap. Additionally, trials' stages can be subdivided in smaller categories (Burdette and Gehan, 1970).

In a situation where evidence from a trial accumulates over time, it is important to monitor results as the study proceeds in order to take action such as early termination or to modify the study design. The interim analysis of accumulating data is motivated by the following reasons: *ethical, administrative* and *economic*. In trials involving human subjects there is the need to frequently monitor results to ensure that humans are not exposed to treatments that are harmful or inferior to standard care. In trials where it appears to be no difference between two treatment it is important to terminate the study early in order to allocate resources to the next most promising treatment in the pipeline. Administrative reasons for interim analysis are the need to ensure that the trial is executed as planned, with subject taken for the relevant population and that eligibility criteria has been satisfied. Sequential analysis methods were originally introduced in order to obtain economic benefits as they exploit the trade-off given by economic costs of running a trial and statistical significance. Sequential methods, when compared

to traditional statistical inference, typically need a smaller sample size, time and costs. When evidence is positive, early termination means that the product can be exploited sooner and, in the case of negative evidence, stopping early involve saving resources (Jennison and Turnbull, 2000).

Many clinical trials are concerned with testing the *equivalence* of the new treatment efficacy and safety to the standard treatment. This is the typical study found in pharmaceutical applications when the objective is to compare two formulations of the same drug (Chow and Liu, 1992). The past 20 years have seen a great increase in the number of studies that make use of the information about cost and effect contained in clinical trials. Most frequently economic evaluation has been incorporated into the drug development process, typically Phase III as well as Phase IV. More than 20 years ago, economic evaluation derived from clinical trials typically were based from primarily from epidemiological data and used only some fey findings from the clinical trial. By the mid-1990s a growing number of trial-based economic evaluations consisted in direct observation of the impact of a therapy on cost and effect. These studies would observe short-term economic impacts and project on the long-term by the use of extrapolation methods (Glick et al., 2007).

Traditional sample size calculations for randomized clinical trials are based on arbitrary rules of inference such as type I and type II errors. Type I error if the failure to accept the null hypothesis when true is often set to $\alpha = 0.05$ regardless of the economic cost of making such error. Type II error is the failure to reject a false null hypothesis and is usually set to $\beta = 0.2$. With such value the resources allocated to the clinical trial will be wasted 20% of the time, even when the true treatment difference is equal to the smallest clinically important one. Type II error, as with type I error is set to a value that does not reflect the economic costs of making such error (William and Pinto, 2005).

In the absence of irreversibility (Palmer and Smith, 2000) or any costs associated with reversing a decision (Eckermann and Willan, 2008), the decision to adopt a technology can be based on expected cost effectiveness. However, if adoption involves large implementa-

tion costs, policy makers cannot switch costlessly between technologies as new evidence becomes available and uncertainty becomes more relevant in the decision making process (Palmer and Smith, 2000).

The theory of real options integrates uncertainty and irreversibility associated with a health-care technology into a unifying theory of economic evaluation that provides the decision maker with a framework to handle the uncertainty inherent in evidence on the cost-effectiveness of a health-care technology. Palmer and Smith (2000) propose the use of real options in order to handle uncertainty in HTA and to show that the degree of irreversibility of actions requires some flexibility in the timing of decisions.

Decision makers, when assessing the central estimates of likely effectiveness produced by sensitivity and/or statistical analyses, are often confronted with the problem of whether the evidence is enough to reject or defer implementation of a technology that appears to be cost effective. The problem is particularly true when the estimated range straddles the critical threshold values. Methods that seek to estimate uncertainty in HTA have been proposed and implemented, however the real option approach differs by i) considering the degree of uncertainty about the future state of the world ii) allowing the investment to have an irreversible commitment of resources iii) considering the case where there exists some discretion about the timing of investment.

Conventional investment decisions consider the case of now or never and little attention is given to the possibility of deferring investment to some later time when better information about costs and benefit of the investment become available. The orthodox investment strategy is to invest when the net present value of the expected investment benefit is greater than zero (Dixit and Pindyck, 1994). There is evidence that business managers, in particular when projects are of large values, do not follow the orthodox investment strategy and delay investment in order to have more information. Similarly, in HTA, deferral is an important and often taken decision as it allows for further information to be gathered.

Part II of the thesis brings together the ideas found in the real option approach, such as deferring investment until more informa-

tion about costs and benefits become available, with sequential hypothesis testing in clinical trials.

1.2 Introduction to Part III

1.2.1 Continuous-time methods in finance

Continuous-time methods in finance can be traced back to the seminal contribution of Merton (1969, 1971, 1973), Samuelson (1965) and Black and Scholes (1973). Merton (1969) formulated the intertemporal consumption and portfolio choice problem in a continuous-time stochastic dynamic programming setting and later showed how this framework can be used to understand equilibrium asset prices. Since the publication of the aforementioned articles, continuous time modelling has become an integral part of financial economics (Sundaresan, 2000). Due to the discrete nature of asset price data, until recently, the development of empirical procedures for the estimation and inference of continuos time models has been slower when compared to the discrete time modelling. The recent availability of reliable and accurate high-frequency assets prices data has given new impetus to continuous-time research and new powerful non-parametric techniques have been developed.

Due to the fact that return volatility plays a crucial role in a number of practical financial management decision such as risk management, asset allocation and option pricing there has been considerable effort in developing models that can produce accurate estimates and forecasts of current and future volatility. In contrast to returns, volatility is non directly observable and common approaches to estimate the unobservable return volatility are based on models that invoke strong parametric assumptions (e.g. ARCH, GARCH), estimated at daily or lower frequency. Alternatively, option prices can be 'reverted' using the appropriate pricing model in order to recover the 'market implied volatility'. However, such procedures are model dependent and include a time-varying volatility risk premium measure which produces biased forecasts (Andersen and Benzoni, 2009). Other methodologies exploit the information found in past volatility by taking a rolling window estimate and assume that such value will give a proxy for current and future

volatility. Such measures, due to the persistence of volatility can provide and indication of current volatility levels but ignore the feature of mean reversion that is found in volatility data series.

Continuos-time diffusion processes have been at the core of theoretical asset and derivative pricing models. However, only recently econometric procedures for their estimation have been developed.

Realised volatility is a non-parametric estimate of return variation that benefit from the information contained in high-frequency data. It was introduced concurrently by Andersen et al. (2001) and Andersen et al. (2003) and by Barnoff-Nielsen and Shepard (2001, 2004). Realised volatility is computed by summing up the intra-day squared returns; if prices do not exhibit micro-structure noise, realised volatility is a consistent estimator of integrated volatility for each trading day.

Part III of the thesis makes use of the realised volatility estimators in order to estimate the volatility and jump parameters driving the stochastic differential equation for stock prices. More specifically, chapter 7 takes the Bollerslev and Zhou (2002) approach and matches the sample moments of the realised volatility to the population moments of the integrated volatility implied by the particular model's structure. Analytical moments are derived and used in a GMM estimator in order to recover the underlying parameters governing the stochastic differential equation for the observed realised volatility. Further, the power and bipower variation measures introduced by Barnoff-Nielsen and Shepard (2004) are employed in order to recover intra-day realised jumps in returns. Barndorff-Nielsen and Shephard showed that by subtracting bipower variation from realised volatility it is possible to estimate the quadratic variation of the jump component. The inclusion of jumps in continuous time models has important implications and finds motivation in a number of empirical studies.

Part I

Mathematical and Statistical Preliminaries

Chapter 2

Mathematical Preliminaries

2.1 Probability spaces, random variables and stochastic processes

Some concepts from general probability theory are recalled before introducing a mathematical model for a random variable.

Definition 2.1. *(Oksendal, 2000) If Ω is a given set, a σ-algebra $\mathcal{F}$ on Ω is a family $\mathcal{F}$ of subsets of Ω with the following properties*

(i) $\emptyset \in \mathcal{F}$

(ii) $F \in \mathcal{F} \Rightarrow F^C \in \mathcal{F}$. *where* $F^C = \Omega \backslash F$ *is the complement of* F *in* Ω.

(iii) $A_1, A_2, \cdots \in \mathcal{F} \Rightarrow A : \cup_{i=1}^{\infty} A_1 \in \mathcal{F}$

The pair $(\Omega, \mathcal{F})$ is called a measurable space. A probability measure P on a measurable space $(\Omega, \mathcal{F})$ is a function $P : \mathcal{F} \to [0, 1]$ such that

(a) $P(\emptyset) = 0$, $P(\Omega) = 1$

(b) if $A_1, A_2, \cdots \in \mathcal{F}$ *and* $\{A_i\}_{i=1}^{\infty}$ *is disjoint (i.e.* $A_i \cap A_j = \emptyset$ *if* $i \neq j$*) then*

$$P\left(\cup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} P(A_i).$$

The triple $(\Omega, \mathcal{F}, P)$ is called a probability space. It is called a complete probability space if $\mathcal{F}$ contains all subsets G of Ω with $P-$outer

measure zero. i.e. with

$$P^*(G) := \inf\{P(F); F \in \mathcal{F}, G \subset F\} = 0$$

A collection $\mathcal{U}$ of all open subsets of a topological space Ω that generates a σ-algebra $\mathcal{H}_\mathcal{U}$, $\mathcal{B} = \mathcal{H}_\mathcal{U}$ is a *Borel σ-algebra* on Ω and elements $B \in \mathcal{B}$ are called *Borel sets*.

The subsets F of Ω belonging to $\mathcal{F}$ are said to be $\mathcal{F}$-measurable. These sets correspond to events and $P(F)$ is the probability that event F occurs.

Definition 2.2. Measurable function *(Williams, 2010) Let $\mathcal{F}$ be a σ-field on Ω. Suppose that $\xi : \Omega \to \mathbb{R}$. For $B \subseteq \mathbb{R}$, define*

$$\xi^{-1}(B) = \{\omega \in \Omega : \xi(\omega) \in B\}.$$

Then ξ is called $\mathcal{F}$-measurable if $\xi^{-1} : \mathcal{B} \to \mathcal{F}$, that is, $\xi^{-1}(B) \in \mathcal{F}, \forall B \in \mathcal{B}$.

Definition 2.3. Borel function*(Capinski and Kopp, 2004) For any interval $I \subset \mathbb{R}$, if all the sets*

$$f^{-1}(I) \in \mathcal{B}$$

we say that f is a Borel *function.*

A random variable X is an $\mathcal{F}-$measurable function $X : \Omega \to \mathbb{R}^n$. Every random variable induces a probability measure μ_X on $(\mathbb{R}, \mathcal{B})$, such as

$$\mu_X(B) = P(X^{-1}(B)).$$

where μ_X is referred to as the distribution of X.

Suppose P and Q are two probability measures on a measurable space $(\Omega, \mathcal{F})$. Then we have the following:

Theorem 2.1. Radon-Nikodym *Suppose Q is absolutely continuous with respect to P (i.e. $P \sim Q \iff [\forall\, A \in \mathcal{F},\ P(A) = 0 \iff Q(A) = 0]$) . Then there exists a random variable f such that*

$$Q(F) = \int_F f\, dP, \quad \forall\, F \in \mathcal{F}. \tag{2.1}$$

The function f is called the *Radon-Nikodym* derivative of Q with respect to P. This can be written as

$$f(\omega) = \frac{dQ}{dP}(\omega).$$

The Theorem tells if and in which way it is possible to change from one probability measure to another.

Definition 2.4. Conditional expectations *(Capinski and Kopp, 2004) For an integrable random variable ξ on a probability space $(\Omega, \mathcal{F}, P)$ and an event $B \in \mathcal{B}$ such that $P(B) \neq 0$ the* conditional expectations *of ξ given B is defined by*

$$\mathbb{E}(\xi \mid B) = \frac{1}{P(B)} \int_B \xi dP.$$

2.2 Stochastic process

Definition 2.5. *A stochastic process is a parametrized collection of random variables*

$$\{X_t\}_{t \in T}$$

defined on a probability space $(\Omega, \mathcal{F}, P)$ and assuming values in $\mathbb{R}^n$.

For each $t \in T$ there is a random variable

$$\omega \mapsto X_t(\omega); \quad \omega \in \Omega$$

where instead fixing $\omega \in \Omega$ gives a path of X_t

$$t \to X_t(\omega); \quad t \in T.$$

Often in stochastic analysis the process is viewed as a function of two variables

$$(t, \omega) \to X(t, \omega)$$

from $T \times \Omega$ into $\mathbb{R}^n$.

In light of the above we now introduce the notion of joint measurability for a stochastic process. The product σ-algebra is generated by the family of sets $\mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2$ and a process X, measurable with respect to the product σ-algebra $\mathcal{B}(\mathbb{R}) \otimes \mathcal{F}$, is said to be *jointly measurable.*

2.2.1 Convergence

When dealing with sequences or families of random variables different notions of convergence apply. One of the simplest notions of convergence is that of *almost everywhere*. Let E be a Borel subset of $\mathbb{R}^n$. For a given sequence (f_n) in $L^p(E)$, $p \geq 1$, the function $f_n \to f$ as $n \to \infty$ converges almost everywhere on E if there is a null set $F \subset E$, where E as a Borel subset of $\mathbb{R}^n$, such that $f_n \to f$ point wise on $E \backslash F$.

For the function $X : \Omega \to \mathbb{R}$, we consider convergence for all $\omega \in \Omega$ without considering those events that are in fact negligible (i.e. with probability zero). This leads to the definition of almost-sure convergence.

Definition 2.6. Almost-sure convergence *(Williams, 2010) Let* $(X_n : n \in N)$ *be a sequence of random variables and let* X *be a random variable on the probability triple* $(\Omega, \mathcal{F}, \mathbb{P})$. *We say that* $X_n \to X$ almost surely *if*

$$P(X_n - X) = 1.$$

The notion of convergence in probability gives a condition on the probability of events when $n \to \infty$.

Definition 2.7. Convergence in probability *(Capinski and Kopp, 2004) A sequence* (X_n) *of random variables on* $(\Omega, \mathcal{F}, P)$ *is said to converge in probability to a random variable* X *if for each* $\epsilon > 0$

$$P(\mid X_n - X \mid > \epsilon) = 0 \quad n \to \infty \tag{2.2}$$

For convergence in probability we consider the probability that $X_n - X$ is at least ϵ away from the limit while for almost sure convergence we consider the whole tail of the sequence $(X_n)_{n \geq k}$. The implication is that convergence almost-surely is stronger than convergence in probability. The last notion of convergence for this section is:

Definition 2.8. Convergence in distribution *(Cont and Tankov, 2003) A sequence* (X_n) *of random variables with values in* E *is said*

to converge in distribution to a random variable X if, for a bounded continuous function $f : E \to \mathbb{R}$

$$\mathbb{E}[f(X_n)] \to \mathbb{E}[f(X)] \quad n \to \infty. \tag{2.3}$$

2.3 Martingales and stopping times

2.3.1 Martingales

As basic datum, a probability space $(\Omega, \mathcal{F}, P)$ is considered and the following definitions (Poor and Hadjiliadsis, 2009) are introduced:

1. A process $X = (X_t; t \in T)$ is called *adapted* if for each t, X_t is $\mathcal{F}_t$-measurable. This means that if X is adapted then the value of X_t is known at time t.
2. A *filtration* $\{\mathcal{F}_t; t \in T\}$ is an increasing sequence of sub-σ-fields of $\mathcal{F}$. A filtration can be viewed as describing the evolution of information as time goes by.
3. A random sequence $\{X_t\}$ on $(\Omega, \mathcal{F}, P)$ is adapted to $\{\mathcal{F}_t\}$ if, for each t, X_t is $\mathcal{F}_t$-measurable.
4. $\{X_t, \mathcal{F}_t\}$ is a *submartingale* if
$$\mathbb{E}\{X_t \mid \mathcal{F}_l\} \geq X_l \quad \forall\, l \leq t \quad \text{a.s.}$$
5. $\{X_t, \mathcal{F}_t\}$ is a *supermartingale* is $\{-X_t, \mathcal{F}_t\}$ is a submartingale.
6. $\{X_t, \mathcal{F}_t\}$ is a *martingale* if
$$\mathbb{E}\{X_t \mid \mathcal{F}_l\} = X_l \quad \forall\, l \leq t \quad \text{a.s.}$$

A supermartingale decreases on average and a submartingale increases on average.

2.3.2 Stopping times

(Williams, 2010) A map $T : \Omega \to \mathbb{R}$ is called a *stopping time* (or Markov time) if

$$\{T \leq n\} = \{\omega : T(\omega) \leq n\} \in \mathcal{F}_n, \quad \forall\, n \leq \infty$$

A stopping time is a random variable taking values in the time set of the filtration. It can assume the value n only on events that are measurable with respect to the filtration at n.

Example 2.2. *Suppose that A_n is an adapted process, with $B \in \mathcal{B}$. Let*

$$T = \inf\{n \geq 0 : A_n \in B\} = \textit{first time of entry of } A \textit{ into } B.$$

If A never enters B, $T = \infty$. We have

$$\{T \leq n\} = \cup_{k \leq n}\{A_k \in B\} \in \mathcal{F}_n$$

so T is a stopping time.

2.3.3 Random walk

A symmetric random walk can be constructed by repeatedly tossing a fair coin. A probability p is assigned to the probability of an event H (head) and $q = 1 - p$ to the probability of an event T (tail), both equal to $1/2$. We consider an infinite sequence of tosses with ω_n denoting the outcome of the n-th toss. Let

$$X_j = 1 \quad if\ \omega_j = H$$
$$X_j = -1 \quad if\ \omega_j = T,$$

and define

$$M_k = \sum_{j=1}^{k} X_j, \quad k = 1, 2, \ldots \tag{2.4}$$

The process M_k, $k = 1, 2, \ldots$ is a *symmetric random walk.*

2.3.4 Brownian motion

Let $M_k,\ k = 1, 2, \ldots$ be a symmetric random walk as in (2.4). The Brownian motion is obtained as the limit of scaled random walks

$$W^n(t) = \frac{1}{\sqrt{n}} M_n(t) \quad \text{as } n \to \infty.$$

Definition 2.9. Brownian motion *Let $(\Omega, \mathcal{F}, P)$ be a probability space. A Brownian motion is a process $W(t), t \geq 0$ such that*

(a) $W(0) = 0$.

(b) The increment has normal distribution with mean $\mathbb{E}[W_{t_{i+1}} - W_{t_i}] = 0$ and variance $Var[W_{t_{i+1}} - W_{t_i}] = t_{i+1} - t_i$.

(c) For all $0 \leq t_1 < t_2 < t_3 < \cdots < t_m$ the increments $W_{t_{n+1}} - W_{t_n},\ n = 1, \ldots, m-1$ are independent.

(d) The paths $t \to W^\omega(t)$ are continuous for almost all ω.

Definition 2.10. p-th variation process *If $X_t(\cdot) \to \mathbb{R}$ is a continuous stochastic process, its p-th variation process,$\langle X, X \rangle_T^{(p)}$ is given by*

$$\langle X, X \rangle_t^{(p)}(\omega) = \lim_{\triangle t_k \to 0} \sum_{t_k \leq t} \mid X_{t_{k+1}}(\omega) - X_k(\omega) \mid^p$$

where $0 = t_1 < t_2 < \cdots < t_n = t$ and $\triangle t_k = t_{k+1} - t_k$. If $p = 1$ this process is referred to as the total variation and in the case $p = 2$ this process is called the quadratic variation.

Any differentiable process has finite total variation, it should be noted that the Brownian motion does not have a finite total variation and its path is thus almost surely not differentiable. The Brownian motion has, however, finite quadratic variation. For a Brownian motion $W_t \in \mathbb{R}$ the quadratic variation process ($p = 2$) is

$$\langle W, W \rangle_t^2(\omega) = t \quad a.s.$$

2.4 Reflection principle

We denote τ_a the first hitting time of the level a for the Brownian motion $(X_t, t \geq 0)$ starting at zero on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The function

$$\mathcal{N}(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{\frac{u^2}{2}} du$$

is the cumulative distribution function for a standard Gaussian distribution $\mathcal{N}(0,1)$. The first hitting time for a level a is

$$\tau_a(X) = \inf\{t \geq 0 \ : \ X_t = a\}.$$

The first time the process X is greater than a is given by

$$\tau_a^+ = \inf\{t \geq 0 \ : \ X_t \geq a\}$$

resp.

$$\tau_a^- = \inf\{t \geq 0 \ : \ X_t \leq a\}.$$

For a Brownian motion starting at zero, $X_0 = x$ and $a > x$ the hitting times $\tau_a^+ = \tau_a$ and $\tau_a^- = 0$. If $a < x$, we have $\tau_a^- = \tau_a$, and $\tau^+ = 0$. In what follows when we refer to hitting time it is intended to be the first hitting time. The running maximum (or respectively the running minimum) is

$$M_t^X = \sup_{s \leq t} X_s, \quad m_t^X = \inf_{s \leq t} X_s.$$

The process M_t is an increasing process with positive values. We consider a pair of random variables (W_t, M_t) where M is the maximum process for the Brownian motion W Next we present the reflection principle

Proposition 2.1. (Consequence of the reflection principle) *(Jeanblanc, Yor and Chesney, 2009) For $y \geq 0, x \leq y$, one has*

$$\mathbb{P}(W_t \leq x, M_t \geq y) = \mathbb{P}W_t \geq 2y - x). \tag{2.5}$$

From the symmetry of the normal distribution, we have

$$\mathbb{P}(W_t \leq x, M_t \geq y) = \mathbb{P}(W_t \geq 2y - x) = \mathcal{N}\left(\frac{x - 2y}{\sqrt{t}}\right) \tag{2.6}$$

The following theorem specifies the joint distribution of W_t and M_t. The result is used in Chapter 3 to compute the fist passage distribution for a Brownian motion.

Theorem 2.3. *(Jeanblanc, Yor and Chesney, 2009) Let W be a Brownian motion starting from zero and $M_t = \sup_{s\leq t} W_s$. Then,*

$$\begin{aligned} y \geq 0, x \leq y, \quad & \mathbb{P}(W_t \leq x, M_t \leq y) = \mathcal{N}\left(\frac{x}{\sqrt{t}}\right) - \mathcal{N}\left(\frac{x-2y}{\sqrt{t}}\right) \\ y \geq 0, x \geq y, \quad & \mathbb{P}(W_t \leq x, M_t \leq y) = \mathbb{P}(M_t \leq y) \\ & = \mathcal{N}\left(\frac{y}{\sqrt{t}}\right) - \mathcal{N}\left(\frac{-y}{\sqrt{t}}\right) \\ & \mathbb{P}(W_t \leq x, M_t \leq y) = 0. \end{aligned} \tag{2.7}$$

The distribution of the pair of random variables (W_t, M_t) is

$$\mathbb{P}(W_t \in dx, M_t \in dy) = \mathbf{1}_{\{y\geq 0\}}\mathbf{1}_{\{x\leq y\}} \frac{2(2y-x)}{\sqrt{2\pi t^3}} \exp\left(-\frac{(2y-x)^2}{2t}\right) dx\, dy \tag{2.8}$$

2.5 Stochastic integration

We start by considering a stochastic process $\{\phi(t,\omega)\}$ and a standard Brownian motion $\{W_t\}$ with $W_0 = 0$. Integrating such function over the Brownian motion

$$\int_0^t \phi_s dW_s \tag{2.9}$$

leads to an integration problem that cannot be dealt by standard integration methods (Lebesgue-Stieltjes). The above integral does not exist for many types of integrands as the paths of the Brownian motion are not of finite variation.

To define this type of integral we set some elementary processes $\{\phi_t\}$ such as

$$\phi_t = \sum_{i=0}^{2^n-1} c_i 1_{\left(\frac{i}{2^n}, \frac{i+1}{2^n}\right)}(t) \tag{2.10}$$

For each $\omega \in \Omega$ we define

$$\int_0^t \phi_s dW_s(\omega) = \sum_{i=0}^{2^n-1} c_i \left[W_{\frac{i+1}{2^n}t} - W_{\frac{i}{2^n}t}\right]. \tag{2.11}$$

When $c_i = W_{\frac{i}{2^n}t}$ the function is elementary and the right hand side of the above equation becomes

$$\mathbb{E}\left[\int_0^t \phi(s,\omega)dW_s(\omega)\right] = \sum_{j\geq 0} \mathbb{E}\left[W_{s_j}(W_{s_{j+1}} - W_{s_j})\right] = 0.$$

On the other hand when $c_i = W_{\frac{i+1}{2^n}t}$ the right hand becomes

$$\mathbb{E}\left[\int_0^t \phi(s,\omega)dW_s(\omega)\right] = \sum_{j\geq 0} \mathbb{E}\left[W_{s_{j+1}}(W_{s_{j+1}} - W_{s_j})\right] = S$$

suggesting that additional requirements are needed for the functions $\phi(t,\omega)$.

Ito suggested an integrand that preserves the martingale properties by choosing an integrand that is not forward looking in time. For each t, $\{\phi_t < u\} \in \mathcal{F}_t = \sigma\{W_s,\ s < t\}$ for any $u \in \mathbb{R}$ i.e. $\{\phi_t\}$ is $\mathcal{F}_t$-adapted.

Definition 2.11. *Let $\mathcal{V} = \mathcal{V}(S,T)$ be the class of function*

$$\phi(t,\omega) : [0,\infty) \times \Omega \to \mathbb{R}$$

such that

(i) $\phi(t,\omega)$ is jointly measurable in t and ω.

(ii) $\phi(t,\omega)$ is $\mathcal{F}_t$-adapted.

(iii) $\mathbb{E}\left[\int_S^T [\phi(t,\omega)]^2 dt\right] < \infty$

Definition 2.12. The Ito integral *(Oksendal, 2000) Let $f \in \mathcal{V}(S,T)$. Then the Ito integral from ϕ is defined as*

$$\int_S^T \phi(t,\omega)dW_t(\omega) = \lim_{n\to\infty} \int_S^T \phi_n(t,\omega)dW_t(\omega) \quad \text{limit in } L^2(P) \tag{2.12}$$

where $\{\phi^{(n)}\}$ is a sequence of elementary functions such that

$$\mathbb{E}\left[\int_S^T (\phi(t,\omega) - \phi_n(t,\omega)^2 dt\right] \to 0 \quad a.s.\ n \to \infty. \tag{2.13}$$

From Definition (2.12) we get the following

Lemma 2.1. The Ito isometry *(Oksendal, 2000)*

$$\mathbb{E}\left[(\phi(t,\omega)dW_t)^2\right] = \mathbb{E}\left[\int_S^T \phi^2(t,\omega)dt\right] \quad \forall\, \phi \in \mathcal{V}(S,T) \tag{2.14}$$

Lemma 2.2. *If* $\phi(t,\omega) \in \mathcal{V}(S,T)$ *and* $\phi^{(n)}(t,\omega) \in \mathcal{V}(S,T)$ *for* $n = 1,2,3,\ldots$ *and* $\mathbb{E}[\int_S^T (\phi^{(n)}(t,\omega) - \phi(t,\omega))^2 dt] \to 0$ *as* $n \to \infty$ *then*

$$\int_S^T \phi^{(n)}(t,\omega) dW_t(\omega) \to \int_S^T \phi(t,\omega) dW_t(\omega) \tag{2.15}$$

Example 2.4. *Assume* $W_0 = 0$. *Then*

$$\int_0^t W_s dW_s = \frac{1}{2} W_t^2 - \frac{1}{2} t^1$$

It is possible to extent the Ito integral $\int \phi \, dW$ for a larger class of integrands ϕ than $\mathcal{V}$ as long there exist a family of σ-algebras $\mathcal{H}_t$ such that W_t is a martingale with respect to $\mathcal{H}_t$ and ϕ_t is adapted to the filtration $\mathcal{H}_t$.

2.5.1 Ito processes

(Poor and Hadjiliadis, 2009) We consider the class of one-dimensional Ito processes of the form

$$X_t = X_0 + \int_0^t \mu_s ds + \int_0^t \sigma_s \, dW_s \tag{2.16}$$

where μ_s and σ_s are $\mathcal{F}_t$-adapted and

$$P\left(\int_0^t \sigma_s^2 \, ds < \infty \right) = 1 \tag{2.17}$$

$$P\left(\int_0^t \mid \mu_s \mid ds < \infty \right) = 1. \tag{2.18}$$

An Ito process is the sum of a finite variation term and a local martingale. For (2.16) the following shorthand notation[2] is used

$$dX_t = \mu_t dt + \sigma_t dW_t.$$

[1] For detail see Example (2.6)

[2] Note that given the non-differentiability of Brownian paths the equation has a meaning only in the integral form of (2.16).

If $\{X_t\}$is an Ito process and $g(t,x)$ is a first order continuously differentiable of the first variable and second order continuously differentiable function of the second variable, $Y_t = g(t, X_t)$ is also an Ito process.

2.6 Ito Formula

Theorem 2.5. Ito Formula. *(Oksendal, 2000) Let X_t be an Ito process given by*

$$dX_t = \mu_t dt + \sigma_t dW_t.$$

Let $g(t,x) \in C^2([0,\infty) \times \mathbb{R})$. Then

$$Y_t = g(t, X_t)$$

is again a Ito process, and

$$dY_t = \frac{\partial g}{\partial t}(t, X_t)dt + \frac{\partial g}{\partial x}(t, X_t)dX_t + \frac{1}{2}\frac{\partial^2 g}{\partial x^2}(t, X_t) \cdot (dX_t)^2 \quad (2.19)$$

where $(dX_t)^2 = (dX_t) \cdot (dX_t)$ is computed according to the rules

$$dt \cdot dt = dt \cdot dW_t = dW_t \cdot dt = 0, \quad dW_t \cdot dW_t = dt. \quad (2.20)$$

Example 2.6. *Take the integral*

$$I = \int_0^t W_s dW_s.$$

Take $X_t = W_t$ and $g(t,x) = \frac{1}{2}x^2$. Then

$$Y_t = g(t, W_t) = \frac{1}{2}W_t^2$$

By Ito formula

$$dY_t = d(\frac{1}{2}W_t^2) = W_t dW_t + \frac{1}{2}dt$$

or in integral form

$$\frac{1}{2}W_t^2 = \int_0^t W_s dW_s + \frac{1}{2}t$$

2.7 Markov process

For a stochastic differential equation such as

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

with $X_t \in \mathbb{R}^n$, $\mu(t,x) \in \mathbb{R}^n$, $\sigma(t,x) \in \mathbb{R}^{n\times m}$ and W_t is $m-$dimensional Brownian motion, μ is the *drift coefficient* and σ the *diffusion coefficient.*

Definition 2.13. Time-homogeneous Ito diffusion *(Oksendal, 2000) This is a stochastic process $X_t(\omega) = X(t,\omega) : [0,\infty) \times \Omega \to \mathbb{R}^n$ satisfying a stochastic differential equation of the form*

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t, \quad t \geq s; \quad X_s = x; \tag{2.21}$$

In a time homogeneous Ito diffusion coefficients μ and σ depend on X_t and not on t. The unique solution for (2.21) is denoted by $X_t = X_t^{s,t};\ t \geq s$. If $s = 0$, X_t^x stands for $X_t^{0,x}$.

Let the probability law for a given Ito diffusion when the initial value is $X_0 = x \in \mathbb{R}^n$ be denoted by Q^x. The expectation with respect to this probability law Q^x is $\mathbb{E}^x[\cdot]$.

Theorem 2.7. Markov property *Let f be a Borel function from $\mathbb{R}^n \to \mathbb{R}$. Then for $t, h \geq 0$*

$$\mathbb{E}^x[f(X_{t+h}) \mid \mathcal{F}_t^{(m)}](\omega) = \mathbb{E}^{X_t(\omega)}[f(X_h)] \tag{2.22}$$

The Markov property holds also for stopping times $\tau(\omega)$. This lead to the following Theorem.

Theorem 2.8. Strong Markov property for Ito diffusions *(Oksendal, 2000) Let f be a bounded Borel function on $\mathbb{R}^n$, τ a stopping time with respect to the filtration $\mathcal{F}_t^{(m)}$ generated by $\{W_s; s \leq t\}$, $\tau < \infty$ a.s. Then*

$$\mathbb{E}^x[f(X_{\tau+h}) \mid \mathcal{F}_\tau^{(m)}] = \mathbb{E}^{X_\tau}[f(X_h)] \quad \forall\ h \geq 0. \tag{2.23}$$

2.8 The generator for an Ito diffusion

Associated to each Ito diffusion one can find a second order partial differential operator A. The *generator* A encodes a great deal of information about the process X_t.

Definition 2.14. Infinitesimal generator *Let $\{X_t\}$ be a (time-homogeneous) Ito diffusion in $\mathbb{R}^n$. For a bounded Borel function f, the (infinitesimal) generator A of X_t is defined by*

$$Af(x) = \lim_{t \downarrow 0} \frac{\mathbb{E}[f(X_t)] - f(x)}{t}, \qquad x \in \mathbb{R}^n$$

The formula for the generator of an Ito diffusion is given by

Theorem 2.9. *(Oksendal, 2000) Let X_t be an Ito diffusion and f a bounded function as in Definition (2.14)*

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t.$$

If $f \in C_0^2(\mathbb{R}^n)$, then

$$\mathcal{A}f(x) = \sum_{i=1}^{n} b_i(x) \frac{\partial f}{\partial x_i} + \frac{1}{2} \sum_{i,j=1}^{n} (\sigma\sigma^T)_{i,j}(x) \frac{\partial^2 f}{\partial x_i \partial x_j}. \tag{2.24}$$

Having defined the generator for an Ito diffusion, it is now possible to introduce some tools that permit to compute the expected value of an Ito diffusion at a stopping time.

2.8.1 Dinkyn formula

Theorem 2.10. Dynkin's formula *Let $f \in C_0^2(\mathbb{R}^n)$. Suppose τ is a stopping time, $\mathbb{E}^x[\tau] < \infty$. Then*

$$\mathbb{E}^x[f(X_\tau)] = f(x) + \mathbb{E}^x \left[\int_0^\tau \mathcal{A}f(X_s)ds \right] \tag{2.25}$$

where $\mathbb{E}^x$ is the expectation w.r.t. the natural probability distribution for an Ito process starting at x.

This differential equation is also called the *characteristic operator.*

2.9 Girsanov's theorem

Lemma 2.3. Bayes' rule *Let μ and ν be two probability measures on a measurable space $(\Omega, \mathcal{G})$ such that $d\nu(\omega) = f(\omega)d\mu(\omega)$ for some $f \in L^1(\mu)$. Let X be a random variable on $(\Omega, \mathcal{G})$ such that*

$$\mathbb{E}_\nu[\mid X \mid] = \int_\Omega \mid X(\omega) \mid f(\omega)d\mu(\omega) < \infty.$$

Let $\mathcal{H}$ be a σ-algebra, $\mathcal{H} \subset \mathcal{G}$. Then

$$\mathbb{E}_\nu[X \mid \mathcal{H}] \cdot \mathbb{E}_\mu[f \mid \mathcal{H}] = \mathbb{E}_\mu[fX \mid \mathcal{H}] \quad a.s. \tag{2.26}$$

Theorem 2.11. Girsanov's Theorem *Oksendal (2010)Let $Y(t) \in \mathbb{R}^n$ be an Ito process of the form*

$$dY(t) = \beta(t,\omega)dt + \theta(t,\omega)dB(t) \quad t \leq T \tag{2.27}$$

where $B(t) \in \mathbb{R}^m$,$\beta(t,\omega) \in \mathbb{R}^n$ and $\theta(t,\omega) \in \mathbb{R}^{n\times m}$. Suppose there exists processes $u(t,\omega) \in \mathcal{W}_{\mathcal{H}}^m$ and $\alpha(t,\omega) \in \mathcal{W}_{\mathcal{H}}^m$ such that

$$\theta(t,\omega)u(t,\omega) = \beta(t,\omega) - \alpha(t,\omega) \tag{2.28}$$

Put

$$M_t = exp\left(-\int_0^t u(s,\omega)dB_s - \frac{1}{2}\int_0^t u^2(s,\omega)ds\right); \quad t \leq T \tag{2.29}$$

and

$$dQ(\omega) = M_T(\omega)dP(\omega) \quad on\ \mathcal{F}_T^{(m)}. \tag{2.30}$$

Assume that M_t is a martingale (w.r.t $\mathcal{F}_T^{(n)}$ and P). then Q is a probability measure on $\mathcal{F}_T^{(m)}$, the process

$$\hat{B}(t) = \int_0^t u(s,\omega)dx + B(t); \quad t \leq T \tag{2.31}$$

is a Brownian motion w.r.t. Q and in terms of $\hat{B}(t)$ the process $Y(t)$ has the stochastic integral representation

$$dY(t) = \alpha(t,\omega)dt + \theta(t,\omega)d\hat{B}(t). \tag{2.32}$$

2.10 Optimal stopping problem

Following Oksendal (2000) we consider an Ito diffusion X_t on $\mathbb{R}^n$ and let g be a reward function on $\mathbb{R}^n$, satisfying

1.
$$g(\xi) \geq 0 \quad \text{for all } \xi \in \mathbb{R}^n \tag{2.33}$$

2. g is continous

We seek to find an optimal stopping time $\tau^* = \tau^*(x, \omega)$ for $\{X_t\}$ such that

$$\mathbb{E}^x[g(X_{\tau^*}] = \sup_\tau \mathbb{E}^x[g(X_\tau)], \quad \forall\, x \in \mathbb{R}^n, \tag{2.34}$$

the supremum being taken over all stopping times τ. We seek the optimal expected reward

$$g^*(x) = \mathbb{E}^x[g(X_\tau)]. \tag{2.35}$$

We introduce some definitions.

Definition 2.15. *(Oksendal, 2000) A measurable function $f : \mathbb{R}^n \to [0, \infty]$ is supermeanvalued (w.r.t X_t) if*

$$f(x) \geq \mathbb{E}^x[f(X_\tau)] \tag{2.36}$$

for all stopping times τ and all $x \in \mathbb{R}^n$. If f is also lower semi-continuous [3] *then f is called l.s.c. superharmonic or just* superharmonic.

Remark 2.1. *If $f \in C^2(\mathbb{R}^n)$ it follows from Dynkin formula that f is super harmonic w.r.t. X_t iff*

$$\mathcal{A}f \leq 0$$

where $\mathcal{A}$ is the characteristic operator of X_t.

[3] Semi continuity is weaker than continuity. A function f is lower semicontinuous at a point y_0 if for any $\epsilon > 0$ there exists a neirborhood U of y_0 such that $f(y) \geq f(y_0) - \epsilon \;\; \forall\, y \in U$. This can be written as $\liminf_{y \to y_0} f(y) \geq f(y_0)$

The next concepts are very important.

Definition 2.16. *(Oksendal, 2000) Let h be a real measurable function on $\mathbb{R}^n$. If f is a super harmonic (supermeanvalued) function $f \geq h$ we say that f is a super harmonic (supermeanvalued) majorant of h (w.r.t. X_t). The function*

$$\bar{h} = \inf_f f(x); \quad x \in \mathbb{R}^n, \tag{2.37}$$

the inf *being taken over all supermeanvalued majorant f of h, is called the least supermeanvalued majorant of h.*

Suppose there exists a function $\hat{h}$ such that:

1. $\hat{h}$ is a super harmonic majorant of h and
2. if f is any other super harmonic majorant of h then $\hat{h} \leq f$.

Then $\hat{h}$ is called the *least superharmonic majorant* of h. If $\bar{h}$ is lower semicontinuos, then $\hat{h}$ exists and $\hat{h} = \bar{h}$. It will be possible to prove that if g is non negative and lower semicontinuous, then $\hat{g}$ exists and $\hat{g} = \bar{g}$ (See Theorem 2.12 below).

Let $g \geq 0$ and let f be a supermeanvalued majorant of g. Then if τ is a stopping time

$$f(x) \geq \mathbb{E}^x[f(X_\tau)] \geq \mathbb{E}^x[g(X_\tau)].$$

So

$$f(x) \geq \sup_\tau \mathbb{E}^x[g(X_\tau)] = g^*(x).$$

Therefore we always have, if $\hat{g}$ exists,

$$\hat{g} \geq g^*(x) \text{ for all } x \in \mathbb{R}^n \tag{2.38}$$

Below we formally state that the converse inequality holds, leading to

$$\hat{g} = g^*.$$

Theorem 2.12. *(***Construction of the least superharmonic majorant***) Let $g = g_0$ be a non negative, lower semicontinuos function on $\mathbb{R}^n$, and define inductively*

$$g_n(x) = \sup_{t \in S_n} \mathbb{E}^x[g_{n-1}(X_t)], \tag{2.39}$$

where $S_n = \{k \cdot 2^{-n}; 0 \leq k \leq 4^n\}, n = 1, 2, 3, \ldots$. Then $g_n \uparrow \hat{g}$ and $\hat{g}$ is the least superharmonic majorant of g. Moreover, $\hat{g} = \bar{g}$.

It is possible to replace the sets S_n with the whole interval $[0, \infty]$:

Corollary 2.1. *Define $h_0 = g$ and inductively*

$$h_n(x) = \sup_{t \geq 0} \mathbb{E}^x[h_{n-1}(X_t)]; \qquad n = 1, 2, 3, \ldots$$

Then $h_n \uparrow \hat{g}$.

2.10.1 Reward functions with negative values

The non-negativity assumption on g given by (2.33) can be relaxed. It can be noted that if g is bounded below, $g \geq -M$ where M is a constant, it is possible to write

$$g_1 = g + M \geq 0$$

and apply the theory to g_1. Since

$$\mathbb{E}^x[g(X_\tau)] = \mathbb{E}^x[g_1(X_\tau)] - M \qquad if\ \tau < \infty\ a.s.$$

we have $g^*(x) = g_1^*(x) - M$ and the problem can be reduced to an optimal stopping problem for the function g_1. If g is not bounded below, problem (2.34) and (2.35) are not well defined unless the following condition holds

$$\mathbb{E}^x[g^-(X_\tau)] < \infty \qquad \text{for all } \tau \tag{2.40}$$

where

$$g^-(x) = -\min(g(x), 0).$$

2.10.2 Connection with variational inequalities

The 'high contact principle' provides useful information in determining g^*. the principle states that, under certain conditions, the solution g^* is a C^1 function on $\mathbb{R}^n$ if $g \in C^2(\mathbb{R}^n)$.

Oksendal (2000) provides a verification theorem for optimal stopping that makes it simpler to verify that a given candidate for g^* is actually equal to g^*.

In what follows, we consider a domain G in $\mathbb{R}^k$ and let

$$dY_t = b(Y_t)dt + \sigma(Y_t)dW_t; \quad Y_0 = y \tag{2.41}$$

be an Ito diffusion in $\mathbb{R}^k$. Define

$$\tau_G = \tau_G(y, \omega) = \inf\{t > 0;\ Y_t(\omega) \notin G\} \tag{2.42}$$

Let $f : \mathbb{R}^k \to \mathbb{R}$, $g :\ \mathbb{R}^k \to \mathbb{R}$ be continuous functions satisfying

(i)
$$\mathbb{E}^y \left[\int_0^{\tau_G} f^-(Y_t)dt \right] < \infty \quad \text{for all } y \in \mathbb{R}^k \tag{2.43}$$

(ii) the family $\{g^-(Y_\tau);\ \tau \text{ stopping time } \tau \leq \tau_G\}$ is uniformly integrable w.r.t. $\mathbb{R}^y$ (the probability law of Y_t), for all $y \in \mathbb{R}^k$ (2.44)

Let $\mathcal{I}$ denote the set of all stopping times $\tau \leq \tau_G$. Consider the following problem: seek $\phi(y)$ and $\tau^* \in \mathcal{I}$ such that

$$\phi(y) = \sup_{\tau \in \mathcal{I}} J^\tau(y) = J^{\tau^*}(y), \tag{2.45}$$

where

$$J^\tau(y) = \mathbb{E}^y \left[\int_0^\tau f(Y_t)dt + g(Y_\tau) \right] \quad \text{for } \tau \in \mathcal{I}. \tag{2.46}$$

Note that $J^0(y) = g(y)$ and so we have

$$\phi(y) \geq g(y) \quad \text{for all} \quad y \in G. \tag{2.47}$$

The partial differential operator

$$L = L_Y = \sum_{i=1}^{k} b_i(y)\frac{\partial}{\partial y_i} + \frac{1}{2}\sum_{i,j=1}^{k}(\sigma\sigma^T)ij(y)\frac{\partial^2}{\partial y_i \partial y_j}$$

coincides with the generator A_Y of Y_T on $C_0^2(\mathbb{R})^k$.

Theorem 2.13. (Variational inequalities for optimal stopping) *(Oksendal, 2000)*

a) *Suppose we can find a function* $\phi : \bar{G} \to \mathbb{R}$ *such that*

i) $\phi \in C^1(G) \cap C(\bar{G})$

ii) $\phi \geq g$ *on* G $lim_{t\to\tau_{\bar{G}}}\ \phi(Y_t) = g(Y_{\tau G})\mathcal{X}_{\{\tau_G<\infty\}}$ *a.s.*

Define

$$D = \{x \in G; \phi(x) > g(x)\}.$$

Suppose Y_t *spends* 0 *time on* ∂D *a.s., i.e.*

iii)

$$\mathbb{E}^y\left[\int_0^{\tau_G} \mathcal{X}_{\partial D}(Y_t)dt\right] = 0 \text{ for all } y \in G$$

and suppose that

iv) ∂D *is a Lipschitz surface, i.e.* ∂D *is locally the graph of a function* $h : \mathbb{R}^{k-1} \to \mathbb{R}$ *such that there exists* $K < \infty$ *with*

$$\mid h(x) - h(y) \mid \leq K \mid x - y \mid \quad \text{for all} \quad x, y$$

Moreover, suppose that following:

v) $\phi \in C^2(G\backslash\partial D)$ *and the second order derivatives of* ϕ *are locally bounded near* ∂D

vi) $L\phi + f \leq 0$ *on* $G\backslash\bar{D}$. *Then*

$$\phi(y) \geq \phi(y) \quad \forall\ y \in G.$$

b) *Suppose in addition to the above, that*

a) $L\phi + f = 0$ *on* D

b) $\tau_D := \inf\{t > 0; Y_t \in D\} < \infty$ *a.s* $\mathbb{R}^y$ $\forall$ $y \in G$

and

c) the family $\{\phi(Y_\tau); \tau \leq \tau_D, \tau \in \mathcal{T}\}$ *is uniformly integrable w.r.t* $\mathbb{R}^y$*, for all* $y \in G$*. Then*

$$\phi(y) = \Phi(y) = sup_{\tau \in \mathcal{T}} \mathbb{E}^y \left[\int_0^\tau f(Y_t)dt + g(Y_\tau) \right]; \quad y \in G \tag{2.48}$$

and

$$\tau^* = \tau_D \tag{2.49}$$

is an optimal stopping time for this problem.

In most applications, including the ones in this thesis, the above theorem is used in the following way. Consider a strongly Markovian stochastic process $(X_t)_{t \geq 0}$, with state space E, and the optimal stopping problem

$$F^*(x) = \sup_\tau E_x \left[e^{-r\tau} F(X_\tau) \right].$$

The state space can be split into two regions: $C = \{x \in E | F^*(x) > F(x)\}$ and $D = \{x \in E | F^*(x) = F(x)\}$. On the set C, which is often called the *continuation region*, the function F^* should satisfy the differential equation $L_X F^* = rF^*$ (the Bellman equation). This guarantees that F^* is the *smallest* superharmonic function dominating F. The function F^* is C^2 on C and should be C^1 on E. So, on the boundary of C it should hold that $F^* = F$ and $\partial F^*/\partial x = \partial F/\partial x$. The latter condition (of then called the *smooth pasting* or *high contact* principle) ensures optimality of the function F^*.

Note that Theorem 2.13 shows that one should also verify that $L_X F^* < rF^*$ on D, i.e. that the function F^* is superharmonic on the entire domain E. This is often neglected in applications of optimal stopping theory.

Chapter 3

Statistical preliminaries

3.1 Bandit problems

This sections aims at briefly outline continuous time bandits and introduces the reader, following Berry and Fristedt (1985), to sequential selections for $k \geq 2$ stochastic processes (or arms, treatments etc.). Multi-armed Bandits are an example of sequential allocation problem with exploration-exploitation trade-off. At each time step, some resource is allocated to an action and some observable payoff is derived. The goal is to maximise gains obtained in the sequence of allocations. Bandits naturally address the trade-off between exploration and exploitation found in sequential experiments. The player needs to balance the exploitation of actions that gave rewards in the past with the exploration of actions that might give larger gains in the future (Bubekc and Cesa-Bianchi, 2012). Bandit problems have been used to allocate resources between different competing projects in large organisations.

Suppose there are two treatment for a certain disease and that patients enter a treatment sequentially. One of the two treatments must be used on a patient and the overall goal is to treat as many patients as effectively as possible. In bandit problems time can be discrete or continuous and the processes are corespondingly discrete or continuous. Typically parameters that characterise the process are unknown and the process selected for observation at any time depends on the previous selection and results. A selection procedure specifies which process is to be selected on the basis of previous observations and selections. Utility for a strategy is defined by averaging over all possible histories resulting from the strategy (Berry

and Fristedt, 1985).

3.1.1 Continuos-time bandits: Brownian motion with unknown drift

Two-point prior

The following section follows closely from Berry and Fristedt (1985). The characteristics of arm 2 are assumed to be known, with constant rate λ, generating the following reward process Y_2 and Y_1

$$Y_2(t) = \lambda t.$$

Arm 1 has an unknown drift and is it generates a Brownian motion

$$Y_1(t) = \theta_1 t + B(t),$$

where $B(t)$ is a standard Wiener process with mean 0 and variance 1 at $t = 1$. The drift θ_1 is random and follows a mixture distribution

$$F = p\delta_a + (1-p)\delta_b,$$

where $b < a$ and F is the distribution of θ_1. It is also assumed that $b < \lambda < a$. The discount function is $e^{-\beta t}$ for some constant $\beta, 0 < \beta < \infty$. The *worth* of a strategy υ is

$$W(F, \lambda; e^{-\beta t}; \upsilon) = \mathbb{E}_\upsilon \int_0^\infty e^{-\beta t} dY_\upsilon(t), \tag{3.1}$$

where $\upsilon(t)$ indicates the arm being observed at time t. For the two arms $\upsilon_1(t) = 1$ and $\upsilon_2(t) = 2$, it follows that

$$\begin{aligned} W(F, \lambda; e^{-\beta t}; \upsilon) &= \mathbb{E}_{\upsilon=1} \int_0^\infty e^{-\beta t} \mathbf{1}_{\{\upsilon(t)=1\}} dY_1(t) \\ &\quad + \mathbb{E}_{\upsilon=2} \int_0^\infty e^{-\beta t} \mathbf{1}_{\{\upsilon(t)=2\}} dY_2(t). \end{aligned} \tag{3.2}$$

The latter term in (3.2) is meaningful provided the random set $\{t : \upsilon(t) = 2\}$ is an a.s. measurable subset of $[0, \infty]$. Given that the integrals in (3.2) are stochastic the assumption of progressive measurability with respect to the stochastic process need to be satisfied. Next the supremum over the strategies is defined as

$$V(F, \lambda; e^{-\beta t}) = \sup_{\upsilon(t)} W(F, \lambda; e^{-\beta t}, \upsilon),$$

There is a constant $C \in (0,1)$ depending on p such that arm 1 is optimal at time t if the current probability that $\theta = a$ is greater then C. Arm 2 becomes optimal when the current probability is less or equal to C. The decision maker, should then choose arm 2 indefinitely into the future when $p = F(\{a\} \leq C)$. When $p > C$ the decision maker should select arm 1 initially and stay with arm 1 until $p(t, Y_1(t)) = C$, where $p(t,y)$ denotes the conditional probability that $\theta_1 = a$ given that $Y_1(t) = y$. At such time the decision maker should select permanently arm 2 (Berry and Fristedt, 1985).

3.2 Sequential probability ratio test

This section follows closely Siegmund (1985) and aims at introducing basic lemmas and the mathematical concepts that underpin sequential analysis. The idea is to introduce the main concepts used in sequential analysis in both the discrete and continuous setting.

Let x denote a random variable with probability density function f. The Neyman-Pearson Lemma for testing a hypothesis $H_0 : f = f_0$ against an alternative $H_1 : f = f_1$, for a constant r and a likelihood ratio $l(x) = f_1(x)/f_0(x)$, is given by

$$\begin{aligned} &\text{Reject } H_0 \quad if \quad l(x) \geq r, \\ &\text{Reject } H_1 \quad if \quad l(x) < r. \end{aligned} \tag{3.3}$$

This class of test is optimal from the point of view of frequentist and Bayesians. With P_i denoting the probability measure conditional on the hypothesis $H_i, i = 0, 1$, any test of H_0 against H_1, which is based on observing x and has a significance level no larger than $\alpha = P_0\{l(x) \geq r\}$, must have power no larger than $P_1\{l(x) \geq r\}$ (Cox and Hinkley (1974), p. 91).

In a sequential probability ratio test there is a third possibility: for intermediate values of $l(x)$, one collects more data. Let $x_1, x_2, \ldots$ be a sequence of random variables with joint density functions

$$P\{x_1 \in d\xi_1, \ldots x_n \in d\xi_n\} = f_n(\xi_1, \ldots .x_n)d\xi_1, \ldots, d\xi_n, \quad (n = 1, 2, \ldots)$$

where each $d\xi_n$ for random variables $i = 1, \ldots, n$ represent a probability measure.

Simple hypothesis $H_0 : f_n = f_{0n}$ for all n are tested against $H_1 : f_n = f_{1n}$ for all n. Let $l_n = l_n(x_1, \ldots, x_n) = f_{1n}(x_1, \ldots, x_n)/f_{0n}(x_1, \ldots, x_n)$. By taking constants $0 < A < B < \infty$ and sampling sequentially the random variables $x_1, x_2, \ldots$ until the random time

$$\begin{aligned} N &= \text{first } n \geq 1 \quad \text{such that } l_n \notin (A, B) \\ &= \infty \text{ if } l_n \in (A, B) \quad \text{for all } n \geq 1 \end{aligned} \tag{3.4}$$

Stop sampling at time N and if $N < \infty$

$$\begin{aligned} \text{Reject } H_0 \qquad & \text{if } l_N \geq B. \\ \text{Accept } H_0 \qquad & \text{if } l_N \leq A. \end{aligned} \tag{3.5}$$

Assuming that the procedure terminates (i.e. $P_i\{N < \infty\} = 1$ for $i = 1, 0$) the size of the test is $P_0\{l_N \geq B\}$ and the power of the test is $P_1\{l_N \leq A\}$.

3.2.1 Approximations for $P_i\{l_N \geq B\}$

In what follows $\alpha = P_0\{l_N \geq B\}$ and $\beta = P_1\{l_N \leq A\}$ is related to A and B. Denote B_k as the subset of n-dimensional space where

$$A < l_k(\xi_1, \ldots, \xi_k) < B$$

for $k = 1, 2, \ldots, n-1$ and $l_n(\xi_1, \ldots, \xi_n) \geq B$. The probability α can be computed as

$$\begin{aligned} \alpha &= P_0\{l_N \geq B\} \\ &= \mathbb{E}_1\left[l_N^{-1}; l_N \geq B\right] \\ &\leq B^{-1} P_1\{l_N \geq B\} = B^{-1}(1 - \beta). \end{aligned} \tag{3.6}$$

Equation (3.6) takes the definition of α involving probability under $i = 0$ and shows the relationship between the two probabilities measures $P_i, i = 1, 0$ and the error probabilities α and β. Next, by interchanging A and B

$$\beta = P_1\{l_N \leq A\} \leq AP_0\{l_N \leq A\} = A(1-\alpha). \tag{3.7}$$

The equalities (3.6) and (3.7) are not exact due to the fact that the l_n does not have to hit the boundary exactly when crossing levels A or B. It is possible to solve the inequalities as

$$\alpha \cong \frac{1-A}{B-A} \tag{3.8}$$

$$\beta \cong A\left(\frac{B-1}{B-A}\right). \tag{3.9}$$

3.2.2 Approximation for $\mathbb{E}_i(N)$

In order to compute the expected value for N the observations x_n are assumed to be independent and identically distributed. The likelihhod is give by

$$l_n = \Pi_{k=1}^{n}\{f_1(x_k)/f_0(x_k)\}.$$

By taking the log

$$\log l_n = \sum_{k=1}^{n} \log\{f_1(x_k)/f_0(x_k)\}$$

a sum of i.i.d variables is obtained. Given constants $a = \log A$ and $b = \log B$

$$\begin{aligned} N &= \text{first } n \geq 1 \quad \text{such that } \log l_n \notin (a,b) \\ &= \infty \text{ if } \log l_n \in (a,b) \quad \text{for all } n \end{aligned} \tag{3.10}$$

Proposition 3.1. Wald's identity *(Siegmund (1985)) Let $y_1, y_2, \ldots$ be i.i.d. with $\mu = \mathbb{E}(y_1)$. Let M be any integer valued random variable such that $\{M = n\}$ is an event determined by conditions on $y_1, \ldots, y_n$ (and independent of $y_{n+1}, \ldots$) for all $n = 1, 2, \ldots$ and assume that $\mathbb{E}(|M|) < \infty$. Then $\mathbb{E}\left(\sum_{k=1}^{M} y_k\right) = \mu\mathbb{E}(M)$.*

By Wald's identity

$$\mathbb{E}_i\{\log l_N\} = \mu_i \mathbb{E}_i(N), \tag{3.11}$$

with

$$\mu_i = \mathbb{E}_i\left[\log\{f_1(x_1)/f_0(x_1)\}\right] \qquad i = 0, 1.$$

By using (3.8-3.9) the $\log l_n$ can be understood as a two-valued random variable taking on values a and b,

$$\mathbb{E}_i\{\log l_N\} \cong aP_i\{l_N \leq A\} + bP_i\{l_N \geq B\}. \tag{3.12}$$

Using (3.8-3.9), (3.11) and (3.12) gives the following approximation

$$\mathbb{E}_1(N) \cong \mu_1^{-1}\{aA(B-1) + bB(1-A)\}/(B-A) \tag{3.13}$$

and

$$\mathbb{E}_0(N) \cong \mu_0^{-1}\{aA(B-1) + bB(1-A)\}/(B-A) \tag{3.14}$$

Remark 3.1. *The expected sample size approximations 3.13 and 3.14 can be expressed in term of the error probabilities α and β. From (3.8), (3.9) and (3.12)*

$$\mathbb{E}_1(N) \cong \mu_1^{-1}\left\{(1-\beta)\log\left(\frac{1-\beta}{\alpha}\right) + \beta\log\left(\frac{\beta}{1-\alpha}\right)\right\} \tag{3.15}$$

and

$$\mathbb{E}_0(N) \cong \mu_0^{-1}\left\{\alpha\log\left(\frac{1-\beta}{\alpha}\right) + (1-\alpha)\log\left(\frac{\beta}{1-\alpha}\right)\right\}. \tag{3.16}$$

3.2.3 Optimality of the Sequential Probability Ratio Test

In what follows a random variable T with values $\{1, 2, \ldots, \infty\}$ is referred to as a stopping time if $\{T = n\} \in \mathcal{F}_n$ for all n. By observing a sequence of random variables $\{x_1, x_2, \ldots, x_n\}$ it is possible to infer whether $T = n$.

A sequential probability ratio test for a simple hypothesis is optimal as it minimises the expected sample size under H_0 and H_1 among all tests having no larger error probabilities (Siegmund, 1985). For a conventional test of $H_0 : f = f_0$ against $H_1 : f = f_1$ with error probabilities $\alpha = P_0\{\text{Reject } \ H_0\}$ and $\beta = P_1\{\text{Accept } \ H_0\}$ approximations (3.15) and (3.16) give the relationship between the expected sample size and the error probabilities. The next Theorem states that these expected sample sizes are the smallest for these error probabilities.

Theorem 3.1. *(Siegmund, 1985) Let T be the stopping time of any test of $H_0 : f = f_0$ against $H_1 : f = f_1$ with error probabilities α,β $(0 < \alpha < 1,\ 0 < \beta < 1)$. Assume $\mathbb{E}_i(T) < \infty$ $(i = 0, 1)$. Then*

$$\mathbb{E}_1(T) \geq \mu_1^{-1}\left\{(1-\beta)\log\left(\frac{1-\beta}{\alpha}\right) + \beta\log\left(\frac{\beta}{1-\alpha}\right)\right\}$$

$$\mathbb{E}_0(T) \geq \ \mu_0^{-1}\left\{\alpha \log\left(\frac{1-\beta}{\alpha}\right) + (1-\alpha)\log\left(\frac{\beta}{1-\alpha}\right)\right\}.$$

where

$$\mu_i = \mathbb{E}_i[\log\{f_1(x_1)/f_0(x_0)\}] \qquad (i = 0, 1)$$

It can be shown, by using Wald likelihood ratio and Wald identity (See Siegmund (1985) pp. 21), that the expected sample sizes are minimal.

3.3 Brownian approximations and truncated tests

3.3.1 Sequential tests for the mean of a Brownian motion

A Brownian motion with drift μ and unit variance is a family of random variables $\{W(t),\ 0 \leq t \leq \infty\}$ with the following:

1. $W(0) = 0$
2. $P_\mu\{W(t) - W(s) \leq x\} = \Phi[(x - \mu(t-s))/(t-s)^{1/2}] \quad 0 \leq s < t < \infty$
3. for all $0 \leq s_1 < t_1 < s_2 \leq \cdots \leq s_n < t_n < \infty,\ n = 2, 3, \ldots$ the random variables $W(t_i) - W(s_i), \quad i = 1, 2, \ldots, n$ are stochastically independent
4. $W(t),\ 0 \leq t < \infty$ is a continuous function of t.

The standard normal cumulative distribution function is denoted by Φ. When the Brownian motion with drift μ and variance σ^2 per unit time is considered, (2) in the list is replaced by

$$P_{\mu,\sigma}\{W(t)-W(s) \leq x\} = \Phi[(x-\mu(t-s))/\sigma(t-s)^{1/2}],\ 0 \leq s < t < \infty$$

.

Proposition 3.2. *For any $-\infty < \mu_i < \infty$ $(i = 0, 1)$, and $t > 0$, the likelihood ratio of $\{W(s), s \leq t\}$ under P_{μ_0} relative to P_{μ_1} is*

$$l(t, W(t); \mu_0, \mu_1) = \exp\left[(\mu_0 - \mu_1)W(t) - \frac{t}{2}(\mu_0^2 - \mu_1^2)\right].$$

Proposition 3.3. *(Likelihood ratio identity). For any $-\infty < \mu_i < \infty$ $(i = 0, 1)$, stopping rule T and $Y \in \mathcal{E}_T$ (the class of random variables prior to T),*

$$\mathbb{E}_{\mu_0}(Y; T < \infty) = \mathbb{E}_{\mu_1}\{Y l(T, W(T); \mu_0, \mu_1); T < \infty\};$$

in particular if $Y = I_A \in \mathcal{E}_T$

$$P_{\mu_0}(A\{T < \infty\}) = \mathbb{E}_{\mu_1}[l(T, W(T); \mu_0, \mu_1); A\{T < \infty\}],$$

with l *given in Proposition 3.2.*

Proposition 3.4. *(Wald's identities). For any stopping rule* T *with* $\mathbb{E}_\mu(T) < \infty$,

$$\mathbb{E}_\mu W(T) = \mu \mathbb{E}_\mu(T)$$

and

$$\mathbb{E}_\mu[(W(T) - \mu T)^2] = \mathbb{E}_\mu(T).$$

Propositions 3.3 and 3.4 are special cases of the discrete time results presented in Section 3.2.

3.3.2 Sequential probability ratio test for the drift of Brownian motion

Let $\{W(t),\ 0 \le t \le \infty\}$ be a Brownian motion with drift μ. For a sequential probability ratio test of $H_0 : \mu = \mu_0$ and $H_1 : \mu = \mu_1$ with constants $A < 1 < B$, the stopping rule is given by

$$T = \inf\{t : l(t, W(t); \mu_1, \mu_0) \notin (A, B)\} \tag{3.17}$$

with decision rule of rejecting H_0 if $l(t, W(t); \mu_1, \mu_0) \ge B$. Due to the fact that $W(t), 0 \le t < \infty$ is continuous and $l(T, W(T); \mu_1, \mu_0)$ equals A or B with probability one and the approximations given in equations (3.8) and (3.9) become equalities. In what follows it is assumed that $\mu_0 < \mu_1$, also the following notation is used: $b = (\mu_1 - \mu_0)^{-1} \log B$, $a = (\mu_1 - \mu_0)^{-1} \log A$ and $\theta = \mu - \frac{1}{2}(\mu_1 + \mu_0)$. The stopping rule (3.17) can be written as

$$T = \inf\{t : W(t) - \frac{1}{2}(\mu_1 + \mu_0)t \notin (a, b)\} \quad (a < 0 < b) \tag{3.18}$$

and under P_μ, $W(t) - \frac{1}{2}(\mu_1 + \mu_0)t$ is a Brownian motion with drift θ.

Theorem 3.2. *(Siegmund, 1985) Let* T *be defined by (3.18) ad set* $\theta = \mu - \frac{1}{2}(\mu_1 + \mu_0)$. *Then*

$$\begin{aligned} P_\mu\{W(T) - \frac{1}{2}(\mu_1 + \mu_0)T = b\} &= \frac{1 - e^{-2a\theta}}{e^{-2b\theta} - e^{-2a\theta}} \qquad (\theta \neq 0) \\ &= \frac{|a|}{|a| + b} \qquad (\theta = 0) \end{aligned}$$

and

$$\begin{aligned}\mathbb{E}_\mu(T) &= [b(1-e^{-2a\theta})+a(e^{-2b\theta}-1)]/\theta(e^{-2b\theta}-e^{-2a\theta}) \quad (\theta \neq 0)\\ &= |a|b \quad (\theta = 0)\end{aligned}$$

For the special case $a = -b$ these become

$$\begin{aligned}P_\mu\{W(T) - \frac{1}{2}(\mu_1+\mu_0)T = b\} &= \frac{1}{1+e^{-2b\theta}}\\ &= 1/2\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}_\mu(T) &= \frac{b}{\theta}\left(\frac{1-e^{-2b\theta}}{1+e^{-2b\theta}}\right)\\ &= b^2\end{aligned}$$

For the stopping rule (T) of any sequential test of $H_0 : \mu = \mu_1$ vs $H_1 : \mu = \mu_1$, the lower bound for $\mathbb{E}_i(T)(i = 0, 1)$ in terms of error probabilities of Theorem 3.1 is unchanged in the continuous case. The relationships in Theorem 3.2 are exact and the sequential probability ratio test achieves the lower bounds on both $\mathbb{E}_{\mu_0}(T)$ and $\mathbb{E}_{\mu_1}(T)$ for error probabilities α and β (Siegmund, 1985).

3.3.3 Truncated sequential tests

Given the absence of an upper bound on the stopping rule, it is natural to consider the effect of truncation on the properties of the test.

In order to introduce a number of concepts in a simple framework, and to keep the discussion contained, in this section a class of one sided stopping rules is introduced.

Following (Siegmund, 1985), let $\{W(t), 0 \leq t < \infty\}$ be a Brownian motion with drift μ and consider the problem of testing $H_0 : \mu \leq \mu_0$ against $H_1 : \mu \geq \mu_0$. It is assumed that for μ much larger than μ_0, the collection of data is expensive which ideally requires a small sample size. It is also assumed that if $\mu \leq \mu_0$ data collection is relatively inexpensive and a large sample size is desirable.

For example in a clinical trial $W(t)$ measures the cumulative difference between patient's response to a new treatment or placebo. In this case $H_0 : \mu = 0$ denotes the null hypothesis of no difference between the new treatment and the placebo, while $H_1 : \mu > 0$ indicates the superior benefit of the new treatment. If $\mu > 0$ it would be advantageous to have a small sample size so more patients can benefit of the new treatment; from the patients point of view the trial could continue indefinitely if $\mu = 0$. A possible approach to this kind of problem is given by a test designed to perform like a sequential ratio test under H_1 and a fixed sample test under H_0.

Let $\mu_1 > \mu_0$ and for $B > 1$ the stopping rule is given by

$$\tau = \inf\{t : l(t, W(t); \mu_1, \mu_0) \geq B\}.$$

With notation $b = (\log B)/(\mu_1 - \mu_0)$ and $\eta = \frac{1}{2}(\mu_0 + \mu_1)$ the definition of τ becomes

$$\tau = \inf\{t : W(t) \geq b + \eta t\} \tag{3.19}$$

For $m > 0$ and $c < b + \eta m$ the following test is considered: $H_0 : \mu \leq \mu_0$ against $H_1 : \mu > \mu_0$. Sampling is stopped at

$$\tau \wedge m = \min(\tau, m),$$

reject the hypothesis H_0 if either $\tau \leq m$ or $\tau > m$ and $W(m) > c$, otherwise do not reject H_0. With c set to $b + \eta m$ the hypothesis H_0 is rejected only if $\tau \leq m$. The power of the test (Siegmund, 1985) is given by

$$\begin{aligned} &P_\mu\{\tau \leq m\} + P_\mu\{\tau > m, W(m) > c\} \\ &= P_\mu\{W(m) > c\} + P_\mu\{\tau < m, W(m) \leq c\}. \end{aligned} \tag{3.20}$$

The following conditional probability is used to compute (3.20) by unconditioning

$$P_\mu\{\tau < m \mid W(m) = \xi\} \quad (\xi < b + \eta m), \tag{3.21}$$

Additionally,

$$P_\mu\{W(m) \in d\xi\} = \Phi[(\xi - m\mu)/m^{1/2}]d\xi/m^{1/2}$$

we have

$$P_\mu\{\tau < m, W(m) \leq c\} \tag{3.22}$$
$$= \int_{-\infty}^{c} P_\mu\{\tau < m \mid \quad W(m) \quad = \xi\}\Phi[(\xi - m\mu)/m^{1/2}]d\xi/m^{1/2}.$$

The first passage distribution is computed by using the reflection principle.

3.4 Bayesian sequential analysis

3.4.1 Introduction

The idea lying at the core of Bayesian sequential analysis is that at every stage of the procedure the posterior Bayes risk of making an immediate decision is compared to the expected posterior risk that will be obtained if more observations are taken (Berger, 1985).

Observation consist in $X_1, X_2, \dots$ random variables. Let $\mathcal{M}_i$ be the sample space of X_i, define

$$\boldsymbol{X}^n = (X_1, X_2, \dots)$$

and assume that $\boldsymbol{X}^n$ has density $f(\boldsymbol{x}^n \mid \theta)$ and distribution function $F(\boldsymbol{x}^n \mid \theta)$ on $\mathcal{M}^n = \mathcal{M}_1 \times \cdots \times \mathcal{M}_n$. The unknown state of nature concerned that is the subject of inference is defined by $\theta \in \Theta$. The random variables X_i are independent and have common density $f(x \mid \theta)$. A sequential sample form such density is given

$$f_n(\boldsymbol{x}^n \mid \theta) = \Pi_{i=1}^n f(x_i \mid \theta). \tag{3.23}$$

The observations are taken sequentially and after observing any number of observations the experimenter can either make an immediate decision or take more observations. There are cost in taking observations, n denotes the number of observations, s denotes the way observations are taken (e.g. group, one at a time) and $a \in \mathcal{A}$ the action taken. The loss (or cost) when θ is the true state of nature is given by

$$L(\theta, n, s, a).$$

The sampling costs can be written as $C(n, s)$ and the sum of decision loss (with linear utility for the decision maker) is $L(\theta, a)$.

3.4.2 Bayesian sequential analysis - Notation

The prior density is $\pi(\theta) \in \Theta$ and X_0 and $\boldsymbol{X}^0$ imply that no observations have been taken. When observations are taken sequentially, with action $a \in \mathcal{A}$, the loss when θ is the true state of nature is denoted by $L(\theta, a, n)$. The loss is assumed to be increasing in n. A *sequential decision procedure* is denoted by

$$\boldsymbol{d} = (\boldsymbol{\tau}, \boldsymbol{\delta}).$$

The sequential decision procedure consist of two components (i) $\boldsymbol{\tau}$ is the *stopping time* and is made up of functions $\tau_0, \tau_1(\boldsymbol{x^1}), \tau_2(\boldsymbol{x^2}), \ldots$ indicating the probability of stopping sampling and make a decision after $\boldsymbol{x^i}$ has been observed (ii) $\boldsymbol{\delta}$ is the *decision rule* and consists of a series of decision functions $\delta_0, \delta(\boldsymbol{x^1}), \delta(\boldsymbol{x^2}), \ldots$, where $\delta_i(\boldsymbol{x^i})$ is the action to be taken if sampling has stopped after observing $\boldsymbol{x^i}$.

The stopping time is the random function of $\boldsymbol{X}$ given by

$$N(\boldsymbol{X}) = min\{n \geq 0 : \tau_n(\boldsymbol{X^n} = 1\}.$$

The *Bayes risk* of a sequential procedure $\boldsymbol{d}$ is defined as

$$r(\pi, \boldsymbol{d}) = \mathbb{E}^{\pi}[R(\theta, \boldsymbol{d})]$$

where $R(\theta, \boldsymbol{d})$ is the *risk function* $R(\theta, \boldsymbol{d}) = \mathbb{E}_{\theta}[L(\theta, \delta_N(\boldsymbol{X}^N), N)]$. A *Bayes sequential procedure* is a sequential procedure that minimises Bayes risk and is denoted by

$$\boldsymbol{d}^{\pi} = (\tau^{\pi}, \delta^{\pi}). \tag{3.24}$$

The Bayes risk of the problem is given by

$$r(\pi) = \inf_{d} r(\pi, \boldsymbol{d}) \tag{3.25}$$

The marginal density is given by

$$m_n(\boldsymbol{x^n}) = \mathbb{E}^{\pi}[f_n(\boldsymbol{x^n} \mid \theta)] = \int_{\theta} f_n(\boldsymbol{x^n} \mid \theta) d\boldsymbol{F}^{\pi}(\theta)$$

and for $m_n(\boldsymbol{x}^n) > 0$ the posterior densities are given by

$$\pi^n(\theta) = \pi(\theta \mid \boldsymbol{x}^n) = \frac{f_n(\boldsymbol{x}^n \mid \theta)\pi(\theta)}{m_n(\boldsymbol{x}^n)}$$

After each observation, there is a new sequential problem starting at that point. Once $\boldsymbol{x}^n$ observation has been observed, the new sequential problems is denoted $\mathcal{E}_n(\boldsymbol{x}^n)$.

Denote the Bayes risk of a procedure $\boldsymbol{d} \in \mathcal{D}^n$ by $r(\pi^n, \boldsymbol{d}, n)$. The objective is to find a procedure that gives minimum Bayes risk

$$r(\pi^n, n) = \inf_{\boldsymbol{d} \in \mathcal{D}^n} r(\pi^n, \boldsymbol{d}, n).$$

The quantity $\boldsymbol{r}(\pi^n, n)$ represents (conditional on $\boldsymbol{x}^n$) Bayes risk of proceeding in optimal fashion at stage n. The quantity $\boldsymbol{r}(\pi^n, n)$ represents the minimum Bayes risk that can be achieved by observing $\boldsymbol{x}^n$. Intuitively, in order to make an immediate decision, the value of $\boldsymbol{r}(\pi^n, n)$ is compared to the Bayesian risk of an immediate decision and continuing obverving data if $\boldsymbol{r}(\pi^n, n)$ is smaller. The posterior expected loss of action a at time n is denoted by

$$r_0(\pi^n, a, n) = \mathbb{E}^{\pi^n}[L(\theta, a, n)].$$

Denote the *posterior Bayesian risk* of an immediate decision at time n as

$$r_0(\pi^n, n) = \inf_{a \in \mathcal{A}} r_0(\pi^n, a, n).$$

With θ having prior π, X with conditional density $f(x \mid \theta)$ and marginal density $m^*(x)$, define for any function $g(x)$,

$$\mathbb{E}^*[g(X)] = \mathbb{E}^{m^*}[g(X)] = \mathbb{E}^{\pi}\mathbb{E}^X_{\theta}[g(X)]$$

where $\mathbb{E}$ is the expectation over X with respect to the marginal density of X.

3.4.3 The sequential probability ratio test (SPRT) as a Bayes procedure

The loss, in a sequential decision-theoretic approach is given by

$$L(\theta, a, n) = L(\theta, a) + nc$$

where $L(\theta, a)$ is "$0 - K_i$" loss. Letting a_i denote accepting H_i, the decision loss take the form of

$$L(\theta_0, a_0) = L(\theta_1, a_1) = 0$$

with

$$L(\theta_0, a_1) = K_1$$

and

$$L(\theta_1, a_1) = K_0.$$

The prior, for a parameter space $\Theta = \{\theta_0, \theta_1\}$, is specified as

$$\pi_0 = \pi(\theta_0) = 1 - \pi_1$$

A Bayesian procedure $\boldsymbol{d}^\pi = (\boldsymbol{\tau}^*, \boldsymbol{\delta}^\pi)$, where $\boldsymbol{\delta}^\pi$ is a Bayesian rule and $\boldsymbol{\tau}^*$ is the stopping rule which stops sampling for the first $n,\ (n = 1, 2, \dots)$ for which

$$\rho_0(\pi^n) = \rho^\infty(\pi^n) \tag{3.26}$$

The posterior Bayes decision risk is goven by $\rho_0(\pi^n)$, while $\rho^\infty(\pi^n)$ satisfies

$$\rho^\infty(\pi^n) = \min\{\rho_0(\pi^n), \mathbb{E}^*[\rho^\infty(\pi^n(\theta \mid X))] + c\} \tag{3.27}$$

the expectation is taken with respect to the marginal density of X induced by $\pi^n(\theta)$. The posterior distribution π^n, are determined by $\pi_0^n = \pi^n(\theta_0)$.

For any prior π, we have

$$p(\pi) = \inf_a \mathbb{E}^\pi[L(\theta, a)] = \min\{\pi_1 K_0, \pi_0 K_1\}. \tag{3.28}$$

The minimum Bayes risk among procedures taking at least one observations is given by

$$\rho^*(\pi) = \inf_{\boldsymbol{d}: N \geq 1} r(\pi, \boldsymbol{d})$$

Lemma 3.1. *(Berger, 1985) The function $p^*(\pi)$ is a concave function of π_0, and is equal to c when $\pi_0 = 1$ or π_0.*

Two cases are considered

1. $p_0(\pi) \leq p^*(\pi)$ for all π_0. When this is the case the Bayes procedure is to immediately make the Bayes decision, taking no observations

2. $p_0(\pi) > p^*(\pi)$ for some $\pi_0 (0 < \pi_0 < 1)$. In this case (see Berger (1985)) $p_0(\pi) > p^*(\pi)$ if $\pi_0' < \pi_0 < \pi_0''$. This leads to the following Theorem.

Theorem 3.3. *(Berger, 1985) The Bayes sequential procedure, $\boldsymbol{d}^\pi$, stops sampling for the first n $(n = 1, 2, \dots)$ for which $\pi_0^n \leq \pi_0'$ or $\pi_0^n \geq \pi_0''$, deciding a_0 if $\pi_0^n \geq \pi_0''$ and a_1 if $\pi_0^n \leq \pi_0'$. The constants $\pi' + 0$ and π_0'' satisfy $\pi' \leq K_0/(K_0 + K_1) \leq \pi_0''$.*

The Bayesian sequential procedure can be written by defining the *likelihood ratio* of θ_1 and $theta_0$ at stage n

$$L = \frac{\Pi_{i=1}^{n} f(x_i \mid \theta_1}{\Pi_{i=1}^{n} f(x_i \mid \theta_0}$$

and noting that

$$\begin{aligned}\pi_0^n = \pi(\theta_0 \mid \boldsymbol{x}^n) &= \frac{\pi(\theta_0)\Pi_{i=1}^{n} f(x_i \mid \theta_0}{\pi(\theta_0)\Pi_{i=1}^{n} f(x_i \mid \theta_0) + \pi(\theta_1)\Pi_{i=1}^{n} f(x_i \mid \theta_1)} \\ &= \frac{1}{1 + (\frac{\pi_1}{\pi_0})\Pi_{i=1}^{n}(f(x_i \mid \theta_1)/f(x_i \mid \theta_0)} \\ &= \frac{1}{1 + (\pi_1/\pi_0)}\end{aligned} \tag{3.29}$$

Corollary 3.1. *If $0 < \pi_0 < 1$, the the Bayes procedure , $\boldsymbol{d}^{\pi}$, is of the following form:*

$$\begin{aligned} &if L_n \leq A, \quad \textit{stop sampling and decide } a_0; \\ &if L_n \geq B, \quad \textit{stop sampling and decide } a_1; \\ &if A < L_n < B, \quad \textit{take another obsevation;}\end{aligned} \tag{3.30}$$

where $A = \pi_0(1 - \pi_0'')/(\pi_1\pi_0'')$ and $B = \pi_0(1 - \pi_0')/(\pi_1\pi_0)$.

In what follows, it is assumed that it is desirable to take at least one observation, equivalently this implies

$$\pi_0' < \pi_0 < \pi_0''. \tag{3.31}$$

When (3.31) is satisfied $A < 1$ and $B > 1$.

Definition 3.1. *(Berger, 1985) The procedure defined in (3.30) with constants $A < 1$ and $B > 1$ is called the sequential probability ratio test (SPRT) with stopping boundaries A and B, and is denoted $\boldsymbol{d}^{A,B}$.*

The Bayesian problem involves choosing $A < 1$ and $B > 1$ to minimize $r(\pi, \boldsymbol{d}^{A,B})$. Let N denote the stopping time of $\boldsymbol{d}^{A,B}$ as

$$N = \min\{n : L_n \leq A \textit{ or } L_n \geq B\},$$

and define probabilities of Type I and Type II error as

$$\alpha_0 = P_{\theta_0} \text{ deciding } a_1 = P_{\theta_0}(L_N \geq B),$$

$$\alpha_1 = P_{\theta_1} \text{ deciding } a_0 = P_{\theta_1}(L_N \leq A),$$

It is assumed that $P_{\theta_1}(N < \infty) = 1$ and $\mathbb{E}_{\theta_1} N < \infty$. Let $\mathbb{E}_{\theta_0}$ and $\mathbb{E}_{\theta_1} N$ denote the expected stopping times under θ_0 and θ_1 respectively. The Bayesian risk is then given by

$$\begin{aligned} r(\pi, \boldsymbol{d}^{A,B}) &= \pi(\theta_0) R(\theta_0, \boldsymbol{d}^{A,B}) + \pi(\theta_1) R(\theta_1, \boldsymbol{d}^{A,B}) \\ &= \pi_0[\alpha_0 K_1 + c\mathbb{E}_{\theta_0} N] + \pi_1[\alpha_1 K_0 + c\mathbb{E}_{\theta_1} N]. \end{aligned} \tag{3.32}$$

The problem reduces to the calculation of α_0,α_1,$\mathbb{E}_{\theta_0} N$ and $\mathbb{E}_{\theta_1} N$ and the minimisation of (3.32) over A and B.

Part II

A Quickest Detection Rule for HTA

Chapter 4

Introduction to Part II

4.1 Clinical trials

In clinical trials the term *endpoint* refers to the response variable (i.e. the outcome measure under study). These can take various forms, in Phase II trials for example, it is not uncommon to see binary response variable (success or fail). A *survival* endpoint refers to the time it takes for some event of interest to occur, this can be for example, the elapsed time before an event such as death, failure or relapse.

In conventional clinical trial designs the total sample size is determined in advance and a single final analysis is performed once data has been observed.

Two sided test for comparing two treatments with normal response and known variance

The difference in response of two treatments, assumed to be a normally distributed random variable with known variance, is denoted θ. The null hypothesis $H_0 : \theta = 0$ states that responses follow the same distribution under both treatments.

The alternative hypothesis, $H_1 : \theta \neq 0$ includes the case $\theta > 0$ and the case $\theta < 0$, corresponding respectively to one treatment to be superior and to one treatment to be inferior. In such case, the standardise test statistics Z is distributed symmetrically around the mean 0 under H_0 and a test rejects H_0 if $\mid Z \mid > c$ for some constant c. When H_0 is rejected the sign of Z indicates which treatment is to be preferred. The Type I error probability is given by

$$\alpha = Pr\{\mid Z \mid > c\} \mid (\theta = 0).$$

The Type I error probability is the probability of wrongly rejecting the null hypothesis. The power of the test for a value of θ is given by

$$Pr\{\mid Z \mid > c\} \mid (\theta = \delta) = Pr\{\mid Z \mid > c\} \mid (\theta = -\delta) = 1 - \beta,$$

and indicates the probability of rejecting the null when it does not hold. The parameter δ indicates the treatment difference. The Type II error is given by β and indicates the probability of failing to reject the null when it is not true.

Fixed sample Clinical Trials

The setup follows Jennison and Turnbull (2000). Let X_{A_i} and $X_{B_i}, i = 1, 2, \ldots$ denote the responses of subjects allocated between two treatments, A and B. Responses are assumed to be independent and normally distributed with $X_A \sim N(\mu_A, \sigma^2)$ for subjects receiving treatment A and $X_B \sim N(\mu_B, \sigma^2)$ for subjects receiving treatment B. Testing the null hypothesis of no treatment difference $H_0 : \mu_A = \mu_B$ against the two sided alternative $\mu_A \neq \mu_B$ with Type I error probability α and, power $1 - \beta$ at $\mu_A - \mu_B = \pm\delta$.

The standardised statistics, for n subjects allocated to each treatment is given by

$$Z = \frac{1}{\sqrt{2n\sigma^2}} \left(\sum_{i=1}^{n} X_{A_i} - \sum_{i=1}^{n} X_{B_i} \right) \sim N((\mu_A - \mu_B)\sqrt{\{n/(2\sigma^2)\}}, 1)$$

and $Z \sim N(0, 1)$ under H_0 and the symmetric two-sided test with Type I error probability α rejects H_0 if

$$\mid Z \mid > \Phi^{-1}(1 - \alpha/2),$$

where Φ denotes the standard normal cumulative distribution function. In order to satisfy the power requirement

$$Pr\{\mid Z \mid > \Phi^{-1}(1 - \alpha/2)\} = 1 - \beta,$$

when $Z \sim N(\pm\delta\sqrt{\{n/2\sigma^2\}}, 1)$ The number of subjects in each treatment arm is given by

$$n_f(\alpha, \beta, \delta, \sigma^2) = \{\Phi^{-1}(1 - \alpha/2) + \Phi^{-1}(1 - \beta)\}^2 2\sigma^2/\delta^2.$$

Fixed-sample designs, are easy to plan and implement but lack flexibility. For example the trial strictly complies to the pre-specified sample size regardless of unforeseen clinical results, such as better than expected or futility or harm (Jennison and Turnbull, 2000). In contrast sequential methods are flexible, regularly assessing data over regular intervals and monitoring possible stopping due to futility or harm while observing the clinical outcomes. For example with sequential methods, if a study shows overwhelming evidence in favour of a treatment, the trial can be stopped allowing patient in the arm with an inferior treatment to receive the more effective treatment (Pocock, 1977; O'Brien and Fleming, 1977).

Sequential methods in clinical trials

The modern theory of sequential analysis stems from the work of Wald (1947) and his sequential probability ratio test (SPRT). For observations that belong to a distribution whose probability density function is known apart from the parameter θ, in its basic form testing for a simple null hypothesis $H_0 : \theta = \theta_0$ versus a simple alternative $H_1 : \theta = \theta_1$ consists in observing successive observations until the likelihood ratio exits a certain interval (a, b) where the appropriate hypothesis is selected; otherwise the experiment is continued.

The SPRT procedure typically leads to lower sample sizes than fixed sample tests. Wald and Wolfowitz (1948) showed that the SPRT has the optimal property that, for pre-specified Type I and Type II error probabilities α and β, it reaches the smallest possible expected sample size or *average sample number* (ASM) when either H_1 or H_0 is true. The SPRT procedure is not bounded, leading to a distribution of sample size that is skewed and displays a large variance. In addition, the ASM can be very large in cases where θ is not equal to those pre specified in the hypotheses. These problems lead to consider curved boundaries for which the boundaries a and b are no longer constant but are a function of the sample size n and that ensure an upper limit n^* on the sample size. Aroian (1968) and Aroian and Robinson (1969) showed, using numerical integration, how to compute a truncated SPRT and ASM curves.

The SPRT is concerned with the problem of selecting one of two

competing hypotheses. The problem of selecting from more than two hypotheses is more challenging and has received less attention. An area that sees the application of such methods is communication theory, where problems include target detection in multiple resolution radar and infrared systems, signal acquisition and pattern recognition (Jennison and Turnbull, 2000). Baum and Veeravalli (1994) developed a procedure that generalises the SPRT to a multi-hypotheses problem.

Armitage (1954, 1958, 1975) and Bross (1952, 1958) implemented frequentist sequential methods in the medical field, with a particular focus on clinical trials. These methods did not find widespread use as continuous assessment of the study was impractical. In response, Cutler et al. (1966) proposed the use of *group sequential methods*. In contrast to fully sequential methods, group sequential tests involve analysing accumulated data at regular intervals rather than after each observation.

The major impetus for group sequential methods came with Pocock (1977), who demonstrated the flexibility of the approach and gave indications for the implementation of group sequential methods attaining Type I error and power requirements. Pocock (1977) analysed accumulating data in a repeated significance test at a nominal significance level. In the group sequential methodology, patients are divided in K equally sized groups containing m subjects on each treatment and the data is analysis at regular intervals after each new group of observations have been observed. In Pocock (1977) test, the critical value $C_p(K, \alpha)$ is chosen to give the overall Type I error probability

$$Pr_{\mu_A-\mu_B=0}\{\text{Reject } H_0 \text{ at analysis } k = 1, k = 2, \ldots, \text{or } k = K\} = \alpha$$

The test is given by

$$\begin{aligned}
&\text{After group} \quad k = 1, 2, \ldots, K-1 \\
&\quad \text{if } \mid Z_k \mid \geq C_p(K, \alpha) \quad \text{stop, reject } H_0 \\
&\quad \text{otherwise} \qquad \text{continue to group } k+1
\end{aligned}$$

$$\begin{aligned}
&\text{After group} \quad K \\
&\quad \text{if } \mid Z_k \mid \geq C_p(K, \alpha) \quad \text{stop, reject } H_0 \\
&\quad \text{otherwise} \qquad \text{stop, accept } H_0
\end{aligned}$$

In Pocock's design testing boundaries are equal throughout the trial. O'Brien and Fleming (1977) proposed a class of group sequential tests based on the truncated SPRT. This test produce conservative stopping boundaries at the early stage of the trial and give a decision rule similar to the fixed sample test in the last trial's stage. In the O'Brien and Fleming's design more stringent levels of significance are allocated at the beginning of the study and alleviates the significant levels towards the end of the trials (Yin, 2012). The O'Brien and Fleming (1977) test consists of

$$\begin{aligned}
&\text{After group} \quad k = 1, 2, \ldots, K-1 \\
&\quad \text{if } \mid Z_k \mid \geq C_p(K, \alpha)\sqrt{K/k} \quad \text{stop, reject } H_0 \\
&\quad \text{otherwise} \qquad \text{continue to group } k+1
\end{aligned}$$

$$\begin{aligned}
&\text{After group} \quad K \\
&\quad \text{if } \mid Z_k \mid \geq C_p(K, \alpha) \quad \text{stop, reject } H_0 \\
&\quad \text{otherwise} \qquad \text{stop, accept } H_0
\end{aligned}$$

Wang and Tsiatis (1987) proposed a family of two sided tests that can produce stopping boundaries of different shapes, including Pocock and Obrien & Fleming as special cases. The two-sided test is indexed by a parameter γ that produces boundaries of different shapes. The test with parameter γ rejects H_0 after group k if

$$\mid Z_k \mid \geq c_k = C_{WT}(K, \alpha, \gamma)(k/K)^{\gamma - 1/2}, \quad k = 1, \ldots, K.$$

The Wang and Tsiatis (1987) test consists of

After group $k = 1, 2, \ldots, K-1$
  if $\mid Z_k \mid \geq C_p(K, \alpha, \gamma)(k/K)^{\gamma - 1/2}$ stop, reject H_0
  otherwise continue to group $k+1$

After group K
  if $\mid Z_k \mid \geq C_p(K, \alpha, \gamma)$ stop, reject H_0
  otherwise stop, accept H_0

In general, the Wand Tsiatis method can produce bounds that give more stringent levels of significance to the end of the trial requiring a lower maximum sample size than the O'Brien & Fleming. With this test the experimenter can choose a suitable value of γ, balancing reductions in expected sample sizes against a high maximum sample size.

4.1.1 Bayesian approaches

The Neyman-Pearson procedure is the frequentist approach that has dominated biostatistics in the last 50 years. In the frequentist approach, parameters are regarded as fixed and not subject to probability distributions. The Bayesian approach, in contrast to the frequentist, considers the true value of a parameter as random variables to which one can assign a probability distribution. What is not known has a probability distribution assigned; what is known is taken as given and what is unknown is given a conditional probability based on the known values. Once the results of the experiment is known, the parameter quantities are taken as known and no longer subject probabilities (Berry, 2006).

The results of experiments are used to update probabilities of parameters: make an observation, update what is known. The updating process has great implications for trial design. The most interesting feature being the ability to quantify what will happen in a trial from any point onward, given the current result. No certain prediction can be made, however, the predictions are assessing

the future with the right amount of uncertainty. These predictive probabilities can be found in any trial and when applied to the final result are very useful in deciding the course of the trial. They are also important at the initial planning stage in assessing the value (in terms of utility or costs) of the trial's design (Berry, 2006).

Berry and Ho (1988) apply the Bayesian decision-theoretic approach of Raiffa and Schlaifer (1961) in which consequences of decisions are considered in terms of the company's assessment of related consequences. Of the many consequences arising from the objective of a clinical trial. Berry and Ho (1988) focus on those that can be measured by financial terms. The aim of the methodology is not to test for statistical hypothesis but rather to allow for the early termination of the clinical trial if the expected loss from pursuing the trial outweighs the expected return. If accumulating evidence shows benefits the trial continues, if the accumulated evidence is sufficiently negative the trial will end early with related saving of resources.

In the Bayesian paradigm, the parameter of interest θ is considered a random variable with a probability distribution. at the start of the trial, the distribution about θ is called a *prior* distribution with density $p_0(\theta)$. As evidence accumulate, this prior is updated of from the *posterior* distribution that summarises the current uncertainty about θ. Bayes law permits the updating and states that the posterior density is given by the normalised product of the prior density and the likelihood. This can be expressed as

$$p(\theta \mid D) \propto L(D \mid \theta) p_0(\theta)$$

where D stands for observed data. At any stage, the posterior distribution can be used to draw inference about θ. It is possible to construct a *credible* set for the parameter of interest with posterior probability equal to some level $1-2\epsilon$. For example, with $\epsilon = 0.025$, a Bayesian interval estimate (θ_L, θ_U) as in Lindley (1965) satisfies

$$Pr\{\theta_L < \theta < \theta_U \mid D) = 0.95.$$

As data is collected the standardised statistic is given by $Z_1, Z_2, \ldots$. At each stage $k = 1, 2, \ldots$ the likelihood for the parameter of interest θ os the normal density with mean θ and variance $\mathcal{I}_k^{-1}$ evaluated

at $\hat{\theta}^{(k)} = \frac{Z_k}{\sqrt{\mathcal{I}_k}}$. The posterior distribution for θ at stage k, for conjugate normal prior distribution $N(\mu_0, \sigma_0)$, is

$$N\left(\frac{\hat{\theta}^{(k)}\mathcal{I}_k + \mu_0\sigma_0^{-2}}{\mathcal{I}_k + \sigma_0^{-2}}, \frac{1}{\mathcal{I}_k + \sigma_0^{-2}}\right) \tag{4.1}$$

Following (Jennison and Turnbull, 2000) the 95% credible set for θ is given by

$$\left(\frac{\hat{\theta}^{(k)}\mathcal{I}_k + \mu_0\sigma_0^{-2}}{\mathcal{I}_k + \sigma_0^{-2}} \pm 1.96\frac{1}{\sqrt{\mathcal{I}_k + \sigma_0^{-2}}}\right).$$

It can be noted that he interval is reduced to the prior mean. In contrast, the fixed sample frequentist confidence intervals are centre around the estimate for the parameter $\hat{\theta}^{(k)}$.

Stopping rules

A stopping rule provides a mechanism for deciding whether to continue or stop an action. For example a stopping time is a random variable that takes value at the occurrence of a certain event. Bayesian inference on termination is easily derived, the problem of design and definition of stopping rules is not so straightforward (Jennison and Turnbull, 2000).

A possible stopping rule (that does not consider costs of utilities) is to stop, at some intermediate stage k, if

$$(i)\ Pr\{\theta < 0 \mid D_k\} < \epsilon \qquad or \qquad (ii)\ Pr\{\theta > 0 \mid D_k\} < \epsilon \tag{4.2}$$

where D_k stands for data observed till time k. A typical value for ϵ is 0.025. The rule is equivalent to stop when $1-\epsilon$ credible regions exclude zero. Berry (1985) follows this type of rule when comparing classical to Bayesian methods and Mehta and Cain (1984) recommended, in the context of Phase II clinical trials, monitoring on the basis of the posterior probability of exceeding a threshold and provide charts that specify the appropriate stopping rules.

The frequentist properties of the Bayesian procedure described in (4.2) show, when not controlled for, inflated Type I error rates. For example, by choosing a value of the prior of $\sigma_0 \to \infty$ in (4.1) and $\epsilon = 0.025$, (4.2) is equivalent to the repeated significance test that stops at the first k that $\mid Z_k \mid > 1.96$. (Jennison and Turnbull,

2000) show that in such case the Type I error rate becomes inflated with K, the maximum number of analysis. When $K \to \infty$ the error rate is equal to 1.00. Some procedures have been suggested to correct what has been termed 'sampling to a forgone conclusion' Cornfield (1966). (Cornfield, 1966) proposed to use a mixed prior, with such approach large number of analysis do not lead to the almost certain rejection of the hypothesis and there is no 'sampling to a forgone conclusion'. Berry (1987), by using data derived from a fully sequential trial on a new treatment for leukaemia that involved missing data points, argued that the Bayesian approach is much more flexible when compared to the classical method. Berger and Berry (1988) argue that the Bayesian approach is better suited to calculate, given some data, 'final probabilities' for hypothesis.

Cornfield (1966) design has not been adopted and current recommendation for clinical trial monitoring are among the lines of (4.2). The concern in the use of Bayesian procedures (with no cost or utilities) is due to the frequentist properties of Bayesian methods and the choice of prior. Pocock and Hughes (1989) argue that the Type I error is a useful tool when it comes to restrict the number of false positives in the medical literature. Regulatory agencies, such as the Food and Drug Association in the U.S.A., that have committees dedicated to reviewing clinical trial's evidence for drug approval, have similar concerns (Jennison and Turnbull, 2000).

Bayesian inference, by making use of posterior probabilities, does not depend on monitoring or stopping rules and consequently, there is no need for maximum sample sizes. Target sample size, useful in practice, can be derived using pre-posterior analyses using predictive distributions. Spiegelhalter and Freedman (1986) for the situation where the decision is based on the credible interval (θ_L, θ_U) and where the decision rule $\theta_L > 0$ recommends the new treatment, $\theta_U < 0$ recommends rejection of the new treatment and $\theta_L \leq 0 \leq \theta_U$ be non-committal. The conditional probability of concluding with recommendation of the new treatment on the basis that is preferable (i.e. $\theta > 0$) is given by

$$\frac{\int_0^\infty Pr\{\theta_L > 0 \mid \theta\} p_0(\theta) d\theta}{\int_0^\infty p_0(\theta) d\theta}$$

Similarly, it is possible to compute the conditional probability

for $\theta_U < 0$ and $\theta < 0$. A target sample size should be such that these probabilities are greater than some level, such as 90%.

Prior distributions

The choice of the prior distribution for the unknown parameters is very important in the context of interim monitoring because data dependent stopping can greatly increase the sensitivity of Bayesian credible intervals to misspecification of the prior (Rosenbaum and Rubin, 1984). While ideally a single defendable prior is desirable, due to the fact that a prior is subjective, it is unlikely that such a prior exists. A typical strategy involves bringing together a collection of analyses that are the result of a collection of priors (Jennison and Turnbull, 2000). A list of priors that might be considered, following Jennison and Turnbull (2000) is listed below.

(i) Clinical priors: represent expert's opinions. Priors are obtains by surveying clinicians knowledgable of the field. A problem with such approach is that chosen experts might be involved in the trials and consequently be optimistic about the new proposed therapy. An alternative approach, is to use meta-analysis studies. This method can also be over-optimistic due to publication biases.

(ii) Sceptical priors: should be selected so they represent an extreme of a range of options.

(iii) Enthusiastic priors: counterbalance a sceptical prior. an enthusiastic prior could be set by centring the distribution to the alternative allowing for a small probability that $\theta < 0$.

(iv) Non-informative priors: represents a lack of any prior opinion. This involves setting the prior probability of θ to lie in a interval around zero and equally in a very wide interval.

(v) 'Handicap' or 'pragmatic Bayes' priors: proposed as a way to control the frequentist properties of aBayesian monitoring procedure. The prior is chosen os that the false positive rate is controlled at a given level α, is there is no planned number of K maximum analysis. However, in the Bayesian paradigm, choosing a prior based on frequentist properties is paradoxical as the procedure does not maintain the properties of independent inference and independence from the sampling scheme Jennison and Turnbull (2000).

Discussion

While many Bayesian approaches have been proposed, the frequentist properties of false positive rates are considered to be very important. Additionally, the determination of sample sizes has been based on frequentist ideas of Type I probability and power of the test. Although Bayesian posterior probability can be very useful in practice, the requirement of inference based on a family of priors can be confusing Jennison and Turnbull (2000). Chevret (2012) reviewed over 300 articles related to Bayesian clinical trials, spamming the last 30 years period. The findings are that Bayesian methods are used more in the analysis than the design of clinical trials and are overall not used much in practice. The challenges associated with Bayesian methods in the frequentist domain, such as the possibly of an inflated Type I error rate or the choice of prior, are of concern to the regulatory agencies and, although Bayesian methods posses many desirable properties and have gained popularity, their use is still limited.

4.1.2 Bayesian decision theory and the value of information

Claxton (1999) proposed to use the Value of Information developed by Pratt et al. (1995) as a way to address how decision makers should interpret the results of probabilistic modelling and to address the question of whether enough evidence has been gathered. In this framework, the expected cost of uncertainty is determined jointly by the probability that a decision based on current evidence will be wrong and the consequences of a wrong decision.

The decision maker must choose between two interventions using only prior information. It is possible to minimise the expected opportunity loss taken as the difference in incremental net benefit between the best choice and the alternative actually chosen. The expected opportunity loss is the expected cost of the uncertainty surrounding the decision problem: this is the expected value of perfect information (EVPI).

With current information, decisions must be made before we know how uncertainties will resolve and the decision maker needs to undertake a decision based on the expected net benefit of each alternatives. If the decision maker has access to perfect information he/she could then undertake decisions for different resolutions of net benefit. The EVPI is the difference between the payoff with perfect and current information (Pratt et al., 1995; Sculpher and Claxton, 2005).

Expected Value of Perfect Information (EVPI)

As in Ades et al. (2004)' set-up, the decision model has unknown parameters θ and the choice is between different treatment j. $\text{NB}(j, \theta)$ is the net benefit of treatment $j = 1, 2, \ldots, J$ for parameters of value θ. The optimal decision, subject to current knowledge, is the one that provides the higher expected net benefit:

$$\max_j \mathbb{E}_\theta \text{NB}(j, \theta). \tag{4.3}$$

Maximising over the possible interventions j in not possible as the true values of θ are unknown. However, it is possible to obtain the expected net benefit of a decision taken with perfect information by averaging equation (6.1) over the joint distribution of θ:

$$\mathbb{E}_\theta \max_j \text{NB}(j, \theta). \tag{4.4}$$

The EVPI is the difference between equation (6.2) and (6.1) , amounting to the difference between the expected value of a decision made with perfect and current information:

$$\text{EVPI} = \mathbb{E}_\theta \max_j \text{NB}(j, \theta) - \max_j \mathbb{E}_\theta \text{NB}(j, \theta). \tag{4.5}$$

Expected Value of Sample Information (EVSI)

In order to establish if the conditions for further research are present and to identify efficient research design there is the need to also consider the expected costs of sample information. The expected value of sample information was introduced as a decision tool for clinical trial design by Claxton and Posnett (1996) and Ades et al. (2004).

The EVPI places an upper bound on returns to further research and provide a necessary but not sufficient condition for conducting further research. If the value of EVPI exceeds the cost of further research it might be worthwhile to gather more information about the problem as a whole or on selected parameters. However, in order to establish if further research will be worthwhile (i.e. net benefits of research are positive) and to identify efficient research design there is the need to consider the marginal benefits and marginal costs of sample information.

Technically efficient research design

The value of information analysis can be extended in order to find the expected value of sample information for particular research design (Ades et al., 2004).

A sample of size n on θ will give a sample result D. If the sample result were known, it would be possible for the decision maker to choose the alternative with the maximum expected payoff. It is possible to compute the expected net-benefit by averaging over the posterior distribution of the net-benefit of each intervention j given the sample result D:

$$\max_j \mathbb{E}_{\theta|D} \text{NB}(j, \theta). \tag{4.6}$$

As the value of D is not known in advance (i.e. the result of the sample is not known), the expected value of a decision taken with sample information is computed by averaging the maximum expected net benefits over the distribution of possible values of D. In other words this amount to compute the expectation over the predictive distribution of the sample results D conditional on θ, averaged over the prior distribution of θ:

$$\mathbb{E}_D \max_j \mathbb{E}_{\theta|D} \text{NB}(j, \theta). \tag{4.7}$$

The EVSI is the difference between the expected value of a decision made with sample information and the expected value with current information:

$$\text{EVSI} = \mathbb{E}_D \max_j \mathbb{E}_{\theta|D} \text{NB}(j, \theta) - \max_j \mathbb{E}_\theta \text{NB}(j, \theta). \tag{4.8}$$

The EVSI proposed in (6.6) is for a single study design and single sample size. In order to establish the optimal sample size for a particular study these computations needs to be repeated for various sample sizes n.

The difference between the EVSI and the cost of acquiring sample information C_s is the expected net benefit of sample information (ENBS). The payoff of research given by

$$\text{ENBS} = \text{EVSI} - C_s \tag{4.9}$$

Where C_s is the cost of obtaining a sample of size n. The optimal value of n is the one that generates the maximum ENBS. (Claxton and Posnett, 1996) state the core of this problem as

$$\frac{\partial \text{EVSI}}{\partial n} = C_m \tag{4.10}$$

where C_m is the marginal cost of sampling. In other words the optimal sample size is the one that satisfies the condition of marginal benefit of sampling to be equal to the marginal sampling cost MB = MC.

Claxton and Thompson (2001) propose to apply a Bayesian decision theoretic approach to the value of information in order to

address i) which clinical decisions are worth addressing through clinical research ii) if a clinical decision problem is worth evaluating, which of the competing alternatives should be considered relevant iii) what is the optimal scale of the prospective research iv) what is the optimal allocation of trial entrants and v) what is the value of the proposed research. The approach found in Claxton and Posnett (1996) and Claxton (1999) are generalised to the analysis of a sequential clinical decision problem. Estimates of the expected net benefit of sample information are used in a dynamic program to establish optimal allocation of trail entrants, optimal sample size, technically efficient research design and the expected net benefit derived from proposed research.

4.1.3 Real option and Investment decisions

A health-care system's objective is to maximise health gains from available resources and a decision of adopting (or rejecting) a technology should be based on its costs, health outcomes and the cost effectiveness threshold. While in the absence of irreversibilities (Palmer and Smith, 2000) or any costs associated with reversing a decision (Eckermann and Willan, 2008), the decision to adopt a technology can be based on expected cost effectiveness, the explicit inclusion of a sunk investment cost becomes important under uncertainty. The implication is that the decision makers need to be reassured that the selected policy is sustainable, as reversing the decision involves an economic cost.

Palmer and Smith (2000) apply Dixit and Pindyck (1994) real option approach to the adjustment (under a certain degree of irreversibility and uncertainty) of the incremental cost-effectiveness ratio for a drug. The conclusion is that for innovations with high uncertainty, large reversal cost and low opportunity costs of delay should be reimbursed at a lower rate than treatments with opposite characteristics.

Forster and Pertile (2012) illustrated through a combined real-option and decision-theoretic approach to HTA that view adoption, treatment and research decision as a single economic project that existing models found in the HTA literature consider only some of the dimensions relevant to optimal decisions, thus leading to

potential efficiency losses in resources allocation.

They refer to the literature on dynamic stochastic optimisation of Dixit and Pindyck (1994) to define two important dimensions that give value to an option: irreversibility and flexibility. The former refers to an action taken at a time t that cannot be undone at $t+1$ and the latter to an action available at time t that can be postponed to time $t+1$.

By using a two-period framework, Forster and Pertile (2012) explore four scenarios under either the present (absence) of irreversibilities and presence (absence) of the flexibility of deferring adoption/abandonment to a later time. Their framework shows that when a technology is costlessly reversible and treatment to patient cannot be deferred, then it is best to treat patients with the treatment that has the highest expected benefit. On the other hand, if there are irreversibilities, adopting on the basis of current evidence might 'lock' patients in a treatment with an inferior treatment. With the present of an option for treatment of patients, it is optimal to post-pone adoption and take advantage of the new evidence as it becomes available.

When adoption treatment and research decisions are viewed as a single economic project, the optimal rule must account for a number of dimensions such as i) the expected costs and benefits of additional research ii) the size of the treatment population over the stages of the project iii) flexibility and irreversibility of actions iv) the dynamic nature of the decision process.

Eckermann and Willam (2007) consider the choices that decision makers face when the new intervention has positive but uncertain net benefits. In such case the decision maker can either i) adopt with no further research ii) adopt and undertake a trial (assuming the decision is reversible) and iii) delay the decision and undertake a trial. In the paper the trade-off between the value and cost of information is considered. By delaying the decision there is an expected opportunity cost for patients receiving the standard intervention outside the trial and when the intervention is adopted, due to reversal costs, the expected value of information is reduced. They suggest that the optimal strategy and trial design needs to consider both the opportunity cost of delay and the costs of rever-

sal.

Eckermann and William (2008) demonstrate that in the case of irreversible decisions, delaying adoption with research is preferred to adoption and no trial when the EVSI is greater than the expected cost of information and the expected opportunity cost of not treating patients with the new technology. When decisions are reversible, adopting the technology with trial becomes an optimal strategy. However, in this case the EVSI is lowered due to the lower probability of reversal being optimal and lower payoffs when reversal is optimal. Decision makers face joint research and reimbursement decisions and such choices should also account for cost of reversal and opportunity costs of delay.

McKenna and Claxton (2011) examine a range of research designs, including length of follow up, sample size, policy options available and opportunity costs related to research. In the paper the authors distinguish the impact that research decisions have on patients enrolled in the trial, not enrolled and at trial end once evidence is reported.

They evaluate adoption and research decisions and for the latter account for a range of opportunity costs attached conducting further research, both under 'approval with research' or 'only in research'. They argue that there is a trade-off between the expected net benefits of early approval and the expected gains if approval is withheld till research reports. If research is possible with approval, the presence of irrecoverable cost associated with this decision might lead to 'only in research' decision as this avoids the commitment of costs until the results of research are known. If the new technology is 'approved with research' and the research reveals that the technology is not as cost effective as initially expected, losses will have been incurred and will not be compensated by later gains. The impact of such decision on an investment time profile will be greater with high uncertainty. The paper highlights the tradeoff between the expected net benefits to current patients from being able to accessing a technology early and the future health benefits to patients that will be realised by withholding approval until new research evidence becomes available.

Claxton et al. (2012) outlines the key principles and assessment

required when considering 'adoption in research' or 'only in research'. Irrecoverable costs become more important (i) when the estimates of cost-effectiveness would alter if the decision would be revised sooner than anticipated (ii) with a higher probability of altering the decision earlier than expected and (iii) the size of the health terms as a proportion of the net health effects. The assessment of irrecoverable costs calls for considerations on how health effects and costs accrue overtime, giving a 'break-even' time point at which the costs are recovered. For example a large negative irrecoverable cost may imply a large negative health effect to be offset by future gains. If the break even time point is not so distant in the future it would indicate that irrecoverable costs are not significant and they would have little influence on the alternative policy payoffs. In contrast, if the break-even time point is in the distant future, the impact of irrecoverable costs must be considered along side with the likelihood that guidance will change.

Griffin et al. (2011) suggested that the decision to adopt or reimburse a technology may damage future prospects of further research being conducted. Under such circumstances there must be a formal assessment of the opportunity loss of immediate adoption (value of forgone research). The omission of the opportunity loss can lead to the adoption of technologies for which current evidence is insufficient.Decisions to adopt or reimburse a new technology based on expected cost-effectiveness can be justified when no irreversibilities are present if its approval has no effect on the prospect of acquiring further evidence that may be needed. The decision made must consider whether the benefits of immediate access to a technology exceed the value of the evidence that maybe forgone.

The investment option approach has been implemented as a watchful waiting regime for diseases with slow progression (Driffield and Smith, 2007). In this case the patient management strategy involves postponing curative treatment. The patient undergoes a period of close observation and periodic tests that monitor the progression of the disease.

The basic ideas is that, by deferring the treatment decision and monitoring the patient, more information is collected and this may obviate the use of expensive, irreversible or risky treatments.

The economic decision structure implies a 3-fold decision at each time point: treat, continue monitoring or discharge. The model is solved using backward induction, with treatment recommendation that follow from the current state. At very low levels of expected net benefits the patient is discharged, at high level the patient is treated immediately. One of the issues reported in the study is that the incorrect modelling of the options can introduce serious biases into treatment advice. However, the technique has the potential to demonstrate that watchful waiting is beneficial in areas where waiting has been traditionally viewed as detrimental.

Discussion

The methods discussed above can inform the decision to adopt or reimburse a technology based on current evidence, indicate if more evidence is needed to reach a decision and be used to identify efficient research design. The question of whether to adopt a technology, whether more evidence is required and how to design future research, depends on economic considerations and traditional approaches based on hypothesis testing that fail to take into account costs attached to decision errors prove not to be an adequate guide for health care technology assessment.

Decision theoretic contributions have considered the expected costs and benefits of additional research by implementing the EVSI approach. These contributions have been extended to account for irreversibility costs such as the forgone value of evidence (Griffin et al., 2011; Claxton et al., 2012; McKenna and Claxton, 2011; Eckermann and Willam, 2007; Eckermann and William, 2008). However, the EVSI is essentially static in nature and, in the value of information decision framework, the issue of dynamics and flexibility of decisions has not yet been yet fully dealt with.

Examples of real options that account for uncertainty, degrees of investment irreversibility and time-dynamics are implemented by Palmer and Smith (2000) and Driffield and Smith (2007). However, these real option models assume that the underlying variable follow a stochastic process and limitations are seen as the expected value of sample information cannot be directly incorporated. Although there are examples of real option approach in the HTA

literature, presently there has been no systematic use of these techniques (Meltzer and Smith, 2012).

A more recent approach that incorporates dynamics, irreversibility and flexibility has been put forward by Pertile *et al.* (2010). The authors extend Chernoff (1961) models of sequential sampling to a multi-period perspective that incorporates the dimensions stated above. The model derive optimal sequential sampling rules for technology adoption and research abandonment decisions. This is a considerable move towards a true dynamic approach in HTA. Sequential sampling can provide a more complex decision making framework that allows for efficiency gains in resource allocation and a better assessment of uncertainty in HTA.

It is now common for clinical trial evaluating medicines, medical devices and clinical procedures to incorporate economic analysis of these interventions. The growing number of clinical-economic trials reflect both interest in economic information for new healthcare technologies and the requirements that many countries have in terms of economic value and clinical efficacy. Trial based cost-effectiveness studies have great appeal because of their high internal validity. It has recently been argued that by improving the quality and uniformity of these studies will be valuable to decision makers who need to consider economic value along with clinical efficacy (Ramsey et al., 2005). A dynamic sequential approach to HTA is likely to be very valuable in this context.

Chapter 5

On a simple quickest detection rule for health-care technology assessment

5.1 Introduction

Health technology assessment (HTA) decisions are based on evidence of relative costs and effectiveness of alternative interventions. Decision makers, when evidence suggests that the incremental net benefit of the new intervention is positive, are faced with the decision of whether to adopt the new intervention over the existing one or, given the uncertainty surrounding the evidence, wait for more information.

The explicit inclusion of a sunk investment cost is important as in the absence of such costs decision makers could switch between technologies as new evidence becomes available. The implication of uncertainty and cost associated with the investment is that the decision makers need to be sufficiently confident that the selected policy is sustainable, as reversing the decision involves an economic cost. The presence of uncertainty and the degree of irreversibility mean that there is economic value in employing a modelling approach that has flexibility in the timing of a decision (Palmer and Smith, 2000).

Forster and Pertile (2012) illustrated through a combined real-

option and decision-theoretic approach to HTA that view adoption, treatment and research decision as a single economic project that existing models found in the HTA literature consider only some of the dimensions relevant to optimal decisions, thus leading to potential efficiency losses in resources allocation. When adoption treatment and research decisions are viewed as a single economic project, the optimal rule must account for a number of dimensions such as i) the expected costs and benefits of additional research ii) the size of the treatment population over the stages of the project iii) flexibility and irreversibility of actions iv) the dynamic nature of the decision process.

More recently Pertile et al. (2013) discussed the use of real options as a way to view adoption, treatment and research decisions as a single economic project. Their approach follows in part from the financial option pricing literature and exploits Bayesian sequential analysis in order to update the beliefs as more evidence is gathered. One of the shortcoming of their approach is due to the requirement of a maximum experiment time N[1]. In their model, once the point N is reached the option expires and becomes worthless and consequently any decision must take either before or at time N. The implication is that the stopping bounds dependent on the maximum time N, influencing the investment/abandonment decision, and it is not clear how one should go about determining the maximum time N. Additionally, the model requires the construction of a computationally intensive grid of values and it's solved by backward induction. In contrast, real option models such as the ones proposed by Dixit and Pindyck (1994) involve solving a set of equations for some unknown values[2] and do not have computationally intensive requirements. While the framework shows the potential of such modelling procedures, presently the real option approach has not been implemented in any systematic way (Meltzer and Smith, 2012).

In this chapter we introduce a sequential value of information (S-

[1]This restriction is embedded in financial options models. In financial markets, option's expiry dates are know and at such time the option's payoff has either positive value or is worthless. American options allow investor to exercise an option before the expiry time and the optimal stopping is dictated by the optimal stopping bounds. This contrast with the real option modelling approach where investment can theoretically be delayed forever.

[2]It is often possible, and it is the case of the model proposed in this paper, to solve for the unknown values in a spreadsheet program such as Excel.

VoI) Bayesian model for the evaluation of health care technologies that allow to find an optimal stopping time at which the decision maker (i) knows that the value of further evidence is zero (i.e. zero value of waiting) and (ii) selects a strategy (either invest or abandon research) that gives maximal health benefit to patients. The S-VoI framework involves observing a trial and at each observation update a Bayesian posterior probability about the effectiveness of the healthcare technology. With only prior information the value of (further) information is a the maximum and it gradually reduces to zero as the trial continues.

In contrast to traditional approaches found in the literature the proposed framework introduces a dynamic sequential Bayesian approach to decision making under uncertainty when the objective of the decision maker (DM) is to maximise health benefit. The proposed model has a number of advantages over existing methodologies (i) by finding an optimal stopping time the decision is taken at the point where there in no value for further waiting (ii) error probabilities can be computed and the decision maker can assess the cost of error (iii) sample size is reduced to the minimum necessary in order to make a decision with minimal error (iv) the model incorporates a penalty for not using the best technology and (v) decision bounds are sample size independent. As a consequence the proposed method maximises expected gains both in terms of health to the population and minimised trial costs. Traditional decision tools such as the expected value of perfect information (EVPI) are based on ex-ante calculation and therefore consider only the deterministic time dimension. The proposed methodology improves decision making by enlarging the strategy space to stopping times.

In the paper clinical evidence is modelled as a noisy process: we start with a discrete binomial tree and, by allowing the number of observation within a time interval to increase, on the limit the random variable's distribution is obtained reflecting the uncertainty surrounding each clinical outcome. The methodology presented in the paper is based upon the work of Shiryaev (1978) and Peskir and Shiryaev (2006). However, there are some crucial difference between our approach and Peskir and Shiryaev's work: (i) we maximise health benefits (i.e. monetary payoffs) rather than minimising

a risk function which has Type I and Type II error probabilities as arguments (ii) we observe an arithmetic Brownian motion and define the likelihood ratio process as the Radon-Nikodyn derivative and while Peskir and Shiryaev (2006) solve the risk function via the posterior probability process, in our approach, given that the likelihood ratio process follows a geometric Brownian motion, it is possible to formulate the solution of the optimal stopping problem in terms of the likelihood ratio (iii) we depart from the traditional rules of statistical inference by incorporating a rate of discounting for the expected payoffs; in this way the optimal stopping problem fully incorporates the economic nature of decision making in HTA.

The paper is organised as follows: section two gives some background, section three deals with the probabilistic environment required for sequential hypothesis testing, section four specifies the decision problem while section five presents the solution to the optimal stopping problem. Section six discusses results implications for the value of information and the irrelevance of inference and section seven presents a case study comparing the model decision with a traditional decision making approach for robot-assisted laparoscopic prostaectomy.

5.2 Clinical trials

The decision maker wishes to test whether a newly produced health technology has effectiveness greater than the minimum required for reimbursement. The decision maker wishes to test if the newly developed health-care technology exceeds the health care system threshold value λ and sets up a a set of tests aimed at uncovering whether the new technology provides the increased effectiveness.

Within such scenario we observe a sequence of outcomes from a clinical trial. The trial evolves through time and at regular points we observe an outcome representing information about the effectiveness of the healthcare technology.

The outcome of a clinical trial is measured in terms of the cumulative benefit to the population over time and is denoted by X_i for each step i. We model the uncertainty of the trial's outcome by allowing X_i to go either up by a factor u or down by a factor d.

Trials evidence is noisy, which implies that trend in the sequence

of observed outcomes cannot be clearly observed. The two factors are given by

$$\begin{aligned} u &= \theta\mu dt + \sigma\sqrt{dt} \\ d &= \theta\mu dt - \sigma\sqrt{dt} \end{aligned} \tag{5.1}$$

where $\theta \in [0,1]$ and dt is obtained by splitting an interval $[0,t]$ into n parts (i.e. $dt = t/n$) and μ is the incremental effectiveness threshold per time period that warrants adoption.

Following the above we model the evolution of the health benefit as binomial tree. The random variable X_i can take values $X_{i-1}+u$ or $X_{i-1}+d$ with equal probability. The factor $\sigma\sqrt{dt}$ determines the size of the noise. The total accumulated evidence after n steps is equal to $X_n = \sum_{i=1}^{n} X_i$. The sequence $X_0, X_1, X_2, \ldots$ describes a stochastic process, where X_0 is the initial value.

Denote $X_t = \lim_{n\to\infty} X_n(t)$ where the limit is understood to be in distribution and $n \to \infty$ implies $dt \downarrow 0$. According to the CLT, the distribution of X_t exists and is given by

$$X_t \sim N(\theta\mu t, \sigma^2 t)$$

implying that in the continuous time limit the process X_t follows the arithmetic Brownian motion

$$X_t = \theta\mu t + \sigma W_t, \tag{5.2}$$

where $(W_t)_{t\geq 0}$ is a standard Brownian motion. Of course, this is only a heuristic argument providing pointwise convergence. In order to prove uniform convergence, a "functional" form of the CLT is required. A fairly straightforward application of Donsker's Invariance Principle (see, for example, Steele, 2001, Theorem 5.4) shows that, indeed, for all $x \in \mathbb{R}$,

$$\lim_{n\to\infty} P\left(X^{(n)}(\cdot) \leq x\right) = P\left(X(\cdot) \leq x\right).$$

The decision maker problem is to find an optimal time at which to make an investment/abandonment decision about the new technology. If the trial's outcome supports the hypothesis H_1 that the effectiveness of the new technology is greater than the health care system minimum requirement there is investment, else, under H_0

research is abandoned and there is no adoption. The problem is then to sequentially test for $H_0 : \theta = 0$ vs $H_1 : \theta = 1$.

5.3 Sequential hypothesis testing

The sequential testing problem of two hypotheses is discussed in Shiryaev (1978) and Peskir and Shiryaev (2006). As in their setup we assume that what follows takes place on a probability space $(\Omega, \mathcal{F}, Q_p)$ and that we are given mutually independent random variables $\theta = \theta(\omega)$ and a standard Wiener process $W = (W_t)_{t \geq 0}$ under the probability measure Q_p.

The probability measure Q_p follows a mixture distribution

$$Q_p = pQ_1 + (1 - p)Q_0 \tag{5.3}$$

for $p \in [0, 1]$.

Since we take a Bayesian viewpoint θ is considered a random variable taking the value of 1 or 0, and Q_p is such that $Q_p\{\theta = 1\} = p$ and $Q_p\{\theta = 0\} = 1 - p$. As outlined above, we observe a process $X = (X_t)_{t \geq 0}$ taking the form

$$X_t = \theta \mu t + \sigma W_t, \tag{5.4}$$

where $\mu > 0$ and $\sigma^2 > 0$ are given and fixed. The conditional distribution of X_t is

$$X_t \mid \ \theta \sim N(\mu \theta t, \sigma^2 t)$$

and thus p and $1 - p$ play the role of a priori probability for the statistical hypotheses

$$H_1 : \ \theta = 1 \quad \text{and} \quad H_0 : \ \theta = 0 \tag{5.5}$$

respectively.

The process X_t generates the filtration $\mathcal{F}_t^X = \sigma(X_s : 0 \leq s \leq t)$, which is augmented with the Q_p-null sets. The likelihood ratio process Λ_t is defined as the Radon-Nikodym derivative

$$\Lambda_t = \frac{d(Q_1 \mid \mathcal{F}_t^X)}{d(Q_0 \mid \mathcal{F}_t^X)} \tag{5.6}$$

Proposition 5.1. *The likelihood ratio process admits the following representation:*

$$\Lambda_t = \exp\left(\frac{\mu}{\sigma^2}\left(X_t - \frac{\mu}{2}t\right)\right), \quad t \geq 0 \tag{5.7}$$

Proof. See Appendix □

Note that under hypotheses H_1 and H_0 the corresponding probability measures are Q_1 and Q_0 respectively. These measures are mutually singular since it holds that

$$\Lambda_t \longrightarrow \begin{cases} 0, & \text{a.s. under } Q_0 \\ \infty, & \text{a.s. under } Q_1 \end{cases} \qquad \text{as } t \to \infty$$

In other words, if we can observe $(X_t)_{t\geq 0}$, as the trial continues the observer will learn the true state of nature and decide between the two hypothesis.

The likelihood ratio process can be expressed as a stochastic differential equation (SDE).

Proposition 5.2. *The likelihood ratio process* $(\Lambda_t)_{t\geq 0}$ *solves the stochastic differential equation*

$$d\Lambda_t = \frac{\mu}{\sigma}\Lambda_t dW_t \tag{5.8}$$

Thus the likelihood ratio Λ follows a geometric Brownian motion on the state space $E = [0, \infty)$. In addition the process $(\Lambda_t)_{t\geq 0}$ is a martingale.

Proof. See Appendix □

Peskir and Shiryaev (2006) express the *posterior probability process* $\pi_t = Q_p(\theta = 1 \mid \mathcal{F}_t^X)$ as a function of the likelihood ratio process using Bayes rule:

$$\pi_t(\Lambda) = \left(\frac{p}{1-p}\Lambda_t\right) \Big/ \left(1 + \frac{p}{1-p}\Lambda_t\right). \tag{5.9}$$

Therefore, we can also write the likelihood ratio process as a function of the prior and the posterior probability process

$$\Lambda_t \;=\; \frac{\pi_t}{1-\pi_t}\frac{1-p}{p}. \tag{5.10}$$

In the remainer we will work with $(\Lambda_t)_{t\geq 0}$ or $(\pi_t)_{t\geq 0}$ interchangeably.

5.4 Decision problem

The observed process $(X_t)_{t\geq 0}$ represents the outcome of the randomised clinical trial (RCT) in terms of cumulative health benefit and expresses the extent of effectiveness of the health care technology. The decision maker seeks to test if the new technology is more effective than the minimum required by the heath care system. The value μ represents the health benefit derived from adopting this new technology. If the new technology is more effective than the threshold λ specified by the health care system the decision maker will invest into this new technology.

The decision maker values payoffs in terms of Quality of Adjusted Life Years (QALY). This is a standard measure[3] for health benefit in health care technology assessments and allow to attach a monetary value to the benefits derived from adopting the technology, conditional on the technology being effective.

We seek to establish an optimal stopping time τ at which the decision takes an investment or abandonment decision about the health care technology. In the model adoption/abandonment decisions are based upon the present net monetary value of QALY gained/lost.

When undertaking the investment the decision maker incurs a sunk cost I. In the investment equation (5.11) below P_1 represents

[3]This is the standard in the UK. Other measures that are specific to the heath care technology can also be used. We use QALY to keep the analysis tractable in terms of monetary benefits/costs.

the monetary benefit from adopting the new health care technology conditional on $\theta = 1$ and P_0 represents the monetary loss of adopting the technology conditional on $\theta = 0$. Thus, $-P_0$ is the opportunity cost of adopting the new technology when this is in fact not better than the standard care, in effect making a type I error.

The table below summarises the various payoffs under investment and abandonment

	$\theta = 1$	$\theta = 0$
Investment	$P_1 > 0$	$P_0 > 0$
Abandonment	$-P_1$	0

The net present value of the investment, denoted by, F_I, is

$$\begin{aligned} F_I(\Lambda) &= \pi(\Lambda)P_1 - (1-\pi(\Lambda))P_0 - I \\ &= \left[\left(\frac{p}{1-p}\Lambda\right) \Big/ \left(1+\frac{p}{1-p}\Lambda\right)\right] P_1 \\ &\quad - \left[\left(1 - \left(\frac{p}{1-p}\Lambda\right) \Big/ \left(1+\frac{p}{1-p}\Lambda\right)\right)\right] P_0 - I, \end{aligned} \tag{5.11}$$

If research is abandoned there is no investment. In equation (5.12) below $-P_1$ describes the monetary loss incurred when research is abandoned conditional on $\theta = 1$, in effect making a type II error. It is assumed that forgone benefits and costs are the same. Therefore the expected payoff of abandoning, denoted by F_A is negative as it identifies the expected QALY loss due to keeping standard care when in fact the new health care technology is more effective. So,

$$F_A(\Lambda) = -\pi(\Lambda)P_1 = -\left[\left(\frac{p}{1-p}\Lambda\right) \Big/ \left(1+\frac{p}{1-p}\Lambda\right)\right] P_1 \tag{5.12}$$

Figure (5.1) shows F_I and F_A as function of the likelihood ratio. It can be noted that both payoffs are non-linear in Λ (even though

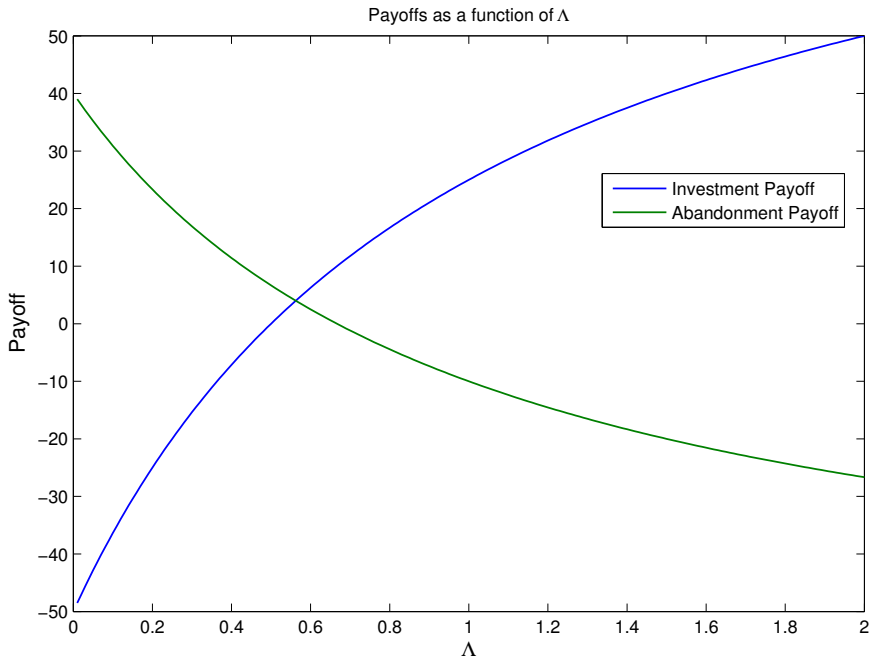

Figure 5.1: Payoffs of investment and abandonment as a function of the likelihood ratio, Λ

they are affine in the posterior probability, π). Additionally the function F_A is concave and the function F_I is convex.

Assuming that all payoffs and trial costs are discounted at a rate $r > 0$ the decision maker needs to find a stopping time τ^* that solves the following optimal stopping problem

$$\begin{aligned} F^*(\Lambda) &= \sup_\tau \mathbb{E}_\Lambda \left[-c \int_0^\tau e^{-rt} dt + e^{-r\tau} \left(\max \left[F_I(\Lambda_\tau), F_A(\Lambda_\tau) \right] \right) \right] \\ &= -\frac{c}{r} + \sup_\tau \mathbb{E}_\Lambda \left[e^{-r\tau} \left[\max \left(F_I(\Lambda_\tau), F_A(\Lambda_\tau) \right) \right] + e^{-r\tau} \frac{c}{r} \right] \\ &= -\frac{c}{r} + \sup_\tau \begin{cases} \mathbb{E}_\Lambda[e^{-r\tau}(F_I(\Lambda_\tau) + \frac{c}{r})] & \text{if } \Lambda_\tau \geq \bar{\Lambda} \\ \mathbb{E}_\Lambda[e^{-r\tau}(F_A(\Lambda_\tau) + \frac{c}{r})] & \text{if } \Lambda_\tau < \bar{\Lambda} \end{cases} \end{aligned} \tag{5.13}$$

where $\bar{\Lambda}$ is the unique point for which $F_I(\bar{\Lambda}) = F_A(\bar{\Lambda})$. The term c represents the cost stream connected to running the trial. This includes sampling costs and the forgone health benefits associated with allocating resources to the trial rather than treating patients. These costs are incurred up to the time at which a decision of investment or abandonment is made.

The solution to (5.13) can intuitively be thought of taking the

following form. The state space will be split in 3 regions. The first one is a region around $\bar{\Lambda}$ where continuation of the trial is optimal, hence called *continuation region*, denoted by

$$C = \{\Lambda \in \mathbb{R}_+ | F^*(\Lambda) > \max(F_A(\Lambda), F_I(\Lambda))\}.$$

When Λ gets large enough we enter the *investment region*, where adoption of the health-care technology is optimal. This region is denoted by

$$D_I = \{\Lambda \in \mathbb{R}_+ | F^*(\Lambda) = F_I(\Lambda)\}.$$

Conversely, when Λ gets low enough, we enter the *abandonment region*, where abandoning the clinical trial is optimal. This region is denoted by

$$D_A = \{\Lambda \in \mathbb{R}_+ | F^*(\Lambda) = F_A(\Lambda)\}.$$

5.5 Problem Solution

The likelihood ratio process $(\Lambda_t)_{t\geq 0}$ follows a geometric Brownian motion for which it is possible to find a solution to the optimal stopping problem (5.13). At the heart of the approach lie functions of the form

$$\varphi(\Lambda) = A\Lambda^{\beta_1} + B\Lambda^{\beta_2}, \tag{5.14}$$

which solve the differential equation

$$\mathcal{A}_\Lambda \varphi = r\varphi. \tag{5.15}$$

Here $\mathcal{A}$ denotes the *generator* (or *characteristic operator*) of the process $(\Lambda_t)_{t\geq 0}$,

$$\mathcal{A}_\Lambda f = \frac{1}{2}\frac{\mu^2}{\sigma^2}\frac{\partial^2 f}{\partial \Lambda^2}, \tag{5.16}$$

A and B are arbitrary constants (to be determined as part of the solution) and $\beta_1 > 1$ and $\beta_2 < 0$ are the roots of the quadratic equation

$$\mathcal{Q}(\beta) = \frac{1}{2}\frac{\mu^2}{\sigma^2}\beta(\beta - 1) - r = 0. \tag{5.17}$$

The following proposition gives sufficient conditions for the existence of a solution to the optimal stopping problem (5.13). For

each pair (Λ_A, Λ_I), $\Lambda_A < \bar{\Lambda} < \Lambda_I$, define the functions

$$\hat{\varphi}(\Lambda) = A\left(\Lambda^{\beta_1} - \Lambda_A^{\beta_1-\beta_2}\Lambda^{\beta_2}\right), \quad \text{and} \quad \check{\varphi}(\Lambda) = B\left(\Lambda^{\beta_2} - \Lambda_I^{\beta_2-\beta_1}\Lambda^{\beta_1}\right).$$

Then define the function φ by

$$\varphi(\Lambda) = \frac{\hat{\varphi}(\Lambda)}{\hat{\varphi}(\Lambda_I)}F_I(\Lambda_I) + \frac{\check{\varphi}(\Lambda)}{\check{\varphi}(\Lambda_A)}F_A(\Lambda_A).$$

It follows from Thijssen (2013, Proposition 6) that

$$\begin{aligned}\varphi(\Lambda) =& \mathbb{E}_\Lambda\left[e^{-r\hat{\tau}(\Lambda_I)}\middle|\hat{\tau}(\Lambda_I) < \check{\tau}(\Lambda_A)\right]Q_\Lambda(\hat{\tau}(\Lambda_I) < \check{\tau}(\Lambda_A))F_I(\Lambda_I)\\ &+\mathbb{E}_\Lambda\left[e^{-r\check{\tau}(\Lambda_A)}\middle|\hat{\tau}(\Lambda_I) > \check{\tau}(\Lambda_A)\right]Q_\Lambda(\hat{\tau}(\Lambda_I) > \check{\tau}(\Lambda_A))F_I(\Lambda_I),\end{aligned}$$

where

$$\hat{\tau}(\Lambda_I) = \inf\{t \geq 0 | \Lambda_t \geq \Lambda_I\},$$

is the first hitting time of Λ_I from below and

$$\check{\tau}(\Lambda_A) = \inf\{t \geq 0 | \Lambda_t \leq \Lambda_A\},$$

is the first hitting time of Λ_A from above.

So, if one defines the function

$$F(\Lambda) = \begin{cases} F_I(\Lambda) & \text{if } \Lambda \geq \bar{\Lambda} \\ F_A(\Lambda) & \text{if } \Lambda < \bar{\Lambda}, \end{cases}$$

and the stopping time $\tau^* = \hat{\tau}(\Lambda_I) \wedge \check{\tau}(\Lambda_A)$, then φ is simply the unconditional expectation of the present value of abandonment or investment, whichever threshold is reached first:

$$\varphi(\Lambda) = \mathbb{E}_\Lambda\left[e^{-r\tau^*}F(\Lambda_{\tau^*})\right].$$

Proposition 5.3. *Suppose that the system of equations*

$$-\frac{\hat{\varphi}'(\Lambda_I, \Lambda_A)}{\hat{\varphi}(\Lambda_I; \Lambda_A)}F_I(\Lambda_I) + F_I'(\Lambda_I) + \frac{\hat{\varphi}'(\Lambda_I, \Lambda_I)}{\hat{\varphi}(\Lambda_A; \Lambda_I)}F_A(\Lambda_A) \tag{5.18}$$

$$-\frac{\check{\varphi}'(\Lambda_A, \Lambda_I)}{\hat{\varphi}(\Lambda_A; \Lambda_I)}F_A(\Lambda_A) + F_A'(\Lambda_A) + \frac{\check{\varphi}'(\Lambda_A, \Lambda_A)}{\hat{\varphi}(\Lambda_I; \Lambda_A)}F_A(\Lambda_A) \tag{5.19}$$

has a solution (Λ_A, Λ_I), with $\Lambda_A < \bar{\Lambda} < \Lambda_I$. Suppose, in addition, that

1. *φ is strictly convex, and*
2. *φ is more convex than F_A on $(0, \bar{\Lambda})$, i.e. $F_A''/F_A' > \varphi''/\varphi'$ on $(0, \bar{\Lambda})$.*

Then the optimal stopping problem (5.13) *has the solution*

$$F^*(\Lambda) = \begin{cases} F_A(\Lambda) & \text{if } \Lambda \leq \Lambda_A \\ \frac{\hat{\varphi}(\Lambda)}{\hat{\varphi}(\Lambda_I)} F_I(\Lambda_I) + \frac{\check{\varphi}(\Lambda)}{\check{\varphi}(\Lambda_A)} F_A(\Lambda_A) & \text{if } \Lambda \in (\Lambda_A, \Lambda_I) \\ F_I(\Lambda) & \text{if } \Lambda \geq \Lambda_I, \end{cases} \tag{5.20}$$

and the optimal stopping time is $\tau^* = \hat{\tau}(\Lambda_I) \wedge \check{\tau}(\Lambda_A)$.

Proof. Note that

$$\mathcal{A}_\Lambda \hat{\varphi} - r\hat{\varphi} = \mathcal{A}_\Lambda \check{\varphi} - r\check{\varphi} = 0,$$

and that

$$\hat{\varphi}(\Lambda_A) = \check{\varphi}(\Lambda_I) = 0.$$

Also, since F_I is concave it is less convex than φ on $[\bar{\Lambda}, \infty)$. Therefore, the result follows immediately from Thijssen (2013, Proposition 7).

□

Figure 5.2 shows the solution for a case with cost of sampling equal to $c = 10$, a prior set to $p = 0.5$, discount rate of $r = 0.15$, payoff of investment $P_1 = 130$, Investment cost of $I = 60$ and losses from adoption when in fact the technology is not more effective than standard care of $P_0 = 60$. The process X_t has standard deviation $\sigma = 0.2$ and mean $\mu = 0.25$. For this base-case scenario it turns out that the conditions of Proposition 5.3 are satisfied for $p \in [0, 0.72]$. For higher values of p, F_A is more convex than φ, which implies that the value function F^* is no longer superharmonic. Since superharmonicity of the value function is a necessary condition for optimal stopping, no solution exists for high values of p. Essentially, for such values it is always optimal to adopt the technology immediately.

Figure 5.3 shows some simulated sample paths for the likelihood ratio process and some hypothetical bounds. Different values for μ and σ in the likelihood ratio process lead to different hitting times.

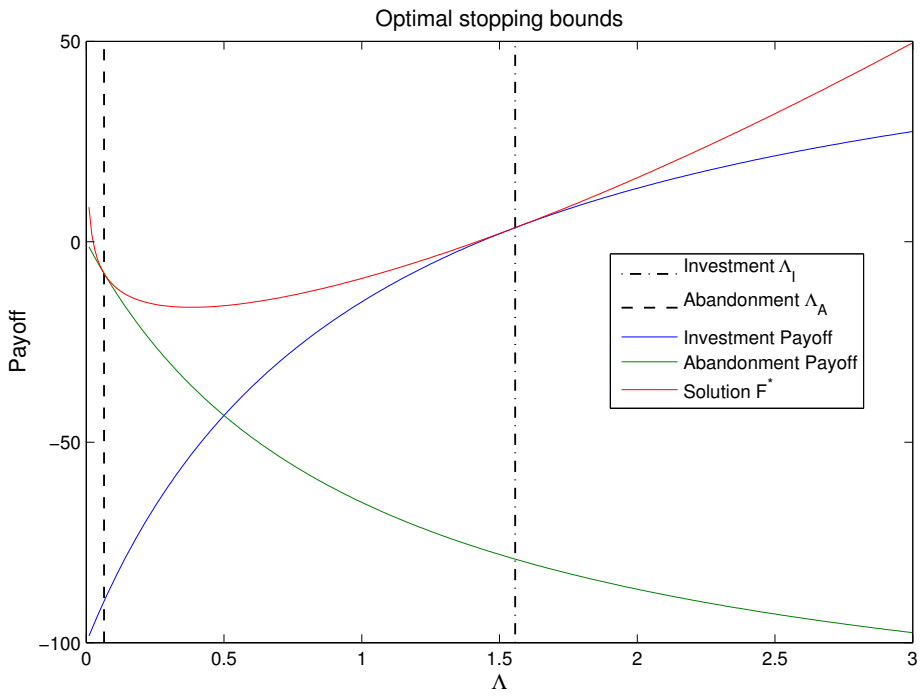

Figure 5.2: Value function F^* and bounds Λ_I, Λ_A for the case with $c = 10$, $p = .5$, $r = .15$, $P_1 = 130$, $P_0 = 60$, $I = 50$, $\mu = .25$, and $\sigma = .2$.

5.6 Analysis of the model

It has been argued (see Claxton (1999)) that classical statistical inference (and its Bayesian counterpart) is arbitrary and irrelevant to clinical decision making. He suggests to use the expected value of perfect information (EVPI) as a way to deal with uncertainty in health-care technology (HCT) assessment. The EVPI is given by the probability that a decision based on mean net benefit is incorrect (i.e. not cost effective) and the size of the opportunity loss of this wrong decision. It should be noted however, that the EVPI represents the *maximum* value of additional information (clinical research) and it is used to decide whether to fund more research. In particular, if the estimated costs of additional research (e.g. another trial) are higher than the EVPI, proposed research should not be undertaken and a decision for adoption by the health care system can be made on existing evidence.

This approach involves checking if sufficient information has

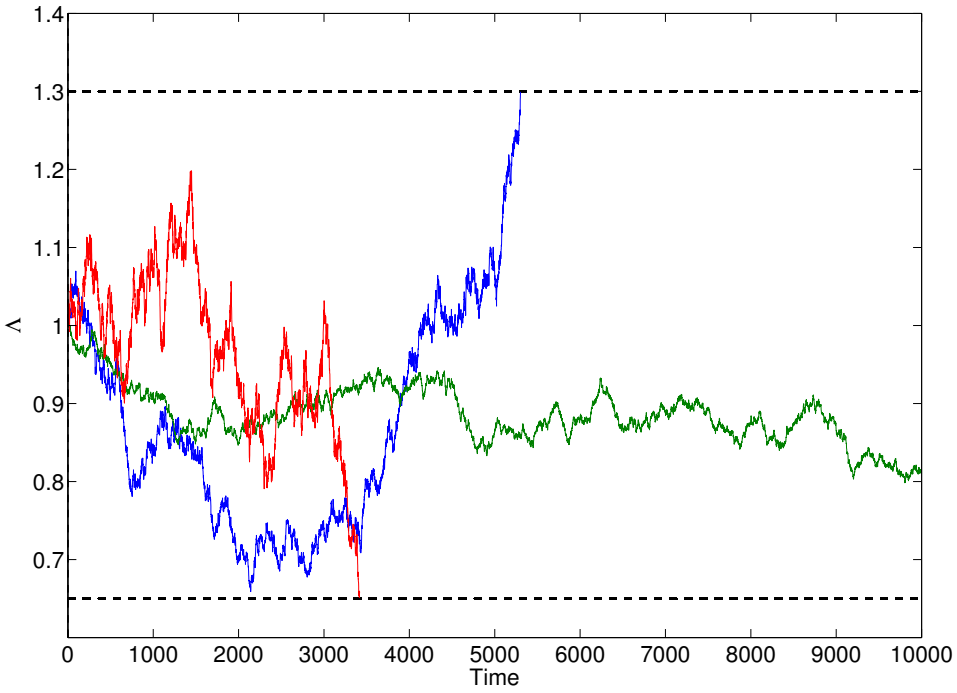

Figure 5.3: Some sample paths of the likelihood ratio process. Parameters are similar as for Figure 5.2 and bounds are fixed to 0.65 for the lower bound and 1.3 for the upper bound

been gathered and belongs to a framework where there is irreversibility of investment and where the decision maker is confronted with a 'invest now or never' type of decision (Pratt et al., 1995). Where reversing policy is costly and the decision maker has the possibility of deferring decision a sequential approach arises naturally.

5.6.1 Option value and waiting for more information

In between the thresholds the solution (5.20) gives the value of the investment / abandonment option *at any point* in the trial. When this value is compared to the investment/abandonment payoff, equation (5.20) reflects the value of waiting for more evidence (i.e. the value of information or the opportunity cost of investment with current evidence).

Figure 5.4 shows the function $\varphi(\Lambda)$ for different values of σ. It can be noted that the value of the investment option (i.e. the option

of investing now or investing later with more evidence) increases with σ. As uncertainty increases, there is more to be gained in waiting, and Figure 5.4 shows that it is possible to quantify the waiting value for different levels of uncertainty. Figure 5.2 shows the investment option value against the investment payoff $F_I(\Lambda)$ and $F_A(\Lambda)$. As the value of waiting for more evidence decreases, at the investment point Λ_I, the value of the investment option $\varphi(\Lambda)$ and the payoff $F_I(\Lambda)$ coincide and the value of waiting goes to zero. Similarly, on the other side, when the value of waiting for more evidence decreases, at the abandonment point Λ_A, the value of the abandonment option $\varphi(\Lambda)$ and the payoff $F_A(\Lambda)$ coincide and the value of waiting goes to zero.

Figure 5.5 shows the value of waiting (i.e. value of information) at different values for Λ. The value of information is at the highest around the initial point $\Lambda = 1$ as at this point the evidence in favour of H_0 and H_1 are equal as there is only prior information available. As Λ increases there is less and less value in waiting and this reaches zero at the optimal adoption/abandonment time τ^*. Outside of the threshold region waiting has no value and the decision maker should act immediately .

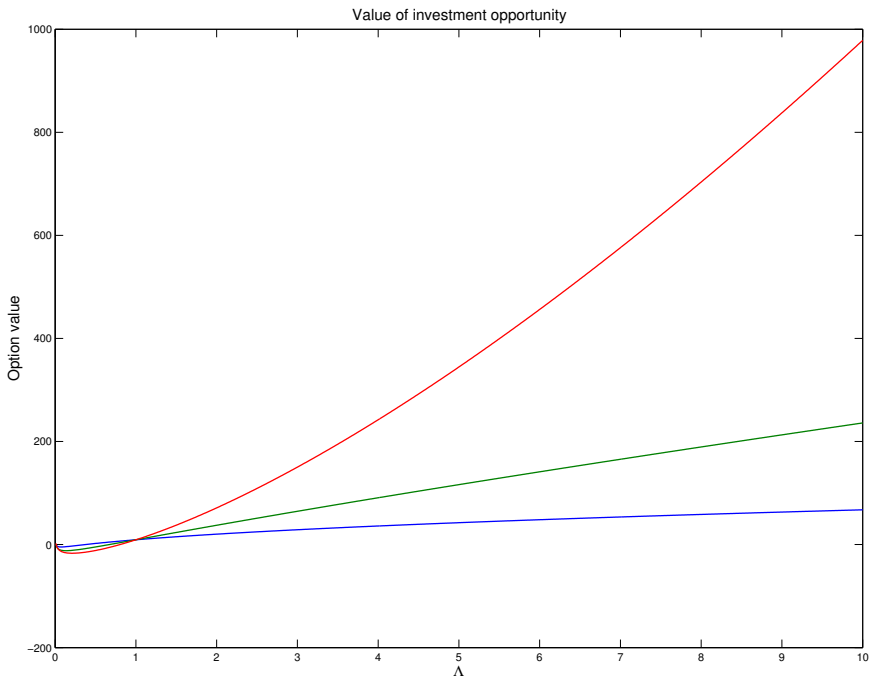

Figure 5.4: $\varphi(\Lambda)$ for different values of σ

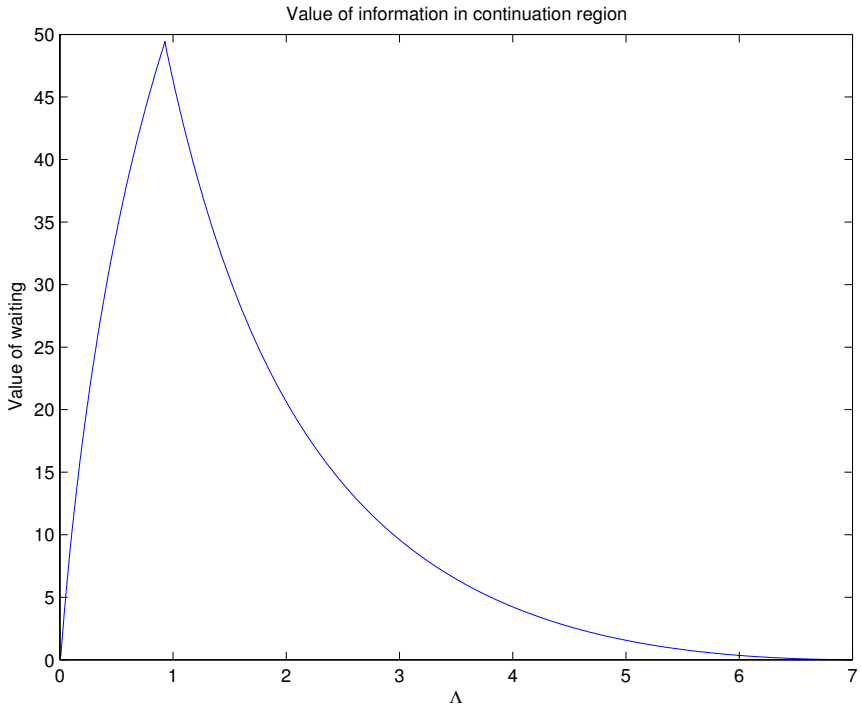

Figure 5.5: Value of information at different Λ

5.6.2 Posterior probability

In the health technology assessment literature one of the relevant decision tools is the probability of a drug being cost-effective (i.e. net benefit to be greater than the cost-effectiveness threshold).

While the standard approach is to compute the probability via simulation methods, in our proposed model the posterior probability $\pi(\Lambda)$ of making a gain of P_1 and the probability $(1-\pi)$ of making loss P_0, can be obtained by looking at the posterior value at the decision bound. In this way it is possible to assess the probability for the heath-care technology to provide a gain P_1 or a loss P_0, in turn allowing to determine the probability for the heath-care technology to be cost effective.

5.6.3 Comparative statics

It is possible to explore the impact of varying parameters on the decision bounds. In this section we explore the comparative statics of the payoffs P_1, P_0, parameters μ, σ, cost c and the discount rate

r. Figure 5.6a-f and Figure 5.7a-f show how the bonds vary due to a change in a single parameter. The prior has been set to a neutral value of 1/2 for all cases.

Figure (5.6a) shows the variation in bounds due to changing the payoff P_1. The payoff P_1 enters both the adoption and the abandonment payoff consequently affecting both upper and lower bounds. As the benefit from adoption increases the loss from not adopting a beneficial technology increases accordingly. It should be noted that as the payoff P_1 increases the upper bound eventually goes below the starting value for the likelihood ratio and the posterior process. As one would expect, holding P_0 and the required initial investment costs constant while increasing substantially the payoff P_1, due to the large gain to the healthcare system, above a certain threshold value it becomes optimal to invest immediately.

Figure (5.6b) shows the bounds variation due to changing the loss P_0. This loss enters the adoption payoff and thus affects only the upper bound. When the loss P_0 increases the adoption payoff decreases forcing the upper bound upwards to reflect the penalty brought in by a larger decision error. A large negative payoff to the healthcare system makes adoption more difficult, as one would expect.

Figure (5.6c) shows the bounds variation due to changing the drift μ. It should be noted that increasing μ and holding σ constant implies that the volatility of the likelihood process given by μ/σ increases. With a higher μ the trial becomes more informative relative to the noise component. This increases the value of waiting for more information. The non-monotoniticy in the expected time until a decision is taken arises because of two opposing forces: on one hand, we get more information per time period, leading to a decision being taken sooner, one the other hand, because the value of waiting increases, we want to make a decision later. It is not clear at priori what effect dominates.

Figure (5.6e) shows the bounds variation due to changing the volatility σ. As for μ, σ also determines the likelihood process volatility. The figure appear counterintuitive as on average, a decision is taken sooner (bands are narrower) for higher levels of noise. This is in contrast with the real option literature, where more uncer-

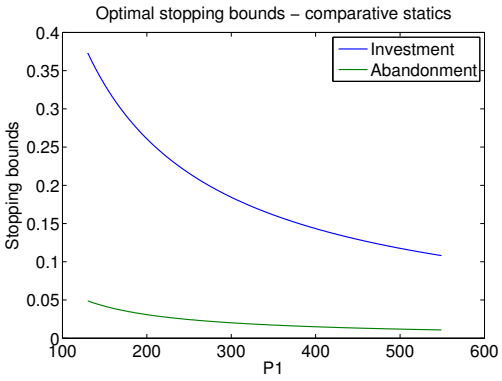

(a) Comparative statics for P_1

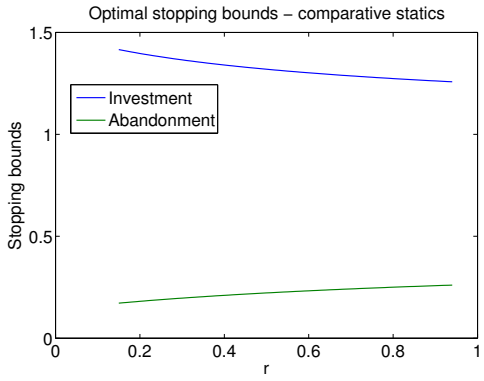

(d) Comparative statics for r

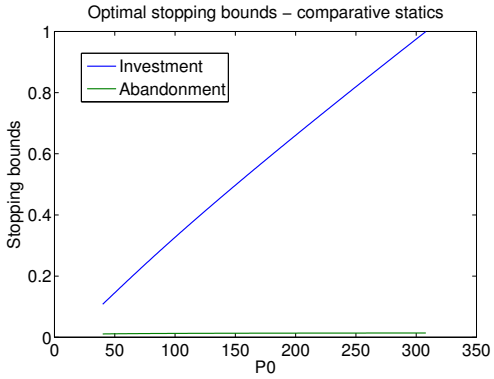

(b) Comparative statics for P_0

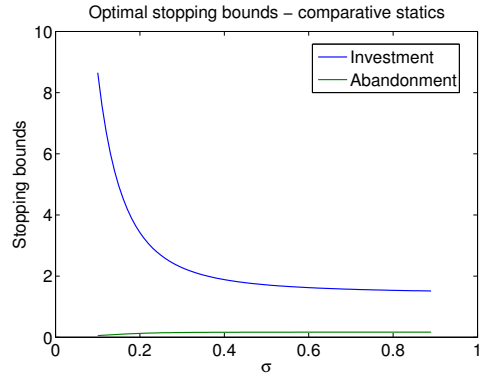

(e) Comparative statics for σ

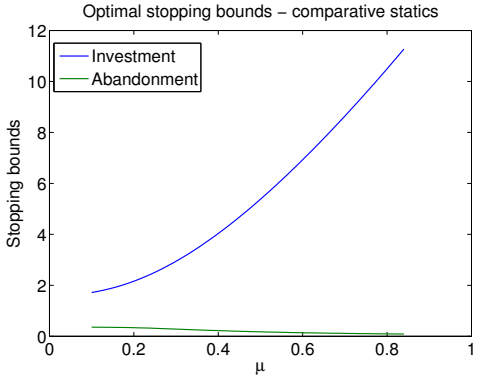

(c) Comparative statics for μ

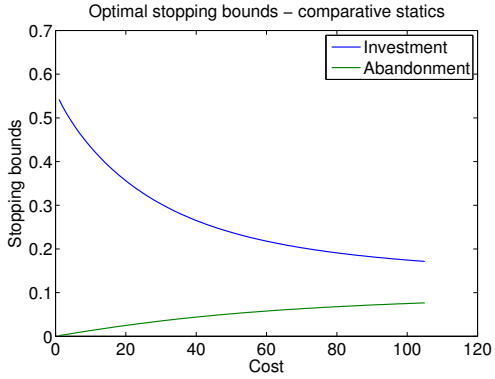

(f) Comparative statics for c

Figure 5.6: Bounds variation for parameter change in terms of the likelihood ratio.

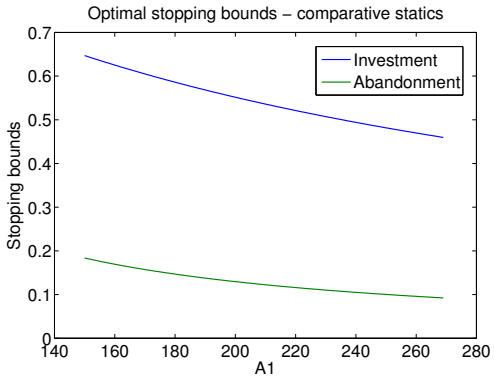

(a) Comparative statics for P_1

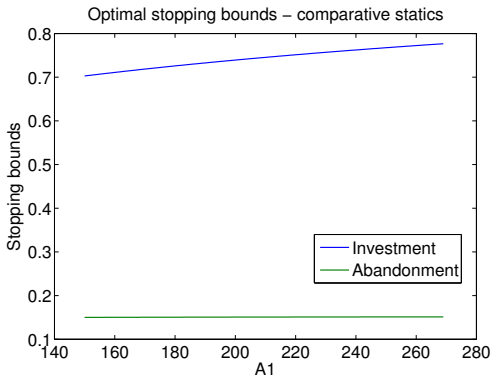

(b) Comparative statics for P_0

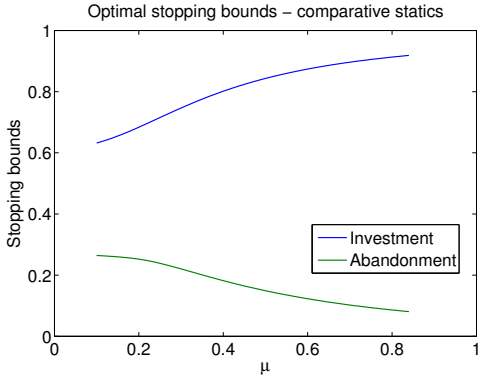

(c) Comparative statics for μ

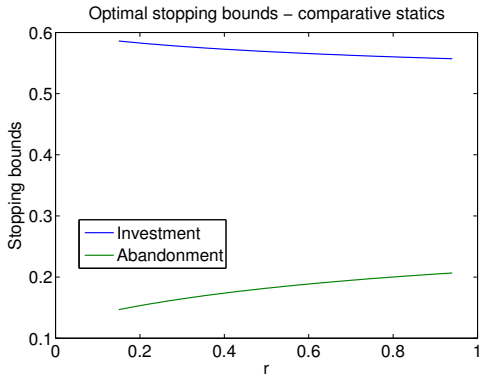

(d) Comparative statics for r

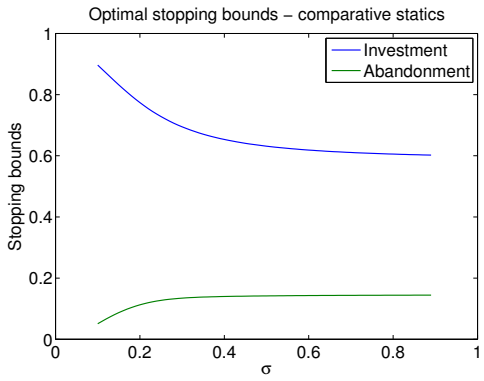

(e) Comparative statics for σ

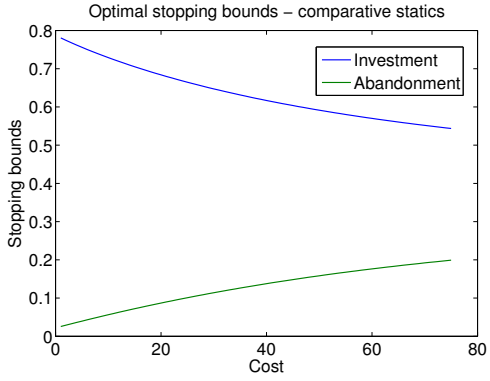

(f) Comparative statics for c

Figure 5.7: Bounds variation for parameter change in terms of the posterior (π).

tainty further delays the decision. The reason for taking a decision sooner is due to the fact that, as σ increases, a trial that provides less information is kept alive. This in turn means that waiting leads to less precise information which reduces the value of waiting. With an expensive trial, one might as well decide sooner. In such case, given that the trial is less informative, the cost of waiting do not outweigh the benefit of evidence.

Figure (5.6d) shows the bounds variation due to changing the discount rate r. The discount rate r enters the payoff functions and a high r decreases the present value of both the benefit and loss. Keeping all other parameters constant, increasing the discount rate r has the effect of correspondingly decreasing both upper and lower bounds. The discount rate affects project's present value and a high discount rate will decrease payoff values affecting decision bounds.

Figure (5.6f) shows the bounds variation due to changing the sampling costs c. Increasing the cost of sampling leads to narrower decision bounds. When the cost of conducting the trial are high the decision bounds become narrower forcing an earlier decision.

5.7 Some Probabilities

5.7.1 Probability of adoption/abandonment

We compute the probability of hitting the adoption or investment bound. The expected discount factor, under the posterior probability $Q_{\pi(\Lambda)}$ (below simply noted as Q_Λ) is given by

$$\begin{aligned}
\mathbb{E}_\Lambda[e^{-r\tau^*}] &= \mathbb{E}_\Lambda[e^{-r\check{\tau}(\Lambda_A)} \mid \tau^* = \check{\tau}(\Lambda_A)]Q_\Lambda(\tau^* = \tau(\Lambda_A)) \\
&+ \mathbb{E}_\Lambda[e^{-r\hat{\tau}(\Lambda_I)} \mid \tau^* = \hat{\tau}(\Lambda_I)]Q_\Lambda(\tau^* = \tau(\Lambda_I)). \\
&= \frac{\hat{\varphi}(\Lambda)}{\hat{\varphi}(\Lambda_I)} + \frac{\check{\varphi}(\Lambda)}{\check{\varphi}(\Lambda_A)} \qquad (5.21)
\end{aligned}$$

Using the fact that

$$Q_\Lambda(\tau^* = \check{\tau}(\Lambda_A)) = 1 - Q_\Lambda(\tau^* = \hat{\tau}(\Lambda_I)) \qquad (5.22)$$

and writing the discount factors in (5.21) as

$$\mathbb{E}_\Lambda[e^{-r\check{\tau}(\Lambda_A)} \mid \tau^* = \check{\tau}(\Lambda_A)] = \left(\frac{\Lambda}{\Lambda_A}\right)^{\beta_2}$$

and

$$\mathbb{E}_\Lambda[e^{-r\hat{\tau}(\Lambda_I)} \mid \tau^* = \hat{\tau}(\Lambda_I)] = \left(\frac{\Lambda}{\Lambda_I}\right)^{\beta_1}$$

we obtain

$$Q_\Lambda(\tau^* = \check{\tau}(\Lambda_A)) = \frac{\mathbb{E}_\Lambda[e^{-r\tau^*}] - \left(\frac{\Lambda}{\Lambda_I}\right)^{\beta_1}}{\left[\left(\frac{\Lambda}{\Lambda_A}\right)^{\beta_2} - \left(\frac{\Lambda}{\Lambda_I}\right)^{\beta_1}\right]}. \tag{5.23}$$

5.7.2 Error Probabilities

We can compute the *ex ante* probabilities that we make an erroneous decision. Following from Shiryaev (1978), the frequentist approach, in contrast to the Bayesian approach above, does not take into account benefits and costs. It is rather concerned with purely inferential concerns such as Type I (falsely rejecting the null) and Type II (falsely rejecting the null) errors.

Under such scenario, the basic problem seeks to find a pair (d, τ), where τ is the time of stopping, d is the decision rule $d : \Omega \to \{0, 1\}$ and accepting H_1 if $d = d_1$ or accepting H_0 if $d = d_0$ such that the probability error of the first and second kind satisfy:

$$Prob(\text{accept } H_1 \mid \text{true } H_0) \leq \alpha$$

$$Prob(\text{accept } H_0 \mid \text{true } H_1) \leq \beta$$

and the mean times of observation $\mathbb{E}_0\tau$ and $\mathbb{E}_1\tau$ are as small as possible. The "payoffs" assigned to these errors are given in Table 5.1.

		State of nature	
		$\{\Theta = 1\}$	$\{\Theta = 0\}$
Decision	Adoption	0	-1
	Abandonment	-1	0

Table 5.1: Payoff matrix of a health technology decision problem.

In the frequentist setting the decision maker wishes to control for these errors. Typical values are $\alpha = 0.05$ and $\beta = 0.2$ with $\alpha + \beta < 1$. The optimal stopping rule turns out to rely on the

likelihood ratio process $(\Lambda_t)_{t\geq 0}$. In fact, the optimal decision rule $(\hat{\tau}, \hat{\delta})$ is (Shiryaev, 1978, Theorem 4.6)

$$\hat{\tau} = \inf\left\{t \geq 0 | \Lambda_t \notin (\hat{A}, \hat{B})\right\}, \quad \text{and} \quad \hat{d} = \begin{cases} 1 & \text{if } \Lambda_t \geq \hat{A} \\ 0 & \text{if } \Lambda_t \leq \hat{B} \end{cases}, \tag{5.24}$$

where the frequentist bounds are given by

$$\hat{A} = \frac{\beta}{1-\alpha}, \quad \text{and} \quad \hat{B} = \frac{1-\beta}{\alpha}. \tag{5.25}$$

In other words, the inferential bounds of the frequentist approach do not change when the decision environment changes.

Finally, for abandonment and adoption bounds Λ_A and Λ_I, the error probabilities are defined as

$$\alpha = Q_0\left(\tau^* = \hat{\tau}(\Lambda_I)\right), \quad \text{and} \quad \beta = Q_1\left(\tau^* = \check{\tau}(\Lambda_A)\right).$$

Such error probabilities, as a function of the investment and adoption bounds and as a corollary of (5.25), are given by the following Wald approximations (See Poor and Hadjiliadis (2009))

$$\alpha = \frac{1-\Lambda_A}{\Lambda_I - \Lambda_A} \quad \text{and} \quad \beta = \Lambda_A \frac{1-\Lambda_A}{\Lambda_I - \Lambda_A} \tag{5.26}$$

5.7.3 Current and future population

The issue of population, such as the tradeoff between the expected net benefits to current patients from being able to accessing a technology early and the future health benefits to patients that will be realised by withholding approval (See McKenna and Claxton (2011)) can be dealt by separately modeling patients that will benefit from the trial in the future and current trial participants. In such case, a rescaled μ would measure the expected net benefit of the new technology in the trial's population while rescaled P_1 and P_0 would account for the population that would benefit from the treatment once the technology is adopted or abandoned.

Choosing the time horizon over which information about a decision problem can be of use is challenging and poses a number of question on it's assessment and integration in a decision model (Philips et al., 2008). The model presented above assumes that the technology once adopted is used forever. However, considerations of a time horizon will impact the cost-effectiveness if investment

costs are not recovered quickly enough. Many real-life investment are finite and expire or became valueless at some point in the future (e.g. patents). In practical applications, where the decision maker needs to limit the investment to a certain time frame and related population, it is necessary to multiply the costs and benefits by an exponential factor

$$1 - \exp(-rT)$$

where T is the life time factor.

5.8 Case study: standard vs robot-assisted laparoscopic prostaectomy

In this section we apply the model developed above to the HTA of robot-assisted and standard laparoscopic prostaectomy from the perspective of the UK national health service using data from a published study (Close et al., 2013). The application of the model developed above to this case study if for illustration purposes only and aims at showing the value that our approach can have for HTA.

Standard laparoscopic prostaectomy and robot-assisted laparoscopic prostaectomy are favoured over the open technique as these cause less bleeding and allow for a quicker return to activities. Robot assisted laparoscopic prostaectomy is increasingly used compared to standard laparoscopic technique. However, the high cost has led authorities to question the value of robotic-assisted procedure to patients and the health care system.

Many of the existing cost studies on prostaectomy techniques do not include cost effectiveness analysis that takes into account the value of relative gains that men achieve if a particular technique has better outcomes.

Close et al. (2013) conduct a cost-utility analysis for two indipendent cohort of 5000 men undertaking respectively robotic or laparoscopic prostaectomy over 10 years. They report that the use of robotic prostaectomy was on average £1412 more costly than laparoscopic prostaectomy and that it was also more effective with mean gain of QALY of 0.08 (95% CI,0.01-0.15) over 10 years for a case load of 200 patients per year. As we take the point of view of the UK health service, we seek to establish if robot prostaectomy

Table 5.2: Cost-effectiveness of standard vs robot-assisted laparoscopi prostaectomy

Parameter	**Description**	**Source**	**Value**
$E_1 - E_0$	Incremental QALY gain	Close *et al*	0.08 QALY
$C_1 - C_0$	Incremental cost	Close *et al*	£1412
σ	Std. deviation	Close *et al*	£1071
μ	Incremental QALY gain	Set as $NIMB > 0$	£1413
p	Prior	Assumed	0.5
r	Discount rate	Close *et al*	3.5%
c	Cost of sampling	Assumed	£10
I	Close *et al*	0	
n	Number of patients	Close *et al*	10000

is cost-effective at the UK NICE threshold of $\lambda = £30,000$ at such threshold value the mean gain is of £2400.

Confidence intervals give a standard deviation σ of £1071 indicating considerable uncertainty. We set the minimum required μ for adoption by the national health service to £1413, just greater than the incremental cost of the robot assisted surgery. In other words we set μ such that the net incremental mean benefit[4] is positive, ensuring a positive gain to the heath service if the technology is adopted. The adoption excess benefit P_1 is set for each patient at £2400 and the cost of wrongly adopting the technology P_0 is set equal to the incremental cost at £1412. The prior is set to a neutral value of $p = 0.5$, the discount rate is set to $r = 0.035$ and initial investment I is assumed to be zero. Having no information on the cost of following patients and reporting the outcome of the procedure, we assume sampling costs for each observation c to be £50. Parameters are summarised in Table 7.1.

Figure 5.8a shows the optimal stopping bounds obtained with such values. The upper bound is $\Lambda_I = 1.30$ and the lower bound $\Lambda_A = 0.0035$. Correspondingly these bounds in terms of the posterior are $\pi_I = 0.56$ and $\pi_A = 0.003$.

The value of the likelihood ratio at the point estimate is $\Lambda = 8.24$ and $\pi = 0.89$, much higher than the required adoption bounds. These value suggest that there is enough evidence to make a investment decision.

[4] $NIMB = (E_1 - E_0)\lambda - (C_1 - C_0)$

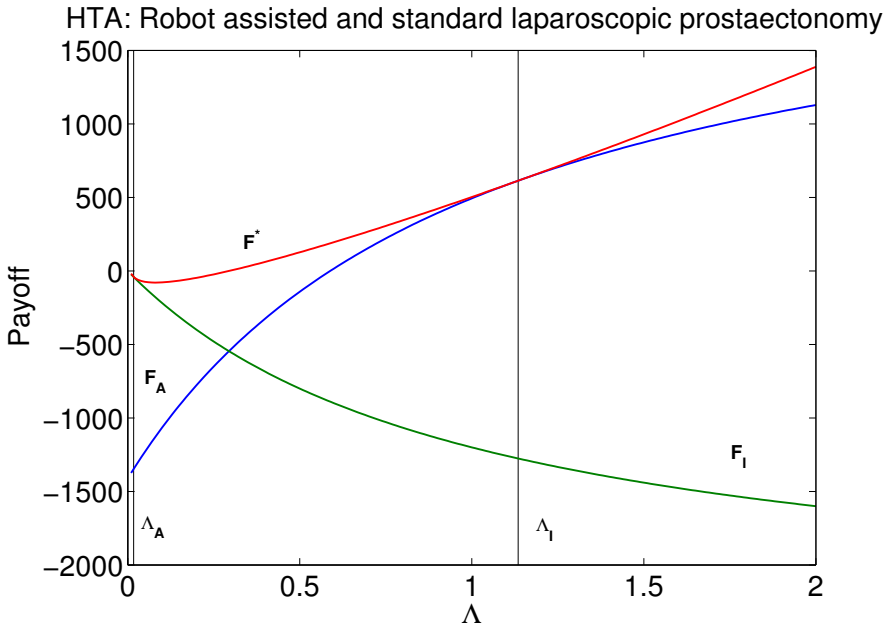

(a) Optimal stopping bounds (Λ)

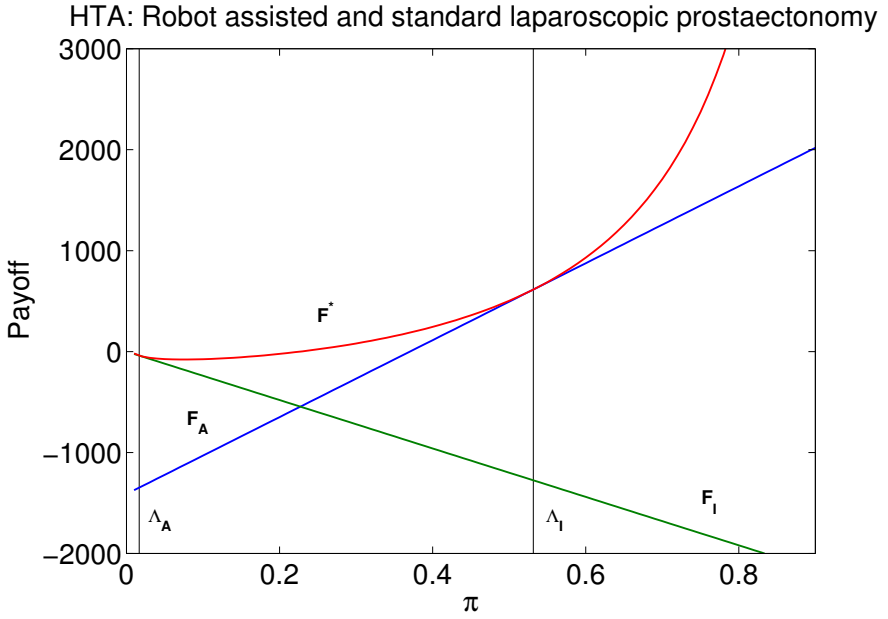

(b) Optimal stopping bounds (π).

Figure 5.8: Solution for Close et al. (2013)

5.8.1 Probability of adoption/abandonment

Using formulas in section (5.7) it is possible to compute the probability of abandonment and the probability of investment. The probability of abandonment using (5.23) and (5.22) is $P_A(\tau^* = \check{\tau}(\Lambda_A)) = 0.69$ and the probability of investment $P_I(\tau^* = \hat{\tau}(\Lambda_I)) = 0.31$.

The probability of committing a type I error is $\alpha = 0.77$ while the probability of committing a type II error is $\beta = 0.0035$. These results seem to go against standard practice of keeping Type I error probabilities low. The reason for a high α and low β in this model is as follows. A Type I error implies that one adopts the technology if it's not effective. This may be costly due to the additional expense related to the technology, but does not harm patients and, therefore, has no impact on patient's health benefits. A Type II error, however, implies not treating with a superior technology. This error carries with it large opportunity costs: the health benefits that *would have been* realised if the technology had been accepted. The model, therefore, does what one would expect: it keeps β low. As a consequence α will be large.

5.9 Conclusion

The Bayesian Sequential Value of Information presented in this paper brings together statistical and economic modelling, allowing for flexible decisions that account for irreversibility costs. The model provides rules that allow the decision maker to take the decision that maximise health benefits and reduce losses on the health care system.

Our novel approach to healthcare technology assessment makes use of Bayes rule in order to compute the posterior probability for the effectiveness of the healthcare technology at each point during the randomised control trial. Decision bounds are a function of uncertainty and prior information and follow from the parameters of the model. Decisions are taken at the moment in which the net benefit of the healthcare technology hits a pre-specified threshold value. At this optimal stopping time there is not more gain to be made by further waiting.

5.10 Appendix

5.10.1 Proof of Lemma 5.1

The Wiener process X_t under P_1 and under P_0 takes the form

$$dX_t = \sigma dW_t \quad P_p = P_0$$

and

$$dX_t = \theta\mu dt + \sigma dW_t \quad P_p = P_1$$

Girsanov's Theorem allows for the change of measure P_1 to P_0.

Define $u(t,\omega) = \frac{-\mu}{\sigma}$ and

$$\begin{aligned}
\Lambda_t(t,\omega) &= exp\left(\frac{\mu}{\sigma}\int_0^t dW_s - \frac{\mu^2}{2\sigma^2}\int_0^t ds\right) \\
&= exp\left(\frac{\mu}{\sigma}\int_0^t \sigma dW_s - \frac{\mu}{2\sigma^2}t\right) \\
&= exp\left(\frac{\mu}{\sigma^2}\left[\int_0^t \sigma dW_s - \frac{\mu}{2}t\right]\right) \\
&= exp\left(\frac{\mu}{\sigma^2}\left(X_t - \frac{\mu}{2}t\right)\right)
\end{aligned}$$

Also, the $(\Lambda_t)_{t\geq 0}$ process is a martingale.

$$\begin{aligned}
\mathbb{E}_{P_0}[\Lambda_t \mid \mathcal{F}_s] &= \mathbb{E}_{P_0}[e^{\frac{\mu}{\sigma^2}(X_t - \frac{\mu}{2}t)} \mid \mathcal{F}_s] \\
&= \mathbb{E}_{P_0}[e^{\frac{\mu}{\sigma^2}[(X_t - X_s) - \frac{\mu}{2}(t-s)]} e^{\frac{\mu}{\sigma^2}(X_s - \frac{\mu}{2}s)} \mid \mathcal{F}_s] \\
&= \Lambda_s \mathbb{E}_{P_0}[e^{\frac{\mu}{\sigma^2}[(X_t - X_s) - \frac{\mu}{2}(t-s)]} \mid \mathcal{F}_s] \\
&= \Lambda_s e^{-\frac{\mu^2}{2\sigma^2}(t-s)} \mathbb{E}_{P_0}[e^{\frac{\mu}{\sigma^2}(X_t - X_s)}] \\
&= \Lambda_s e^{-\frac{\mu^2}{2\sigma^2}(t-s)} e^{\frac{\mu^2}{2\sigma^2}(t-s)} \\
&= \Lambda_s \qquad (5.27)
\end{aligned}$$

■

5.10.2 Proof of Lemma 5.2

Apply Ito's lemma to $\Lambda_t = \exp\left(\frac{\mu}{\sigma}\left(X_t - \frac{\mu}{2}t\right)\right)$ gives

$$\begin{aligned} d\Lambda_t &= \frac{\partial \Lambda}{\partial t}dt + \frac{\partial \Lambda}{\partial x}dx + \frac{1}{2}\frac{\partial^2 \Lambda}{\partial x^2}dx^2 \\ &= -\frac{1}{2}\frac{\mu^2}{\sigma^2}\Lambda_t + \frac{\mu}{\sigma^2}\Lambda_t \sigma dW_t + \frac{1}{2}\frac{\mu^2}{\sigma^4}\Lambda_t \sigma_t^2 dt \\ &= \frac{\mu}{\sigma}\Lambda dW_t \end{aligned}$$

■

Chapter 6

Don't Stop 'Til You Get Enough: a quickest detection approach to HTA

6.1 Introduction

When uncertainty about the net benefits of a health-care technology is present there is a positive probability of making an incorrect decision. The expected value of information developed by Raiffa and Schlaifer (See Pratt et al. (1995)) and later applied to the case of health technology assessment (HTA) and clinical research design by Claxton and Posnett (1996) and Claxton (1999) can be used to quantify the expected opportunity loss associated with this uncertainty. When the expected opportunity loss is less than the cost of a new study the information is deemed to be sufficient and a decision can be made. When this static decision making approach is implemented to clinical research design it suggests to compute an ex-ante optimal (fixed) sample size deemed to be sufficient for the purposes of decision making. Claxton (1999) put forward the idea that inference is irrelevant to decision making and suggested that the question of whether more evidence is needed should be determined by the value of information framework developed by Raiffa and Schaifer (See Pratt et al. (1995)).

Recently, William and Pinto (2005) suggested a method for computing the ex-ante sample size n^* for a clinical trial that maximises the difference between the cost of a trial and the expected value of

the results using the incremental net benefit as the main outcome for the trial. William and Kowgier (2008) developed the model of William and Pinto (2005) to a multistage adaptive-design involving an early termination rule based on the expected net gain from the trial computed for each stage j. If the EVSI in the next stage $j+1$ is less than the total cost at $j+1$ then the trial terminates at the end of the jth stage and the decision rule can be applied. Although it is theoretically possible to construct a purely multistage model that jointly determines the value of n_j^* for all j maximizing the expected net gain, due to its complexity William and Kowgier (2008) suggest to proceed in two-stages steps where at each stage j the (ex-ante) two stage calculation is repeated and the maximisation process is repeated at each j. Another early termination approach is found in Berry and Ho (1988) who take the point of view of a pharmaceutical company that wishes to maximise profits and uses a one-sided decision-theoretic approach in order to determine if experimentation of a newly developed drug should be stopped early in case of negative evidence.

In recent years the literature has seen the application of the real option approach to investment decisions in health technology assessment (Palmer and Smith, 2000). This literature aimed at incorporating the dynamic nature of the decision process and considers the role of flexibility and irreversibility of investment. More recently Pertile et al. (2013) solved the dynamic problem of the economic valuation of a new health technology in the content of the optimal stopping under sequential sampling literature developed by Chernoff (1961). Forster and Pertile (2012) discuss the use of real options analysis as a way to view adoption, treatment and research decisions as a single economic project and argue that the dynamic approach to HTA can provide efficiency gains in resource allocation.

In this paper we present a comparison between the traditional value of information framework as found in Claxton (1999) and a dynamic decision theoretic approach. We adopt a sequential value of information (S-VoI) rule (see Bregantini and Thijssen (2013)) as this helps the user to reach a decision between two hypotheses after a minimal number of experiments. This method, in contrast to Pertile et al. (2013), does not involve an estimation problem for the

unknown net incremental mean benefit but specifies some bounds at which a decision can be taken for a given hypothesis on the mean benefit level. When the cumulative net incremental mean benefit hits one of the bounds the observed sample size is sufficient and the decision, either for investment or for abandonment, can be undertaken with minimal error. In Pertile et al. (2013)'s approach, a shortcoming is given by the requirement of a-priori specified maximum experiment time N. The implication is that the stopping bounds dependent on the maximum time N, influencing the investment/abandonment decision, and it is not clear how one should go about determining the maximum time N. The requirement of a pre-specified experiment time is avoided in the S-VoI model.

In contrast to the static approach, the S-VoI model does not force a decision after observing n observations no matter the information contained in the observed sample. In particular, in the case of a fixed sample, the size can be dangerously small or redundantly large for making a reasonably good inference on which of the two hypotheses is true.

With sequential testing on the other hand, no observations are wasted. In fact, as soon as we can declare that one of the two hypotheses is true with reasonable certainty, we stop taking observations. For this reason, in the presence of sampling costs, it is clear that sequential testing is a method of testing that is less costly on average than its competitor fixed sample size testing (Poor and Hadjiliadis, 2009)).

Consistent with Claxton (1999), and in contrast to traditional sample size calculations for randomised clinical trials based on type I and type II probabilities rules that do not account for the monetary cost or making the wrong decisions, the S-VoI focuses on expected payoff and aims at maximising health benefits with minimum error probability. In the sequential setting the implication for the irrelevance of inference suggested in Claxton (1999) is that fixing the sample size ex-ante is not optimal and, as with rules based on type I and type II error minimisation, can lead to choices that do not maximise health benefits with minimum error probability.

The paper begins by outlining the static value of information approach and the sequential value of information. These are then

followed by an illustrative example of research design based on simulations for the two models. Finally we report and contrast, in terms of monetary value of gained health benefit, the expected research design outcome for the Value of Information (VoI) approach to HTA found in Claxton (1999) and Claxton and Posnett (1996) and the S-VoI.

6.2 Static decision rules

We begin by introducing the main tools of the Value of Information approach as found in Raiffa and Shlaifer (see Pratt et al. (1995)) and adapted to the case of health technology assessment by Claxton and Posnett (1996) and Claxton (1999).

6.2.1 Expected Value of Perfect Information

Claxton (1999) propose to use the EVPI as a way to address how decision makers (DM) should interpret the results of probabilistic modelling and to address the question of whether enough evidence has been gathered. This approach mirrors the sequential nature of decision making: making an initial decision; deciding to gather evidence; revising decisions in the light of this new information; and again considering whether more information is required. It also ensures that the type of information acquired through research is driven by the objectives of the health care system and is valued in a way which is consistent with the budget constraint on service provision. In this framework, the expected cost of uncertainty is determined jointly by the probability that a decision based on current evidence will be wrong and the consequences of a wrong decision.

As in Ades et al. (2004)' set-up, the decision model has unknown parameters θ and the choice is between different treatment j. $\text{NB}(j, \theta)$ is the net benefit of treatment $j = 1, 2, \ldots, J$ for parameters of value θ. The optimal decision, subject to current knowledge, is the one that provides the higher expected net benefit:

$$\max_j \mathbb{E}_\theta \text{NB}(j, \theta). \tag{6.1}$$

Maximising over the possible interventions j in not possible as the true values of θ are unknown. However, it is possible to obtain the

expected net benefit of a decision taken with perfect information by averaging equation (6.1) over the joint distribution of θ:

$$\mathbb{E}_\theta \max_j \text{NB}(j, \theta). \tag{6.2}$$

The EVPI is the difference between equation (6.2) and (6.1) , amounting to the difference between the expected value of a decision made with perfect and current information:

$$\text{EVPI} = \mathbb{E}_\theta \max_j \text{NB}(j, \theta) - \max_j \mathbb{E}_\theta \text{NB}(j, \theta). \tag{6.3}$$

Expected Value of Sample Information

The value of information analysis can be extended in order to find the expected value of sample information for particular research design (Ades et al., 2004). In order to establish if the conditions for further research are present and to identify efficient research design there is the need to also consider the expected costs of sample information. The expected value of sample information was introduced as a decision tool for clinical trial design by Claxton and Posnett (1996) and Ades et al. (2004).

The EVPI places an upper bound on returns to further research and provides a necessary but not sufficient condition for conducting further research. If the value of EVPI exceeds the cost of further research it might be worthwhile to gather more information about the problem as a whole or on selected parameters. However, in order to establish if further research will be worthwhile (i.e. net benefits of research are positive) and to identify efficient research design there is the need to consider the marginal benefits and marginal costs of sample information.

Technically efficient research design

The EVSI can be calculated for a particular sample size from the prior information and the estimate of the sample variance (σ^2/n). The EVSI is then

A sample of size n on θ will give a sample result D. If the sample result were known, it would be possible for the decision maker to choose the alternative with the maximum expected payoff. It is possible to compute the expected net-benefit by averaging over the

posterior distribution of the net-benefit of each intervention j given the sample result D:

$$\max_j \mathbb{E}_{\theta|D}\text{NB}(j,\theta). \tag{6.4}$$

As the value of D is not known in advance (i.e. the result of the sample is not known), the expected value of a decision taken with sample information is computed by averaging the maximum expected net benefits over the distribution of possible values of D. In other words this amount to compute the expectation over the predictive distribution of the sample results D conditional on θ, averaged over the prior distribution of θ:

$$\mathbb{E}_D \max_j \mathbb{E}_{\theta|D}\text{NB}(j,\theta). \tag{6.5}$$

The EVSI is the difference between the expected value of a decision made with sample information and the expected value with current information:

$$\text{EVSI} = \mathbb{E}_D \max_j \mathbb{E}_{\theta|D}\text{NB}(j,\theta) - \max_j \mathbb{E}_\theta \text{NB}(j,\theta). \tag{6.6}$$

The EVSI proposed in (6.6) is for a single study design and single sample size. In order to establish the optimal sample size for a particular study these computations needs to be repeated for various sample sizes n.

The expected net benefit of sampling is the difference between the total benefit and the total variable cost for a particular sample size:

$$\text{ENBS}_{(n)} = \text{EVSI}_{(n)} - C_{s(n)} \tag{6.7}$$

The subscript n indicates the step in the trial and the cost $C_{s(n)}$ is the total trial cost at step n. The optimal sample size n^* maximises $\text{ENBS}_{(n)}$. The optimal value of n is given by the following condition

$$\frac{\partial \text{EVSI}}{\partial n} = C_n \tag{6.8}$$

As for the EVPI, simulation methods have been proposed in order to deal with the non-linearities and non-normal distribution of the net benefit (See Ades et al. (2004). However, the solution to the

decision problem in the value of information approach, as noted in equation (6.8), remains a static one: the maximisation of the EVSI is computed ex-ante, it computes a single value for the optimal sample size n and does not take into account any information that arises during the trial, in effect making the choice of n reasonable before the trail actually starts, but as we show in section 6.5.1, suboptimal at any point $n > 0$.

Cost

The EVSI does not account for costs different than those directly associated in running the trial. There are no health losses connected in delaying the decision and not treating patients with a more effective technology. The issue of forgone value of information has been introduced by Griffin et al. (2011), however, the value of information remained a static decision framework. A dynamic approach to research design has been undertaken by Claxton and Thompson (2001) where the approach found in Claxton and Posnett (1996) and Claxton (1999) are generalised to the analysis of a sequential clinical decision problem.

Claxton (1999) advocates that deciding which alternatives should be chosen, given existing information, and deciding whether more information should be required are two simultaneous but conceptually separate steps. The VoI provides a way to distinguish between these two concepts.

6.3 Sequential Value of Information (S-VoI)

In a sequential value of information (S-VoI) decision making model developed in Bregantini and Thijssen (2013), the DM is faced with a two-sided decision: either invest in the health care technology or abandon the health care project. The S-VoI is a quickest detection model that allows to test for the hypothesis with the minimum number of observation required, maximising payoff with minimum error probability.

In contrast to models that propose simulation based solutions, the S-VoI model uses continuous time mathematics that allows to fully understand the modelling results. The use of continuous over

discrete time modelling enables to access a mathematical toolbox that provides analytical solutions. While discrete time realise on a very large sample size to ensure convergence, in continuous time such requirement is avoided as convergence is guaranteed even in steps of tiny size. During the trial, the decision maker observes the net benefit for each patient as a sequence of outcomes. The net benefit over a small time interval is given by μdt. The decision maker however, cannot clearly observe the net benefit due to a noise element $\sqrt{\sigma}$. The evolution of the sequence is described by the following tree diagram

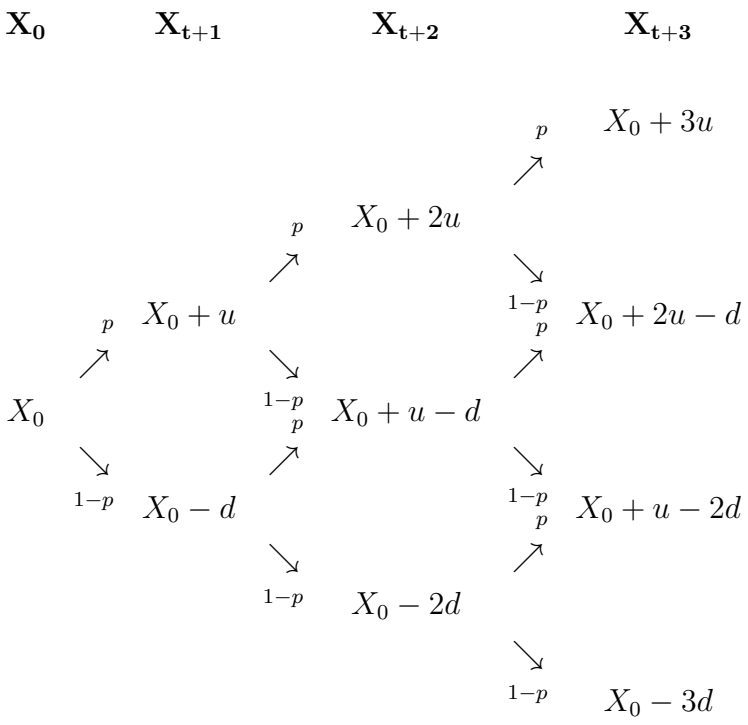

In the above tree the initial value X_0 can increase by a factor $u = \theta\mu dt + \sigma\sqrt{dt}$ or decrease by a factor $d = \theta\mu dt - \sigma\sqrt{dt}$. with probability $p = 0.5$. The term θ can be equal to 1 or 0 and will be used below for hypotheses testing. The sequence of random variable is

$$X_i = \begin{cases} \theta\mu dt + \sqrt{dt} & \text{with } pr = 1/2 \\ \theta\mu dt - \sqrt{dt} & \text{with } pr = 1/2 \end{cases}$$

At each point in the sequence the value of $X_{(n)}$ is given by

$X_{(n)} = \sum_{i=0}^{n} X_i$ and as the time interval between steps decreases we denote $X_t = \lim_{n\to\infty} X_n(t)$ where the limit is understood to be in distribution and $n \to \infty$ implies $dt \downarrow 0$. According to the CLT[1], the limit X_t exists in distribution and is given by

$$X_t \sim N(\theta\mu t, \sigma^2 t)$$

implying that in the continuous time limit the process X_t follows the arithmetic Brownian motion

$$X_t = \theta\mu t + \sigma W_t. \tag{6.9}$$

Equation (6.9) describes the net benefit as a continuous time sequence of random variables. In the equation, θ represents the hypothesis that the health care technology is effective and provides the claimed net benefit $\mu > 0$. With $\theta = 1$ the technology is effective and when $\theta = 0$ the technology is no better than standard care (in which case $\theta\mu t = 0$).

We consider the case where the decision maker is interested in testing the claim from a manufacturer that seeks reimbursement for a newly developed health care technology that should provide excessive benefit μ. The claim could also be related to the minimum effectiveness required for cost-effectiveness (i.e. μ such that net incremental mean benefit (NIMB) is positive) as part a cost-effectiveness trial by a health care manufacturer. Such test will allow the manufacturer to provide stronger evidence in support for government reimbursement.

We consider research design for a project that has an irreversible fixed cost I and net present value of adoption given as function of the posterior probability π. The investment payoff is $F_I(\pi)$ and the abandonment payoff $F_A(\pi)$. At each point in the sequence, Bayes rule allows to compute a posteriori probability process π_t as a function of (i) the prior probability assigned to the likelihood of the technology being more effective than standard care and (ii) a likelihood process $\Lambda_t(X_t)$ as a function of the trial sequence X_t (For details see Bregantini and Thijssen (2013)).

In this way the posterior probability π_t that the new technology is more effective than standard care in continuously updated via

[1] See Chapter 5 for the functional argument

Bayes rule. In order to reflect the possibility of investment when the technology is not effective (i.e. a type I error) and the possibility of abandoning the project when the technology is better than standard care (i.e. type II error) the payoffs are specified as follows:

$$F_I(\pi) = \pi P_1 - (1-\pi)P_0 - I \qquad P_1 > 0,\ P_0 > 0 \qquad (6.10)$$

$$F_A(\pi) = -\pi P_1 \qquad P_1 > 0 \qquad (6.11)$$

The term P_1 represents the monetary benefit to the healthcare system of investment in the new healthcare technology conditional on $\theta = 1$ and $-P_0$ represents the monetary loss of new healthcare technology conditional on $\theta = 0$.

Subject to sampling costs c and discount rate r, the problem is to find an optimal stopping time τ^* at which a decision can be taken, payoffs maximised and the value of waiting for an additional sample is zero. At the optimal stopping time τ^* the likelihood process $\Lambda_t(X_t)$ hits either the upper investment bound Λ_I or the lower abandonment bound Λ_A. The likelihood process provides evidence for hypothesis $H_1 : \theta = 1$ or $H_0 : \theta = 0$. At τ^* the DM stops sampling and an optimal decision can be taken, either for investment with payoff F_I (i.e. supporting H_1) or abandonment with negative payoff F_A (i.e. supporting H_0).

In the optimal stopping model, the decision to invest/abandon or continue research, in contrast to the VoI approach, is subject to the information generated by the random variable X_t. As the trial continues, information about the net benefit X increases, and consequently uncertainty about the true net benefit decreases. The optimal decision is taken at the time τ^*, when the value of waiting for a further sample is zero.

6.4 Quickest detection decision rules

The S-VoI model specifies investment and abandonment bounds that aim at maximising payoff. The following simulation study shows hitting times τ^* (times at which it is optimal to make a

Table 6.1: Simulation

Simulation	σ	mode$_{\tau^*}$	μ_{τ^*}	σ_{τ^*}	min$_{\tau^*}$	max$_{\tau^*}$
1	$\sigma = 0.1$	192	314	197.7	10	2619
2	$\sigma = 0.15$	195	435	321	34	3925
3	$\sigma = 2$	175	548	478	36	6803
4	$\sigma = 2.5$	179	667	660	25	8181

decision) for $100,000$ simulated sample path for some hypothetical levels of uncertainty. The table 6.1 below reports statistics for a simulation study based on the following values: $P_1 = 130$, $P_0 = 60$, $r = 15\%$, $I = 40$, $\mu = 0.15$. The value for σ is increased in small steps for each simulation in order to show the consequence of different degrees of uncertainty on the distribution of hitting times τ^*.

Figure (6.1) below shows the distribution of τ^* for different values of μ. As it can be noted in Figure (6.1a) the distribution of τ^* is centred around the mode[2] value of $\tau^* = 192$ with few events that occur after the $\tau^* = 1250$ region. Figure (6.1c) τ^* has a much thicker tail after $\tau^* = 1250$, indicating that there is a greater number of τ^* events after this value than in the previous model. For $\sigma = 0.15$ the mean $\mu_{\tau^*} = 435$ and $\sigma_{\tau^*} = 321$ with a minimum hitting time of 36 and a maximum hitting time of 3803. A substantial increase from simulation 1. In simulation 3, as σ increases, the statistical values for μ_{τ^*}, σ_{τ^*}, min$_{\tau^*}$, max$_{\tau^*}$ increase. It can be noted in Figure 6.1b that the number of events occurring after $\tau^* = 1200$ is much greater than in simulation 1 and 2 with some extreme events occurring well inside the far right tail of the distribution with a maximum of $\tau^* = 6803$. Figure 6.1d shows the distribution hitting time τ^* for simulation 4. Of the simulated models, this is the most extreme case with $\sigma = 0.25$. The mean is $\mu_{\tau^*} = 667$ with $\sigma_{\tau^*} = 660$ with a the maximum $\tau^* = 8181$ with most of the events occurring before $\tau = 4000$.

The above results suggests that while the mode for the hitting times do not vary much, distribution varies considerably and the dispersion for hitting times (or decision times) varies considerably given different levels of uncertainty. A consequence of this, for

[2]Here the mode is reported as the mean is affected by very extreme values and does not fully characterise the large concentrations of hitting times

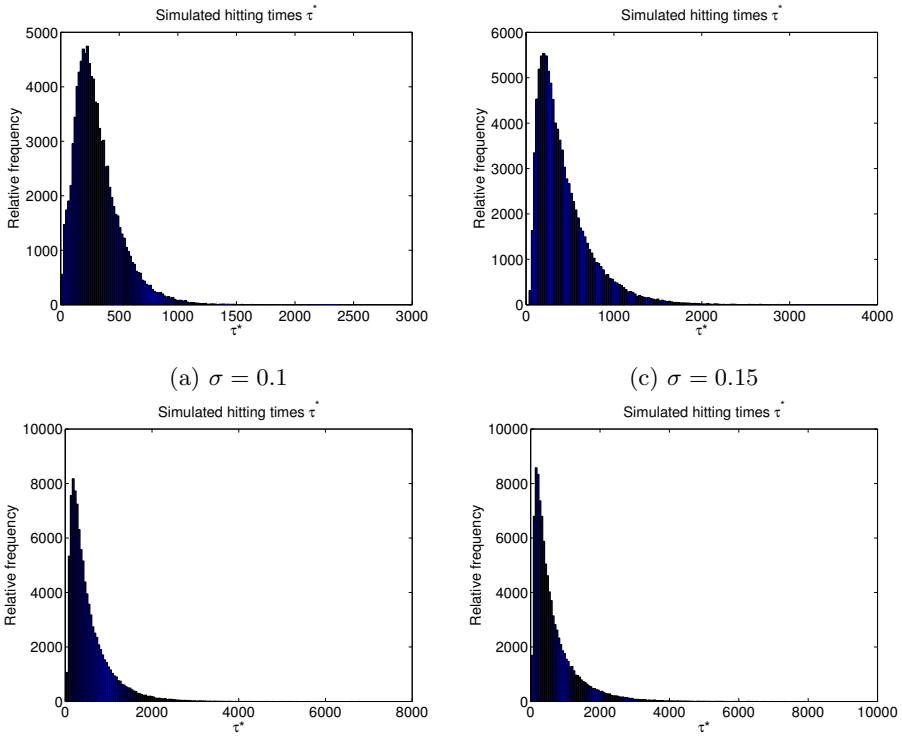

(a) $\sigma = 0.1$

(c) $\sigma = 0.15$

(b) $\sigma = 0.2$

(d) $\sigma = 0.25$

Figure 6.1: Simulated τ^* for different values of σ

models that use the total cost of new research (or stage of a trial) as a rule to determine if a new study should be undertaken (e.g. EVPI), is that when uncertainty is high it is difficult to correctly asses the cost of a new trial due to the uncertainty surrounding the optimal stopping time τ^*. For these ex-ante models, when estimating and comparing the average net benefit of new research with its costs, once should account for the uncertainty surrounding the optimal stopping time τ^*.

Additionally, another important consequence of uncertainty is that by having a rule that specifies ex-ante a fixed sample size for a trial, decisions might be taken at points where information is not sufficient or alternatively decision might be taken later than necessary with corresponding costs for the health care system. The cost of employing such ex-ante rules, for the specific case of EVSI, is discussed in the next section.

Table 6.2: Cost-effectiveness of standard vs robot-assisted laparoscopi prostaectomy

Parameter	Description	Source	Value
$E_1 - E_0$	Incremental QALY gain	Close *et al*	0.08 QALY
$C_1 - C_0$	Incremental cost	Close *et al*	£1412
σ	Std. deviation	Close *et al*	£1071
μ	Incremental QALY gain	Set as $NIMB > 0$	£1413
p	Prior	Assumed	0.5
r	Discount rate	Close *et al*	3.5%
c	Cost of sampling	Assumed	£50
I	Investment	Close *et al*	0
λ	QALY value	Close *et al*	£30000
n	Number of patients	Close *et al*	10000

6.5 EVSI vs quickest detection rules

In this section we aim at showing, with a simple illustrative example, the difference in a research design application between the two approaches. We simulate for a number of cases the stopping time produced by the dynamic Bayesian model and compare this to the optimal sample size given by the static value of information approach. For illustration purposes data is taken from Close et al. (2013)'s study of cost-effectiveness of standard vs robot-assisted laparoscopi prostaectomy. Data is shown in Table 6.2.

The EVSI predicts that the optimal sample size is $n^* = 91$. Figure (6.2) shows the relative frequency of hitting time τ^* for 10,000 simulated sample paths with sampling fixed at one new patient added per day (i.e. we assume 365 patients per year). When comparing the simulated hitting times with a static approach it can be noted that this last is likely to overestimate the sample size, nonetheless it can also underestimate the sample size for a good number of cases. In the analysis that follows τ and n represent the same values[3]. In the case of the S-VoI the mean hitting time τ^* is $\tau^* = 216$, just over 7 months. As there simulation displays some very high values for τ^*, the mode is reported as it avoids the influ-

[3]Since the same sampling scale is used we use τ to denote the dynamic optimal stopping model and retain n for the traditional EVSI approach. However, τ and n are equivalent and represent the same values.

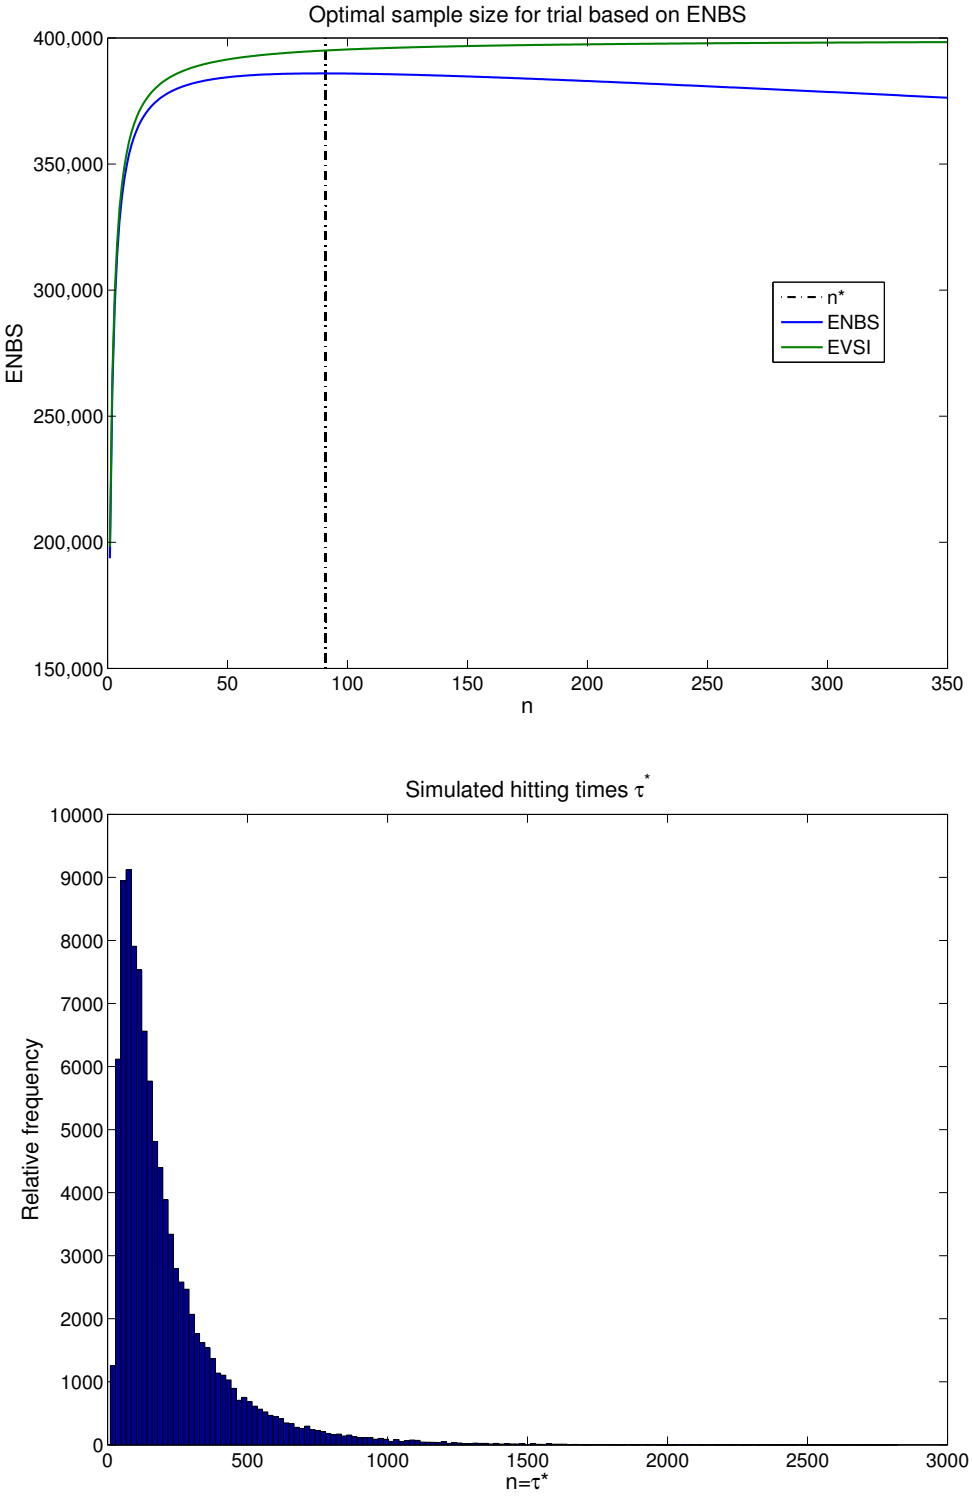

Figure 6.2: EVSI optimal sample size n^* and simulated τ^*

Table 6.3: Expected payoffs

Decision model	**Optimal sample size**	**Expected payoff £**
EVSI	Fixed (n^*)	207.7
S-VoI	Flexible (τ^*)	564.3
Health gain from S-VoI		356.6

ence of the very large hitting times and captures the most frequent τ^*. The mode is $\tau^* = 60$ (2 months), indicating that the most frequent decision time is lower than the EVSI. This seem to suggest that the S-VoI is quicker (in frequency) in making a decision than the EVSI for the values used in the case study.

The EVPI describes the advantage of full information over partial information. The EVSI involves computing the opportunity loss of making a decision based on prior information solely. In sharp contrast, in the S-VoI each decision is based on the appropriate information set generated by the random variable X_t.

The S-VoI approach provides a way to undertake the quickest decision that minimises expected opportunity loss both in terms of forgone health benefits to patients or resources allocated to the trial.

6.5.1 Cost of non optimal decisions

For the case study above, we compare the sequential-VoI and the EVSI payoffs. This gives an estimate of the costs involved in taking decisions based on a fixed, deterministic rule versus a sequential flexible rule.

Table (6.3) shows the expected value obtained from 100,000 simulated paths for a trial sequence X_t based on the Close et al. (2013) prostaectomy study reported above. The expected payoff is obtained by taking the maximum payoff value at the optimal decision point in the sequence X_t. For each simulated sample path, the expected payoffs are computed at the EVSI and the S-VoI decision point. This point is given by the optimal sample size (n^*) for the EVSI and at the optimal stopping time (τ^*) for the S-VoI[4]. The expected payoff at the optimal sample size predicted by the EVSI provides a low value, indicating that (n^*) is suboptimal when com-

[4]The discount rate r provides the link between the continuous and discrete time models.

pared to the S-VoI approach that instead selects the trial size (i.e. optimal stopping time τ^*) that maximises expected payoff. The S-VoI total expected gain for the health care system when compared to the EVSI approach is £356,600 per 1000 patients.

6.6 Further considerations and research

While the above example aims at showing the potential of the S-VoI, some additional aspects, relevant for clinical research and HTA, would need to be considered in practical applications. McKenna and Claxton (2011) highlights the tradeoff between the expected net benefits to current patients from being able to accessing a technology early and the future health benefits to patients that will be realised by withholding approval until new research evidence becomes available. The issue of population can be dealt by the model by separately modeling patients that will benefit from the trial in the future and current trial participants. In such case, a rescaled μ would measure the expected net benefit of the new technology in the trial's population while rescaled[5] P_1 and P_0 would account for the population that would benefit from the treatment once the technology is adopted or abandoned.

The value of information framework must consider the future population benefiting from the information derived from research. Choosing the time horizon over which information about a decision problem can be of use is challenging and poses a number of question on it's assessment and integration in a decision model (Philips et al., 2008). While the S-VoI assumes that the technology once adopted is used forever, considerations of a time horizon will impact the cost-effectiveness if investment costs are not recovered quickly enough. Many real-life investment are finite and expire or became valueless at some point in the future (e.g. patents). In the real option literature such scenario is dealt by allowing for a downward drift in the net benefit or by allowing the net present value of the project to jump to zero (Schwartz and Trigeorgis, 2001). Such time horizon adjustments for the S-VoI model can be implemented by

[5] A simple rescaling is given by $\mu * pop\ trial$ and $P_1 * number\ of\ future\ patients$ and $P_0 * number\ of\ future\ patients$

multiplying the costs and benefits by an exponential factor

$$1 - \exp(-rT)$$

where T is the life time factor.

6.7 Conclusion

Within the context of the value of information approach we compare deterministic versus dynamic rules for research design in HTA. The value of information approach selects the optimal trial length based on the prior information available and it produces a decision rule that proves to be inefficient for a great majority of cases. The reason is to be found in that under uncertainty evidence is accumulated over time and the point at which sufficient information is reached is not known at the start of the trial.

We show that this optimal decision point is reached at a random time that is optimal under some payoff based rules. As this optimal stopping time cannot be predicted at the start of the trial the research design advocated by the EVSI is inefficient brings losses to the health care system.

Part III

Inference for Stochastic Processes

6.8 Introduction to Part III:

6.8.1 Realised volatility

Realised volatility, in contrast to models that rely on strong parametric assumptions, for continuously observed prices allows for the measurement of return variation along with returns. The estimator was introduced concurrently by Andersen et al. (2001, 2003) and Barnoff-Nielsen and Shepard (2001, 2002). In the realised volatility framework the instantaneous return is decomposed into a logarithmic price process that features a martingale innovator. Consequently the expectation of the price process for high frequency data is zero with variance given by the quadratic variation of the local martingale. Realised volatility is computed by summing the intraday squared returns; with the absence of micro-structure noise realised volatility will provide a consistent estimator of the integrated volatility for each trading day.

Realistic models describing asset price dynamics consist in a continuous stochastic diffusion component and a discontinuous jump component (Andersen et al., 2002). Under jumps, the quadratic variation of the local martingale can be decomposed into a continuous and a discontinuous component. With realised volatility, when jumps are present in the data, the estimates are the sum of the quadratic variation of these two components. Barnoff-Nielsen and Shepard (2004) proposed the use of a partial generalisation of the quadratic variation called the realised bipower variation. This measure estimates the integrated variance isolating the continuous component of the process driving returns. The difference between realised volatility and bipower variation permits to isolate the jumps component.

Since the work of Barnoff-Nielsen and Shepard (2004) the econometrics for jumps identification and testing has seen a burst. Barnoff-Nielsen and Shepard (2006) show how the realised bipower variation can be used to construct statistical tests for the presence of jumps and derive the appropriate asymptotic distributional theory. Huang and Tauchen (2005) evaluate the Barnoff-Nielsen and Shepard (2006) jump detection test and propose a z-statistics that avoids the over-rejection of the null found in Barnoff-Nielsen and

Shepard (2006). The z-statistics makes use of the tripower quarticity measure developed by Andersen et al. (2005)

6.8.2 Statistical inference

Statistical inference, based on realised volatility measure, that test for the correct specification of the functional form of the volatility process can be found in a number of articles(e.g. see Corradi and Distaso (2006), Bollerslev and Zhou (2002), Todorov (2009)). Bollerslev and Zhou (2002) derive analytically the first two conditional moments for the integrated volatility and estimate the parameters governing the stochastic differential equation driving volatility using the generalised method of moments. Corradi and Distaso (2006) make use of the eigenfunction stochastic volatility model class found in Meddahi (2001), while Todorov (2009) estimates a Levy-driven Continuous Auto-regressive Moving Average model (CARMA) model.

Empirical evidence from high-frequency data analysis supports the hypothesis that there is jump component in the stochastic differential equation driving returns (e.g. see Barnoff-Nielsen and Shepard (2006), Huang and Tauchen (2005), Andersen et al. (2007)).

A moment-based econometric technique looks promising in the contest of estimation of the parameters for the stochastic differential equation that drive returns and for shedding further light on the size and frequency of stochastic jumps.

Chapter 7

Moment-Based estimation of Stochastic Volatility

7.1 Introduction

Estimation and forecasting of current and future volatility of returns is of practical importance in a number of financial applications. Over the past decade the development of areas such as risk management and option pricing has led to significant growth in financial market volatility research. A number of difficulties are encountered when estimating volatility. For example a major issue is that while daily raw returns are observable, volatility is latent. To deal with this latency a common approach is to conduct inference through strong parametric assumptions (e.g. ARCH, GARCH models) or to adopt a proxy for the unobserved volatility by inverting observed derivative prices and obtain market-based forecasts of implied volatility. The drawback of such methods is that they are model dependent and the measure can incorporate a time-varying volatility risk premium thus providing biased forecasts of the underlying asset's volatility (Andersen and Benzoni, 2009).

The main goal of this paper is to provide statistical inference for stochastic volatility models and to identify a class of jump-diffusion models that are successful in approximating S&P 500 high-frequency intra-day dynamics for the period 1997-2011 and could therefore be used as an adequate basis for continuous-time asset pricing applications. Further, the application of the power and bipower variation allows the empirical recovery of the realised

quadratic variation of the jump component on the S&P 500 index over each given day. Observing the jump process is of particular interest given the recent period of market turmoil.

There a number of articles that make use of realised measures to test for the correct specification of the functional form for the volatility process (e.g. see Todorov (2009), Corradi and Distaso (2006) and Bollerslev and Zhou (2002)). While Corradi and Distaso (2006) focus on the class of eigenfunction stochastic volatility models developed by Meddahi (2001) and Todorov (2009) focuses on a Levy-driven Continuous Auto-Regressive Moving Average (CARMA) model, Bollerslev and Zhou (2002) analytically derive the first two conditional moments for the latent integrated volatility, the realisation of which is given by the realised volatility estimator.

This paper contributes to the existing literature in a number of ways. First, moments are developed for the stochastic volatility model allowing for the leverage effect and for finite-activity jumps in returns. These are tested for their accuracy in the content of realised volatility. Second, in order to better understand the properties of the finite sample GMM estimators Monte Carlo experiments are presented for all the above extensions revealing estimation biases present in current modelling practices. Finally, applications of these methods are made on the equity asset class over the recent period of market turmoil giving new insights over this exceptionally high period of financial markets' volatility. The paper is particularly insightful for the jump dynamics.

The paper is structured as follows: section two describes the theoretical framework while section three states the stochastic volatility model and derives the two conditional moments that are later used in the GMM estimation. Section four reports the results of the Monte Carlo study for the GMM estimator and section five describes the data set used and reports data analysis for the series under study. Finally section six reports the GMM estimates while section seven and section eight describe model's jump in returns and leverage extensions.

7.2 Theoretical Framework

7.2.1 Quadratic Variation

The logarithm of an asset price $p_t = log\,(P_t)$ is assumed to be a semi-martingale in order to rule out arbitrage opportunities (Back, 1991). This is defined by

$$dp_t = \mu_t dt + \sigma_t dW_t$$

where μ_t and σ_t are predictable process independent of the Brownian motion dW_t. The drift μ_t is of finite variation, while σ_t is strictly positive and square integrable (i.e. $\mathbb{E}\left(\int_0^t \sigma_s^2 ds\right) < \infty, \forall t$).

Merton (1980) developed an estimator for integrated volatility that involves the sum of intraday squared returns. He showed that choosing increasingly finer observation intervals increases the accuracy of the estimator. Andersen et al. (2001) formally introduced realised volatility as a sum of intraday squared returns.

The continuously compounded return for a trading day $T - t$ is given by

$$r\,(T-t) = p_T - p_t = \int_t^T \mu_\tau d\tau + \int_t^T \sigma_\tau dW_\tau \tag{7.1}$$

and its quadratic variation $QV\,(T,t)$is defined as

$$QV\,(T,t) = \int_t^T \sigma_\tau^2 d\tau. \tag{7.2}$$

Innovations in the drift do not affect the sample path variation of the return; intuitively, the mean term μ_t, when cumulated over many high-frequency returns over short horizons can be neglected Andersen and Benzoni (2009). The diffusive sample path variation over the interval $T - t$ is the integrated variance $IV\,(T,t)$

$$IV\,(T,t) = \int_t^T \sigma_\tau^2 d\tau \tag{7.3}$$

Equations (7.2) and (7.3) coincide when no jumps are present[1]. With no micro-structure noise[2] and measurement error, the quadratic variation for the return process can be approximated by cumulative

[1]The distinction between QV and IV is in the approximation of the observation frequencies. IV belong to the pure continuous time case.

[2]This requirement can be relax. See section 7.5.1

squared return process. This was formally introduced by Andersen *et al* (2001)[3] and is referred to as realised volatility. The realised volatility estimator, for a partition $\left\{t + \frac{j}{n}, j = 1 \dots n \cdot k\right\}$, with R defined as in equation (7.1) is given by

$$RV\left(T,t\right) = \sum_{j=1}^{n \cdot k} r\left(t + \frac{j}{n}, \frac{1}{n}\right)^2. \tag{7.4}$$

When the sampling frequency increases, semi-martingale theory ensures that realised volatility (7.4) converges to the quadratic variation (7.2). Formally

$$RV\left(T,t,n\right) \xrightarrow{p} QV\left(T,t\right) \quad \text{as } n \to \infty.$$

7.2.2 Bipower Variation

In the presence of jumps, as pointed out by Andersen *et al* (2001), realised volatility becomes

$$RV\left(T,t\right) = \int_t^T \sigma_\tau^2 d\tau + \sum_{i=1}^{N(t)} J_i^2 \tag{7.5}$$

where N is a finite activity ($N(t) < \infty$) simple counting process. Barnoff-Nielsen and Shepard (2004) introduced the generalised power variation and focused on the $1,1$-order bipower variation. This is defined as

$$BV\left(T,t\right) = \mu^{-2} \sum_{j=1}^{M-1} \mid r_j \mid\mid r_{j+1} \mid \tag{7.6}$$

where $\mu = \mathbb{E} \mid u \mid$ is the mean of the absolute value of a standard normally distributed random variable, u . This amounts to $\mathbb{E} \mid u \mid = \sqrt{2/\pi} \simeq 0.79788$. Barnoff-Nielsen and Shepard (2004) then show that as $n \to \infty$

$$BV\left(T,t\right) \xrightarrow{p} \int_t^T \sigma^2\left(u\right) du. \tag{7.7}$$

In the stochastic volatility semi-martingale case this implies that in the limit the difference between the realised volatility and the quadratic variation amounts to

$$RV\left(T,t\right) - BV\left(T,t\right) \xrightarrow{p} 0 \tag{7.8}$$

[3]See Appendix A for details.

In constrast, in the presence of jumps, the realised power variation is unaffected by jumps and the difference in Eq. (7.8) converges in probability to a positive, finite quantity

$$RV\left(T,t\right)-BV\left(T,t\right)\overset{p}{\rightarrow}\sum_{i=1}^{N(t)}J^{2}. \tag{7.9}$$

This is the core theoretical insight that enabled for the empirical estimation of jumps in the paper. Andersen et al. (2007) point out that it is possible for the jump quantity $\sum J^2$ in Eq. (7.9) to be negative in finite samples and suggest a zero truncation scheme

$$J_{(T,t)}^{2}=max\left[RV_{(T,t)}-BV_{(T,t)},0\right] \tag{7.10}$$

to ensure non negative daily estimates. The result given by (7.10) will be later used in order to construct a test for jumps.

7.3 Stochastic volatility diffusion and GMM estimation

The generic continuous time stochastic volatility model can be stated as

$$dp_{t} = \mu(p_{t},V_{t}t)dt+v\left(p_{t},V_{t}\right)dB_{t} \tag{7.11}$$

$$dV_{t} = k\left(p_{t},V_{t}t\right)dt+\sigma\left(p_{t},V_{t}t\right)dW_{t} \tag{7.12}$$

where p_t is the time t logarithmic price for some asset and B_t and W_t are (possibly correlated) Brownian motions. Consistent with empirical evidence provided by Andersen and Bollerslev (1997), indicating that there is little predictable variation in the mean for high-frequency returns, the drift term is set to zero $\mu(p_t, V_t) = 0$. By the theory of quadratic variation given above, we have

$$\underset{n\rightarrow\infty}{lim}\sum_{i=1}^{2^{N}}\left[p_{t+i/2^{N}(T-t)}-p_{t+(i-1)/2^{N}(T-t)}\right]^{2}\overset{a.s}{\rightarrow}\int_{t}^{T}v^{2}\left(p_{s},V_{s}\right)ds\equiv\mathcal{V}_{t,T} \tag{7.13}$$

where $\mathcal{V}_{t,T}$ denotes integrated volatility from time t to T. The point in time volatility $v\left(p_t, V_t\right)$ is in general unobservable. By taking the sum of increasingly finer sampled squared high-frequency returns,

it is possible to obtain estimates of the integrated volatility process. In the limit, it is possible to observe the integrated volatility.

In practice, continuously sampled observations are unavailable and integrated volatility is not truly observable. GMM modelling can be implemented by exploiting the assumption that the number of observations in the sample moments converges to infinity at a slower rate than the almost sure convergence rate of the quadratic variation (Bollerslev and Zhou, 2002).

7.3.1 Baseline stochastic volatility (SV) model

The first model to be estimated is a stochastic volatility model with no drift given by

$$\begin{aligned} dp_t &= \sqrt{V_t} dB_t \\ dV_t &= k\left(\theta - V_t\right) dt + \sigma \sqrt{V_t} dW_t \end{aligned} \tag{7.14}$$

where V_t is the scalar latent volatility process. This model with correlated Brownian motions was presented by Heston (1993) and has found wide applications in empirical finance. In the model θ determines the long-run (unconditional) mean of the volatility, k is the mean reversion parameter, while σ denotes the local variance (volatility of volatility) parameter.

The process is well defined with $\theta > 0$ (non negativity), $k > 0$ (stationarity in the mean), and $\sigma^2 \leq 2k\theta$ (stationarity in volatility). The deterministic part of process is asymptotically stable if $k > 0$ and the condition $\sigma^2 \leq 2k\theta$ ensures that the variance process cannot reach zero Feller (1951).

A distinction is drawn between two information sets: the continuous σ-algebra $\mathcal{F}_t = \sigma\left\{V_s; s \leq t\right\}$, generated by the point-in-time volatility and the discrete σ-algebra $\mathcal{G}_t = \sigma\left\{\mathcal{V}_{t-s-1,t-s}; s = 0, 1, 2, \ldots, \infty\right\}$ generated by the integrated volatility series. The coarser filtration is a subset of the finer filtration ($\mathcal{G}_t \subset \mathcal{F}_t$) and applying the law of iterated expectations leads to $\mathbb{E}\left[\mathbb{E}\left(\cdot \mid \mathcal{F}_t\right) \mid \mathcal{G}_t\right] = \mathbb{E}\left(\cdot \mid \mathcal{G}_t\right)$.

Conditional mean

Next conditional moments are derived. Following from Cox et al. (1985), the conditional mean of the point in time volatility is given by

$$\mathbb{E}\left(V_T \mid \mathcal{F}_t\right) = \alpha_{T-t} V_t + \beta_{T-t} \tag{7.15}$$

where $\alpha_{T-t} = e^{k(T-t)}$ and $\beta_{T-t} = \theta\left(1 - e^{-k(T-t)}\right)$ are function of k, θ and $T - t$. Bollerslev and Zhou (2002) express the conditional mean of the stochastic process as a linear function of the point in time volatility giving

$$\mathbb{E}\left(\mathcal{V}_{t,T} \mid \mathcal{F}_t\right) = \mathbb{E}\left(\int_t^T V_s ds \mid \mathcal{F}_t\right) = a_{T-t} V_t + b_{T-t} \tag{7.16}$$

where $a_{T-t} = 1/k\left(1 - e^{k(T-t)}\right)$ and $b_{T-t} = \theta\left(T - t\right) - (\theta/k)\left(1 - e^{-k(T-t)}\right)$ denote functions of the drift parameters and the sampling interval. Using equations (7.15) and (7.16) for a one-day horizon (thus $a = a_1, b = b_1, \alpha = \alpha_1$ and $\beta = \beta_1$ gives

$$\begin{aligned}
\mathbb{E}\left[\mathbb{E}\left(\mathcal{V}_{t+1,t+2} \mid \mathcal{F}_{t+1}\right) \mid \mathcal{F}_T\right] &= a\mathbb{E}\left(V_{t+1} \mid \mathcal{F}_t\right) + b \\
&= a\left(\alpha V_t + \beta\right) + b \\
&= \alpha\left[\mathbb{E}\left(\mathcal{V}_{t,t+1} \mid \mathcal{F}_t\right) - b\right] + a\beta + b
\end{aligned}$$

simplifying to

$$\mathbb{E}\left(\mathcal{V}_{t+1,t+2} \mid \mathcal{F}_t\right) = \alpha\mathbb{E}\left(\mathcal{V}_{t,t+1} \mid \mathcal{F}_t\right) + \beta.$$

This equation can be conditioned on the coarser information set $\mathcal{G}_t$, giving the link between the first moment of the integrated volatility and lagged integrated volatility. This gives

$$\mathbb{E}\left[\mathbb{E}\left(\mathcal{V}_{t+1,t+2} \mid \mathcal{F}_t\right) \mid \mathcal{G}_t\right] = \mathbb{E}\left(\mathcal{V}_{t+1,t+2} \mid \mathcal{G}_t\right) = \alpha\mathbb{E}\left(\mathcal{V}_{t,t+1} \mid \mathcal{G}_t\right) + \beta. \tag{7.17}$$

Conditional second moment

The derivation of the second moment follows in similar fashion to the above. Again, following Cox et al. (1985)

$$\begin{aligned}
\mathbb{E}\left(V_T^2 \mid \mathcal{F}_t\right) &= Var\left(V_T \mid \mathcal{F}_t\right) + \left[\mathbb{E}\left(V_T \mid \mathcal{F}_t\right)\right]^2 \\
&= C_{T-t} V_t + D_{T-t} + \left[\alpha_{T-t} V_t + \beta_{T-t}\right]^2.
\end{aligned}$$

In the same fashion as for the first moment, the conditional variance of the integrated volatility is expressed as a function of the point in time volatility. An application of Itō's Lemma leads to

$$Var\left(\mathcal{V}_{t,T} \mid \mathcal{F}_t\right) = A_{T-t}V_t + B_{T-t} \tag{7.18}$$

where A_{T-t} and B_{T-t} are functions of the parameters[4]. By combining (7.16) and (7.18) it is possible to obtain, for the one day horizon

$$\begin{aligned} \mathbb{E}\left(\mathcal{V}_{t,t+1}^2 \mid \mathcal{F}_t\right) &= Var\left(\mathcal{V}_{t,t+1} \mid \mathcal{F}_t\right) + \left[\mathbb{E}\left(\mathcal{V}_{t,t+1} \mid \mathcal{F}_t\right)\right]^2 \\ &= a^2V_t^2 + \left(2ab + A\right)V_t + \left(b^2 + B\right). \end{aligned} \tag{7.19}$$

By application of the Law of Iterated Expectation on different information sets and substituting expressions it is possible to obtain

$$\begin{aligned} \mathbb{E}\left[\mathbb{E}\left(\mathcal{V}_{t+1,t+2} \mid \mathcal{F}_t\right) \mid \mathcal{G}_t\right] &= \mathbb{E}\left(\mathcal{V}_{t+1,t+2}^2 \mid \mathcal{G}_t\right) \\ &= H\mathbb{E}\left(\mathcal{V}_{t,t+1}^2 \mid \mathcal{G}_t\right) + I\mathbb{E}\left(\mathcal{V}_{t,t+1} \mid \mathcal{G}_t\right) + J \end{aligned} \tag{7.20}$$

H, I and J are functions[5].

7.3.2 Conditional moments

GMM estimation requires the specification of moment conditions. Zhou (2001), in a Monte Carlo application to the square-root diffusion process of Cox *et al* (1985), constructed a GMM estimator with lag-one augmented moments.

Bollerslev and Zhou (2002) employ the analytical solutions for the conditional mean given by (7.17) and the conditional second moment given by (7.20), augmented by the lag-one and the lag-one squared realised volatility to construct a standard GMM estimator. This leads to the following six moments:

[4] See Appendix A for details.
[5] See Appendix A for details.

$$f_t(\xi) \equiv \begin{bmatrix} \mathbb{E}\left[V_{t+1,t+2} \mid \mathcal{G}_t\right] - V_{t+1,t+2} \\ \mathbb{E}\left[V^2_{t+1,t+2} \mid \mathcal{G}_t\right] - V^2_{t+1,t+2} \\ \mathbb{E}\left[V_{t+1,t+2}V_{t-1,t} \mid \mathcal{G}_t\right] - V_{t+1,t+2}V_{t-1,t} \\ \mathbb{E}\left[V^2_{t+1,t+2}V_{t-1,t} \mid \mathcal{G}_t\right] - V^2_{t+1,t+2}V_{t-1,t} \\ \mathbb{E}\left[V_{t+1,t+2}V^2_{t-1,t} \mid \mathcal{G}_t\right] - V_{t+1,t+2}V^2_{t-1,t} \\ \mathbb{E}\left[V^2_{t+1,t+2}V^2_{t-1,t} \mid \mathcal{G}_t\right] - V^2_{t+1,t+2}V^2_{t-1,t} \end{bmatrix} \tag{7.21}$$

The moment's one period lag implies a $MA(1)$ error structure[6]. The augmented moments involve the product of the related error term with the lagged (and the squared lag) realised volatility. The true parameters for the process are given by ξ_0 and following Hansen (1982) formulation the set of orthogonality conditions are of the form

$$\mathbb{E} f_t(\xi_0 \mid \mathcal{G}_t) = 0. \tag{7.22}$$

The idea behind GMM is to chose ξ in order to make the sample mean of the moment conditions,

$$g_T(\xi) = \frac{1}{T}\sum_{t=1}^{T-2} f_t(\xi) \tag{7.23}$$

as close as possible to the population moment of zero Hamilton (1994). The corresponding GMM estimate $\hat{\xi}_T$ is the value of ξ that minimises

$$g_T(\xi)' \hat{S}_T^{-1} g_T(\xi) \tag{7.24}$$

where S_T^{-1} is the inverse of the asymptotic variance matrix. $\hat{S}_T$, is an estimate of

$$S = \lim_{T\to\infty} (1/T)\sum_{t=1}^{T}\sum_{v=-\infty}^{\infty} \mathbb{E}\left\{[g_T(\xi)]\,[g_T(\xi)]'\right\}$$

In the asset price simulation and the empirical estimates the heteroschedasticity and autocorrelation robust covariance matrix estimator with Bartlett-Kernel and lag length of five is used.

[6]The first moment can be written as $\mathcal{V}_{t+1,t+2} - (\alpha\mathcal{V}_{t,t+1} + \beta) = MA(1)$ while the second moment can be written as $\mathcal{V}^2_{t+1,t+2} - \left(H\mathcal{V}^2_{t,t+1} + I\mathcal{V}_{t,t+1} + J\right) = MA(1)$.

7.4 Monte Carlo study

As in later sections the conditional moments developed above will be used to estimate realised volatility obtained from high-frequency data, a Monte Carlo study is constructed in order to test for the accuracy of the GMM estimator.

To simulate the asset price, the SDE (7.14) is discretized using the Euler-Maruyama approximation with 78 artificial five-minute intervals per trading day and further partition to 780 observations per day. This last finer partitioning is implemented in order to emulate the "continuous-time record" and to provide a comparison for the GMM estimation between the five-minutes quadratic variation and the true integrated volatility. Figure 7.1 below shows a sample path for the SDE.

Integrated volatility is approximated by the application of (7.13) to the simulated sample. This involves summing squared intraday returns in order to obtain a series of realised volatility. The study design, in order to achieve a comparable set of results, follows Bollerslev and Zhou (2002); two series with respective sample size of T=1000 and T=4000 are simulated in order to check for long-span asymptotics. This amount to around 4 and 16 years of daily observations. The total number of Monte Carlo replications is 1000.

7.4.1 Simulation results

The simulation concentrates on a single scenario where $k = 0.03, \theta = 0.25$ and σ=0.1 and statistics for the Monte Carlo estimation are given in Table 7.1. This reports the true parameter values, the mean and median for the simulated values and RMSE for the GMM estimation.

7.4.2 Parameter estimates

From the simulation estimates, it can be noted that as the sample size increases, the mean and median get closer to the true value, thus reducing the estimation bias. The GMM estimates for the parameter k are slightly upward biased, while for the parameter θ there is a small downward bias. It can be seen that for the small

Table 7.1: Monte Carlo GMM estimation

	True value	**Mean**		**Median**		**RMSE**	
	T	1000	4000	1000	4000	1000	4000
GMM with integrated volatility							
Observations per day: 780							
k	0.03	0.0384	0.0326	0.0375	0.0323	0.0092	0.0087
θ	0.25	0.2340	0.2451	0.2283	0.2437	0.0544	0.0265
σ	0.10	0.0997	0.1004	0.0999	0.1004	0.0042	0.0050
GMM with quadratic variation from high-frequency returns							
Observations per day: 78							
k	0.03	0.0349	0.0314	0.0337	0.0310	0.0072	0.0058
θ	0.25	0.2407	0.2447	0.2361	0.2428	0.0529	0.0262
σ	0.10	0.0979	0.0994	0.1009	0.1023	0.0025	0.0043

sample size of 1000, k has a mean value of 0.0384 compared to the true value of 0.030 while the mean for θ is 0.2340 compared to the true value of 0.25 for the integrated volatility case. For the quadratic variation, k has a mean value of 0.0349 compared to the true value of 0.030 while the mean for θ is 0.2407 compared to the true value of 0.25.

Finally for the volatility-of-volatility parameter σ the bias is very small. The RMSE, in the case of integrated volatility, for the parameters k, θ decreases with the long-span and has little variation for the σ parameter. Similarly, for the case with high-frequency returns, k, θ decreases with the long-span and has little variation for the σ parameter. This can be noted when comparing the values for the sample size of 1000 to the sample size of 4000. The mean value for θ improves respectively from 0.2340 to 0.2407 and for the parameter k this improves from 0.0384 to 0.0349 while the volatility parameter σ changes from 0.0997 to 0.0979.

Bollerslev and Zhou (2002) found that for the smaller sample size the standard GMM J-test for over-identifying restrictions exhibits a small bias and that this tends to disappear once the sample size increases from 1000 to 4000.

7.5 Data

Data for the stock price index is obtained from Market Grain Research (MGR) in Marietta, Georgia, USA, and consists of continuously recorded five-minute prices on the S&P 500. The full sample starts from August 1, 1997 through June 20, 2011 and consists of 271,092 high-frequency observation for a total of 3,494 days.

7.5.1 Sampling Frequency

While finely sampled data would minimise the estimation error, the presence of market micro-structure issues such as price discreteness, bid-ask spread and non-synchronous trading effects, imply that the underlying semi-martingale assumption is violated at very high sampling frequency Bollerslev et al. (2008). There is thus the need to strike a balance between noise and sampling frequency. Hansen and Lunde (2006) empirical estimates suggest that for highly liquid assets a five-minute sampling frequency provides an adequate compromise.

7.5.2 Calendar effects and other adjustments

Following Andersen et al. (2001) a number of days have been explicitly excluded in order to avoid calendar effects arising from various holiday periods and the related reduced trading hours. Several holidays have been removed, including Christmas (December 24-26), New Year's (December 31 and January 1-2), and July Fourth. In addition, the following moving holidays have been excluded from the sample, Good Friday, Easter Monday, Memorial Day, July Fourth (when officially on the 3rd), Labor Day, Thanks-giving and the day after.

As argued by Andersen et al. (2001) these adjustments do not completely eliminate all holiday market slowdowns, however, they eliminate the most important calendar effects. The data set, once holidays has been removed, is reduced to 267,129 observations consisting of 3,428 days. The typical trading day is characterised by 78 observations, however, the data set is in-homogenous as in places exhibits some missing price observations due to interrupted data feed. This is particularly true for the year 2000, where on average a

Table 7.2: Data summary statistics

S&P 500 realized volatility	
RV	S&P500
Period	1997-2011
μ	1.065
σ	2.017
σ_{annual}	15.93
Skewness	9.49
Kurtosis	154.33
Min	0.0004
Max	0.507

datapoint goes missing every two trading days for a total of about 100 data points.

In order to deal with in-homogenous high-frequency data Andersen et al. (2001) deleted trading days containing the 15 longest DM/$ zero runs. For the S&P 500, days with missing data are in general characterised very small lapses of 1 and more rarely 2 observations. By adopting a similar approach to Andersen et al. (2001), and by taking into account the shorter trading day consisting of 78 observations, days with less than 70 observations per day are deleted. This amounts to deleting a total of eight trading days, taking the sample to 3420.

Table 7.2 summarised data statistics for the S& P 500 over the period. This compares well with data obtained using Shepard and Shepard (2010) methodology.

7.6 Empirical estimation

7.6.1 Estimation results

This section presents GMM estimates for equation (7.14) for the periods 1st August 1997 to the 20th June 2011. Table 7.3 reports estimates for the stochastic volatility model (SV)

$$\begin{aligned} dp_t &= \sqrt{V_t} dB_t \\ dV_t &= k\left(\theta - V_t\right) dt + \sigma\sqrt{V_t} dW_t \end{aligned} . \tag{7.25}$$

Reported estimates confirm the data features highlighted in the summary statistics. The long-run mean parameter $\theta \times 100$ is close

Table 7.3: GMM Estimation

SV model estimates	
	S&P 500
	1997-2011
k	0.1135
	(0.0353)
θ	0.0103
	(0.0011)
σ	0.0840
	(0.0104)
GMM test of overidentifying restrictions	
$J-stat$	0.5919
$P\left[\mathcal{X}^2 > J\right]$	0.8983
$2k\theta \geq \sigma^2$	Reject

to the sample mean of the realised volatility in Table 7.2. The speed adjustment parameter k is slightly smaller than the one obtained in the exchange rate estimates of Bollerslev and Zhou (2002) and has also a large standard error[7]. When the results for the S&P 500 are compared to the estimates of Bollerslev and Zhou (2002) for the DM/$ exchange rate, it can be noted that the volatility generated by the equity index displays a much lower volatility-of-volatility σ and average level of spot variance θ. For the equity index the speed of variance mean reversion k is the dominating parameter. This indicates that the S&P 500 is less volatile than the DM/$ exchange rate but reverts to the mean level in similar fashion to the exchange rate.

The test for over-identifying restrictions does not reject the model for the samples analysed, however, the stationarity condition $2k\theta \geq \sigma^2$ is violated by the data. This means that the model is asymptotically stable ($k > 0$), however, the variance can reach zero.

Overall the estimates display strong mean reversion, confirming what can be inferred by visual inspection (See Figure 7.2).

[7]No direct comparison with previous studies is possible as the methodology has never been applied to the S&P500. Here the results are compared to the exchange rate as a way to check for reasonable estimates

7.7 Stochastic volatility with leverage

7.7.1 Leverage parameter (SVL Model)

In this section Heston (1993)'s model is extended by allowing for correlated Wiener processes. In model (7.14) this would involve

$$\begin{aligned} dp_t &= \sqrt{V_t} dB_t \\ dV_t &= k\left(\theta - V_t\right) dt + \sigma\sqrt{V_t} dW_t^1 \\ dB_t &= \rho dW_t^1 + \sqrt{1-\rho^2} dW_t^2 \end{aligned} \tag{7.26}$$

where $W_t^1, W_t^2$are independent Wiener processes and the leverage parameter is given by $dB_t dW_t^1 = \rho dt$. Various moments aimed at estimating the leverage parameter have been proposed[8]. By applying Itō's lemma Garcia et al. (2011) develop the following moment involving the return over a given day $r_{t,T} = p_T - p_t$ for the estimation of the leverage parameter.

$$\mathbb{E}\left[r_{t,T}\frac{\mathcal{V}_{t+1,t+2}-b}{a} \mid \mathcal{G}_t\right] = \frac{1}{e^k k}\left[\rho\sigma\left(\frac{\mathbb{E}\left(\mathcal{V}_{t,t+1} \mid \mathcal{G}_t\right)-b}{a}k + \theta\left(-1+e^k-k\right)\right)\right] \tag{7.27}$$

This is directly implementable.

Table 7.4 shows the Monte Carlo estimation results. The RMSE for the estimated coefficient have similar magnitude than estimates presented in Garcia et al. (2011). As with their estimates for the leverage parameter, the mean and median values for $\hat{\rho}$ indicate a downward bias. This amounts to about 15%, a value close to the one reported in Garcia et al. (2011)'s study.

7.7.2 Estimation results

Table 7.5 displays estimates for the stochastic volatility model with leverage (SVL)

The estimates for k, and σ are in line with Garcia et al. (2011), while θ is much lower. The leverage parameter ρ is bigger as Garcia *et al* finds -0.165 and indicates a much higher leverage effect for the period 1997-2011. It should be noted that the above results are not directly comparable to Garcia et al. (2011) as in their study

[8]Ishida *et al* 2011 report poor performance of the leverage moment proposed by Bollerslev and Zhou (2002) and suggest as an alternative a leverage moment based upon Corradi and Distaso (2006) realised correlation estimator.

Table 7.4: Monte Carlo estimation for SVL Model

	True value	**Mean**		**Median**		**RMSE**	
	T	1000	4000	1000	4000	1000	4000
GMM with integrated volatility							
Observations per day: 780							
k	0.03	0.0378	0.0326	0.0363	0.0322	0.0139	0.0057
θ	0.25	0.2362	0.2445	0.2307	0.2424	0.0540	0.0268
σ	0.10	0.0998	0.1003	0.0996	0.1004	0.0043	0.0022
ρ	-0.5	-0.441	-0.4764	-0.4234	-0.4697	0.2738	0.1485
GMM with quadratic variation from high-frequency returns							
Observations per day: 78							
k	0.03	0.0356	0.0313	0.0342	0.0310	0.0130	0.0054
θ	0.25	0.2392	0.2457	0.2341	0.2445	0.0547	0.0262
σ	0.10	0.1009	0.1023	0.1011	0.1023	0.0083	0.0048
ρ	-0.5	-0.4615	-0.4684	-0.4322	-0.4673	0.5892	0.1543

parameters are obtained through joint estimation of high frequency data and daily option prices.

The test for over-identifying restriction is not rejected, however, as in the previous estimation the stationary condition $2k\theta \geq \sigma^2$ is violated.

7.8 Realised volatility, jumps and Microstructure noise correction

Using the theory developed earlier (See Section 7.2.2), the jump component is separated from the continuous part of the semimartingale. Figure 7.4 displays the squared sum of realised jumps ($\sum_{i=1}^{N(t)} J_i^2$) in each day for the sample period. What can be observed is that large squared sums corresponds to large deviation in the variability of returns with the largest jumps occurring during the recent period of market turmoil.

Graphical inspection indicates that the jump squared sums are not distributed evenly across the sample but rather are clustered around periods of high volatility. The jumps are obtained using Equation (7.10).

Before proceeding with the development of a GMM estimator that would allow to recover the parameter of the distribution of

Table 7.5: SVL model estimates

Stochastic volatility model estimates	
	S&P 500
	1997-2011
k	0.1135
	(0.0354)
θ	0.0103
	(0.0011)
σ	0.0834
	(0.0104)
ρ	-0.4625
	(0.1401)
GMM test of overidentifying restrictions	
$J-stat$	0.5538
$P\left[\mathcal{X}^2 > J\right]$	0.9069
$2k\theta \geq \sigma^2$	Reject

the Poisson process driving jumps, the next section will describe some asymptotic distribution theory that can be used in order to accurately test for the presence of jumps.

7.8.1 Asymptotic Distribution theory

So far the analysis relied on nonparametric jump estimates defined by the difference between realised volatility and bipower variation. This is theoretical justified when the partition $\triangle \to 0$, however, when the sampling frequency allows for $\triangle > 0$, the implementation of such nonparametric estimate is subject to measurement error. The non-negativity truncation in Equation (7.10) eliminates theoretical infeasible negative values for the squared jumps, however, it also likely allows for many small measurement errors to enter the estimate.

Equation (7.9) can be effectively used to develop a test for the presence of jumps. Huang and Tauchen (2005) used the following relative jump measure

$$RJ\left(T,t\right) = \frac{RV\left(T,t\right) - BV\left(T,t\right)}{RV\left(T,t\right)}$$

as an indicator of the contribution of jumps to the total realised

variance for a given day. Barnoff-Nielsen and Shepard (2004, 2006) found that under sufficient regularity, frictionless markets and in the absence of jumps in the price path

$$\sqrt{\triangle}\frac{RV\left(T,t\right)-BV\left(T,t\right)}{\left[\left(\mu_1^{-4}+2\mu_1^{-2}-5\right)\int_t^T\sigma^4\left(s\right)ds\right]^{1/2}}\xrightarrow{d}N\left(0,1\right),$$

as $\triangle\to 0$. This implies that for an abnormally large value of the standardised difference $RV\left(T,t\right)-BV\left(T,t\right)$ should be interpreted as a significant jump in the T,t time interval. In order to obtain the statistics there is the need to estimate the integrated quarticity. Andersen et al. (2005) suggest that integrated quarticity can be consistently estimated even in the presence in jumps by the realised tripower quarticity measure,

$$TP\left(T,t\right)=\triangle^{-1}\mu_{4/3}^{-3}\sum_{j=3}^{M}\mid r_{t+j\triangle,\triangle}\mid^{4/3}\mid r_{t+(j-1)\triangle,\triangle}\mid^{4/3}\mid r_{t+(j-2)\triangle,\triangle}\mid^{4/3} \tag{7.28}$$

where $\mu_{4/3}=2^{2/3}\cdot\Gamma\left(7/6\right)\cdot\Gamma\left(1/2\right)^{-1}=\mathbb{E}\left(\mid Z\mid^{4/3}\right)$ and $\Gamma\left(\cdot\right)$ is the Gamma function. It can be shown (Andersen et al., 2005) that

$$TP\left(T,t\right)\xrightarrow{p}\int_t^T\sigma_s^4ds. \tag{7.29}$$

The W_t statistics was found to over-reject the null of no jumps for large critical values. Huang and Tauchen (2005) after extensive simulation found that the z-test statistics gives reasonable power against several realistic stochastic volatility jump diffusion models. This ratio statistics can be written as

$$z_t=\triangle^{1/2}\frac{RV\left(T,t\right)-BV\left(T,t\right)}{\sqrt{\left(\mu_1^{-4}+2\mu_1^{-2}-5\right)max\left(1,\frac{TP(T,t)}{BV^2(T,t)}\right)}}, \tag{7.30}$$

which is very closely approximated by a standard normal distribution. Jumps are identified as the values of z_t in excess of the critical value Φ_α for some level of significance α, jumps are given by

$$J_\alpha^2\left(T,t\right)=\mathbb{I}\left[z_t>\Phi_\alpha\right]\left[RV\left(T,t\right)-BV\left(T,t\right)\right] \tag{7.31}$$

where $\mathbb{I}\left[\cdot\right]$ is an indicator function. Andersen et al. (2005) suggest, in order to ensure that the total variation is the sum of the

estimated continuous sample path and the jump component, the following component

$$C_\alpha(T,t) = \mathbb{I}[z_t \leq \Phi_\alpha] RV(T,t) + \mathbb{I}[z_t > \Phi_\alpha] BV(T,t) \quad (7.32)$$

for $\Phi_\alpha > 0$, equations (7.31) and (7.32) automatically guarantee that $J_\alpha^2(T,t)$ and $C_\alpha(T,t)$ are positive. The non-negativity truncation of Equation (7.10) corresponds directly to $\alpha = 0.5$ or $J_{t,0.5}$

7.9 Stochastic Volatility with Jumps in Returns (SVJ) Model

Realistic models for asset prices should consist of a continuous stochastic volatility component plus a jump component (Andersen et al., 2002). Several models have been introduced in the literature: Bates (1996) introduced one of the most popular models of this class where the Heston stochastic volatility model is extended by adding an independent jumps component. In the case of no drift the stochastic volatility model (7.14) with jumps returns is given by

$$\begin{aligned} dp_t &= \sqrt{V_t} dB_t + JdN(\lambda t) \\ dV_t &= k(\theta - V_t)dt + \sigma\sqrt{V_t} dW_t \end{aligned} \quad (7.33)$$

where the Brownian motion dB_t and the Poisson process $JdN(\lambda t)$ are independent, with jumps arriving at the exponential rate of λdt with jump size J, determined by the normal distribution (μ_J, σ_J^2). The jump parameters are identified by the following two moment conditions:

$$\begin{aligned} \mathbb{E}\left[J_{t+1}^2 \mid \mathcal{G}_t\right] &- J_t^2 \\ \mathbb{E}\left[J_{t+1}^4 \mid \mathcal{G}_t\right] &- J_t^4 \end{aligned} \quad (7.34)$$

The two moments for Gaussian jumps $J \sim N(0, \sigma_J^2)$, are introduced by using the moment generating function. With the assumption of $\mu = 0$ these are given by

$$\mathbb{E}\left[J^2 \mid \mathcal{F}_t\right] = \lambda \mathbb{E}\left(J^2 \mid \mathcal{F}_t\right) = \lambda \sigma_J^2 \quad (7.35)$$

$$\mathbb{E}\left[J^4 \mid \mathcal{F}_t\right] = \lambda \mathbb{E}\left(J^4 \mid \mathcal{F}_t\right) = 3\lambda \sigma_J^4 \quad (7.36)$$

Table 7.6: Monte Carlo estimation for SVJ model (T=1000)

	T	**Mean**				RMSE			
True value	1000	$\alpha_{0.5}$	$\alpha_{0.05}$	$\alpha_{0.01}$	$\alpha_{0.001}$	$\alpha_{0.5}$	$\alpha_{0.05}$	$\alpha_{0.01}$	$\alpha_{0.001}$
Estimation using high frequency data									
Observations per day: 82									
k	0.03	0.034	0.035	0.034	0.035	0.012	0.013	0.012	0.012
θ	0.25	0.236	0.240	0.247	0.245	0.055	0.053	0.0514	0.052
σ	0.10	0.098	0.098	0.099	0.100	0.020	0.020	0.018	0.019
λ	0.10	0.076	0.123	0.100	0.086	0.018	0.033	0.027	0.021
σ_{jump}	0.50	0.050	0.450	0.472	0.486	0.066	0.077	0.054	0.057

Table 7.7: Monte Carlo Estimation for SVJ model (T=4000)

	T	**Mean**				RMSE			
True value	4000	$\alpha_{0.5}$	$\alpha_{0.05}$	$\alpha_{0.01}$	$\alpha_{0.001}$	$\alpha_{0.5}$	$\alpha_{0.05}$	$\alpha_{0.01}$	$\alpha_{0.001}$
Estimation using high frequency data									
Observations per day: 82									
k	0.03	0.031	0.031	0.031	0.031	0.005	0.005	0.005	0.005
θ	0.25	0.245	0.245	0.247	0.248	0.026	0.027	0.026	0.027
σ	0.10	0.101	0.101	0.102	0.102	0.013	0.010	0.010	0.008
λ	0.10	0.172	0.120	0.096	0.082	0.075	0.023	0.010	0.019
σ_{jump}	0.50	0.417	0.455	0.480	0.498	0.087	0.054	0.035	0.035

7.9.1 Monte Carlo study

Tables 7.6 and 7.7 show Monte Carlo simulation for $\lambda = 0.1$ and $\sigma_{jumps} = 0.5$. These values are also used in Huang and Tauchen (2005). There is no current agreement for the value of α (Tauchen and Zhou, 2011), for example Andersen et al. (2007) employ various values for α in the same study. The following Monte Carlo experiments show the ability of extrapolating jumps from realised variance using the techniques discussed in section (7.2). As noted by Andersen et al. (2007) the level of $\alpha = 0.5$ corresponds to the non-negativity truncation of equation (7.10) .

Table 7.6 and 7.7 report Monte Carlo estimation for the SVJ model for the sample size of 1000 and 4000 respectively.

The Monte Carlo simulations show that the RMSE and mean of the estimated parameters varies for different values of α. It can be observed that the RMSE for λ and σ_{jump} is minimised in the region

Table 7.8: Stochastic volatility model with jumps (SVJ)

Stochastic volatility model (SVJ) **estimates**	
	S&P 500
	1997-2011
k	0.1153
	(0.0389)
θ	0.0097
	(0.0010)
σ	0.0796
	(0.0112)
$\lambda_{\alpha=0.01}$	0.1762
	(0.0319)
$\sigma_{J,\alpha=0.01}$	0.0545
	(0.0060)
GMM test of overidentifying restrictions	
$J-stat$	0.7282
$P\left[\mathcal{X}^2 > J\right]$	0.8666
$2k\theta \geq \sigma^2$	Reject

of $\alpha = 0.01$ and $\alpha = 0.001$, or correspondingly a 99% and 99.9% confidence level with a slightly bigger bias for the case of $\alpha = 0.001$. Consequently, in the empirical estimates that follows, α will be set to $\alpha = 0.01$.

7.9.2 Estimation results

Estimation for the stochastic volatility model with jumps is given by Table 7.8.

Estimates in Table 7.8 confirms mean reversion in k. The jump intensity λ is in the same order of magnitude [9] as for the maximum likelihood estimation in Andersen et al. (2002).

Figure 7.2, 7.3 and 7.4 shows the estimated Realised variance, the component series $C_{\alpha=0.01}$ and the jump's quadratic variation. It can be noted that the S&P 500 displays jump clustering around the recent market turmoil and that jumps tend to occur in periods of high volatility.

[9] These are not directly comparable as Andersen et al. (2002) obtain the value for λ from option prices. To the author's knowledge there are no previous studies for which a one-to-one comparison can be made.

From Figure 7.3. it is possible to note that an increase in jump activity is associated with an increase in realised variance: this in turn could suggest a stochastic volatility factor with jumps as in Eraker (2004) and Eraker *et al* (2003). It also provides support to Jacod and Todorov (2010) findings that most stock market jumps in the S&P 500 stock index are associated with volatility jumps.

7.10 Result summary

Table (7.9) displays the estimated parameters for the SV, SVL, SVLJ[10] models. The leverage effect is confirmed by the data and confirming previous studies as ρ is negative for the sample. It should be noted that jumps in returns change the size of the leverage effect. Jumps contribute to the leverage parameter for about 30% of its value and the ρ values suggest that jumps are correlated with returns.

Once the leverage parameter is estimated using realised volatility generated by the bipower variation estimator the leverage value falls. The leverage parameter ρ, when obtain using Bipower variation rather than quadratic variation takes value $\rho = -0.3516$ instead of $\rho = -0.4625$.

For all models the J-test is not rejected by the data while the stability condition is rejected. The implication of $2\theta k \geq \sigma^2$ for a Feller (1951) type of process is that the volatility process can reach zero.

7.11 Conclusion

The paper focuses on modelling realised volatility in the context of high-frequency data. The quadratic variation for the jump component is obtained by subtracting bipower variation from the realised volatility measure. Monte Carlo estimation shows that the proposed moments work well and the GMM estimation is able to recover the parameters driving the stochastic differential equation (SDE) with low RMSE for all proposed stochastic volatility models.

The parameters obtained are in line to the ones obtained in the literature, both in terms of signs and magnitude. The leverage

[10]The SVLJ is the SVJ model with added leverage moment as in equation (7.27) .

Table 7.9: Summary

Estimation summary				
	SV	SVL	SVJ	SVLJ
	S&P 500	S&P 500	S&P 500	S&P 500
	1997-2011	1997-2011	1997-2011	1997-2011
k	0.1135	0.1135	0.1153	0.1153
	(0.0353)	*(0.0354)*	*(0.0389)*	*(0.0389)*
θ	0.0103	0.0103	0.0097	0.0097
	(0.0011)	*(0.0.0011)*	*(0.0010)*	*(0.0010)*
σ	0.0840	0.0834	0.0796	0.0796
	(0.0104)	*(0.0104)*	*(0.0112)*	*(0.0112)*
ρ		-0.4625		-0.3510
		(0.1401)		*0.1027*
λ			0.1762	0.1762
			(0.0319)	*(0.0319)*
σ_J			0.0545	0.0545
			(0.0060)	*(0.0060)*
GMM test of overidentifying restrictions				
$J-stat$	0.5919	0.5538	0.7282	0.7280
$P\left[\mathcal{X}^2 > J\right]$	0.8983	0.9069	0.8666	0.8666
$2k\theta \geq \sigma^2$	Reject	Reject	Reject	Reject

effect displays a lower value when bipower variation is applied to the series, indicating that a large proportion of the leverage effect is due to jumps. This finding additionally provides evidence that jumps are negatively correlated with returns.

Estimation for the stochastic volatility model, the stochastic volatility model with leverage and the stochastic volatility model with jumps in returns are rejected due to the stability condition not being satisfied.

The common feature for all tested model is a high value for the parameter k. This high mean reversion is potentially the cause for the rejection of the stability condition.

The instability in the SDE that drives volatility could potentially indicate a two factor model. While a volatility process that is driven by two SDEs would be a natural extension to the work reported above, it should be noted that due to the number of parameters involved, estimation of a two factor model, requires very efficient algorithms not usually available with standard software and in even such cases it is not possible to obtain sensible estimates using standard GMM techniques [11]. Potential future work will involve methodologies that avoid the 'hill climbing algorithms' commonly employed in classical estimation. A possibility is the Laplace type estimators based on Markov Chain Monte Carlo proposed by Chernozhukov and Hong (2003).

The S&P 500 displays jump clustering around the recent market turmoil with large jumps that tend to occur in periods of high volatility. This finding provides support for jump clustering, time-varying jump frequency and jump size distribution. Although possible time-varying jumps have been recently indicated by Tauchen and Zhou (2011), the estimates obtained above provide much stronger support to this argument. The size of the leverage effect, in particular, suggests large negative jumps occur more often than positive jumps.

As increases in jump activity are associated with an increase in realised variance, further modelling should involve a jump compo-

[11] The two factor model proposed by Bollerslev and Zhou (2002) and estimated by the author using MATLAB ©did not produce any sensible estimates. Although such two-factor model estimation is theoretically possible, it is likely that estimation based on GMM and analytical moments is too complex and standard algorithms such as the Newton-Raphson method would fail to converge.

nent in the volatility process that is correlated with the jump component present in the SDE driving returns. One possible choice is a stochastic volatility factor that is Levy-driven such as the Barnoff-Nielsen and Shepard (2001, 2002) model.

7.12 Figures

Figure 7.1: Simulated Path for SDE

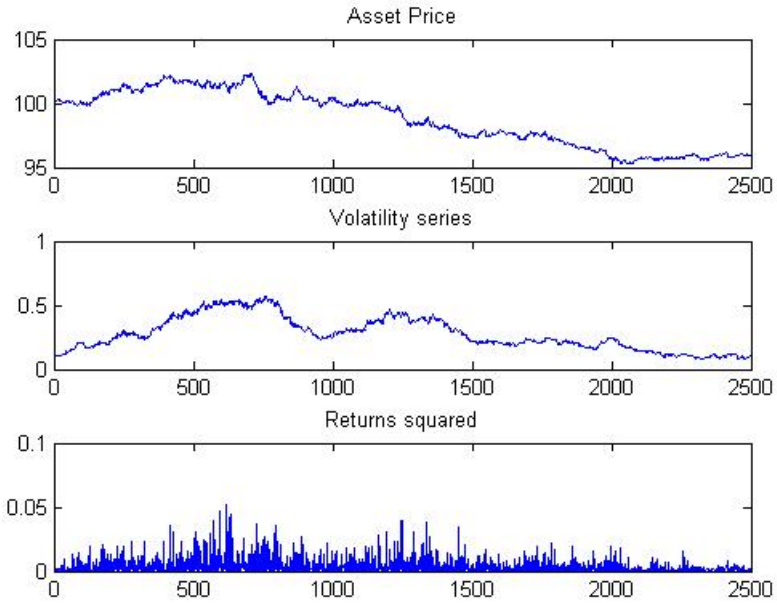

Figure 7.2: Realized Variance of S&P 500

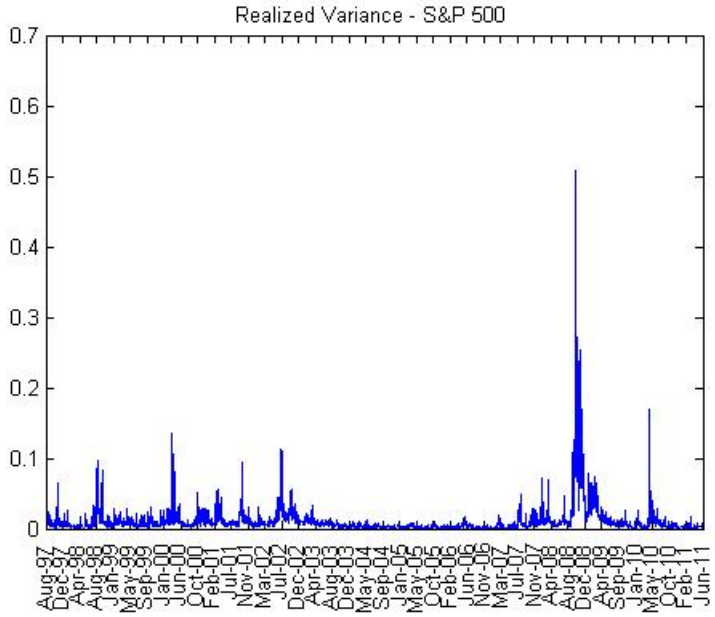

Figure 7.3: Bipower Variation and Realized Variance Component

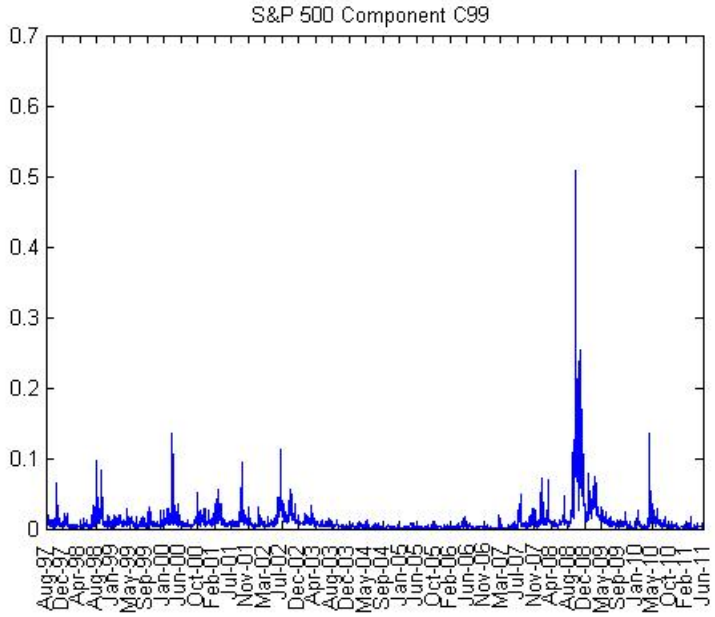

Figure 7.4: Realized jumps $\alpha = 0.01$

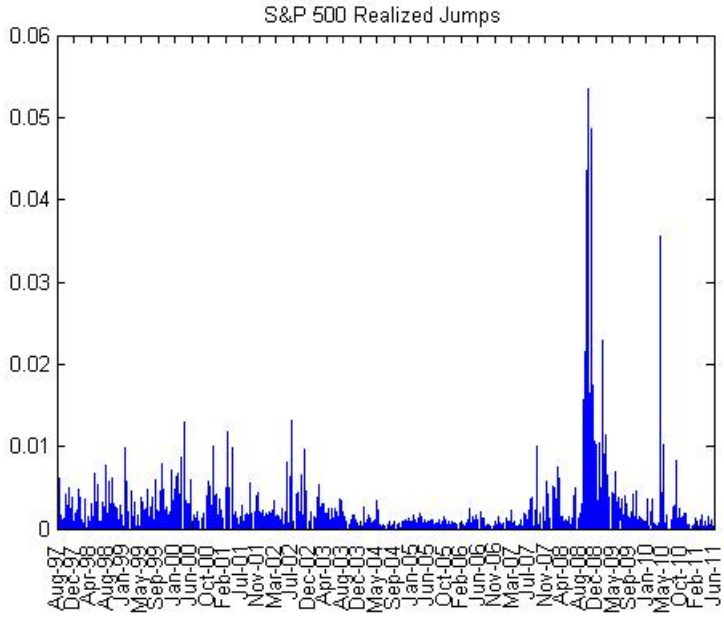

Bibliography

Ades, A.E.; Lu, G., and Claxton, K. Expected value of sample information calculations in medical decision modelling. *Medical decision making*, 24:207–27, 2004.

Andersen, T. G and Benzoni, L. *Handbook of Financial Time Series*. Springer-Verlag, Berlin, Germany, 2009.

Andersen, T. G. and Bollerslev, T. Intraday periodicity and volatility persistence in financial markets. *Journal of empirical finance*, 4:115–158, 1997.

Andersen, T. G.; Bollerslev, T.; Diebold, F.X., and Labys, P. The distribution of realised exchange rate volatility. *Journal of the American Statistical Association*, 96:42–55, 2001.

Andersen, T. G; Benzoni, L., and Lund, J. An empirical investigation of continuous-time equity return models. *Journal of Finance*, 57:1239–1284, 2002.

Andersen, T. G.; Bollerslev, T.; Diebold, F.X., and Labys, P. Modelling and forecasting realized volatility. *Econometrica*, 71:579–626, 2003.

Andersen, T. G.; Bollerslev, T., and Diebold, F.X. *Handbook of Financial Econometrics*. North-Holland, 2005.

Andersen, T. G.; Bollerslev, T., and Diebold, F.X. Roughing it up: including jump components in the measurement, modelling and forecasting of return volatility. *The review of economics and statistics*, 89:701–720, 2007.

Armitage, P. Sequential tests in prophylactic and therapeutic trials. *Quarterly J. Med*, 23:255–274, 1954.

Armitage, P. Numerical studies in the sequential estimation of a binomial parameter. *Biometrika*, 45:1–15, 1958.

Armitage, P. *Sequential medical trials*. Oxford, Blackwell, 1975.

Aroian, L.A. Sequential analysis, direct method. *Technometrics*, 10:125–132, 1968.

Aroian, L.A. and Robinson, D.E. Direct methods for exact truncated sequential tests of the mean for a normal distribution. *Technometrics*, 11:661–675, 1969.

Back, K. Asset pricing for general processes. *Journal of mathematical economics*, 20:371–395, 1991.

Barnoff-Nielsen, O.E. and Shepard, N. Non-gaussian ou based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B*, 63:167–241, 2001.

Barnoff-Nielsen, O.E and Shepard, N. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of royal statistical society: Series B*, 64(2):253–280, 2002.

Barnoff-Nielsen, O.E. and Shepard, N. Power and bipower variation with stochastic volatility and jumps. *Journal of financial econometrics*, 2:1–37, 2004.

Barnoff-Nielsen, O.E and Shepard, N. Impact of jumps on returns and realised variances: econometric analysis of time-deformed levy process. *Journal of econometrics*, 131:217–252, 2006.

Bates, D.S. Jumps and stochastic volatility: exchange rate process implicit in deutche mark options. *The review of financial studies*, 9:69–107, 1996.

Baum, C.W. and Veeravalli, V.V. A sequential procedure for multihypothesis testing. *IEEE Trans. Inform. Theory*, 40:1994–2007, 1994.

Berger, J. O. *Statistical decision theory and Bayesian analisys*. Springer-Verlag. New York, 1985.

Berger, J.O. and Berry, D.A. Statistical analysis and the illusion of objectivity. *American scientist*, 76(2):159–165, 1988.

Berry, D. A. Bayesian clinical trials. *Nature review drug discovery*, 5:27–36, 2006.

Berry, D. A. and Fristedt, B. *Bandit Problems*. Springer-Verlag. Berlin, Heidelberg, 1985.

Berry, D.A. Interim analysis in clinical trials: Classical vs. bayesian approaches. *Statistics in medicine*, 4:521–6, 1985.

Berry, D.A. Interim analysis in clinical trials: the role of the likelihood principle. *The American statistician*, 41:117–122, 1987.

Berry, D.A. and Ho, C-H. One-sided sequential stopping boundaries for clinical trials: a decision-theoretic approach. *Biometrics*, 44: 219–227, 1988.

Black, F. and Scholes, M. The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81(3):637–654, 1973.

Bollerslev, T. and Zhou, H. Estimating stochastic volatility diffusion using conditional moments of integrated volatility. *Journal of econometrics*, 109:33–65, 2002.

Bollerslev, T.; Tzuo, H., and Tauchen, G. Risk, jumps and diversification. *Journal of econometrics*, 144:234–256, 2008.

Bregantini, D. and Thijssen, J. J. J. On a simple quickest detection rule for health-care technology assessment. *University of York working paper*, 2013.

Bross, I. Sequential medical plans. *Biometrics*, 8:188–205, 1952.

Bross, I. Sequential clinical trials. *Journal of chronic diseases*, 1958.

Bubekc, S. and Cesa-Bianchi, N. Regret analysis of stochastic and nonstochastic mlti-armed problems. *Foundations and trends in machine learning*, 5(1):1–122, 2012.

Burdette, W.J. and Gehan, E.A. *Planning and analysis of clinical studies*. Charles C. Thomas, 1970.

Capinski, M. and Kopp, E. *Measure, integral and probability*. Springer, London, 2004.

Chernoff, H. *Sequential tests for the mean of a normal distribution*. Proc. 4th Berkley Symp. Mathematical Statistics and Probability, 1961.

Chernozhukov, V. and Hong, H. An mcmc approach to classical estimation. *Journal of econometrics*, 115:293–346, 2003.

Chevret, S. Bayesian adaptive trials: a dream for statisticians only ? *Statistics in medicine*, 31:1002–1013, 2012.

Chow, S-C and Liu, J-P. *Design and analysis of bioavailability and bioequivalence studies*. New York: Marcel Dekker, 1992.

Claxton, K. The irrelevance of inference: a decision-making approach to the stochastic evaluation of health care technologies. *Journal of health economics*, 18:341–364, 1999.

Claxton, K. and Posnett, J. An economic approach to clinical trial design and research priority setting. *Health economics*, 5:513–524, 1996.

Claxton, K. and Thompson, K. M. A dynamic programming approach to the efficient design of clinical trials. *Journal of health economics*, 20:797–822, 2001.

Claxton, K.; Longworth, S. Palmer L.; Bojke, L.; Griffin, S.; McKenna, C.; Soares, M.; Spackman, E., and Youn, J. Informing a decision framework for when nice should recommend the use of health technologies only in the context of an appropriately designed programme of evidence development. *Health technology assessment*, 16(46):NA, 2012.

Close, A.; Robertson, C.; Rushton, S.; Shirley, M.; Vale, L.; Ramsay, C., and Pickard, R. Comparative cost-effectiveness of robot-assisted and standard laparoscopic prostatectomy as alternatives to open radical prostatectomy for treatment of men with localised prostate cancer. *European urology*, 64(3):361–369, 2013.

Cornfield, J. Sequential trial, sequential analysis and the likelihood principle. *The American statistician*, 20:18–23, 1966.

Corradi, V. and Distaso, W. Semiparametric comparison of stochastic volatility models using realised measures. *Review of economic studies*, 73:635–667, 2006.

Cox, D.R. and Hinkley, D.V. *Theoretical statistics*. Chapman and Hall, 1974.

Cox, J.; Ingersoll, J., and Ross, S. An intertemporal general equilibrium model of asset prices. *Econometrica*, 53:363–384, 1985.

Cutler, S. J.; Greehouse, S.W.; Cornfield, J., and Schneiderman, M.A. The role of hypothesis testing in clinical trials. *Journal of chronic diseases*, 19:857–882, 1966.

Dixit, A. K. and Pindyck, R. S. *Investment under uncertainty*. Princeton University Press, 1994.

Driffield, T. and Smith, P.C. A real option approach to watchful waiting: theory and illustration. *Medical decision making*, 27(2): 178–188, 2007.

Eckermann, S. and Willam, A. Expected value of information and decision making in hta. *Health Economics*, 16:195–209, 2007.

Eckermann, S. and William, A. R. The option value of delay in health technology assessment. *Medical decision making*, 28:300–305, 2008.

Feller, W. Two singular diffusion problems. *Annals of mathematics*, 54:173–182, 1951.

Forster, M. and Pertile, P. Optimal decision rules for HTA under uncertainty: a wider, dynamic perspective. *Health Economics*, 22 (12):1507–14, 2012.

Garcia, R.; Lewis, M.; Pastorello, S., and Renault, E. Estimation of objective and risk-neutral distributions based on moments of integrated volatility. *Journal of econometrics*, 160(22-32), 2011.

Glick, H. A.; Doshi, J. A., and Polsky, S. S. Sonnad D. *conomic evaluation in clinical trials*. Oxford University Press, 2007.

Griffin, S. C.; Claxton, K. P.; Palmer, J. S., and Schulpher, J. M. Dangerous omissions: the consequences of ignoring decision uncertainty. *Health economics*, 20:212–224, 2011.

Hamilton, J.D. *Time series analysis*. Princeton University Press, 1994.

Hansen, L.P. Large sample properties of generalised method of moments estimators. *Econometrica*, 50(1029-1054), 1982.

Hansen, L.P. and Lunde, A. Realised variance and market microstructure noise. *Journal of business and economic statistics*, 24(127-218), 2006.

Heston, S. A closed-form solution for options with stochastic volatility with application to bond and currency options. *Review of financial studies*, 6(327-343), 1993.

Hjelmgren, J.; Berggren, F., and Andersson, F. Health economic guidelines-similarities, differences and some implication. *Value in health*, 4(3):225–250, 2001.

Huang, X. and Tauchen, G. The relative contribution of jumps to total price variance. *Journal of financial econometrics*, 3:456–499, 2005.

Jacod, J. and Todorov, V. Do price and volatility jump together. *The annals of applied probability*, pages 1425–1429, 2010.

Jennison, C. and Turnbull, B. W. *Group sequential methods with application to clinical trials*. Chapman, 2000.

Lindley, D.V. *Introduction to probability and statistics from a Bayesian viewpoint; part 2, inference*. Cambrdige University Press. Cambridge, UK., 1965.

McKenna, C. and Claxton, K. Addressing adoption and research design decisions simultaneusly: the role of value of sample information analysis. *Medical decision making*, 31:853–865, 2011.

Meddahi, N. An eigenfunction approach for volatility modelling. *The university of Montreal - Manuscript*, 2001.

Mehta, C.R. and Cain, K.C. Charts for early stopping of pilot studies. *Journal of clinical oncology*, 2:679–682, 1984.

Meltzer, D. O and Smith, P. C. *Handbook of health economics*, volume 2, chapter vii, page 433. Elsevier, 2012.

Merton, R.C. Lifetime portfolio selection under uncertainty. *Review of economic and statistics*, 51:247–257, 1969.

Merton, R.C. Optimum consumpiton and portfolio rules in continuous time model. *Journal of economic theory*, 3:373–413, 1971.

Merton, R.C. An intertemporal capital asset pricing model. *Econometrica*, 41:8678–87, 1973.

Merton, R.C. Estimating the expected return on the market. *Journal of financial economics*, 8:1–39, 1980.

O'Brien, P.C. and Fleming, T.R. A multiple testing procedure for clinical trials. *Biometrics*, 40:549–556, 1977.

Oksendal, B. *Stochastic differential equations: an introduction with application*. Springer, New York, 2010.

Palmer, S. and Smith, P. C. Incorporating option value into the economic evaluation of health care technologies. *Journal of Health Economics*, 19:755–766, 2000.

Pertile, P.; Forster, M., and LaTorre, D. Optimal bayesian sequential sampling rules for the economic evaluation of health technologies. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, pages n/a–n/a, 2013.

Peskir, G. and Shiryaev, A. N. *Optimal stopping and free boundary problems*. Birkhauser Verlag, Switzerland., 2006.

Philips, Z.; Claxton, K., and Palmer, S. The half-life of truth: what are appropiate time horizons for research decisions. *Health economics*, May-June(28):287–299, 2008.

Pocock, S.J. 1977. *Biometrika*, 64:191–199, 1977.

Pocock, S.J. and Hughes, M.D. Practical problems in interim analyses, with particular regard to estimation. *Controlled clinical trials*, 10:209S–221S, 1989.

Poor, V. H. and Hadjiliadis, O. *Quickest detection*. Cambridge University Press, 2009.

Pratt, W. J.; Raiffa, H., and Schlaifer, R. *Statistical decision theory*. MIT Press, 1995.

Raiffa, H. and Schlaifer, R. *Applied statistical decision theory*. Clinton press. Boston, 1961.

Ramsey, S.; Wilke, R.; A., Briggs; Brown, R.; Buxton, M.; Chawla, A.; Cook, J.; Glick, H.; Lijas, B.; Petitti, D., and reed, S. Good

research practices for cost-effectivness analysis alongside clinical trials. *Value in health*, 8(5):521–33, 2005.

Rosenbaum, P.R. and Rubin, D.B. Sensitivity of bayes inference with data dependent sampling rules. *The American statistician*, 38:106–109, 1984.

Samuelson, P. Rational theory of warrant pricing. *Industrial management*, 6:13–31, 1965.

Schwartz, E.S. and Trigeorgis, L. *Real options and investment under uncertainty*. MIT Press, 2001.

Sculpher, M. and Claxton, K. Establishing the cost-effectivness of new pharmaceuticals under conditions of uncertainty-when is there sufficient evidence. *Value in health*, 8(4):433–445, 2005.

Shepard, K. and Shepard, N. Realising the future: forecasting with high frequency based volatility (heavy) models. *Journal of applied econometrics*, 25:197–231, 2010.

Shiryaev, A. N. *Optimal stopping rules*. Springer-Verlag. Berlin, Heidelberg, 1978.

Siegmund, D. *Sequential anlysis*. Springer-Verlag. New York, 1985.

Spiegelhalter, D.J. and Freedman, L.S. A predictive approach to selecting the size of a clinical trial, based on subjective clinical optinion. *Statistics for medicine*, 5:1–13, 1986.

Steele, Michael J. *Stochastic calculus and financial applications*. Springer-Verlag New York, Inc, 2001.

Sundaresan, S.M. Continuous-time methods in finance: a review and assessment. *Journal of finance*, 55(4):1569–1622, 2000.

Tauchen, G. and Zhou, H. Relised jumps on financial markets and predicting credit spreads. *Journal of econometrics*, 160:102–108, 2011.

Thijssen, J. J. J. An easy-to-use toolkit for solving optimal stopping problems. *Working paper University of York*, 2013.

Todorov, V. Estimation of continuous-time stochastic volatility models with jumps using high-frequency data. *Journal of econometrics*, 148:131–148, 2009.

Wald, A. *Sequential analysis*. New York: Wiley, 1947.

Wald, A and Wolfowitz, J. Optimum character of the sequential probability ratio test. *Ann. Math. Statist.*, 19:326–339, 1948.

Wang, S.K. and Tsiatis, A.A. Approximately optimal one-parameter boundaries for group sequential trials. *Biometrics*, 43: 193–200, 1987.

William, A. and Kowgier, M. Determining optimal sample sizes for multi-stage randomized clinical trials using value of information methods. *Clinical Trials*, 5:289–300, 2008.

William, A. R. and Pinto, E. M. The value of information and optimal clinical trial design. *Statistics in medicine*, 24:1791–1806, 2005.

Williams, D. *Probability with martingales*. Cambridge University Press, 2010.

Yin, G. *Clinical trial design: Bayesian and frequentist adaptive methods*. Wiley, U.S.A, 2012.

Zhou, H. A study of finite sample properties of emm, gmm, qmle, and mle for square-root interest rate diffusion model. *Journal of computational finance*, 5(2):89–122, 2001.