Continuous-Time Stochastic Analysis of Optimal Experimentation and of Derivative Asset Pricing

Sven Rady

London School of Economics and Political Science

Thesis submitted for the degree of Doctor of Philosophy

Abstract

This thesis applies continuous-time stochastic techniques to problems in economics of information and financial economics.

The first part of the thesis uses non-linear filtering and stochastic control theory to study a continuous-time model of optimal experimentation by a monopolist who faces an unknown demand curve subject to random changes. It is shown that different probabilities of a demand curve switch can lead to qualitatively very different optimal behaviour. Moreover, the dependence of the optimal policy on these switching probabilities is discontinuous. This suggests that a market or an economy embedded in a changing environment may alter its behaviour dramatically if the volatility of the environment passes a critical threshold.

The second part of the thesis studies continuous-time models of derivative asset pricing. First, a review of the so-called direct approach to debt option pricing emphasises the principal modelling problems of this approach and highlights the shortcomings of certain models proposed in the literature.

Next, the connection between martingale measures and numeraire portfolios is exploited in problems of option pricing with strict upper and lower bounds on the underlying financial variable. This leads to a new decomposition of option prices in terms of exercise probabilities calculated under particular martingale measures and allows a simple proof of certain generalisations of the Black-Scholes option price formula.

Finally, martingale methods are used to examine pricing formulae for general contingent claims, yielding a new method for inferring state prices from a given pricing formula. It is shown that if price processes are continuous semimartingales and the pricing formula is sufficiently regular, then the latter uniquely determines the risk-neutral law of the underlying asset price.

Contents

List of Figures

Acknowledgements

I could not have written this thesis without the support that I enjoyed over the last years, first at the University of Bonn, then at the LSE.

I am particularly indebted to Lucien Foldes, my supervisor at the LSE, for his advice, encouragement and help. I benefited immensely from his teaching and his research which gave me the tools for all my subsequent work. I would also like to thank him for numerous stimulating discussions on finance and economics that have influenced my thinking on many issues.

The LSE Financial Markets Group (FMG) offered an ideal environment, both for my work and for me personally. I am especially grateful to its director, David Webb, for the generous help he gave me wherever he could. My work benefited particularly from the FMG student seminar, run by Margaret Bray, which provided an invaluable forum for the discussion of ideas at an early stage.

I am also indebted to Urs Schweizer for encouraging me to join the European Doctoral Program in Quantitative Economics (EDP). During my first year in the program, I was based in the Department of Statistics at the University of Bonn, and I would like to thank its director, Dieter Sondermann, and his assistant, Klaus Sandmann, for introducing me to the theory of option pricing.

Chapter 1 and 2 of the thesis grew out of joint projects with Godfrey Keller (also from the FMG) and Klaus Sandmann, respectively. Their enthusiasm was a great inspiration to me, and I would like to thank them for their excellent cooperation. Chapter 1 of the thesis benefited greatly from comments and suggestions by Patrick Bolton, Margaret Bray, Phil Dybvig and John Hardman Moore. I am particularly grateful to Phil Dybvig for a wonderful stay at Washington University in St. Louis during which the work on Chapter 1 advanced considerably. Chapter 2 benefited from the suggestions of an anonymous referee, and Chapter 3 from a conversation with Steve Heston. Bruce Grundy and Stanley Pliska gave helpful comments on Chapter 4.

Over a period of two years, I worked part-time in the Swaps and Options Research Team at Paribas Capital Markets, London. I would like to thank Bruno Dupire and his colleagues for interesting discussions and the opportunity to confront theoretical models with the reality of the derivatives business.

I am grateful for financial support from the Deutsche Forschungsgemeinschaft (Sonderforschungsbereich 303), the German Academic Exchange Service (DAAD), the COLONIA Studienstiftung, the LSE Financial Markets Group, and the Economic and Social Research Council (ESRC).

I would like to thank Dirk Schoenmaker, who shared an office with me during four long years, for his patience, Denis Gromb for introducing me to LaTeX, and all my friends at the LSE for the good times we had.

Finally, this is the moment to thank Birgit for her wonderful understanding and encouragement. This thesis is dedicated to her.

Introduction

Many economic or financial activities are characterised by the need to make intertemporal decisions in the presence of continuing uncertainty. A firm which plans to introduce a new product, for example, is typically uncertain about the demand curve it will face. Even after the product has been introduced, uncertainty will persist as the demand can vary over time due to the emergence of rival products, a change in tastes, and many other factors. An indvidial's decision how much of his wealth to consume now, how much to save, and how to allocate his savings across different investment opportunities, is another example. This decision is affected by uncertainty about the returns that the investments will generate. On a larger scale, the pension fund to which the individual may contribute is confronted with the same uncertainty, and so are numerous other financial institutions. While examples obviously abound, it is nevertheless possible to identify certain basic types of intertemporal decision problems under uncertainty. Two of these problems are studied in this thesis.

Part I of the thesis consists of a single chapter which is concerned with optimal learning by experimentation in a multi-period setting, a problem which has attracted considerable interest in the microeconomic literature. The experimentation problem arises whenever an agent who has to choose a sequence of actions is uncertain about the distribution of the rewards that his actions will generate. In such a situation, a rational agent will take into account not only the current reward that an action is likely to produce, but also the potential information content of the outcome. Some actions, for instance, might yield a relatively low expected payoff, but reveal relatively precise information about the underlying distribution of rewards.

This information would allow the agent to make better decisions and reap higher payoffs in the future. The agent might therefore want to experiment, that is, to choose actions which are suboptimal in terms of expected current revenue alone. Clearly, this involves an opportunity cost in the form of current revenue forgone. The problem of optimal experimentation is to find the optimal trade-off between this opportunity cost and the long-term benefit resulting from the information acquired.

The seller in the above example faces an experimentation problem. In each period, he puts some quantity of the good on the market and earns a profit which is random due to factors beyond his control. The distribution of this profit depends on the unknown demand curve for the good. In this situation, a myopic seller would just try to maximise current expected profit in each period. The truly optimal strategy, by contrast, may require the seller to give up some expected profit now in order to learn more about the demand curve and then achieve higher profits in the future. For the sake of concreteness, our analysis of optimal experimentation in Part I is cast in terms of this particular problem.

Two approaches to the problem of learning by experimentation in a single-agent setting have emerged in the literature. One approach is to restrict the analysis to a two-period framework, and then to determine in which way the ability to gather information in the first period (which will be useful in the second period) affects the agent's behaviour. Examples of this approach can be found in Mirman, Samuelson and Urbano (1993) and other papers referenced in their introduction.

The second approach formulates an infinite-horizon model, in which case it is natural to look at the limiting behaviour of the agent. The first such model in the economics literature is due to Rothschild (1974), and has subsequently been extended in a number of different directions; see for example McLennan (1984), Easley and Kiefer (1988), Kiefer (1989), and Aghion, Bolton, Harris and Jullien (1991). A common result of these models is that the agent's beliefs about the underlying distribution of rewards converge, in which case experimentation will cease and no further information will be gathered – in the limit, the agent will learn everything that is *worth* knowing. The question then arises as to whether the beliefs converge to a one-point distribution at the true parameter value. Typically, the answer is that with positive probability they do not.

We take the second approach as our starting point. However, whereas the above papers assume there to be an unknown but *fixed* distribution of rewards, we follow Kiefer (1991) and allow this distribution to change randomly over time. Despite the fact that these variations are highly stylised abstract models of information acquisition, it does seem more realistic to allow for the possibility that new data continues to be pertinent. In this case, beliefs may continue to evolve, and the agent is not doomed to take the same action for ever more.

We depart from Kiefer's model (and the overwhelming majority of similar models in the economics literature) by formulating the problem in continuous time.[1] This is motivated by the following considerations. The experimentation problem can be reformulated as a problem of optimal control with the agent's posterior belief as state variable. Working in continuous time allows us to apply techniques from non-linear filtering theory and optimal control theory for diffusion processes. This leads us via the Bellman equation to an ordinary differential equation for the value function. Although there is little hope for closed-form solutions, solving the differential equation numerically proves to be easier than calculating a numerical solution to the fixed point problem for the Bellman operator which arises in discrete-time settings.

Our main finding is that, in certain scenarios, the optimal behaviour depends *qualitatively* on the switching probabilities. More precisely, a small variation in the likelihood of switches of the demand curve can cause a *discontinuous* change in the optimal policy. This phenomenon is novel in the economic literature on experimentation.

The result is potentially of great significance for economic theory. It suggests that a market or an economy embedded in a changing environment will alter its behaviour dramatically if, in the eyes of the economic agents, the volatility of the environment passes a critical threshold. Thus, a slight increase in variability may

[1]Previously, such a formulation has only been adopted by Bolton and Harris (1993) and Felli and Harris (1994) who study multi-agent learning problems with a fixed distribution of rewards.

not just lead to a moderate reduction in information gathering activities – it could in fact provoke a near cessation of these activities, with drastic consequences for the process of information aggregation.

Part II of the thesis, consisting of Chapters 2 - 4, addresses an intertemporal decision problem that has been at the centre of modern finance: the hedging and pricing of derivative securities.[2]

The need to hedge, that is, to cover part or all of a liability by an offsetting asset position, arises in many financial institutions. A pension fund, for example, must pay out pensions at predetermined dates and will gear its portfolio towards meeting these obligations. Similarly, an investment bank that has sold derivative securities to corporate clients will try to hedge its position by building up a portfolio of securities that produces enough cash flows to satisfy the clients' claims. A *perfect* hedge requires that a claim be matched exactly by the payoff of the hedging portfolio. The construction of such a perfect hedge is the central problem of derivative assets analysis.

In fact, if the hedging problem is solved, then the problem of finding the fair premium for a derivative asset is solved as well. More precisely, a claim that admits a perfect hedge must trade at a price equal to the current value of the hedging portfolio. Every other price for the claim would lead to arbitrage opportunities which rational agents could exploit by selling the claim and buying the portfolio if the claim is more expensive, and vice versa. Of course, perfect hedges only exist in an idealised world without market frictions such as transaction costs or borrowing constraints. However, the perfect markets paradigm provides a useful benchmark in theory as well as in practical applications, and we shall maintain it throughout the second part of the thesis.

Another prerequisite for the hedging argument to work is that the primitive securities themselves do not give rise to arbitrage opportunities. An arbitrage opportunity is a portfolio strategy with negative initial investment, but non-negative value later on, thus providing a gain today without creating any liabilities to-

[2]A derivative security is an asset whose payoff is entirely determined by the prices or payoffs of some underlying securities. Options on stocks are prominent examples of derivatives.

morrow. Together with suitable restrictions on the set of portfolio strategies that investors can use, the existence of a so-called martingale measure is a sufficient condition for the absence of arbitrage.[3] This is a new probability measure under which the prices of all primitive securities, expressed in units of some numeraire asset, are martingales. Under such a measure, the value of a hedging portfolio, measured in units of the numeraire, is again a martingale. In particular, the initial portfolio value equals the expectation of the terminal portfolio value under the martingale measure. Since taking expectations preserves non-negativity, a portfolio strategy with non-negative final value must have a non-negative initial investment. In other words, if there exists a martingale measure, arbitrage opportunities are precluded. Moreover, the value of a derivative that admits a perfect hedge can be calculated by taking the expectation of its discounted payoff under a martingale measure.

Starting with Black and Scholes (1973) and Merton (1973), derivative assets analysis has constantly relied on continuous-time models. On the one hand, these models have the advantage of not prescribing trading dates in advance, which makes them a much better approximation of real market activities than discrete-time models. On the other hand, continuous time allows the use of mathematical results which do not exist, or are not as powerful, in discrete time, such as Itô's change-of-variable formula, certain martingale representation theorems, and Girsanov's theorem on changes of the probability measure. Applying these results, the literature on derivative securities has developed a multitude of models in which derivatives can be perfectly hedged and, by the above arbitrage argument, priced. Often the prices of the underlying securities are modelled as diffusion processes. In this case, the price of a derivative asset will be given by a pricing formula, that is, a deterministic function of the underlying security prices and time. Moreover, this function can be calculated as a solution to a certain partial differential equation of parabolic type, the so-called fundamental valuation equation. The formula of Black and Scholes (1973) for options on stocks is the first and foremost example of such a pricing formula.

[3] Cf. Harrison and Kreps (1979), Harrison and Pliska (1981), Dybvig and Huang (1988).

Chapter 2 of the thesis critically reviews the so-called direct approach to debt option pricing which attempts to extend the Black-Scholes framework to the valuation of options on bonds. In fact, the valuation of these options has been approached from two different angles. The term structure approach regards bonds and bond options as interest rate dependent derivatives. Consequently, it formulates the pricing problem in the framework of a term structure model, that is, a model which describes the evolution of the term structure of interest rates over time.[4] This requires a consistent specification of the price processes of *all* traded zero-coupon bonds, usually as a function of one or more state variables such as the short term interest rate.[5] The direct approach, by contrast, regards the bond on which the option is written as a primitive security in its own right, and tries to follow the Black-Scholes paradigm as closely as possible. In particular, the direct approach models only those securities which are of immediate relevance to the pricing problem at hand, without relating them to the term structure as a whole or to a system of state variables. Thus, the stock price process of the Black-Scholes model is just replaced by a bond price process.

Our review of the direct approach emphasises its main modelling problems: first, the problem of specifying bond price processes that reach par value at maturity with probability one; second, the problem of precluding negative bond yields and negative implied forward rates; third, the problem of ensuring the absence of arbitrage opportunities between the bonds.

Our analysis highlights the shortcomings of some of the models proposed in the literature. In fact, Bühler and Käsler (1989) construct the only model within the direct approach that solves the three problems mentioned above, and still provides closed-form expressions for option prices. Lognormal bond price models, such as Ball and Torous (1983) and Kemna, de Munnik and Vorst (1989), fulfil the par value condition and lead to option price formulae of the same type as in Merton (1973), but they clearly assign a positive probability to the occurrence of negative interest rates. In view of this, Schöbel (1986) proposes a method of modifying Merton

[4]Cf. Vasicek (1977), Cox, Ingersoll and Ross (1985), Heath, Jarrow and Morton (1992).

[5]See Duffie and Kan (1992), El Karoui, Myneni and Viswanathan (1992), Cheyette (1994).

type option price formulae by imposing an additional boundary condition which is necessary for non-negative interest rates.[6] We show that his derivation of the modified formulae is unduly complicated. A more serious flaw of his method is that he does not develop an underlying bond price model in which the proposed option prices could indeed be derived by the standard hedging argument. As a first step towards identifying such a model, we follow Breeden and Litzenberger (1978) and calculate the Arrow-Debreu or state prices implicit in these option prices. We find a positive Arrow-Debreu price for the event that a certain interest rate vanishes at the expiry date of the option. Economic intuition suggests that this event ought to have probability zero. This indicates that Schöbel, while "correcting" for negative interest rates on the level of option prices, implicitly accepts a highly implausible behaviour of bond prices. Further analysis in Chapter 4 will support this view.

Chapter 3 investigates the valuation of options when the underlying financial variable has the following two characteristics: (i) the process has natural upper and lower boundaries; (ii) its diffusion coefficient is quadratic in the current value of the variable. The bond price model of Bühler and Käsler (1989) falls into this category, and so does an exchange rate model which Ingersoll (1989a, b) develops to price currency options in a perfectly credible target zone regime. This type of model can be regarded as a generalisation of the Black-Scholes model; in fact, the latter is obtained by choosing 0 and $+\infty$ as the lower and upper boundaries for the underlying variable.

It is remarkable that this generalisation preserves one of the most attractive features of the Black-Scholes model, namely the existence of closed-form expressions for the prices of European call and put options. Ingersoll (1989a, b) and Bühler and Käsler (1989) compute these expressions by applying a change of variable to the corresponding fundamental valuation equation. Our solution to the pricing problem, by contrast, relies on a probabilistic change-of-numeraire technique which goes back to El Karoui and Rochet (1989).

The first result of the chapter is a new decomposition of call and put prices

[6]See also Briys, Crouhy and Schöbel (1991).

that holds whenever the underlying variable has strict upper and lower bounds. In fact, these option prices can be decomposed in terms of two particular numeraire portfolios (whose definition reflects the presence of the bounds) and the martingale measures associated with these numeraires. The second contribution is a new proof of the pricing formulae for standard options in models with a quadratic diffusion term.

Derivative assets analysis usually takes a model of the underlying price processes as given and attempts to value derivatives relative to that model. Chapter 4 addresses the converse question: given some set of derivatives prices, what can we say about the price processes of the underlying securities? More precisely, we suppose that we have information about the price of a derivative asset in the form of a pricing formula, and investigate the restrictions such a formula imposes on the underlying price dynamics.

Assuming that asset prices are continuous semimartingales, we consider pricing formulae that satisfy a variant of the fundamental valuation equation. We show that such a formula implies a complete characterisation of the behaviour of the underlying asset price under a martingale measure. This characterisation takes the form of a stochastic differential equation. The law of the underlying price process under a martingale measure is completely determined by the pricing formula, and is the same for all martingale measures. Consequently, all claims contingent on the price path of the underlying security can be perfectly hedged and hence priced by arbitrage. In particular, the pricing formula implies a unique set of Arrow-Debreu prices for events that can be defined in terms of the underlying asset price.

While similar in spirit to Breeden and Litzenberger's (1978) calculation of state prices implicit in option prices, our approach uses rather different mathematical tools, based mainly on semimartingale calculus. The main result follows directly from a characterisation theorem for continuous local martingales which extends work by McGill, Rajeev and Rao (1988) on Brownian motion.

As an application of our result, we return to the analysis of pricing formulae for debt options. We show that the pricing formulae of Schöbel (1986) and Briys, Crouhy and Schöbel (1991) imply a positive probability for a certain forward rate process to be absorbed at its lower bound 0 during the life of the option. Thus, the event that this interest rate vanishes at the expiry date of the option indeed has positive probability, hence a positive Arrow-Debreu price. This explains the findings of Chapter 2.

Part I

Optimal Experimentation

Chapter 1

Optimal Experimentation in a Changing Environment

In this chapter, we study a situation in which an economic agent can learn by experimenting. Experimentation typically entails an opportunity cost, which is weighed against the long-term benefit resulting from the information acquired. Specifically, the agent is a monopolist facing an unknown demand curve, and an experiment results in a noisy observation of a point on it, leading the agent to revise his beliefs about the underlying demand.

The problem of optimal experimentation by a monopolist has attracted considerable interest in the microeconomic literature on learning. Starting with Rothschild (1974), this literature focusses mainly on the limiting behaviour of the monopolist in an infinite-horizon setting with an unknown but fixed demand curve.[1] Rothschild's analysis has been extended in a number of different directions by McLennan (1984), Easley and Kiefer (1988), Kiefer (1989), and Aghion, Bolton, Harris and Jullien (1991). A common result of these papers is that the monopolist's beliefs converge, so experimentation ceases in the limit. This means that in the long run, the monopolist will learn everything that is *worth* knowing. The question then arises as to whether he will learn the truth, that is, whether the beliefs converge to a one-point distribution at the true parameter value. Typically, the answer is that with positive probability they do not.

In contrast to this literature, we follow Kiefer (1991) and allow the demand

[1] See Mirman, Samuelson and Urbano (1993) and the references therein for an analysis of optimal monopoly experimentation in a two-period framework.

curve to change randomly. In our model, as in Kiefer's, there are two states, each characterised by a linear expected demand curve. Transitions between these states are governed by a Markov process. The monopolist knows the slope and intercept of each expected demand curve and the transition probabilities of the Markov process. However, he does not know which demand curve he faces. At each moment of time, he chooses a quantity from a given interval of feasible quantities, and observes a price which is the expected price for the current state plus noise. Given this noisy signal of the underlying demand curve, the monopolist then updates his belief about the current state in a Bayesian fashion. The monopolist's objective is to maximise the expected discounted value of profits over an infinite horizon. He experiments whenever he chooses a quantity different from the so-called myopic quantity which, given his current belief, would maximise expected current profit.

We depart from Kiefer's model (and the overwhelming majority of similar models in the economics literature) by formulating the problem in continuous time.[2] Solving the model, we are ultimately lead to the problem of finding a solution to an ordinary differential equation. While we canot offer closed-form solutions, numerically solving the differential equation proves to be easier and far less time-consuming than calculating a numerical solution to the fixed point problem which arises in Kiefer's setting.

The chapter proceeds as follows. The model is presented in Section 1.1. The two subsequent sections are devoted to a steady development in which some important results from the literature on non-linear filtering and stochastic control are spelt out. Section 1.2 deals with the stochastic process of beliefs, leading on to a reformulation of the agent's decision problem as a diffusion control problem with the current belief as the state variable. Section 1.3 introduces the corresponding Bellman equation. We recall that the value function is a generalised solution of this equation and formulate a verification theorem. Using the Bellman equation, we examine two necessary conditions for experimentation: experimentation must be informative, and information must be useful. These are essentially the same as the

[2]Such a formulation has previously been adopted by Bolton and Harris (1993) and Felli and Harris (1994) who study multi-agent learning problems. The distribution of rewards, however, is *fixed* in their models.

conditions found in Mirman *et al.* (1993), despite the differences in the approach taken. Finally, we use the Bellman equation to solve for the optimal quantities as a function of the current belief and the second derivative of the value function.

In Section 1.4, we discuss scenarios in which there is one-sided experimentation, meaning that the monopolist either produces always more than the myopic quantity, or always less. One-sided experimentation arises if the expected demand curves do not intersect at a quantity inside the range of the myopic policy function, which is a closed interval spanned by the quantities that maximise one-period profit for each of the two expected demand curves. Again our findings are comparable with those in Mirman *et al.*, namely, the monopolist deviates from the myopic action by moving in the direction of wider spreads between the demand curves, thereby making the observations more informative. Experimentation remains moderate in the sense that the optimal quantity always lies inside the range of the myopic policy function. Moreover, experimentation is qualitatively the same for all parameter values. We obtain a full characterisation of the value function as a solution to a two-point boundary value problem. Even without a closed-form solution, we are still able to derive analytical results for the comparative statics of the monopolist's behaviour with respect to some of the model parameters.

In Section 1.5, we then look at situations that give rise to two-sided experimentation: for some beliefs, the monopolist produces more than the myopic quantity, for others, less. This happens whenever the expected demand curves intersect in the interior of the range of the myopic policy function. Then there is exactly one belief at which a myopic agent would choose the so-called confounding quantity, that is, the quantity which leads to the same expected price in either state and thus to a completely uninformative price signal. For sufficiently high discount rate, noise intensity and switching probabilities, even a fully optimizing monopolist chooses the confounding quantity at this particular belief. In this case, the optimal policy function is continuous, and experimentation is again moderate in the sense defined above. For brevity, we therefore call this a scenario of *moderate experimentation.* For sufficiently low discount rate, noise intensity and switching probabilities, on the other hand, we find *extreme experimentation* in the sense that the optimal quantity as a function of the belief exhibits a jump from one boundary of the interval of feasible quantities to the other. Thus, the monopolist's behaviour depends *qualitatively* on the parameters describing his environment.

Due to mathematical complications, our analytical results for the two-sided case are somewhat weaker than those obtained in the one-sided case. We therefore have to rely to a larger extent on numerical results. Some of these are presented in Section 1.6. The advantage of working in continuous time becomes evident, allowing far simpler numerical computations than the discrete-time approach predominant in the literature. These numerical results, whilst not being a substitute for analytical arguments, can be used as evidence for conjectures about the way in which the value function and the optimal policy vary with the discount rate, the signal-to-noise ratio, and the likelihood of changes of the demand curve. In particular, they indicate a *discontinuous* switch from extreme to moderate experimentation as the likelihood of a demand curve change increases through some critical level.

A summary and concluding remarks are given in Section 1.7. Some technical results relating to the case of one-sided experimentation (Section 1.4) can be found in an appendix.

1.1 The Model

The time parameter is continuous. There are two states indexed by k, $k \in \{0,1\}$. The state changes according to the transition probabilities

$$\begin{aligned}
\Pr(k_{t+\Delta t} = 0 \mid k_t = 0) &= 1 - \lambda_0 \Delta t + o(\Delta t), \\
\Pr(k_{t+\Delta t} = 1 \mid k_t = 0) &= \lambda_0 \Delta t + o(\Delta t), \\
\Pr(k_{t+\Delta t} = 0 \mid k_t = 1) &= \lambda_1 \Delta t + o(\Delta t), \\
\Pr(k_{t+\Delta t} = 1 \mid k_t = 1) &= 1 - \lambda_1 \Delta t + o(\Delta t)
\end{aligned}$$

with $\lambda_i \geq 0$ for $i = 0, 1$. In particular,

$$\Pr\left(k_s = i \;\forall s \in [t, t+\Delta t] \mid k_t = i\right) = \exp(-\lambda_i \Delta t);$$

see Karlin and Taylor (1981, p.146).

In state k, the expected demand curve (price as a function of quantity) is given by

$$p = \alpha_k - \beta_k q$$

where the α_k, β_k are fixed, known and strictly positive. The realised price is the expected price plus some noise term.

Most of the time, we shall assume that the slopes of the two demand curves are different, that is, $\beta_0 \neq \beta_1$. In this case, let $(\hat{q}, \hat{p})$ be the point in the (q,p)-plane where the two expected demand curves intersect. This point is easily seen to have the co-ordinates

$$\hat{q} = \frac{\Delta\alpha}{\Delta\beta}, \quad \hat{p} = \frac{\alpha_0\beta_1 - \alpha_1\beta_0}{\Delta\beta}$$

where $\Delta\alpha = \alpha_1 - \alpha_0$ and $\Delta\beta = \beta_1 - \beta_0$. The location of this intersection will turn out to be crucial. Of course, only if $\hat{q} \geq 0$ and $\hat{p} \geq 0$ can they be interpreted as a quantity and a price, respectively.

At each moment of time, the monopolist chooses a quantity q_t from an interval $Q = [q_{\min}, q_{\max}]$ of feasible quantities, and observes a price which is a noisy signal of whether $k_t = 0$ or 1. The monopolist's subjective probability at time t that $k_t = 1$ is denoted by π_t, and he updates this belief in a Bayesian fashion.

We assume that there is no cost to production, hence revenue equals profit. Working in continuous time, we model the revenue flow as

$$dR_t = q_t \,[(\alpha_{k_t} - \beta_{k_t} q_t)\, dt \;+\; \sigma dZ_t]$$

where Z is a standard Wiener process independent of the process k, and $\sigma > 0$, fixed and known. Alternatively, we can write this as $dR_t = q_t \, dP_t$ with the cumulative price process P given by

$$dP_t = (\alpha_{k_t} - \beta_{k_t} q_t)\, dt \;+\; \sigma dZ_t.$$

This is the process which the agent observes.

Consequently, the belief π_t is the conditional probability that $k_t = 1$ given the history of the process P. In the same way, admissible strategies $\mathbf{q} = \{q_t\}$ for the monopolist are such that the action taken at time t depends only on the price history up to that time. To make these ideas more precise, assume that the Brownian motion Z and the Markov process k are given on some complete probability space and are both adapted to the filtration $\{\mathcal{F}_t\}$. Let $\mathcal{Q}_0$ denote the set of all processes $\mathbf{q} = \{q_t\}$ which take values in Q, the interval of feasible quantities, and are adapted to the aforementioned filtration. Each $\mathbf{q} \in \mathcal{Q}_0$ gives rise to a unique cumulative price process $P^{\mathbf{q}}$. The information contained in prices is summarised by $\{\mathcal{F}_t^{\mathbf{q}}\}$, the filtration generated by $P^{\mathbf{q}}$. A process $\mathbf{q} \in \mathcal{Q}_0$ is an admissible strategy for the monopolist if q_t is adapted to the filtration $\{\mathcal{F}_t^{\mathbf{q}}\}$. The set of admissible strategies will be denoted by $\mathcal{Q}$. If the monopolist follows the strategy $\mathbf{q} \in \mathcal{Q}$, the posterior probability at time t that $k_t = 1$ is $\pi_t = \Pr(k_t = 1 \mid \mathcal{F}_t^{\mathbf{q}}) = \mathrm{E}[k_t \mid \mathcal{F}_t^{\mathbf{q}}]$.

Given the initial belief $\pi_0 = \pi$, the agent's objective is to choose a strategy $\mathbf{q} \in \mathcal{Q}$ so as to maximise

$$\begin{aligned} u^{\mathbf{q}}(\pi) &= \mathrm{E}_\pi\left[\int_0^\infty r\, e^{-r\,t} dR_t\right] \\ &= \mathrm{E}_\pi\left[\int_0^\infty r\, e^{-r\,t}\; q_t\left[(\alpha_{k_t} - \beta_{k_t} q_t)\, dt \;+\; \sigma dZ_t\right]\right] \end{aligned}$$

where $r > 0$ is the discount rate, fixed and known. Up to the multiplication by r, which expresses the payoff in per period terms, $u^{\mathbf{q}}(\pi)$ is the expected present value of the revenue flow from strategy $\mathbf{q}$. Note that we can also write

$$u^{\mathbf{q}}(\pi) = \mathrm{E}_\pi\left[\int_0^\infty r\, e^{-r\,t}\; q_t\left[\alpha_{k_t} - \beta_{k_t} q_t\right] dt\right]$$

since the stochastic integral with respect to the Wiener process Z has zero expectation.

1.1.1 The Value Function

As usual, the value function for the monopolist's decision problem is defined as

$$u(\pi) = \sup_{\mathbf{q}\in\mathcal{Q}}\; u^{\mathbf{q}}(\pi) \tag{1.1}$$

for $\pi \in [0, 1]$. A strategy $\mathbf{q} \in \mathcal{Q}$ is optimal for initial belief π if it attains the supremum in (1.1). Given $\epsilon > 0$, a strategy $\mathbf{q} \in \mathcal{Q}$ is called ϵ-optimal for initial belief π if $u^{\mathbf{q}}(\pi) \geq u(\pi) - \epsilon$.

The value function is clearly bounded. Moreover:

Proposition 1.1.1 *The value function u is continuous and convex, and possesses one-sided derivatives D_-u and D_+u which are bounded.*

PROOF: For fixed $\mathbf{q} \in \mathcal{Q}$, $u^{\mathbf{q}}$ is affine in π. Indeed,

$$\begin{aligned} u^{\mathbf{q}}(\pi) &= \pi \, \mathrm{E}_{k_0=1}\left[\int_0^\infty r\, e^{-rt}\, q_t\, [\alpha_{k_t} - \beta_{k_t} q_t]\, dt\right] \\ &\quad + (1-\pi)\, \mathrm{E}_{k_0=0}\left[\int_0^\infty r\, e^{-rt}\, q_t\, [\alpha_{k_t} - \beta_{k_t} q_t]\, dt\right]. \end{aligned}$$

For $\pi = \eta\, \pi_1 + (1-\eta)\, \pi_2$ with $0 \leq \eta \leq 1$, we therefore have

$$\begin{aligned} u^{\mathbf{q}}(\pi) &= \eta\, u^{\mathbf{q}}(\pi_1) + (1-\eta)\, u^{\mathbf{q}}(\pi_2) \\ &\leq \eta\, u(\pi_1) + (1-\eta)\, u(\pi_2) \end{aligned}$$

by the definition of the value function. Taking the supremum on the left hand side proves convexity. A convex function is continuous on the interior of its domain, so we only have to show continuity at $\pi = 0$ and $\pi = 1$. Suppose for example that the value function is not continuous at $\pi = 0$. Due to convexity, this can only mean $u(0) > u(0+)$. By definition of the value function, there exists a policy $\mathbf{q} \in \mathcal{Q}$ such that $u^{\mathbf{q}}(0) > u(0+)$. But then $u^{\mathbf{q}}(\pi) > u(\pi)$ for small $\pi > 0$, which is a contradiction. The right boundary $\pi = 1$ is dealt with in the same way. Now, convexity implies the existence of a left-hand derivative D_-u on $]0,1]$ and a right-hand derivative D_+u on $[0,1[$, both being non-decreasing functions, the former left-continuous, the latter right-continuous, with $D_-u \leq D_+u$ on their common domain. We want to show that they are bounded. From the above representation of the payoff function $u^{\mathbf{q}}$ we see that there is a constant $K > 0$ such that $|(u^{\mathbf{q}})'(\pi)| \leq K$ for all $\mathbf{q} \in \mathcal{Q}$ and all π. Now, suppose that $(D_-u)(\pi_1) < -K$ for some belief $\pi_1 > 0$. Then there is a $\pi_2 < \pi_1$ such that $u(\pi_1) - u(\pi_2) < -K\,(\pi_1 - \pi_2)$, i.e., $u(\pi_2) > u(\pi_1) + K\,(\pi_1 - \pi_2)$. By definition of the value function, we can find a strategy $\mathbf{q} \in \mathcal{Q}$ with $u(\pi_2) \geq u^{\mathbf{q}}(\pi_2) > u(\pi_1) + K\,(\pi_1 - \pi_2)$. But then the linearity of $u^{\mathbf{q}}$ implies $u^{\mathbf{q}}(\pi_1) \geq u^{\mathbf{q}}(\pi_2) - K\,(\pi_1 - \pi_2) > u(\pi_1)$, which is a contradiction. Using a similar argument for the right-hand derivative, we obtain $-K \leq D_-u \leq D_+u \leq K$ on $]0,1[$. Due to left and right continuity, respectively, this also proves that both $(D_-u)(1)$ and $(D_+u)(0)$ are bounded in absolute value by K. ■

The convexity of u expresses the fact that information is valuable to the agent.[3] If the prior belief is π_1 with probability η and π_2 with probability $1-\eta$, then the agent can only gain from being told which before choosing a strategy.

On a purely mathematical level, convexity implies further regularity properties of the value function which will be of use later on.[4]

1.1.2 Strategies Depending on Beliefs

Consider a strategy $\mathbf{q} \in \mathcal{Q}$ and the associated processes of cumulative prices P_t and beliefs π_t. By the law of iterated expectations, we have

$$u^{\mathbf{q}}(\pi) = \mathrm{E}_\pi \left[\int_0^\infty r\, e^{-r\,t}\, \mathrm{E}_\pi[q_t\,(\alpha_{k_t} - \beta_{k_t} q_t) \mid \mathcal{F}_t^{\mathbf{q}}]\, dt \right].$$

Note that $\mathrm{E}_\pi[q_t\,(\alpha_{k_t} - \beta_{k_t} q_t) \mid \mathcal{F}_t^{\mathbf{q}}]$ is the expected revenue, given the observed price history, for quantity q_t. By definition of π_t and the $\mathcal{F}_t^{\mathbf{q}}$-measurability of q_t, this expected revenue equals $q_t\,[(1-\pi_t)\alpha_0 + \pi_t\alpha_1 - ((1-\pi_t)\beta_0 + \pi_t\beta_1)q_t]$. To simplify the notation, we introduce the functions

$$\begin{aligned} \alpha(\pi) &= (1-\pi)\alpha_0 + \pi\alpha_1, \\ \beta(\pi) &= (1-\pi)\beta_0 + \pi\beta_1 \end{aligned}$$

which describe the expected intercept and slope parameter of the demand curve given the belief π. The expected revenue from setting quantity q is then

$$R(\pi, q) = q\,[\alpha(\pi) - \beta(\pi)q].$$

Thus,

$$u^{\mathbf{q}}(\pi) = \mathrm{E}_\pi \left[\int_0^\infty r\, e^{-r\,t}\, R(\pi_t, q_t)\, dt \right]. \tag{1.2}$$

This expression for $u^{\mathbf{q}}(\pi)$ does not involve the stochastic variable k_t any more; instead, the payoff relevant quantities are described as functions of π_t alone.

This suggests looking at strategies based exclusively on the information contained in beliefs. To make this precise, let $\{\mathcal{G}_t^{\mathbf{q}}\}$ be the filtration generated by the process of beliefs corresponding to $\mathbf{q} \in \mathcal{Q}$. By construction, $\mathcal{G}_t^{\mathbf{q}} \subseteq \mathcal{F}_t^{\mathbf{q}}$. The set of

[3] That is, unless the value function is linear.

[4] See the proof of Corollary 1.3.1.

all admissible strategies $\mathbf{q}$ which are adapted to the filtration $\{\mathcal{G}_t^{\mathbf{q}}\}$ will be denoted by $\tilde{\mathcal{Q}}$. Strategies in this class depend merely on the history of beliefs.

Finally, one can consider strategies that depend only on the *current* belief. A strategy $\mathbf{q} \in \tilde{\mathcal{Q}}$ is a (stationary) Markov strategy if there exists a policy function $q : [0,1] \to Q$ such that $q_t = q(\pi_t)$. The strategies calculated in subsequent sections will be of this type.

1.2 The Evolution of Beliefs

It follows from Liptser and Shiryayev (1977, Chapters 8 and 9) that given a choice of quantities $\mathbf{q} \in \mathcal{Q}$, the beliefs evolve according to the stochastic differential equation

$$d\pi_t = \lambda(\pi_t)\,dt + \sigma^{-1}\pi_t(1-\pi_t)(\Delta\alpha - \Delta\beta\, q_t)\,dZ_t^{\mathbf{q}} \tag{1.3}$$

where

$$\lambda(\pi) = (1-\pi)\lambda_0 - \pi\lambda_1$$

and

$$dZ_t^{\mathbf{q}} = \sigma^{-1}\left(\underbrace{(\alpha_{k_t} - \beta_{k_t}q_t)\,dt \;+\; \sigma dZ_t}_{\text{realised price}} \;-\; \underbrace{\mathrm{E}_\pi[\alpha_{k_t} - \beta_{k_t}q_t \mid \mathcal{F}_t^{\mathbf{q}}]\,dt}_{\text{expected price}}\right).$$

Note that $\mathrm{E}_\pi[\alpha_{k_t} - \beta_{k_t}q_t \mid \mathcal{F}_t^{\mathbf{q}}]$ is the expected price, given the observed price history, for quantity q_t. By definition of π_t and the $\mathcal{F}_t^{\mathbf{q}}$-measurability of q_t, this expected price equals $\alpha(\pi_t) - \beta(\pi_t)\,q_t$, hence

$$dZ_t^{\mathbf{q}} = \sigma^{-1}\left(\alpha_{k_t} - \beta_{k_t}q_t \;-\; [\alpha(\pi_t) - \beta(\pi_t)q_t]\right)\,dt \;+\; dZ_t. \tag{1.4}$$

Liptser and Shiryayev show that the process $Z^{\mathbf{q}}$ is a Wiener process with respect to the filtration $\{\mathcal{F}_t^{\mathbf{q}}\}$. This result, together with (1.3), is the key fact on which we can base the mathematical analysis of our model.

Equation (1.3) emphasises the two separate forces which drive the updating. The drift term $\lambda(\pi_t)\,dt$ takes account of the possibility that the state may change over the next infinitesimal period of time. If $\lambda_0 + \lambda_1 \neq 0$, the linear function λ is downward sloping and vanishes at

$$\tilde{\pi} = \frac{\lambda_0}{\lambda_0 + \lambda_1}.$$

We can write

$$\lambda(\pi) = -(\lambda_0 + \lambda_1)(\pi - \tilde{\pi})$$

which shows that, via the drift term in (1.3), state switches introduce mean reversion into the evolution of beliefs.

The diffusion term $\sigma^{-1}\pi_t(1-\pi_t)(\Delta\alpha - \Delta\beta\, q_t)\, dZ_t^{\mathbf{q}}$, on the other hand, captures the influence of the observed price signal on the evolution of beliefs. $Z^{\mathbf{q}}$ being a Wiener process, this part of the updating is completely unpredictable. Intuitively, this expresses the fact that the current belief already incorporates everything that there is to know, so any change must come as a surprise. The lower σ, and the greater the spread $\Delta\alpha - \Delta\beta\, q_t$ between the two demand curves, the more informative is the price signal, and the more pronounced is the change of beliefs after the signal is observed. This holds of course only if the agent is not subjectively certain of the current state. For $\pi = 0$ or 1, the agent rules out any possibility of learning from the price signal, so the diffusion term vanishes no matter which action is taken.

To get some insights into the boundary behaviour of the stochastic process of beliefs, let us consider Markov strategies. The results obtained will apply to the policies which we shall encounter in subsequent sections.

Proposition 1.2.1 *Let* $\mathbf{q} \in \tilde{\mathcal{Q}}$ *be a Markov strategy with the policy function* $q : [0,1] \to Q$. *Suppose that* q *is continuous at* $\pi = 0$ *and* $\pi = 1$ *with* $q(0) \neq \hat{q} \neq q(1)$. *Then* (1.3) *defines a diffusion process with the following boundary behaviour:*

(a) The boundaries $\pi = 0$ *and* $\pi = 1$ *are both unattainable, that is, they cannot be reached in finite time from the interior of the interval of possible beliefs.*

(b) For $\lambda_0 = 0$ $(\lambda_1 = 0)$, *the process of beliefs, if started at the boundary* $\pi = 0$ $(\pi = 1)$, *remains there forever.*

(c) For $\lambda_0 > 0$ $(\lambda_1 > 0)$, *the process of beliefs, if started at the boundary* $\pi = 0$ $(\pi = 1)$, *moves immediately into the interior of the interval of possible beliefs.*

In other words, $\pi = i$ $(i = 0, 1)$ *is a natural boundary if* $\lambda_i = 0$, *and an entrance boundary if* $\lambda_i > 0$.

PROOF: To prove the proposition, one has to verify certain integral criteria which can be found for example in Karlin and Taylor (1981, Chapter 15). This is a straightforward but rather tedious exercise, and therefore omitted. ■

For a Markov strategy $\mathbf{q} \in \tilde{\mathcal{Q}}$ with policy function $q : [0,1] \to Q$, it can be shown[5] that the payoff function $u^{\mathbf{q}}$ solves the backward equation

$$\frac{1}{2\sigma^2}\pi^2(1-\pi)^2[\Delta\alpha - \Delta\beta\, q(\pi)]^2\, u''(\pi) + \lambda(\pi)\, u'(\pi) - r\, u(\pi) + r\, R(\pi, q(\pi)) = 0 \quad (1.5)$$

corresponding to (1.3) at every point $\pi \in]0,1[$ such that (i) the "demand spread" quadratic $\pi \mapsto [\Delta\alpha - \Delta\beta\, q(\pi)]^2$ is continuous and non-zero at π; (ii) the expected revenue function $\pi \mapsto R(\pi, q(\pi))$ is continuous at π. Moreover, the payoff function $u^{\mathbf{q}}$ has a continuous first derivative on every interval where $\Delta\alpha - \Delta\beta\, q(\pi) \neq 0$. Finally, discontinuities of the above quadratic functions are "absorbed" by the second derivative of $u^{\mathbf{q}}$.

1.3 The Bellman Equation

The representation (1.2) for the payoff $u^{\mathbf{q}}(\pi)$, the stochastic differential equation (1.3) for the evolution of beliefs and the fact that $Z^{\mathbf{q}}$ is a Wiener process allow us to consider the monopolist's decision problem as a problem of optimal control of a diffusion process, the diffusion in question being the process of beliefs.[6]

The present section deals with the associated Bellman equation. We shall first quote a result from Krylov (1980) stating that the value function is a generalised solution of the Bellman equation. After interpreting the economics behind this equation, we shall then present a so-called verification theorem, that is, sufficient conditions for a given function of beliefs to be the value function. Next, necessary conditions for experimentation to occur at the optimum will be given along the lines of Mirman, Samuelson and Urbano (1993). Finally, we shall derive the quantities that solve the maximisation problem in the Bellman equation.

[5]See for example Morton (1971).

[6]Standard references on controlled diffusions are Fleming and Rishel (1975) and Krylov (1980).

1.3.1 The Value Function as a Solution of the Bellman Equation

To state the following result, we need to introduce the concept of a generalised derivative. Functions $u_1, u_2 :]0,1[\to I\!R$ are called generalised first and second derivatives of a given function u if

$$\int_0^1 \phi(\pi)\, u_1(\pi)\, d\pi = -\int_0^1 \phi'(\pi)\, u(\pi)\, d\pi$$

and

$$\int_0^1 \phi(\pi)\, u_2(\pi)\, d\pi = \int_0^1 \phi''(\pi)\, u(\pi)\, d\pi$$

for all functions ϕ that are infinitely differentiable and of compact support in $]0,1[$. Generalised derivatives are defined only up to changes on null sets. Moreover, if a function is once or twice differentiable in the classic sense, then the true derivatives are also generalised derivatives.

Proposition 1.3.1 *In the interior* $]0,1[$ *of the set of beliefs, the value function* u *has two generalised derivatives,* u_1 *and* u_2*, which are locally bounded. With these generalised derivatives, the value function* u *satisfies the Bellman equation*

$$\sup_{q \in Q} \left\{ \frac{1}{2\sigma^2}\pi^2(1-\pi)^2(\Delta\alpha - \Delta\beta\, q)^2\, u_2(\pi) + \lambda(\pi)\, u_1(\pi) - r\, u(\pi) + r\, R(\pi, q) \right\} = 0 \tag{1.6}$$

almost everywhere on $]0,1[$.

PROOF: This follows directly from Theorem 6, p.289, of Krylov (1980). ■

Using the convexity of u, we can strengthen this result as follows.

Corollary 1.3.1 *The value function* u *is continuously differentiable on* $[0,1]$ *and has a generalised second derivative* $u_2 \geq 0$ *on* $]0,1[$ *such that*

$$u'(\pi_2) - u'(\pi_1) = \int_{\pi_1}^{\pi_2} u_2(\pi)\, d\pi \tag{1.7}$$

for all π_1 *and* π_2 *in the open unit interval. Moreover,*

$$\sup_{q \in Q} \left\{ \frac{1}{2\sigma^2}\pi^2(1-\pi)^2(\Delta\alpha - \Delta\beta\, q)^2\, u_2(\pi) + \lambda(\pi)\, u'(\pi) - r\, u(\pi) + r\, R(\pi, q) \right\} = 0 \tag{1.8}$$

almost everywhere on $]0,1[$.

PROOF: u is convex by Proposition 1.1.1. As its left-hand derivative $D_- u$ is left-continuous and non-decreasing, one can define a measure μ on $]0,1[$ via $\mu[\pi_1,\pi_2[= (D_- u)(\pi_2) - (D_- u)(\pi_1)$. This measure represents the second derivative of u in the sense of a distribution:

$$\int_0^1 \phi''(\pi)\, u(\pi)\, d\pi = \int_0^1 \phi(\pi)\, d\mu(\pi)$$

for every function ϕ that is infinitely differentiable and of compact support in $]0,1[$. Moreover, this property characterises μ uniquely.[7] Comparing it with the definition of the generalised second derivative u_2, we conclude that $d\mu = u_2\, d\pi$. In particular,

$$(D_- u)(\pi_2) - (D_- u)(\pi_1) = \int_{\pi_1}^{\pi_2} u_2(\pi)\, d\pi\,,$$

so $D_- u$ is continuous, and u is continuously differentiable on the open unit interval. By Proposition 1.1.1, $u'(\pi)$ has a continuous extension on the whole of $[0,1]$. Finally, we can replace u_1 by u' in the Bellman equation. ■

1.3.2 Interpretation

Some economic insights can be gained from rewriting the Bellman equation as[8]

$$\begin{aligned} u(\pi) \;=\;& \underbrace{\frac{\lambda(\pi)}{r}\, u'(\pi)}_{\text{value of state changes}} \\ &+ \sup_{q\in Q}\left\{ \underbrace{\frac{1}{2r\sigma^2}\pi^2(1-\pi)^2(\Delta\alpha - \Delta\beta\, q)^2\, u''(\pi)}_{\text{value of information}} + \underbrace{R(\pi,q)}_{\text{myopic payoff}} \right\}, \qquad (1.9) \end{aligned}$$

showing the trade-off between information gathering and myopic profit maximisation.

Indeed, the discussion after equation (1.3) above shows that $\sigma^{-2}\pi^2(1-\pi)^2(\Delta\alpha - \Delta\beta\, q)^2$ provides a measure for the informativeness of the price signal obtained from setting the quantity q. This informativeness is valued with the shadow price $u''(\pi)/2r$.

[7] Cf. Krylov (1980, p.49).

[8] In this subsection, we shall use the standard notation for derivatives despite the fact that the theory so far only guarantees the existence of *generalised* second derivative.

The term "myopic payoff" needs some explanation. We refer here to a strong form of myopia which assumes that the current belief will persist forever. If this were correct, setting the quantity q forever would indeed yield

$$\int_0^\infty r\, e^{-r\,t}\, R(\pi, q)\; dt = R(\pi, q).$$

Alternatively, we can think of this as the expected payoff for $r = \infty$. As r tends to infinity, the distribution on $\mathbb{R}_+$ with density $r\, e^{-r\,t}$ degenerates to a point mass at $t = 0$, and the agent becomes myopic in so far as he does not care for the future any more: for $r = \infty$, $u^{\mathbf{q}}(\pi) = R(\pi, q_0)$.

According to (1.3), $\lambda(\pi)$ indicates the direction in which the belief is likely to move due to possible state changes. This piece of information has the shadow price $u'(\pi)/r$. The contribution to the value function which we called "value of state changes" is positive if the mean reversion force works in the direction of value increases.

Finally, note that the optimisation in (1.9) is formally the same with or without state switches. However, the value function will of course be affected by the possibility of change.

1.3.3 A Verification Theorem

The following proposition provides a sufficient condition for a solution of the Bellman equation to be the value function, and for a Markov strategy to be ϵ-optimal. In the literature on stochastic control, a result of this type is usually called a verification theorem.

Proposition 1.3.2 *Let u be a continuously differentiable function on $[0, 1]$ with a generalised second derivative $u_2 \geq 0$ on $]0, 1[$ such that (1.7) holds for all π_1 and π_2, and $\pi^2\,(1-\pi)^2\, u_2(\pi) \to 0$ as $\pi \to 0$ and $\pi \to 1$, respectively. If u solves the Bellman equation (1.8) on $[0, 1]$, then:*

(a) $u(\pi) \geq u^{\mathbf{q}}(\pi)$ for all $\mathbf{q} \in \mathcal{Q}$.

(b) *Let $\epsilon \geq 0$. If $\tilde{\mathbf{q}} \in \tilde{\mathcal{Q}}$ is a Markov strategy with policy function $q: [0,1] \to Q$ such that for any π*

$$\frac{1}{2r\sigma^2}\pi^2(1-\pi)^2[\Delta\alpha - \Delta\beta\, q(\pi)]^2\, u_2(\pi) + \frac{\lambda(\pi)}{r}\, u'(\pi) - u(\pi) + R(\pi, q(\pi)) \geq -\epsilon\,,$$

then $u^{\tilde{\mathbf{q}}}(\pi) \geq u(\pi) - \epsilon$ for all π. In particular, $\tilde{\mathbf{q}}$ is ϵ-optimal (in $\mathcal{Q}$).

(c) *If there is a Markov strategy as in (b) for any $\epsilon > 0$, then u is the value function: $u(\pi) = \sup_{\mathbf{q}\in\mathcal{Q}} u^{\mathbf{q}}(\pi)$.*

(d) *If $\mathbf{q}^* \in \tilde{\mathcal{Q}}$ is a Markov strategy with policy function $q^*: [0,1] \to Q$ such that for every π, the supremum in (1.8) is attained at $q^*(\pi)$, then $\mathbf{q}^*$ is optimal: $u(\pi) = \max_{\mathbf{q}\in\mathcal{Q}} u^{\mathbf{q}}(\pi) = u^{\mathbf{q}^*}(\pi)$.*

Note that under (c) or (d), it is enough to consider belief-dependent Markov strategies, as indicated at the beginning of this section.

PROOF: Let the initial belief be $\pi_0 = \pi$. For an arbitrary policy $\mathbf{q} \in \mathcal{Q}$ consider the stochastic process $M^{\mathbf{q}}$ given by

$$M_T^{\mathbf{q}} = \int_0^T r\, e^{-r\,t}\; R(\pi_t, q_t)\, dt \;+\; e^{-r\,T}\; u(\pi_T).$$

By a generalisation of Itô's lemma,[9]

$$\begin{aligned} M_T^{\mathbf{q}} \;=\; M_0^{\mathbf{q}} \;+ \int_0^T e^{-r\,t}\Big\{&\frac{1}{2\sigma^2}\pi_t^2(1-\pi_t)^2(\Delta\alpha - \Delta\beta\, q_t)^2\, u_2(\pi_t) \\ &+ \lambda(\pi_t)\, u'(\pi_t) - r\, u(\pi_t) + r\, R(\pi_t, q_t)\Big\}\, dt \\ + \;\sigma^{-1} \int_0^T & e^{-r\,t}\pi_t(1-\pi_t)(\Delta\alpha - \Delta\beta\, q_t)\, dZ_t^{\mathbf{q}}. \end{aligned}$$

(1.8) implies that the expression under the first integral is non-positive, so $M^{\mathbf{q}}$ is a supermartingale. In other words, $\mathrm{E}_\pi[M_T^{\mathbf{q}}] \leq M_0^{\mathbf{q}}$ or

$$u(\pi) \geq \mathrm{E}_\pi\left[\int_0^T r\, e^{-r\,t}\; R(\pi_t, q_t)\, dt\right] + e^{-r\,T}\; \mathrm{E}_\pi[u(\pi_T)].$$

Letting T go to infinity, we see that the first term on the right hand side becomes $u^{\mathbf{q}}(\pi)$, while the second term tends to zero. This proves part (a). Now let $\epsilon \geq 0$.

[9]Cf. Lemma IV.45.9, p.105, of Rogers and Williams (1987).

If $\mathbf{q}$ is a Markov strategy such that

$$\frac{1}{2\sigma^2}\pi^2(1-\pi)^2[\Delta\alpha - \Delta\beta\, q(\pi)]^2\, u_2(\pi) + \lambda(\pi)\, u'(\pi) - r\, u(\pi) + r\, R(\pi, q(\pi)) \geq -r\epsilon$$

for all π, then

$$\mathrm{E}_\pi[M_T^{\mathbf{q}}] \geq M_0^{\mathbf{q}} - \epsilon \int_0^T r\, e^{-rt}\, dt.$$

Letting $T \to \infty$ yields $u^{\mathbf{q}}(\pi) \geq u(\pi) - \epsilon$. Parts (b) - (d) follow immediately. ■

1.3.4 Necessary Conditions for Experimentation

The agent is said to *experiment* if he deviates from the action that maximises expected current reward. That is, given the belief π, the monopolist experiments if he sets a quantity different from the myopic quantity

$$q^m(\pi) = \arg\max_{q \in Q} R(\pi, q) = \frac{\alpha(\pi)}{2\beta(\pi)}.$$

Experimentation entails an opportunity cost in the form of a loss in current payoff. In fact, while the myopic optimum (or maximum expected revenue) is

$$m(\pi) = \max_{q \in Q} R(\pi, q) = R(\pi, q^m(\pi)) = \frac{\alpha(\pi)^2}{4\beta(\pi)},$$

the expected revenue from setting an arbitrary quantity q is

$$R(\pi, q) = m(\pi) - \beta(\pi)[q - q^m(\pi)]^2,$$

which decreases strictly as the distance between q and $q^m(\pi)$ increases.

In a two-period framework, Mirman, Samuelson and Urbano (1993) identify two necessary conditions for experimentation to occur at the optimum:

- Experimentation must be informative.
- Information must be valuable.

Guided by their analysis, we now study two special cases of our model where one of the above conditions is violated.

The case of uninformative experimentation. Assume that the two demand curves have the same slope parameter, $\beta_0 = \beta_1$. Thus, the monopolist faces an unknown and possibly changing intercept.[10] As the demand curves are parallel, a change in output does not affect the spread between the two possible price distributions. Thus, the choice of a quantity has no impact on the informativeness of the price signal. This renders experimentation uninformative, and the myopic quantity is optimal. Indeed, for $\Delta\beta = 0$, equation (1.9) reduces to

$$u(\pi) = \frac{\lambda(\pi)}{r} u'(\pi) + \frac{\Delta\alpha^2}{2r\sigma^2}\pi^2(1-\pi)^2 u''(\pi) + \sup_q R(\pi, q),$$

implying the optimal quantity $q^*(\pi) = q^m(\pi)$ for all π.

The case of worthless information. Suppose that the two demand curves intersect exactly on the quantity axis, that is, $\hat{p} = 0$ or

$$\frac{\alpha_0}{\beta_0} = \frac{\alpha_1}{\beta_1}.$$

Then, one and the same quantity is optimal under either demand curve: $q^m(0) = q^m(1)$, which we denote by q_0. In this situation, information is clearly worthless, and we expect the constant policy $q^*(\pi) = q_0$ to be optimal. Indeed, it is straightforward to verify that the corresponding linear payoff function $u^{q_0}(\pi)$ solves the Bellman equation.

For the rest of the paper, we rule out these two cases by making the following

Assumption *The abovementioned necessary conditions for experimentation are satisfied, that is,*

- *Experimentation is informative:* $\Delta\beta \neq 0$
- *Information is valuable:* $\hat{p} \neq 0$

More precisely, we assume without loss of generality that the demand curve in state 1 *is steeper than the demand curve in state* 0*:* $\beta_1 > \beta_0$*, that is,* $\Delta\beta > 0$.

For later reference, we note that this assumption implies strict monotonicity of the myopic policy function q^m. Indeed,

$$(q^m)'(\pi) = \frac{\alpha_1\beta_0 - \alpha_0\beta_1}{2\beta(\pi)^2} = -\frac{\Delta\beta\,\hat{p}}{2\beta(\pi)^2},$$

[10] Of course, we assume $\alpha_0 \neq \alpha_1$ to avoid trivialities.

hence q^m is strictly increasing if $\hat{p} < 0$, and strictly decreasing if $\hat{p} > 0$.

Due to strict monotonicity, the image of the myopic policy function q^m is a closed interval, denoted by Q^m, with boundaries $q^m(0)$ and $q^m(1)$.[11] Moreover, there are exactly two cases. Either $\hat{q} \notin Q^m$, or there is a unique belief $\hat{\pi}$ such that $q^m(\hat{\pi}) = \hat{q}$. We shall see that these two scenarios lead to very different outcomes.

1.3.5 Optimal Quantities

The Bellman equation can be used to determine the optimal quantity as a function of π and $u_2(\pi)$, the generalised second derivative of the value function. We introduce the notation

$$V(\pi) = \frac{\Delta\beta^2}{2r\sigma^2}\pi^2(1-\pi)^2\,u_2(\pi)$$

and rewrite (1.8) as

$$u(\pi) - \frac{\lambda(\pi)}{r}\,u'(\pi) - m(\pi) = \sup_{q\in Q} \Phi[V(\pi), \pi, q] \tag{1.10}$$

with the function

$$\Phi[V, \pi, q] = V\,[q-\hat{q}]^2 - \beta(\pi)\,[q - q^m(\pi)]^2.$$

Now, fix a belief π and a value of $V(\pi)$, and write $\Phi(q)$ for $\Phi[V(\pi), \pi, q]$. The optimal quantities are determined as follows.

Strictly concave case: $V(\pi) < \beta(\pi)$. In this case, the first order condition determines the unique optimal quantity

$$q^* = \begin{cases} q_{\max} & \text{if } \Phi'(q_{\max}) \geq 0; \\ q_{\min} & \text{if } \Phi'(q_{\min}) \leq 0; \\ \frac{\beta(\pi)}{\beta(\pi)-V(\pi)}\,q^m(\pi) - \frac{V(\pi)}{\beta(\pi)-V(\pi)}\,\hat{q} & \text{else.} \end{cases} \tag{1.11}$$

[11]We assume of course that the these two myopic quantities are admissible, hence $Q^m \subseteq Q$.

Linear case: $V(\pi) = \beta(\pi)$. In this case, $\Phi'(q) \equiv 2\beta(\pi)(q^m(\pi) - \hat{q})$, and the optimal quantities are

$$\arg\max_{q \in Q} \Phi(q) = \begin{cases} q_{\max} & \text{if } \Phi'(q) > 0, \text{ i.e. } q^m(\pi) > \hat{q}; \\ \\ q_{\min} & \text{if } \Phi'(q) < 0, \text{ i.e. } q^m(\pi) < \hat{q}; \\ \\ [q_{\min}, q_{\max}] & \text{if } \Phi'(q) = 0, \text{ i.e. } q^m(\pi) = \hat{q}. \end{cases} \tag{1.12}$$

Note that $\Phi(q) \equiv 0$ in the third sub-case.[12]

Strictly convex case: $V(\pi) > \beta(\pi)$. Here, we necessarily have corner solutions. The question as to which corner produces the higher value can be decided by looking at $\Phi'(q_c)$ where

$$q_c = \frac{q_{\max} + q_{\min}}{2}$$

is the centre point of the interval of admissible quantities:

$$\arg\max_{q \in Q} \Phi(q) = \begin{cases} q_{\max} & \text{if } \Phi'(q_c) > 0; \\ \\ q_{\min} & \text{if } \Phi'(q_c) < 0; \\ \\ \{q_{\min}, q_{\max}\} & \text{if } \Phi'(q_c) = 0. \end{cases} \tag{1.13}$$

The sign of $\Phi'(q_c)$ indicates where the symmetry axis of the graph of $\Phi(q)$, a parabola, lies relative to the midpoint of the interval $Q = [q_{\min}, q_{\max}]$. A positive sign, for instance, means that the symmetry axis lies to the left of q_c, which implies that the value at the right corner is higher than the value at the left corner.

Inserting the optimal quantities into the Bellman equation, one can derive an ordinary differential equation for the value function. This is the starting point for the analysis presented in the following two sections.

1.4 One-Sided Experimentation

In this section, we study scenarios where the two demand curves do not intersect in Q^m, that is, $\hat{q} \notin Q^m$. These cases allow a simple characterisation of the value

[12]Of course, this can only occur if $\hat{q} \in Q^m$ in the first place.

function and the optimal policy.

In particular, it will turn out that the value function is twice differentiable and that the optimal quantity is always an inner solution. As indicated at the end of the previous section, we can find an ODE for the value function by substituting the optimal quantity into the Bellman equation. In the case of an inner solution, $V(\pi) < \beta(\pi)$ and

$$q^*(\pi) = \frac{\beta(\pi)}{\beta(\pi) - V(\pi)}\, q^m(\pi) - \frac{V(\pi)}{\beta(\pi) - V(\pi)}\, \hat{q}$$

is optimal. Using (1.10) and simplifying, we obtain

$$\begin{aligned} u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi) - m(\pi) &= V(\pi)\,[q^*(\pi) - \hat{q}]^2 - \beta(\pi)[q^*(\pi) - q^m(\pi)]^2 \\ &= \frac{\beta(\pi)V(\pi)}{\beta(\pi) - V(\pi)}[q^m(\pi) - \hat{q}]^2. \end{aligned}$$

Note that $\beta(\pi)[q^m(\pi) - \hat{q}]^2 = m(\pi) - \hat{m}$ where $\hat{m} = \hat{p}\hat{q} = R(\pi, \hat{q})$.[13] Thus,

$$u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi) - m(\pi) = \frac{V(\pi)}{\beta(\pi) - V(\pi)}[m(\pi) - \hat{m}]$$

which, for $u(\pi) - \lambda(\pi)u'(\pi)/r - \hat{m} > 0$, is equivalent to the ordinary differential equation

$$\frac{\Delta\beta^2}{2r\sigma^2}\pi^2(1-\pi)^2\, u''(\pi) - \beta(\pi)\frac{u(\pi) - \lambda(\pi)u'(\pi)/r - m(\pi)}{u(\pi) - \lambda(\pi)u'(\pi)/r - \hat{m}} = 0. \tag{1.14}$$

The value function for $\hat{q} \notin Q^m$ can be fully characterised as the unique solution to (1.14) with values bounded below by the myopic payoff function m and bounded above by the straight line joining the myopic payoffs $m(0)$ and $m(1)$. This line has the equation

$$\overline{m}(\pi) = (1-\pi)m(0) + \pi m(1).$$

By strict convexity of the function m, we have $m < \overline{m}$ on $]0,1[$.

Theorem 1.4.1 *The value function $u(\pi)$ is strictly convex, once continuously differentiable on $[0,1]$, and analytic on $]0,1[$. It is the unique solution of* (1.14) *such that*

$$m(\pi) < u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi) < \overline{m}(\pi)$$

[13]If $\hat{q} \geq 0$ and $\hat{p} \geq 0$, $\hat{m}$ can be interpreted as the expected revenue, given any belief π, from choosing the quantity $\hat{q}$. Moreover, if $\hat{q} \in Q^m$, then $\hat{m} = m(\hat{\pi})$.

on $]0,1[$. *In particular,*

$$\begin{aligned} u(0) - \frac{\lambda(0)}{r}\, u'(0) &= m(0), \\ u(1) - \frac{\lambda(1)}{r}\, u'(1) &= m(1). \end{aligned}$$

The optimal policy takes values in Q^m *and is given by*

$$q^*(\pi) = q^m(\pi) + \frac{u(\pi) - \lambda(\pi)u'(\pi)/r - m(\pi)}{m(\pi) - \hat{m}}\, [q^m(\pi) - \hat{q}]\,.$$

PROOF: It is shown in the appendix to this chapter that there exists a solution u of (1.14) with the desired properties. Let us verify that any such u satisfies the Bellman equation. From (1.14), we see immediately that the corresponding function V satisfies $V(\pi) < \beta(\pi)$, so we are in the case where the maximand $\Phi[V(\pi), \pi, q]$ as defined earlier is strictly concave in q. We fix a belief π in $]0,1[$ and write $\Phi(q)$ for $\Phi[V(\pi), \pi, q]$. For the sake of concreteness, assume that $\hat{q}$ lies to the left of Q^m. Then $\Phi'(q^m(\pi)) = 2V(\pi)\,[q^m(\pi) - \hat{q}] > 0$, so the optimal quantity satisfies $q^*(\pi) > q^m(\pi)$. Next, $\hat{q} < Q^m$ implies that the myopic policy function q^m is strictly decreasing, thus $q^m(0)$ is the maximum of Q^m. We shall prove that $\Phi'(q^m(0)) < 0$, hence $q^*(\pi) < q^m(0)$. Using (1.14), we see that $\Phi'(q^m(0)) < 0$ if and only if

$$[u(\pi) - \frac{\lambda(\pi)}{r}\, u'(\pi) - m(\pi)]\,[q^m(\pi) - \hat{q}] < [q^m(0) - q^m(\pi)]\,[m(\pi) - \hat{m}].$$

As $m(\pi) - \hat{m} = \beta(\pi)[q^m(\pi) - \hat{q}]^2$ and $m(\pi) = R(\pi, q^m(0)) + \beta(\pi)[q^m(0) - q^m(\pi)]^2$, this inequality is equivalent to

$$u(\pi) - \frac{\lambda(\pi)}{r}\, u'(\pi) < R(\pi, q^m(0)) + \beta(\pi)[q^m(0) - q^m(\pi)][q^m(0) - \hat{q}].$$

But $u(\pi) - \lambda(\pi)u'(\pi)/r < \overline{m}(\pi)$, so a sufficient condition for $q^m(\pi) < q^m(0)$ is

$$\overline{m}(\pi) < R(\pi, q^m(0)) + \beta(\pi)[q^m(0) - q^m(\pi)][q^m(0) - \hat{q}].$$

This condition holds for all $\pi > 0$. In fact, both sides of the inequality are linear in π.[14] They coincide at $\pi = 0$, so it is enough to show that the inequality holds at

[14]Note that $\beta(\pi)\, q^m(\pi) = \alpha(\pi)/2$.

$\pi = 1$, which is straightforward. This proves that the optimal quantity is an inner solution, given by the third line of (1.11). As we have seen before, this implies that the Bellman equation reduces to the differential equation (1.14), so the solution u satisfies the Bellman equation. Moreover, we can use (1.14) to replace V in (1.11), which yields the policy function q^* as stated in the theorem. The case where the demand curves intersect to the right of Q^m can be dealt with in exactly the same way.

Next, we want to show that u is the value function. It follows from an extension to the standard existence theorems that the stochastic differential equation

$$\begin{aligned} d\pi_t &= \Big\{\lambda(\pi_t) \\ &\quad + \frac{\Delta\beta}{\sigma^2}\,\pi_t\,(1-\pi_t)\,[\hat{q} - q^*(\pi_t)]\,\Big(\alpha_{k_t} - \beta_{k_t} q^*(\pi_t) \;-\; [\alpha(\pi_t) - \beta(\pi_t) q^*(\pi_t)]\Big)\Big\}\,dt \\ &\quad + \frac{\Delta\beta}{\sigma}\,\pi_t\,(1-\pi_t)\,[\hat{q} - q^*(\pi_t)]\,dZ_t \end{aligned} \tag{1.15}$$

which is obtained from combining (1.3) and (1.4) has a solution π^* for any given starting value $\pi_0 \in [0,1]$.[15] Now, define the strategy $\mathbf{q}^*$ by $q_t^* = q^*(\pi_t^*)$ and consider the associated price process $dP_t^* = (\alpha_{k_t} - \beta_{k_t}\, q_t^*)\,dt + \sigma\,dZ$. Section 1.2 implies that the corresponding process of beliefs $\pi_t^{**} = \mathrm{E}[k_t \mid \mathcal{F}_t^{\mathbf{q}^*}]$ also solves (1.15) with initial value π_0. By the uniqueness result in Theorem 9.2 of Liptser and Shiryayev (1977), the processes π^* and π^{**} coincide, so π^* is indeed the process of beliefs associated with the strategy $\mathbf{q}^*$. The latter is therefore a Markov policy in $\tilde{\mathcal{Q}}$ with policy function $q^* : [0,1] \to Q$ such that, for every π, the supremum in the Bellman equation is attained at $q^*(\pi)$. Thus, u is the value function by Proposition 1.3.2. In particular, the solution u must be unique. ■

Let

$$\overline{q}(\pi) = q^m(\pi) + \frac{\overline{m}(\pi) - m(\pi)}{m(\pi) - \hat{m}}\;[q^m(\pi) - \hat{q}]\,.$$

This defines a function with the same monotonicity properties as q^m, that is, strictly

[15]Cf. Liptser and Shiryayev (1977, p.330). This is in fact a *strong* solution. A *weak* solution would be enough for our purposes.

increasing for $\hat{p} < 0$, and strictly decreasing for $\hat{p} > 0$.[16] Note that $\overline{q}$ and q^* coincide at either end of the unit interval, so the range of $\overline{q}$ equals Q^m.

Corollary 1.4.1 *If $\hat{q} < Q^m$, then $q^m(\pi) < q^*(\pi) < \overline{q}(\pi)$ on $]0,1[$. In particular, the agent experiments by increasing quantity. On the other hand, if $\hat{q} > Q^m$, then $\overline{q}(\pi) < q^*(\pi) < q^m(\pi)$ on $]0,1[$, so the agent experiments by reducing quantity.*

The intuition behind quantity expansion or reduction is straightforward. The monopolist deviates from the myopic quantity by moving in the direction of wider spreads between the two possible demand curves, thus making price observations more informative. It is less obvious, though, that experimentation should always be moderate in the sense that the optimal quantity lies inside the interval of myopic quantities.

Next, we turn to the comparative statics of the monopolist's behaviour.

Proposition 1.4.1 *Experimentation decreases with a rise in the discount rate. More precisely, $|q^*(\pi) - q^m(\pi)|$ is strictly decreasing in r for all $\pi \in]0,1[$.*

PROOF: We show in the appendix that $u(\pi) - \lambda(\pi)u'(\pi)/r$ is strictly decreasing in r for all $\pi \in]0,1[$. The proposition follows therefore from the representation of $q^*(\pi)$ given in Theorem 1.4.1. ■

The intuition behind this result is again clear. As the discount rate increases, the future becomes less important to the agent. The value of information falls,[17] and with it the agent's willingness to sacrifice current revenue for potential future gains from experimentation.

We expect (and our simulations strongly suggest) a similar decrease in experimentation for an increase, *ceteris paribus*, in the noise parameter σ or the transition

[16]The derivative of $\overline{q}$ is

$$\overline{q}'(\pi) = -\frac{[q^m(0)-\hat{q}]\,[q^m(1)-\hat{q}]\,\Delta\beta\,\hat{p}}{2\,\beta(\pi)^2\,[q^m(\pi)-\hat{q}]^2} = \frac{[q^m(0)-\hat{q}]\,[q^m(1)-\hat{q}]}{[q^m(\pi)-\hat{q}]^2}\,(q^m)'(\pi);$$

as $\hat{q} \notin Q^m$, this is well defined and of the same sign as the derivative of q^m.

[17]In fact, (1.14) shows that a decrease in $u(\pi) - \lambda(\pi)u'(\pi)/r$ is tantamount to a decrease in $V(\pi)$ and hence in the value of information.

intensities λ_0 and λ_1.[18] A higher level of noise renders the price signal less informative, while a higher probability of state switches increases the risk of information becoming obsolescent. Both are bound to reduce the monopolist's incentive to experiment.

Conjecture 1.4.1 $|q^*(\pi) - q^m(\pi)|$ *is strictly decreasing in* σ *for all* $\pi \in \,]0,1[$. *So experimentation decreases with an increase in the noise parameter.*

Conjecture 1.4.2 *Holding* $\tilde{\pi}$ *fixed,* $|q^*(\pi) - q^m(\pi)|$ *is strictly decreasing in* $\lambda_0 + \lambda_1$ *for all* $\pi \in \,]0,1[$. *Hence experimentation decreases as the transition intensities increase.*

We have not been able so far to prove these conjectures in full generality. One particular case, however, allows for an analytic proof similar to the one given for Proposition 1.4.1: we show in the appendix that Conjecture 1.4.1 holds in the absence of state switches, that is, for $\lambda_0 = \lambda_1 = 0$.

1.5 Two-Sided Experimentation

Assume now that $\hat{q} > 0$ lies in the interior of Q^m, i.e., that it separates the optimal quantities corresponding to demand curves 0 and 1.[19] Given our standing assumption about $\Delta\beta$, this means $q^m(1) < \hat{q} < q^m(0)$ and the existence of a unique belief $\hat{\pi}$ such that $q^m(\hat{\pi}) = \hat{q}$. Calculation reveals that

$$\hat{\pi} = \frac{\alpha_0}{\Delta\alpha} - 2\,\frac{\beta_0}{\Delta\beta}.$$

This is the belief which would lead the agent to choose quantity $\hat{q}$ if he were to place no value on information. It is easily seen that $m(\hat{\pi}) = \hat{p}\hat{q} = \hat{m}$ is the minimum over all π of the myopic payoff function m.

This scenario is more complicated than the previous ones. First of all, the quantity $\hat{q}$ is special in so far as the expected price for this quantity equals $\hat{p}$ regardless

[18] *Ceteris paribus* means here in particular that $\tilde{\pi}$ is held fixed. In other words, λ_0 and λ_1 are increased by the same factor.

[19] We have not studied the border-line cases $\hat{q} = q^m(0)$ and $\hat{q} = q^m(1)$. We expect the results of the previous section to remain valid, with quantity expansion in the first case, and quantity reduction in the second.

of the state of demand. Choosing $\hat{q}$ therefore leads to a completely uninformative price signal.[20] As this constitutes a confounding action in the sense of Easley and Kiefer (1988), we shall refer to $\hat{q}$ as the confounding quantity. Lacking a better name, we shall sometimes call $\hat{\pi}$ the confounding belief. In a static environment ($\lambda_0 = \lambda_1 = 0$), the existence of a confounding quantity opens up the possibility of a cessation of learning. Indeed, if the monopolist finds it optimal to set quantity $\hat{q}$ at some stage,[21] his belief will not change over the next instant, so the confounding quantity will again be optimal. This pattern will repeat itself forever, and no more information is gathered.

A second complication arises from the fact that there is no longer an unambiguous direction of increasing informativeness of the price signal. Assume for example that the current belief is slightly higher than $\hat{\pi}$, so the myopically optimal quantity is slightly below $\hat{q}$. The true optimum will usually involve some deviation from the myopic quantity, motivated by the desire to render observed prices more informative. Following the logic of the one-sided experimentation encountered in the previous section, the monopolist might wish to reduce quantity. However, it could also make sense to *increase* quantity beyond $\hat{q}$ and thus achieve a wider spread between the two possible price distributions.

For beliefs close to the boundaries of the unit interval, on the other hand, we naturally expect the same experimentation behaviour with respect to the myopic quantity as in the one-sided scenarios of the previous section, that is, increasing quantity for beliefs π close to 0, and decreasing quantity for beliefs π close to 1. The optimal policy as a function of beliefs will then have to move downward past $\hat{q}$ as π increases. This raises the following question: does the optimal quantity change continuously, or is there a jump?

In either case, we can no longer expect the ODE (1.14) to characterise the value function on the entire unit interval. We therefore have to go back to the Bellman equation and the description of the optimal policy obtained earlier. We know from

[20] Note that the diffusion coefficient in (1.3) vanishes for $q = \hat{q}$.

[21] We shall see below that this can only happen at belief $\hat{\pi}$.

Section 1.3.5 that with the notation

$$V(\pi) = \frac{\Delta\beta^2}{2r\sigma^2}\pi^2(1-\pi)^2\, u_2(\pi)$$

and

$$v(\pi) = u(\pi) - \frac{\lambda(\pi)}{r}\, u'(\pi),$$

we have the Bellman equation

$$v(\pi) = m(\pi) + \sup_{q\in Q} \Phi[V(\pi), \pi, q]$$

where

$$\Phi[V,\pi,q] = V\,[q-\hat{q}]^2 - \beta(\pi)\,[q-q^m(\pi)]^2.$$

The calculation of the optimal quantity for a given belief depends on whether Φ as a function of q is linear, convex or concave.

1.5.1 Analysing the Bellman Equation

We first convince ourselves that the confounding quantity can be optimal only for the belief $\hat{\pi}$.

Lemma 1.5.1 *The confounding quantity $\hat{q}$ is strictly suboptimal at any belief $\pi \neq \hat{\pi}$ where the value function u satisfies the Bellman equation.*

PROOF: $\Phi(\hat{q}) < 0 \leq \Phi(q^m(\pi))$, so $\hat{q}$ is strictly dominated. ■

Thus, if the optimal policy function q^* is continuous, it will satisfy $q^*(\hat{\pi}) = \hat{q}$ and $q^*(\pi) \neq \hat{q}$ else. Moreover, the Bellman equation then suggests the interior boundary condition $v(\hat{\pi}) = \hat{m}$.[22] Conversely, a value function with $v(\hat{\pi}) > \hat{m}$ indicates an optimal policy function that jumps past the confounding quantity without ever assuming it.

To study this possibility in more detail, consider now the open strip $]0,1[\times I\!R$ with generic element (π, v). We are concerned with the area $\mathcal{A}$ strictly above the curve $v = m(\pi)$:

$$\mathcal{A} = \big\{(\pi,v) \in\ ]0,1[\ \times I\!R:\ v > m(\pi)\big\}.$$

[22] In the static case ($\lambda_0 = \lambda_1 = 0$), this condition follows independently of the Bellman equation. If $q^*(\hat{\pi}) = \hat{q}$, then $u(\hat{\pi}) = \int_0^\infty r\,e^{-r\,t}\, R(\hat{\pi},\hat{q})\,dt = R(\hat{\pi},\hat{q}) = \hat{m}$ as the belief $\hat{\pi}$ and the quantity $\hat{q}$ will prevail forever.

We shall show that this area can be divided into three regions, each with two sub-regions, by rays emanating from the point $(\hat{\pi}, \hat{m})$, which is the lowest point on the curve $v = m(\pi)$.

The two major rays correspond to the linear case $V(\pi) = \beta(\pi)$, one for π to the left of $\hat{\pi}$ and one for π to the right of $\hat{\pi}$. The region between these major rays is associated with the convex case $V(\pi) > \beta(\pi)$, in which the maximisation problem has a corner solution, and it can be further sub-divided by another ray, the two sub-regions being associated with the optimal quantities $q^* = q_{\max}$ and $q^* = q_{\min}$. The other two regions are both associated with the concave case $V(\pi) < \beta(\pi)$ and they can also be further sub-divided by minor rays.

In each case one sub-region is associated with the problem having an interior solution, and the other with the problem having a corner solution.

To summarise, moving clockwise from the left, we shall have

- Φ concave:
  - o interior solution
  - o corner solution $q_{\max}$
- Φ convex:
  - o corner solution $q_{\max}$
  - o corner solution $q_{\min}$
- Φ concave:
  - o corner solution $q_{\min}$
  - o interior solution.

For the derivations which follow, it is more convenient to work with the function Ψ instead of Φ, where $\hat{m} + \Psi[V, \pi, q] \equiv m(\pi) + \Phi[V, \pi, q]$. Simple algebra involving the relationships $m(\pi) = \hat{m} + \beta(\pi)\,[q^m(\pi) - \hat{q}]^2$ and $q^m(\pi) - \hat{q} = -\frac{1}{2}\,\Delta\alpha(\pi - \hat{\pi})/\beta(\pi)$ leads to the following expressions for Ψ and its derivative as a function of q:

$$
\begin{aligned}
\Psi[V, \pi, q] &= (V - \beta(\pi))\,[q - \hat{q}]^2 - \Delta\alpha\,[q - \hat{q}]\,(\pi - \hat{\pi}), \\
\Psi'(q) &= 2(V - \beta(\pi))\,[q - \hat{q}] - \Delta\alpha\,(\pi - \hat{\pi}).
\end{aligned}
$$

Note that $\Psi'(q) = \Phi'(q)$ by construction. The Bellman equation (1.10) now becomes

$$v(\pi) = \hat{m} + \sup_{q \in Q} \Psi[V(\pi), \pi, q].$$

We know from Section 1.3.5 that the optimal quantity and hence the value of $v(\pi)$ are found by evaluating $\Psi'(q)$ for various values of q.

Let us first determine the regions in $\mathcal{A}$ where $\Psi(q)$ is strictly concave, linear, or strictly convex.

Lemma 1.5.2 *Let π be a belief at which u satisfies the Bellman equation with $v(\pi) > m(\pi)$. Then:*

(i) Strictly concave case:

$$V(\pi) < \beta(\pi) \iff \begin{cases} \pi < \hat{\pi} \quad \textit{and} \quad v(\pi) < \hat{m} - \Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}); \\ \\ \pi > \hat{\pi} \quad \textit{and} \quad v(\pi) < \hat{m} + \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}). \end{cases}$$

(ii) Linear case:

$$V(\pi) = \beta(\pi) \iff \begin{cases} \pi < \hat{\pi} \quad \textit{and} \quad v(\pi) = \hat{m} - \Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}); \\ \\ \pi > \hat{\pi} \quad \textit{and} \quad v(\pi) = \hat{m} + \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}). \end{cases}$$

(iii) Strictly convex case:

$$V(\pi) > \beta(\pi) \iff \begin{cases} \pi \leq \hat{\pi} \quad \textit{and} \quad v(\pi) > \hat{m} - \Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}); \\ \\ \pi \geq \hat{\pi} \quad \textit{and} \quad v(\pi) > \hat{m} + \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}). \end{cases}$$

PROOF: We first prove the implications "$\Rightarrow$", starting with (ii). If $V(\pi) = \beta(\pi)$, then $\Psi(q) = -\Delta\alpha\,[q - \hat{q}]\,(\pi - \hat{\pi})$. Suppose $\pi = \hat{\pi}$. Then $\Psi(q) = 0$ and $v(\pi) = \hat{m}$, which is excluded by our requirement that $v(\pi) > m(\pi)$. So we necessarily have $\pi \neq \hat{\pi}$. When $\pi < \hat{\pi}$, $\Psi(q)$ is maximised at $q_{\max}$, hence $v(\pi) = \hat{m} - \Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi})$. On the other hand, when $\pi > \hat{\pi}$, $\Psi(q)$ is maximised at $q_{\min}$, so $v(\pi) = \hat{m} + \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi})$. Moving on to (i), we see that the case $\pi = \hat{\pi}$ is again excluded. Indeed, for $\pi = \hat{\pi}$, $\Psi(q) = (V - \beta(\hat{\pi}))\,[q - \hat{q}]^2$ is maximised at $\hat{q}$ with $\Psi(\hat{q}) = 0$ and hence $v(\pi) = \hat{m}$. Now, if $\pi < \hat{\pi}$, then $\Psi(q) \leq -\Delta\alpha\,[q - \hat{q}]\,(\pi - \hat{\pi}) \leq -\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi})$ for all admissible q. The first inequality is strict for $q \neq \hat{q}$, the second for $q \neq q_{\max}$, so $v(\pi) < \hat{m} - \Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi})$. If, on the other hand, $\pi > \hat{\pi}$, then $\Psi(q) \leq \Delta\alpha\,[\hat{q} - q]\,(\pi - \hat{\pi}) \leq \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi})$ and $v(\pi) < \hat{m} + \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi})$ by the same argument. As for (iii), one obviously has $\sup_q \Psi(q) \geq \Psi(q_{\max}) > -\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi})$ and $\sup_q \Psi(q) \geq \Psi(q_{\min}) > \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi})$. This yields two strict lower bounds for $v(\pi)$. Note that the second bound is redundant for $\pi \leq \hat{\pi}$, the first for $\pi \geq \hat{\pi}$. This completes the proof of the sufficiency part. The implications "$\Leftarrow$" now follow from the fact that the conditions to the left and to the right of the equivalence sign fully exhaust the set of possible values of $V(\pi)$ and the area $\mathcal{A}$, respectively. ■

Thus, the two major rays that separate the region of convexity from the regions of concavity are

$$v = \hat{m} - \Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}),$$

going up and to the left from $(\hat{\pi}, \hat{m})$, and

$$v = \hat{m} + \Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}),$$

going up and to the right.

Next, we sub-divide the regions where $\Psi(q)$ is strictly concave.

Lemma 1.5.3 *Let π be a belief at which u satisfies the Bellman equation with $v(\pi) > m(\pi)$. Assume that $V(\pi) < \beta(\pi)$. Then the optimal policy can be characterised as follows.*

(i) Corner solution $q_{\max}$:

$$q^* = q_{\max} \iff \pi < \hat{\pi} \quad \textit{and} \quad v(\pi) \geq \hat{m} - \tfrac{1}{2}\,\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}).$$

(ii) Interior solutions:

$$q^* \in \,]q_{\min}, q_{\max}[\, \iff \begin{cases} \pi < \hat{\pi} \quad \textit{and} \quad v(\pi) < \hat{m} - \tfrac{1}{2}\,\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}); \\ \\ \pi > \hat{\pi} \quad \textit{and} \quad v(\pi) < \hat{m} + \tfrac{1}{2}\,\Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}). \end{cases}$$

(iii) Corner solution $q_{\min}$:

$$q^* = q_{\min} \Longleftrightarrow \quad \pi > \hat{\pi} \quad \textit{and} \quad v(\pi) \geq \hat{m} + \tfrac{1}{2}\,\Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}).$$

PROOF: As in the last proof, it is enough to show the implications "$\Rightarrow$", as we have again exhaustive sets of conditions on either side of the equivalence sign. We first deal with (i) and (iii). According to (1.11), $q_{\max}$ is optimal only if $\Psi'(q_{\max}) \geq 0$. This is equivalent to $V(\pi) - \beta(\pi) \geq \frac{1}{2}\,\Delta\alpha\,(\pi - \hat{\pi})/[q_{\max} - \hat{q}]$, and now $V(\pi) < \beta(\pi)$ confirms $\pi < \hat{\pi}$. Substituting the inequality for $V(\pi) - \beta(\pi)$ into the expression for $\Psi(q_{\max})$ leads to $v(\pi) \geq \hat{m} - \frac{1}{2}\,\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi})$. On the other hand, the corner solution $q_{\min}$ occurs only if $\Psi'(q_{\min}) \leq 0$, that is, $V(\pi) - \beta(\pi) \geq -\frac{1}{2}\,\Delta\alpha\,(\pi - \hat{\pi})/[\hat{q} - q_{\min}]$. In this case, $V(\pi) < \beta(\pi)$ confirms that $\pi > \hat{\pi}$. Substituting the inequality for $V(\pi) - \beta(\pi)$ into $\Psi(q_{\min})$ yields the bound $v(\pi) \geq \hat{m} + \frac{1}{2}\,\Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi})$. As for the remaining sub-case (ii), an inner solution requires $\Psi'(q_{\max}) < 0$ and $\Psi'(q_{\min}) > 0$. We only have to consider beliefs $\pi \neq \hat{\pi}$. If $\pi < \hat{\pi}$, the second inequality is redundant, while the first one implies $V(\pi) - \beta(\pi) < \frac{1}{2}\,\Delta\alpha\,(\pi - \hat{\pi})/[q_{\max} - \hat{q}]$ and

$$\Psi(q) \leq \frac{\frac{1}{2}\,\Delta\alpha\,(\pi - \hat{\pi})}{[q_{\max} - \hat{q}]}\,[q - \hat{q}]^2 - \Delta\alpha\,(\pi - \hat{\pi})\,[q - \hat{q}].$$

This inequality is strict unless $q = \hat{q}$. Its right hand side has a strict maximum in $q_{\max}$, so $\Psi(q) < -\frac{1}{2}\,\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi})$ which leads to the desired inequality for $v(\pi)$. A similar argument can be given for $\pi > \hat{\pi}$. ■

Thus, the two minor rays that sub-divide the regions of concavity are

$$v = \hat{m} - \tfrac{1}{2}\,\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}),$$

corresponding to the condition $\Psi'(q_{\max}) = 0$, and

$$v = \hat{m} + \tfrac{1}{2}\,\Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}),$$

corresponding to the condition $\Psi'(q_{\min}) = 0$. Both leave from $(\hat{\pi}, \hat{m})$ with half the slope of their respective major rays. Alternatively, we can write

$$v \;=\; m(\pi) + \beta(\pi)\,[q_{\max} - q^m(\pi)]\,[q^m(\pi) - \hat{q}]$$

for the first minor ray and

$$v = m(\pi) + \beta(\pi)\,[q^m(\pi) - q_{\min}]\,[\hat{q} - q^m(\pi)]$$

for the second. With the equations in this form it is easily seen that the rays intersect the axis $\pi = 0$ above $m(0)$, and the line $\pi = 1$ above $m(1)$.

Finally, the region where $\Psi(q)$ is strictly convex consists also of two sub-regions, defined by the optimality of $q_{\min}$ or $q_{\max}$, respectively. Recall that $q_c = \frac{1}{2}(q_{\min} + q_{\max})$ defines the centre of the interval $[q_{\min}, q_{\max}]$.

Lemma 1.5.4 *Let π be a belief at which u satisfies the Bellman equation with $v(\pi) > m(\pi)$. Assume that $V(\pi) > \beta(\pi)$. Then the optimal quantities can be characterised as follows.*

$$q^* = \begin{cases} q_{\max} & \Leftrightarrow \quad \pi < \hat{\pi} + \frac{2}{\Delta\alpha}\frac{q_c - \hat{q}}{[q_{\max}-\hat{q}][\hat{q}-q_{\min}]}\,(v(\pi) - \hat{m})\ ; \\ \{q_{\min}, q_{\max}\} & \Leftrightarrow \quad \pi = \hat{\pi} + \frac{2}{\Delta\alpha}\frac{q_c - \hat{q}}{[q_{\max}-\hat{q}][\hat{q}-q_{\min}]}\,(v(\pi) - \hat{m})\ ; \\ q_{\min} & \Leftrightarrow \quad \pi > \hat{\pi} + \frac{2}{\Delta\alpha}\frac{q_c - \hat{q}}{[q_{\max}-\hat{q}][\hat{q}-q_{\min}]}\,(v(\pi) - \hat{m})\ . \end{cases}$$

PROOF: From (1.13), a necessary condition for both extreme quantities to be optimal is that $\Psi'(q_c) = 0$ or $2\,(V(\pi) - \beta(\pi))\,[q_c - \hat{q}] - \Delta\alpha\,(\pi - \hat{\pi}) = 0$. Combining this with the fact that $v(\pi) = \hat{m} + \Psi(q_{\max})$, one obtains

$$\pi = \hat{\pi} + \frac{2}{\Delta\alpha}\,\frac{q_c - \hat{q}}{[q_{\max} - \hat{q}][\hat{q} - q_{\min}]}\,(v(\pi) - \hat{m})\ .$$

The stated inequalities for π when only one extreme quantity is optimal follow immediately. This proves the implications "$\Rightarrow$". The same argument as in the previous proofs yields the implications "$\Leftarrow$". ∎

Our third minor ray is therefore parameterised by

$$\pi = \hat{\pi} + \frac{2}{\Delta\alpha}\,\frac{q_c - \hat{q}}{[q_{\max} - \hat{q}]\,[\hat{q} - q_{\min}]}\,(v - \hat{m})\ .$$

For $q_c \neq \hat{q}$, this can also be written as

$$v = \hat{m} + \tfrac{1}{2}\Delta\alpha\frac{[q_{\max} - \hat{q}]\,[\hat{q} - q_{\min}]}{q_c - \hat{q}}\,(\pi - \hat{\pi})\ .$$

The ray goes up and to the right if $q_c > \hat{q}$, and up and to the left if $q_c < \hat{q}$. For $q_c = \hat{q}$, the ray is simply the vertical line $\pi = \hat{\pi}$. In any case, $q_{\max}$ is optimal to the left of the ray, $q_{\min}$ is optimal to the right, and along the ray itself, either extreme quantity is optimal. If there are jumps in the optimal policy, we expect them to occur at beliefs π such that the corresponding point $(\pi, v(\pi))$ lies on the third minor ray. Moreover, we expect jumps to occur to the right of $\hat{\pi}$ if $q_c > \hat{q}$, and to the left of $\hat{\pi}$ if $q_c < \hat{q}$. Only if $q_c = \hat{q}$ should we expect jumps at exactly $\hat{\pi}$.

1.5.2 A Differential Equation for the Value Function

The last three lemmata show how to construct the optimal policy corresponding to a solution u of the Bellman equation. Inserting the optimal quantities into the Bellman equation, we obtain an equation linking the generalised second derivative $u_2(\pi)$ to π and $v(\pi)$. We find the ODE (1.14) again as long as the point $(\pi, v(\pi))$ lies below the first or the second minor ray. For $(\pi, v(\pi))$ above these two rays, we have versions of the backward equation (1.5) for constant quantity $q_{\max}$ or $q_{\min}$, respectively. More precisely, if the value function satisfies the Bellman equation at a belief π with $v(\pi) > m(\pi)$, then

$$u_2(\pi) = G(\pi, v(\pi)) \tag{1.16}$$

where the function G is defined as follows:

$$G(\pi, v) = \begin{cases} G_{\text{int}}(\pi, v) & \text{if} \quad m(\pi) < v \leq \hat{m} - \frac{1}{2}\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}) \\ & \text{or} \quad m(\pi) < v \leq \hat{m} + \frac{1}{2}\Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi})\ ; \\ & \\ G_{\max}(\pi, v) & \text{if} \quad v \geq \hat{m} - \frac{1}{2}\Delta\alpha\,[q_{\max} - \hat{q}]\,(\pi - \hat{\pi}) \\ & \text{and} \quad \pi \leq \hat{\pi} + \frac{2}{\Delta\alpha}\,\frac{q_c - \hat{q}}{[q_{\max} - \hat{q}]\,[\hat{q} - q_{\min}]}\,(v - \hat{m})\ ; \\ & \\ G_{\min}(\pi, v) & \text{if} \quad v \geq \hat{m} + \frac{1}{2}\Delta\alpha\,[\hat{q} - q_{\min}]\,(\pi - \hat{\pi}) \\ & \text{and} \quad \pi \geq \hat{\pi} + \frac{2}{\Delta\alpha}\,\frac{q_c - \hat{q}}{[q_{\max} - \hat{q}]\,[\hat{q} - q_{\min}]}\,(v - \hat{m}) \end{cases}$$

with

$$G_{\mathrm{int}}(\pi, v) = \frac{2r\sigma^2}{\Delta\beta^2} \frac{\beta(\pi)}{\pi^2(1-\pi)^2} \frac{v - m(\pi)}{v - \hat{m}} ,$$

$$G_{\mathrm{max}}(\pi, v) = \frac{2r\sigma^2}{\Delta\beta^2} \frac{1}{\pi^2(1-\pi)^2} \frac{v - R(\pi, q_{\mathrm{max}})}{[q_{\mathrm{max}} - \hat{q}]^2} ,$$

$$G_{\mathrm{min}}(\pi, v) = \frac{2r\sigma^2}{\Delta\beta^2} \frac{1}{\pi^2(1-\pi)^2} \frac{v - R(\pi, q_{\mathrm{min}})}{[\hat{q} - q_{\mathrm{min}}]^2} .$$

Using the different representations of the three minor rays, it is straightforward to check that the values of the functions G_{int}, G_{max} and G_{min} match along the relevant rays, so G is continuous on $\mathcal{A}$.[23] By Corollary 1.3.1 the right-hand side of (1.16) is therefore continuous at least on $]0,1[-\{\hat{\pi}\}$, and on the whole open unit interval if $v(\hat{\pi}) \neq \hat{m}$. As (1.16) holds almost everywhere, this means that the generalised derivative u_2 has a version which is continuous at any π such that $v(\pi) > m(\pi)$. By Corollary 1.3.1 again, this implies that u is in fact twice continuously differentiable whenever $v(\pi) > m(\pi)$. We thus have

Proposition 1.5.1 *The value function solves the ordinary differential equation*

$$u''(\pi) = G(\pi, v(\pi)) \tag{1.17}$$

on the open set $\{\pi \in]0,1[: v(\pi) > m(\pi)\}$. *In particular, the value function is twice continuously differentiable on this set.*

We expect of course the same boundary conditions for the value function as in the previous section. Indeed:

[23] As for differentiability of G, the partial derivatives of G_{int} and G_{max} with respect to π and v coincide along the first minor ray, and the same is true for G_{int} and G_{min} along the second minor ray. Along the central minor ray, however, the partial derivatives of G_{max} and G_{min} do not match, so G is not differentiable there. (The derivatives with respect to v match if $q_c = \hat{q}$. The derivatives with respect to π are always different.) Everywhere else, by contrast, G is continuously differentiable.

Corollary 1.5.1 *The value function satisfies the boundary conditions $v(0) = m(0)$ and $v(1) = m(1)$.*

PROOF: Suppose for example that $v(0) > m(0)$. Using the continuity of v and (1.17), we can find $K > 0$ and $\epsilon > 0$ such that $u''(\pi) \geq K\pi^{-2}$ for $0 < \pi \leq \epsilon$. Then

$$u'(\pi) = u'(\epsilon) - \int_{\pi}^{\epsilon} u''(\xi)\, d\xi \leq u'(\epsilon) - K \int_{\pi}^{\epsilon} \frac{d\xi}{\xi^2} = u'(\epsilon) - \left[\frac{1}{\pi} - \frac{1}{\epsilon}\right] \longrightarrow -\infty$$

as $\pi \to 0$. This, however, contradicts the boundedness of u'. ■

1.5.3 The Static Case

As a benchmark for subsequent results, we first examine the case without state switching, i.e., $\lambda_0 = \lambda_1 = 0$. Our results are essentially the same as those obtained in the discrete-time learning literature. More specifically, our model with unknown but fixed expected demand curve provides a continuous-time version of the model in Kiefer (1989).

In order to state the main result for this case, we define

$$\overline{m}_0(\pi) = \frac{\hat{\pi} - \pi}{\hat{\pi}}\, m(0) + \frac{\pi}{\hat{\pi}}\, \hat{m}$$

for $0 \leq \pi \leq \hat{\pi}$, and

$$\overline{m}_1(\pi) = \frac{1 - \pi}{1 - \hat{\pi}}\hat{m} + \frac{\pi - \hat{\pi}}{1 - \hat{\pi}}\, m(1)$$

for $\hat{\pi} \leq \pi \leq 1$. These functions describe the rays joining $(\hat{\pi}, \hat{m})$ with $(0, m(0))$ and $(1, m(1))$, respectively. By strict convexity of the function m, we have $m < \overline{m}_0$ on $]0, \hat{\pi}[$, and $m < \overline{m}_1$ on $]\hat{\pi}, 1[$.

Proposition 1.5.2 *Let $\lambda_0 = \lambda_1 = 0$. The value function is continuously differentiable on $[0,1]$ and strictly convex. There are two cases, depending on the value of u at $\hat{\pi}$:*

If $u(\hat{\pi}) = \hat{m}$, then the value function solves the ordinary differential equation

$$u''(\pi) = G_{\text{int}}(\pi, u(\pi)) \tag{1.18}$$

with

$$m(\pi) < u(\pi) < \overline{m}_0(\pi)$$

on $]0,\hat{\pi}[$ *and*

$$m(\pi) < u(\pi) < \overline{m}_1(\pi)$$

on $]\hat{\pi},1[$. *The optimal policy function is continuous and assumes the value* $\hat{q}$ *at* $\hat{\pi}$.

If $u(\hat{\pi}) > \hat{m}$, *then* u *solves*

$$u''(\pi) = G(\pi, u(\pi)) \tag{1.19}$$

with

$$m(\pi) < u(\pi) < \overline{m}(\pi)$$

on $]0,1[$. *In this case, the optimal policy jumps from* $q_{\max}$ *to* $q_{\min}$ *as* u *crosses the third minor ray, and* $\hat{q}$ *is never optimal.*

The exact form of the optimal policy function is easily obtained from the partition of the set $\mathcal{A}$.

PROOF: We first convince ourselves that $u > m$ on $]0,1[-\{\pi\}$. Fix an initial belief π_0 in this set; without loss of generality, $\pi_0 < \hat{\pi}$. Consider the Markov strategy $\mathbf{q}$ generated by the following policy function: $q(\pi) = q^m(\pi)$ for $\pi < \pi_0$, and $q(\pi) = q^m(\pi_0)$ for $\pi \geq \pi_0$.[24] Obviously, $R(\pi, q(\pi)) \geq R(\pi, q^m(\pi_0))$, and this inequality is strict for $\pi < \pi_0$. Moreover, there is a non-zero probability that the associated belief process reaches the interval $[0, \pi_0[$. Thus, using the martingale property of the belief process and the linearity of $R(\pi, q^m(\pi_0))$ in π,

$$\begin{aligned} u^{\mathbf{q}}(\pi_0) &> E_{\pi_0}\left[\int_0^\infty r\, e^{-r\,t}\, R(\pi_t, q^m(\pi_0))\, dt\right] \\ &= \int_0^\infty r\, e^{-r\,t}\, R(\pi_0, q^m(\pi_0))\, dt \;=\; R(\pi_0, q^m(\pi_0)) \;=\; m(\pi_0). \end{aligned}$$

This implies $u(\pi_0) > m(\pi_0)$. As π_0 was arbitrary, we have the desired strict lower bound. By Proposition 1.5.1, this implies that u solves (1.19) on $]0,\hat{\pi}[$ and $]\hat{\pi},1[$. In particular, $u''(\pi) > 0$ on these two subintervals, hence strict convexity. The statement for the case $u(\hat{\pi}) > \hat{m}$ follows now immediately. If $u(\hat{\pi}) = \hat{m}$, strict convexity implies $u(\pi) < \overline{m}_0(\pi)$ on $]0,\hat{\pi}[$ and $u(\pi) < \overline{m}_1(\pi)$ on $]\hat{\pi},1[$. But the rays joining $(\hat{\pi},\hat{m})$ with $(0,m(0))$ and $(1,m(1))$ lie below the first and second minor ray, respectively. So we are in the case of an inner solution, hence the ODE (1.18). ■

[24]This is indeed a well-defined strategy (cf. the proof of Theorem 1.4.1).

It is straightforward to prove verification theorems for the two cases, i.e., statements that a solution u to the ODE (1.19) with regularity properties and bounds as in Proposition 1.5.2 is the value function.[25]

The learning literature suggests that the case $u(\hat{\pi}) = \hat{m}$ arises for high discount rates.[26]

Proposition 1.5.3 *Let $\lambda_0 = \lambda_1 = 0$. For given parameters α_0, α_1, β_0, β_1 and σ, there is a unique discount rate r^* such that the value function satisfies $u(\hat{\pi}) = \hat{m}$ if $r > r^*$, and $u(\hat{\pi}) > \hat{m}$ if $r < r^*$.*

PROOF: Let $u[r]$ denote the value function for discount rate r, and define $S = \{r > 0 : \; u[r](\hat{\pi}) > \hat{m}\}$. It can be shown that $u[r]$ converges to $\overline{m}$ as $r \to 0$, so S is non-empty. Using the same technique as in the one-sided case,[27] one easily shows that if $r \in S$ and $r' < r$, then $r' \in S$. Therefore, $r^* = \sup S$ has the property that $u[r](\hat{\pi}) > \hat{m}$ if $r < r^*$, and $u[r](\hat{\pi}) = \hat{m}$ if $r > r^*$. Moreover, it can be shown that $u[r]$ converges to m as $r \to \infty$, so r^* is a finite number. ■

Thus the optimal policy depends qualitatively on the discount rate r. For $r < r^*$, there is *extreme experimentation* over a certain range of beliefs. The value of information at $\hat{\pi}$ is so large that it pays the monopolist to avoid the confounding quantity $\hat{q}$ and choose one of the extreme quantities $q_{\min}$ or $q_{\max}$ instead. For $r > r^*$, by contrast, a scenario of *moderate experimentation* arises, with optimal quantities that lie always inside the range of the myopic policy function. In particular, the value of information at $\hat{\pi}$ is so small that it is not worth while experimenting at all: the loss of current revenue outweighs the potential gains from the information acquired. So the monopolist chooses the (myopically optimal) confounding quantity $\hat{q}$ at $\hat{\pi}$. But this implies that he ceases to experiment and learn: choosing $\hat{q}$ makes the diffusion term in (1.3) vanish, so the belief will not change any more, and the monopolist is caught in a trap. It can be shown that the belief process π converges with positive probability to $\hat{\pi}$. Thus there is a positive probability of *incomplete learning* even in the long run.

[25]Cf. the corresponding results for the general case presented in the following section.

[26]See for instance Easley and Kiefer (1988).

[27]Cf. the appendix.

This incomplete learning scenario has been extensively described in the discrete-time learning literature.[28] While our model reproduces this result for a static environment, we are mainly interested in the effect of state switching on the monopolist's optimal behaviour.

1.5.4 The General Case

Our results on one-sided experimentation (Theorem 1.4.1) suggest the following "invariance principle" for the introduction of state switching:

> "Results valid for the case without state switching ($\lambda_0 = \lambda_1 = 0$) carry over to the general case if, in all expressions, we formally replace $u(\pi)$ by $v(\pi) = u(\pi) - \lambda(\pi)u'(\pi)/r$, and leave the rest unchanged."

In particular, we would expect Proposition 1.5.2 to generalise in this way.

We have, however, not been able to prove such a generalisation so far. The missing links are spelt out in the following two conjectures. Both are supported by our numerical results.

Conjecture 1.5.1 *$v > m$ on $]0,1[-\{\hat{\pi}\}$.*

Conjecture 1.5.2 *v is strictly convex.*

If these conjectures are true, then the generalisation of Proposition 1.5.2 follows directly from Proposition 1.5.1.

Lacking a proof of these conjectures, we now formulate verification theorems giving sufficient conditions for the two scenarios we expect: a scenario of moderate experimentation where $v(\hat{\pi}) = \hat{m}$, which should occur for high switching intensities; and a scenario of extreme experimentation with $v(\hat{\pi}) > \hat{m}$ that should arise for low switching probabilities (provided, of course, that $r < r^*$).

We turn to the case $v(\hat{\pi}) = \hat{m}$ first.

Lemma 1.5.5 *There is a function u continuous on $[0,1]$ and once continuously differentiable on $[0,1] - \{\hat{\pi}\}$ which solves (1.14) on $]0,1[-\{\hat{\pi}\}$ with*

$$m(\pi) < u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi) < \overline{m}_0(\pi)$$

[28] Cf. Easley and Kiefer (1988), Aghion *et al.* (1991) and the references given in these papers.

on $]0,\hat{\pi}[$ *and*

$$m(\pi) < u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi) < \overline{m}_1(\pi)$$

on $]\hat{\pi},1[$.[29]

PROOF: Repeat the steps described in the appendix on the subintervals $[0,\hat{\pi}]$ and $[\hat{\pi},1]$. One shows in particular that $\overline{m}_0$ and $\overline{m}_1$ are strict supersolutions on the respective open subinterval. ■

For u as in the lemma, we introduce the following policy function q^*:

$$q^*(\pi) = \begin{cases} q^m(\pi) + \frac{u(\pi)-\lambda(\pi)u'(\pi)/r-m(\pi)}{m(\pi)-\hat{m}} \, [q^m(\pi) - \hat{q}] & \text{for} \quad \pi \neq \hat{\pi} \\ \\ \hat{q} & \text{for} \quad \pi = \hat{\pi}. \end{cases}$$

This is a candidate for the optimal policy.

Proposition 1.5.4 *Let u be as in the previous lemma. Suppose in addition that u is once continuously differentiable on the whole of $[0,1]$. If the corresponding function q^* is continuous and possesses (finite) one-sided derivatives at $\hat{\pi}$, then u is the value function and q^* defines an optimal policy.*

In particular, there is at most one such function u.

PROOF: For $\pi \neq \hat{\pi}$, (1.14) implies that we are in the strictly concave case $V(\pi) < \beta(\pi)$. As the graphs of $\overline{m}_0$ and $\overline{m}_1$ lie below the first and second minor ray, respectively, the optimal quantity is an inner solution, given by the above function q^*, and u satisfies the Bellman equation. Moreover, the difference $q^*(\pi) - \hat{q}$ for $\pi \neq \hat{\pi}$ can be written as

$$q^*(\pi) - \hat{q} = \frac{\beta(\pi)}{\beta(\pi) - V(\pi)} [q^m(\pi) - \hat{q}] = -\frac{\Delta\alpha}{2} \frac{\pi - \hat{\pi}}{\beta(\pi) - V(\pi)}.$$

By assumption, the one-sided limits $\lim_{\pi\to\hat{\pi}\pm} [q^*(\pi) - \hat{q}]/(\pi - \hat{\pi})$ exist, so V has one-sided limits $V(\hat{\pi}\pm) < \beta(\hat{\pi})$, and the one-sided limits $u''(\hat{\pi}\pm)$ exist as well. One can use $u''(\hat{\pi}+)$ or $u''(\hat{\pi}-)$ to define the value $u_2(\hat{\pi})$ of a generalised second derivative. As the corresponding function V satisfies $V(\hat{\pi}) < \beta(\hat{\pi})$, the confounding quantity

[29] In particular, $u(\pi) - \lambda(\pi)\, u'(\pi)/r$ coincides with $m(\pi)$ at $\pi = 0$, $\hat{\pi}$ and 1.

is optimal at $\hat{\pi}$, and the Bellman equation reduces to the condition $v(\hat{\pi}) = \hat{m}$. Thus, the given function u satisfies the Bellman equation at $\hat{\pi}$ as well. The same arguments as in the proof of Theorem 1.4.1 now yield the result.[30] ■

We expect functions u as in the proposition to exist for sufficiently high switching probabilities.[31]

Whenever the proposition applies, we have a scenario of moderate experimentation with quantity expansion for beliefs $\pi < \hat{\pi}$ and quantity reduction for $\pi > \hat{\pi}$, looking just like a combination of two one-sided scenarios. As in the static case for $r > r^*$, the value of information at $\hat{\pi}$ is too small to warrant any experimentation at all. Unless we are in the knife-edge case $\hat{\pi} = \tilde{\pi}$, however, experimentation will not cease at $\hat{\pi}$, since the possibility of a state switch still leads the agent to update his belief.

The lower the switching intensities, the more experimentation we expect. It is therefore unlikely that $\hat{q}$ remains optimal at $\hat{\pi}$ for low λ_0 and λ_1. For this case, we predict $v > m$ throughout the open unit interval. Moreover, the optimal policy is expected to jump from $q_{\max}$ to $q_{\min}$ as v crosses the third minor ray.

We now formulate a sufficient condition for this case, involving a solution of the ODE (1.17) with $v(\hat{\pi}) > \hat{m}$.

Proposition 1.5.5 *Let u be a solution of* (1.17) *with $v > m$ on* $]0,1[$. *Then it is the value function.*

Again, there can be at most one such solution.

PROOF: Using the partition of $\mathcal{A}$ and the corresponding lemmata, one easily verifies that the given function u solves the Bellman equation. Let q^* be an optimal policy

[30] The existence of finite one-sided derivatives of q^* at $\hat{\pi}$ implies Lipschitz-continuity which is a prerequisite for the existence result in Liptser and Shiryayev (1977, chapter 9). If u is as in the proposition, but the corresponding policy function q^* has one-sided slope $-\infty$ at $\hat{\pi}$, the existence of a strong solution to the stochastic differential equation (1.15) is no longer guaranteed. Nor can we use Theorem 2.6.1 in Krylov (1980) to establish existence of a weak solution since $q^*(\hat{\pi}) = \hat{q}$ implies a violation of Krylov's non-degeneracy condition for the diffusion term. However, if we were able to find for any $\epsilon > 0$ an ϵ-optimal Markov strategy given by a Lipschitz-continuous policy function, then the above proposition would hold without restrictions on the slope of $\mathbf{q}^*$ at $\hat{\pi}$.

[31] One might be able to prove this using the implicit function theorem on a suitable Banach space.

function.[32] To show that u is indeed the value function, we follow the logic of the proof of Theorem 1.4.1 and use the stochastic differential equation (1.15) to construct the process of beliefs for the Markov strategy determined by q^*. Given any initial value π_0, Theorem 2.6.1 in Krylov (1980) implies that the stochastic differential equation

$$d\pi_t = \lambda(\pi_t) + \frac{\Delta\beta}{\sigma}\pi_t(1-\pi_t)\left[\hat{q} - q^*(\pi_t)\right]dZ_t \tag{1.20}$$

has a weak solution (π^*, Z^0) with Z^0 a Wiener process. We extend the corresponding filtered probability space in such a way that it supports an independent Markov process $\{k_t\}$ taking values in $\{0,1\}$ with transition probabilities as in Section 1.1. Consider the bounded process

$$\eta_t = \frac{1}{\sigma}\Big(\alpha_{k_t} - \beta_{k_t}q^*(\pi_t^*) \;-\; \left[\alpha(\pi_t^*) - \beta(\pi_t^*)q^*(\pi_t^*)\right]\Big).$$

By Girsanov's theorem,[33] there is a new measure under which

$$Z_t = Z_t^0 - \int_0^t \eta_s\, ds$$

is a Wiener process. In other words, (π^*, Z) is a weak solution to the stochastic differential equation (1.15). To complete the proof, we now use the same arguments as in the proof of Theorem 1.4.1. ∎

We expect functions u as in the proposition to exist for sufficiently low switching probabilities.[34]

More precisely, we expect that there is a critical value of the switching probabilities at which the optimal behaviour changes discontinuously from extreme to moderate experimentation.

Conjecture 1.5.3 *Let parameters α_0, α_1, β_0, β_1, σ and $r < r^*$ be given. Fix $\tilde{\pi}$ and consider all pairs of switching intensities (λ_0, λ_1) such that $\lambda_0/(\lambda_0+\lambda_1) = \tilde{\pi}$. There is a unique $\lambda_0^* > 0$ such that Proposition* 1.5.4 *applies for $\lambda_0 > \lambda_0^*$ and Proposition* 1.5.5 *applies for $\lambda_0 < \lambda_0^*$.*

[32]One has to make a choice between $q_{\max}$ and $q_{\min}$ whenever both are optimal quantities in the Bellman equation.

[33]Cf. Revuz and Yor (1991).

[34]Again, it may be possible to prove this using the implicit function theorem on a suitable Banach space.

The numerical results presented in the next section support this conjecture.

1.6 Some Numerical Results

We report some numerical results for the two-sided case. They were obtained by calculating an approximate solution to the two-point boundary value problem in question, namely the ODE (1.17) subject to the boundary conditions $v(0) = m(0)$ and $v(1) = m(1)$. By Propositions 1.5.2, 1.5.4 and 1.5.5, these numerical solutions are approximations of the value function.

The graphs illustrate the impact, *ceteris paribus*, of changes in the discount rate or the probability of state switches on the monopolist's behaviour:

- The monopolist experiments more as the discount rate decreases.
- The monopolist experiments more as the probability of state switches decreases.[35]

The graphs also show the discontinuous change in the optimal policy as the discount rate or the switching probabilities cross their critical levels.

We used the following demand curve parameters:

- $\alpha_0 = 40$, $\beta_0 = 2/3$;
- $\alpha_1 = 60$, $\beta_1 = 3/2$.

They imply $q^m(0) = 30$, $q^m(1) = 20$, $\hat{q} = 24$ and $\hat{\pi} = 0.4$.

We further set the range of feasible quantities

- $q_{\min} = 13\frac{1}{3}$, $q_{\max} = 40$.

The values for $q_{\min}$ and $q_{\max}$ are derived from the points where the demand curves cross the axes. As shown in Figure 1.1, $q_{\max}$ is the smaller of the two values where the demand curves cross the quantity axis – a quantity larger than this generates a negative expected price in state 1 (where the demand curve is steeper); $p_{\min}$ is the expected price corresponding to $q_{\max}$ in state 0. Similarly, $p_{\max}$ is the lower of

[35]Here, in particular, *ceteris paribus* means that $\tilde{\pi}$ is held fixed.

the two values where the demand curves cross the price axis – a price higher than this generates a negative expected quantity in state 0 (where the demand curve is flatter); $q_{\min}$ is the expected quantity corresponding to $p_{\max}$ in state 1.

Finally, we chose the noise parameter

- $\sigma = 5$.

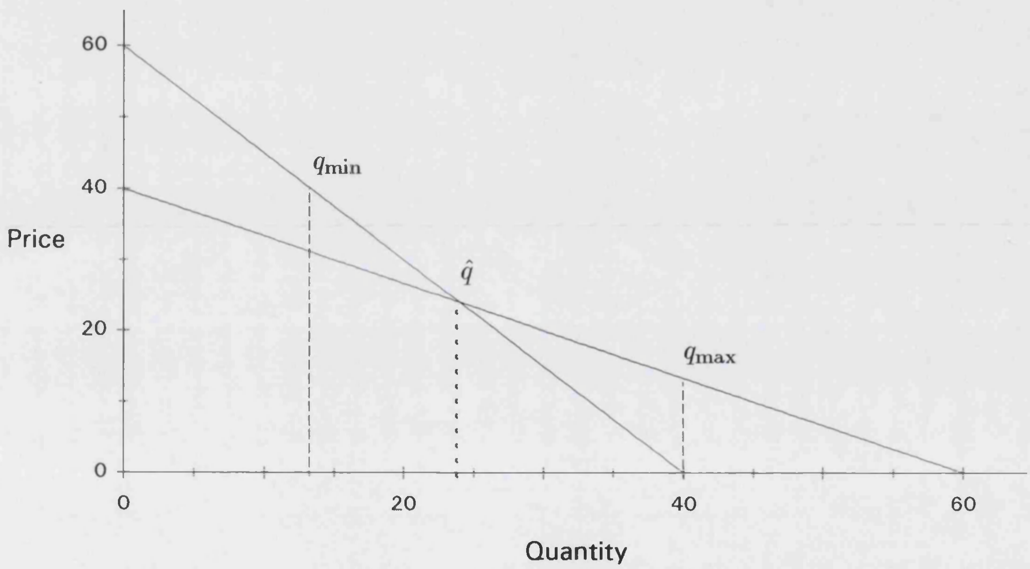

Figure 1.1: The two demand curves used in the simulations

1.6.1 The Value Function and Optimal Policy

We first consider cases without state switching.

No state switching, moderate experimentation. Figure 1.2 shows that the relatively high discount rate $r = 0.5$ implies moderate experimentation.[36] We have $u(\hat{\pi}) = \hat{m}$ and the optimal policy function is very close to the myopic policy, going continuously through $\hat{q}$. If we reduce the discount rate (not shown), the incentive to experiment increases, and both the differences $u - m$ and $q^* - q^m$ grow. This change is gradual until we reach the critical discount rate r^* where optimal behaviour switches suddenly from moderate to extreme experimentation.

No state switching, extreme experimentation. The discount rate in Figure 1.3, $r = 0.1$, is already some way below the critical level.[37] The value function has lifted off $\hat{m}$ at $\hat{\pi}$, and the optimal policy function jumps from $q_{\max}$ to $q_{\min}$ just after $\hat{\pi}$. (It is *after* $\hat{\pi}$ because q_c, the midpoint of the interval of feasible quantities, lies above $\hat{q}$.) A further reduction of the discount rate (not shown) increases the distance between u and m. At the same time, the extreme quantities $q_{\max}$ and $q_{\min}$ become optimal over a larger range of beliefs: as u goes up, its intersection with the left minor ray moves to the left and its intersection with the right minor ray moves to the right. Also, its intersection with the central minor ray moves to the right, so the jump occurs further to the right of $\hat{\pi} = 0.4$.

Introducing state switches and increasing the switching probabilities has an adverse effect on experimentation. This is seen in the next results where $r = 0.1$ and $\tilde{\pi} = 0.5$ are held fixed while $\lambda_0 = \lambda_1$ assume different non-zero values.

Slow state switching, extreme experimentation. With $\lambda_0 = \lambda_1 = 0.025$, we still have a case of extreme experimentation, with $v > m$ on the open unit interval as shown in Figure 1.4.[38] The range where one of the extreme quantities is optimal has shrunk compared to the case without state switches. Note also that u

[36]The bold line in the upper panel is the graph of the function v (also u in case $\lambda_0 = \lambda_1 = 0$), the thin line that of the myopic payoff function m. In the lower panel, the bold line is the optimal policy function q^*, while the thin line is the myopic policy q^m.

[37]The upper panel now also shows the three minor rays introduced in the previous section, and the lower panel has an enlarged vertical axis to accommodate extreme quantities. It can be seen that as v crosses each ray in turn, q^* first reaches $q_{\max}$, then jumps, and finally leaves $q_{\min}$.

[38]We now show v and u separately in the top figure, u being plotted as a bold dashed line.

has become flatter, which illustrates the lower value of information. An increase in the switching intensities to $\lambda = 0.05$ reduces experimentation further (not shown). While v is still above m, it has moved down, and the extreme quantities are optimal only in a very small region around the confounding belief.

Fast state switching, moderate experimentation. A further increase in the switching intensities to 0.075 leads to Figure 1.5. Now v touches m at $\hat{\pi}$: we have moderate experimentation with quantity expansion for $\pi < \hat{\pi}$ and quantity reduction for $\pi > \hat{\pi}$. With switching intensities yet higher, we get a further reduction in the distance between v and m, and in particular, the policy function q^* becomes less steep around $\hat{\pi}$.

Note that our numerical results support the conjectures stated in Section 1.5.4. The numerically calculated functions v seem to satisfy $v(\pi) > m(\pi)$ on the subintervals $]0, \hat{\pi}[$ and $]\hat{\pi}, 1[$, and they appear to be strictly concave. Moreover, the results suggest a critical level for $\lambda_0 = \lambda_1$ just above 0.05 at which the optimal behaviour changes discontinuously from extreme to moderate experimentation.

Technical notes

The numerical results were obtained by using the method of *relaxation*.[39] Beliefs were discretised with a step size 10^{-3}, decreasing to 10^{-5} around the critical values where the optimal policy takes extreme values and jumps. The iterative procedure was deemed to have converged when the maximum pointwise difference between successive approximations to the value function and its first derivative were less than 0.0001%. Convergence was quite rapid, varying from 5 iterations for the high discount rate without switching (e.g. Figure 1.2), through 8 iterations for the low discount rate without switching (Figure 1.3) and the low discount rate with high switching intensity (e.g. Figure 1.5), to 18 iterations for the low discount rate with an intermediate switching intensity close to the critical level (not shown).

The procedure was implemented on a VAX minicomputer under VMS v5.4 (*not* a supercomputer) and as each iteration took approximately 19 seconds of CPU time, the numerical solutions each took between only 1.5 and 6 minutes to calculate. This further highlights the advantage of our approach.

[39] See Press *et al.* (1988), Chapter 16.

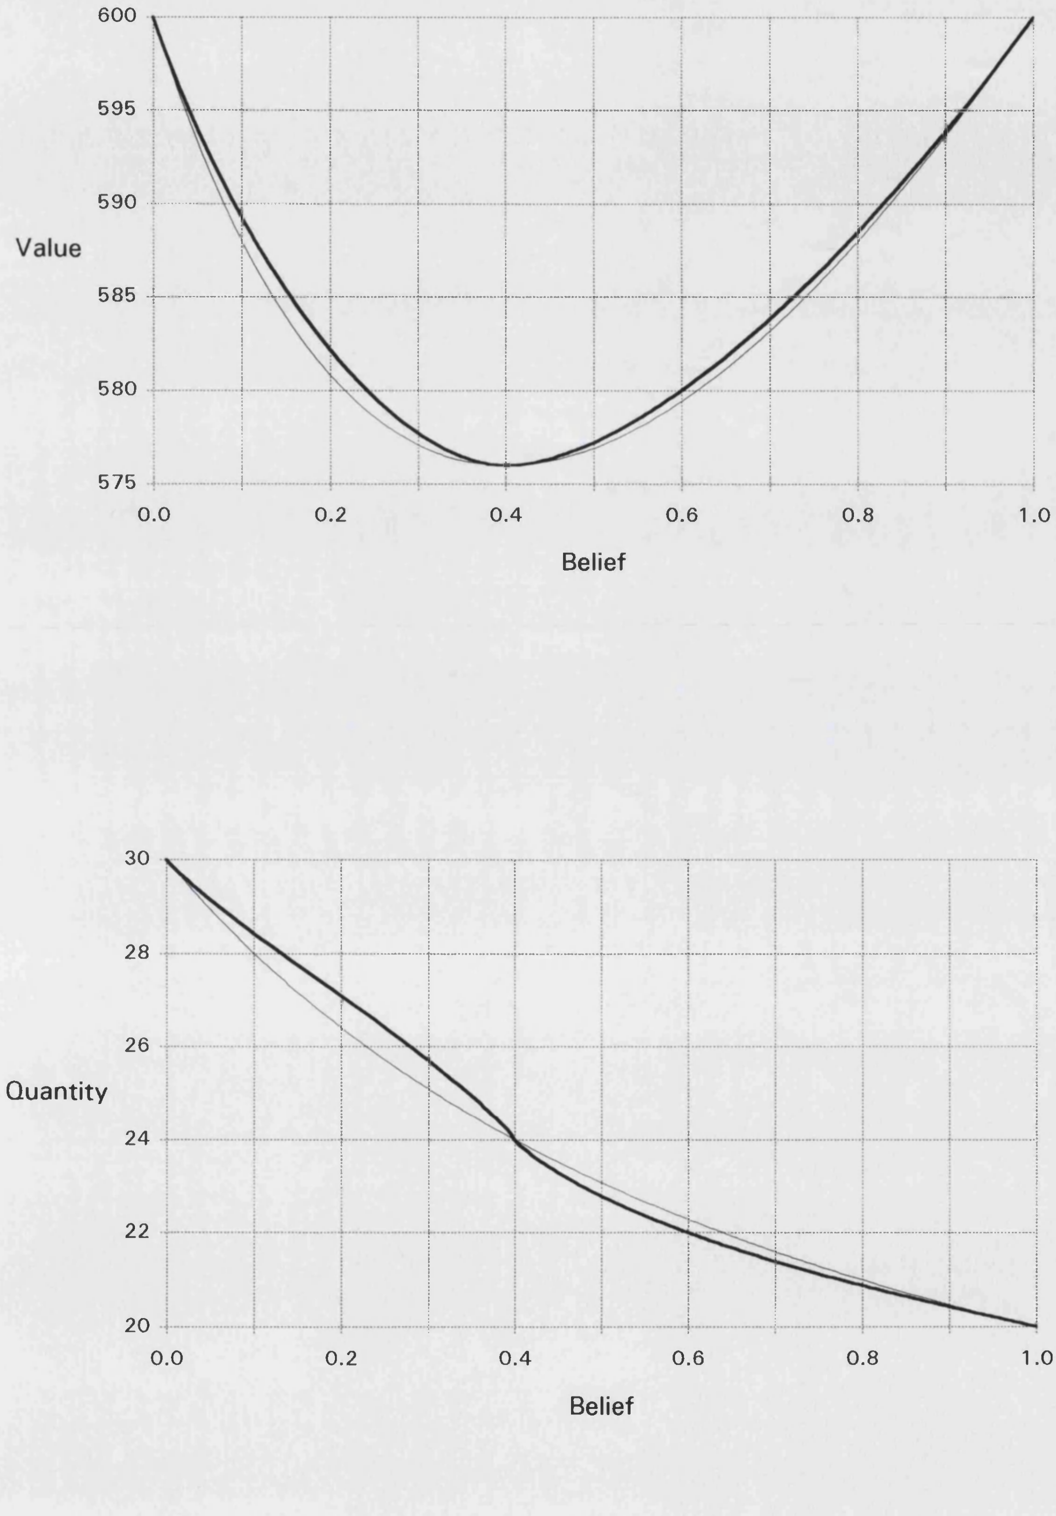

Figure 1.2: Value function and optimal policy for $r = 0.5$, $\lambda_0 = \lambda_1 = 0$

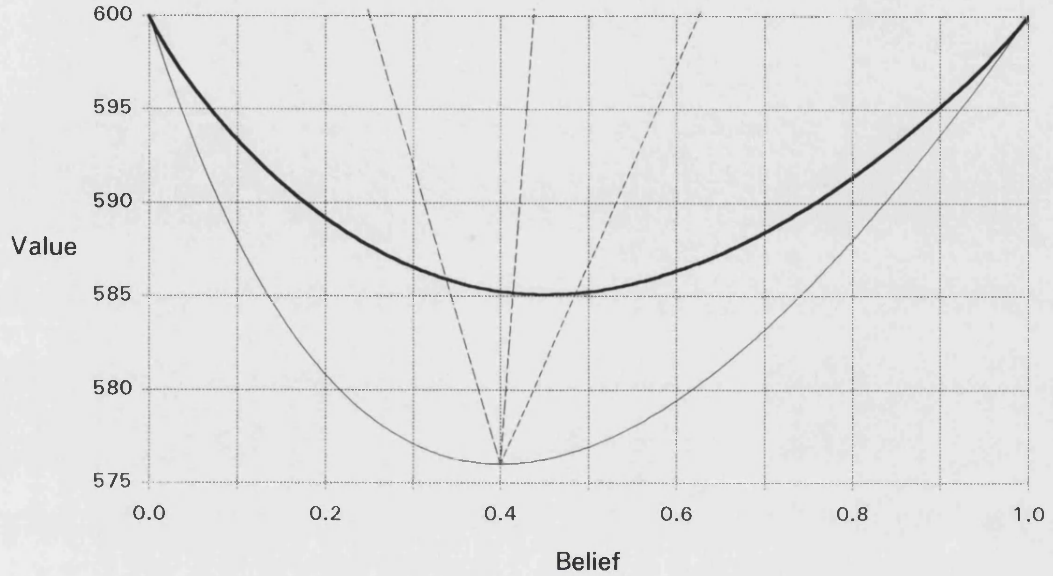

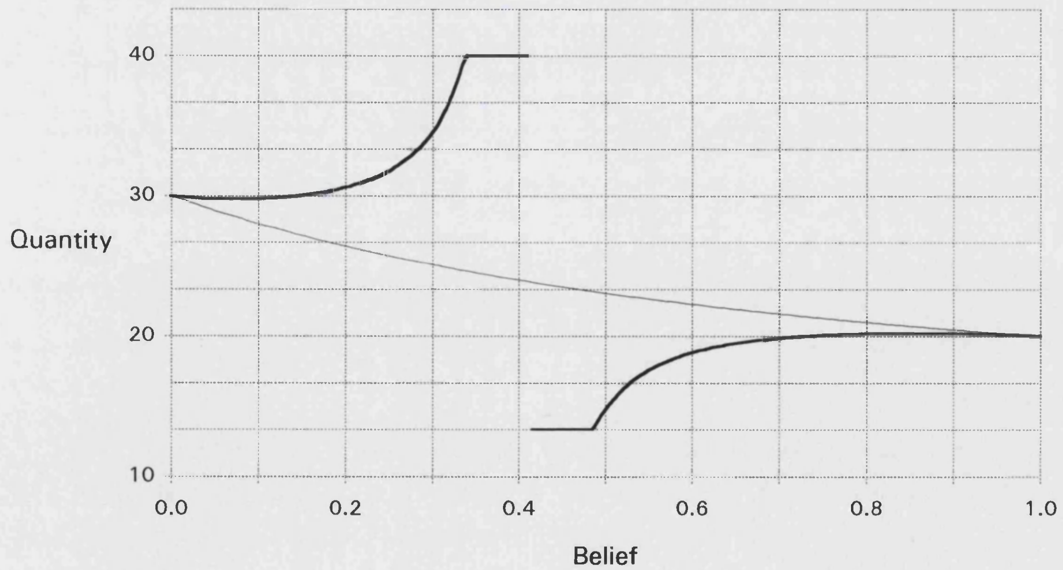

Figure 1.3: Value function and optimal policy for $r = 0.1$, $\lambda_0 = \lambda_1 = 0$

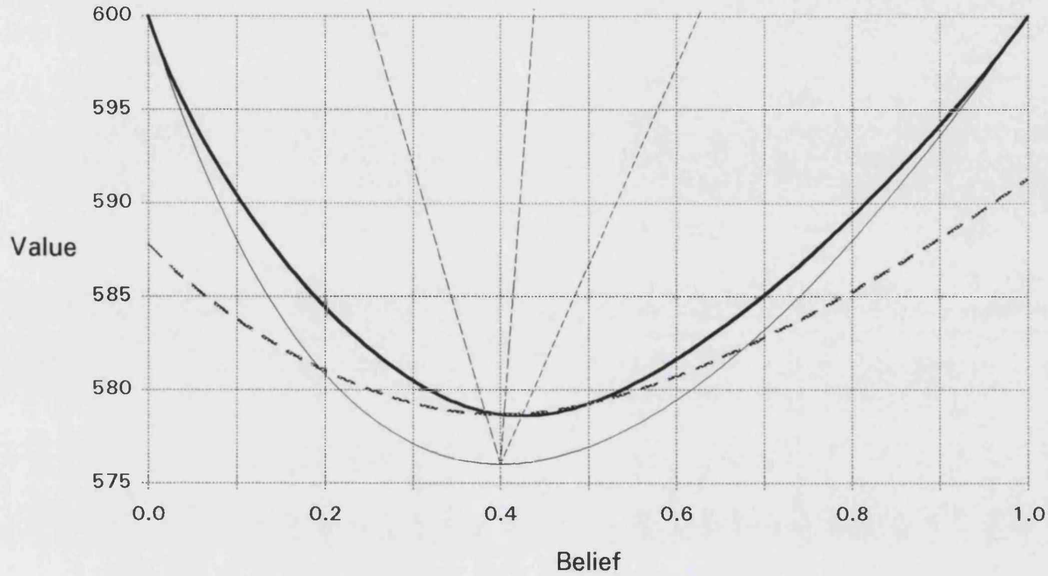

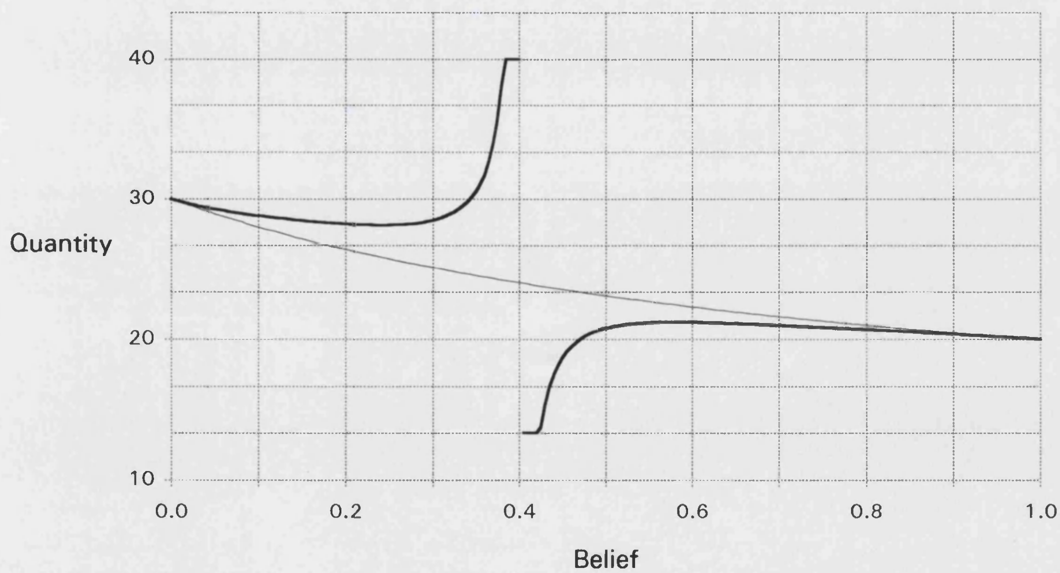

Figure 1.4: Value function and optimal policy for $r = 0.1$, $\lambda_0 = \lambda_1 = 0.025$

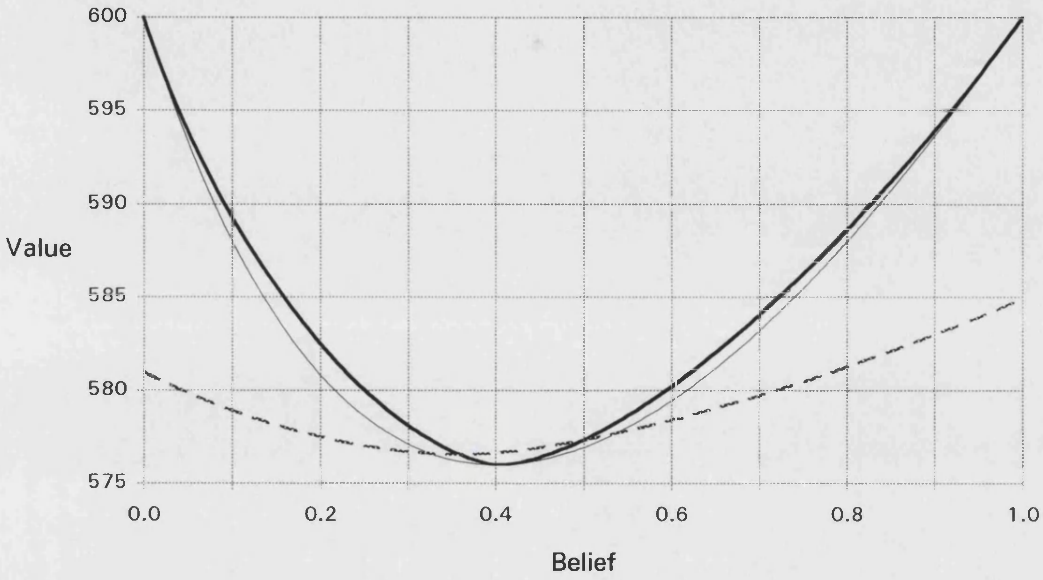

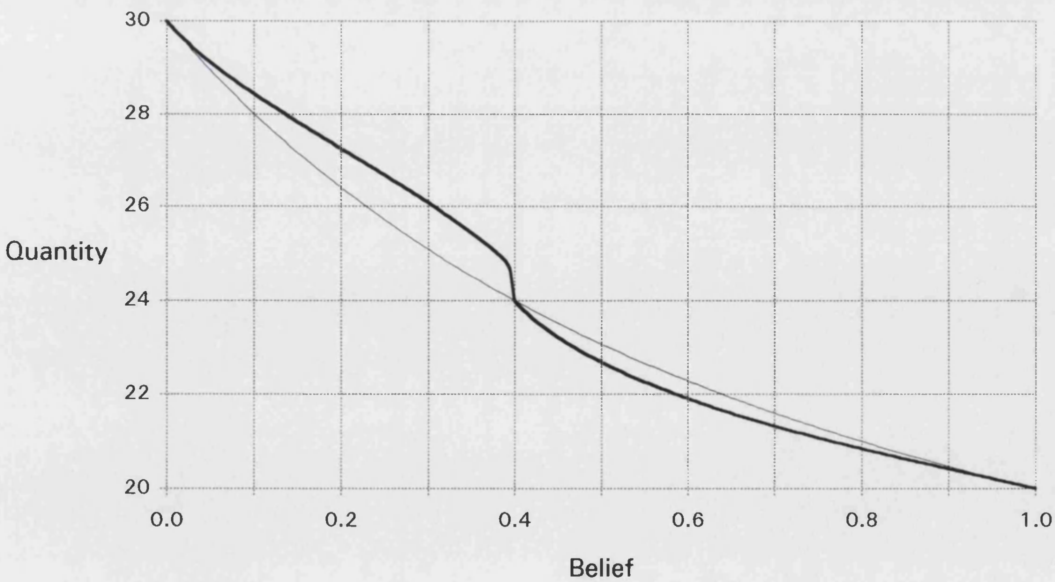

Figure 1.5: Value function and optimal policy for $r = 0.1$, $\lambda_0 = \lambda_1 = 0.075$

1.6.2 Sample Paths

The next set of results are examples of sample paths of beliefs and associated quantities. First we choose an initial state and an initial belief. One iteration then consists of: (a) calculate the optimal quantity given the current belief (using the above numerical results); (b) introduce a shock; (c) update the belief using a slightly different formulation of equation (1.3) in its discrete form, namely

$$\begin{aligned}\delta\pi_t = \lambda(\pi_t)\,\delta t \;&+\; \sigma^{-2}\pi_t(1-\pi_t)(k_t-\pi_t)(\Delta\alpha - \Delta\beta\, q_t)^2\,\delta t \\ &+\; \sigma^{-1}\pi_t(1-\pi_t)(\Delta\alpha - \Delta\beta\, q_t)\,\delta Z_t;\end{aligned}$$

(d) update the state if required (depending on the transition probabilities λ_0 and λ_1). These four steps are then repeated to generate a succession of beliefs and quantities.

It is worthwhile noting the circumstances which make the various terms in the above equation either 0 (or arbitrarily small) or unambiguously non-zero. First, when there is no state switching, the first term vanishes because $\lambda(\pi_t) = 0$, but when there is state switching this term becomes a mean reversion force to $\tilde{\pi}$ (which equals 0.5 when $\lambda_0 = \lambda_1$). Secondly, when there is only moderate experimentation, $q^*(\hat{\pi}) = \hat{q}$ so that for beliefs close to $\hat{\pi}$ the last two terms become arbitrarily small; but in cases of extreme experimentation, $q^*(\hat{\pi}) = q_{\max}$ or $q_{\min}$ so that when beliefs are close to $\hat{\pi}$ these last two terms are quite large. Third, the factor $k_t - \pi_t$ in the middle term means that the agent's belief is pulled towards the truth; and when the belief is close to the truth, the middle term becomes small. Finally, the last two terms are small whenever the belief is close to 0 or 1.

We begin with cases without state switching, first with a high discount rate, implying moderate experimentation, then with a low discount rate, implying extreme experimentation.[40]

No state switching, moderate experimentation. As $\delta\pi$ becomes vanishingly small at $\hat{\pi}$ in this case, π_t cannot cross $\hat{\pi}$. If the initial belief is between the

[40] The upper panel shows the evolution of the agent's belief; the lower panel shows the associated quantity. In regions of moderate experimentation, the graphs of beliefs and of quantities are almost mirror images of each other. When extreme quantities are seen, we enlarge the scale of the vertical axis in the lower panel. Note that in these cases there are large quantity swings whenever the belief is close to the confounding belief.

truth and $\hat{\pi}$, it can converge either to the truth or to $\hat{\pi}$; but if the initial belief is on the "wrong" side of $\hat{\pi}$ compared with the true state, it will converge to $\hat{\pi}$. These possibilities are shown in Figure 1.6, where the true state is $k = 0$. In one sample when $\pi_0 = 0.25$, the agent's belief is drawn towards the truth – after about $t = 20$ there appears to be little adjustment. (However, if the vertical axes were magnified, we would see that updating has not in fact ceased.) In a second sample path with $\pi_0 = 0.25$, we use the same shocks but apply them with the opposite sign. Now, the noise drives the agent's belief towards $\hat{\pi}$ where it becomes trapped, and again there appears to be little movement after about $t = 20$. In the third sample, starting with $\pi_0 = 0.75$, the belief is pulled towards the truth, but is stuck on the "wrong" side of the boundary $\hat{\pi}$, and there is negligible movement soon after about $t = 25$.

No state switching, extreme experimentation. In contrast with the above case, $\delta\pi$ is quite large at $\hat{\pi}$, so π_t can move smoothly through $\hat{\pi}$. Figure 1.7 uses the same true state, $k = 0$, the same initial beliefs, and the same shocks as above, to aid comparison. In the first sample when $\pi_0 = 0.25$, the agent's belief approaches the truth more rapidly – after only about $t = 10$ there seems to be little adjustment. In the second sample when $\pi_0 = 0.25$ (with less benign shocks, when the agent's belief became trapped near $\hat{\pi}$), the noise now simply has the effect that convergence towards the truth is slightly retarded. In the third sample, when $\pi_0 = 0.75$, $\hat{\pi}$ is no longer a boundary, and the agent's belief moves through $\hat{\pi}$ to the "right" side, and, after being driven away again by the noise, settles down and is very close to the truth by about $t = 25$.

When we introduce state switching, beliefs will never settle down close to the truth (or indeed $\hat{\pi}$) because there is always a pull towards $\tilde{\pi} = 0.5$. Nevertheless, with low switching intensity, there is still a strong incentive to experiment, and the state might be unchanged for sufficiently long to exert a significant attraction on the agent's belief. For high switching intensities, on the other hand, both the incentive to experiment and the average length of time between two switches are small, so we cannot expect the belief to track the true state. This is what we see in the final two graphs, where the discount rate is low, the initial state is 0, and the initial belief is 0.25.[41]

Slow state switching, extreme experimentation. For $\lambda_0 = \lambda_1 = 0.025$, Figure 1.8 shows that, by the time of the first state change, the agent's belief has predominantly been between 0 and 0.2. After the state change, the mean reversion force and the attraction to the new true state are initially in the same direction, and as this is a case of high experimentation (which is extreme around $\hat{\pi}$) the agent learns about the change. This pattern is repeated after each occasion the state switches, and the true state is tracked quite well.

Fast state switching, moderate experimentation. With $\lambda_0 = \lambda_1 = 0.075$, we find moderate experimentation. At $\hat{\pi}$, only the mean reversion force is operative, and so, once the agent's belief has moved through $\hat{\pi}$ in the direction of $\tilde{\pi}$, it is stuck on that side (see Figure 1.9). We see very little evidence of the true state being tracked, though a generous viewer may conclude that when the state is 0 the agent's belief is usually between the barrier $\hat{\pi} = 0.4$ and $\tilde{\pi} = 0.5$, whereas when the state is 1 the agent's belief is more often between $\tilde{\pi} = 0.5$ and the truth.

Note the dramatic qualitative change in sample paths. First, with no state switching and a high discount rate, there is a positive probability that the agent will not learn the true state, even when the initial belief is broadly correct. As the discount rate falls through its critical value, extreme experimentation kicks in, and it is generic for the agent's belief to converge to the truth. The introduction of a low rate of state switching implies that the gathering of information does not cease: we see large quantity differences shortly after a change of state, because the agent finds that it is worthwhile to track the true state fairly closely. However, when the switching intensity is increased beyond the critical level (for the prevailing discount rate), experimentation falls away and the path followed by the agent's belief is by and large confined to a narrow band around the long-term mean.

[41]With state switching, we show just one sample path on each panel, the bold dashed line representing the true state. Again the quantity axis is appropriately scaled. Also, the time axis is now extended so that the number of state switches we show is "reasonable" for the parameter values under consideration.

Technical notes

State switching was implemented by repeatedly drawing a number from the uniform distribution on the unit interval. Let $\lambda = \lambda_0 = \lambda_1$ denote the common switching intensity. If the number drawn is less than $1 - \exp(-\lambda)$, then the state remains unchanged, else it switches. Over a time interval of 100, we expect to see 10 switches for $\lambda = 0.1$. For other values of λ, the time interval is "stretched" accordingly, so for $\lambda = 0.025$, for example, we expect these 10 switches to occur by the time $t = 400$.

The shocks were generated by repeated draws from the standard normal distribution.[42] For given time increment δt, the shock δZ was taken to be $\sqrt{\delta t}$ times the draw from the standard normal distribution.

In order to maintain a reasonable approximation to the continuous case that we are modelling, we must ensure that each $\delta\pi$ is not so large that the agent's belief can jump to (or past) 0, 1, or $\hat{\pi}$. To achieve this, the time variable was incremented by 0.05 in each discrete period, i.e. $\delta t = 0.05$. (This means that in the graphs illustrating the cases without state switching there are several hundred iterations, and in those with state switching there are a few thousand.)

[42]The full support of the normal distribution brings with it the usual problem of the possibility of observing a negative price. With the parameters we have been using, the particular danger area is when a quantity near to $q_{\max}$ is optimal in state 1, in which case the expected price is close to 0. The fact that we noted only one such price in the several thousand iterations performed is evidence for the notion that the probability of such an occurrence is extremely small.

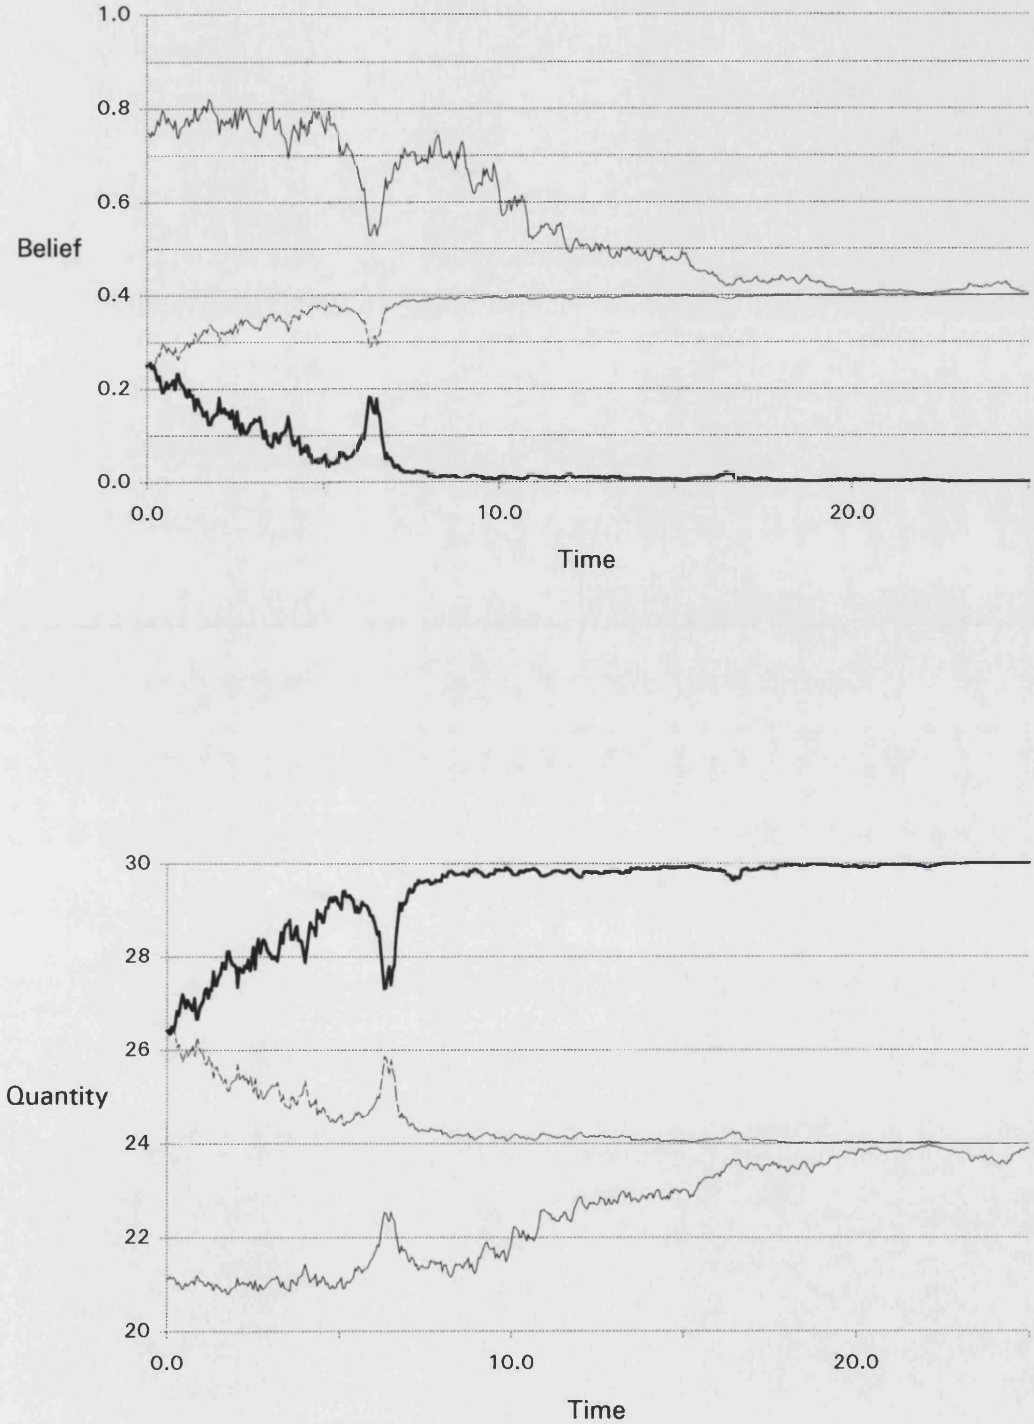

Figure 1.6: Sample paths for $r = 0.5$, $\lambda_0 = \lambda_1 = 0$

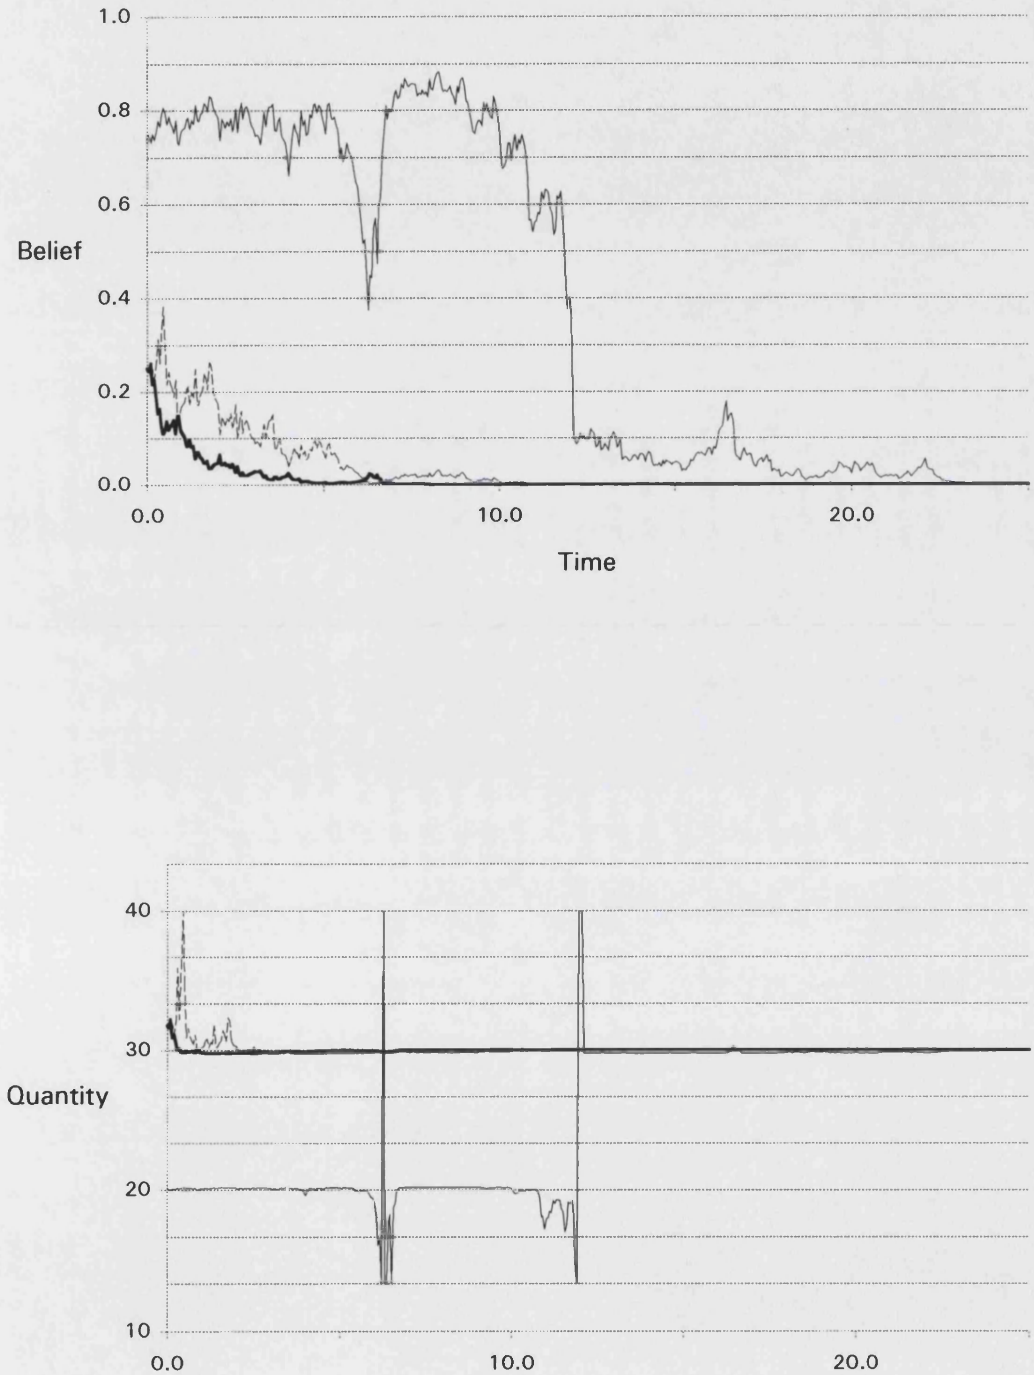

Figure 1.7: Sample paths for $r = 0.1$, $\lambda_0 = \lambda_1 = 0$

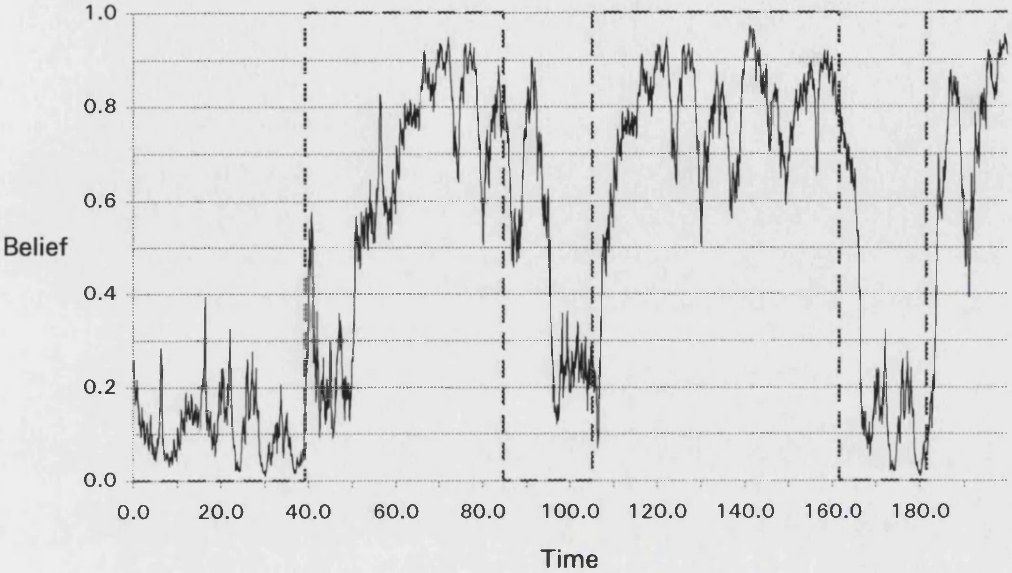

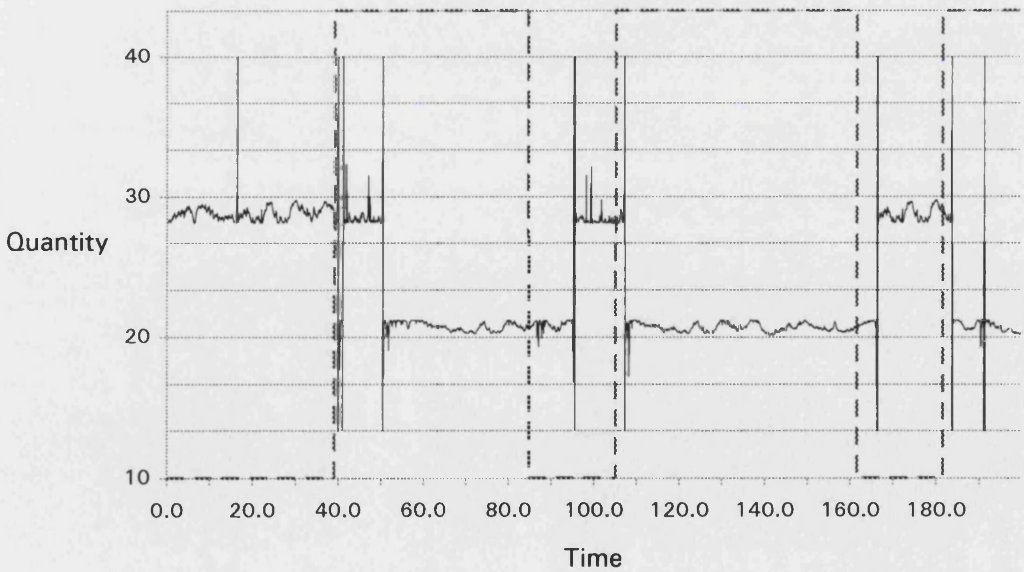

Figure 1.8: Sample paths for $r = 0.1$, $\lambda_0 = \lambda_1 = 0.025$

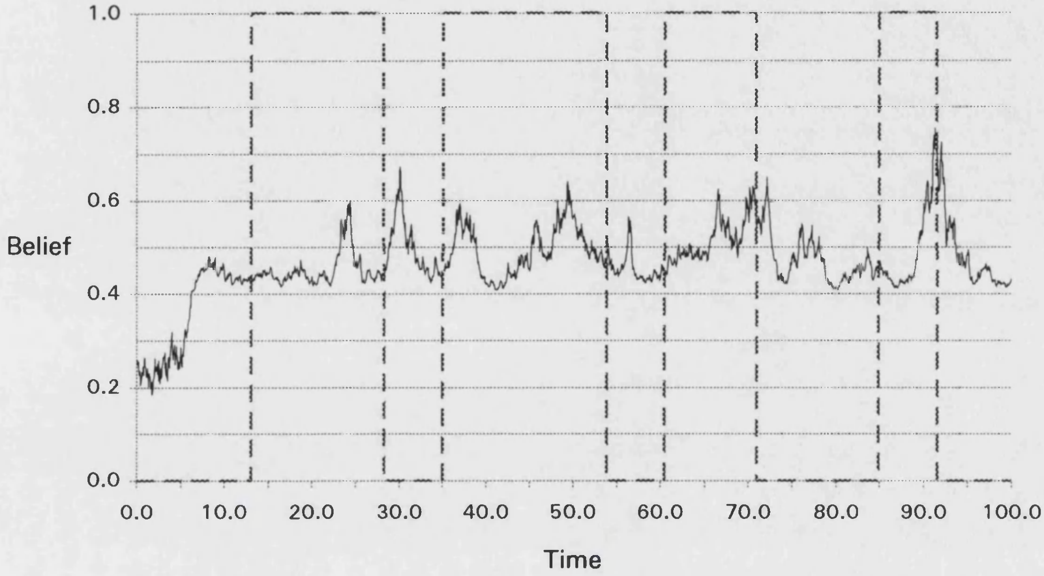

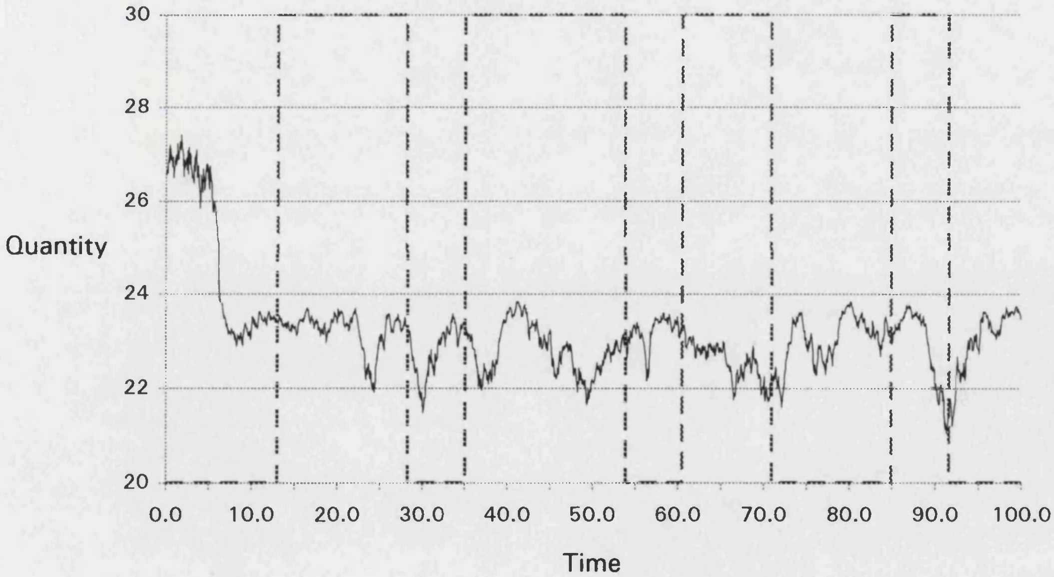

Figure 1.9: Sample paths for $r = 0.1$, $\lambda_0 = \lambda_1 = 0.075$

1.7 Conclusion

We have studied the behaviour of a monopolist facing a demand curve which switches at random, and who receives noisy signals by choosing a stream of quantities and observing the prices they generate. This causes him to update his beliefs about the current demand, and given that the environment is changing, the agent continues to experiment even in the long run.

We formulate the problem in continuous time, which leads us via the Bellman equation to an ordinary differential equation for the value function. The advantages of this approach are two-fold: (a) even though a closed-form solution is generally not obtainable, certain properties of the value function can be established analytically (e.g. convexity and differentiability); furthermore, some comparative statics results can be obtained, even without an explicit solution, allowing us to demonstrate how the value function and optimal policy vary with, say, the discount rate; (b) using numerical methods, it is a far easier task to solve the ODE than it is to determine the fixed point of the Bellman operator which arises in a discrete-time setting.

After discussing the evolution of beliefs and the Bellman equation for our problem, we consider two broad cases. In the case of one-sided experimentation, the agent deviates from myopic behaviour by experimenting towards wider spreads between the demand curves, where the price observations are more informative. This experimentation remains, however, moderate in a well-defined sense and is qualitatively the same for all parameter values.

In the two-sided case (where the demand curves intersect at a feasible quantity and price) there is a confounding quantity $\hat{q}$ and a confounding belief $\hat{\pi}$, and the monopolist's behaviour depends *qualitatively* on the discount rate r and the probability of a change of demand. For high discount rates, the value function touches its myopic counterpart at the confounding belief. The optimal policy moves continuously through the confounding quantity at $\hat{\pi}$ – there is no experimentation at the confounding belief. Moreover, experimentation remains *moderate* overall. For low discount rates and low probabilities of a demand curve switch, on the other hand, the value function is higher than its myopic counterpart everywhere, and the optimal policy involves *extreme* experimentation around $\hat{\pi}$, exhibiting a jump from $q_{\max}$ to $q_{\min}$ in this region. Notably, this discontinuity is "absorbed" by the value function, which is still smooth.[43]

We have some numerical solutions which demonstrate graphically the salient differences in the above scenarios, with regard to the value function and optimal policy. In particular, they show the *discontinuous* change in the optimal behaviour as the discount rate or the switching intensities cross their critical levels. The numerical solutions were then used to construct a variety of sample paths for beliefs and corresponding optimal quantities in order to illustrate how these vary with high and low discount rates and with fast and slow switching between demand curves.

Our main finding, the discontinuous change in the optimal behaviour, is a novel phenomenon in the economic literature on experimentation. It suggests that agents in a changing environment might reduce their information gathering activities dramatically if the rate of change passed a critical threshold. Thus, a slight increase in the variability of the environment could cause a dramatic change in the behaviour of a market or an economy.

[43]Contrast this with the results in Kiefer (1989, 1991) where the value function has a kink for parameter values corresponding to a low discount rate.

Appendix

Existence of a Solution in the One-Sided Case

We first prove the existence of a solution to the differential equation (1.14) as in Theorem 1.4.1. The technique used will also lead to a proof of Proposition 1.4.1 and the special case of Conjecture 1.4.1.

We fix a discount rate r, a noise parameter σ, and transition intensities λ_0 and λ_1. For any function u which is differentiable on the interval $]0,1[$, we define

$$v(\pi) = u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi).$$

Using this substitution, we can write the differential equation (1.14) as $u''(\pi) = G(\pi, v(\pi))$ with the function

$$G(\pi, v) = \frac{2r\sigma^2}{\Delta\beta^2} \frac{\beta(\pi)}{\pi^2(1-\pi)^2} \frac{v - m(\pi)}{v - \hat{m}}.$$

Differentiating $v(\pi)$ twice, each time substituting $G(\pi, v(\pi))$ for $u''(\pi)$, we obtain the ODE

$$v''(\pi) = \left(1 + 2\,\frac{\lambda_0 + \lambda_1}{r}\right) G(\pi, v(\pi)) - \frac{\lambda(\pi)}{r} \frac{d}{d\pi} G(\pi, v(\pi)). \tag{A.1}$$

Thus, if u solves (1.14), then v solves (A.1). Conversely, we shall show that any solution v of (A.1) leads to a solution u of (1.14). This is of course trivial for $\lambda_0 + \lambda_1 = 0$. Therefore, assume $\lambda_0 + \lambda_1 > 0$, and let

$$\rho = \frac{r}{\lambda_0 + \lambda_1},$$

that is, $\rho^{-1} = -\lambda'(\pi)/r$.

Proposition A.1 *Let v be continuous on $[0,1]$ and solve (A.1) on $]0,1[$. Then*

$$u(\pi) = \rho\,|\pi - \tilde{\pi}|^{-\rho}\,\mathrm{sign}(\pi - \tilde{\pi}) \int_{\tilde{\pi}}^{\pi} |\xi - \tilde{\pi}|^{\rho - 1}\, v(\xi)\, d\xi$$

defines a continuously differentiable function on $[0,1]$ such that

$$u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi) = v(\pi).$$

On $]0,1[$, u is analytic and solves (1.14).

PROOF: Note first that v is analytic on $]0,1[$ by the Cauchy-Kowalewski theorem. In particular, as $\rho - 1 > -1$, the integral $\int_{\tilde{\pi}}^{\pi} |\xi - \tilde{\pi}|^{\rho-1} v(\xi)\, d\xi$ exists and defines a continuously differentiable function of π on $[0,1] - \{\tilde{\pi}\}$. It is now straightforward to check that u satisfies the linear differential equation

$$u(\pi) - \frac{\lambda(\pi)}{r} u'(\pi) = v(\pi)$$

on $[0,1] - \{\tilde{\pi}\}$. Applying the Cauchy-Kowalewski theorem once more, we see that u is analytic on $]0,1[-\{\tilde{\pi}\}$. Next, we can use the analyticity of v to show that u is also well defined and analytic at $\tilde{\pi}$ with $u(\tilde{\pi}) = v(\tilde{\pi})$ and derivatives $u^{(k)}(\tilde{\pi}) = \rho v^{(k)}(\tilde{\pi})/(\rho + k)$. In particular, u solves the above linear ODE on the whole of $[0,1]$.[44] Differentiating twice, we obtain

$$(1 + 2\rho^{-1})\, u''(\pi) - \frac{\lambda(\pi)}{r} u'''(\pi) = v''(\pi)$$

on $]0,1[$. Combining this with the fact that v solves (A.1), we see that $w(\pi) = u''(\pi) - G(\pi, v(\pi))$ with G as before is an analytic solution on $]0,1[$ of the homogeneous linear differential equation

$$(1 + 2\rho^{-1})\, w(\pi) - \frac{\lambda(\pi)}{r} w'(\pi) = 0.$$

But the only regular solution of this ODE is $w = 0$, so u solves (1.14). ∎

Our problem is therefore reduced to finding a solution of (A.1) with $m < v < \overline{m}$ on the interior of the unit interval. We shall apply an existence theorem which can be found in Bernfeld and Lakshmikantham (1974). To this end, we first show that m and $\overline{m}$ are a subsolution and a supersolution[45] of (A.1), respectively.

Lemma A.1 *The myopic payoff function m is a strict subsolution of* (A.1) *on* $]0,1[$.

[44] In fact, u is the only such solution. The standard method of variation of constants shows that the general solution of this ODE on either of the subintervals $[0,\tilde{\pi}[$ or $]\tilde{\pi},1]$ is $u(\pi) + C|\pi - \tilde{\pi}|^{-\rho}$ for some constant C. Boundedness as π tends to $\tilde{\pi}$ requires $C = 0$. Of course, it was by this method that we obtained the representation of u in the first place.

[45] For our purposes, these concepts can be defined as follows. Consider an ordinary second-order differential equation $y'' = f(x,y,y')$. A *subsolution (supersolution)* of this ODE is a function y of class C^2 such that $y'' \geq f(x,y,y')$ ($y'' \leq f(x,y,y')$). We speak of a *strict* subsolution (supersolution) if $y'' > f(x,y,y')$ ($y'' < f(x,y,y')$).

PROOF: We have $G(\pi, m(\pi)) = 0$ on $]0,1[$. On the other hand, $m'' > 0$. ■

Lemma A.2 *The linear function $\overline{m}$ is a strict supersolution of* (A.1) *on* $]0,1[$.

PROOF: $\overline{m}'' = 0$, so we have to show that the right hand side of (A.1) with v replaced by $\overline{m}$ is positive. It will be convenient to rewrite the ODE (A.1) in a more explicit form. To this end, set

$$H(\pi, v) = \frac{v - m(\pi)}{v - \hat{m}} = 1 - \frac{m(\pi) - \hat{m}}{v - \hat{m}}.$$

Note that $G(\pi, v) = 2r\sigma^2\, \beta(\pi)\, H(\pi, v)/(\Delta\beta^2\, \pi^2\, (1-\pi)^2)$. We make this substitution in (A.1) and differentiate. Collecting terms in r, λ_0 and λ_1, and simplifying, we can finally rewrite (A.1) as

$$\begin{aligned} v''(\pi) = \frac{2\sigma^2}{\Delta\beta^2}\, \frac{1}{\pi^2(1-\pi)^2} \Bigg\{ & r\, \beta(\pi) H(\pi, v(\pi)) \\ & + \lambda_0\, (1-\pi) \left[\frac{\beta_0 + \beta(\pi)}{\pi}\, H(\pi, v(\pi)) - \beta(\pi)\, \frac{d}{d\pi} H(\pi, v(\pi)) \right] \\ & + \lambda_1\, \pi \left[\frac{\beta_1 + \beta(\pi)}{1-\pi}\, H(\pi, v(\pi)) + \beta(\pi)\, \frac{d}{d\pi} H(\pi, v(\pi)) \right] \Bigg\}. \end{aligned} \tag{A.2}$$

When we replace $v(\pi)$ by $\overline{m}(\pi) = (1-\pi)m(0) + \pi m(1)$ on the right hand side of this equation, the coefficient of r is clearly always positive on $]0,1[$. The expressions in square brackets associated with λ_0 and λ_1 simplify to $f_0(\pi)/(\overline{m}(\pi) - \hat{m})^2$ and $f_1(\pi)/(\overline{m}(\pi) - \hat{m})^2$ respectively, where f_0 and f_1 are quadratics in π:

$$\begin{aligned} f_0(\pi) &= K \left[m(1) - \hat{m} + [m(0) - m(1)]\, (1-\pi)^2 \right], \\ f_1(\pi) &= K \left[m(0) - \hat{m} + [m(1) - m(0)]\, \pi^2 \right] \end{aligned}$$

with $K = (\Delta\beta\, \hat{p})^2/(4\beta_0\beta_1)$. Thus, $f_0(0) = f_1(0) = K\,[m(0) - \hat{m}]$ and $f_0(1) = f_1(1) = K\,[m(1) - \hat{m}]$, hence f_0 and f_1 are both non-negative at each end of the unit interval. As the two quadratics are strictly monotonic on $[0,1]$, they are both non-negative over the entire unit interval. ■

The existence theorem we want to apply relies on an *a priori* bound for the right hand side of the ODE (A.1). The following lemma provides this bound. To state the result in the most convenient way, we write (A.1) as

$$v''(\pi) = \frac{1}{\pi^2\, (1-\pi)^2}\, F[\pi, v(\pi), v'(\pi)]. \tag{A.3}$$

The function F can be obtained directly from (A.2):

$$
\begin{aligned}
F[\pi, v_0, v_1] = \frac{2\sigma^2}{\Delta\beta^2} \Bigg\{ & r\,\beta(\pi) H(\pi, v_0) \\
& + \lambda_0\,(1-\pi) \left[\frac{\beta_0 + \beta(\pi)}{\pi}\, H(\pi, v_0) - \beta(\pi)\, H_1[\pi, v_0, v_1] \right] \\
& + \lambda_1\,\pi \left[\frac{\beta_1 + \beta(\pi)}{1-\pi}\, H(\pi, v_0) + \beta(\pi)\, H_1[\pi, v_0, v_1] \right] \Bigg\}
\end{aligned}
$$

with

$$
H_1[\pi, v_0, v_1] \;=\; \frac{m(\pi) - \hat{m}}{(v_0 - \hat{m})^2}\, v_1 \;-\; \frac{m'(\pi)}{v_0 - \hat{m}} \,.
$$

Lemma A.3 *Let $J \subset\,]0,1[$ be a closed interval. Then there is a constant C_J depending only on J such that $|F[\pi, v_0, v_1]| \leq C_J\,(1 + |v_1|)$ for all $\pi \in J$, $\underline{m}(\pi) \leq v_0 \leq \overline{m}(\pi)$ and $v_1 \in I\!R$.*

PROOF: All terms involving H are clearly bounded on J. The terms involving H_1, on the other hand, are bounded in absolute value by $c_0 + c_1|v_1|$ for some constants c_0 and c_1. ∎

Proposition A.2 *There exists a solution v of* (A.1) *on* $]0,1[$ *with $\underline{m} < v < \overline{m}$.*

PROOF: We shall apply Theorem 1.7.2 of Bernfeld and Lakshmikantham (1974) which is formulated for problems on the real line $I\!R$ rather than on a finite interval. We therefore make the change of variables

$$
t = \log\frac{\pi}{1-\pi} \qquad \text{or} \qquad \pi = \frac{1}{1+e^{-t}}
$$

and set $x(t) = v(\pi)$. Then

$$
\begin{aligned}
\frac{dt}{d\pi} &= \frac{1}{\pi\,(1-\pi)}\,, \\
\frac{d^2t}{d\pi^2} &= \frac{2\pi - 1}{\pi^2\,(1-\pi)^2}\,,
\end{aligned}
$$

hence

$$
\begin{aligned}
v'(\pi) &= \frac{1}{\pi\,(1-\pi)}\, x'(t)\,, \\
v''(\pi) &= \frac{1}{\pi^2\,(1-\pi)^2}\, x''(t) + \frac{2\pi - 1}{\pi^2\,(1-\pi)^2}\, x'(t)\,.
\end{aligned}
$$

Noting that

$$\frac{1}{\pi\,(1-\pi)} \;=\; (1+e^{t})\,(1+e^{-t})$$

and

$$2\pi - 1 \;=\; \frac{1-e^{-t}}{1+e^{-t}}\,,$$

we see that (A.3) transforms into

$$x''(t) \;=\; \frac{1-e^{t}}{1+e^{t}}\,x'(t) \;+\; F\left[\frac{1}{1+e^{-t}}\,,\; x(t)\,,\; (1+e^{t})\,(1+e^{-t})\,x'(t)\right]. \qquad \text{(A.4)}$$

Of course, m and $\overline{m}$ transform into a subsolution $\underline{x}$ and a supersolution $\overline{x}$ of (A.4), respectively. Now consider any interval $[-a,a]$. Using the last lemma above, we can find a constant C_a depending only on a such that

$$\left|\frac{1-e^{t}}{1+e^{t}}\,x_1 \;+\; F\left[\frac{1}{1+e^{-t}}\,,\; x_0\,,\; (1+e^{t})\,(1+e^{-t})\,x_1\right]\right| \le C_a\,(1\;+|x_1|)$$

for all $t \in [-a,a]$, $\underline{x}(t) \le x_0 \le \overline{x}(t)$ and $x_1 \in I\!R$. By Bernfeld and Lakshmikantham (1974, Theorem 1.7.2, p.45), we can now conclude that (A.4) possesses a solution x defined on the entire real line with $\underline{x} \le x \le \overline{x}$. Changing variables back from t to π, we obtain the existence of a solution v to (A.1) on $]0,1[$ satisfying the weak inequalities $m \le v \le \overline{m}$. To show that these inequalities hold actually in the strict sense, we use the fact that m and $\overline{m}$ are a strict subsolution and supersolution, respectively. Assume for example that there is a belief $\breve{\pi} \in\,]0,1[$ such that $v(\breve{\pi}) = m(\breve{\pi})$. Then the function $v - m$ has a local minimum at $\breve{\pi}$, so $v'(\breve{\pi}) = m'(\breve{\pi})$ and $v''(\breve{\pi}) \ge m''(\breve{\pi})$. Yet

$$v''(\breve{\pi}) \;=\; \frac{1}{\breve{\pi}^2\,(1-\breve{\pi})^2}\,F[\breve{\pi}, v(\breve{\pi}), v'(\breve{\pi})] \;=\; \frac{1}{\breve{\pi}^2\,(1-\breve{\pi})^2}\,F[\breve{\pi}, m(\breve{\pi}), m'(\breve{\pi})] \;<\; m''(\breve{\pi})$$

— a contradiction. The strict inequality $v < \overline{m}$ can be shown in the same way. ■

We can finally state the result which was the starting point in the proof of Theorem 1.4.1.

Proposition A.3 *There exists a function u on $[0,1]$ with the following properties: u is strictly convex and once continuously differentiable on $[0,1]$; on $]0,1[$, u is analytic and solves* (1.14) *with*

$$m(\pi) < u(\pi) - \frac{\lambda(\pi)}{r}\,u'(\pi) < \overline{m}(\pi).$$

PROOF: The function v of Proposition A.2 can obviously be extended to a continuous function on the closed unit interval. By Proposition A.1, there is a solution u of (1.14) with the stated regularity properties. Finally, $v > m$ implies $u'' > 0$ by (1.14), so u is strictly convex. ∎

Using the same techniques, it is now relatively easy to prove the result underlying Proposition 1.4.1.

Proposition A.4 *Fix σ, λ_0 and λ_1, and consider two interest rates $r_1 < r_2$. Let $u[r_1]$ and $u[r_2]$ be the corresponding value functions, and define the respective functions $v[r_1]$ and $v[r_2]$ as above. Then $v[r_1] > v[r_2]$ on $]0,1[$.*

PROOF: The representations (A.2) or (A.3) of our ODE (A.1) show that the coefficient of r on the right hand side is positive whenever $v(\pi) > m(\pi)$. The function $v[r_1]$, which solves (A.1) for the interest rate r_1, is therefore a strict supersolution on $]0,1[$ of the corresponding ODE for r_2. Proceeding as in the proof of Proposition A.2, we obtain a solution v of the ODE for r_2 such that $m < v < v[r_1]$ on the open unit interval. Now, the uniqueness part of Theorem 1.4.1 implies $v = v[r_2]$. ∎

Along these lines, we can also verify Conjecture 1.4.1 for the case without state transitions.

Proposition A.5 *Let r be given, and assume $\lambda_0 = \lambda_1 = 0$. Consider two noise parameters $\sigma_1 < \sigma_2$ and the corresponding value functions $u[\sigma_1]$ and $u[\sigma_2]$. Then $u[\sigma_1] > u[\sigma_2]$ on $]0,1[$.*

PROOF: The value function $u[\sigma_1]$, which solves (1.14) for σ_1, is a strict supersolution on $]0,1[$ of the corresponding ODE for σ_2. ∎

Finally, let us mention sufficient conditions under which Conjectures 1.4.1 and 1.4.2 would hold in full generality. If the solution v from Proposition A.2 is strictly convex, one can prove the result regarding a change in σ in the same way as we proved the comparative statics for r. Moreover, if $v'' > u''$ on $]0,1[$, then the result for a change in $\lambda_0 + \lambda_1$ follows by the same technique.[46]

[46] This last condition is satisfied in all the numerical examples that we have computed. We do not show graphs for the one-sided case here, but v'' obviously exceeds u'' in our figures for the two-sided case.

Part II

Derivative Asset Pricing

Chapter 2

The Direct Approach to Debt Option Pricing[1]

The valuation of debt options has occupied a central place in the literature on contingent claim pricing and the term structure of interest rates. While most commonly traded debt options are written on coupon bonds, there has also been considerable interest in the valuation of European options on zero-coupon bonds. On the one hand, this is motivated by the fact that discount bond options provide simple building blocks for other more complex interest rate dependent claims.[2] On the other hand, zero-coupon bonds very closely resemble stocks which pay no dividends; the fact that the bond price has to reach par value eventually will only have an impact close to maturity. This makes the valuation of options on discount bonds one of the few areas in derivative assets analysis where one can expect closed-form solutions that are as tractable and elegant as the famous Black and Scholes (1973) formula for stock options.

In this chapter, we review the continuous-time literature on the so-called direct approach to debt option pricing. Trying to follow the Black-Scholes framework as closely as possible, this approach specifies bond prices directly, without relating them to the term structure as a whole or to state variables such as the short term

[1]The material in this chapter has been published in the *Review of Futures Markets*; see Rady and Sandmann (1994).

[2]It is well known, for instance, that FRAs (forward rate agreements) can be re-interpreted as options on zero-coupon bonds, and caps and floors as strings of such options; see for instance Sandmann (1991) or Briys, Crouhy and Schöbel (1991). In some circumstances, it is possible to write an option on a coupon bond as a sum of options on discount bonds; cf. Jamshidian (1987, 1989) and El Karoui and Rochet (1989).

interest rate. This literature starts with Ball and Torous (1983) where the stock option pricing model of Merton (1973), an extension of Black and Scholes' work to stochastic interest rates, is adapted to debt options. The main contribution of Ball and Torous consists in replacing the Brownian motion which drives the Black-Scholes and Merton stock price model by a Brownian bridge process. Thus, they succeed in modelling the main difference between stocks and bonds: under absence of default risk, bonds reach a predetermined face value at their maturity whereas stocks have no such target value. However, the constant volatility[3] of bond prices in the Ball-Torous model turns out to imply a highly implausible bond yield behaviour. By introducing bond price processes with time dependent volatility into the Merton framework, Kemna, de Munnik and Vorst (1989) are able to avoid this problem.

The two models mentioned so far both specify bond prices as lognormal variables. While this leads to closed-form option price formulae of the same type as in Merton (1973), it also means that negative bond yields and negative forward yields occur with positive probability. This problem has been addressed by Schöbel (1986). He derives boundary conditions for discount bond options under the assumption that yields do not become negative. He then proposes a method to modify option price formulae like that of Ball and Torous (1983) in accordance with these boundary conditions. Yet Schöbel does not develop a bond price model in which yields would indeed remain non-negative. Bühler and Käsler (1989) are the first to achieve this within the direct approach. With a very ingenious formulation of bond prices, their model guarantees positive bond yields as well as positive forward yields, and still has the advantage of providing analytic solutions for option prices.

While the above papers deal exclusively with discount bonds, Schaefer and Schwartz (1987) and Bühler (1988) use the direct approach to price options on coupon bonds. Both papers let the volatility of the underlying bond depend on the bond's duration. In such a setting, bond option prices must be calculated numerically. Unfortunately, both papers have to make extremely restrictive assumptions

[3]Practitioners as well as academic researchers have used the term "volatility" to denote various quantities that measure the riskiness of an asset. We adopt the following convention: "volatility" is synonymous with "instantaneous standard deviation of returns".

in order to keep the numerical complexity of the valuation problem at a reasonable level.

The preceding paragraphs have already mentioned two of the main modelling problems encountered by the direct approach: first, the problem of specifying bond price processes that reach par value at maturity with probability one; second, the problem of modelling bond prices in a way that precludes negative yields. A third problem has to do with the internal consistency of models: bond price processes must be specified such that no arbitrage opportunities between the bonds arise. Under suitable restrictions on the trading strategies that investors can use, the existence of a so-called martingale measure is a sufficient condition for the absence of arbitrage opportunities.[4] Cheng (1991) shows that there is no such measure for the Ball-Torous model. Reacting to Cheng's work, de Munnik (1990) proves the existence of a martingale measure for the model of Kemna, de Munnik and Vorst (1989). Bühler and Käsler (1989) provide the most elegant solution. While de Munnik's arguments are technically rather intricate, Bühler and Käsler are able to invoke a general result that immediately implies the existence of a martingale measure for their model.

The aim of this chapter is to emphasise the above modelling problems and to discuss in detail the different solutions proposed in the literature. It is organised as follows. Section 2.1 gives a short introduction to the principal features of discount bonds and discount bond options. Section 2.2 presents the portfolio duplication argument that underpins derivative asset pricing and sets out the general framework of the direct approach. Sections 2.3 and 2.4 discuss the lognormal models of Ball and Torous (1983) and Kemna, de Munnik and Vorst (1989), respectively. In Section 2.5, we analyse the modified pricing formulae proposed by Schöbel (1986). The model of Bühler and Käsler (1989) is presented in Section 2.6. In Section 2.7, we briefly discuss the valuation models for options on coupon bonds by Schaefer and Schwartz (1987) and Bühler (1988). Section 2.8 contains concluding remarks. Some technical details are given in an appendix.

[4] See for example Harrison and Pliska (1981) or Müller (1985).

2.1 Discount Bonds and European Options

A *zero-coupon* or *discount bond* is a nominal security that pays its owner a predetermined amount of money, the face value, at a predetermined maturity date. We consider only bonds without any default risk, such as treasury bills. Moreover, face values are normalised to 1 without loss of generality. Writing $B_{t,T}$ for the time t price of a zero-coupon bond which matures at $T \geq t$, we thus have the terminal value condition

$$B_{T,T} = 1 \tag{2.1}$$

for all maturities T.

Bond price models can be classified according to whether they generate negative interest rates or not.[5] By definition, the *yield to maturity* $Y_{t,T}$ satisfies

$$B_{t,T} = \exp(-(T-t)\, Y_{t,T})$$

for $t < T$. In other words, $Y_{t,T}$ is the continuously compounded interest rate at time t for a loan repayable at T. The *forward yield* Y_{t,T_1,T_2} is defined by

$$\frac{B_{t,T_2}}{B_{t,T_1}} = \exp(-(T_2 - T_1)\, Y_{t,T_1,T_2})$$

for $t \leq T_1 < T_2$. It is the interest rate as seen at time t for a loan starting at T_1 and repayable at T_2.

Suppose that investors can hold cash.[6] If they are rational and prefer more to less, they will not accept to lend (or lend forward) at a negative interest rate, so bond prices must satisfy

$$B_{t,T} \leq 1 \qquad \text{for all } t < T \tag{2.2}$$

and

$$B_{t,T_2} \leq B_{t,T_1} \qquad \text{for all } t \leq T_1 < T_2 \ . \tag{2.3}$$

In the following, a model that violates (2.2) or (2.3) with positive probability will be said to *generate negative yields.*

[5] We use the terms "negative" and "positive" in the strict sense, meaning "< 0" and "> 0", respectively. A quantity satisfying the weak inequality "≥ 0" is called "non-negative", etc.

[6] This is certainly a reasonable assumption in most practical applications.

A *European call option* on a zero-coupon bond of maturity T is the right to buy the bond at some specified date $\tau < T$ for some predetermined amount K. If the price of the bond at the exercise date τ is higher than the exercise price K, the net cash flow of the call will be the difference $B_{\tau,T} - K$, otherwise the net cash flow is zero. Therefore, the call is worth

$$[B_{\tau,T} - K]^+$$

at τ.[7] A *European put option* is the right to sell a bond for some predetermined amount K. Its net cash flow at the exercise date is

$$[K - B_{\tau,T}]^+.$$

Obviously, the final payoffs of these two options are related via

$$[K - B_{\tau,T}]^+ = [B_{\tau,T} - K]^+ - B_{\tau,T} + K.$$

By a simple arbitrage argument, this implies the so-called *put-call parity* between time t prices of European call and put options with common exercise price K and common exercise date τ:[8]

$$P_t = C_t - B_{t,T} + K B_{t,\tau}.$$

As the value of a put option is always non-negative, put-call parity yields the following lower bound on the call price:

$$C_t \geq [B_{t,T} - K\, B_{t,\tau}]^+. \tag{2.4}$$

An upper bound for the price of a call is

$$C_t \leq B_{t,T} \tag{2.5}$$

since the call cannot be worth more than the underlying security itself.[9]

An additional upper bound holds when there are no negative yields. In this case only exercise prices K between 0 and 1 are of interest, and the maximal payoff

[7] By definition, $[x]^+ = \max\{x, 0\}$ for $x \in I\!R$.

[8] See for example Stoll (1968).

[9] Conditions (2.4) and (2.5) were derived by Merton (1973). For the role of condition (2.5) in option pricing see also Gleit (1978).

of a call is $1 - K$. The call price is therefore bounded from above by the present value of $1 - K$:

$$C_t \leq (1 - K)\, B_{t,\tau}. \tag{2.6}$$

Combining (2.4) and (2.6), Schöbel (1986) obtains the following condition:

$$C_t = (1 - K)\, B_{t,\tau} \quad \text{whenever} \quad B_{t,T} = B_{t,\tau}. \tag{2.7}$$

Put-call parity leads to a similar result for put options.

So far, no assumptions have been made on the stochastic behaviour of bond prices. But it is already obvious that the price of an option will not only depend on its underlying bond, but also on the price of a zero-coupon bond whose maturity coincides with the exercise date of the option. The direct approach to bond option pricing studies models in which these two bonds suffice to "span" the option and hence to determine its price.

2.2 Option Pricing by Portfolio Duplication

This section presents the standard portfolio duplication argument which is the basis of derivative asset pricing. We shall focus on the case of a European call option written on a discount bond with face value 1 and maturity T. The option is assumed to have exercise date $\tau < T$ and strike price K. For $0 \leq t \leq T$, let B_t denote the time t price of the bond on which the option is written (the "underlying bond"). The price of the "reference bond", a zero-coupon bond of maturity τ, is denoted by R_t, $0 \leq t \leq \tau$. The value of the option will depend on the properties of the stochastic processes B and R. For the moment, we only assume that they are continuous Itô processes.[10] The main idea is to construct a dynamically adjusted portfolio in the two bonds that yields the same final cash flow as the option. To make this more precise, we need some definitions:[11]

[10] An introduction to the theory of such processes and their use in financial models can be found in Duffie (1992). For the sake of simplicity, technical requirements such as integrability conditions will not be made explicit here.

[11] In this chapter, we shall not give a precise definition of a space of admissible portfolio strategies. However, the strategies we shall deal with can be checked to have the relevant properties; see Duffie (1992).

A *portfolio strategy* is a two-dimensional predictable stochastic process $\theta = (\theta^0, \theta^1)$ on the time interval $[0, \tau]$ such that the stochastic integrals $\int \theta^0 \, dR$ and $\int \theta^1 \, dB$ exist. Think of θ^0_t and θ^1_t as the number of reference and underlying bonds held at time t. Predictability means that the decision how many bonds to hold at t is based only on information available before t. The stochastic integrals above can be interpreted as the gains or losses from bond trade according to the strategy θ.

The *value process* V^θ of a strategy θ is given by

$$V^\theta_t = \theta^0_t R_t + \theta^1_t B_t.$$

A strategy θ is called *self-financing* if V^θ has the stochastic differential

$$dV^\theta_t = \theta^0_t \, dR_t + \theta^1_t \, dB_t \, .$$

This means that after the initial investment V^θ_0 is made, the adjustment of the portfolio is financed without injecting or withdrawing any funds. Changes in the portfolio value are exclusively due to gains or losses from bond trade.

We say a self-financing strategy *generates the option* if the terminal portfolio value equals the cash flow of the option, i.e.,

$$V^\theta_\tau = [B_\tau - K]^+,$$

and V^θ_t respects at any time the lower and upper bounds derived in Section 2.1, that is, either

$$[B_t - K\, R_t]^+ \leq V^\theta_t \leq B_t \tag{2.8}$$

or, if the bond price model precludes negative yields,

$$[B_t - K\, R_t]^+ \leq V^\theta_t \leq \min\{B_t, (1-K)R_t\}. \tag{2.9}$$

To rule out arbitrage opportunities,[12] the option price must then coincide with the portfolio value, i.e.,

$$C_t = V^\theta_t$$

[12]Formally, an arbitrage opportunity can be defined as a self-financing portfolio strategy with negative initial investment, but non-negative final value; see for instance Duffie (1992). Thus, an arbitrage opportunity is a trading strategy that provides a gain today without creating any liabilities in the future.

for all $t \leq \tau$. This determines the *arbitrage price* of the option.

The construction of a generating strategy can be simplified in the following way. Instead of the bond price process (R, B), we consider the normalised process $(1, \hat{B})$ where

$$\hat{B}_t = \frac{B_t}{R_t}$$

(we assume that R^{-1} is also an Itô process). In the same way, we obtain the normalised value process $\hat{V}^\theta$ of a strategy θ:

$$\hat{V}_t^\theta = \frac{V_t^\theta}{R_t} = \theta_t^0 + \theta_t^1 \hat{B}_t .$$

These definitions may seem purely formal, but they are easily interpreted in economic terms. (R, B) is a model of the *spot* markets, so a portfolio strategy θ describes spot trading with corresponding portfolio value V^θ. Suppose now that investors can also trade on *forward* markets. Then $\hat{B}_t$ is just the time t forward price of the underlying bond for delivery at τ (obviously, the corresponding forward price of the reference bond is always 1). Moreover, if the strategy θ is implemented on the forward markets, the resulting forward value process is just $\hat{V}^\theta$.

We call θ *self-financing on the forward markets* if

$$d\hat{V}_t^\theta = \theta_t^1 \, d\hat{B}_t,$$

and we say such a self-financing strategy θ *generates the option on the forward markets* if

$$\hat{V}_\tau^\theta = [\hat{B}_\tau - K]^+$$

and $\hat{V}_t^\theta$ respects the bounds resulting from division of (2.8) and (2.9) by R_t. The strategy θ then determines the arbitrage forward price of the option:

$$\hat{C}_t = \hat{V}_t^\theta$$

for $t \leq \tau$.

Intuitively, the renormalisation of prices should have no economic effects. The following lemma confirms this intuition.

Lemma 2.2.1 *A portfolio strategy is self-financing on the spot markets if and only if it is self-financing on the forward markets. Furthermore, a strategy generates the* *option on the spot markets if and only if it generates the option on the forward markets.*

PROOF: The proof of the first part consists essentially in an application of Itô's formula and is given in Müller (1985) for a more general framework. The second part then follows trivially. ∎

In order to make further progress, we have to specify bond price processes more precisely. Aiming to construct a generating strategy on the forward markets, we shall in fact start from an explicit description of the forward price process $\hat{B}$. We assume that this process satisfies

$$d\hat{B}_t = \alpha_t \, \hat{B}_t \, dt + \nu(t, \hat{B}_t) \, \hat{B}_t \, dW_t \tag{2.10}$$

where α is some stochastic process, $\nu(t, x)$ a continuous function and W a standard Wiener process. We call α the *drift rate process* and ν the *volatility function* of the forward bond price, interpreting them as the instantaneous expectation and standard deviation, respectively, of the infinitesimal rate of return $d\hat{B}/\hat{B}$. Thus, (2.10) restricts the volatility of the forward bond to be a deterministic function of the current forward bond price and time. This restriction, which rules out more complicated dependence of the forward bond volatility on current or past bond prices, will enable us to determine the arbitrage forward price of the option as a deterministic function of t and $\hat{B}_t$.[13]

Our second lemma shows how to construct generating strategies. Let the interval I denote the state space of the forward price process $\hat{B}$, and $\overline{I}$ its closure. We assume that either $I =]0, \infty[$ or $I =]0, 1[$.[14] In view of Lemma 2.2.1, we do not specify the markets where the strategy is implemented.

Lemma 2.2.2 *Let $u(t, x)$ be continuous on $[0, \tau] \times \overline{I}$ and a solution of the partial differential equation*

$$u_t + \tfrac{1}{2} \, \nu^2 \, x^2 \, u_{xx} = 0 \tag{2.11}$$

[13]See Jamshidian (1990) for a formulation of this result in a term structure model.

[14]This covers all the models we shall deal with except Schöbel (1986).

on $[0,\tau[\times I$. *Then the strategy* θ *defined by*

$$\theta_t^1 = u_x(t,\hat{B}_t), \quad \theta_t^0 = u(t,\hat{B}_t) - u_x(t,\hat{B}_t)\,\hat{B}_t \tag{2.12}$$

is self-financing.

Moreover, suppose that u *has the terminal value* $u(\tau,x) = [x-K]^+$ *and satisfies*

$$[x-K]^+ \leq u(t,x) \leq x \qquad \textit{if } I =]0,\infty[$$

or

$$[x-K]^+ \leq u(t,x) \leq \min\{x, 1-K\} \qquad \textit{if } I =]0,1[\;.$$

Then θ *generates the call option.*

PROOF: (2.12) implies $\hat{V}_t^\theta = u(t,\hat{B}_t)$. By Itô's lemma and (2.10),

$$d\hat{V}_t^\theta = \left[u_t(t,\hat{B}_t) + \tfrac{1}{2}\,\nu^2(t,\hat{B}_t)\,\hat{B}_t^2\,u_{xx}(t,\hat{B}_t)\right]\,dt + u_x(t,\hat{B}_t)\,d\hat{B}_t.$$

By (2.11) and (2.12), this reduces to $d\hat{V}_t^\theta = \theta_t^1\,d\hat{B}_t$, so θ is self-financing. The rest is easy to check. ■

A generating strategy as in Lemma 2.2.2 yields the arbitrage forward price

$$\hat{C}_t = u(t,\hat{B}_t)$$

and the arbitrage spot price

$$C_t = R_t\,\hat{C}_t = R_t\,u(t,\hat{B}_t)$$

for the European call. In accordance with Merton's (1973) theory of rational option pricing, the spot price is homogeneous of degree one in B_t and R_t.

Note that only the volatility function ν appears in the partial differential equation (2.11). The drift rate α in (2.10) has therefore no effect on the functional relationship between the arbitrage price of the option and the bond prices B and R.

It would, however, be wrong to conclude that the drift is completely irrelevant for option pricing. When deriving the above option price, we simply *postulated* absence of arbitrage opportunities between traded securities. The drift of the forward bond price emerges as an important factor when we start to look for conditions that *guarantee* the internal consistency of the bond price model (R, B).[15] Together with suitable restrictions on admissible portfolio strategies, the existence of a so-called *martingale measure* for the forward bond price is a sufficient condition for the absence of arbitrage opportunities. This is a new probability measure that has the same null sets as the original measure and makes the forward price a martingale, which means that the current forward price is an unbiased forecast of future forward prices. Under such a measure, the forward value process of a self-financing portfolio strategy is again a martingale. In particular, the initial investment required to set up a self-financing portfolio strategy equals the expectation of the strategy's terminal value under the martingale measure. As taking expectations preserves non-negativity, a trading strategy with non-negative final value must have a non-negative initial investment. In other words, if there exists a martingale measure, arbitrage opportunities are precluded.

In the setting described by equation (2.10), a martingale measure exists if and only if the quotient of the drift and the volatility of $\hat{B}$,

$$\frac{\alpha_t}{\nu(t, \hat{B}_t)}$$

satisfies certain integrability conditions.[16] Thus, the internal consistency of a bond price model depends indeed on both the drift and the volatility of the forward bond price.

There is a second important reason why the drift term matters in option pricing. When applying an option pricing model, we need estimates for the volatility parameters which enter the valuation formula. It is in general impossible to estimate these parameters from historical price data without taking into account the drift.[17]

Let us conclude this section with an example of how the above lemmata are

[15] In the following, we try to convey only the main ideas. For a thorough discussion including technical details, see for instance Duffie (1992).

[16] See Harrison and Kreps (1979), Harrison and Pliska (1981) and Müller (1985).

[17] A treatment of this estimation problem is beyond the scope of this survey. We therefore refer the reader to Lo (1986, 1988) and references given there. De Munnik (1992) applies Lo's methodology to the model of Kemna, de Munnik and Vorst (1989). Practitioners often use an "implied volatility approach" to avoid the estimation problem altogether; inverting the option price formula, they calculate volatility parameters from observed option prices.

applied. Consider bond price processes that have the stochastic differentials

$$dB_t = \alpha_t^B B_t\, dt + \sigma_B(t)\, B_t\, dW_t^B, \tag{2.13}$$

$$dR_t = \alpha_t^R R_t\, dt + \sigma_R(t)\, R_t\, dW_t^R \tag{2.14}$$

with stochastic drift rate processes α^B, α^R, and volatility functions σ_B, σ_R whicht depend only on t. W^B and W^R are assumed to be Wiener processes having infinitesimal correlation

$$d\langle W^B, W^R\rangle_t = \rho dt$$

with constant $\rho \in [-1,1]$.[18] After applying Itô's formula to calculate $d\hat{B}_t$, it is easy to verify that there exists a Brownian motion W such that (2.10) holds with volatility function $\nu : [0,\tau] \to \mathbb{R}_+$ given by

$$\nu(t) = \sqrt{\sigma_B{}^2(t) - 2\rho\sigma_B(t)\sigma_R(t) + \sigma_R{}^2(t)}.$$

In fact, W can be defined by

$$dW_t = \frac{\sigma_B(t)}{\nu(t)} dW_t^B - \frac{\sigma_R(t)}{\nu(t)} dW_t^R \,.$$

The state space is $I =]0,\infty[$. The unique solution of (2.11) satisfying the terminal value condition and the bounds specified in Lemma 2.2.2 is well known:[19]

$$u(t,x) = x\,\Phi\left(\frac{1}{\sqrt{s(t)}}\left[\log\frac{x}{K} + \frac{s(t)}{2}\right]\right) - K\,\Phi\left(\frac{1}{\sqrt{s(t)}}\left[\log\frac{x}{K} - \frac{s(t)}{2}\right]\right) \tag{2.15}$$

where Φ denotes the standard normal distribution function and

$$s(t) = \int_t^\tau \nu^2(\xi)\, d\xi.$$

This yields the familiar formula for the arbitrage price of a call:

$$C_t = B_t\,\Phi(d_t^+) - K\,R_t\,\Phi(d_t^-) \tag{2.16}$$

[18]This is the framework common to Ball and Torous (1983) and Kemna, de Munnik and Vorst (1989). The models that Black and Scholes (1973) and Merton (1973) used for stock option pricing can also be seen as special cases of (2.13) - (2.14). The correlation coefficient ρ could of course be made time dependent as well.

[19]The growth condition $0 \leq u(t,x) \leq x$ guarantees uniqueness of the solution; see Gleit (1978).

with

$$d_t^{\pm} = \frac{1}{\sqrt{s(t)}}\left[\log\frac{B_t}{KR_t} \pm \frac{s(t)}{2}\right]. \tag{2.17}$$

It is easy to verify that a generating strategy for the option is

$$\theta_t^1 = \Phi(d_t^+), \quad \theta_t^0 = -K\,\Phi(d_t^-).$$

The arbitrage price of a European put option can be determined by put-call parity:

$$P_t = -B_t\,\Phi(-d_t^+) + K\,R_t\,\Phi(-d_t^-). \tag{2.18}$$

The strategy

$$\theta_t^1 = -\Phi(-d_t^+), \quad \theta_t^0 = K\,\Phi(-d_t^-)$$

generates the put option.

As Merton (1973) was first to derive option price formulae of this type for a stochastic reference bond, we shall call (2.16) and (2.18) the *Merton call and put price formula*, respectively, for volatility function ν.[20]

2.3 Constant Volatility: The Brownian Bridge

The first paper using the direct approach to price call and put options on zero-coupon bonds is Ball and Torous (1983). Their analysis starts from the following observation. The Black and Scholes (1973) model of stock price movements, a geometric Brownian motion

$$S_t = S_0\,\exp\big([\mu - \sigma^2/2]t + \sigma W_t\big) \tag{2.19}$$

with constants μ and σ, cannot be reinterpreted as a model of bond prices since this process specification is incompatible with a terminal value condition of the form (2.1). In fact, the variance of the process is strictly increasing with time.

[20]Formulae of this type also hold in so-called linear Gaussian models of the term structure of interest rates. Examples are Vasicek (1977) and its extension by Hull and White (1990). A systematic analysis of Gaussian models, as well as derivations of the pricing formulae we are considering here, can be found in El Karoui and Rochet (1989), El Karoui, Myneni and Viswanathan (1992), and Jamshidian (1991). The deterministic volatility examples in Heath, Jarrow and Morton (1992) belong also to this category.

Ball and Torous incorporate the terminal value condition by replacing the Wiener process in (2.19) with a standard Brownian bridge, i.e., a continuous Gaussian process $\eta_{t,T}$ $(0 \leq t \leq T)$ with $\eta_{0,T} = \eta_{T,T} = 0$, zero mean, and covariance function $\mathrm{E}[\eta_{t,T}\,\eta_{s,T}] = t\,(T-s)/T$ for $t \leq s$. This bridge process can be constructed as the solution of the stochastic differential equation

$$d\eta_{t,T} = -\frac{\eta_{t,T}}{T-t}\,dt + dW_t \tag{2.20}$$

where W is a Wiener process. Note how the drift pulls the process back to zero. The pull-back force, $-1/(T-t)$, becomes stronger as time goes by and eventually pulls the process towards its fixed endpoint.[21]

More precisely, Ball and Torous model the price process of the underlying bond as the lognormal variable

$$B_t = B_0\,\exp(\mu_B t + \sigma_B \eta_{t,T}) \tag{2.21}$$

with constant $\sigma_B > 0$, and choose

$$\mu_B = -\frac{\log B_0}{T} = Y_{0,T}$$

in order to fulfil the terminal value condition (2.1). The bond price process consists of two parts: the price path that would occur if there were no uncertainty,

$$B_0\,\exp(\mu_B t) = B_0^{\frac{T-t}{T}},$$

and a stochastic term driven by $\eta_{t,T}$ that characterises the random fluctuations around the deterministic path. As the distribution of $\eta_{t,T}$ is symmetric around 0, the deterministic path describes the median bond price.

For option pricing, a reference bond with maturity τ equal to the exercise date of the option is needed. Ball and Torous suppose that the price process of the reference bond is of type (2.21) as well. This leads to the following model:

$$\begin{aligned} B_t &= B_0^{\frac{T-t}{T}}\,\exp(\sigma_B \eta_{t,T}), \\ R_t &= R_0^{\frac{\tau-t}{\tau}}\,\exp(\sigma_R \eta_{t,\tau}) \end{aligned}$$

[21]See Karlin and Taylor (1981) for a detailed treatment of the Brownian bridge.

with

$$
\begin{aligned}
d\eta_{t,T} &= -\frac{\eta_{t,T}}{T-t}\,dt + dW_t^B, \\
d\eta_{t,\tau} &= -\frac{\eta_{t,\tau}}{\tau-t}\,dt + dW_t^R.
\end{aligned}
$$

The instantaneous correlation coefficient between the Wiener processess W^B and W^R is assumed to be a constant ρ. The forward price $\hat{B} = B/R$ has the form

$$\hat{B}_t = \hat{B}_0 \exp([\mu_B - \mu_R]t + \sigma_B \eta_{t,T} - \sigma_R \eta_{t,\tau}).$$

Being the quotient of two lognormally distributed variables, it is itself lognormal. Therefore, negative forward yields also occur with positive probability.

By Itô's lemma,

$$
\begin{aligned}
dB_t &= \left(\frac{{\sigma_B}^2}{2} - \frac{\log B_t}{T-t}\right) B_t\,dt + \sigma_B B_t\,dW_t^B, \\
dR_t &= \left(\frac{{\sigma_R}^2}{2} - \frac{\log R_t}{\tau-t}\right) R_t\,dt + \sigma_R R_t\,dW_t^R.
\end{aligned}
$$

This is an example of the general specification (2.13)-(2.14). We can apply the results of Section 2.2 with a constant forward price volatility equal to

$$\sqrt{{\sigma_B}^2 - 2\rho\sigma_B\sigma_R + {\sigma_R}^2}$$

which we shall denote by σ. Substituting this into (2.16) gives the arbitrage price of a European call in the Ball-Torous setting:

$$C_t = B_t\,\Phi(d_t^+) - K\,R_t\,\Phi(d_t^-) \tag{2.22}$$

where

$$d_t^\pm = \frac{1}{\sigma\sqrt{\tau - t}}\left[\log\frac{B_t}{KR_t} \pm \tfrac{1}{2}\sigma^2\,(\tau - t)^2\right].$$

This call price formula has exactly the same structure as the Black-Scholes formula.[22] This may be surprising at first sight: after all, within the common framework of (2.13)-(2.14), the above bond price model differs considerably from

[22] Formally, the Black-Scholes call price formula is obtained from (2.22) by setting $\sigma_R = 0$, i.e., by assuming the reference bond to have a constant yield, and by replacing B_t with the stock price.

the model of Black and Scholes (1973). Yet we have seen that only the volatility of the forward price of the underlying asset enters the option price formula. As both models assume this volatility to be constant, the similarity of the resulting pricing relationships is easily explained.

Despite this similarity, the formulae (2.22) and the Black-Scholes formula have a very different theoretical status. While the Black-Scholes model possesses a martingale measure[23] and hence satisfies sufficient conditions for the absence of arbitrage opportunities, the Ball-Torous model admits no martingale measure. Cheng (1991) shows that the drift term of the Brownian bridge which forces the process towards a fixed endpoint is incompatible with the requirements for the existence of a martingale measure. However, this does not necessarily imply that there are arbitrage opportunities in the Ball-Torous model: the existence of a martingale measure is sufficient, but in general not necessary for the absence of arbitrage.[24] To stress the difference between the two models, we might say that pricing in the Black-Scholes model proceeds safely from sufficient conditions for no arbitrage, whereas pricing in the Ball-Torous setting is merely based on *necessary* conditions for no arbitrage: all we have shown is that *if* the Ball-Torous model is arbitrage-free, then the price of a call option must be given by equation (2.22).

On a less theoretical level, one can criticise the Ball-Torous bond price model for the unrealistic yield behaviour that it implies. This problem, together with a possible solution, will be addressed in the following section.

2.4 Time Dependent Volatility

Using a Brownian bridge, Ball and Torous succeed in specifying a bond price process that reaches par value at maturity. It is instructive to examine the resulting yield

[23] See for example Müller (1985).

[24] Existence of a martingale measure and absence of arbitrage are equivalent if the number of trading dates is finite; see for example Harrison and Pliska (1981). While the sufficiency part still goes through in models with an infinite number of trading dates, the necessity part breaks down. In fact, Back and Pliska (1991) construct an example of a securities market which is arbitrage-free, but has no martingale measure.

process. (2.21) implies

$$Y_{t,T} = \mu_B - \frac{\sigma_B}{T-t}\,\eta_{t,T}$$

which is normally distributed with mean μ_B and variance ${\sigma_B}^2 t/(T-t)T$. Note that the latter increases without bounds as t increases. We can analyse this further by looking at yield changes over infinitesimal time periods. The stochastic differential of $Y_{t,T}$ is

$$dY_{t,T} = -\frac{\sigma_B}{T-t}\,dW_t^B$$

by Itô's lemma and (2.20). Thus, the diffusion coefficient (instantaneous standard deviation) of the yield process explodes as t tends to T. Kemna, de Munnik and Vorst (1989) show that this causes almost every yield path to reach negative values. In other words, negative yields to maturity are generated with probability one!

This highlights the serious drawbacks of the Ball-Torous model. One possible way to avoid them is to replace the Brownian bridge $\eta_{t,T}$ by a process of the form

$$\tilde{\eta}_{t,T} = k(t,T)\,W_t^B$$

where $k(t,T)$, a continuously differentiable function defined for $t \in [0,T]$, is positive for $0 < t < T$ and zero for $t = T$. Defining μ_B as before and setting

$$B_t = B_0 \exp(\mu_B t + \sigma_B \tilde{\eta}_{t,T}) = B_0^{\frac{T-t}{T}} \exp(\sigma_B \tilde{\eta}_{t,T}), \qquad (2.23)$$

one obtains again a bond price model that satisfies the terminal value condition. The distributions of B_t and $Y_{t,T}$ are lognormal and normal, respectively. More precisely,

$$Y_{t,T} = \mu_B - \frac{\sigma_B}{T-t}\,\tilde{\eta}_{t,T}$$

has mean μ_B and variance ${\sigma_B}^2\, k^2(t,T)\, t/(T-t)^2$ which stays bounded as t tends to T if and only if $k(t,T)/(T-t)$ does. This is also the condition for the diffusion coefficient of $Y_{t,T}$ to stay bounded, as we can see by applying Itô's lemma:

$$d\tilde{\eta}_{t,T} = \frac{k'(t,T)}{k(t,T)}\,\tilde{\eta}_{t,T}\,dt + k(t,T)\,dW_t^B$$

and

$$\begin{aligned} dY_{t,T} &= -\frac{\sigma_B}{(T-t)^2}\,\tilde{\eta}_{t,T}\,dt - \frac{\sigma_B}{T-t}\,d\tilde{\eta}_{t,T} \\ &= \left[\frac{1}{T-t} + \frac{k'(t,T)}{k(t,T)}\right] (Y_{t,T} - \mu_B)\,dt - \frac{\sigma_B\, k(t,T)}{T-t}\,dW_t^B. \end{aligned}$$

We have no empirical argument for a special form of the function $k(t,T)$. Kemna, de Munnik and Vorst (1989) propose $k(t,T) = (T-t)/T$. The resulting yield process is simply a Brownian motion starting at μ_B. This model succeeds where the Ball-Torous model fails. First, yields to maturity have bounded variance. Second, while negative yields occur with positive probability, as is the case in any model with lognormal bond prices, this probability is far smaller than one for reasonable parameter values. Third, de Munnik (1992) shows that this model admits a martingale measure and hence precludes arbitrage opportunities.

Turning to the valuation of bond options when bond prices are of the form (2.23), we use Itô's formula once more to calculate the stochastic differential of the process B. The result is

$$dB_t = \alpha_t^B B_t\, dt + \sigma_B\, k(t,T)\, B_t\, dW_t^B$$

with drift rate process

$$\alpha_t^B = \mu_B + \sigma_B\, k'(t,T)\, W_t^B + \tfrac{1}{2}{\sigma_B}^2\, k^2(t,T).$$

Let R_t, the price of the reference bond, also be of the form (2.23), i.e.,

$$R_t = R_0\, \exp(\mu_R t + \sigma_R \tilde{\eta}_{t,\tau})$$

with $\tilde{\eta}_{t,\tau} = k(t,\tau)\, W_t^R$, and assume that the instantaneous correlation coefficient ρ of the Wiener processes W^B and W^R is constant. This is again a special case of (2.13)-(2.14). The volatility of the forward bond price is time dependent:

$$\nu(t) = \sqrt{{\sigma_B}^2\, k^2(t,T) - 2\rho\sigma_B\sigma_R\, k(t,T)\, k(t,\tau) + {\sigma_R}^2\, k^2(t,\tau)}.$$

The arbitrage price for a European call in this situation has again the familiar form (2.16)-(2.17) with the function $s(t)$ now given by

$$s(t) = {\sigma_B}^2 \int_t^T k^2(\xi,T)\, d\xi - 2\rho\sigma_B\sigma_R \int_t^T k(\xi,T)\, k(\xi,\tau)\, d\xi + {\sigma_R}^2 \int_t^T k^2(\xi,\tau)\, d\xi.$$

While two flaws of the Ball-Torous model, namely the exploding variance of the yield to maturity and the non-existence of a martingale measure, can be remedied by specifying bond price processes with time dependent volatility, a major problem remains unsolved. In all the models considered so far, yields to maturity and forward yields can take negative values. This in turn distorts option prices: a call option on a zero-coupon bond with exercise price equal to the bond's face value has a positive price in these models. Schöbel (1986) and Bühler and Käsler (1989) propose solutions to this problem. We shall analyse them in the following two sections.

2.5 Correcting for Negative Yields: An Additional Boundary Condition

We have seen in Section 2.1 that in a bond price model with non-negative forward yields, the price of a call with exercise price $0 < K < 1$ satisfies condition (2.7). In terms of forward prices, this says that $\hat{C}_t = 1 - K$ whenever $\hat{B}_t = 1$.

The pricing formulae derived in lognormal models such as Ball and Torous (1983) or Kemna, de Munnik and Vorst (1989) are easily seen to violate condition (2.7), which reflects the fact that these models do generate negative forward yields. More generally, call prices calculated in any model where (2.10) holds with at most time dependent volatility function $\nu : [0, \tau] \to I\!R_+$ will violate (2.7). In fact, the Merton call price formula (2.16) implies that for $\hat{B}_t = 1$,

$$\hat{C}_t = \Phi\left(\frac{-\log K + s(t)/2}{\sqrt{s(t)}}\right) - K\,\Phi\left(\frac{-\log K - s(t)/2}{\sqrt{s(t)}}\right)$$

which is different from $1 - K$ in general.

In view of this, Schöbel (1986) and Briys, Crouhy and Schöbel (1991) propose to replace (2.16) with the modified call price formula

$$C_t = R_t\, u^*(t, \hat{B}_t) \tag{2.24}$$

where $u^* : [0, \tau] \times [0, 1] \to I\!R_+$ solves again the partial differential equation (2.11) for the given volatility function ν, but now subject to the conditions

$$\begin{aligned} u^*(\tau, x) &= [x - K]^+ , \\ u^*(t, 0) &= 0 , \\ u^*(t, 1) &= 1 - K . \end{aligned}$$

The first equation is the usual terminal value condition. The second is a boundary condition derived from (2.5). The third condition is new; it imposes property (2.7).

Schöbel solves this problem by transforming it into a heat conduction problem on the non-negative real half-axis.[25] This transformation is rather complicated and, as it turns out, unnecessary. To see this, let $u(t, x; K)$ be the Merton solution (2.15) for arbitrary $K > 0$. For all $0 \leq t \leq \tau$, we then have

$$u(t, 1; K) = \Phi\left(\frac{-\log K + s(t)/2}{\sqrt{s(t)}}\right) - K\,\Phi\left(\frac{-\log K - s(t)/2}{\sqrt{s(t)}}\right)$$

and

$$u(t, 1; K^{-1}) = \Phi\left(\frac{\log K + s(t)/2}{\sqrt{s(t)}}\right) - K^{-1}\,\Phi\left(\frac{\log K - s(t)/2}{\sqrt{s(t)}}\right)$$

where

$$s(t) = \int_t^\tau \nu^2(\xi)\, d\xi$$

as usual. This implies

$$\begin{aligned} u(t, 1; K) - K\, u(t, 1; K^{-1}) &= \Phi\left(\frac{-\log K + s(t)/2}{\sqrt{s(t)}}\right) - K\,\Phi\left(\frac{-\log K - s(t)/2}{\sqrt{s(t)}}\right) \\ &\quad - K\,\Phi\left(\frac{\log K + s(t)/2}{\sqrt{s(t)}}\right) + \Phi\left(\frac{\log K - s(t)/2}{\sqrt{s(t)}}\right) \\ &= 1 - K \end{aligned}$$

since $\Phi(-z) + \Phi(z) = 1$. Therefore, the function $u^* : [0, \tau] \times [0, 1] \to I\!R_+$ defined by

$$u^*(t, x) = u(t, x; K) - K\, u(t, x; K^{-1})$$

for $0 < K < 1$ clearly solves (2.11) with $u^*(t, 0) = 0$ and $u^*(t, 1) = 1 - K$. Moreover,

$$u^*(\tau, x) = [x - K]^+ - K\,\left[x - K^{-1}\right]^+ = [x - K]^+$$

as $x \leq 1 < K^{-1}$. Computing (2.24) with this function now yields the call price formula

$$C_t = B_t\,\Phi(d_t^+) - K\,R_t\,\Phi(d_t^-) - K\,B_t\,\Phi(\tilde{d}_t^+) + R_t\,\Phi(\tilde{d}_t^-) \tag{2.25}$$

where $d_t^\pm$ is given by (2.17) and

$$\tilde{d}_t^\pm = \frac{1}{\sqrt{s(t)}}\left[\log\frac{K B_t}{R_t} \pm \frac{s(t)}{2}\right].$$

[25] See also Schöbel (1987).

The first two terms of (2.25) coincide of course with formula (2.16). Schöbel (1986) writes the last two terms of (2.25) as $-[K\,B_t\,\Phi(\tilde{d}_t^+) - R_t\,\Phi(\tilde{d}_t^-)]$ and interprets the expression in square brackets as the Merton price of a European call with exercise price 1 written on a discount bond of face value K.[26] He calls this the "antioption". Our derivation of the pricing formula suggests a simpler interpretation. In the case of constant volatility, for example, the expression in square brackets is simply the Merton price of K calls with exercise price K^{-1} written on the original underlying bond.

For the European put option, put-call parity yields

$$P_t = -B_t\,\Phi(-d_t^+) + K\,R_t\,\Phi(-d_t^-) - K\,B_t\,\Phi(\tilde{d}_t^+) + R_t\,\Phi(\tilde{d}_t^-), \qquad (2.26)$$

which is the put price (2.18) minus the same "antioption" price as in (2.25).

Neither Schöbel (1986) nor Briys, Crouhy and Schöbel (1991) specify a bond price model such that portfolio duplication would lead to formulae (2.25) and (2.26). As a first step towards identifying such a model, we follow Breeden and Litzenberger (1978) and calculate so-called *Arrow-Debreu* or *state prices* implied by the above option prices. Assume that time 0 bond prices are B_0 and R_0 with $B_0 \leq R_0$. As $R_\tau = 1$, the states of the world at time τ can be identified with the possible realisations of B_τ. If there are no negative yields, we thus have the continuum of states $]0,1]$. We look for a distribution function F with $F(0) = 0$ and $F(1) = 1$ such that time 0 bond and option prices are discounted expected values of time τ payoffs with respect to F.[27] This means

$$B_0 = R_0 \int_0^1 x\,dF(x), \qquad (2.27)$$

and

$$P_0^K = R_0 \int_0^1 [K-x]^+\,dF(x) = R_0 \int_0^K (K-x)\,dF(x) \qquad (2.28)$$

where P_0^K denotes the initial price of the put for exercise price $0 \leq K \leq 1$. In the usual way, $R_0\,F(x)$ can be interpreted as the price of the Arrow-Debreu security $1_{\{B_\tau \leq x\}}$ paying 1 if $B_\tau \leq x$, and 0 otherwise. (2.27) and (2.28) express

[26]In fact, Schöbel (1986) deals only with the case of constant volatility, but his arguments obviously extend to the time dependent case.

[27]F can be interpreted as the distribution function of B_τ under a martingale measure.

the consistency of these Arrow-Debreu prices with "actual" prices of bonds and options.

Integration by parts yields

$$\int_0^K (K-x)\,dF(x) = [(K-x)\,F(x)]_0^K + \int_0^K F(x)\,dx = \int_0^K F(x)\,dx,$$

hence

$$P_0^K = R_0 \int_0^K F(x)dx.$$

The put price P_0^K has derivatives of all orders with respect to K in $]0,1[$, and

$$F(K) = \frac{1}{R_0}\,\frac{\partial}{\partial K} P_0^K$$

for $0 < K < 1$, so F inherits differentiability on the open unit interval. To determine F more precisely, we calculate the derivative[28]

$$\begin{aligned}\frac{\partial}{\partial K} P_0^K &= -\frac{B_0\,\phi(-d_0^+)}{K\,\sqrt{s(0)}} + R_0\,\Phi(-d_0^-) + \frac{R_0\,\phi(-d_0^-)}{\sqrt{s(0)}} \\ &\quad + \frac{R_0\,\phi(\tilde{d}_0^-)}{K\,\sqrt{s(0)}} - B_0\,\Phi(\tilde{d}_0^+) - \frac{B_0\,\phi(\tilde{d}_0^+)}{\sqrt{s(0)}} \\ &= R_0\,\Phi(-d_0^-) - B_0\,\Phi(\tilde{d}_0^+).\end{aligned}$$

Thus,

$$F(K) = \Phi(-d_0^-) - \hat{B}_0\,\Phi(\tilde{d}_0^+). \tag{2.29}$$

Note that $F(K) \to 0$ for $K \downarrow 0$, so F is continuous at 0. For $K \uparrow 1$, however,

$$F(K) \to \Phi\left(-\frac{\log \hat{B}_0 - s(0)/2}{\sqrt{s(0)}}\right) - \hat{B}_0\,\Phi\left(\frac{\log \hat{B}_0 + s(0)/2}{\sqrt{s(0)}}\right) < 1,$$

so F has a jump at 1.

In other words, the Arrow-Debreu security $1_{\{B_\tau = 1\}}$ commands a positive price, whereas all the other securities $1_{\{B_\tau = x\}}$ with $x < 1$ have price zero. This must mean that due to boundary condition (2.7), the probability mass which the original bond price model placed on outcomes $B_\tau \geq 1$ is now concentrated in the state $B_\tau = 1$, so this state occurs with positive probability. In particular, a bond price model consistent with formulae (2.25) and (2.26) must assign positive probability to the

[28] Let ϕ denote the standard normal density function. We use the following two facts in the calculation: $K\,R_0\,\phi(-d_0^-) = B_0\,\phi(-d_0^+)$ and $K\,B_0\,\phi(\tilde{d}_0^-) = R_0\,\phi(\tilde{d}_0^+)$.

event that the yield $Y_{\tau,T}$ becomes zero. Thus, with positive probability, there will be no reward for holding the underlying bond from τ to T.

Using a different approach, we shall show in Chapter 4 that in any arbitrage-free bond price model with non-negative forward yields that supports the option price formulae (2.25) and (2.26), the forward yield process $Y_{t,\tau,T}$ necessarily has an absorbing boundary at 0, and reaches this boundary with positive probability. In other words, at each time $0 < t \leq \tau$, there is a positive probability that $B_t = R_t$, and once this has happened, the bond prices coincide until τ. Thus, while satisfying condition (2.7), the proposed valuation formulae imply a rather implausible bond price and forward yield behaviour. A more satisfactory model will be presented in the following section.

2.6 State Dependent Volatility

In Section 2.4, we considered models of the type

$$R_t = m_R(t) \exp\bigl(\sigma_R(t)\, W_t^R\bigr)$$

$$B_t = m_B(t) \exp\bigl(\sigma_B(t)\, W_t^B\bigr)$$

with functions m_R, m_B, σ_R and σ_B being at most time dependent.[29] Such a model postulates that after taking the logarithm of bond prices, i.e., after applying the bijective mapping

$$\Lambda : \mathbb{R}^2_{++} \to \mathbb{R}^2 \ : \ \begin{pmatrix} r \\ b \end{pmatrix} \mapsto \begin{pmatrix} \log r \\ \log b \end{pmatrix},$$

we are dealing with Gaussian processes. More precisely, the image of the bond prices under Λ is equal to the image of the deterministic components plus a Wiener process term with time dependent coefficients:

$$\Lambda \begin{pmatrix} R_t \\ B_t \end{pmatrix} = \Lambda \begin{pmatrix} m_R(t) \\ m_B(t) \end{pmatrix} + \begin{pmatrix} \sigma_R(t)\, W_t^R \\ \sigma_B(t)\, W_t^B \end{pmatrix} \tag{2.30}$$

The main argument against this approach is that such a model generates negative yields. Indeed, to ensure positive yields to maturity and positive forward yields,

[29] In fact, we had for instance $m_R(t) = R_0^{\tau - t/\tau}$ and, with a slight abuse of notation, $\sigma_R(t) = \sigma_R\, k(t, \tau)$.

the vector of bond prices $\binom{R_t}{B_t}$ ought to take values in the triangle

$$D = \left\{ \begin{pmatrix} r \\ b \end{pmatrix} \in \,]0,1[^{\,2} \; : \; r > b \right\}.$$

Given any bijective mapping $\Psi : D \to I\!R^2$, we can construct a model that has positive yields by rewriting (2.30) with Ψ rather than Λ, i.e., by postulating that bond prices satisfy

$$\Psi \begin{pmatrix} R_t \\ B_t \end{pmatrix} = \Psi \begin{pmatrix} m_R(t) \\ m_B(t) \end{pmatrix} + \begin{pmatrix} \sigma_R(t)\, W_t^R \\ \sigma_B(t)\, W_t^B \end{pmatrix}.$$

The bond prices themselves can then be recovered by means of the inverse mapping $\Psi^{-1} : I\!R^2 \to D$.

However, which transformation Ψ should one use? There is no obvious choice. Ideally, it would be a simple mapping that leads to a tractable bond price distribution and closed-form solutions for option pricing. In fact, these goals are achievable, as Bühler and Käsler (1989) prove with the very ingenuous choice of the mapping[30]

$$\Psi : D \to I\!R^2 : \begin{pmatrix} r \\ b \end{pmatrix} \mapsto \begin{pmatrix} \log \frac{r}{1-r} \\ \log \frac{b}{r-b} \end{pmatrix}.$$

Its inverse is given by

$$\Psi^{-1} : I\!R^2 \to D : \begin{pmatrix} w_1 \\ w_2 \end{pmatrix} \mapsto \begin{pmatrix} [1 + e^{-w_1}]^{-1} \\ [1 + e^{-w_1}]^{-1}\, [1 + e^{-w_2}]^{-1} \end{pmatrix}.$$

The resulting bond prices are

$$R_t = \left[1 + \frac{1 - m_R(t)}{m_R(t)} \exp\bigl(-\sigma_R(t)\, W_t^R\bigr) \right]^{-1} \tag{2.31}$$

and

$$B_t = R_t \left[1 + \frac{m_R(t) - m_B(t)}{m_B(t)} \exp\bigl(-\sigma_B(t)\, W_t^B\bigr) \right]^{-1}. \tag{2.32}$$

Thus the price of the underlying bond depends explicitly on the price of the reference bond. In particular, both sources of uncertainty, W^B and W^R, have an impact on the price process of the underlying bond. By contrast, the forward bond price has a relatively simple representation, involving only the Wiener process W^B:

$$\hat{B}_t = \left[1 + \frac{m_R(t) - m_B(t)}{m_B(t)} \exp\bigl(-\sigma_B(t)\, W_t^B\bigr) \right]^{-1}.$$

[30]See also Käsler (1991). A one-dimensional variant of this mapping was first used by Bühler (1988) to model the price process of a coupon bond. See below for a brief discussion of this model.

Finally, note that $m_R(t)$ is the median of R_t, while $m_B(t)/m_R(t)$ is the median of $\hat{B}_t$.

Bühler and Käsler (1989) develop this model for constant σ_R and σ_B; the generalisation to time dependent parameters presented here is trivial. They suggest estimating these parameters from the current term structure, but do not propose any method for doing so.

The model (2.31)-(2.32) fulfils all the natural requirements discussed in Section 2.1: the terminal value condition (2.1); equation (2.2) which precludes negative yields to maturity; and equation (2.3) which rules out negative forward yields.

The distributions of R_t and $\hat{B}_t$ and the conditional distribution of B_t given R_t belong to a class of distributions studied already in Johnson (1946, 1949).[31] It is easy to calculate their density functions. The bond price R_t, for example, has the density function

$$\frac{1}{\sqrt{2\pi\, t}\,\sigma_R(t)\, x\,(1-x)}\;\exp\left(-\frac{1}{2t\,\sigma_R{}^2(t)}\left[\log\frac{x}{1-x}-\log\frac{m_R(t)}{1-m_R(t)}\right]^2\right)$$

for $0 < x < 1$. Johnson has shown that random variables with density functions of this type have finite moments, but there are no closed-form expressions for them. One can show, however, that the expected value of R_t is bounded:[32]

$$\left[1+\frac{1-m_R(t)}{m_R(t)}\;\exp\!\left(\tfrac{1}{2}\sigma_R{}^2(t)\,t\right)\right]^{-1} \le \mathrm{E}[R_t] \le \left[1+\frac{1-m_R(t)}{m_R(t)}\exp\!\left(-\tfrac{1}{2}\sigma_R{}^2(t)\,t\right)\right]^{-1}.$$

A similar relationship holds for $\hat{B}_t$.

For option pricing, we need to calculate the stochastic differential of the forward price process $\hat{B}$. Granted differentiability of the functions m_B, m_R and σ_B, Itô's lemma yields[33]

$$\begin{aligned} d\hat{B} \;&=\; \left[\frac{m_B{}'\,m_R-m_B\,m_R{}'}{m_B\,[m_R-m_B]}+\sigma_B{}'\,W^B+\sigma_B{}^2\left(\tfrac{1}{2}-\hat{B}\right)\right]\hat{B}\,(1-\hat{B})\,dt \\ &\quad+\;\sigma_B\,\hat{B}\,(1-\hat{B})\,dW^B. \end{aligned}$$

[31] Johnson constructs classes of distributions by applying the "method of translation" to a standard normal variable Z. The class of lognormal distributions, for instance, is obtained by means of the exponential transformation $Z \mapsto \exp(\gamma+\delta Z)$. The transformation $Z \mapsto [\gamma+\delta\exp(\theta Z)]^{-1}$ defines a class which Johnson denotes by S_B. This is the type of distributions we are dealing with in the Bühler-Käsler model.

[32] Cf. Rady and Sandmann (1994).

[33] To simplify the notation, the time variable t has been omitted.

The volatility of the forward bond price is time and state dependent:

$$\nu(t,x) = \sigma_B(t)\,(1-x)$$

in the notation of Section 2.2. The state space of $\hat{B}$ is $]0,1[$. In view of Lemma 2.2.2, we want to solve

$$u_t + \tfrac{1}{2}\,\sigma_B{}^2\,x^2\,(1-x)^2\,u_{xx} = 0$$

subject to the terminal value condition

$$u(\tau,x) = [x-K]^+$$

and the bounds

$$[x-K]^+ \le u(t,x) \le \min\{x, 1-K\}.$$

The appendix shows how to solve this problem by transforming it into a heat conduction problem on the real axis.[34] The solution is

$$\begin{aligned} u(t,x) &= (1-K)\,x\,\Phi\!\left(\frac{1}{\sqrt{s(t)}}\left[\log\frac{(1-K)\,x}{K\,(1-x)} + \frac{s(t)}{2}\right]\right) \\ &\quad - K\,(1-x)\,\Phi\!\left(\frac{1}{\sqrt{s(t)}}\left[\log\frac{(1-K)\,x}{K\,(1-x)} - \frac{s(t)}{2}\right]\right) \end{aligned} \tag{2.33}$$

with

$$s(t) = \int_t^\tau \sigma_B{}^2(\xi)\,d\xi.$$

Consequently, the arbitrage price of the European call is

$$C_t = (1-K)\,B_t\,\Phi(e_t^+) - K\,(R_t - B_t)\,\Phi(e_t^-) \tag{2.34}$$

with

$$e_t^\pm = \frac{1}{\sqrt{s(t)}}\left[\log\frac{(1-K)\,B_t}{K\,(R_t-B_t)} \pm \frac{s(t)}{2}\right].$$

This is the formula proposed by Bühler and Käsler (1989). The strategy

$$\theta_t^1 = (1-K)\,\Phi(e_t^+) + K\,\Phi(e_t^-), \quad \theta_t^0 = -K\,\Phi(e_t^-)$$

generates the option.

[34] A probabilistic approach to the calculation of the option price will be presented in Chapter 3.

The model (2.31)-(2.32) is unique within the direct approach in as much as it guarantees positive yields to maturity as well as positive forward yields and still produces a closed-form solution for the arbitrage price of a standard European option. Moreover, the existence of a martingale measure is easily demonstrated for this model.[35]

2.7 Options on Coupon Bonds

Prior to Bühler and Käsler (1989), bond price models with state dependent volatility have first been proposed by authors using the direct approach to price options on coupon bonds.[36]

Schaefer and Schwartz (1987) assume that the price process of the underlying coupon bond satisfies

$$dB_t = \alpha_t^B \, B_t \, dt + k \, B_t^\ell \, D(t, B_t) \, dW$$

where k and ℓ are constants and $D(t, B_t)$ is the duration of the bond.[37] This specification of volatility reflects the fact that bond returns become less variable as the maturity date approaches. The authors leave the drift rate process α^B unspecified because they are mainly interested in the connection between duration and the variability of bond returns, and because the drift rate does not enter the valuation equation (2.11) anyway. Due to the complicated volatility function, there are in general no analytic solutions for option prices, so a numerical procedure is needed to solve (2.11). Schaefer and Schwartz further assume that the reference bond has a constant rate of return r, that is,

$$dR_t = r \, R_t \, dt.$$

[35]It was said in Section 2.2 that a martingale measure exists if the quotient of the drift rate and the volatility of $\hat{B}$ satisfies certain integrability conditions. In the model of Bühler and Käsler (1989), this quotient is a bounded process and hence fulfils those conditions trivially. See Section 3.4 below for an explicit construction of a martingale measure.

[36]The holder of a coupon bond obtains a fixed amount F at the maturity date T plus a sequence of interest payments c_i at dates $T_1 < T_2 < \ldots < T_N \leq T$.

[37]The duration of a bond is a weighted average of the dates at which its cash flows occur; see Cox, Ingersoll and Ross (1979) for details. The duration of a zero-coupon bond is just its time to maturity. Thus, if $\ell = 1$ and the underlying bond pays no coupons, one obtains the same bond price volatility as in Kemna, de Munnik and Vorst (1989).

This assumption is of course inconsistent with a stochastic price for the underlying bond — it merely serves to keep the numerics of the valuation problem as simple as possible. Neither the terminal value condition nor the question of negative yields are addressed in this paper.

Bühler (1988), also using duration to describe bond volatility, proposes a more sophisticated alternative to Schaefer and Schwartz (1987). In his model, the price process of the underlying coupon bond fulfils the terminal value condition, and the bond yield remains always positive.[38] He starts from the following observation. Let $B_{\max}(t)$ be the par value plus the undiscounted coupon payments from t on. Then the yield of the bond at time t is positive if and only if $B_t < B_{\max}(t)$. Bühler goes on to construct a bond price process with this property[39] and obtains the following bond price dynamics:

$$dB_t = -\frac{\log B_t}{T-t}\, B_t\, dt + k\, B_t\, \frac{B_{\max}(t) - B_t}{B_{\max}(t) - 1 + \delta}\, D(t, B_t)\, dW_t$$

with constants k and δ. The drift term pulls the process towards the par value (which we have normalised to one) and away from the boundaries of the state space, 0 and $B_{\max}(t)$. Again, option prices must be calculated numerically. Rather than imposing a constant rate of return for the reference bond, Bühler simplifies the numerical procedure by specifying

$$dR_t = r(B_t)\, R_t\, dt$$

where $r(B_t)$ is the yield of the underlying bond multiplied by a time dependent factor. This supposes perfect instantaneous correlation between the bond yields, which, though far less restrictive than the assumption made by Schaefer and Schwartz, is still problematic.

It may well be that by relaxing the restrictive assumptions made by Schaefer and Schwartz or Bühler, the direct approach could eventually provide a fully satisfactory valuation model for options on coupon bonds. To judge from the Bühler

[38] The yield of a coupon bond is defined as that constant interest rate which makes the current price of the bond equal to the present discounted value of its future cash flows.

[39] This is the first example of the transformation method described at the beginning of Section 2.6. Bühler uses a monotonic mapping to transform a process with values in $\mathbb{R}$ in such a way that the resulting process has the desired features.

model, however, such an attempt would necessarily involve considerable technical complications. The term structure approach seems therefore more appropriate for the pricing of coupon bond options. Describing simultaneously the discount bonds of all maturities, this approach can treat coupon bonds simply as linear combinations of discount bonds. Thus, one encounters no particular modelling difficulties when moving from zero-coupon to coupon bonds.[40]

2.8 Conclusion

In this chapter, we have given a detailed survey of the so-called direct approach to debt option pricing. This approach specifies bond price processes directly, without relating them to the term structure as a whole or to state variables such as the short term interest rate. This approach is attractive for two reasons. First of all, it is *parsimonious* in that only those securities which are relevant to the pricing problem at hand have to be modelled. Moreover, modelling a small set of securities imposes fewer restrictions than modelling the whole term structure, say. Therefore, the direct approach is more *flexible*, for example in specifying the correlation between bonds.

Our presentation of the portfolio duplication technique in Section 2.2 stresses the fact that the volatility of the forward bond price is the crucial model characteristic for the calculation of option prices. Therefore, the chapter has been structured according to the specification of volatility, moving from constant volatility (Ball and Torous, 1983) over time dependent volatility (Kemna, de Munnik and Vorst, 1989) to time and state dependent volatility (Bühler and Käsler, 1989).

Focussing primarily on zero-coupon bonds, we have emphasised the main modelling problems encountered by the direct approach: first, the problem of specifying bond prices that fulfil the terminal value condition, i.e., that reach par value at

[40]Furthermore, there are well-known term stucture models which ensure positive yields and possess a martingale measure, e.g. Cox, Ingersoll and Ross (1985) or Heath, Jarrow and Morton (1992, Section 7). As for closed-form solutions, Jamshidian (1987, 1989) and El Karoui and Rochet (1989) showed that one-factor models of the term structure provide tractable formulae for the prices of European options on coupon bonds. In these models, the price of a coupon bond option can be written as the sum of the prices of discount bond options.

maturity; second, the problem of precluding negative yields to maturity and negative forward yields; third, the problem of ensuring an arbitrage-free bond price model.

The model of Bühler and Käsler (1989) is the only one to solve all three problems and still allow closed-form solutions for option prices. While lognormal models such as Ball and Torous (1983) and Kemna, de Munnik and Vorst (1989) also lead to explicit pricing formulae, their common weakness is that negative yields to maturity and negative forward yields occur with positive probability. Schöbel (1986) and Briys, Crouhy and Schöbel (1991) therefore propose modified pricing formulae which they obtain by imposing an additional boundary condition. Our analysis, which is to be completed in Chapter 4, indicates a serious flaw in Schöbel's method: in fact, he implicitly assumes that the forward yield process has an absorbing boundary at zero.

We have not tested the above models, nor have we dealt with the problem of parameter estimation. These issues are of course crucial for the choice of a model and its implementation. In practice, a simple model with some theoretical weaknesses will be preferred to a theoretically more satisfactory alternative if its parameters are easy to estimate and its weaknesses negligible for realistic parameter values. De Munnik (1992) argues along these lines when discussing the model of Kemna, de Munnik and Vorst (1989). He asserts that the model is easier to estimate than the Bühler-Käsler model, and shows that the probability of negative yields is small for realistic parameter values. Thus, the theoretical flaw of lognormal models may be irrelevant in practice. Since these models have the genuine advantage of providing familiar option price formulae, they will continue to play a role in practical applications.

As for the valuation of options on coupon bonds, we briefly discussed Schaefer and Schwartz (1987) and Bühler (1988). In the absence of closed-form solutions, these papers make very restrictive assumptions on the reference bond in order to simplify the numerical computation of option prices. In view of this, we concluded that the term structure approach is more appropriate for the pricing of options on coupon bonds.

Appendix

Solution of the Bühler-Käsler Terminal Value Problem

Consider the terminal value problem

$$u_t + \tfrac{1}{2}\,\sigma_B{}^2\,x^2\,(1-x)^2\,u_{xx} = 0,$$
$$u(\tau, x) = f(x)$$

for general f. This problem on $[0,\tau]\times]0,1[$ is transformed by introducing the new time variable

$$s = \int_t^\tau \sigma_B{}^2(\xi)\,d\xi,$$

the new space variable

$$z = \log\frac{x}{1-x} \qquad \text{or} \qquad x = \frac{1}{1+e^{-z}},$$

and by setting

$$u(t,x) = a(z)\,b(s)\,h(s,z).$$

Differentiable functions a and b are to be chosen in such a way that any solution h of the heat conduction equation yields a solution u of the original partial differential equation.

One easily calculates the derivatives

$$\begin{aligned}
z_x &= \frac{1}{x(1-x)}, \\
z_{xx} &= \frac{2x-1}{x^2(1-x)^2}, \\
u_x &= [a_z h + a h_z]\, b\, z_x\,, \\
u_{xx} &= \{a_{zz}h + 2a_z h_z + a h_{zz} + (2x-1)[a_z h + a h_z]\}\,\frac{b}{x^2(1-x)^2}, \\
u_t &= -a\,[b_s h + b h_s]\,\sigma_B{}^2.
\end{aligned}$$

Inserting this in the above PDE, dividing through by σ_B^2 and using

$$2x - 1 = \frac{e^z-1}{e^z+1} = \frac{e^{z/2}-e^{-z/2}}{e^{z/2}+e^{-z/2}} = \tanh\frac{z}{2},$$

we get

$$ab\left[\frac{h_{zz}}{2} - h_s\right] + \left(a_z + \frac{a}{2}\tanh\frac{z}{2}\right) bh_z + \left[\frac{1}{2}\left(a_{zz} + \frac{a_z}{2}\tanh\frac{z}{2}\right) b - ab_s\right] h = 0.$$

In order to make the h_z-term vanish, a has to solve the linear differential equation

$$a_z + \frac{a}{2}\tanh\frac{z}{2} = 0\ .$$

Separation of variables leads to the solutions

$$a(z) = \frac{c}{e^{z/2} + e^{-z/2}}\ ;$$

we choose $c = 1$.

Using the equation for a, one obtains

$$a_{zz} + a_z \tanh\frac{z}{2} = -\frac{a}{4}\,.$$

Therefore, the h-term vanishes if

$$\frac{b}{8} + b_s = 0.$$

We choose the solution

$$b(s) = e^{-s/8}\,.$$

Thus, setting

$$u(t,x) = \frac{e^{-s/8}}{e^{z/2} + e^{-z/2}}\ h(s,z),$$

we obtain the following transformed problem on $[0,\tau] \times I\!R$:

$$\begin{aligned}
&\tfrac{1}{2}h_{zz} - h = 0,\\
&h(0,z) = \left(e^{z/2} + e^{-z/2}\right) f\left(\frac{1}{1+e^{-z}}\right).
\end{aligned}$$

The solution of this problem is

$$h(s,z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h(0, z + \xi\sqrt{s}) e^{-\xi^2/2}\, d\xi.$$

We omit the corresponding formula for $u(t,x)$.

Valuation of a Call Option in the Bühler-Käsler Model

For

$$h(0,z) = (e^{z/2} + e^{-z/2}) \left[\frac{1}{1+e^{-z}} - K \right]^+ ,$$

the solution is given by

$$\begin{aligned} h(s,z) &= \frac{1}{\sqrt{2\pi}} \int_{\frac{1}{\sqrt{s}}\left[\log \frac{K}{1-K} - z\right]}^{\infty} \left(e^{\frac{1}{2}[z+\xi\sqrt{s}]} + e^{-\frac{1}{2}[z+\xi\sqrt{s}]} \right) \left(\frac{1}{1+e^{-[z+\xi\sqrt{s}]}} - K \right) e^{-\xi^2/2}\, d\xi \\ &= (1-K)\, I_1 - K\, I_2 \end{aligned}$$

with

$$I_1 = \frac{1}{\sqrt{2\pi}} \int_{\frac{1}{\sqrt{s}}\left[\log \frac{K}{1-K} - z\right]}^{\infty} e^{\frac{1}{2}[z+\xi\sqrt{s}]} e^{-\xi^2/2}\, d\xi = e^{z/2}\, e^{s/8}\, \Phi\left(\frac{1}{\sqrt{s}} \left[z + \log \frac{1-K}{K} + \frac{s}{2} \right] \right)$$

and

$$I_2 = \frac{1}{\sqrt{2\pi}} \int_{\frac{1}{\sqrt{s}}\left[\log \frac{K}{1-K} - z\right]}^{\infty} e^{-\frac{1}{2}[z+\xi\sqrt{s}]} e^{-\xi^2/2}\, d\xi = e^{-z/2}\, e^{s/8}\, \Phi\left(\frac{1}{\sqrt{s}} \left[z + \log \frac{1-K}{K} - \frac{s}{2} \right] \right).$$

Therefore,

$$\begin{aligned} u(t,x) &= \frac{e^{-s/8}}{e^{z/2} + e^{-z/2}}\, h(s,z) \\ &= (1-K) \underbrace{\frac{e^{z/2}}{e^{z/2} + e^{-z/2}}}_{=\,x} \Phi\left(\frac{1}{\sqrt{s}} \left[\log \frac{x(1-K)}{(1-x)K} + \frac{s}{2} \right] \right) \\ &\quad - K \underbrace{\frac{e^{-z/2}}{e^{z/2} + e^{-z/2}}}_{=\,1-x} \Phi\left(\frac{1}{\sqrt{s}} \left[\log \frac{x(1-K)}{(1-x)K} - \frac{s}{2} \right] \right), \end{aligned}$$

which is (2.33).

Chapter 3

Option Pricing with a Quadratic Diffusion Term

In the option pricing model of Black and Scholes (1973), the price of the underlying asset is a random variable with full support on the positive half-axis. This makes it difficult to apply the Black-Scholes model in situations where the underlying financial variable has bounded support. We have already seen an example in Chapter 2: if interest rates are positive, then spot and forward prices of bonds are clearly bounded. Bühler and Käsler (1989) therefore construct a model where the forward price of the underlying discount bond has a strict upper bound.

Exchange rates in a credible target zone regime also have bounded support. To price currency options in such an environment, Ingersoll (1989a, b) develops an exchange rate model with strict upper and lower stabilisation bounds.

The mathematical structure of these two models is exactly the same. In both cases, the underlying financial variable is assumed to be a diffusion process with the following characteristics: (i) the process has natural upper and lower boundaries; (ii) its diffusion coefficient is quadratic in the current value of the variable. This specification is easily seen to generalise the Black-Scholes model; in fact, the latter is obtained by choosing 0 and $+\infty$ as the lower and upper bounds.

It is remarkable that this generalisation preserves one of the most attractive features of the Black-Scholes model, namely the existence of closed-form solutions for the prices of European call and put options. Ingersoll (1989a, b) and Bühler and Käsler (1989) compute these solutions by applying a judicious change of variable to the corresponding fundamental valuation equation.[1] The present chapter, by contrast, applies a probabilistic change-of-numeraire technique which goes back to El Karoui and Rochet (1989). This technique makes the different steps in the calculation of the option price more transparent and easier to interpret. Moreover, it elucidates the structure of the pricing formula by decomposing it in terms of two particular numeraire portfolios and the risk-neutral probabilities associated with these.

The chapter is organised as follows. Section 3.1 sets out the framework of our analysis and introduces the change-of-numeraire technique. Section 3.2 presents a general expression for the price of a call option in the presence of strict upper and lower bounds on the underlying relative price. Applying this result, Section 3.3 calculates the call price in models where the underlying relative price has a quadratic diffusion term. Section 3.4 then shows how the general result applies to the models of Bühler and Käsler (1989) and Ingersoll (1989a, b). Section 3.5 concludes the chapter.

3.1 Martingale Measures and Numeraires

Fix a finite time interval $\mathcal{T} = [0, \tau]$, a probability space $(\Omega, \mathcal{F}, P)$ and a filtration $(\mathcal{F}_t)_{t \in \mathcal{T}}$ satisfying the usual conditions. $\mathcal{F}_0$ is assumed to be almost trivial, and $\mathcal{F}_\tau = \mathcal{F}$.

Consider a financial market with continuous and frictionless trade in two primitive assets, labelled 0 and 1, which pay no dividends in $\mathcal{T}$. Let their price processes S^i $(i = 0, 1)$ be positive semimartingales on $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t \in \mathcal{T}})$. Relative security prices are given by the process $X = S^1/S^0$.

A probability measure Q equivalent to P is called a *martingale measure with respect to asset* 0 if X is a Q-martingale, i.e., if each X_t is Q-integrable and

$$X_t = E^Q \left[X_\tau | \mathcal{F}_t \right]$$

[1] In fact, there is a slight difference in the approach taken. Ingersoll transforms the fundamental PDE into the standard Black-Scholes PDE and then uses the Black-Scholes solution. Bühler and Käsler, by contrast, transform the fundamental PDE directly into the heat equation which they solve in the usual way; see Käsler (1991). This is also the approach adopted in the appendix to Chapter 2.

for all $t \in \mathcal{T}$. Alternatively, such a measure Q is said to be *risk-neutral with respect to asset* 0. Let $\mathbb{P}_0$ denote the set of these measures.

Assumption (M) *$\mathbb{P}_0$ is non-empty.*

One element of $\mathbb{P}_0$, denoted by Q_0 and called the *reference measure*, will be held fixed throughout the chapter.

As in Harrison and Pliska (1983), a vector process $\theta = (\theta^0, \theta^1)$ is called an *admissible trading strategy* if the following properties (i) – (iv) hold:

(i) θ is predictable.

This expresses the informational restriction that trades can be based only on information obtained prior to trading. To formulate the remaining two conditions, let

$$V_t^\theta = \theta_t^0 S_t^0 + \theta_t^1 S_t^1$$

denote the *value process* corresponding to θ.

(ii) V^θ is non-negative.

(iii) θ^1 is integrable with respect to X, and the normalised value process V^θ / S^0 satisfies

$$\frac{V_t^\theta}{S_t^0} = \frac{V_0^\theta}{S_0^0} + \int_0^t \theta_s^1 \, dX_s.$$

(iv) The normalised value process V^θ / S^0 is a Q_0-martingale.

Condition (ii) rules out negative portfolio values. Condition (iii) states that all changes in portfolio value are due to the assets' performance rather than to injection or withdrawal of funds. In other words, admissible strategies are self-financing.[2] Condition (iv) says that there are no expected gains from trade. It rules out arbitrage opportunities and certain foolish strategies that throw away money.[3] The space of admissible strategies will be denoted by Θ.

[2]A straightforward integration-by-parts argument shows that (iii) implies the more intuitive representation

$$V_t^\theta = V_0^\theta + \int_0^t \theta_s^0 \, dS_s^0 + \int_0^t \theta_s^1 \, dS_s^1$$

for the value process, provided the integrals exist.

[3]Note that (iv) is the only condition that might depend on the choice of reference measure.

A positive process N is called a *numeraire* if there is a trading strategy $\theta \in \Theta$ such that $N = V^\theta$. Extending our previous definition, we call a probability measure Q equivalent to P a *martingale measure for numeraire* N (or *risk-neutral with respect to* N) if V^θ/N, the portfolio value expressed in units of the numeraire, is a Q-martingale for any strategy $\theta \in \Theta$. We shall write $\mathbb{P}_N$ for the set of all such measures, and $\mathbb{P}_1$ if $N = S^1$.

Given the measure Q_0 and a numeraire N, define a probability measure Q_N equivalent to Q_0 (and hence to P) via the Radon-Nikodym derivative

$$\frac{dQ_N}{dQ_0} = \frac{S_0^0}{N_0}\frac{N_\tau}{S_\tau^0}\,. \tag{3.1}$$

Note that N/S^0 is a Q_0-martingale by definition, so the right hand side of (3.1) has indeed expectation equal to one under Q_0. In case $N = S^1$, we shall write Q_1 for the measure defined by (3.1).

Lemma 3.1.1 *Let N be a numeraire and Y a random variable with* $\mathrm{E}^{Q_0}[|Y|/S_\tau^0] < \infty$. *Then*

$$\mathrm{E}^{Q_N}\left[\frac{Y}{N_\tau}\middle|\mathcal{F}_t\right] = \frac{S_t^0}{N_t}\,\mathrm{E}^{Q_0}\left[\frac{Y}{S_\tau^0}\middle|\mathcal{F}_t\right]$$

for all $t \in \mathcal{T}$.

Proof: The expectation on the left hand side is clearly well-defined and, by a version of Bayes' rule,

$$\mathrm{E}^{Q_N}\left[\frac{Y}{N_\tau}|\mathcal{F}_t\right] = \frac{\mathrm{E}^{Q_0}\left[\frac{dQ_N}{dQ_0}\frac{Y}{N_\tau}\middle|\mathcal{F}_t\right]}{\mathrm{E}^{Q_0}\left[\frac{dQ_N}{dQ_0}\middle|\mathcal{F}_t\right]}\,.$$

Using (3.1) and the fact that $\mathrm{E}^{Q_0}[N_\tau/S_\tau^0|\mathcal{F}_t] = N_t/S_t^0$ completes the proof. ■

Applying this lemma to $Y = V_\tau^\theta$, we see immediately that $Q_N \in \mathbb{P}_N$. We call it the martingale measure obtained from Q_0 by *change of numeraire.* If Q_N and $Q_{\tilde N}$ are obtained from Q_0 by changing the numeraire to N and $\tilde N$, respectively, then (3.1) implies

$$\frac{dQ_N}{dQ_{\tilde N}} = \frac{\tilde N_0}{N_0}\frac{N_\tau}{\tilde N_\tau}\,. \tag{3.2}$$

Equations (3.1) and (3.2) are at the heart of the change-of-numeraire technique in derivative asset pricing which goes back to El Karoui and Rochet (1989).[4]

A *contingent claim* is a non-negative random variable Γ on $(\Omega, \mathcal{F})$ such that Γ/S^0_τ is Q_0-integrable. A contingent claim is *attainable* if there exists a trading strategy $\theta \in \Theta$ that replicates the claim, i.e., that satisfies $V^\theta_\tau = \Gamma$. In this case, the portfolio value V^θ_t determines the time t *arbitrage price* $\pi_t(\Gamma)$ of the claim. By property (iv) above, this price can be calculated as

$$\pi_t(\Gamma) = S^0_t\, E^{Q_0}\left[\left.\frac{\Gamma}{S^0_\tau}\right| \mathcal{F}_t\right],$$

that is, without reference to the replicating strategy. More generally, consider an arbitrary measure $Q \in I\!P_0$ under which Γ/S^0_τ is integrable. Independent of whether Γ is attainable or not,

$$\pi^Q_t(\Gamma) = S^0_t\, E^{Q}\left[\left.\frac{\Gamma}{S^0_\tau}\right| \mathcal{F}_t\right]$$

is called the *price under* Q of the claim at time t.[5]

3.2 European Call Options

Consider an option to receive at time τ one unit of asset 1 in exchange for $K > 0$ units of asset 0. This is a slight generalisation of a classical European call option. Indeed, the latter is just the special case where asset 0 is a default-free zero-coupon bond of maturity τ.

The option has the following value at the exercise date:

$$\Gamma = \left[S^1_\tau - K S^0_\tau\right]^+$$

or, equivalently,

$$\Gamma = (S^1_\tau - K S^0_\tau)\, 1_{\mathcal{E}}$$

[4]For a more detailed examination of the relationship between numeraires and martingale measures see Conze and Viswanathan (1991).

[5]Jacka (1992) shows that a contingent claim Γ is attainable if and only if it has the same initial price $\pi_0(\Gamma)$ under all $Q \in I\!P_0$ for which both dQ_0/dQ and dQ/dQ_0 are bounded. Moreover, he shows that for bounded Γ/S^0_τ, the attainability of the claim does not depend on which reference measure Q_0 was used to define the space of admissible trading strategies.

where

$$\mathcal{E} = \{\omega \in \Omega : \ S^1_\tau(\omega) > K S^0_\tau(\omega)\}$$

is the event that the option ends "in the money" and is exercised.

It is well known that the price of a European option can be expressed in terms of exercise probabilities calculated under certain martingale measures. El Karoui and Rochet (1989) were the first to derive a variant of the following result.[6]

Proposition 3.2.1 *The option price under Q_0 is*

$$\pi^{Q_0}_t(\Gamma) = S^1_t \, Q_1 \left(\mathcal{E}|\mathcal{F}_t\right) - K S^0_t \, Q_0 \left(\mathcal{E}|\mathcal{F}_t\right)$$

where $Q_1 \in I\!P_1$ is the measure obtained from Q_0 by changing the numeraire to asset 1.

PROOF: By definition,

$$\pi^{Q_0}_t(\Gamma) = S^0_t \, \mathrm{E}^{Q_0}\left[\frac{\Gamma}{S^0_\tau}\bigg|\, \mathcal{F}_t\right] = S^0_t \, \mathrm{E}^{Q_0}\left[\frac{S^1_\tau}{S^0_\tau} 1_{\mathcal{E}}\bigg|\, \mathcal{F}_t\right] - K\, S^0_t \, \mathrm{E}^{Q_0}[1_{\mathcal{E}}|\mathcal{F}_t]\,.$$

Lemma 3.1.1 implies that

$$S^0_t \, \mathrm{E}^{Q_0}\left[\frac{S^1_\tau}{S^0_\tau} 1_{\mathcal{E}}\bigg|\, \mathcal{F}_t\right] = S^1_t \, \mathrm{E}^{Q_1}[1_{\mathcal{E}}|\mathcal{F}_t]\,,$$

hence the proposition. ■

A different decomposition of the option price can be obtained when the relative price $X = S^1/S^0$ is bounded.

Assumption (B) *There are constants $0 \le \ell < u \le +\infty$ such that*

$$\ell S^0_t < S^1_t < u S^0_t$$

for all $t \in \mathcal{T}$.

Consider two portfolios, the first of which is long one unit of asset 0 and short u^{-1} units of asset 1, while the second is long one unit of asset 1 and short ℓ units of asset 0.[7] Let

$$U = S^0 - u^{-1} S^1$$

[6]See also Geman, El Karoui and Rochet (1991).

[7]Of course, u^{-1} is understood to be zero if $u = +\infty$.

and

$$L = S^1 - \ell S^0$$

denote the corresponding value processes. Under Assumption (B), these are positive processes, hence numeraires.

Proposition 3.2.2 *Under Assumption* (B), *the option price under* Q_0 *is*

$$\pi_t^{Q_0}(\Gamma) = \frac{1}{1-u^{-1}\ell}\left\{(1-u^{-1}K)\, L_t\, Q_L(\mathcal{E}|\mathcal{F}_t) - (K-\ell)\, U_t\, Q_U(\mathcal{E}|\mathcal{F}_t)\right\}$$

where $Q_U \in I\!P_U$ *and* $Q_L \in I\!P_L$ *are the measures obtained from* Q_0 *by changing the numeraire to* U *and* L*, respectively.*

PROOF: It is straightforward to check that

$$S_\tau^1 - K S_\tau^0 = \frac{(1-u^{-1}K)\, L_\tau - (K-\ell)\, U_\tau}{1-u^{-1}\ell}\,.$$

Thus,

$$\begin{aligned}\pi_t^{Q_0}(\Gamma) &= \frac{1}{1-u^{-1}\ell}\left\{(1-u^{-1}K)\, S_t^0\, \mathrm{E}^{Q_0}\left[\frac{L_\tau}{S_\tau^0}1_{\mathcal{E}}\middle|\, \mathcal{F}_t\right]\right.\\ &\qquad\left. - (K-\ell)\, S_t^0\, \mathrm{E}^{Q_0}\left[\frac{U_\tau}{S_\tau^0}1_{\mathcal{E}}\middle|\, \mathcal{F}_t\right]\right\}.\end{aligned}$$

Lemma 3.1.1 now implies

$$S_t^0\, \mathrm{E}^{Q_0}\left[\frac{L_\tau}{S_\tau^0}1_{\mathcal{E}}\middle|\, \mathcal{F}_t\right] = L_t\, \mathrm{E}^{Q_L}[1_{\mathcal{E}}|\mathcal{F}_t]$$

and

$$S_t^0\, \mathrm{E}^{Q_0}\left[\frac{U_\tau}{S_\tau^0}1_{\mathcal{E}}\middle|\, \mathcal{F}_t\right] = U_t\, \mathrm{E}^{Q_U}[1_{\mathcal{E}}|\mathcal{F}_t]\,.$$

This is the desired result. ■

We have again expressed the call price as a function of certain exercise probabilities, this time evaluated under martingale measures associated with the numeraires U and L.

The exercise event $\mathcal{E}$ can be characterised in terms of the random variable $Y_\tau = L_\tau/U_\tau$:

$$\mathcal{E} = \left\{\omega \in \Omega :\ Y_\tau(\omega) > \frac{K-\ell}{1-u^{-1}K}\right\}.$$

Ingersoll (1989a, b) and Bühler and Käsler (1989) propose models where the law of the process $Y = L/U$ under both Q_U and Q_L is very simple, so that the above exercise probabilities are easy to determine.

3.3 Models with a Quadratic Diffusion Term

The following assumption postulates that after a change of measure, relative asset prices follow a diffusion process with quadratic diffusion coefficient. We shall see later that the models of Ingersoll (1989a, b) and Bühler and Käsler (1989) are of this type. Let constants $\sigma > 0$ and $0 \leq \ell < u \leq +\infty$ be given.

Assumption (Q) *There exists a Q_0-Wiener process W^0 such that the process of relative asset prices $X = S^1/S^0$ solves the stochastic differential equation*

$$dX_t = \sigma\,(X_t - \ell)(1 - u^{-1}X_t)\,dW_t^0$$

with initial value $\ell < X_0 < u$.

Standard results from the theory of stochastic processes imply that the above stochastic differential equation has in fact a solution. This solution is unique both in the strong and weak sense, satisfies Assumption (B) and is a martingale; see for example Revuz and Yor (1991) and Karlin and Taylor (1981). In particular, Q_0 is indeed risk-neutral with respect to asset 0.

Note that the lognormal dynamics of Black and Scholes (1973) and Merton (1973) are obtained as the special case where $\ell = 0$ and $u = +\infty$.

3.3.1 Characterisation

It turns out that Assumption (Q) can be formulated equivalently in terms of the process $Y = L/U$. Let $Q_U \in \mathbb{P}_U$ be the measure obtained from Q_0 by changing the numeraire to U, and define $\hat{\sigma} = (1 - u^{-1}\ell)\sigma$.

Lemma 3.3.1 *Assumption* (Q) *holds if and only if there exists a Q_U-Wiener process W^U such that Y solves*

$$dY_t = \hat{\sigma}\,Y_t\,dW_t^U$$

with initial value $Y_0 > 0$.

PROOF: Suppose Assumption (Q) holds. By Itô's lemma and some algebra,[8]

$$dY_t = \hat{\sigma} Y_t \left\{ dW_t^0 + \tilde{\sigma}\,(X_t - \ell)\, dt \right\}$$

where $\tilde{\sigma} = u^{-1}\sigma$. Define a process W^U by $dW_t^U = dW_t^0 + \tilde{\sigma}\,(X_t - \ell)\, dt$ with $W_0^U = 0$. We want to show that W^U is a Wiener process under Q_U. By equation (3.2),

$$\frac{dQ_U}{dQ_0} = \frac{S_0^0}{U_0}\frac{U_\tau}{S_\tau^0} = \frac{1 - u^{-1}X_\tau}{1 - u^{-1}X_0}\,.$$

On the other hand,

$$\frac{d[1 - u^{-1}X_t]}{1 - u^{-1}X_t} = -\tilde{\sigma}\,(X_t - \ell)\, dW_t^0,$$

hence, by the formula for the martingale exponential,

$$1 - u^{-1}X_t = (1 - u^{-1}X_0)\,\exp\!\left(-\tilde{\sigma}\int_0^t (X_s - \ell)\, dW_s^0 - \tfrac{1}{2}\,\tilde{\sigma}^2 \int_0^t (X_s - \ell)^2\, ds\right).$$

In particular,

$$\frac{dQ_U}{dQ_0} = \exp\!\left(-\tilde{\sigma}\int_0^\tau (X_s - \ell)\, dW_s^0 - \tfrac{1}{2}\,\tilde{\sigma}^2 \int_0^\tau (X_s - \ell)^2\, ds\right).$$

Girsanov's theorem now implies that W^U is indeed a Q_U-Wiener process; cf. Revuz and Yor (1991).

Conversely, suppose we have W^U as in the lemma. Itô's lemma and some straightforward computations yield

$$dX_t = \sigma\,(X_t - \ell)(1 - u^{-1}X_t)\left\{ dW_t^U - \tilde{\sigma}\,(X_t - \ell)\, dt \right\}.$$

Let W^0 be the process defined by $dW_t^0 = dW_t^U - \tilde{\sigma}\,(X_t - \ell)\, dt$ with $W_0^0 = 0$. As

$$\frac{d[1 - u^{-1}X_t]}{1 - u^{-1}X_t} = -\tilde{\sigma}\,(X_t - \ell)\, dW_t^U + \tilde{\sigma}^2\,(X_t - \ell)^2\, dt,$$

[8]The following facts are used in the calculations. If

$$y = \frac{x - \ell}{1 - u^{-1}x}\,,$$

then

$$\frac{dy}{dx} = \frac{1 - u^{-1}\ell}{(1 - u^{-1}x)^2} \quad\text{and}\quad \frac{d^2y}{dx^2} = \frac{2u^{-1}(1 - u^{-1}\ell)}{(1 - u^{-1}x)^3}\,.$$

Moreover,

$$\frac{dx}{dy} = \frac{(1 - u^{-1}x)^2}{1 - u^{-1}\ell} \quad\text{and}\quad \frac{d^2x}{dy^2} = \frac{-2u^{-1}(1 - u^{-1}x)^3}{(1 - u^{-1}\ell)^2}\,.$$

the formula for the martingale exponential now implies

$$1 - u^{-1}X_t = (1 - u^{-1}X_0)\,\exp\left(-\tilde{\sigma}\int_0^t (X_s - \ell)\,dW_s^U + \tfrac{1}{2}\,\tilde{\sigma}^2\int_0^t (X_s - \ell)^2\,ds\right)$$

and

$$\frac{dQ_0}{dQ_U} = \frac{1 - u^{-1}X_0}{1 - u^{-1}X_\tau} = \exp\left(\tilde{\sigma}\int_0^\tau (X_s - \ell)\,dW_s^U - \tfrac{1}{2}\,\tilde{\sigma}^2\int_0^\tau (X_s - \ell)^2\,ds\right).$$

By Girsanov's theorem, W^0 is a Wiener process under Q_0. ■

Thus, Assumption (Q) holds if and only if there is a change of measure that makes the process Y a driftless geometric Brownian motion with volatility $\hat{\sigma}$. This is the key to our calculation of the option price.

3.3.2 The Option Price

Let $(\mathcal{G}_t)_{t\in\mathcal{T}}$ be the filtration generated by the process X, and set $\mathcal{G} = \mathcal{G}_\tau$. The following result is well known.

Proposition 3.3.1 *Under Assumption* (Q), *any contingent claim* Γ *with* $\mathcal{G}$*-measurable normalised payoff* Γ/S_τ^0 *is attainable.*

PROOF: The proposition is an immediate consequence of the martingale representation property of X on $(\Omega, \mathcal{G}, Q_0, (\mathcal{G}_t)_{t\in\mathcal{T}})$; see Revuz and Yor (1991). ■

In particular, this guarantees the attainability of the option to receive one unit of asset 1 in exchange for K units of asset 0, as its normalised payoff $[S_\tau^1 - K S_\tau^0]^+/S_\tau^0 = [X_\tau - K]^+$ is clearly measurable with respect to $\mathcal{G}$.[9]

Proposition 3.3.2 *Under Assumption* (Q), *the option to receive one unit of asset* 1 *in exchange for* K *units of asset* 0 *is attainable. For* $\ell < K < u$, *its time* t *arbitrage price is*

$$\pi_t(\Gamma) = \frac{1}{1 - u^{-1}\ell}\left\{(1 - u^{-1}K)\,(S_t^1 - \ell S_t^0)\,\Phi(e_t^+) - (K - \ell)\,(S_t^0 - u^{-1}S_t^1)\,\Phi(e_t^-)\right\}$$

[9]Moreover, the normalised payoff of the option is bounded, so attainability does not depend on which reference measure was chosen to define the space of admissible trading strategies; see Jacka (1992).

where

$$e_t^{\pm} = \frac{1}{\hat{\sigma}\sqrt{\tau - t}}\left[\log\frac{S_t^1 - \ell S_t^0}{S_t^0 - u^{-1}S_t^1} - \log\frac{K-\ell}{1-u^{-1}K} \pm \tfrac{1}{2}\hat{\sigma}^2(\tau - t)\right]$$

and $\hat{\sigma} = (1 - u^{-1}\ell)\sigma$.

PROOF: We want to apply Proposition 3.2.2, so let Q_U and Q_L be the measures obtained from Q_0 by changing the numeraire to U and L, respectively. To calculate the probability of exercise under Q_U, let W^U be a Q_U-Wiener process as in Lemma 3.3.1. By the formula for the martingale exponential,

$$Y_\tau = Y_0 \exp\bigl(\hat{\sigma}\, W_\tau^U - \tfrac{1}{2}\hat{\sigma}^2\tau\bigr).$$

The properties of the Wiener process W^U now imply

$$\begin{aligned}
Q_U(\mathcal{E}|\mathcal{F}_t) &= Q_U\left(Y_\tau > \frac{K-\ell}{1-u^{-1}K}\,\middle|\, Y_t\right) \\
&= Q_U\left(\log Y_\tau - \log Y_t > \log\frac{K-\ell}{1-u^{-1}K} - \log Y_t\right) \\
&= Q_U\left(\hat{\sigma}\,(W_\tau^U - W_t^U) > \log\frac{K-\ell}{1-u^{-1}K} - \log Y_t + \tfrac{1}{2}\hat{\sigma}^2(\tau - t)\right) \\
&= \Phi\left(\frac{1}{\hat{\sigma}\sqrt{\tau - t}}\left[\log Y_t - \log\frac{K-\ell}{1-u^{-1}K} - \tfrac{1}{2}\hat{\sigma}^2(\tau - t)\right]\right).
\end{aligned}$$

Next, define a process W^L by $dW_t^L = dW_t^U - \hat{\sigma}\, dt$ with $W_0^L = 0$. As

$$\frac{dQ_L}{dQ_U} = \frac{U_0}{L_0}\frac{L_\tau}{U_\tau} = \frac{Y_\tau}{Y_0} = \exp\bigl(\hat{\sigma}\, W_\tau^U - \tfrac{1}{2}\hat{\sigma}^2\tau\bigr),$$

Girsanov's theorem implies that W^L is a Wiener process under Q_L. By construction,

$$Y_\tau = Y_0 \exp\bigl(\hat{\sigma}\, W_\tau^L + \tfrac{1}{2}\hat{\sigma}^2\tau\bigr),$$

hence

$$\begin{aligned}
Q_L(\mathcal{E}|\mathcal{F}_t) &= Q_L\left(\hat{\sigma}\,(W_\tau^L - W_t^L) > \log\frac{K-\ell}{1-u^{-1}K} - \log Y_t - \tfrac{1}{2}\hat{\sigma}^2(\tau - t)\right) \\
&= \Phi\left(\frac{1}{\hat{\sigma}\sqrt{\tau - t}}\left[\log Y_t - \log\frac{K-\ell}{1-u^{-1}K} + \tfrac{1}{2}\hat{\sigma}^2(\tau - t)\right]\right).
\end{aligned}$$

This completes the proof. ∎

Standard arguments[10] show that the trading strategy

$$\begin{aligned}\theta_t^0 &= \frac{1}{1-u^{-1}\ell}\Big\{-(1-u^{-1}K)\,\ell\,\Phi(e_t^+)-(K-\ell)\,\Phi(e_t^-)\Big\},\\ \theta_t^1 &= \frac{1}{1-u^{-1}\ell}\Big\{(1-u^{-1}K)\,\Phi(e_t^+)+(K-\ell)\,u^{-1}\,\Phi(e_t^-)\Big\}\end{aligned}$$

is admissible and replicates the option.

For $\ell = 0$ and $u = +\infty$, we of course obtain the option price formula of Black and Scholes (1973) and Merton (1973) with $\hat\sigma = \sigma$. Setting $u = +\infty$ but $\ell > 0$ leads to a formula proposed by Rubinstein (1983).

The result is easily extended to allow a time dependent, but deterministic, parameter function $\sigma(t)$ in Assumption (Q). One only has to replace $\hat\sigma$ in Proposition 3.3.2 by

$$(1-u^{-1}\ell)\,\sqrt{\frac{1}{\tau-t}\int_t^\tau \sigma^2(s)\,ds}.$$

The price of a generalised put option, that is, an option to *give up* one unit of asset 1 in exchange for K units of asset 0, can be calculated in the same way. Alternatively, one can use a version of put-call parity.

3.4 Examples

This section shows how the models of Bühler and Käsler (1989) and Ingersoll (1989a, b) fit into the framework developed in the previous sections.

3.4.1 Options on Zero-Coupon Bonds

Fix dates $T > \tau > 0$ and let assets 0 and 1 be pure discount bonds without default risk, maturing at τ and T, respectively. Without loss of generality, their face values can be normalised to 1, i.e., $S_\tau^0 = 1$ and $S_T^1 = 1$. Consider a standard European call option written on bond 1 with exercise price K and exercise date τ. As $S_\tau^0 = 1$, this call can be considered as an option to receive one unit of bond 1 in exchange for K units of bond 0.

[10]See for instance Harrison and Pliska (1981).

Bühler and Käsler (1989) propose a model where the bond prices satisfy $S_t^0 < 1$ for $t < \tau$ and $S_t^1 < S_t^0$ for $t \in \mathcal{T}$. In particular, Assumption (B) holds with $u = 1$ and $\ell = 0$. More specifically, the relative price $X_t = S_t^1/S_t^0$ has the form

$$X_t = \left[1 + \frac{1 - m(t)}{m(t)}\, e^{-\sigma W_t}\right]^{-1}$$

where $m : \mathcal{T} \to]0,1[$ is a continuously differentiable function, σ a positive constant and W a standard Wiener process under the measure P.[11] X_t is the time t forward price of bond 1 for delivery at time τ, while $m(t)$ is the median value of this forward price.

We want to show that this model satisfies Assumption (Q). Itô's lemma yields

$$dX_t = \sigma\, X_t\,(1 - X_t)\ \{\alpha_t\, dt + dW_t\}$$

with the bounded process

$$\alpha_t = \frac{m'(t)}{\sigma m(t)[1 - m(t)]} + \sigma\,\left(\tfrac{1}{2} - X_t\right).$$

Define a process W^0 by

$$dW_t^0 = \alpha_t\, dt + dW_t$$

and $W_0^0 = 0$, and let Q_0 be the measure obtained via the Radon-Nikodym derivative

$$\frac{dQ_0}{dP} = \exp\!\left(-\int_0^{\tau} \alpha_s\, dW_s - \tfrac{1}{2}\int_0^{\tau} \alpha_s^2\, ds\right).$$

Girsanov's theorem implies that W^0 is a Wiener process under Q_0. By construction, $dX_t = \sigma\, X_t\,(1 - X_t)\, dW_t^0$, so Assumption (Q) holds.[12]

By Proposition 3.3.2, the arbitrage price of the call option with exercise price K between 0 and 1 is

$$\pi_t(\Gamma) = (1 - K)\, S_t^1\, \Phi(e_t^+) - K\,(S_t^0 - S_t^1)\,\Phi(e_t^-)$$

[11] $m(t)$ corresponds to $m_B(t)/m_R(t)$ in Section 2.6 and is the median value of the forward price X_t. The parameter σ here corresponds to σ_B there. As in Chapter 2, there would be no difficulty in allowing σ to depend on time.

[12] Alternatively, one could construct a measure under which the process

$$Y_t = \frac{X_t}{1 - X_t} = \frac{m(t)}{1 - m(t)}\, e^{\sigma W_t}$$

is a martingale and then use Lemma 3.3.1.

with

$$e_t^{\pm} = \frac{1}{\sigma\sqrt{\tau - t}} \left[\log \frac{S_t^1}{S_t^0 - S_t^1} - \log \frac{K}{1-K} \pm \tfrac{1}{2}\sigma^2 (\tau - t) \right].$$

This is the pricing formula derived in Bühler and Käsler (1989).

3.4.2 Currency Options in a Target Zone Regime

Consider an option to buy at some future date τ one unit of a foreign currency for K units of the domestic currency. If asset 0 is a default-free domestic discount bond paying one domestic currency unit at time τ, and asset 1 its foreign counterpart, then the currency option can be interpreted as the right to receive one unit of asset 1 in exchange for K units of asset 0. Note that S^1, the *domestic* price of asset 1, is the product of two factors: the spot exchange rate, s, giving the number of domestic currency units needed to purchase one unit of the foreign currency, and $S^{1,f}$, the price of asset 1 in *foreign* units. Assuming for simplicity that the domestic interest rate r_d and the foreign interest rate r_f are constant, we clearly have

$$S_t^0 = e^{-r_d(\tau - t)}, \quad S_t^{1,f} = e^{-r_f(\tau - t)} \quad \text{and} \quad S_\tau^1 = s_t\, e^{-r_f(\tau - t)}.$$

By covered interest rate parity, $X_t = S_t^1/S_t^0$ is now just the time t forward rate for currency exchange at time τ.

Ingersoll (1989a) models a perfectly credible target zone regime by imposing the condition

$$\xi(t) < s_t < \Xi(t)$$

with deterministic functions ξ and Ξ. He shows that not every pair of boundary functions is admissible. Given $\xi(0)$ and $\Xi(0)$, the tightest possible bounds are in fact

$$\begin{aligned} \xi(t) &= \xi(0)\, e^{(r_d - r_f)t}, \\ \Xi(t) &= \Xi(0)\, e^{(r_d - r_f)t}. \end{aligned}$$

For these functions, the above condition translates into

$$\xi(\tau)S_t^0 < S_t^1 < \Xi(\tau)S_t^0,$$

that is, Assumption (B) with $\ell = \xi(\tau)$ and $u = \Xi(\tau)$.

As for the spot rate dynamics, one of the models studied in Ingersoll (1989a) has

$$ds_t = \mu_t s_t\, dt + \sigma\, [s_t - \xi(t)][1 - s_t/\Xi(t)]\, dW_t$$

with an unspecified drift rate process μ, a Wiener process W, and the above boundary functions. By Itô's lemma, the corresponding forward rate dynamics are

$$dX_t = (\mu_t + r_f - r_d) X_t\, dt + \sigma\, [X_t - \xi(\tau)][1 - X_t/\Xi(\tau)]\, dW_t,$$

which, under suitable conditions on μ, implies Assumption (Q). If so, the arbitrage price of the currency option is given by Proposition 3.3.2 and can be written as

$$\pi_t(\Gamma) = [s_t - \xi(t)]\, S_t^{1,f}\, \frac{1 - K/\Xi(\tau)}{1 - \xi(\tau)/\Xi(\tau)}\, \Phi(e_t^+) - [1 - s_t/\Xi(t)]\, S_t^0\, \frac{K - \xi(\tau)}{1 - \xi(\tau)/\Xi(\tau)}\, \Phi(e_t^-)$$

with

$$e_t^{\pm} = \frac{1}{\hat{\sigma}\sqrt{\tau - t}} \left[\log \frac{s_t - \xi(t)}{1 - s_t/\Xi(t)} - \log \frac{K - \xi(\tau)}{1 - K/\Xi(\tau)} \pm \tfrac{1}{2}\hat{\sigma}^2 (\tau - t) \right]$$

and $\hat{\sigma} = [1 - \xi(\tau)/\Xi(\tau)]\sigma$. This is the same result as in Ingersoll (1989a).

An extension of this analysis to "futures-style" options (futures contracts on option payoffs) is presented in Ingersoll (1989b). Assuming a quadratic diffusion term for the underlying futures price, Ingersoll calculates valuation formulae similar to the one above. Again, the results of Sections 3.2 and 3.3 apply.

3.5 Conclusion

We have studied the pricing of an option to exchange one asset for another in the presence of strict upper and lower bounds on the relative price of these assets. Our first result shows how to decompose the option price in terms of two particular numeraire portfolios and the probabilities of exercise calculated under the martingale measures for these numeraires. This decomposition is particularly useful in models where the relative asset price has a quadratic diffusion coefficient. The second contribution of the chapter is a new derivation of option prices in this class of models.

Chapter 4

State Prices Implicit in Valuation Formulae for Derivative Assets

Derivative assets analysis usually takes a model of the underlying price processes as given and attempts to value derivatives relative to that model. Chapters 2 and 3 are examples of this approach. The present chapter addresses the converse question: given some set of derivatives prices, what can we say about the price processes of the underlying securities? More precisely, we assume that we have information about the price of a derivative asset in the form of a pricing formula, that is, a deterministic function of the underlying security prices and time, and investigate the restrictions such a formula imposes on the underlying price dynamics.

We restrict ourselves to the simplest possible setting with a riskless cash account, one risky security, and one derivative. Assuming that asset prices are continuous semimartingales, we consider pricing formulae that satisfy a variant of the fundamental valuation equation which is familiar from derivative asset pricing in a diffusion setting.[1] We show that such a pricing formula uniquely determines the law of the underlying asset price under a martingale measure, thus implying a unique system of state prices for payoffs contingent on the price path of the underlying security.

While similar in spirit to Breeden and Litzenberger's (1978) calculation of state prices implicit in option prices, our approach uses rather different mathematical tools, based mainly on semimartingale calculus. The main result follows directly from a characterisation theorem for continuous local martingales which extends

[1] Equation (2.11) is in fact a special case of this PDE.

work by McGill, Rajeev and Rao (1988) on Brownian motion.

To illustrate our approach, we return to the analysis of pricing formulae for options on zero-coupon bonds which we began in Chapter 2. A direct application of our result confirms that Merton type formulae are inconsistent with non-negative interest rates. We then show that the pricing formulae of Schöbel (1986) and Briys, Crouhy and Schöbel (1991) imply a positive probability for the forward yield process to hit the lower bound 0 during the life of the option. Moreover, this lower bound is shown to be an absorbing barrier for the forward yield. This explains our finding in Section 2.5 – a positive price for the Arrow-Debreu security that pays one unit if and only if the forward yield at the exercise date is zero.

The rest of the chapter is organised as follows. After introducing our setup, Section 4.1 states and interprets the main result. The mathematical theorem which underlies this result is presented in Section 4.2. Section 4.3 applies our approach to pricing formulae for options on zero-coupon bonds. Section 4.4 concludes the chapter.

4.1 The Main Result

We fix a finite time interval $\mathcal{T} = [0, \tau]$, a probability space $(\Omega, \mathcal{F}, P)$ and a filtration $(\mathcal{F}_t)_{t \in \mathcal{T}}$ satisfying the usual conditions. $\mathcal{F}_0$ is assumed to be almost trivial, and $\mathcal{F}_\tau = \mathcal{F}$.

Consider three securities, labelled 0, 1 and 2. We make the following assumptions:

- Trade in these securities is continuous and frictionless.
- The securities pay no dividends.
- Security 0 has a constant price $X_t^0 \equiv 1$.
- The price processes of securities 1 and 2, denoted by X^1 and X^2, are positive continuous semimartingales.

Security 0 can be thought of as a riskless cash account with zero interest. Alternatively, we can interpret (X^0, X^1, X^2) as a normalised price system, expressed in units of some numeraire asset.

Note that we do not assume that the filtration $(\mathcal{F}_t)_{t \in \mathcal{T}}$ is generated by the price processes of assets 0, 1 and 2. We allow the filtration to contain more information than just past prices of these three assets. This additional information could be the price history of other securities or, more generally, non-price information about arbitrary economic variables. For later use, let $(\mathcal{G}_t^i)_{t \in \mathcal{T}}$ be the completion of the filtration generated by X^i, and write $\mathcal{G}^i = \mathcal{G}_\tau^i$.

A *martingale measure* for the price system (X^0, X^1, X^2) is defined as a probability measure Q equivalent to P such that both processes X^i $(i = 1, 2)$ are Q-martingales. As seen in Section 3.1, the existence of a martingale measure ensures absence of arbitrage opportunities in a suitably chosen space of admissible trading strategies. Such a measure, if it exists, is in general not unique.[2]

We say that the price of asset 2 is given by a *pricing formula* if there is some function $u(t, x)$ such that

$$X_t^2 = u(t, X_t^1)$$

for all $t \in \mathcal{T}$. The literature on the valuation of derivative assets has calculated pricing formulae for a variety of securities. Adopting for a moment the perspective of derivative assets analysis, think of assets 0 and 1 as primitive securities, and of asset 2 as a derivative with payoff depending on the price of asset 1 at the terminal date. Given the price processes of the primitive assets, the task is to determine the fair price of asset 2. Typically, this involves the following steps.[3] First, one establishes the existence of a martingale measure for the system (X^0, X^1) of primitive asset prices. Next, one proves that the derivative claim is attainable, i.e., that it can be replicated by a dynamically adjusted self-financing trading strategy in the primitive assets. The price of the derivative asset must then be equal to the value of the replicating portfolio; any deviation would lead to arbitrage opportunities. Moreover, the price of the derivative is again a martingale under the given mar-

[2]Uniqueness of the martingale measure corresponds to completeness of the securities market. See Harrison and Pliska (1983) or Müller (1985).

[3]We have in fact gone through these steps in Chapter 3 already.

tingale measure, so it can be calculated without reference to a replicating strategy, just by taking expectations of the final payoff under the martingale measure. Finally, if the primitive asset prices have the Markov property, then the solution of the valuation problem indeed takes the form of a pricing formula.

Assume for example that (X^0, X^1) has a martingale measure Q such that X^1 solves the stochastic differential equation

$$dX_t^1 = \sigma(t, X_t^1)\, dW_t \tag{SDE}$$

with $\sigma(t, x)$ sufficiently regular and W a Wiener process under Q. Then one has the following well-known result.[4] Asset 2 is attainable, and its price process is of the form $X_t^2 = u(t, X_t^1)$ with $u(t, x)$ being a solution of the partial differential equation

$$u_t + \tfrac{1}{2}\sigma^2 u_{xx} = 0. \tag{PDE}$$

Thus, if X^1 has a martingale measure under which it is a diffusion satisfying (SDE), we get pricing formulae for derivatives involving solutions to the valuation equation (PDE). Our aim is to prove a converse to this statement.

Returning to the general setup, let us assume that (X^0, X^1, X^2) has a martingale measure, and let the price of asset 2 be given by a pricing formula $X_t^2 = u(t, X_t^1)$ where u is once continuously differentiable with respect to t and twice with respect to x. Fix a time t and a realisation x of the random variable X_t^1. Suppose that $u_{xx}(t, x) > 0$, say, so u is strictly convex in its second argument around (t, x). By Jensen's inequality, the holder of asset 2 can therefore expect a gain from the random movements of X^1 over a short time interval. The existence of a martingale measure, however, precludes such a gain. To balance the Jensen effect, the passing of time must therefore have a tendency to reduce the value of asset 2, in other words, $u_t(t, x) < 0$. By the same argument, $u_{xx}(t, x) < 0$ implies $u_t(t, x) > 0$.[5]

[4] Lemma 2.2.2 is a special case of this result.

[5] A mathematically precise argument runs as follows. For $T > t$, Itô's lemma implies

$$X_T^2 - X_t^2 - \int_t^T u_x(s, X_s^1)\, dX_s^1 = \int_t^T u_t(s, X_s^1)\, ds + \tfrac{1}{2} \int_t^T u_{xx}(s, X_s^1)\, d\langle X^1 \rangle_s\,.$$

Under a martingale measure, the left hand side is a continuous local martingale while the right hand side is of finite variation, so both must vanish identically. This requires $u_t(t, X_t^1)$ and $u_{xx}(t, X_t^1)$ to be of opposite sign whenever the latter expression is non-zero.

Provided $u_{xx}(t,x) \neq 0$, we can now define

$$\sigma(t,x) = \sqrt{-\frac{2\,u_t(t,x)}{u_{xx}(t,x)}}$$

and thereby satisfy (PDE) at the given point (t,x). In this sense, (PDE) is just a consequence of a simple "no expected gain" argument, and does not impose restrictions on the underlying process X^1. In the theorem below, we shall therefore make the additional assumption that the above function $\sigma(t,x)$ which we defined point by point on a subset of the domain of u has in fact a *continuous* extension to the whole of that domain.

For a similar reason, we shall also stipulate that u be sufficiently non-linear, i.e., that u_{xx} does not vanish too often. Clearly, a linear pricing formula will not restrict the underlying process at all – a statement like "two shares cost twice the price of one share" will not tell us anything about the underlying stock price model.

We are now ready to formulate the main result of this chapter.

Theorem 4.1.1 *Let the price system* (X^0, X^1, X^2) *on* $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t\in\mathcal{T}})$ *satisfy*

$$X_t^2 = u(t, X_t^1)$$

where $u(t,x)$ *is a solution of* (PDE) *with a continuous function* $\sigma(t,x)$. *Assume that* $\{t \in \mathcal{T} : \sigma(t, X_t^1) = 0\}$ *has Lebesgue measure zero almost surely.*[6] *In addition, suppose that at least one of the following two conditions holds:*

$$u_{xx}(t, X_t^1) \neq 0 \textit{ for all } t \in \mathcal{T} \textit{ almost surely,}$$

or

$$\{t \in \mathcal{T} : u_{xx}(t, X_t^1) = 0\} \textit{ is almost surely a Lebesgue null set}$$
$$\textit{and } u(t,x) \textit{ is analytic.}$$

Finally, let Q *be a martingale measure for this price system. Then there is a Wiener process* W *under* Q *such that the price process* X^1 *satisfies* (SDE).

[6] This assumption is inessential. See Remark 4.2.5 below.

PROOF: The theorem is a direct consequence of Theorem 4.2.1 below, a somewhat more general mathematical result which we will prove in the following section. ■

According to Theorem 4.1.1, a pricing formula satisfying (PDE) under the stated conditions completely characterises the behaviour of the price of asset 1 under the martingale measure Q.[7] Indeed, (SDE) together with the fact that W is a Wiener process completely determine the law of X^1 under Q. As a first consequence, note that the pricing formula implies the Markov property for X^1 under Q.[8] (PDE) is then just the associated backward equation.

More important, the law of X^1 is the same under *all* martingale measures. In other words, all martingale measures coincide on $\mathcal{G}^1$. By a theorem of Jacka (1992), this implies that all $\mathcal{G}^1$-measurable contingent claims are attainable, hence priced by arbitrage.[9] This holds in particular for the Arrow-Debreu security with time τ payoff 1_A where $A \in \mathcal{G}^1$. The pricing formula thus implies a unique system of Arrow-Debreu or state prices for events in $\mathcal{G}^1$. As usual, these prices are obtained by taking the expectation of the corresponding Arrow-Debreu payoffs under any martingale measure.

The idea of extracting state prices from derivative prices goes back at least to Breeden and Litzenberger (1978).[10] In the present setting, their argument can be rendered as follows.[11] Assume that we have a securities market with assets 0 and 1 as before but, instead of asset 2, European call options written on asset 1 for any strike price and exercise date. Let $C_0^{\tau,K}$ denote today's call price for exercise date τ and strike price K. Assume that there exists a martingale measure for this securities market, and let F_τ be the corresponding distribution function for the random variable X_τ^1. Call prices must satisfy

$$C_0^{\tau,K} = \int_K^\infty (x - K)\, dF_\tau(x)$$

[7]We assume in the following that σ satisfies the regularity conditions for uniqueness of weak solutions of (SDE). See for example Karatzas and Shreve (1988).

[8]Recall that we merely assumed this price to be a continuous semimartingale.

[9]Alternatively, this result follows from the martingale representation property of X^1 with respect to the smaller filtration $(\mathcal{G}_t^1)_{t\in\mathcal{T}}$; cf. Revuz and Yor (1991).

[10]See also Ross (1976).

[11]We went through a variant of this argument in Section 2.5 where we started from put rather than call prices.

by definition of a martingale measure. Integration by parts yields

$$C_0^{\tau,K} = \int_K^\infty (1 - F_\tau(x))\, dx$$

and hence

$$F_\tau(K) = 1 + \frac{\partial}{\partial K} C_0^{\tau,K}.$$

Thus the distribution function F_τ is uniquely determined by the given option prices.[12]

Dupire (1994) takes this analysis one step further. He assumes the existence of a martingale measure under which X^1 is a diffusion process satisfying (SDE) for some unknown function $\sigma(t,x)$. Using the forward equation associated with such a diffusion, he shows that given the call price $C_0^{\tau,K}$ for all τ and K, it is possible, under certain regularity conditions, to back out the function σ from the distribution functions F_τ. Therefore, the law of the process is completely determined by these call prices, and we have again a unique set of Arrow-Debreu prices for events in $\mathcal{G}^1$.[13]

Our approach, as summarised in Theorem 4.1.1, and Dupire's approach can be regarded as "dual" to each other. This feature appears most clearly in the analysis of call option prices. Suppose that time t call prices are given by some function $u(t,x;\tau,K)$ where x is the concurrent price of the underlying asset, τ the exercise date, and K the exercise price. Dupire's result means that a unique set of state prices can be extracted from the values $u(0,x_0;\tau,K)$ where the initial price of the underlying asset is fixed, while τ and K are variable. Theorem 4.1.1, on the other hand, determines state prices on the basis of the values $u(t,x;\overline{\tau},\overline{K})$ for fixed option characteristics, but variable t and x.[14] Thus, Dupire's result and Theorem 4.1.1 are "dual" in the sense that the former varies the "forward variables" (τ, K), and the latter the "backward variables" (t,x).

[12]The value $F_\tau(K)$ is the price of the Arrow-Debreu claim $1_{X_\tau^1 \leq K}$. Differentiating once more, where possible, we get the *state price density* $f_\tau(K) = \frac{\partial^2}{\partial K^2} C_0^{\tau,K}$. This is Breeden and Litzenberger's original result.

[13]Dupire's work is one of the first contributions to a recent literature on "implied trees"; see Rubinstein (1994) and the references therein. This literature tries to construct models which, in contrast to the Black-Scholes model, are consistent with the market prices of standard European options. These models are then used to hedge and price "exotic" over-the-counter derivatives.

[14]The choice of $\overline{\tau}$ and $\overline{K}$ is irrelevant, of course, since Theorem 4.1.1 does not depend on the particular form of the derivative's terminal value.

4.2 A Characterisation Theorem

In this section, we state and prove the mathematical result which underlies Theorem 4.1.1. We extend the work of McGill, Rajeev and Rao (1988) on Brownian motion to a larger class of continuous local martingales.[15] Throughout the section, we consider a finite time interval $\mathcal{T}$ as before and a filtered probability space satisfying the usual conditions.

Theorem 4.2.1 *Let X_t be a continuous local martingale which fulfils the following condition.*

> (A) *There exist a continuous function $\sigma(t,x)$ and a solution $u(t,x)$ of*
>
> $$u_t + \tfrac{1}{2}\sigma^2 u_{xx} = 0 \tag{PDE}$$
>
> *such that*
>
> > (A1) *the process $u(t,X_t)$ is a local martingale;*
> >
> > (A2) *$\{t \in \mathcal{T} : \sigma(t,X_t) = 0\}$ is almost surely a Lebesgue null set;*
> >
> > (A3) *$u_{xx}(t,X_t) \neq 0$ for all $t \in \mathcal{T}$ almost surely.*

Then there exists a Wiener process W such that

$$X_t = X_0 + \int_0^t \sigma(s,X_s)\,dW_s. \tag{4.1}$$

Moreover, this continues to hold if (A3) *is replaced with the two conditions*

> (A4) *$\{t \in \mathcal{T} : u_{xx}(t,X_t) = 0\}$ is almost surely a Lebesgue null set;*
>
> (A5) *$u(t,x)$ is analytic.*

Remark 4.2.1 Note that if (PDE) holds, conditions (A2) and (A4) together are equivalent to the condition that $\{t \in \mathcal{T} : u_t(t,X_t) = 0\}$ is almost surely a Lebesgue null set. This is the condition used in McGill, Rajeev and Rao (1988).

Remark 4.2.2 McGill, Rajeev and Rao (1988) study the case $\sigma(t,x) \equiv 1$ with infinite time horizon, i.e., $\mathcal{T} = \mathbb{R}_+$. In this case, (PDE) is just the heat equation,

[15] I am grateful to Lucien Foldes for having drawn my attention to McGill, Rajeev and Rao (1988) after I had obtained a weaker version of Theorem 4.2.1 independently.

(A5) is automatic, and a continuous local martingale satisfying condition (A) or its variant is a Brownian motion in accordance with (4.1). Lévy's characterisation of Brownian motion is recovered as the special case where the solution of the heat equation is taken to be $u(t,x) = x^2 - t$.

Remark 4.2.3 To obtain Theorem 4.1.1, let Q be a martingale measure for the price system (X^0, X^1, X^2) and apply Theorem 4.2.1 to the martingales X^1 and $u(t, X_t^1) = X_t^2$ on $(\Omega, \mathcal{F}, Q, (\mathcal{F}_t)_{t\in\mathcal{T}})$.

Remark 4.2.4 Obviously, Theorem 4.1.1 holds as well for the larger set of *local martingale measures*, that is, for measures Q equivalent to P such that the price processes X^i are local martingales. Such measures have been studied for example by Schweizer (1992) and Babbs and Selby (1993).

The proof of Theorem 4.2.1 is given in a sequence of lemmata.

Lemma 4.2.1 *Let X_t be a continuous local martingale with quadratic variation process*

$$\langle X\rangle_t = \int_0^t \sigma^2(s, X_s)\, ds. \tag{4.2}$$

Assume that (A2) *holds. Then there exists a Wiener process W satisfying* (4.1).

PROOF: If (4.2) holds, we can define a process W by setting

$$W_t = \int_0^t \phi_s \, dX_s$$

where $\phi_s = \sigma(s, X_s)^{-1}$ if $\sigma(s, X_s) \neq 0$, and $\phi_s = 0$ otherwise. W satisfies (4.1) and has quadratic variation $\langle W\rangle_t = \int_0^t 1_{\{\sigma(s,X_s)\neq 0\}}\, ds$. (A2) implies $\langle W\rangle_t = t$, and the assertion follows from Lévy's characterisation theorem. ■

Remark 4.2.5 If (A2) is not satisfied, i.e., if $\{t \in \mathcal{T} : \sigma(t, X_t) = 0\}$ is not a null set, (4.2) still implies the representation (4.1). However, W is then no longer a Wiener process on the original filtered probability space, but on an extension of it. See Ikeda and Watanabe (1989, Theorem 7.1′ on page 90) for details.

Lemma 4.2.2 *Let X_t be a continuous local martingale. If there exist a continuous function $\sigma(t,x)$ and a solution $u(t,x)$ of* (PDE) *such that* (A1) *and* (A3) *are fulfilled, then X_t has quadratic variation given by* (4.2).

PROOF: Itô's rule, (PDE) and (A1) imply

$$\int_0^t u_{xx}(s,X_s)(d\langle X\rangle_s - \sigma^2(s,X_s)\,ds) = 0.$$

(4.2) follows by (A3). ■

This completes the proof of Theorem 4.2.1 for condition (A). The case where we replace (A3) by (A4) and (A5) is covered in the following lemma. Its proof builds on the arguments in McGill, Rajeev and Rao (1988).

Lemma 4.2.3 *Let X_t be a continuous local martingale. If there exist functions $\sigma(t,x)$ and $u(t,x)$ such that* (PDE), (A1), (A4) *and* (A5) *are fulfilled, then the quadratic variation process of X_t satisfies* (4.2).

PROOF: We start from the obvious equation

$$\begin{aligned}\langle X\rangle_t &= \int_0^t \sigma^2(s,X_s)\,ds \\ &\quad + \int_0^t 1_{\{u_{xx}(s,X_s)\neq 0\}}\,[d\langle X\rangle_s - \sigma^2(s,X_s)\,ds] \\ &\quad + \int_0^t 1_{\{u_{xx}(s,X_s)=0\}}\,[d\langle X\rangle_s - \sigma^2(s,X_s)\,ds].\end{aligned}$$

As in the proof of the previous lemma, one obtains

$$\int_0^t u_{xx}(s,X_s)\,[d\langle X\rangle_s - \sigma^2(s,X_s)\,ds] = 0$$

and hence

$$\int_0^t 1_{\{u_{xx}(s,X_s)\neq 0\}}\,[d\langle X\rangle_s - \sigma^2(s,X_s)\,ds] = 0.$$

On the other hand, (A4) implies

$$\int_0^t 1_{\{u_{xx}(s,X_s)=0\}}\;\sigma^2(s,X_s)\,ds = 0.$$

Thus

$$\langle X\rangle_t = \int_0^t \sigma^2(s,X_s)\,ds + \int_0^t 1_{\{u_{xx}(s,X_s)=0\}}\;d\langle X\rangle_s,$$

and (4.2) holds if

$$\int_0^T 1_{\{u_{xx}(s,X_s)=0\}}\;d\langle X\rangle_s = 0. \tag{4.3}$$

Consider a new time variable $\xi \geq 0$ and set $t_\xi = \inf\{t \in \mathcal{T} : \langle X \rangle_t > \xi\}$ if this set is not empty, and $t_\xi = \tau$ otherwise. Then, after extending the filtered probability space, there is a Brownian motion $(B_\xi)_{\xi \geq 0}$ such that $X_t = B_{\langle X \rangle_t}$; see Ikeda and Watanabe (1989, Theorem 7.2′ on page 91). Write $\overline{\xi} = \langle X \rangle_\tau$. (4.3) is equivalent to $\{\xi \leq \overline{\xi} : \ u_{xx}(t_\xi, B_\xi) = 0\}$ being a Lebesgue null set. Consider now the stopped continuous semimartingale $Y_\xi = u_{xx}(t_\xi, B_{\xi \wedge \overline{\xi}})$. As a direct consequence of the occupation density formula for semimartingale local time, we have $\int 1_{\{Y_\xi = 0\}} \, d\langle Y \rangle_\xi = 0$ and hence

$$\int_0^{\overline{\xi}} 1_{\{u_{xx}(t_\xi, B_\xi) = 0\}} \, u_{xxx}^2(t_\xi, B_\xi) \, d\xi = 0$$

which means

$$\{\xi \leq \overline{\xi} : \ u_{xx}(t_\xi, B_\xi) = 0\} \subseteq \{\xi \leq \overline{\xi} : \ u_{xxx}(t_\xi, B_\xi) = 0\}$$

up to a Lebesgue null set. Assuming that (4.3) does not hold and arguing inductively, one shows that there exists at least one point (t_0, x_0) where u_{xx} and all its space derivatives vanish. Next, using (PDE) and another induction argument, one can easily show that *all* partial derivatives of u_{xx} vanish at (t_0, x_0). But then, due to the analyticity of u_{xx} postulated in (A5), condition (A4) is violated. Thus (4.3) must hold. ∎

4.3 Bond Options and Implied Forward Yields

In this section, we use our results to analyse pricing formulae for European options. As in Chapter 2, we study the case of options on zero-coupon bonds.

We fix an interval $\mathcal{T} = [0, \tau]$ and a filtered probability space $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t \in \mathcal{T}})$ as in Section 4.1. Let S^0 be the price process of a default-free discount bond maturing at τ, i.e., satisfying $S^0_\tau = 1$ almost surely. This bond, which we call the reference bond, will serve as numeraire. Consider a second bond, called the underlying bond, that matures at a time $T > \tau$. Let S^1 be its price process up to time τ. The third security is a European call option written on the underlying bond with exercise date τ and strike price K. Its price process is denoted by S^2. By definition, the terminal value of the option is $S^2_\tau = [S^1_\tau - K]^+$. We make the following assumptions:

- Trade in the bonds and the option is continuous and frictionless.
- The price processes of the bonds and the option are positive continuous semimartingales.

In order to obtain the setting studied in Section 4.1, we define the processes

$$X^i = \frac{S^i}{S^0} \qquad (i = 0,\, 1,\, 2)$$

which describe asset prices in units of the reference bond. In less abstract terms, X_t^1 and X_t^2 are just the time t forward prices of the underlying bond and the option for delivery at τ. These processes are again positive continuous semimartingales.

Recall from Section 2.1 that the forward yield $Y_{t,\tau,T}$ implied by the bond prices S_t^0 and S_t^1 is

$$Y_{t,\tau,T} = -\frac{\log S_t^1 - \log S_t^0}{T-\tau} = -\frac{\log X_t^1}{T-\tau}.$$

This is the continuously compounded interest rate as seen at time t for a loan which starts at τ and is repaid at T. We shall write Y_t for $Y_{t,\tau,T}$. The bond price model (S^0, S^1) is said to *generate negative forward yields* if

$$P\left(\{\omega \in \Omega : \; \exists t \in \mathcal{T} \;\; Y_t(\omega) < 0\}\right) > 0;$$

otherwise, the bond price model satisfies

$$P\left(\{\omega \in \Omega : \; \forall t \in \mathcal{T} \;\; Y_t(\omega) \geq 0\}\right) = 1$$

and is said to have *non-negative forward yields.* Finally, we say that the bond price model has *positive forward yields* if

$$P\left(\{\omega \in \Omega : \; \forall t \in \mathcal{T} \;\; Y_t(\omega) > 0\}\right) = 1.$$

4.3.1 Merton Type Option Prices

This section deals with the type of valuation formulae going back to Black and Scholes (1973) and Merton (1973). Let a positive continuous function $\nu : \mathcal{T} \to \mathbb{R}_{++}$ be given. Extending the terminology used in Chapter 2, we say that the price system (S^0, S^1, S^2_t) satisfies the *Merton call price formula for volatility function* ν if

$$S^2_t = S^0_t \, u(t, X^1_t)$$

or, equivalently,

$$X^2_t = u(t, X^1_t)$$

with u defined by equation (2.15):

$$u(t,x) = x\,\Phi\left(\frac{1}{\sqrt{s(t)}}\left[\log\frac{x}{K} + \frac{s(t)}{2}\right]\right) - K\,\Phi\left(\frac{1}{\sqrt{s(t)}}\left[\log\frac{x}{K} - \frac{s(t)}{2}\right]\right)$$

where

$$s(t) = \int_t^\tau \nu^2(\xi)\,d\xi.$$

Several authors have obtained pricing relationships of this type for options on zero-coupon bonds, either in the context of a term structure model that specifies the evolution of bond prices for a continuum of maturities,[16] or in the "minimalist" framework of Chapter 2 where only the prices of the underlying bond and the reference bond are modelled.

The common feature of all these models is that they allow for negative interest rates. A direct application of Theorem 4.1.1 confirms that option price formulae of the Merton type are indeed inconsistent with non-negative interest rates.

Proposition 4.3.1 *Assume that the price system (S^0, S^1, S^2) on $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t\in\mathcal{T}})$ satisfies the Merton call price formula for volatility function ν. Let Q be a martingale measure for this price system. Then the forward price of the underlying bond solves the stochastic differential equation*

$$dX^1_t = \nu(t)\, X^1_t \, dW_t \tag{4.4}$$

where W is a Wiener process under the measure Q.

PROOF: Recall from Section 2.2 that u solves (PDE) with $\sigma(t,x) = \nu(t)\,x$. Moreover, u is strictly convex in x for all $t < \tau$. So Theorem 4.1.1 applies. ■

[16] See the references in the footnote at the end of Section 2.2.

By the formula for the martingale exponential, (4.4) is equivalent to

$$X_t^1 = X_0^1 \exp\left(\int_0^t \nu(s)\, dW_s - \tfrac{1}{2}\int_0^t \nu^2(s)\, ds\right). \tag{4.5}$$

Thus, up to the time change

$$t \mapsto \int_0^t \nu^2(s)\, ds,$$

the forward price process X^1 is a driftless geometric Brownian motion under any martingale measure Q, and the forward yield Y is simply a Brownian motion with drift. This implies in particular that Q and, by equivalence of measures, P assign a positive probability to the event $\{\omega : \exists t\ Y_t(\omega) < 0\}$. Thus, we obtain the well-known

Result 4.3.1 *A bond price model in which an option formula of the Merton type holds necessarily generates negative forward yields.*

4.3.2 An Upper Bound on the Forward Bond Price

Next consider a price system (S^0, S^1, S^2) where forward yields remain non-negative, i.e., $Y_t \geq 0$ and $X_t^1 \leq 1$ for all t. Assume that the strike price of the call option satisfies $0 < K < 1$; as $S_\tau^1 = X_\tau^1 \leq 1$, only these exercise prices are of interest. Schöbel's (1986) condition (2.7) now reads as follows:

$$S_t^2 = (1-K)\, S_t^0 \quad \text{whenever} \quad S_t^1 = S_t^0,$$

that is,

$$X_t^2 = 1 - K \quad \text{whenever} \quad X_t^1 = 1.$$

Thus, the forward call price assumes the deterministic value $1-K$ when the forward yield Y_t is at its lower bound 0.

We know that Merton call prices violate this condition. Schöbel proposes the modified pricing formula

$$X^2 = u^*(t, X_t^1)$$

with

$$u^*(t,x;K) = u(t,x;K) - K\, u(t,x;K^{-1})$$

where $u(t, x; K)$ denotes the Merton call price function for strike price K.[17]

Briys, Crouhy and Schöbel (1991) use a formula of this type to value interest rate caps and floors. They see the second term in u^* as a price correction which ensures consistency of bond option prices with non-negative interest rates. Moreover, they interpret the additional boundary condition as the effect of an absorbing barrier, but do not clarify the nature of the absorption phenomenon.

Following a different approach, Sondermann (1988) obtains an option price formula of the same type. Starting with the discrete-time binomial approximation to the Black-Scholes model, he imposes an upper bound (which may be any positive number) on the forward price of the underlying asset. He shows the existence and uniqueness of a martingale measure for the bounded process and calculates the corresponding option price. On letting the grid size go to zero, Sondermann obtains essentially the above pricing formula for constant volatility and an arbitrary upper bound. He notes that in his discrete-time approximations, the martingale measure makes the upper bound an absorbing barrier for the forward price process. Interested mainly in the limit of the valuation formula as the grid size tends to zero, he does not study the limit forward price process itself. It seems obvious, though, that this limit should be the Black-Scholes forward price process with an absorbing boundary.

We shall use the method developed in the previous section to confirm this intuition. Focussing on bond options, we present only the case of an upper bound equal to 1, but the results carry over to the setting of Sondermann (1988). Before applying the technique underlying Theorems 4.1.1 and 4.2.1, let us first point out that absorption of the forward price and the forward yield at their respective boundaries is indeed the only behaviour consistent with the absence of arbitrage.

Lemma 4.3.1 *Let (S^0, S^1) be a bond price model on $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t \in \mathcal{T}})$ with a martingale measure Q. Assume that the model has non-negative forward yields and consider the hitting time $\chi = \inf\{t \in \mathcal{T} : X_t^1 = 1\}$. Then $X_t^1 = 1$ on $[\chi, \tau]$ almost surely under either measure.*

[17] See Section 2.5.

PROOF: Apply Theorem 4.16 of Elliott (1982) to the Q-martingale $1 - X^1$. ■

There is a simple intuition behind Lemma 4.3.1. Assume that in a world of non-negative forward yields, the underlying bond and the reference bond have the same price at some time t_0. A portfolio short one underlying bond and long one reference bond costs nothing at t_0. After t_0, the portfolio value cannot fall below zero since the underlying bond will never cost more than the reference bond. On the other hand, the portfolio cannot rise in value either, otherwise it would certainly trade at a positive price now. Therefore, the two bond prices must coincide for ever, that is, until the shorter lived bond expires. By the same token, forward bond prices and forward yields are absorbed at their upper and lower bound, respectively. Any other boundary behaviour, for example reflection, would lead to arbitrage opportunities.

As a corollary, we get the following simple classification.

Proposition 4.3.2 *Let (S^0, S^1) be a bond price model admitting a martingale measure. Then exactly one of the following statements holds true:*

(i) *The model generates negative forward yields.*

(ii) *The model has non-negative forward yields, the probability that the forward yield reaches its lower bound 0 is positive, and 0 is an absorbing barrier for the forward yield.*

(iii) *The model has positive forward yields.*

We have seen that bond price models consistent with a Merton type formula belong to category (i). As for models with non-negative yields in which an option formula of the Schöbel type holds, we have to establish which of the two properties (ii) and (iii) is satisfied, that is, whether the bound is reached with positive probability or not. The following proposition does more than that: it gives a complete description of the behaviour of the forward bond price under a martingale measure.

Proposition 4.3.3 *Assume that the price system (S^0, S^1, S^2) on $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{t \in \mathcal{T}})$ has non-negative forward yields and satisfies the Schöbel call price formula for volatility function ν. Let Q be a martingale measure for this price system. Then* *there is a Wiener process $\overline{W}$ on an extension $(\overline{\Omega}, \overline{\mathcal{F}}, \overline{Q}, (\overline{\mathcal{F}}_t)_{t\in\mathcal{T}})$ of $(\Omega, \mathcal{F}, Q, (\mathcal{F}_t)_{t\in\mathcal{T}})$ such that the forward price of the underlying bond satisfies*

$$dX_t^1 = 1_{\{t\le\chi\}}\, \nu(t)\, X_t^1\, d\overline{W}_t \tag{4.6}$$

with $\chi = \inf\{t \in \mathcal{T} :\ X_t^1 = 1\}$.

PROOF: Note the following properties of the Schöbel call price function: u^* solves (PDE) with $\sigma(t,x) = \nu(t)\,x$, is strictly convex in x for $x < 1$ and satisfies $u_t^*(t,1) = u_{xx}^*(t,1) = 0$ for all t. Using these properties and Lemma 4.3.1, one shows as in the proof of Lemma 4.2.2 that

$$\langle X^1\rangle_t = \int_0^t 1_{\{s\le\chi\}}\, \nu^2(s)\, (X_s^1)^2\, ds.$$

The proposition now follows directly from Ikeda and Watanabe (1989, Theorem 7.1′ on page 90). The process $\overline{W}$ is constructed as

$$\overline{W}_t = \int_0^t \frac{1_{\{s\le\chi\}}}{\nu(s)\, X_s^1}\, dX_s^1 + \int_0^t 1_{\{s>\chi\}}\, dW_s'$$

where W' is a Wiener process on some filtered probability space $(\Omega', \mathcal{F}', Q', (\mathcal{F}_t')_{t\in\mathcal{T}})$. The extension $(\overline{\Omega}, \overline{\mathcal{F}}, \overline{Q}, (\overline{\mathcal{F}}_t)_{t\in\mathcal{T}})$ is obtained by taking the products $\overline{\Omega} = \Omega \times \Omega'$, $\overline{Q} = Q \otimes Q'$ and $\overline{\mathcal{F}}_t = \mathcal{F}_t \otimes \mathcal{F}_t'$. ■

Let Q and $\overline{W}$ be as in the proposition. By the formula for the martingale exponential, (4.6) implies that X^1 is the stopped process

$$X_t^1 = \overline{X}_{t\wedge\chi}$$

with

$$\overline{X}_t = X_0^1 \exp\left(\int_0^t \nu(s)\, d\overline{W}_s - \tfrac{1}{2}\int_0^t \nu^2(s)\, ds\right).$$

Thus, the forward bond price process implied by Schöbel's option price formula is obtained by imposing an absorbing barrier at 1 on $\overline{X}$, a forward price process of the Merton type. $\overline{Q}$ assigns positive probability to the event that $\overline{X}_t = 1$ for some $t \in \mathcal{T}$. Under both Q and P, the forward bond price therefore reaches its upper boundary with positive probability.

Result 4.3.2 *A bond price model with non-negative forward yields which satisfies an option price formula of the Schöbel type assigns positive probability to the event that the forward yield reaches its lower bound 0, where it is absorbed.*

In particular, the Arrow-Debreu security paying one unit if and only if this yield is zero at the expiry date of the option commands a positive price.

4.4 Conclusion

The validity of a pricing formula for a derivative asset has strong implications for the risk-neutral behaviour of the underlying asset prices. In a simple setting with continuous asset price processes, we have studied valuation formulae that depend on only one underlying price and satisfy a certain partial differential equation. We have shown that such a formula implies a complete characterisation of the behaviour of the underlying asset price under a martingale measure. This characterisation takes the form of a stochastic differential equation. The law of the underlying price process under a martingale measure is completely determined by the pricing formula, and is the same for all martingale measures. Consequently, all claims contingent on the price path of the underlying asset are attainable and hence priced by arbitrage. In particular, the pricing formula implies a unique set of Arrow-Debreu prices for events which are determined by the price path of the underlying asset.

As an illustration of our main result, we have analysed certain pricing formulae for European options on discount bonds. This analysis has shown that the modified valuation formulae proposed by Schöbel (1986) and Briys, Crouhy and Schöbel (1991) imply an implausible behaviour of the forward yield, involving absorption of this yield at zero.

References

AGHION, P., P. BOLTON, C. HARRIS, and B. JULLIEN (1991): "Optimal Learning by Experimentation," *Review of Economic Studies*, 58, 621-654.

BABBS, S.H., and M.J.P.SELBY (1993): "Pricing by Arbitrage in Incomplete Markets," preprint, Financial Options Research Center, University of Warwick.

BACK, K., and S.R. PLISKA (1991): "On the Fundamental Theorem of Asset Pricing with an Infinite State Space," *Journal of Mathematical Economics*, 20, 1-18.

BALL, C.A., and W.N. TOROUS (1983): "Bond Price Dynamics and Options," *Journal of Financial and Quantitative Analysis*, 18, 517-531.

BERNFELD, S.R., and V. LAKSHMIKANTHAM (1974): *An Introduction to Nonlinear Boundary Value Problems.* New York and London: Academic Press.

BLACK, F., and M. SCHOLES (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-654.

BOLTON, P., and C. HARRIS (1993): "Strategic Experimentation," STICERD Discussion Paper No. TE/93/261, London School of Economics.

BREEDEN, D., and R. LITZENBERGER (1978): "Prices of State Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621-651.

BRIYS, E., M. CROUHY, and R. SCHÖBEL (1991): "The Pricing of Default-Free Interest Rate Cap, Floor, and Collar Agreements," *Journal of Finance*, 46, 1879-1892.

BÜHLER, W. (1988):"Rationale Bewertung von Optionsrechten auf Anleihen," *Zeitschrift für betriebswirtschaftliche Forschung*, 10, 851-883.

BÜHLER, W., and J. KÄSLER (1989): "Konsistente Anleihepreise und Optionen auf Anleihen," preprint, Universität Dortmund, Germany.

CHENG, S. (1991):"On the Feasibility of Arbitrage-Based Option Pricing when Stochastic Bond Price Processes are Involved," *Journal of Economic Theory*, 53, 185-198.

CHEYETTE, O. (1994): "Markov Representation of the Heath-Jarrow-Morton Model," preprint, BARRA Inc., Berkeley.

CONZE, A., and R. VISWANATHAN (1991): "Probability Measures and Numeraires," preprint, Universities Paris Dauphine and Paris VI.

COX, J.C., J.E. INGERSOLL, and S.A. ROSS (1979): "Duration and the Measurement of Basic Risk," *Journal of Business*, 52, 51-61.

COX, J.C., J.E. INGERSOLL, and S.A. ROSS (1985): "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 384-408.

DUFFIE, D. (1992): *Dynamic Asset Pricing Theory.* Princeton, New Jersey: Princeton University Press.

DUFFIE, D., and R. KAN (1992): "A Yield-Factor Model of Interest Rates," preprint, Graduate School of Business, Stanford University.

DUPIRE, B. (1994):"Pricing with a Smile," *RISK*, 7, January, 18-20.

DYBVIG, P.H., and C. HUANG (1988): "Nonnegative Wealth, Absence of Arbitrage, and Feasible Consumption Plans," *Review of Financial Studies*, 1, 377-401.

EASLEY, D., and N.M. KIEFER (1988): "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56, 1045-1064.

EL KAROUI, N., and J.-C. ROCHET (1989): "A Pricing Formula for Options on Coupon Bonds," SEEDS Working Paper 72.

EL KAROUI, N., R. MYNENI, and R. VISWANATHAN (1992): "Arbitrage Pricing and Hedging of Interest Rate Claims with State Variables," preprint, Université de Paris VI and Stanford University.

ELLIOTT, R.J. (1982): *Stochastic Calculus and Applications.* New York: Springer-Verlag.

FELLI, L., and C. HARRIS (1994): "Job Matching, Learning and the Distribution of Surplus," STICERD Discussion Paper No. TE/94/277, London School of Economics.

FLEMING, W.H., and R.W. RISHEL (1975): *Deterministic and Stochastic Control.* New York: Springer-Verlag.

GEMAN, H., N. EL KAROUI, and J.-C. ROCHET (1991): "Changes of Numéraire, Changes of Probability Measure and Option Pricing," preprint, to appear in *Journal of Applied Probability*, June 1995.

GLEIT, A. (1978): "Valuation of General Contingent Claims: Existence, Uniqueness and Comparison of Solutions," *Journal of Financial Economics*, 6, 71-87.

HARRISON, J.M., and D.M. KREPS (1979): "Martingales and Arbitrage in Multiperiod Security Markets," *Journal of Economic Theory*, 20, 381-408.

HARRISON, J.M., and S.R. PLISKA (1981): "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and their Applications*, 11, 215-260.

HARRISON, J.M., and S.R. PLISKA (1983): "A Stochastic Calculus Model of Continuous Trading: Complete Markets," *Stochastic Processes and their Applications*, 15, 313-316.

HEATH, D., R. JARROW, and A. MORTON (1992): "Bond Pricing and the Term Structure of Interest Rates: A New Methodology For Contingent Claim Valuation," *Econometrica*, 60, 77-105.

HULL, J., and A. WHITE (1990): "Pricing Interest-Rate Derivative Securities," *Review of Financial Studies*, 3, 573-590.

IKEDA, N., and S. WATANABE (1989): *Stochastic Differential Equations and Diffusion Processes.* Amsterdam: North-Holland.

INGERSOLL, J.E. (1989a): "Contingent Foreign Exchange Contracts," preprint, Yale School of Management.

INGERSOLL, J.E. (1989b): "Valuing Foreign Exchange Derivatives with a Bounded Exchange Rate Process," preprint, Yale School of Management.

JACKA, S.D. (1992): "A Martingale Representation Result and an Application to Incomplete Financial Markets," *Mathematical Finance*, 2, 239-250.

JAMSHIDIAN, F. (1987): "Pricing of Contingent Claims in the One-Factor Term Structure Model," preprint, Merrill Lynch Capital Markets, New York.

JAMSHIDIAN, F. (1989): "An Exact Bond Option Formula," *Journal of Finance*, 44, 205-209.

JAMSHIDIAN, F. (1990): "The Preference-Free Determination of Bond and Option Prices from the Spot Interest Rate," *Advances in Futures and Options Research*, 4, 51-67.

JAMSHIDIAN, F. (1991): "Bond and Option Valuation in the Gaussian Interest Rate Model," *Research in Finance*, 9, 131-170.

JOHNSON, N.L. (1946): "System of Frequency Curves Generated by Methods of Translation," *Biometrica*, 149-176.

JOHNSON, N.L. (1949): "Bivariate Distribution Based on Simple Translation Systems," *Biometrica*, 297-304.

KARATZAS, I., and S. SHREVE (1988): *Brownian Motion and Stochastic Calculus.* New York: Springer-Verlag.

KARLIN, S., and H.M. TAYLOR (1981): *A Second Course in Stochastic Processes.* New York: Academic Press.

KÄSLER, J. (1991): "Optionen auf Anleihen," PhD thesis, Universität Dortmund, Germany.

KEMNA, A.G.Z., J.F.J. DE MUNNIK, and A.C.F. VORST (1989): "On Bond Price Models with a Time-Varying Drift Term," preprint, Erasmus Universiteit Rotterdam, The Netherlands.

KIEFER, N.M. (1989): "A Value Function Arising in the Economics of Information," *Journal of Economic Dynamics and Control*, 13, 201-223.

KIEFER, N.M. (1991): "A Dynamic Model of Optimal Learning with Obsolescence of Information," working paper, Cornell University.

KRYLOV, N.V. (1980): *Controlled Diffusion Processes.* New York: Springer-Verlag.

LIPTSER, R.S., and A.N. SHIRYAYEV (1977): *Statistics of Random Processes I.* New York: Springer-Verlag.

LO, A.W. (1986): "Statistical Tests of Contingent-Claim Asset-Pricing Models: A New Methodology," *Journal of Financial Economics*, 17, 143-173.

LO, A.W. (1988): "Maximum Likelihood Estimation of Generalized Itô Processes with Discretely Sampled Data," *Econometric Theory*, 4, 231-247.

MCGILL, P., B. RAJEEV, and B.V. RAO (1988): "Extending Lévy's Characterisation of Brownian Motion," in *Séminaire de Probabilité XXII*, eds. J. Azéma, P.A. Meyer, M. Yor. New York: Springer-Verlag, 163-165.

MCLENNAN, A. (1984): "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 7, 331-347.

MERTON, R.C. (1973): "Theory of Rational Option Pricing," *Bell Journal of Economics and Management*, 4, 141-183.

MIRMAN, L.J., L. SAMUELSON, and A. URBANO (1993): "Monopoly Experimentation," *International Economic Review*, 34, 549-564.

MORTON, R. (1971): "On the Optimal Control of Stationary Diffusion Processes with Natural Boundaries and Discounted Costs," *Journal of Applied Probability*, 8, 561-572.

MÜLLER, S. (1985): *Arbitrage Pricing of Contingent Claims.* New York: Springer-Verlag.

DE MUNNIK, J.F.J. (1992): "The Valuation of Interest Rate Derivative Securities," PhD thesis, Tinbergen Institute Research Series No.40, The Netherlands.

PRESS, W.H., B.P. FLANNERY, S.A. TEUKOLSKY, and W.T. VETTERLING (1986): *Numerical Recipes: the Art of Scientific Computing.* Cambridge and New York: CUP.

RADY, S., and K. SANDMANN (1994): "The Direct Approach to Debt Option Pricing," *Review of Futures Markets*, 13(2), 461-514.

REVUZ, D., and M. YOR (1991): *Continuous Martingales and Brownian Motion.* New York: Springer-Verlag.

ROGERS, L.C.G., and D. WILLIAMS (1987): *Diffusions, Markov Processes and Martingales, Vol.2: Itô Calculus.* Chichester: Wiley.

ROSS, S.A. (1976): "Options and Efficiency," *Quarterly Journal of Economics*, 90, 75-89.

ROTHSCHILD, M. (1974): "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9, 185-202.

RUBINSTEIN, M. (1983): "Displaced Diffusion Option Pricing," *Journal of Finance*, 38(1), 213-217.

RUBINSTEIN, M. (1994): "Implied Binomial Trees," *Journal of Finance*, 49(3), 771-818.

SANDMANN, K. (1991): *Arbitrage und die Bewertung von Zinssatzoptionen.* Heidelberg: Physika-Verlag.

SCHAEFER, S.M., and E.S. SCHWARTZ (1987): "Time-Dependent Variance and the Pricing of Bond Options," *Journal of Finance*, 5, 1113-1128.

SCHÖBEL, R. (1986): "On a Difference between Options on Stocks and Default-Free Debt Options: The Effect of the Antioption," working paper, Universität Lüneburg, Germany.

SCHÖBEL, R. (1987): *Zur Theorie der Rentenoption.* Berlin: Duncker & Humblot.

SCHWEIZER, M. (1992): "Martingale Densities for General Asset Processes," *Journal of Mathematical Economics*, 21, 177-188.

SONDERMANN, D. (1988): "Option Pricing with Bounds on the Underlying Security," in *Bankpolitik, finanzielle Unternehmensführung und die Theorie der Finanzmärkte*, eds. B. Rudolph and J. Wilhelm. Berlin: Duncker & Humblot, 421-442.

STOLL, H.R. (1969): "The Relationship between Put and Call Prices," *Journal of Finance*, 24, 802-824.

VASICEK, O.A. (1977): "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177-188.

Errata

Page 9, line 6: The new sentence should start with 'An individual's decision as to how much ...'.

Page 11, line 9: Replace 'ever more' by 'evermore'.

Page 31, footnote 8: Insert 'a' after 'the existence of'.

Page 36, footnote 11: Delete 'the'.

Page 54: Replace the proof of Proposition 1.5.3 by:

PROOF: Let $u[r]$ denote the value function for discount rate r, and define $S = \{r > 0 : \; u[r](\hat{\pi}) > \hat{m}\}$. Using the same technique as in the one-sided case (cf. the appendix to Chapter 1), one easily shows that if $r \in S$ and $r' < r$, then $r' \in S$. Therefore, $r^* = \sup S$ has the property that $u[r](\hat{\pi}) > \hat{m}$ if $r < r^*$, and $u[r](\hat{\pi}) = \hat{m}$ if $r > r^*$. It remains to be shown that r^* is a finite positive number.

It is straightforward to show that for

$$r < \underline{r} = \frac{\Delta\beta^2}{2\sigma^2} \, \frac{\hat{\pi}^2 \, (1 - \hat{\pi})^2 \, m''(\hat{\pi})}{\beta(\hat{\pi})},$$

the myopic payoff function m is a strict subsolution of the Bellman equation on the open unit interval, i.e.,

$$\sup_{q \in Q} \left\{ \frac{1}{2\sigma^2} \pi^2 (1 - \pi)^2 (\Delta\alpha - \Delta\beta \, q)^2 \, m''(\pi) - r \, m(\pi) + r \, R(\pi, q) \right\} > 0$$

for $\pi \in]0, 1[$. Similarly, one shows that for

$$r \geq \overline{r} = \frac{2\Delta\beta^2}{\sigma^2} \, \frac{\check{\pi}^2 \, (1 - \check{\pi})^2 \, m''(\check{\pi})}{\beta(\check{\pi})}$$

with

$$\check{\pi} = \frac{\beta_0}{\beta_0 + \beta_1},$$

the function $\check{m} = 2m - \hat{m}$ is a supersolution of the Bellman equation, that is,

$$\sup_{q \in Q} \left\{ \frac{1}{2\sigma^2} \pi^2 (1 - \pi)^2 (\Delta\alpha - \Delta\beta \, q)^2 \, \check{m}''(\pi) - r \, \check{m}(\pi) + r \, R(\pi, q) \right\} \leq 0$$

for all π. Now, results presented in Fleming and Soner (1993, Chapters IV and V) imply that $u > m$ on $]0,1[$ if $r < \underline{r}$, and $u \leq \check{m}$ if $r \geq \overline{r}$. Thus $\underline{r} \leq r^* \leq \overline{r}$. ■

Page 56, line 6: Replace 'subinterval' by 'subintervals'.

Page 87, line 12: The correct reference is de Munnik (1992).

Page 88, line -7: Replace 'accept' by 'agree'.

Page 91, line 4: Insert 'as to' after 'decision'.

Page 96: The last word in the line after equation (2.14) should be 'which'.

Page 108, line 12: The fifth word from the end of the line should be 'ingenious'.

Page 137: In the statement of Theorem 4.1.1, the condition

$$\{t \in \mathcal{T} : \ u_{xx}(t, X_t^1) = 0\} \text{ is almost surely a Lebesgue null set}$$
$$\text{and } u(t,x) \text{ is analytic}$$

must be replaced by

$$\{t \in \mathcal{T} : \ u_{xx}(t, X_t^1) = 0\} \text{ is almost surely a Lebesgue null set,}$$
$$u(t,x) \text{ is analytic, and } \sigma^2(t,x) \text{ has partial derivatives of all orders.}$$

Page 140: In the statement of Theorem 4.2.1, condition (A5) should read

(A5) $u(t,x)$ is analytic, and $\sigma^2(t,x)$ has partial derivatives of all orders.

Page 153: Add the following entry to the list of references:

FLEMING, W.H., and H.M. SONER (1993): *Controlled Markov Processes and Viscosity Solutions.* New York: Springer Verlag.