Pricing and Hedging of Derivative Securities:

Some Effects of Asymmetric Information and Market Power

by

Alexander Stremme

London School of Economics and Political Science

Thesis submitted in partial fulfillment of the requirements of the degree of Doctor of Philosophy at the University of London

January 1999

Abstract

This thesis consists of a collection of studies investigating various aspects of the interplay between the markets for derivative securities and their respective underlying assets in the presence of market imperfections. The classic theory of derivative pricing and hedging hinges on three rather unrealistic assumptions regarding the market for the underlying asset. Markets are assumed to be perfectly elastic, complete and frictionless. This thesis studies some effects of relaxing one or more of these assumptions. Chapter 1 provides an introduction to the thesis, details the structure of what follows, and gives a selective review of the relevant literature. Chapter 2 focuses on the effects that the implementation of hedging strategies has on equilibrium asset prices when markets are imperfectly elastic. The results show that the feedback effect caused by such hedging strategies generates excess volatility of equilibrium asset prices, thus violating the very assumptions from which these strategies are derived. However, it is shown that hedging is nonetheless possible, albeit at a slightly higher price. In Chapter 3, a model is developed which describes equilibrium asset prices when market participants use technical trading rules. The results confirm that technical trading leads to the emergence of speculative price "bubbles". However, it is shown that although technical trading rules are irrational ex-ante, they turn out to be profitable ex-post. In Chapter 4, a general framework is developed in which the optimal trading behaviour of a large, informed trader can be studied in

an environment where markets are imperfectly elastic. It is shown how the optimal trading pattern changes when the large trader is allowed to hold options written on the traded asset. In Chapter 5, the results of the preceding chapter are used to study the interplay between options markets and the markets for the underlying assets when prices are set by a market maker. It turns out that the existence of the option creates an incentive for the informed trader to manipulate markets, which implies that equilibrium on both markets can only exist when option prices are adjusted to reflect this incentive. This requirement of price alignment explains the "smile" pattern of implied volatility, an empirically observed phenomenon that has recently been the focus of extensive research. Chapter 6 finally addresses recent proposals by some researchers suggesting that central banks should issue options in order to stabilise exchange rates. The argument, in line with the findings of Chapter 2, is based on the fact that hedging a long option position requires countercyclical trading that would reduce volatility. However, the results of Chapter 6 show that the option creates an incentive for market manipulation which, rather than protecting against speculative attacks, in fact creates an additional vehicle for such attacks. Chapter 7 concludes.

Acknowledgements

First and foremost, I would like to express my gratitude to Lucien Foldes, my supervisor. His guidance and encouragement have been invaluable. I would also like to extend my thanks to Hans Föllmer, who introduced me to the field of financial economics and guided me through the first year of work on this thesis. While Hans Föllmer taught me most of what I know about the mathematics, making the most abstract concepts seem intuitive, Lucien Foldes taught me the art of economic thought and showed me that financial economics was more than simply applied mathematics. I am also deeply indebted to David Webb for providing me with a place in the Financial Markets Group and supporting me in so many ways. Access to the facilities and personnel of the FMG have played a very large part in my completion of this project.

Throughout my tenure at the Financial Markets Group I have drawn upon the wisdom and intellect of numerous researchers who have passed through or, in some cases, set down roots. In addition to the invaluable intellectual stimulus, I am equally grateful for the warm welcome the group extended to me. Many of the group's current and former members have become good friends. While an exhaustive list would fill several pages, I would like to mention Alvaro Almeida, Simon Emrich, Sylvain Friederich, Urs Haegler, Marc Henry, Harry Kinzler, Bob Nobay, Sven Rady, Jörn

Rothe, and Daniel Sommer. It is with particular pleasure however that I acknowledge the endurance and tolerance demonstrated by Ward Brown and Richard Payne as my office-mates over the past few years. They were great colleagues and have become dear friends. Finally, the Financial Markets Group would not be what it is without the fabulous effort of its past and present administrative staff.

I am indebted to Rüdiger Frey and Asbjørn Hansen for the fruitful collaboration which ultimately lead to the papers that constitute Chapters 2 and 5 of this thesis, and also for many useful discussions and suggestions regarding other parts of my work. Also, I would like to thank Sudipto Bhattacharya, Patrick Bolton, Margaret Bray, Darrell Duffie, Gérard Genotte, Stanley Pliska, and participants of seminars and conferences at the London School of Economics and elsewhere for many helpful comments.

Financial support by the Financial Markets Group, the German Academic Exchange Service (DAAD), the London School of Economics Economica Scholarship, and the Economic and Social Research Council (ESRC) is gratefully acknowledged.

Finally, I would like to thank my family for their support, both financial and emotional, and my wife, Raquel, for the patience with which she endured my moods, and for all her love, support and encouragement without which this thesis would not have been completed. It is appreciated very deeply.

"If you have a closer look at all aspects of life,
everything has an option embedded in it."
Anonymous Student

"These ones are very small,
the ones out there are far away."
Father Ted

Foreword

The chapters of this thesis are derived from several research papers, some of which are the result of joint work, and some of which have subsequently been published. Chapter 2 is an extended version of a working paper entitled "Portfolio Insurance and Volatility", written jointly by Rüdiger Frey and myself. An updated version of this paper has been published in *Mathematical Finance* under the title "Market Volatility and Feedback Effects from Dynamic Hedging". Rüdiger Frey was also partly involved in the very early stages of the development of Chapter 3. Chapter 5 is the result of joint work with Asbjørn Hansen and has been submitted to the *Review of Financial Studies*. All other chapters are my own work.

As a result of the attempt to keep each chapter self-contained, some arguments which are needed in more than one chapter are at least in part repeated accordingly. This applies in particular to the proof of Theorem 4.3.1, variants of which appear in Sections 5.3.3 and 6.3.3. The reader is encouraged to skip those sections if appropriate.

London, January 1999

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis consists of a collection of studies investigating various aspects of the interplay between the markets for derivative securities and their respective underlying assets in the presence of market imperfections. Classic theory of derivatives pricing and hedging hinges on three key assumptions regarding the market for the underlying asset. Markets are assumed to be perfectly elastic (traders can trade arbitrarily large quantities without affecting prices), frictionless (there are no transaction costs), and complete (all contingent claims can be replicated). Clearly, this is a very stylised view of real financial markets, where these assumptions are satisfied only to a very limited extent. This is why in recent years, a rapidly growing strand of literature has focused on studying the implications of relaxing one or more of these assumptions.

Throughout this thesis, we assume that markets for the underlying asset are imperfectly elastic, i.e. agents' trading activities have an effect on equilibrium prices. We investigate the manner in which the dynamic properties of equilibrium asset prices are affected by different types of trading behaviour. More specifically, we consider cases in which market participants implement dynamic trading strategies designed to hedge option positions, follow technical trading rules, or strategically exploit their market power in order to maximise expected profits.

Classic option pricing theory is based on the concept of replication. Based on the stochastic law governing the dynamics of the underlying asset's price, a self-financing, dynamic trading strategy is constructed which replicates the option's pay-offs. However, in imperfectly elastic markets, the implementation of such strategies is likely to affect the underlying equilibrium price process, thus perturbing the very model they are derived from. The Group of Ten (1993) for example reports that the strategies related to option hedging or replication do indeed have significant impact on prices, in particular on volatility. Moreover, typical hedging strategies require selling into falling markets and buying into rising markets and are thus likely to have destabilising effects on prices. In fact, hedging has been blamed as one of the contributing factors in the October 1987 stock market crash.

Similar arguments have been put forward with regards to the implications of technical trading. While economic theory generally classifies technical trading as irrational, empirical evidence underlines its importance, in particular in foreign exchange markets; see Allen and Taylor (1992). Like hedging, technical trading creates positive feedback effects of past returns on current prices, and is hence seen as one of the contributing factors in recent market crashes. Moreover, in contrast to the predictions of most of the theoretical literature, empirical evidence seems to prove that technical trading rules are indeed profitable; see Levich and Thomas (1993), Brock, Lakonishok, and LeBaron (1992), and Sweeney (1986).

The implications of finitely elastic markets and the inherent opportunities for market manipulation become particularly important in the presence of asymmetric information. While extensive literature has arisen in recent years which studies the effects of asymmetric information on asset prices, comparatively little has been said about the implications of information asymmetries for the pricing of options. In classic models of optimal trading under asymmetric information such as Kyle (1985), the informed trader always has an incentive to move prices towards the asset's expected true value, thus partially revealing her private information. However, if the informed trader holds an option, written on the underlying asset, this may no longer be true. Intuitively, such an option would create an incentive for the informed trader to manipulate prices, thus pushing the option into the money. The gains from this manipulation will be traded off against potential losses made in the underlying market. This manipulation incentive is likely to have important implications on the amount of information revealed by the informed trader's actions. Moreover, the value of the option to the informed trader will incorporate potential gains from manipulation and will hence be different from the option's arbitrage price. This has important consequences for the pricing of options in the context of asymmetric information and market power.

The effects outlined in the preceding sections are likely to have important implications in many other situations. For example, recent suggestions for central banks to use options as exchange rate policy instruments will have to be re-evaluated in light of the market manipulation incentives created by such options. Empirical studies suggest that the main objectives of central bank intervention on the foreign exchange market is to stabilise nominal rates and keep them in a certain "target zone"; see for instance Edison (1993) or Almekinders (1995). The use of *options* as instruments of exchange rate policy was first proposed by Taylor (1995). He suggests that central banks should buy *put* options, written on the domestic currency. When the domestic currency depreciates as a result of a speculative attack, Taylor argues, the option is deep in the money allowing the central bank to buy foreign currency at deflated rates, which can then be sold in the spot market to support the domestic currency.

Breuer (1997) identifies five main drawbacks of the approach suggested by Taylor, the most important of which is the destabilising effect of "delta hedging". Market participants who have written the options bought by the central bank would typically hedge their position by appropriate strategies in the spot market. Following the arguments made earlier, this type of trading activity is likely to have destabilising effects on equilibrium exchange rates. Empirical evidence supporting the importance of such effects has been reported, amongst others, by Malz (1995). In view of this, Wiseman (1996) and Breuer (1997) propose an alternative way of utilising options as exchange rate policy instruments. Rather then *buying* put options, they argue, the central bank should *write* call options. As a result, when hedging their long option position, market participants would be required to sell in rising markets and buy in falling markets, thus stabilising exchange rates. However, this argument rests on the assumption that the buyers of such options indeed content themselves with hedging their position. In light of the arguments made earlier, the following question arises: While it is true that the option induces extremely risk-averse market participants to hedge their positions thus stabilising rates, the situation might change drastically if instead the option is held by speculators who, being risk-neutral, have no incentive to hedge risk but to maximise expected profits.

In Chapter 2, we study the manner in which the additional demand generated by strategies designed to hedge an option position affects the underlying asset's equilibrium price, and in particular its volatility structure. There have been a number of studies on the impact of dynamic hedging on the price of the underlying asset. Grossman (1988) focuses on informational differences between buying an option and running the corresponding replicating strategy. Gennotte and Leland (1990) study the effects of hedging in an asymmetric information model similar to the one considered by Grossman and Stiglitz (1980). Jarrow (1994) analyzes, in a discrete-time model, how standard option pricing theory can be extended to the situation where hedging does affect the underlying price process. The results have subsequently been extended to the continuous-time case in Frey (1996). Platen and Schweizer (1998) use the feedback effect of portfolio insurance to explain the "smile pattern" of implied volatility. Their model relies to a large extent on the model outlined in Chapter 2, and previously published in Frey and Stremme (1997). The study by Brennan and Schwartz (1989) is the one most closely related to the one presented in Chapter 2. However, while in their model optimal trading is solely based on exogenous information on the long-term, fundamental prospects of the underlying assets, we allow expectations to be affected by current prices. This approach is supported by empirical evidence, as reported for example by the Group of Ten (1993), Allen and Taylor (1992), or De Long, Shleifer, Summers, and Waldmann (1990a).

As analytical framework for our analysis, we first construct a discrete-time equilibrium model in which myopic, profit-maximising agents interact with "program traders" who follow dynamic hedging strategies. To facilitate the qualitative analysis, we then pass to the continuous-time limit. This approach to the construction of diffusion models for asset prices was first proposed by Föllmer and Schweizer (1993). While in the absence of hedging, equilibrium prices follow a geometric Brownian motion as in the classic Black-Scholes model, the introduction of program traders causes market volatility to increase and become time and price dependent. The magnitude of this effect is much greater than that observed by Brennan and Schwartz (1989). Our findings are in line with empirical observations, as reported for example by the Group of Ten (1993).

The transformation of volatility as predicted by our model has important implications for practical applications. While replicating strategies can still be shown to exist, they can no longer be calculated explicitly. This is why most practioners rely on the simple strategies prescribed by the Black-Scholes formula. This bears the obvious question as to how well these strategies perform if in fact the feedback effect caused by their implementation violates the very assumptions they are derived from. Using an idea of El Karoui, Jeanblanc-Picqué, and Shreve (1998) we show that while perfect replication is impossible, simple Black-Scholes strategies can nonetheless be used to "super-replicate" the option's pay-off.

In Chapter 3, we consider an equilibrium model for exchange rates in which some agents follow technical trading rules. There have been a number of theoretical studies analysing the effects of technical trading. In Day and Huang (1990), rational "information traders" interact with technical traders, and equilibrium is determined by a market maker. Unlike in the model considered here, Day and Huang do not introduce any source of noise into the market. As a result, prices in their model follow deterministic chaotic dynamics. In Lux (1995), traders act upon rumours and fads rather than technical analysis. Contagious mood swings cause cyclical fluctuations of equilibrium prices.

As analytical framework for our study, we first construct a discrete-time equilibrium model in which technical traders interact with pure "noise traders". To facilitate the analysis of the resulting equilibrium exchange rate process, we then pass to the continuous-time limit. It turns out that, as long as technical traders are active, bubbles will always emerge. Exchange rate fluctuations caused by noise are picked up by technical analysis and amplified by the buy or sell signals generated by the corresponding trading rules. However, we also show that any bubble will always burst within finite time.

We demonstrate that, while the expected profit from simple buy-and-hold strategies is zero, technical trading rules themselves are profitable on average. In other words, technical trading can be seen as a type of "self-fulfilling prophecy". These results are in line with the empirical findings of Levich and Thomas (1993), Brock, Lakonishok, and LeBaron (1992), and Sweeney (1986).

In Chapter 5, we investigate the manner in which an option, held by a large, informed trader, changes the nature of equilibrium prices for the underlying asset, and how this feeds back into option prices. As analytical framework, we choose the model developed in Chapter 4, which in turn is based on continuous-time Kyle (1985) model. A large trader, who receives a private signal about the fundamental value of the traded asset, interacts with pure noise traders. Prices are set by a risk-neutral competitive market maker. In this type of model, the informed trader typically drives prices towards the expected value of the asset, thus successively revealing the private information to the market maker. We extend the Back (1992) model by introducing an over-the-counter market on which options can be traded prior to trading in the underlying asset.

Unsurprisingly, it turns out that the presence of the option creates an incentive for the informed trader to manipulate prices of the underlying asset in order to increase the option's pay-off. Uninformed traders might hence face a price mark-up over the expected true value of the asset. In other words, the informed trader creates a price bubble at the expense of uninformed traders. It is the potential existence of such bubbles which constitutes the main difference between our model and the model by Back (1993), in which this kind of manipulation is precluded by the existence of a liquid options market with additional noise.

In our model, the only feasible equilibrium requires that the market maker does not believe the informed trader to have an incentive to manipulate prices. This belief can only be rational if equilibrium option prices are "synchronised" with prices on the underlying market. As a consequence, option prices are non-linear in quantity. This is in line with the findings in Jarrow (1994) and Frey (1996), where option prices are determined by the cost of replication, and thus non-linearity arises as a consequence of finite elasticity of the underlying market. It is also consistent with empirical observations made in the OTC market for options, in particular for large quantities. We derive an explicit expression for implied volatility. We show that the implied volatility pattern generated by equilibrium prices in our model displays the famous "skew smile pattern", which has been observed in most financial markets ever since the October 1987 crash.

There have been a number of studies that explain the smile pattern of implied volatility. While most of these explain the smile pattern by exogenously altering the volatility structure of the underlying price process, no assumptions regarding the underlying price volatility are needed in our model. The smile pattern rather occurs endogenously as a consequence of the market structure. To the best of our knowledge the only other paper in which the smile has been obtained endogenously is the paper by Platen and Schweizer (1998). Their analysis is based on a modified version of the feedback model developed in Chapter 2 and published previously in Frey and Stremme (1997). However, in order to explain the smile pattern, Platen and Schweizer have to assume an upward-sloping demand curve for the underlying asset, which implies that the equilibrium they obtain is highly unstable.

The emergence of the skew smile pattern in our model is a result of the attempt to prevent market breakdowns. Empirically, the skewness is a feature which has only been observed *after* the October 1987 stock market crash. The explanation offered by our analysis could be that after the crash, market participants and regulators implemented measures aimed at preventing similar events in the future. This might have brought option prices more in line with the markets for the underlying asset.

According to our findings, such an increased level of price synchronicity would indeed result in the implied volatility pattern that has been observed since the crash.

In Chapter 6, we analyse the possible implications of using options as central bank exchange rate policy instrument. We model the situation in which a speculator, having bought the option issued by the central bank, instead of simply hedging her position, strategically exploits the leveraged position in the market provided by the option. As theoretical framework, we chose the continuous-time model developed in Chapter 4, with prices given by Walrasian equilibrium. The risk-neutral speculator interacts with "information traders", who base their demand for foreign exchange on the flow of fundamental "news". We assume that the central bank's objective is to stabilise exchange rates and to keep them within a given "target zone", which is in line with the empirical findings of Edison (1993). In order to achieve their objective, the central bank can intervene using spot transactions or by issuing options. In addition to trading in the spot market, the speculator may purchase the options issued by the central bank.

It turns out that in the absence of options, the speculator has no incentive to manipulate markets. In fact, in this case it is optimal not to trade at all. An option issued by the central bank however creates an alternative way for the speculator to realise the gains from market manipulation. Upon arrival of "bad news", when information traders cause the domestic currency to depreciate, the option creates an incentive for the speculator to squeeze exchange rates even further. Scenarios can arise in which without speculation, the central bank's foreign currency reserves would have been sufficient to support the domestic currency, while they cannot sustain the additional pressure arising from the speculator's manipulation. In other words, instead of *protecting* against speculative attacks, options issued by the central bank in fact *create* an additional vehicle for such attacks.

Chapter 2

Feedback Effects from Dynamic Hedging

2.1 Introduction

Standard derivative pricing theory is based on arbitrage arguments, which in turn rest on three key hypotheses about the markets for the underlying asset. Markets are assumed to be complete, frictionless and perfectly elastic. Clearly, this is a very stylised view of real financial markets, in which these assumptions are satisfied only up to a certain extent. This is why a rapidly growing literature has concentrated on the implications of relaxing one or more of them. In this chapter we drop the elasticity assumption and study the manner in which dynamic hedging strategies affects the underlying asset's equilibrium price, in particular its volatility structure.

Hedging strategies are derived from specific assumptions on the stochastic law that governs the underlying price dynamics. In practice they are seen both as theoretical concept for option pricing and, more importantly for our analysis, as device to manage risk as incurred for instance by selling OTC derivative contracts.

We believe an analysis of the feedback effects caused by dynamic hedging in imperfectly elastic markets to be important for a number of reasons. To begin with, when carried out on a large scale, dynamic hedging is most likely to perturb the very stochastic law it is based upon. We ask how hedging strategies perform when we allow the underlying price process to be affected by their implementation, even if this effect is not fully taken into account in designing them. Moreover, hedging is mostly used to replicate payoffs that are convex functions of the underlying asset's price, requiring the investor to sell shares of the underlying asset when its price declines and to buy when its price goes up.[1] Therefore one should expect an increase of market volatility under the impact of such trading behaviour. Dynamic hedging is thus likely to have a destabilising effect on prices.

There have been a number of studies on the impact of dynamic hedging on the price of the underlying asset. Grossman (1988) focuses on informational differences between buying an option and running the corresponding replicating strategy. Gennotte and Leland (1990) study the effects of portfolio insurance in a model with asymmetric information similar to the one considered by Grossman and Stiglitz (1980). They find that the better this activity is understood by the other market participants, the weaker is the effect of hedging.

Jarrow (1994) analyses, in a discrete-time model, how standard option pricing theory can be extended to a situation where there are feedback effects from the demand of a "large trader" on the underlying price process. Platen and Schweizer (1998) use

[1] This kind of trading behaviour is also referred to as "Portfolio Insurance".

the feedback effect of portfolio insurance to explain the "smile pattern" of implied volatilities that is observed in practice. Their model relies to a large extent on the model developed here and previously published in (Frey and Stremme 1994) and (Frey and Stremme 1997).

Brennan and Schwartz (1989) address an issue very similar to the one discussed here. They analyse the transformation of market volatility under the impact of portfolio insurance in a finite-horizon economy in which securities are traded continuously but consumption takes place only at the terminal date. Agents are hence only concerned about the long-term prospects of the asset. Since the risky asset's terminal value is entirely determined by an exogeneously given random variable, which is interpreted as the fundamental value of the asset, agents' expectations are solely driven by the successively revealed information about the value of this state variable. In particular, agents do not alter their expectations in reaction to changes in current price. Markets are thus very liquid, causing the feedback effect of hedging on volatility to be relatively small.

Empirical evidence however suggests that in many situations there is "an enormous amount of short-term position taking", whereas the funds dedicated to long-term investment are limited by uncertainty and agents' risk aversion; see for instance Goodhart (1988). A theoretical justification for this kind of trading behaviour is given by De Long, Shleifer, Summers, and Waldmann (1990b). Moreover, when making trading decisions for very short periods, like in intra-day dealing, investors seem to rely more on the information conveyed by current price movements than on the long-term fundamental prospects of the assets. This "Keynesian" view of investment is supported by evidence reported by the Group of Ten (1993) or Allen and Taylor (1992). In fact there is evidence for a positive feedback effect of current price changes on expectations, see De Long, Shleifer, Summers, and Waldmann (1990a).

In the present chapter we develop a framework in which the effect of dynamic hedging on the underlying asset's price process can be studied. We start by constructing a general discrete-time temporary equilibrium model in which short-term investment can be modelled. To get a clearer picture of the equilibrium price process, and in particular of its volatility, we then pass to a limiting continuous-time diffusion. This approach to the construction of diffusion models for asset prices was first proposed by Föllmer and Schweizer (1993).

As a special case we consider an economy populated by traders whose preferences over future wealth exhibit constant relative risk-aversion as in the Brennan-Schwartz study. Agents take changes in current prices as signals for future price movements. If they were the only traders in the market the equilibrium price process would be a geometric Brownian motion as in the classic Black-Scholes option pricing model. However, if they interact with program traders who are running dynamic hedging strategies, the structure of the limiting equilibrium price process becomes more complex: while it still can be represented as an Itô process, its volatility increases and becomes time and price dependent. A comparison reveals that the increase in volatility is much more pronounced in our study than that observed by Brennan and Schwartz (1989). This finding underlines the importance of agents' expectations in determining the liquidity of the market. Moreover it illustrates that there exist realistic scenarios in which the effect of hedging is far larger than predicted by Brennan and Schwartz.

We derive an explicit expression for the transformation of market volatility under the impact of hedging. We use this transformation rule to study in particular the feedback effects generated by the strategies derived from the classic Black-Scholes formula. It also allows us to study the importance of different payoff structures being hedged. We show that increasing heterogeneity of the distribution of hedged contracts reduces both level and price sensitivity of volatility.

As reported by the Group of Ten (1993) the effects predicted by our analysis are indeed observed in practice:

> "[T]he existence of options and related dynamic hedging could increase volatility, especially in the smaller and less liquid currency segments, as the spot exchange rate approaches the strike price. When strike prices and/or option maturities are highly concentrated, a large volume of one-way hedging could occur in a short period. Market participants reported that sharp [...] movements in spot prices were frequently observed as a result of such concentrations."

Price dependent volatility, as results from hedging in our model, causes problems in practical application. While hedging strategies can still be shown to exist they can no longer be calculated explicitly. This is why most practioners rely on the classic Black-Scholes formula. Using an idea of El Karoui, Jeanblanc-Picqué, and Shreve (1998) we are able to show that simple strategies derived from a constant-volatility Black-Scholes model are still sufficient to completely hedge the risk incurred by selling OTC derivatives. This remains true even if their implementation causes the actual volatility to be price dependent, as is the case in our model. However, the misperception of the feedback effect on volatility generates a "tracking error": the terminal value of the hedge portfolio might exceed the payoff it was supposed to replicate, thus requiring an initial "over-investment" in the strategy. But again, heterogeneity proves to be beneficial: the tracking error and hence the over-investment diminishes with increasing heterogeneity of the distribution of hedged payoffs.

The remainder of the chapter is organised as follows. In Section 2.2 we develop the general discrete-time temporary equilibrium framework. We then specify a concrete sample economy in which agents whose preferences exhibit constant relative risk-aversion interact with program traders who run dynamic hedging strategies. Section 2.3 is devoted to the passage to the continuous-time limit. We characterise the limiting price process as the solution of an Itô type Stochastic Differential Equation, thus obtaining an explicit expression for its volatility. In Section 2.4 we conduct a detailed study of the feedback effect caused by the implementation of Black-Scholes hedging strategies and directly compare our findings to those of Brennan and Schwartz (1989). Section 2.5 concludes.

2.2 The Discrete-Time Model

We consider a sequence of discrete-time infinite-horizon economies. More precisely, for each $n = 1, 2, \ldots$ there is a sequence of times $0 = t_0^n < t_1^n < \ldots < t_k^n < \ldots$ at which trading takes place on a Walrasian market. In view of the intended passage to a continuous-time limit diffusion we assume that

$$\Delta^n := \sup_k (t_{k+1}^n - t_k^n) \longrightarrow 0 \qquad \text{as } n \to \infty.$$

TRADED ASSETS:

There are two assets in the economy, a riskless one (typically a bond or money market account), and a risky one (typically a stock or foreign exchange rate). We take the riskless asset as numéraire, thereby making interest rates implicit in our model. Moreover, we assume the market for the riskless security to be perfectly elastic. This is an idealisation of the fact that money markets are far more liquid than those for the typical risky asset considered here. The equilibrium price at time t_k^n of the risky asset, accounted in units of the numéraire, is denoted by X_k^n.

In the present paper we are only interested in the feedback effect of hedging on the underlying asset's volatility, and not in developing a *pricing* theory. In finitely liquid markets it is no longer obvious how to derive option prices from the prices of the underlying. To avoid the price inconsistencies that can arise from an inadequate modelling of the relationship between stock and options markets we assume that there is no liquid market for options on the risky asset.

AGGREGATE DEMAND SCHEDULE:

At any time t_k^n the aggregate demand function for the risky asset is assumed to be given by a smooth function $G^n : [0, \infty) \times \mathbb{R}_+^2 \longrightarrow \mathbb{R}$ in the form

$$x \mapsto G^n(t_k^n, F_k^n, x). \tag{2.1}$$

Here, x is the (proposed) Walrasian price. $(F_k^n)_{k=0,1,\dots}$ is a stochastic process describing the current state of the economy, to be specified in more detail later. Note that the above form of the demand function implies in particular that all the information necessary for the investors to form their demand can be summarised in F_k^n and x.

EQUILIBRIUM:

We normalise total supply of the risky asset to one, hence the equilibrium price X_k^n at time t_k^n is determined by the *market clearing equation*

$$G^n(t_k^n, F_k^n, X_k^n) \equiv 1. \tag{2.2}$$

The following assumptions are of technical nature. They ensure existence and uniqueness of equilibria and guarantee convergence of the equilibrium price processes (see Section 2.3). We will see later how these assumption can be achieved in a more concrete specification of the economy (see Section 2.2.1 and Corollary 2.3.4).

Assumption 2.2.1 *The functions G^n are smooth, and the sequence $\{G^n\}_{n=1,2,\dots}$ converges uniformly on compacts to a smooth function $G : [0, \infty) \times \mathbb{R}_{++}^2 \to \mathbb{R}$. Moreover,*

1. *For every* $(t,f) \in [0,\infty) \times \mathbb{R}_{++}$ *the equations* $G^n(t,f,x) = 1$ *and* $G(t,f,x) = 1$ *have exactly one solution in* x*, denoted by* $\psi^n(t,f)$ *and* $\psi(t,f)$*, respectively.*

2. *For every compact set* $K \subset\subset \mathbb{R}^2_{++}$ *the sequence* $\{\psi^n\}_{n=1,2,\ldots}$ *is relatively compact in the space of all bounded functions on* K *endowed with the supremum norm.*

3. *For any fixed* t *and* f*, the derivatives of* G^n *satisfy "in equilibrium":*

$$\frac{\partial G^n}{\partial x}(t,f,x)\bigg|_{x=\psi^n(t,f)} < 0 \quad \text{and} \quad \frac{\partial G^n}{\partial f}(t,f,x)\bigg|_{x=\psi^n(t,f)} > 0,$$

and the analogous statements hold for the limit function G.

REMARKS: Note that (1) guarantees that there is a unique solution to the market clearing equation in each discrete-time economy characterised by G^n as well as in the continuous-time limit economy characterised by G. The first inequality in (3) together with the Implicit Function Theorem implies in particular that the solution in the limit economy depends smoothly on f, i.e. the function $\psi : [0,\infty) \times \mathbb{R}_{++} \longrightarrow \mathbb{R}$ is smooth. The second inequality in (3) implies in addition that for fixed t the mapping $f \mapsto \psi(t,f)$ is invertible, i.e. there exists a smooth function $\psi^{-1} : [0,\infty) \times \mathbb{R}_{++} \longrightarrow \mathbb{R}$ such that $\psi^{-1}(t,\psi(t,f)) = f$ and $\psi(t,\psi^{-1}(t,x)) = x$ for all t, f and x. Finally note that differentiability of the G^n together with (3) implies that the ψ^n are differentiable, too. The Arzéla-Ascoli Theorem then implies that (2) holds if the ψ^n and their first derivatives are uniformly bounded on compacts.

AGENTS:

There are two groups of agents in the market, called "reference traders" and "program traders", respectively. The economy in which there are only reference traders active constitutes the benchmark case for our analysis. It will be compared with the case in which reference traders interact with program traders, who are running dynamic hedging strategies. To measure the relative market weights of the two groups we introduce the parameter $\rho \in [0,1]$, which denotes the *relative market share* of program traders, thus leaving a share of $1-\rho$ to the reference traders. The total aggregate demand function can then be written in the form

$$G^n(t,f,x) = (1-\rho)D^n(f,x) + \rho\phi^n(t,x), \tag{2.3}$$

where $D^n(f,x)$ and $\phi^n(t,x)$ are *normalised*[2] aggregate demand functions for the reference traders and program traders, respectively. We neglect the aggregation problem and specify a representative reference trader whose (normalised) demand function takes the form $x \mapsto D^n(F_k^n, x)$. A more detailed specification in which the reference trader's preferences exhibit constant relative risk-aversion will be given in Section 2.2.1.

A typical program trader might be a bank hedging a portfolio of written OTC contracts by running a dynamic trading strategy in the underlying asset. Since the majority of the demand for such contracts is motivated by considerations beyond the scope of our model,[3] we take the extreme view that the hedging objectives of our program traders are exogenously given. Moreover, the hedging strategy for a portfolio of payoffs is just the portfolio of the hedging strategies for the individual payoffs. Thus we can concentrate on a representative program trader, whose (normalised) demand function takes the form $x \mapsto \phi^n(t_k^n, x)$. We make the following assumptions on the strategy functions ϕ^n:

[2]Normalised in the sense that, if reference resp. program traders were the only agents in the market, D^n resp. ϕ^n would be their demand function. When both groups interact, their demand functions are weighted by their respective market weights, $1-\rho$ and ρ.

[3]The Group of Ten (1993) for example reports that "derivative instruments were primarily used for risk hedging purposes".

Assumption 2.2.2 *The functions ϕ^n are smooth, and the sequence $\{\phi^n\}_{n=1,2,\ldots}$ converges uniformly on compacts to a smooth function $\phi : [0,\infty) \times \mathbb{R}_{++} \longrightarrow \mathbb{R}$. Moreover,*

1. *for every compact set $K \subset [0,\infty) \times \mathbb{R}_{++}$ we have*
$$\sup_n \sup_{(t,x)\in K} \left| \frac{\partial \phi^n}{\partial t}(t,x) \right| < \infty,$$
2. *ϕ^n is increasing in the underlying price, i.e.*
$$\frac{\partial \phi^n}{\partial x}(t,x) > 0 \quad \text{for all } t \geq 0 \text{ and all } x \in \mathbb{R}_{++},$$
3. *and ϕ^n is normalised in such a way that*
$$\sup_{t,x} |\phi^n(t,x)| = 1.$$

Finally, the limit ϕ satisfies the analogous conditions to (2) and (3).

Note that by (3) ρ also determines the fraction of the total supply of the risky asset that is subject to portfolio insurance. Also note that (2) reflects the fact that a typical hedging strategy requires that shares of the underlying be sold when its price has declined and vice versa, as was mentioned in the introduction.

Assumption (2.2.2) is satisfied in particular if ϕ^n and ϕ are mixtures of hedging strategies as given by the Black-Scholes formula. It is also possible to consider strategies ϕ^n derived from the discrete state-space model of Cox, Ross, and Rubinstein (1979). Convergence of such strategies to their continuous-time counterpart ϕ follows, for instance, from results by He (1990).

2.2.1 A Case with Constant Relative Risk Aversion

In this section we provide a concrete specification of the preferences and beliefs of the representative reference trader. The assumptions introduced in this section ensure that the aggregate demand functions of this particular economy satisfy Assumption (2.2.1) (see Corollary 2.3.4 below). The model considered here is closely related to the kind of temporary equilibrium models discussed by De Long, Shleifer, Summers, and Waldmann (1990b). Consider an overlapping generations model without bequests, in which agents live for two periods. When young, the representative reference trader receives an exogenous stochastic income F_k^n which she invests in the available assets. When old, she just consumes all her wealth and then disappears from the market. Thus, at any time t_k^n the young agent chooses the number d of shares of the risky asset she wants to hold in order to maximise expected utility of next period's wealth. Given her income is f, her demand function will be

$$D^n(f,x) = \arg\max_{d \geq 0} E\left[\, u\left(f + d \cdot (\tilde{X}_{k+1}^n(x) - x)\right)\,\right],$$

where u is her von Neumann-Morgenstern utility function and $\tilde{X}_{k+1}^n(x)$ the agent's belief about next period's price. Note here that we allow the expected future price to depend explicitly on the current price x, i.e. agents may update their expectations in reaction to changes in current prices.

Of course this overlapping generations scenario must not be taken literally. It is a stylised model of a market where agents' investment decisions are made sequentially over time and where each decision is determined mainly by myopic optimisation. For example, market participants might be managers of investment funds who are managing a stochastically fluctuating amount of funds. Typically fund managers are (at least partly) compensated according to the performance of their portfolio, evaluated at certain predetermined dates. Therefore their investment decisions are often predominantly aimed at the next evaluation date.

Assumption 2.2.3 *Reference traders' beliefs and preferences and the evolution of their income over time are assumed to be given as follows:*

1. *The representative reference trader's preferences exhibit constant relative risk aversion, i.e. her von Neumann-Morgenstern utility function u satisfies $u'(z) = z^{-\gamma}$ for some $\gamma > 0$.*

2. *Given current price x, the agent believes next period's price $\tilde{X}^n_{k+1}(x)$ to be of the form $\tilde{X}^n_{k+1}(x) = x \cdot \xi^n_k$ for some random variable ξ^n_k. We assume $(\xi^n_k)_{k=0,1,\ldots}$ to be serially independent and independent of x, and that $E[\xi^n_k] \geq 1$.*

3. *Given current income F^n_k, next period's income F^n_{k+1} is given by $F^n_{k+1} = F^n_k \cdot \zeta^n_k$ for some random shock $\zeta^n_k > 0$. We assume $(\zeta^n_k)_{k=0,1,\ldots}$ to be serially independent.*

Note that by (2) there is a *positive feedback* from current price x into agents' expectations: after a rise of x they anticipate a rise in future prices and in the case of a price decline they expect future prices to fall as well. A list of striking observations which emphasise the importance of such extrapolative expectations on financial markets has been compiled by De Long, Shleifer, Summers, and Waldmann (1990a). We will see in Section 2.4.3 that this way of expectation formation leads to destabilising effects of dynamic hedging which are much larger than those observed in the framework of Brennan and Schwartz (1989).

By Assumption (2.2.3) the solution to the agent's utility maximisation problem, given income f and proposed price x, is uniquely determined by the first order condition

$$\begin{aligned} 0 &= E\left[u'\left(f + d \cdot (\tilde{X}^n_{k+1}(x) - x)\right) \cdot (\tilde{X}^n_{k+1}(x) - x) \right] \\ &= E\left[(f + d \cdot x \cdot (\xi^n_k - 1))^{-\gamma} \cdot x \cdot (\xi^n_k - 1) \right]. \end{aligned} \tag{2.4}$$

As an immediate consequence of this characterisation we get

Lemma 2.2.4 *Under Assumption (2.2.3) the representative reference trader's demand function D^n satisfies the following homogeneity properties:*

1. $D^n(\alpha f, x) = \alpha D^n(f, x)$ *for all* α, *and*
2. $D^n(\alpha f, \alpha x) = D^n(f, x)$ *for all* α.

In particular, using these homogeneities we find

$$D^n(f, x) = D^n(x \cdot \frac{f}{x}, x) = \frac{f}{x} \cdot D^n(1, 1) =: \frac{f}{x} \cdot D^n_*. \tag{2.5}$$

Note that, due to homogeneity, the *weighted* demand function $(1-\rho)D^n(F^n_k, x)$ can also be written as $D^n((1-\rho)F^n_k, x)$. It can hence be interpreted as the demand of a representative reference trader who has only a fraction $1-\rho$ of the aggregate wealth F^n_k at her disposal. Thus ρ plays the same role in our model as does the parameter α in the model of Brennan and Schwartz (1989).

EQUILIBRIUM WITHOUT PROGRAM TRADERS:

In the absence of program traders the market clearing equation takes the form

$$1 = D^n_* \cdot \frac{F^n_k}{X^n_k} \qquad \Longrightarrow \qquad X^n_k = D^n_* \cdot F^n_k. \tag{2.6}$$

Using Assumption (2.2.3) (3) this implies $X^n_{k+1} = D^n_* \cdot F^n_{k+1} = D^n_* \cdot F^n_k \cdot \zeta^n_k = X^n_k \cdot \zeta^n_k$. We summarise this in the following

Lemma 2.2.5 *In the economy specified in Assumption (2.2.3), the equilibrium price process* $(X_k^n)_{k=0,1,\ldots}$ *in the absence of program traders is given by*

$$X_{k+1}^n = X_k^n \cdot \zeta_k^n. \tag{2.7}$$

In particular, expectations are rational if and only if $\xi_k^n \equiv \zeta_k^n$ *for all* k.

As an example of income dynamics consider the sequence

$$\zeta_k^n = \exp\left((\mu - \frac{1}{2}\eta^2)(t_{k+1}^n - t_k^n) + \eta\sqrt{t_{k+1}^n - t_k^n} \cdot \varepsilon_{k+1}^n\right),$$

where $(\varepsilon_k^n)_{k=0,1,\ldots}$ is an i.i.d. sequence of random variables. If, for instance, the ε_k^n are standard-normally distributed, then prices follow a discretised geometric Brownian Motion. If, on the other hand, the ε_k^n are just the increments of a random walk, then prices are given by a geometric random walk as in Cox, Ross, and Rubinstein (1979).

EQUILIBRIUM WITH PROGRAM TRADERS:

In the presence of program traders the market clearing equation becomes

$$1 = (1-\rho) \cdot D_*^n \cdot \frac{F_k^n}{X_k^n} + \rho \cdot \phi^n(t_k^n, X_k^n). \tag{2.8}$$

In order to ensure existence of a unique equlibrium we have to make an additional technical assumption. It will turn out later that this is mainly a restriction on the market weight ρ of program traders (see Section 2.4 below).

Assumption 2.2.6 *There exists a constant* $\kappa > 0$ *such that for each* n

$$1 - \rho \cdot \phi^n(t,x) - \rho \cdot x \frac{\partial \phi^n}{\partial x}(t,x) \geq \kappa \qquad \text{for all } t \geq 0 \text{ and } x \in \mathbb{R}_{++}. \tag{2.9}$$

Moreover, we require (2.9) also to hold for the limit function ϕ.

Proposition 2.2.7 *Given Assumption (2.2.6), there exists for every market weight $\rho \in (0,1)$ a unique equilibrium in the economy specified in Assumptions (2.2.3) and (2.2.2).*

REMARKS: As the proof of Proposition 2.2.7 below shows, condition (2.9) guarantees that in equilibrium the aggregate demand function is strictly *decreasing* in x, such that equilibrium prices will depend smoothly on the reference trader's income F_k^n. This rules out price jumps and "crashes" of the kind discussed by Gennotte and Leland (1990) and Schönbucher (1993). Moreover, (2.9) ensures that the equilibrium is stable under the usual Walrasian tatônnement process. We believe that this is an important feature of our model, which contrasts with the analysis of Platen and Schweizer (1994). In order for them to explain volatility smiles by feedback effects from dynamic hedging they have to assume an (excess) demand function for their reference traders that is *increasing* in price. Not only does this result in the equilibrium being unstable, it even gives rise to arbitrage opportunities for the program traders in the sense of Jarrow (1994).

PROOF: Existence follows from continuity of the demand functions and

$$\begin{aligned} \limsup_{x\to\infty}\left((1-\rho)D^n(f,x)+\rho\phi^n(t,x)\right) &\leq \rho < 1, \\ \lim_{x\to 0}\left((1-\rho)D^n(f,x)+\rho\phi^n(t,x)\right) &= +\infty \end{aligned}$$

for all $t, f > 0$. For uniqueness it is sufficient to show that whenever t, f and x solve the market clearing equation (2.8), we have $\frac{\partial}{\partial x}\left((1-\rho)D^n(f,x)+\rho\phi^n(t,x)\right) < 0$. A direct computation using Lemma 2.2.4 gives

$$\frac{\partial}{\partial x}\left((1-\rho)D^n(f,x)+\rho\phi^n(t,x)\right) = -\frac{1}{x}\cdot\left((1-\rho)D^n(f,x)-\rho x\frac{\partial\phi^n}{\partial x}(t,x)\right).$$

The term in the brackets on the right-hand side equals $1-\rho\phi^n(t,x)-\rho x\frac{\partial\phi^n}{\partial x}(t,x)$ in equilibrium, which is positive by Assumption (2.2.6). □

2.3 The Continuous-Time Model

In order to get a clearer picture of the equilibrium price process and in particular of its volatility structure, we will now pass to the limiting continuous-time model. This also brings us closer to the original Black-Scholes model. To maintain a maximal level of generality we turn back to the situation described in the beginning of Section 2.2. That is, we assume a general aggregate demand function of the form (2.1) satisfying Assumption (2.2.1). For each $n = 1, 2, \dots$ let $(X_k^n)_{k=0,1,\dots}$ be the unique equilibrium price process, i.e. solution to the market clearing equation (2.2), which we know exists by Assumption (2.2.1) (1).

In order to formulate our results we need a common base space on which the distributions of all these processes can be compared. Let $\mathcal{D}^d[0,\infty)$ denote the d-dimensional *Skorohod-Space*; cf. Jacod and Shiryaev (1987). We identify any sequence $\xi_0^n, \xi_1^n, \dots$ defined for times $t_0^n, t_1^n, \dots$ with the RCLL function

$$\xi_t^n := \sum_{k=0}^{\infty} \xi_k^n \cdot 1_{\{t_k^n \le t < t_{k+1}^n\}}.$$

Let $(X_t^n)_{t\ge0}$ and $(F_t^n)_{t\ge0}$ denote the RCLL versions of $(X_k^n)_{k=0,1,\dots}$ and $(F_k^n)_{k=0,1,\dots}$, respectively. For the passage to the limit we require the state variable processes $(F_k^n)_{k=0,1,\dots}$ to converge to a continuous-time limit. Remember that for the CRRA case without program traders the equilibrium price process $(X_k^n)_{k=0,1,\dots}$ is proportional to $(F_k^n)_{k=0,1,\dots}$; cf. Lemma 2.2.5. Since our objective is to study the effect of hedging in a Black-Scholes type environment we assume that the limiting state variable process is a geometric Brownian Motion.[4]

[4]The main result of this section, Theorem 2.3.2, easily carries over to more general diffusion processes.

Assumption 2.3.1 *Suppose that the sequence $\{F^n\}_{n=1,2,\ldots}$ of state variable processes converges in distribution to a geometric Brownian Motion with constant drift and volatility parameters μ and η, respectively.*

We are now ready to state the main result of this section:

Theorem 2.3.2 *Suppose the sequences $\{G^n\}_{n=1,2,\ldots}$ and $\{F^n\}_{n=1,2,\ldots}$ of aggregate demand functions and state variable processes satisfy Assumptions (2.2.1) and (2.3.1), respectively.*

Then the sequence $\{X^n\}_{n=1,2,\ldots}$ of equilibrium price processes converges in distribution, and the limit distribution is uniquely characterised as the law of the solution $(X_t)_{t\geq 0}$ of the SDE

$$\begin{aligned} X_t &= X_0 - \int_0^t \left(\frac{\frac{\partial G}{\partial f}(s, F_s, X_s)}{\frac{\partial G}{\partial x}(s, F_s, X_s)} \eta F_s \right) dW_s \qquad (2.10) \\ &- \int_0^t \left(\frac{\frac{\partial G}{\partial t}(s, F_s, X_s)}{\frac{\partial G}{\partial x}(s, F_s, X_s)} + \frac{\frac{\partial G}{\partial f}(s, F_s, X_s)}{\frac{\partial G}{\partial x}(s, F_s, X_s)} \mu F_s - \frac{1}{2} H(s, F_s, X_s) \eta^2 F_s^2 \right) ds, \end{aligned}$$

where $(W_t)_{t\geq 0}$ is a standard Wiener process, $(F_t)_{t\geq 0}$ is just short for $F_t = \psi^{-1}(t, X_t)$, and H is a smooth function that depends only on first and second order derivatives of G. In particular, the instantaneous volatility of the process $(X_t)_{t\geq 0}$ at time t is given by

$$v(t, X_t) \cdot \eta := -\frac{\frac{\partial G}{\partial f}(t, \psi^{-1}(t, X_t), X_t)}{\frac{\partial G}{\partial x}(t, \psi^{-1}(t, X_t), X_t)} \cdot \frac{\psi^{-1}(t, X_t)}{X_t} \cdot \eta. \qquad (2.11)$$

Proof: By Assumption (2.2.1) (1) the equilibrium price process in economy n is given by $X_t^n = \psi^n(\tau_t^n, F_t^n)$, where $\tau_k^n = t_k^n$ if $t_k^n \leq t < t_{k+1}^n$. We first show that the ψ^n converge:

Lemma 2.3.3 *The sequence $\{\psi^n\}_{n=1,2,\ldots}$ converges uniformly on compacts to the smooth function ψ defined by the relation $G(t,f,\psi(t,f)) \equiv 1$.*

Smoothness of ψ follows from the Implicit Function Theorem and Assumption (2.2.1) (3). Convergence is shown in Appendix 2.A. By Assumption (2.3.1) the sequence $\{F^n\}_{n=1,2,\ldots}$ converges in distribution by to a process $(F_t)_{t\geq 0}$ satisfying

$$F_t = F_0 + \int_0^t \eta F_s \, dW_s + \int_0^t \mu F_s \, ds \tag{2.12}$$

for some standard Wiener process $(W_t)_{t\geq 0}$. A version of the Continuous Mapping Theorem then implies convergence in distribution of the triplet $(\tau_t^n, F_t^n, \psi^n(\tau_t^n, F_t^n))$ to $(t, F_t, \psi(t, F_t))$. To characterise the limiting distribution on $\mathcal{D}^3\,[0,\infty)$ we apply Itô's Lemma to $\psi(t, F_t)$ and use (2.12) to obtain

$$\begin{aligned} \psi(t,F_t) &= \psi(0,F_0) + \int_0^t \left(\frac{\partial\psi}{\partial f}(s,F_s)\eta F_s\right) dW_s \\ &\quad + \int_0^t \left(\frac{\partial\psi}{\partial t}(s,F_s) + \frac{\partial\psi}{\partial f}(s,F_s)\mu F_s + \frac{1}{2}\frac{\partial^2\psi}{\partial f^2}(s,F_s)\eta^2 F_s^2\right) ds. \end{aligned}$$

By differentiating the defining equation $G(t,f,\psi(t,f)) \equiv 1$, the derivatives of ψ can be expressed in terms of those of G, proving that $X_t = \psi(t,F_t)$ indeed solves the SDE (2.10). Expression (2.11) for the volatility is obtained by simply plugging $F_t = \psi^{-1}(t,X_t)$ back into (2.10). To complete the proof note that the drift and dispersion functions in equation (2.10) are smooth by assumption and thus locally Lipschitz. This implies pathwise uniqueness of (2.10) and hence uniqueness in distribution. □

We now relate the concrete specification with CRRA utility as outlined in Assumptions (2.2.3) and (2.2.2) to the general situation characterised by Assumption (2.2.1), and deduce the shape of the volatility in this special case.

Corollary 2.3.4 *Suppose that the sequence of discrete-time economies specified in Assumptions (2.2.3) and (2.2.2) satisfies in addition Assumption (2.2.6). If the D_*^n from (2.5) converge to some D_* as $n \longrightarrow \infty$, then the corresponding sequence $\{G^n\}_{n=1,2,\dots}$ of demand functions satisfies (2.2.1) with a limiting demand function G of the form*

$$G(t, f, x) = (1-\rho) \cdot D_* \cdot \frac{f}{x} + \rho \cdot \phi(t, x).$$

Hence, under (2.3.1) the corresponding sequence of price processes $\{X^n\}_{n=1,2,\dots}$ by Theorem 2.3.2 converges in distribution to a continuous-time diffusion process $(X_t)_{t\geq 0}$ whose instantaneous volatility at any time t is given by

$$v(t, X_t) \cdot \eta := \frac{1 - \rho\phi(t, X_t)}{1 - \rho\phi(t, X_t) - \rho X_t \frac{\partial \phi}{\partial x}(t, X_t)} \cdot \eta. \tag{2.13}$$

In particular, volatility is increasing *in the market weight ρ of program traders, bounded below by the "reference volatility" η and bounded above by η/κ. Note also that in the absence of program traders, i.e. when $\rho = 0$, $v(t, X_t) \equiv 1$. Market volatility then equals the volatility η of the exogenous state variable process $(F_t)_{t\geq 0}$.*

PROOF: The convergence of G^n to G is obvious by the assumed convergence of D_*^n to D_* and (2.2.2). It is easily seen that (2.2.6) implies (2.2.1) (1) and (3). Finally (2.2.1) (2) follows since the ψ^n and their first derivatives are uniformly bounded on compacts by (2.2.6) and (2.2.2) (1). The form of the volatility function is derived in Appendix 2.A. □

REMARKS: Of course, when taken literally, the overlapping generations scenario of Section 2.2.1 does no longer make sense when passing to the continuous-time limit, since the lifespan of every generation becomes infinitesimal. However, in the specific case considered there, we view the discrete-time model as the main focus of our interest, and the continuous-time model merely as a useful tool for analysing its properties. The reason for taking this view is that we want to contrast the study of Brennan and Schwartz (1989) by a temporary equilibrium model in which investors' decisions are driven by myopic considerations. One could, for instance, fix a time discretisation, i.e. an index $N \in \mathbb{N}$, and keep the reference traders' demand functions fixed: $D^n \equiv D^N$ for all $n \in \mathbb{N}$. The convergence results, Theorem 2.3.2 and Corollary 2.3.4, provide the necessary link between the discrete-time and the continuous-time model, thus permitting the use of the limiting diffusion as a tool for the analysis. Note that the above construction does not necessarily imply that we restrict the program traders to the same fixed discretisation. Their strategy functions may converge to continuous trading as $n \longrightarrow \infty$.

In our model, the Brownian Motion driving the dynamics of the state variable process $(F_t)_{t \geq 0}$ is the only source uncertainty. However, had we allowed the aggregate hedging function ϕ or the weight ρ to depend on some exogenous uncertainty, we would typically have ended up with a "Stochastic Volatility Model". Such models recently have become a focus of attention, see for example Hull and White (1987), Föllmer and Schweizer (1990), or Ball and Roma (1994).

2.4 Feedback-Effects from Black-Scholes Trading

Price dependent volatility—as generated by dynamic hedging in our model—causes major problems in practical applications of option pricing theory. Although hedging strategies may still be shown to exist they can in most cases no longer be calculated explicitly. This is why in practice most investors base their trading on the classical Black-Scholes Formula, which postulates constant volatility. In this section we therefore study in more detail the feedback effect generated by the corresponding strategies and analyse the extent to which they are still appropriate, when the effect of their implementation on prices is taken into account.

We work directly in the limiting diffusion model, because the explicit expression for the volatility faciliates the analysis. Therefore we contend ourselves with specifying properties only of the limiting demand function G which we assume to be of the form $G = (1-\rho)D + \rho\phi$ with ϕ being a mixture of Black-Scholes trading strategies. We remark, however, that under Assumption (2.2.2) the weak convergence of the equilibrium price processes implies the convergence of the corresponding gains from trade; see Duffie and Protter (1992). Hence our results on the performance of hedge strategies are also meaningful for the discrete-time models of Section 2.2.

2.4.1 Hedge Demand Generated by Black-Scholes Strategies

First we want to specify the strategy used by the representative program trader in more detail. As shown by Leland (1980), every convex payoff can be represented as the terminal value of a portfolio consisting of a mixture of European call options and a static position in the underlying. Therefore we concentrate on such portfolios. Consider first the problem of replicating the payoff of one single call option with strike price K and maturity date T. As was shown in the seminal paper by Black and Scholes (1973), if the trading decisions are based on the assumption of the underlying asset price following a geometric Brownian Motion with constant volatility σ, the corresponding price at any time t is given by the solution $c(t, X_t)$ of the terminal value problem

$$\left(\frac{\partial}{\partial t} + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2}{\partial x^2}\right) c(t,x) \equiv 0, \qquad c(T,x) = [x-K]^+, \tag{2.14}$$

and the corresponding strategy is $\frac{\partial c}{\partial x}(t, X_t)$. We denote the price and strategy functionals for a fixed contract (K,T) by $C(\sigma, K, T-t, x)$ and $\varphi(\sigma, K, T-t, x)$, respectively.[5] The terminal value problem (2.14) is explicitly solvable and the strategy function is given by

$$\varphi(\sigma, K, \tau, x) := \mathcal{N}\left(\frac{\log x - \log K}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau}\right), \tag{2.15}$$

where $\mathcal{N}$ is the standard normal distribution function and τ denotes the time to maturity $T - t$. In the sequel we will refer to σ also as the *input volatility*, since it is the volatility used for the computation of hedging strategies.

We assume that the aggregate demand of our representative program trader is independent of time. This models the scenario of many program traders entering and leaving the market at random times, so that the average composition of payoffs being replicated is constant over time. Formally, the representative program trader's demand is given by $\rho\phi(\sigma, x)$ where ρ is the market weight and

$$\phi(\sigma, x) = a + \int_{\mathbb{R}^2_+} \varphi(\sigma, K, \tau, x)\, \nu(dK \otimes d\tau). \tag{2.16}$$

Here, a represents the static position in the underlying and ν is a measure on $\mathbb{R}^2_+$ that describes the distribution of strike prices K and times-to-maturity τ in the portfolio. For convenience we define $\Gamma(\sigma, x) := x\frac{\partial\phi}{\partial x}(\sigma, x)$.[6] We want to show that the feedback effects of portfolio insurance are mitigated if the distribution of strike prices and times to maturity is relatively heterogeneous. To this end we concentrate on the following extreme case:

---

[5]Note that both value and strategy function depend on current time t and maturity time T only via their difference, the time-to-maturity $\tau = T - t$.

[6]In the standard option pricing theory the price derivative of a strategy function is known as the strategy's "gamma", which motivates this notation.

Assumption 2.4.1 *ν has a smooth density with respect to the Lebesgue-measure, i.e. ν is of the form $\nu(dK \otimes d\tau) = g(K, \tau)\, dK \otimes d\tau$, where $g : \mathbb{R}_+ \times [0, \infty) \longrightarrow \mathbb{R}_+$ is a smooth density function having compact support in $\mathbb{R}_+ \times [0, \infty)$.*[7]

Next we want to verify that Assumption (2.4.1) ensures that for ρ sufficiently small there is a unique equilibrium in the economy with CRRA agents (i.e. that Assumption (2.2.6) holds). Observe first that on the single contract level the function $x \frac{\partial \varphi}{\partial x}(\sigma, K, \tau, x)$ explodes when $x \to K$ and $\tau \to 0$. This corresponds to the well-known fact that option hedging strategies require extremely large changes of the hedge portfolio when the option is at the money and close to maturity. Surprisingly, this problem disappears in the aggregate, if the distribution ν is non-singular. The following proposition shows that bounds on $\Gamma(x, \sigma)$ can be found that depend only on the degree of heterogeneity of the distribution ν:

Proposition 2.4.2 *Suppose that $\sigma > \underline{\eta}$ for some $\underline{\eta} > 0$. Under Assumption (2.4.1) we have the following estimates for all $x \in \mathbb{R}_{++}$:*

(i)
$$|\Gamma(\sigma, x)| \leq \int_0^\infty \int_0^\infty \left| \frac{\partial}{\partial K} \left(K g(K, \tau)\right) \right| dK\, d\tau,$$

(ii)
$$\left| \frac{\partial}{\partial \sigma} \Gamma(\sigma, x) \right| \leq \frac{2}{\underline{\eta}} \cdot \int_0^\infty \int_0^\infty \left| \frac{\partial^2}{\partial \tau \partial K} \left(\tau K g(K, \tau)\right) \right| dK\, d\tau.$$

The proof is given in Appendix 2.B. Because of (i) we can achieve Assumption (2.2.6) without further restrictions on the distribution ν simply by requiring the porfolio insurance weight ρ not to be too large. Calculations with some sample density functions have shown that any reasonable value for ρ can be permitted.

---

[7] Note that we explicitly allow $g(K, 0) > 0$ for some K, i.e. arbitrarily small times to maturity.

REMARK: By $dK = Kd(\log K)$ the expression on the right-hand side of (i) above can be interpreted as a measure of the heterogeneity of the distribution of logarithmic strike prices, averaged over the time to maturity τ. An inspection of equation (2.17) below reveals that the feedback effect of dynamic hedging on market volatility mainly manifests itself through the appearence of $\Gamma(\sigma, x)$ in the denominator of $v(\sigma, x)$. Hence by (i) we see that this "disturbance" is controlled by the degree of heterogeneity of ν. This is most apparent in the economy with CRRA agents. Here we get from Corrollary 2.3.4

$$v(\sigma, x) = \frac{1 - \rho\phi(\sigma, x)}{1 - \rho\phi(\sigma, x) - \rho\Gamma(\sigma, x)} \leq \frac{1 - \rho}{1 - \rho - \rho \sup \Gamma(\sigma, x)},$$

such that even the maximal increase in volatility is controlled by the degree of heterogeneity of ν.

2.4.2 Rational Black-Scholes Trading

We now investigate the extent to which the Black-Scholes formula is still appropriate for the design of hedge strategies in our setting. We work with a limiting demand function of the form[8]

$$G(\sigma, f, x) = (1 - \rho) \cdot D(f, x) + \rho \cdot \phi(\sigma, x),$$

where $\phi(\sigma, x)$ is as in (2.16). Throughout the rest of this section we require only Assumptions (2.2.1) and (2.4.1) to hold. In particular we do not require the specific CRRA utility. By Theorem 2.3.2 the volatility of the limiting diffusion is

$$v(\sigma, x) \cdot \eta = -\frac{(1 - \rho)\frac{\partial D}{\partial f}(\psi^{-1}(x), x)}{x \cdot (1 - \rho)\frac{\partial D}{\partial x}(\psi^{-1}(x), x) + \rho\Gamma(\sigma, x)} \cdot \psi^{-1}(x) \cdot \eta. \tag{2.17}$$

[8]In the subsequent analysis σ is a parameter which does not vary with f or x. Its appearence does not alter the validity of Theorem 2.3.2.

As a first step, we use an idea of El Karoui, Jeanblanc-Picqué, and Shreve (1998) to derive a formula for the "tracking error". This number measures the difference between the *actual* and the *theoretical* value of a self-financing hedge portfolio for a European call calculated from the Black-Scholes formula with constant volatility σ. Recall that the theoretical value is given by $C_t := C(\sigma, K, T-t, X_t)$. The actual value V_t of the self-financing portfolio defined by initially investing $V_0 = C_0$ and holding $\varphi(\sigma, K, T-t, X_t)$ shares of the underlying at any time $t \le T$ is given by the cumulative gains from trade, i.e.

$$V_t = V_0 + \int_0^t \varphi(\sigma, K, T-s, X_s)\, dX_s.$$

The *tracking error* e_t is then defined as the difference between actual and theoretical value:

$$e_t := V_t - C_t.$$

Since $C_T = [X_T - K]^+$, e_T measures the deviation of the hedge portfolio's terminal value from the payoff it is supposed to replicate. In particular, if the tracking error is always positive, the terminal value of the hedge portfolio of an investor following the strategy $\varphi(\sigma, K, T-t, X_t)$ always completely covers the option's payoff. The following is a simplified version of Theorem 6.2 in El Karoui, Jeanblanc-Picqué, and Shreve (1998).

Proposition 2.4.3 *Suppose that the underlying asset's price follows a diffusion with volatility (2.17). Then the tracking error for a single option is given by*

$$e_t = \frac{1}{2}\int_0^t \left(\sigma^2 - \eta^2 v^2(\sigma, X_s)\right) \cdot X_s^2 \frac{\partial^2 C}{\partial x^2}(\sigma, K, T-s, X_s)\, ds. \tag{2.18}$$

In particular, if $\sigma \ge \eta v(\sigma, x)$, the tracking error is always positive.

PROOF: By Itô's Lemma,

$$C_t = \underbrace{C_0 + \int_0^t \varphi(\sigma, K, T-s, X_s) dX_s}_{=V_t}$$
$$+ \int_0^t \left(\frac{\partial C}{\partial t}(\sigma, K, T-s, X_s) + \frac{1}{2}\eta^2 v^2(\sigma, X_s) \cdot X_s^2 \frac{\partial^2 C}{\partial x}(\sigma, K, T-s, X_s) \right) ds.$$

Substituting (2.14) into the above equation yields the desired expression for the tracking error. Moreover, $C(\sigma, K, T-t, x)$ being convex in x, its second derivative is always positive. Hence by (2.18) the sign of the tracking error is entirely determined by the sign of the volatility difference $\sigma^2 - \eta^2 v^2(\sigma, x)$. □

Proposition 2.4.3 shows that in the presence of the feedback effect, simple Black-Scholes hedging strategies, (i.e. strategies based on the assumption of constant volatility,) are no longer sufficient to perfectly replicate a derviative security's payoff: the tracking-error is almost surely non-zero. Given this observation, there are then two possible directions in which one could proceed. First, retaining the restriction to the class of Black-Scholes strategies, one can study the tracking error and its determinants, thus addressing the question how well hedging performs *within* this class. Second, one can ask whether the class of strategies can be enlarged in such a way that perfect hedging becomes possible again. Or, in other words, whether rational agents who understand the feedback effect caused by their trading, could fully take it into account when designing their hedging strategies.

In this paper we take the first route, for the following reasons. The replication problem in finitely elastic markets becomes very complex, since the feedback effect on volatility introduces non-linearities in the partial differential equation that has to be solved in order to obtain hedging strategies. Not only does solving this problem exceed the scope of this paper, it is also true that practitioners rely almost exclusively on Black-Scholes type models. We therefore believe it a valid and important question to ask how well Black-Scholes hedging can perform in the presence of a feedback effect, even if perfect replication is impossible. However, solving the replication problem in finitely elastic markets is a very promising area for future research.

As we have just seen, if the terminal value of the hedge portfolio is to be no smaller than the payoff it is supposed to cover, the input volatility σ used for the computation of the hedging strategy must be no smaller than the actual market volatility. On the other hand, following the Black-Scholes strategy corresponding to a certain input volatility σ requires an initial investment of $C(\sigma, K, T, X_0)$. Since C is increasing in σ, to keep the initial "over-investment" as low as possible, investors should seek to find the smallest such σ. This motivates the following

Definition 2.4.4 *The constant $\overline{\sigma}$ is called a* super-volatility *if $\overline{\sigma}$ is the smallest positive solution of the equation*

$$\sigma = \sup\{\, \eta v(\sigma, x) \,:\, x \in \mathbb{R}_{++} \,\}. \tag{2.19}$$

Note the recursive structure: Since the choice of the input volatility and hence of the trading strategy affects the actual volatility, σ appears on both sides of (2.19). It will turn out that sufficient for the existence of a super-volatility is the following

Assumption 2.4.5 *The volatility function (2.17) has the following properties:*

1. *There are constants $0 < \underline{\eta} < \overline{\eta} < \infty$ so that $\underline{\eta} \leq v(\sigma, x) \leq \overline{\eta}$, $\forall\, \sigma \in [\underline{\eta}, \overline{\eta}]$, $\forall\, x \in \mathbb{R}_{++}$.*
2. *There is a positive constant β so that $\eta \cdot \frac{\partial v}{\partial \sigma}(\sigma, x) \leq 1 - \beta$, $\forall\, \sigma \in [\underline{\eta}, \overline{\eta}]$, $\forall\, x \in \mathbb{R}_{++}$.*

Essentially (ii) means that variations in the input volatility do not affect the actual market volatility too much. Since $\frac{\partial}{\partial \sigma}\Gamma(\sigma, x)$ is bounded according to Proposition 2.4.2, this assumption holds as long as ρ is not too large. Of course to check if (i) is satisfied one has to know the function D. In case of the economy with CRRA agents (i) is implied by Assumption (2.2.6).

Proposition 2.4.6 *Suppose Assumption (2.4.5) holds. Then the super-volatility in the sense of Definition 2.4.4 exists and is given by*

$$\overline{\sigma} := \sup\{\sigma^*(x) \, : x \, \in \mathbb{R}_{++}\},$$

where $\sigma^*(x)$ is the unique solution of the fixed point problem $\eta v(\sigma, x) = \sigma$.

PROOF: Assumption (2.4.5) implies that the equation $\eta v(\sigma, x) - \sigma = 0$ has a unique solution $\sigma^*(x)$ for each $x > 0$. Now by definition $\overline{\sigma} \geq \sigma^*(x)$ for any $x \in \mathbb{R}_{++}$. Since by Assumption (2.4.5) (ii) the mapping $\sigma \mapsto \eta v(\sigma, x) - \sigma$ is strictly decreasing, we get

$$\eta v(\overline{\sigma}, x) - \overline{\sigma} \leq \eta v(\sigma^*(x), x) - \sigma^*(x) = 0$$

for all $x \in \mathbb{R}_{++}$, hence $\overline{\sigma}$ is indeed an upper bound for $\eta v(\sigma, x)$. It is also the smallest such bound. This is obvious in the case when the supremum in (2.19) is attained for some x, and it follows from Assumption (2.4.5) (ii) in the general case. To prove that $\overline{\sigma}$ is also the smallest solution to (2.19), simply note that by definition $\sigma^*(x)$ gets arbitrarily close to $\overline{\sigma}$. Hence for $\tilde{\sigma} < \overline{\sigma}$ there is always an x with $\tilde{\sigma} < \sigma^*(x)$ such that by monotonicity

$$\eta v(\tilde{\sigma}, x) - \tilde{\sigma} > \eta v(\sigma^*(x), x) - \sigma^*(x) = 0.$$

□

REMARK: From a practioner's viewpoint it is reasonable to use a Black-Scholes strategy based on the supervolatility $\overline{\sigma}$ for hedging purposes as long as the difference between $\overline{\sigma}$ and $\inf\{\eta v(\overline{\sigma}, x),\ x \in \mathbb{R}_{++}\}$ and hence the possible initial overinvestment is relatively small. As we will see in the simulations reported below, this issue, and hence the robustness of the Black-Scholes Formula with respect to the feedback from dynamic hedging, depends largely on the heterogeneity of the insured payoffs.

2.4.3 Comparison with the Brennan and Schwartz Study

Using explicit numerical computations (see below for details) we can compare our results to those obtained by Brennan and Schwartz (1989). Table 2.1 lists the ratios of the volatility in the presence of program traders to that in the benchmark economy without program traders. While the utility functions, and hence notably the risk aversion, of the reference traders are identical in both models, an inspection of Table 2.1 shows that the effects of program trading on volatility are much stronger in our model. As was explained in the introduction this is due to the different expectation formation of the reference traders in the two models.

2.4.4 Numerical Computations

First we computed the resulting volatility function $\eta h(\overline{\sigma}, x)$ as a function of x, using as hedging input the super-volatility $\overline{\sigma}$ for a variety of different weights ρ, different reference volatilities η, and different levels of heterogeneity of ν. Figures 2.1 and 2.2 show the dramatic effect of heterogeneity. Here, we graphed the reference volatility "o" and the resulting volatility "•" against the current price using a value of $\rho =$

10%.[9] All numerical results, including those not featured in this paper, support our findings from the qualitative analysis: Volatility increases with the market share ρ of portfolio insurance as well as with reference volatility η. Both the level of increase and the price dependency are reduced by heterogeneity.

We then ran Monte Carlo simulations to generate sample price paths and used the tracking error formula (2.18) to compute the terminal value of a hedging portfolio based on the super-volatility $\overline{\sigma}$. We compared the results to the payoff of the option it was supposed to duplicate, again for a variety of different parameter constellations. Figures 2.3 and 2.4 again capture the striking effect of heterogeneity. For every sample path we graphed the terminal value "$\bullet$" of the portfolio against the terminal price of the underlying. The straight line depicts the exact option payoff. Again all results we obtained strongly support our qualitative findings: The tracking error is largest around the option's striking price and almost vanishes as the option gets deeper in the money or out of the money. We see that even a comparatively low level of heterogeneity is sufficient for the super-hedging portfolio to duplicate the option's payoff almost perfectly.

---

[9]The fraction of the aggregate equity value subject to formal portfolio insurance prior to the events of October 1987 was approximately 5%. However, one should bear in mind that the amount of "informal" portfolio insurance may have amounted to considerabely more than this. Moreover, part of the aggregate equity supply is held because of the associated control rights and not for speculative reasons, such that the actual "supply" should be considered smaller than aggregate equities. Hence the "actual" market weight of program trading might be larger than just these 5%.

2.5 Conclusions

In this paper we have analysed the feedback effect of dynamic hedging strategies on the equilibrium price process of the underlying asset in an economy where the market for the latter is only finitely elastic. We gave an explicit expression for the transformation of market volatility which allowed us to carry out a detailed study of the feedback effects caused by dynamic hedging. A comparison with the analysis of Brennan and Schwartz revealed the importance of agents' expectations in determining market liquidity and hence the amplitude of the feedback effect on volatility. Adding to the existing literature, we identified heterogeneity of the distribution of hedged contracts as one of the key determinants for the transformation of volatility. Moreover, we showed that simple hedging strategies derived from the assumption of constant volatility may still be appropriate even though their implementation obviously violates this assumption. However, investors might have to "over-invest" in their hedging strategies. To sum up, we find that classical Black-Scholes theory is quite robust with respect to the feedback effects discussed, as long as the distribution of different payoff claims being hedged does not become too homogeneous. Nonetheless future research is needed to extend the work of Jarrow (1994) on option pricing theory in an economy where agents' hedging strategies affect the underlying asset's price process.

2.A Appendix: Complements to Section 2.3

PROOF OF LEMMA 2.3.3: First we prove the pointwise convergence of the ψ^n to ψ. For t and f fixed, every subsequence of the sequence $x_n := \psi^n(t, f)$ contains a further subsequence $\{x_{n_j}\}_{j=1,2,...}$ which converges to some $x \in \mathbb{R}_+$ by Assumption

(2.2.1) (2). Now we have the following estimate:

$$\begin{aligned}
&\left|G^{n_j}(t,f,x_{n_j}) - G(t,f,x)\right| \\
&\leq \left|G^{n_j}(t,f,x_{n_j}) - G(t,f,x_{n_j})\right| + \left|G(t,f,x_{n_j}) - G(t,f,x)\right|.
\end{aligned}$$

The first term on the right-hand side tends to zero for $j \longrightarrow \infty$ because G^n converges to G uniformly on compacts. The second term tends to zero because of the uniform continuity of G on compacts. It follows that $G(t,f,x) = \lim_j G^{n_j}(t,f,x_{n_j}) = 1$ and hence $x = \psi(t,f)$. Now by Assumption (2.2.1) (2) the sequence $\{\psi^n\}_{n=1,2,\ldots}$ also converges uniformly on compacts to ψ. □

PROOF OF COROLLARY 2.3.4: Using (2.11) and the particular form of the limiting demand function together with $F_t = \psi^{-1}(t, X_t)$ we get

$$v(t,X_t) = -\frac{\frac{\partial G}{\partial f}(t,\psi^{-1}(t,X_t),X_t)}{\frac{\partial G}{\partial x}(t,\psi^{-1}(t,X_t),X_t)} \cdot \frac{\psi^{-1}(t,X_t)}{X_t} = \frac{-(1-\rho)D_* \cdot \frac{F_t}{X_t}}{-(1-\rho)D_* \cdot \frac{F_t}{X_t} + \rho X_t \frac{\partial \phi}{\partial x}(t,X_t)}.$$

Using the market clearing equation $1 \equiv (1-\rho)D_* \frac{F_t}{X_t} + \rho\phi(t,X_t)$ to substitute for $(1-\rho)D_* \frac{F_t}{X_t}$ in the above expression gives the desired shape of the volatility of the limiting diffusion. □

2.B Appendix: Complements to Section 2.4

ESTIMATES FOR $\Gamma(\sigma, x)$: Observe first that by (2.15) we have

$$x \frac{\partial}{\partial x} \varphi(\sigma, K, \tau, x) = -K \frac{\partial}{\partial K} \varphi(\sigma, K, \tau, x),$$

which implies:

$$\begin{aligned}
\Gamma(\sigma, x) &= \int_0^\infty \int_0^\infty x \frac{\partial}{\partial x} \varphi(\sigma, K, \tau, x) g(K, \tau)\, dK\, d\tau \\
&= -\int_0^\infty \int_0^\infty K \frac{\partial}{\partial K} \varphi(\sigma, K, \tau, x) g(K, \tau)\, dK\, d\tau \\
&= \int_0^\infty \int_0^\infty \varphi(\sigma, K, \tau, x) \frac{\partial}{\partial K} \left(K g(K, \tau)\right) dK\, d\tau
\end{aligned}$$

by partial integration and the assumption of g having compact support. But since $0 \leq \varphi \leq 1$ by (2.15), this implies

$$|\Gamma(\sigma, x)| \leq \int_0^\infty \int_0^\infty \left| \frac{\partial}{\partial K} \left(K g(K, \tau)\right) \right| dK\, d\tau.$$

ESTIMATES FOR $\frac{\partial}{\partial \sigma}\Gamma(\sigma, x)$: Observe first that by (2.15) we have

$$\sigma \frac{\partial}{\partial \sigma} \varphi(\sigma, K, \tau, x) = 2\tau \frac{\partial}{\partial \tau} \varphi(\sigma, K, \tau, x),$$

which together with the results from the previous paragraph implies

$$\begin{aligned}
\frac{\partial}{\partial \sigma}\Gamma(\sigma, x) &= \int_0^\infty \int_0^\infty \frac{\partial}{\partial \sigma} \varphi(\sigma, K, \tau, x) \frac{\partial}{\partial K} \left(K g(K, \tau)\right) dK\, d\tau \\
&= \int_0^\infty \int_0^\infty 2 \frac{\tau}{\sigma} \frac{\partial}{\partial \tau} \varphi(\sigma, K, \tau, x) \frac{\partial}{\partial K} \left(K g(K, \tau)\right) dK\, d\tau \\
&= -\frac{2}{\sigma} \int_0^\infty \int_0^\infty \varphi(\sigma, K, \tau, x) \frac{\partial^2}{\partial \tau \partial K} \left(\tau K g(K, \tau)\right) dK\, d\tau
\end{aligned}$$

again by partial integration and the assumption of g having compact support. But since $0 \leq \varphi \leq 1$ by (2.15) and furthermore $\sigma \geq \underline{\eta}$, this implies

$$\left| \frac{\partial}{\partial \sigma}\Gamma(\sigma, x) \right| \leq \frac{2}{\underline{\eta}} \cdot \int_0^\infty \int_0^\infty \left| \frac{\partial^2}{\partial \tau \partial K} \left(\tau K g(K, \tau)\right) \right| dK\, d\tau.$$

	Program Trader's Market Weight					
	1 %		5 %		10 %	
Price	F-S	B-S	F-S	B-S	F-S	B-S
80	1.01	1.00	1.06	1.02	1.13	1.04
90	1.02	1.00	1.08	1.02	1.18	1.04
100	1.02	1.00	1.09	1.02	1.19	1.05
110	1.02	1.00	1.08	1.02	1.18	1.05
120	1.01	1.00	1.06	1.02	1.14	1.05

Table 2.1: Comparison with the Results of Brennan and Schwartz.

This table shows the ratios of actual to reference volatility in our model (F-S) and the Brennan-Schwartz model (B-S) for different market weights of the program trader.

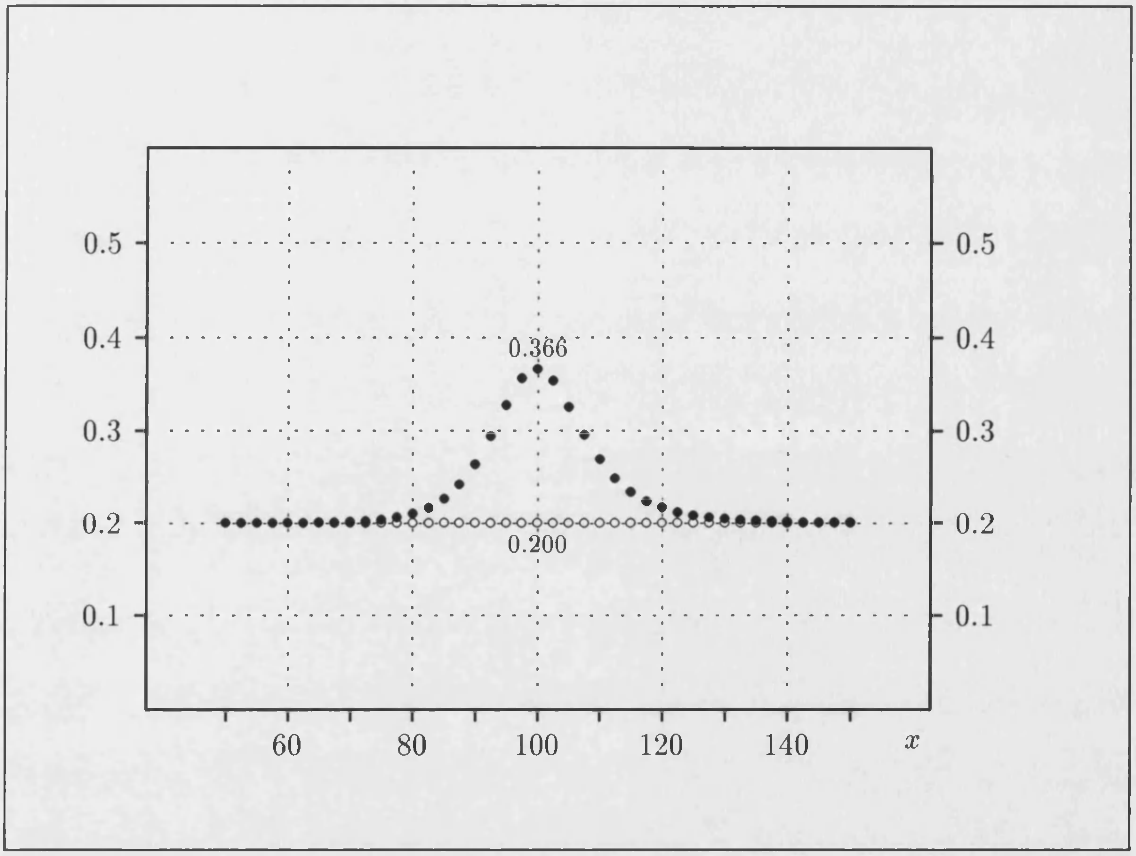

Figure 2.1: Volatility if Hedged Contracts are Concentrated.

This graph shows the actual market volatility (•) compared to reference volatility (∘) when hedging is based on the super-volatility $\overline{\sigma}$. *The variances of strike prices and times-to-maturity are* 0.0625, *i.e. the distribution is highly concentrated. The market weight of portfolio insurance is* $\rho = 10\%$.

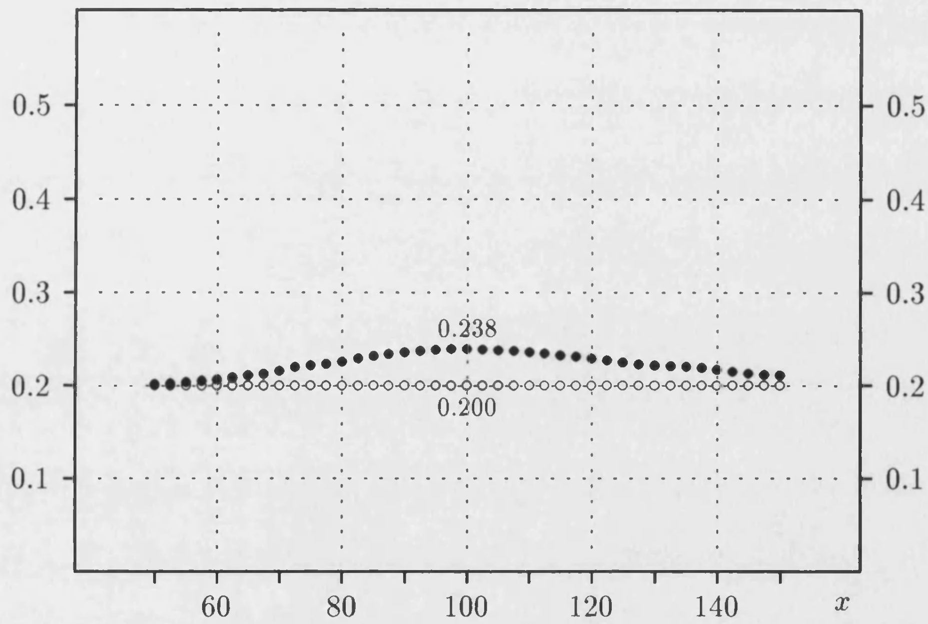

Figure 2.2: Volatility if Hedged Contracts are Heterogeneous.

This graph shows the actual market volatility (•) compared to reference volatility (○) when hedging is based on the super-volatility $\overline{\sigma}$. *The variances of strike prices and times-to-maturity are* 0.5, *i.e. the distribution is relatively heterogeneous. The market weight of portfolio insurance is* $\rho = 10\%$.

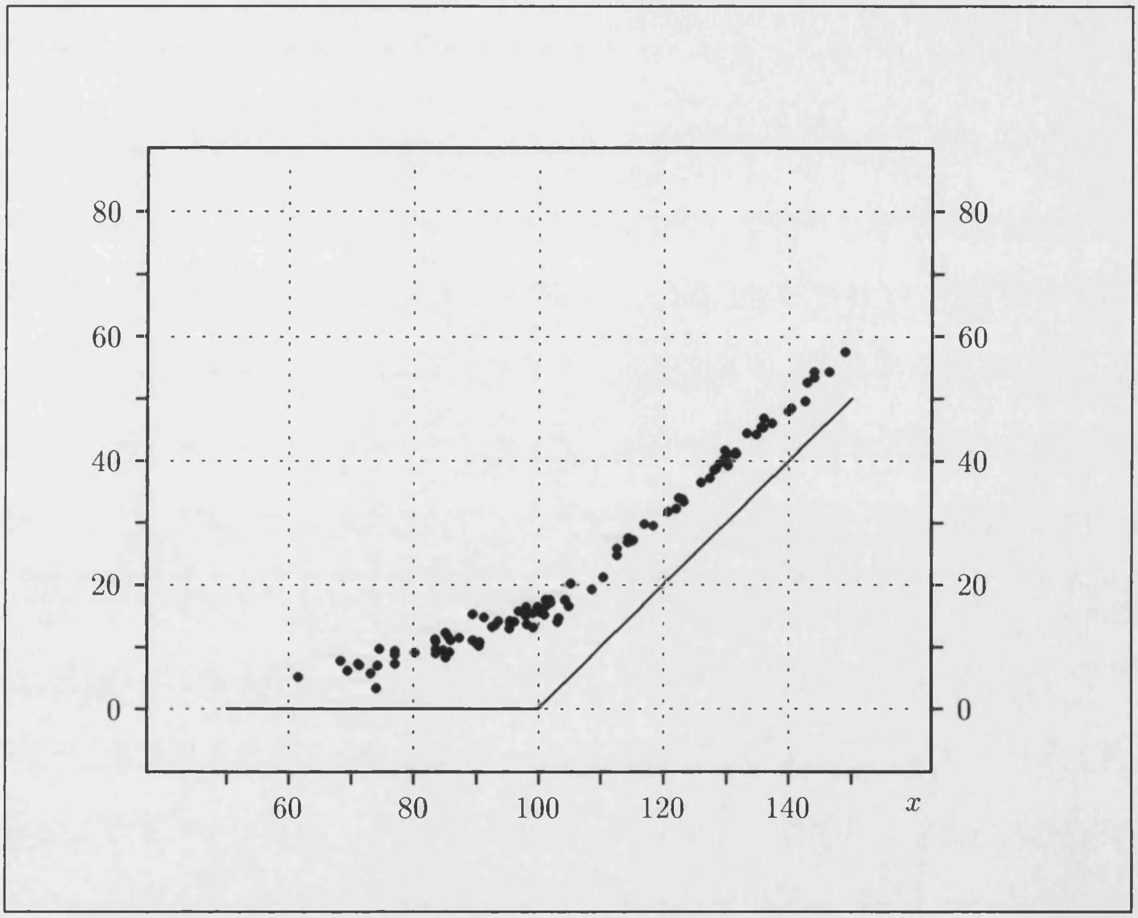

Figure 2.3: Tracking Error if Hedged Contracts are Concentrated.

This graph shows for 100 *sample paths the terminal value (•) of a portfolio designed to hedge a European call with strike price* 100.0*, based on the super-volatility* $\overline{\sigma}$*, compared to the option's pay-off itself. The variances of strike prices and times-to-maturity are* 0.0625*, i.e. the distribution is highly concentrated. The market weight of portfolio insurance is* $\rho = 10\%$.

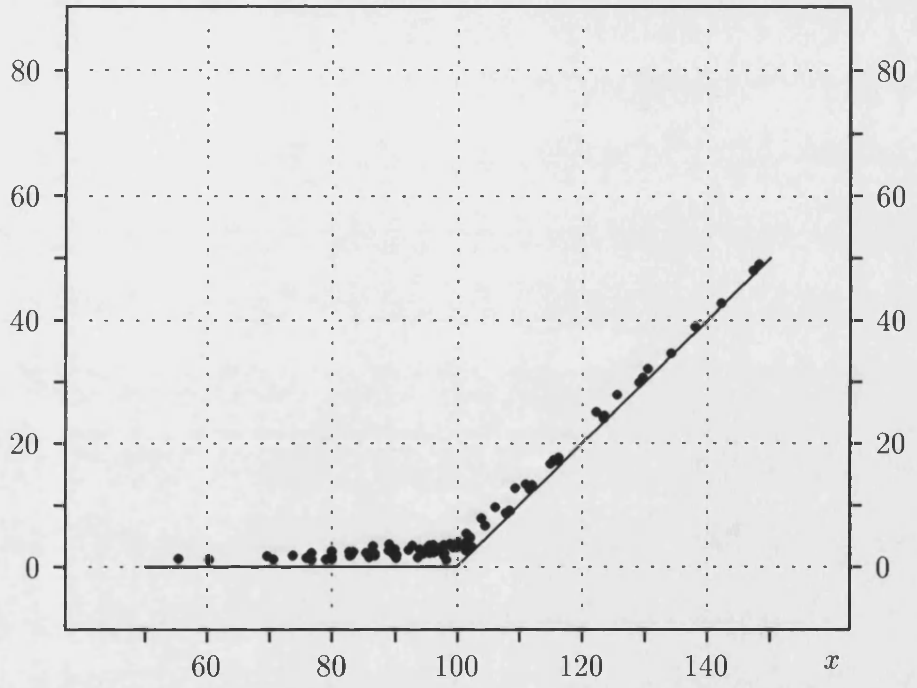

Figure 2.4: Tracking Error if Hedged Contracts are Heterogeneous.

This graph shows for 100 *sample paths the terminal value (•) of a portfolio designed to hedge a European call with strike price* 100.0*, based on the super-volatility* $\overline{\sigma}$*, compared to the option's pay-off itself. The variances of strike prices and times-to-maturity are* 0.5*, i.e. the distribution is relatively heterogeneous. The market weight of portfolio insurance is* $\rho = 10\%$.

Chapter 3

Technical Trading in Foreign Exchange Markets

3.1 Introduction

Technical Trading has always enjoyed a higher reputation amongst practitioners than amongst economists. This is so mainly because the theory of efficient markets leaves no scope for profits to be made on basis of technical analysis of historical price patterns. If money can be made by predicting future returns based on past prices only, one should expect market participants to take advantage of this opportunity, causing current prices to adjust thus eradicating any such profit opportunity. However, two important facts have brought technical trading into the focus of attention of a growing strand of economic research. First, technical trading and its inherent positive feedback effect of past returns on current prices have been blamed for causing or at least aggravating some of the market crashes in recent history; as reported amongst others by the Group of Ten (1993). Second, empirical evidence shows that technical trading rules are indeed widely used by market participants, in particular on foreign exchange markets and, more importantly, the trading activity triggered by such rules has significant effects on prices; see for example Allen and Taylor (1992).

Two main questions arise in this context. First, if market participants do follow technical trading rules, what is the effect on equilibrium exchange rates? Second, can the use of technical trading rules be rationalised? Empirical evidence seems to suggest that profits can indeed be made by following technical trading rules, see Levich and Thomas (1993), Brock, Lakonishok, and LeBaron (1992), and Sweeney (1986). This chapter addresses these questions from a theoretical viewpoint. We consider an equilibrium model for exchange rates in which some agents follow technical trading rules. As analytical framework, we choose the discrete-time equilibrium model developed in Föllmer and Schweizer (1993). Technical traders interact with pure "noise traders", and exchange rates are determined by Walrasian equilibrium. To analyse the properties of the resulting equilibrium exchange rate process in more detail, we then pass to the continuous-time limit. It turns out that, as long as technical traders are active, bubbles will always emerge. In other words, exchange rate fluctuations caused by noise will be picked up by technical analysis and amplified by the buy or sell signals generated by the corresponding trading rules. However, whether these bubbles can grow into market crashes depends on how far traders are willing to go in following their technical trading rules. If there is a certain amount of "fundamentalism" in the market, i.e. traders who expect a trend reversal when exchange rates have moved too far from their fundamental level, bubbles will always burst.

There have been a number of theoretical studies analysing the effects of technical trading. In Day and Huang (1990), rational "information traders" interact with technical traders. Equilibrium is determined by a market maker who sets asset prices such as to maintain a target inventory. Unlike the model considered here, Day and Huang do not introduce noise. As a result, equilibrium prices in their model follow deterministic "chaotic" dynamics. The qualitative behaviour of equilibrium prices however is similar to that obtained here. In Lux (1995), traders act upon rumours and fads rather than technical analysis. Two types of agents, "optimistic" and "pessimistic" traders interact with one-another. Contagious mood swings cause equilibrium prices to fluctuate cyclically. While both these papers explain the emergence of price bubbles, neither addresses the rationality of agent's behaviour and the profitability of the resulting trading strategies.

Interestingly, despite that fact that technical trading obviously introduces a substantial amount of autocorrelation of returns, it can be shown that expected returns over any finite horizon are nonetheless zero. However, while no profits can be made using simple buy-and-hold strategies, the existence of technical trading causes technical trading rules themselves to be profitable. In other words, technical trading can be seen as a kind of "self-fulfilling prophecy": If market participants believe in technical trading rules, then they are indeed profitable. That is, although markets are efficient ex-ante, the use of technical trading rules can be rationalised by their ex-post profitability. These results are in line with the empirical findings of Levich and Thomas (1993), Brock, Lakonishok, and LeBaron (1992), and Sweeney (1986). Profitability of support and resistance rules was shown in Curcio and Goodhart (1992).

The remainder of this chapter is organised as follows. In Section 3.2 we introduce the model and the mathematical framework for our analysis. In the following section we formalise the demand arising from technical trading rules and discuss its properties. Section 3.4 analyses the structure of equilibria and the stochastic properties of the resulting equlibrium exchange rate process. The profitability of technical trading rules is studied in 3.4.4. Section 3.5 concludes.

3.2 The Model

We employ the discrete-time framework introduced in Föllmer and Schweizer (1993). We will model the dynamics of equilibrium exchange rates in discrete time and then pass to the continuous-time limit to facilitate the qualitative analysis (see Section 3.4). More specifically, we consider a sequence of discrete-time economies for $n = 1, 2, \dots$ with trading intervals $\Delta t^n > 0$. In a given economy n, trading takes place at times $t_k^n := k \cdot \Delta t^n$, $(k = 0, 1, \dots)$. In view of the intended passage to the continuous-time limit we furthermore assume that

$$\Delta t^n \longrightarrow 0 \qquad \text{as } n \to \infty.$$

TRADED ASSETS:

There are two assets in the economy, a riskless one (typically a domestic bond or money market account), and a risky one. We normalise the price of the riskless asset to one, thereby making interest rates implicit in our model. Since our focus in this chapter is the foreign exchange market, we think of the risky asset as a foreign currency, and refer to its price sometimes as *exchange rate*. Note that this is the inverse of the notion of exchange rate as it is used in the U.K., where it specifies the amount of foreign currency one can buy with one unit of domestic currency. We would like to stress that although we interpret our analysis in the context of foreign exchange markets, the results are also applicable to other scenarios such as for example equity markets. The price of the foreign currency at any trading date t_k^n is denoted X_k^n.

AGENTS:

There are two types of agents active in the market; *noise* traders and *information* or *technical* traders. At any time t_k^n, noise traders' aggregate demand for the foreign currency is given by a random variable δ_k^n, which we assume to be serially independent and independent of the price history.

TECHNICAL TRADERS:

There are m technical traders in the economy, denoted $i = 1, \ldots, m$. Each agent i is characterised by an element $\theta_i \in \Theta \subseteq \mathbb{R}^p$ of some p-dimensional parameter space, to be made more specific later. The θ_i are assumed to be drawn independently from a probability distribution μ on Θ. Agent i's trading behaviour is expressed in terms of a demand function for the foreign currency. More specifically, if at time t_k^n an exchange rate of X is quoted, agent i will demand an amount

$$\alpha(\theta_i)e_k^n(X;\theta_i)$$

of foreign currency. Here, $\alpha(\theta_i)$ is the "market weight" of agent i, and $e_k^n(X,\theta_i)$ is agent i's normalised demand function. Note that e_k^n can (and will) depend on the history of past prices, $X_{k-1}^n, X_{k-2}^n, \ldots$. The aggregate demand from all technical traders is then

$$\sum_{i=1}^{m} \alpha(\theta_i)e_k^n(X;\theta_i).$$

LOG-LINEAR DEMAND:

In order to be able to model explicitly the demand arising from technical trading rules, we assume individual agents' demand functions to be *log-linear*:

$$e_k^n(X;\theta_i) = \log \frac{\hat{X}_k^n(\theta_i)}{X}.$$

We interpret $\hat{X}_k^n(\theta_i)$ as agent i's view on what the exchange rate *should* be. The above demand specification then says that agent i will *buy* the foreign currency if its price is *below* $\hat{X}_k^n(\theta_i)$ and *sell* if it is above. We will refer to $\hat{X}_k^n(\theta_i)$ as agent i's *reference* *level* at time t_k^n. Note that, by the very nature of technical trading rules, $\hat{X}_k^n(\theta_i)$ will depend on the price history $X_{k-1}^n, X_{k-2}^n, \ldots$. We will formalise this in more detail later in Section 3.3. Although the assumption of log-linear demand functions seems arbitrary, it is very common in Monetary Economics, see for example Gourieroux, Laffont, and Monfort (1982), and is thus in line with our focus on foreign exchange markets. It will be convenient to express agents' reference levels $\hat{X}_k^n(\theta_i)$ in terms of *return projections*. More precisely, we define agent i's return projection $\hat{R}_k^n(\theta_i)$ at time t_k^n implicitly so that we can write

$$\hat{X}_k^n(\theta_i) = X_{k-1}^n \cdot \exp\left(\hat{R}_k^n(\theta_i) \cdot \Delta t^n\right).$$

Note that, by explicitly including the time interval Δt^n, we have effectively normalised $\hat{R}_k^n(\theta_i)$ to reflect the projected return *per unit of time*. We can now re-write the model in logarithmic terms. Defining $Y_k^n := \log X_k^n$, $\hat{Y}_k^n(\theta_i) := \log \hat{X}_k^n(\theta_i)$, etc., we can write agent's i's excess demand function in logarithmic form as

$$e_k^n(Y;\theta_i) = \left(Y_{k-1}^n - Y + \hat{R}_k^n(\theta_i) \cdot \Delta t^n\right).$$

Note that, in slight abuse of notation, we use the same symbol e_k^n for the demand function in logarithmic terms.

EQUILIBRIUM:

Denote by $\rho \in [0,1]$ the relative share of total demand attributable to technical trading. In the context of foreign exchange markets, it is appropriate to assume that the total supply of foreign currency is zero, reflecting the fact that all that matters is the *net* balance of foreign currency held by domestic traders versus domestic currency held by traders in the foreign country. The *equilibrium exchange rate* at any time t_k^n is then given as the solution X_k^n of the *market clearing equation*

$$\rho \cdot \sum_{i=1}^{m} \beta^m(\theta_i) e_k^n(X_k^n;\theta_i) + (1-\rho) \cdot \delta_k^n \equiv 0. \tag{3.1}$$

Here, $\beta^m(\theta_i)$ is the *normalised* relative market weight of agent i, defined as

$$\beta^m(\theta_i) := \frac{\alpha(\theta_i)}{\sum_{i=1}^{m} \alpha(\theta_i)}.$$

With the log-linear demand specification, the market clearing equation takes the form

$$\Delta Y_k^n = \sum_{i=1}^{m} \beta^m(\theta_i)\hat{R}_k^n(\theta_i)\Delta t^n + \frac{1-\rho}{\rho} \cdot \delta_k^n. \tag{3.2}$$

Note that δ_k^n is as yet unspecified, so that by suitable rescaling we can assume that the term $(1-\rho)/\rho$ equals one. In order to write the equilibrium equation in form of a Stochastic Difference Equation, we define the cumulative noise demand process as

$$Z_k^n := \delta_0^n + \cdots + \delta_k^n.$$

The market clearing equation then takes the form

$$\Delta Y_k^n = \sum_{i=1}^{m} \beta^m(\theta_i)\hat{R}_k^n(\theta_i) \cdot \Delta t^n + \Delta Z_k^n. \tag{3.3}$$

Since $\hat{R}_k^n(\theta_i)$ depends on the price history $Y_{k-1}^n, Y_{k-2}^n, \dots$, equation (3.3) for the (logarithmic) price process $(Y_k^n)_{k=0,1,\dots}$ can be seen as a Stochastic Delay Equation.

REMARK: For large m, the Law of Large Numbers implies

$$\sum_{i=1}^{m} \beta^m(\theta_i)\hat{R}_k^n(\theta_i) \longrightarrow \overline{R}_k^n := \int_\Theta \alpha(\theta)\hat{R}_k^n(\theta)\,\mu(d\theta)$$

We will later make use of this limiting argument to analyse the behaviour of the equilibrium exchange rate process. In the following Section 3.3, we will give examples of some explicit specifications of $\hat{R}_k^n(\theta_i)$ for a variety of technical trading rules, and derive a set of basic properties. In Section 3.4 we will then use these properties to characterise the dynamic behaviour of the equilibrium exchange rate process and analyse the profitability of different trading strategies. In the limit, for very large numbers of technical traders, the equilibrium equation hence takes the simple form

$$\Delta Y_k^n = \overline{R}_k^n \cdot \Delta t^n + \Delta Z_k^n.$$

In particular, in the absence of technical trading, we have $\Delta Y_k^n = \Delta Z_k^n$, i.e. exchange rate fluctuations are driven by noise alone. This implies

$$E_{k-1}[\, X_k^n \,] = X_{k-1}^n \cdot E[\, \exp(\Delta Z_k^n)\,].$$

In order to focus on the effect of technical trading, we assume that in the absence of technical trading, the expected return on the exchange rate is zero, $E_{k-1}[\, \Delta Z_k^n \,] = 1$. Thus, exchange rates in this case satisfy the *Efficient Market Hypothesis* in its weak form, i.e. $E_{k-1}[\, X_k \,] = X_{k-1}$.

3.3 Technical Trading Rules

In this section, we will formalise the demand arising from technical trading rules. Consider for a moment the situation where all agents believe in the *Efficient Market Hypothesis* in the weak form, i.e. the current price is the best estimate of next period's price. Formally, this would be expressed as $\hat{R}_k^n(\theta_i) = 0$ for all i. In this case, the market clearing equation (3.3) implies

$$\Delta Y_k^n = \Delta Z_k^n$$

In other words, $X_k^n = X_{k-1}^n \cdot \exp(\Delta Z_k^n)$, i.e. changes in price are purely driven by noise. By assumption, $E_{k-1}[\, \exp(\Delta Z_k^n)\,] = 1$, so that

$$E_{k-1}[\, X_k^n \,] = X_{k-1}^n$$

In other words, if agents believe in market efficiency, then it turns out that markets *are* in fact efficient. In view of this, "taking a position" in the market means diverting from the assumption of efficiency, which is formalised by a non-zero return projection $\hat{R}_k^n(\theta_i)$.

BOUNDED MEMORY:

Technical trading rules are typically based on past price patterns. Therefore, the return projection $\hat{R}_k^n$ at any time t_k^n depends on the *history* of past prices, $Y_{k-1}^n, Y_{k-2}^n, \ldots$ To facilitate the analysis, we will however require $\hat{R}_j^n$ to display *bounded memory*, i.e. to depend on price observations only up to a certain time in the past. We formalise this as follows:

Assumption 3.3.1 *We assume there exists a fixed depth of memory $T > 0$ and a* deterministic *function $R^n : \Theta \times \mathbb{R}^l \longrightarrow \mathbb{R}$ such that for any $\theta_i \in \Theta$ the return projection $\hat{R}_k^n(\theta_i)$ can be expressed as*

$$\hat{R}_k^n(\theta_i) = R^n\left(\theta_i; Y_{k-l}^n, \ldots, Y_{k-1}^n\right),$$

where l is the largest integer such that $l \cdot \Delta t^n \leq T$.

Note that l depends on n, but for notational convenience we omit the superscript n. It is also worth mentioning that we assume the function R^n to be independent of k, i.e. trading rules are only allowed to depend on price patterns but not directly on time. This assumption is made for computational convenience only and could easily be relaxed.

In other words, we assume that the return projection at time t_k^n depends only on past price realisations no longer than T in the past. Note in particular that we assume demand memory to be bounded uniformly accross all agents. Note also that we assume T to be independent of n, so that also the trading rules in the continuous-time limit will have a memory bounded by T. For notational convenience, we denote the relevant price history at time t_k^n by

$$\mathfrak{Y}_k^n := \left(Y_{k-l}^n, Y_{k-l+1}^n, \dots, Y_{k-1}^n\right)$$

so that we can write $\hat{R}_k^n(\theta_i) = R^n(\theta_i; \mathfrak{Y}_k^n)$. For the remainder of the chapter, we will express everything in terms of the generic function R^n, and we will omit the argument $\mathfrak{Y}_k^n$ whenever it is unambiguous.

3.3.1 Generic Trading Rules

Generically, the majority of technical trading rules can be described as a combination of a *signal* rule and a set of *actions*. The signal rule determines, as a function of past price patterns, when an action is to be taken. The action taken typically consists of taking a position in the market, i.e. buying or selling a certain quantity of foreign currency when the price is below or above some projection. For example, a simple Trend Chasing rule generates a "buy" signal whenever past price patterns seem to show an upward trend, and a "sell" signal in case of a downward trend.

Specifically, let j indicate a specific trading rule. Denote by P_j^n the return projection corresponding to the application of rule j, i.e. an agent i who bases her trading exclusively on rule j would be characterised by $R^n(\theta_i) = P_j^n$. Note that in practice, the majority of traders would not follow any one rule exclusively, but rather implement a "mix" of different rules according to the price patterns observed. We will model this type of behaviour below, after discussing some examples of specific trading rules most widely used in practice. For further reference regarding specific technical trading rules see the survey by Neely (1997).

3.3.2 Support and Resistance Levels

Support and Resistance (S&R) rules are among the most widely used technical trading rules. Resistance levels are thresholds set above the maximum of past exchange rates over a certain period, while the support levels are set below the minimum of past exchange rates. The S&R rule generates a buy signal when exchange rates rise beyond the resistance level, and a sell signal when rates fall below the support level. Formally,

$$\begin{aligned} &\text{if} \quad Y_{k-1}^n > \max_{i=2,\dots,l} Y_{k-i}^n + M && \text{then} \quad P_j^n = \gamma M > 0, \\ &\text{if} \quad Y_{k-1}^n < \min_{i=2,\dots,l} Y_{k-i}^n - M && \text{then} \quad P_j^n = -\gamma M < 0, \end{aligned}$$

for some positive constant γ.

3.3.3 Moving Average Rules

Besides support and resistence levels, the moving average (MA) rule is probably the most commonly used technical trading rule in foreign exchange markets. A trader following an MA rule tracks a short-term and a long-term moving average of past exchange rates. Whenever the short-term average moves above the long-term average, the rule generates a buy signal. Analogously, a sell signal is produced when the short-term average falls below the long-term average. Formally,

$$\begin{aligned} &\text{if} \quad \sum_{i=1}^{l} w_i Y_{k-i}^n > M && \text{then} \quad P_j^n = \gamma M > 0, \\ &\text{if} \quad \sum_{i=1}^{l} w_i Y_{k-i}^n < -M && \text{then} \quad P_j^n = -\gamma M < 0, \end{aligned}$$

for some positive constant γ. Here, $w_1 \geq \dots \geq w_l$ are weights with $w_1 > 0 > w_l$.

3.3.4 Other Rules

There are many other technical trading rules. Oscillator rules for example generate buy or sell signals whenever the *difference* between the short and the long moving average reverses its trend. Charting techniques on the other hand generate buy or sell signals when certain price patterns are detected. Examples include so-called "Head-and-Shoulders", "V-Base" or "Saucer" patterns. Instead of formalising all possible rules explicitly, we summarise their key features in the follwing

Assumption 3.3.2 *The P_j^n are uniformly bounded, and for fixed $Y_{k-l}^n, \ldots, Y_{k-2}^n$,*

$$\begin{array}{llll} \textit{as} & Y_{k-1}^n \longrightarrow +\infty, & P_j^n > 0 & \textit{eventually} \\ \textit{as} & Y_{k-1}^n \longrightarrow -\infty, & P_j^n < 0 & \textit{eventually} \end{array}$$

In other words, technical trading rules generate buy signals when current rates increase by a sufficient margin beyond their past maximum, and sell signals when current rates fall by a sufficient margin below their past minimum.

3.3.5 Fundamentalism

Most practitioners trust technical analysis only up to a certain point. When exchange rates move too far away from what is perceived as their "fundamental" level, many traders anticipate a trend reversal. To model this kind of behaviour, we introduce an additional trading rule $j = 0$ as

$$\begin{array}{llll} \text{if} & \min_{i=1,\ldots,l} Y_{k-i}^n > F + M & \text{then} & P_0^n = -\gamma M < 0, \\ \text{if} & \max_{i=1,\ldots,l} Y_{k-i}^n < F - M & \text{then} & P_0^n = \gamma M > 0, \end{array}$$

for some positive constant γ. In other words, when exchange rates move too far away from the "fundamental" level F, traders expect the trend to revert towards that level.

In practice, most traders will be neither pure "fundamentalists" nor pure technical traders. Therefore, we allow agents to "mix" different rules. Formally, we assume that agent i's total return projection is of the form

$$R^n(\theta_i) = \sum_j S_j^n(\theta_i) \cdot P_j^n.$$

For any given rule j, the indicator $S_j^n(\theta_i)$ determines whether to apply rule j or not, contingent on the observed pattern of past rates. Note that we allow $S_j^n(\theta_i)$ to depend on θ_i, i.e. the choice of trading rule may vary accross agents. Moreover, we do not restrict $S_j^n(\theta_i)$ to take values only in $\{0,1\}$, i.e. an individual agent may apply a weighted mix of different rules. For example, agents might follow pure technical rules as long as rates stay within reasonable distance of their fundamental level, while gradually shifting to fundamentalism as rates move further away.

3.4 Equilibrium

3.4.1 Aggregation

To facilitate our analysis, we now move to the limiting case which is characterised by an infinity of agents, i.e. we consider the aggregate demand of technical traders as their number, m, tends to infinity. Assuming sufficient regularity, the Law of Large Numbers implies

Theorem 3.4.1 *For any given $\mathfrak{y} \in \mathbb{R}^l$,*

$$\sum_{i=1}^{m} \beta^m(\theta_i) R^n(\mathfrak{y};\theta_i) \longrightarrow \overline{R}^n(\mathfrak{y}) := \int_\Theta \alpha(\theta) R^n(\mathfrak{y};\theta)\,\mu(d\theta) \qquad \text{as } m \longrightarrow \infty.$$

Although an individual agent's return projection $R^n(\theta_i)$ may well be a discontinuous function, we assume for what follows that in the aggregate, $\overline{R}^n : \Theta \times \mathbb{R}^l \longrightarrow \mathbb{R}$ is sufficiently smooth. For what follows, we will work with the limit economy characterised by $\overline{R}^n$. This should be seen as an approximation of the situation where there is a comparatively large number of heterogeneous technical traders active in the market.

3.4.2 Discrete-Time Equilibrium

In the limit economy, we can then write the equilibrium equation for the (logarithmic) exchange rate process as

$$\Delta Y_k^n = \overline{R}(\mathfrak{Y}_k^n) \cdot \Delta t^n + \Delta Z_k^n. \tag{3.4}$$

In order to derive qualitative results regarding the dynamic behaviour of the price process, we will focus our attention on the $\mathbb{R}^l$-valued process $(\mathfrak{Y}_k^n)_{k=0,1,\ldots}$ of *price histories*, rather than on the price process itself. It is obvious that, since $\overline{R}^n(\mathfrak{Y}_k^n)$ depends on the price history, the exchange rate process $(Y_k^n)_{k=0,1,\ldots}$ itself cannot be Markovian. However, it can be shown that the process $(\mathfrak{Y}_k^n)_{k=0,1,\ldots}$ of price histories is a Markov process, see Scheutzow (1984). We will now define the concepts of recurrence and transience, which we will use to characterise the dynamic behaviour of the equilibrium exchange rate process.

Definition 3.4.2 *The process* $(\mathfrak{Y}_k^n)_{k=0,1,\ldots}$ *is called*

(i) recurrent *if for all bounded Borel sets* $A \subseteq \mathbb{R}^l$ *with non-zero Lebesgue-measure, there exists a sequence* $k_i \longrightarrow \infty$ *such that*

$$\mathfrak{Y}_{k_i}^n \in A \qquad \textit{for all } i$$

almost surely; and

(ii) transient *if for all bounded Borel sets* $A \subseteq \mathbb{R}^l$ *with non-zero Lebesgue-measure,*

$$\mathfrak{Y}_k^n \notin A \qquad \textit{for sufficiently large } k$$

almost surely.

In other words, the process is recurrent if it visits any set $A \subseteq \mathbb{R}^l$ of price histories infinitely often. In particular, a recurrent process cannot converge as $k \longrightarrow \infty$. A transient process on the other hand eventually leaves any bounded set $A \subseteq \mathbb{R}^l$, i.e. transient processes must diverge as $k \longrightarrow \infty$.

Theorem 3.4.3 *Let* $(Y_k^n)_{k=0,1,\ldots}$ *be a solution to the Stochastic Delay Equation (3.4). If the function* $\overline{R} : \mathbb{R}^l \longrightarrow \mathbb{R}$ *is locally bounded and measurable, then* $(\mathfrak{Y}_k^n)_{k=0,1,\ldots}$ *is* either *recurrent* or *transient. Morover, in the recurrent case there exists a unique stationary probability distribution* π^* *on* $\mathbb{R}^l$.

PROOF: Scheutzow (1984, Theorem 10). □

This remarkable theorem says that the exchange rate process $(Y_k^n)_{k=0,1,\ldots}$ can only display *either* of two radically different types of asymptotic behaviour. Exchange rates either fluctuate randomly, with bubbles emerging and bursting, or they drift off to infinity. The theorem also implies that, however strong the fundamentalism is in the market, it can never keep exchange rates within a certain region, unless traders have unlimited funds. It is intuitively clear that whether the exchange rate process is recurrent or transient will depend on the relative weight of technical trading versus fundamentalism. In order to analyse the dynamic behaviour of the equilibrium exchange rate process, in particular to derive sufficient conditions for recurrence, we will now pass to the continuous-time limit as the trading intervals, Δt^n, tend to zero.

3.4.3 Continuous-Time Equilibrium

In order to facilitate the passage to the limit, we have to choose a common base space on which the distributions of the equilibrium exchange rate processes can be studied for every n. Denote by $\mathcal{D}^d[0,T]$ and $\mathcal{D}^d[0,\infty)$ the d-dimensional *Skorohod Space* of $\mathbb{R}^d$-valued RCLL functions on $[0,T]$ respectively $[0,\infty)$; cf. Jacod and Shiryaev (1987). To simplify notation, we write simply $\mathcal{D}[0,T]$ and $\mathcal{D}[0,\infty)$ in the case where $d=1$. With any given sequence $\xi_0^n, \xi_1^n, \ldots$ of numbers we associate the RCLL function

$$\xi_t^n := \xi_k^n \quad \text{for } t_k^n \leq t < t_{k+1}^n.$$

In this fashion, we identify the discrete-time processes t_k^n, Y_k^n and Z_k^n with their respective continuous-time counterparts τ_t^n, Y_t^n, and Z_t^n, all having paths in $\mathcal{D}[0,\infty)$. In the same way, we identify the price history $\mathfrak{Y}_k^n$ with an element $\mathfrak{Y}_t^n$ in $\mathcal{D}[0,T]$. Finally, we need to extend in the obvious way the aggregate return projection $\overline{R}^n$ to be defined on $\mathcal{D}[0,T]$. To ensure the convergence of the equilibrium exchange rate processes as $n \longrightarrow \infty$, we have to make the following

Assumption 3.4.4

(i) There exists a continuous, uniformly bounded function $\overline{R} : \mathcal{C}[0,T] \longrightarrow \mathbb{R}$ *with the following property: Whenever* $\mathfrak{y}^n \in \mathcal{D}[0,T]$ *and* $\mathfrak{y} \in \mathcal{C}[0,T]$ *are such that* $\mathfrak{y}^n \longrightarrow \mathfrak{y}$ *in* $\mathcal{D}[0,T]$*, then*

$$\overline{R}^n(\mathfrak{y}^n) \longrightarrow \overline{R}(\mathfrak{y}) \qquad \text{as } n \longrightarrow \infty.$$

(ii) The sequence of processes $\{Z^n\}_{n=1,2,\ldots}$ *converges in distribution on* $\mathcal{D}[0,\infty)$ *to a continuous semi-martingale* Z.

Part (i) of this assumption merely states that technical traders react in a similar fashion to a given exchange rate pattern in continuous time as they would to a discrete-time approximation of the same pattern. We are now ready to state the desired convergence result.

Theorem 3.4.5 *Suppose that Assumption 3.4.4 holds. Then the triplet* (τ_t^n, Z_t^n, Y_t^n) *of processes is tight in* $\mathcal{D}^3[0,\infty)$*, and any limit* (t, Z_t, Y_t) *is a weak solution to the Stochastic Delay Equation*

$$dY_t = \overline{R}(\mathfrak{Y}_t)\, dt + dZ_t, \tag{3.5}$$

where $\mathfrak{Y}_t \in \mathcal{C}[0,T]$ *is defined in the obvious way. Moreover, if (3.5) admits a unique weak solution, then* (τ_t^n, Z_t^n, Y_t^n) *converges in distribution to* (t, Z_t, Y_t) *as* $n \longrightarrow \infty$.

PROOF: Mémin and Slominski (1991, Theorem 2.10). □

In order to study the qualitative behaviour of the equilibrium exchange rate process, we will now define the concepts of recurrence and transience in continuous time.

Definition 3.4.6 *The process $(\mathfrak{Y}_t)_{t\geq 0}$ is called*

(i) recurrent *if for all bounded Borel sets $A \subseteq \mathcal{C}[0,T]$ with non-zero Lebesgue-measure, there exists a sequence $t_i \longrightarrow \infty$ such that*

$$\mathfrak{Y}_{t_i} \in A \qquad \text{for all } i$$

almost surely; and

(ii) transient *if for all bounded Borel sets $A \subseteq \mathcal{C}[0,T]$ with non-zero Lebesgue-measure,*

$$\mathfrak{Y}_t \notin A \qquad \text{for sufficiently large } t$$

almost surely.

The interpretation of recurrence and transience is essentially the same as in discrete time. A recurrent process will always return to any given set $A \subseteq \mathcal{C}[0,T]$ of price histories, while a transient process eventually leaves every bounded set.

In order to be able to explicitly study the dynamic behaviour of the equilibrium exchange rate process, we will have to assume a more concrete specification of the cumulative demand of noise traders. As is common in the literature, we will assume that noise traders' demand is actually driven by white noise, so that the cumulative demand process is a generalised Brownian motion. More specifically, we assume that the process Z_t satisfies a Stochastic Differential Equation of the form

$$dZ_t = \lambda\, dt + \sigma\, dW_t,$$

where W_t is a standard Brownian motion, and λ and σ are *constants*. The equation describing the dynamics of logarithmic exchange rates, Y_t, then becomes:

$$dY_t = \left(\lambda + \overline{R}\,(\mathfrak{Y}_t)\right) dt + \sigma\, dW_t. \tag{3.6}$$

Itô's Lemma implies that the *actual* exchange rate process itself, $X_t := \exp(Y_t)$, satisfies

$$dX_t = \left(m + \overline{R}(\mathfrak{Y}_t)\right) X_t\, dt + \sigma X_t\, dW_t, \qquad \text{with} \quad m := \lambda + \frac{1}{2}\sigma^2. \tag{3.7}$$

Note in particular that the exchange rate model we have thus obtained is in fact dynamically *complete* in the sense of Harrison and Pliska (1981). In particular, although the exchange rate in our model may display cyclical fluctuations and bubbles, from the point of view of derivatives pricing, it "looks" just like the classic Black-Scholes model.

Theorem 3.4.7 *Let $(Y_t)_{t\geq 0}$ be a solution to the Stochastic Delay Equation (3.6). If the function $\overline{R} : \mathcal{C}[0,T] \longrightarrow \mathbb{R}$ is locally bounded and measurable, then $(\mathfrak{Y}_t)_{t\geq 0}$ is* either *recurrent* or *transient. Moreover, in the recurrent case there exists a unique stationary probability distribution π^* on $\mathcal{C}[0,T]$, and for any initial distribution π_0 of $\mathfrak{Y}_0$ the distribution π_t of $\mathfrak{Y}_t$ converges in total variation to π^*.*

PROOF: Scheutzow (1984, Theorem 3). □

We are now in the position to derive sufficient conditions for the recurrence of the equilibrium exchange rate process. For any element $\mathfrak{y} \in \mathcal{C}[0,T]$, define

$$\underline{\mathfrak{y}} := \min_{0\leq t\leq T} \mathfrak{y}(t); \qquad \overline{\mathfrak{y}} := \max_{0\leq t\leq T} \mathfrak{y}(t);$$

Theorem 3.4.8 *Suppose the function $\overline{R} : \mathcal{C}[0,T] \longrightarrow \mathbb{R}$ has the following properties:*

$$\left(\lambda + \overline{R}(\mathfrak{y})\right) \cdot \overline{\mathfrak{y}} \longrightarrow -\infty \qquad \text{as } \underline{\mathfrak{y}} \longrightarrow +\infty,$$

$$\left(\lambda + \overline{R}(\mathfrak{y})\right) \cdot \underline{\mathfrak{y}} \longrightarrow -\infty \qquad \text{as } \overline{\mathfrak{y}} \longrightarrow -\infty.$$

Then the process $(\mathfrak{Y}_t)_{t\geq 0}$ is recurrent.

PROOF: Scheutzow (1984, Theorem 5). □

Note that the recurrence on $\mathbb{R}$ of the logarithmic exchange rate process Y_t obviously implies recurrence on $\mathbb{R}_+$ of the exchange rate process X_t itself. In light of Sections 3.3.4 and 3.3.5, it is obvious that the condition for recurrence translates into a requirement regarding the relative balance of fundamentalism versus pure technical trading. More precisely, if on average fundamentalism becomes dominant eventually when rates become arbitrarily high or low, then the resulting equilibrium exchange rate process will be recurrent. For what follows, we will assume that the hypotheses of Theorem 3.4.8 are satisfied.

3.4.4 Profitability of Trading Strategies

We will now study the profitability of different trading strategies in an environment in which technical traders are active. For the purpose of doing so, we consider a *small* investor, i.e. one whose trading does not affect equilibrium prices. Let $\phi(t)$ be a trading strategy, i.e. at time t, the amount of foreign currency held is given by $\phi(t)$. We assume that $\phi(t)$ is adapted and predictable with respect to the filtration generated by the exchange rate process X_t. Note that predictability merely reflects the intuition that the decision on the amount of currency held over any given period must be based on the information available at the *beginning* of that period. However, beyond the intuition, predictability is also an important technical requirement in order for the gains from trade to be well-defined; see Harrison and Pliska (1981) for a detailed elaboration on the issue of admissibility of trading strategies. For any given trading strategy $\phi(t)$, we can now define the *gains from trade* as

$$G_t(\phi) := \int_0^t \phi(s)\, dX_s.$$

Note that, by running a suitable strategy in the domestic money market, we can always make the trading strategy *self-financing*. Hence, since we have normalised the price of the domestic bond to one, the process $G_t(\phi)$ indeed describes the gain made from following the trading strategy $\phi(t)$.

BUY-AND-HOLD:

Let use first consider a simple "buy-and-hold" strategy, i.e. $\phi \equiv 1$. Recall that from Theorem 3.4.8 we know that the exchange rate process X_t is *ergodic*. The Ergodic Theorem then implies that the expected gain over any finite period is the same as the long-term average gain along any given trajectory. More precisely,

$$E\,[\,G_t(\phi)\,] = \lim_{n\to\infty} \frac{1}{n}\sum_{k=0}^{n-1} G_t(\phi)\circ\vartheta_{kt}, \qquad \text{alsmost surely.}$$

Here, ϑ_{kt} is the *shift operator*, i.e. $X_t \circ \vartheta_{kt} = X_{(k+1)t}$. We are now in the position to study the profitability of the buy-and-hold strategy:

Proposition 3.4.9 *Over any finite period t, the expected gain from the buy-and-hold strategy is zero,*

$$E\,[\,G_t(\phi)\,] = 0.$$

PROOF: For the buy-and-hold strategy, the gains from trade over any finite period t are simply $G_t(\phi) = X_t - X_0$. Hence we obtain

$$\begin{aligned} E\,[\,G_t(\phi)\,] &= \lim_{n\to\infty} \frac{1}{n}\sum_{k=0}^{n-1}\left(X_{(k+1)t} - X_{kt}\right) \\ &= \lim_{n\to\infty} \frac{1}{n}\left(X_{nt} - X_0\right) = \lim_{n\to\infty} \frac{1}{n} X_{nt}. \end{aligned}$$

Hence, the limit of X_{nt}/n as $n \longrightarrow \infty$ must exist. But since X_t is recurrent, this limit can only be zero, since otherwise X_t would have to diverge to infinity which would contradict recurrence. $\square$

The interpretation of this result is interesting. Even though we allow the noise demand to impose a non-zero trend on the exchange rate process, and even though one would expect that such a trend will be amplified by technical trading, the expected return on investment in the foreign currency is nonetheless zero.

TECHNICAL TRADING:

In this section, we will investigate the profitability of technical trading rules themselves. More specifically, suppose an investor simply "mimics" the trading behaviour of the "average trader",

$$\phi(t) = \left(m + \overline{R}(\mathfrak{Y}_t)\right).$$

Proposition 3.4.10 *The expected gains from the average market strategy are positive,*

$$E\left[\, G_t(\phi) \,\right] > 0.$$

PROOF: Using the definition of the gains from trade process together with the equilibrium equation (3.7) for the exchange rate process, we obtain

$$G_t(\phi) = \int_0^t \sigma\phi(s)\, dW_s + \int_0^t \left(m + \overline{R}(\mathfrak{Y}_s)\right) \phi(s)\, ds.$$

Upon taking expectations, the integral with respect to the Wiener process vanishes, and we obtain

$$E\left[\, G_t(\phi) \,\right] = \int_0^t E\left[\, \left(m + \overline{R}(\mathfrak{Y}_s)\right)^2 \,\right] ds > 0$$

□

3.5 Conclusions

We have studied the dynamic properties of equilibrium exchange rates when market participants follow technical trading rules. By its very nature, technical trading leads to feedback effects similar to those inherent in the implementation of dynamic hedging strategies. As a consequence, technical trading causes exchange rate fluctuations of a magnitude far beyond the level justified by changes in the economic fundamentals. In other words, technical trading leads to the emergence of irrational price bubbles. However, if there is a sufficient level of fundamentalism in the market, it can be shown that bubbles will always burst in finite time.

Moreover, our results demonstrate that, while being ex-ante irrational, their very existence can make technical trading rules ex-post profitable. In other words, technical trading can be seen as a kind of "self-fulfilling prophecy". These results are in line with the empirical evidence.

Chapter 4

Optimal Trading for a Large Trader

4.1 Introduction

The purpose of the present chapter is twofold. First, it provides a unified generic framework in which the optimal trading pattern of a large trader who might possess private information can be studied. The framework is general enough to encompass a wide variety of standard models, most notably classic micro structure models such as Kyle (1985), as well as general equilibrium models. Second, it analyses the manner in which the optimal trading pattern changes if the large trader holds an option, written on the traded asset.

The equilibrium price setting mechanism is specified in reduced-form, encompassing as special cases price-setting by a competitive market maker as in Kyle (1985), Back (1992) or Chapter 5, as well as Walrasian equilibrium models like the ones considered in Chapters 2 and 6. The results of this chapter will be used in the subsequent chapters to study the interplay between options markets and the market for the underlying asset in the presence of asymmetric information, and to analyse the implications of market power for the use of options as exchange rate policy instruments.

The remainder of this chapter is organised as follows. In Section 4.2, we introduce the model and the mathematical framework for our analysis. The next section gives a detailed account of the large trader's optimisation problem. The main result of this section is the characterisation of the large trader's optimal trading strategy in the presence of an option, (Theorem 4.3.1). Section 4.4 analyses the value of the option to the informed trader and compares it to its arbitrage price.

4.2 The Model

We will outline the model here in its most general form. More detailed specifications and interpretations of its components will be given later, see Chapters 5 and 6. The structure of the model is closely related to the model developed in Back (1992) and Back and Pedersen (1996). Their model in turn is essentially a continuous-time extension of the classic Kyle (1985) model. While prices in these models are set by a competitive *market maker*, we consider in this chapter a reduced-form specification in which prices are given by a generic *reaction function*. Obviously, price setting by a market maker is included in this setup as a special case, but our formulation allows also for alternative price setting mechanisms such as for example Walrasian equilibrium.

Traded Assets:

Two assets are traded continuously over the period from time zero to some final date T. The first asset is a riskless bond or money market account. We use the bond as numéraire, normalising its price to one, thus making interest rates implicit in our model. The second asset is risky, typically a stock or a foreign currency. After time T, the risky asset yields a random pay-off V_T. While the distribution of V_T is common knowledge, its *actual* value is unknown ex ante.

Agents:

There is a single, risk-neutral, large trader active in the market. For the sake of generality, we specify the model in reduced form from the point of view of the large trader. At any time $t \in [0, T]$, we denote the total demand of the large trader by X_t. Trades at time t take place at price P_t. In the reduced form specification, we assume that prices are given by a *reaction function* in the form

$$P_t = H(t, X_t, Z_t).$$

Here, Z_t is the value of some exogenous stochastic factor process, to be made more specific later. In lack of anything better, we will refer to Z_t as the "noise" process. We will see in the applications how this type of reduced form pricing can be obtained as a result of a market equilibrium in which the large trader interacts with other economic agents and prices are determined by a market maker or a Walrasian auction mechanism, see Chapters 5 and 6, and Section 4.2.2. Note that the fact that the large trader's actions may affect prices is reflected by the fact that X_t appears as an argument of the pricing rule. More specifically, we will assume that

$$\frac{\partial}{\partial x} H(t, x, z) > 0; \qquad \frac{\partial}{\partial z} H(t, x, z) > 0. \tag{4.1}$$

In other words, we assume that the price of the risky asset is positively related to both the factor Z_t as well as the large trader's demand, X_t. In addition to observing the process Z_t, the large trader may also receive a *signal* S_t. The large trader's expectation at time t of the terminal value of the risky asset, based on the information available, is denoted by $F(t, S_t, Z_t)$. Note that the signal will be informative for the large trader only in some of the applications considered below. In these cases, we will refer to the large trader also as the *informed trader*. Based on the information available at time t, the large trader forms his or her demand X_t for the risky asset. For reasons of expositional simplicity, we restrict the large trader to using trading strategies $\theta(t)$ so that the demand X_t is given by $dX_t = \theta(t)\,dt$. Note in particular that this assumption forces the demand process X_t to be *absolutely continuous*. Although this assumption might seem unnecessarily restrictive, it can be shown that even if the large trader is allowed to choose from a more general set of strategies, the optimal strategy will turn out to be absolutely continuous; see Back (1992).

The main difference to the Back and Pedersen (1996) model is the following; In addition to trading in the market for the risky asset, the large trader may hold an option which, if exercised, pays off an amount $\varphi(P_T)$ at time T. Note that we assume the pay-off to depend on the terminal *price* of the asset, P_T, and not its true value. A simple example would be a cash-settled European call option with strike K, in which case $\varphi(P_T) = (P_T - K)$. In addition to the trading strategy $\theta(t)$, the large trader hence has to choose an exercise policy I for the option, which we assume to take on the value 1 if the option is exercised, and 0 if not. The large trader's objective is then to choose a trading strategy $\theta(t)$ and an exercise policy I for the option such as to maximise expected terminal wealth. We will formalise the large trader's problem formally in the next section. Note that we do not model the options market here, since this chapter focuses primarily on the large trader's optimal trading strategy.

The way in which the results developed here relate to equilibrium on the options market and what implications they have for the pricing of options is discussed in detail in Chapter 5.

We will now discuss briefly some examples of equilibrium models which would lead to the reduced form pricing rule given above.

4.2.1 Market Maker

In the models developed in Kyle (1985) and Back (1992), and further extended in Chapter 5, the large trader interacts with noise traders whose demand for the risky asset, Z_t, is exogenously given. A competitive market maker observes aggregate order flow, $Y_t = X_t + Z_t$, and sets prices according to a pricing rule of the form

$$P_t = H(t, Y_t).$$

It is obvious that from the large trader's point of view, this specification constitutes a special case of the generic framework introduced in the preceding section. In the next chapter, we analyse the interplay between equilibrium on the underlying market and the options market within this framework. We show that the existence of the option creates an incentive for the large trader to manipulate prices away from the expected true value. As a consequence, the only feasible equilibrium requires option prices to be alighned with the prices for the underlying asset in a way which gives rise to the famous "smile pattern" of implied volatility, see Chapter 5 for details.

4.2.2 Walrasian Equilibrium

As a second example, we consider the case in which the large trader interacts with another group of agents, and prices are set by a Walrasian auction mechanism. Without being specific about the nature of these other traders, we assume that their demand at time t for the risky asset, given a quoted price p, is given by the demand function

$$p \mapsto D(t, Z_t, p). \tag{4.2}$$

Here, the process Z_t could be interpreted either as the randomly fluctuating income of these traders, as a noise term driving liquidity demand, or as a process describing some fundamental information on which traders base their expectations about the terminal value of the asset. We will give more explicit examples of such demand functions $D(t, Z_t, p)$ below. At any time t the equilibrium price is the solution P_t of the *market clearing equation*

$$D(t, Z_t, P_t) + X_t \equiv 1. \tag{4.3}$$

Here, we have normalised the total supply of the risky asset to one. In other words, in this scenario the pricing function $H(t, x, z)$ is defined implicitly by the relation

$$D(t, x, H(t, x, z)) + x \equiv 1.$$

Differentiating this relation it is obvious that sufficient for Assumption (4.1) to hold is that the demand function $D(t, z, p)$ is differentiable with

$$\frac{\partial}{\partial z} D(t, z, p) > 0, \qquad \frac{\partial}{\partial p} D(t, z, p) < 0. \tag{4.4}$$

In this scenario, the value at time T of the risky asset to the large trader is given by the price at which the terminal position X_t can be "unloaded" in the market, i.e. $V_T = h(0, Z_T)$. Note that in this setting, the signalling process S_t does not play any role at all. In Chapter 6, I use the framework developed in this section to study the use of currency options as exchange rate policy instruments.

4.2.3 Mathematical Setup and Notation

For expositional simplicity, we will assume that both the signalling process, S_t, as well as the noise process, Z_t, are given by generalised Brownian motions. More specifically, we assume that S_t and Z_t are solutions to the stochastic differential equations

$$dS_t = \sigma_S(t)\, dW_t^S \qquad \text{and} \qquad dZ_t = \sigma_Z(t)\, dW_t^Z, \tag{4.5}$$

where W_t^S and W_t^Z are independent standard Brownian motions, and $\sigma_S(t)$ and $\sigma_Z(t)$ are deterministic functions. Note however that the results of this chapter can easily be extended to more general diffusion processes. A trading strategy is a process $\theta(t)$, adapted to the filtration generated by S_t and Z_t, such that the corresponding demand process

$$X_t^\theta := X_0^\theta + \int_0^t \theta(s)\, ds \tag{4.6}$$

is well-defined.

Let $P_t^{s,x,z}$ be a weak solution to (4.5) and (4.6) conditional on $S_t = s$, $Z_t = z$ and $X_t^\theta = x$, defined on some suitable measurable space $(\Omega, \mathcal{F})$. We will omit the superscripts s, x or z whenever there is no ambiguity, and we also write $P^{s,x,z}$ for $P_0^{s,x,z}$. Denote by $(\mathcal{F}_t)_{t\geq 0}$ the filtration on $(\Omega, \mathcal{F})$ generated by S_t and Z_t, augmented to satisfy the "usual conditions", see Karatzas and Shreve (1988, Section 1.2) for details. Finally, for notational convenience we introduce the following convention: whenever we are given a function which depends on the state variables *and* time t, we will use lowercase letters to denote the value of this function at time T. For example, we write $f(s,z) := F(T,s,z)$ and $h(x,z) := H(T,x,z)$.

4.3 The Large Trader's Problem

We develop a precise formulation of the large trader's maximisation problem in Sections 4.3.1, 4.3.2 and 4.3.3. The main result is stated at the end of Section 4.3.3 (Theorem 4.3.1). An intuitive derivation is provided in Sections 4.3.4 ... 4.3.7, while a rigorous proof is given in Appendix 4.A.

4.3.1 Wealth Dynamics

Suppose first that the large trader changes her position X_t in the risky asset only at discrete times $0 = t_0 < t_1 < \ldots < t_n = T$. Denote by X_k the number of shares held over the period from t_k to t_{k+1}, and by B_k the number of bonds. Trade at time t_k in the risky asset takes place at price P_k. The budget identity for the large trader is hence

$$B_{k-1} = B_k + P_k \cdot \Delta X_k,$$

where $\Delta X_k := X_k - X_{k-1}$. That is, the wealth from previous period's bond position, B_{k-1}, is used to finance the change in the position in the risky asset, $P_k \cdot \Delta X_k$, and any remaining funds are re-invested in bonds, B_k. We can rewrite this as

$$\Delta B_k = -P_k \cdot \Delta X_k = -P_{k-1} \cdot \Delta X_k - \Delta P_k \cdot \Delta X_k. \tag{4.7}$$

The continuous-time equivalent of this equation is

$$dB_t = -P_t^- \, dX_t - d[P, X]_t, \tag{4.8}$$

where P^- denotes the left-continuous version of P. Note that, under the working assumption of absolute continuity of the large trader's strategy, $dX_t = \theta(t)\, dt$, the covariation term vanishes and the budget dynamics become $dB_t = -P_t^- \theta(t)\, dt$.

4.3.2 Objective Function

The large trader's objective is to maximise expected terminal wealth. Ignoring the option for the moment, the large trader's terminal wealth is given by

$$B_T + V_T \cdot X_T = B_0 + V_T \cdot X_0 + \int_0^T \left(V_T - P_t^-\right) \theta(t)\, dt. \tag{4.9}$$

Since B_0 and X_0 are exogenously given and thus irrelevant for the optimisation problem, we will assume $B_0 = X_0 = 0$ and focus on the last term in the above expression. Given any strategy $\theta(t)$ and the associated demand process $dX_t^\theta = \theta(t)\, dt$, prices are set according to the rule $P_t = H(t, X_t^\theta, Z_t)$. Note here that since X_t^θ and Z_t are by definition continuous, so is the price process P_t. Recall also that given a final value S_T of the signalling process, the expected value of V_T is given by $f(S_T, Z_T)$. Using the law of iterated expectations, we can hence write the expected terminal wealth of the large trader for given starting values $S_0 = s$, $X_0 = x$, $Z_0 = z$, and given strategy $\theta(t)$ as

$$\begin{aligned} & E^{s,x,z}\left[\int_0^T (V_T - P_t)\, \theta(t)\, dt\right] \\ = \quad & E^{s,x,z}\left[\int_0^T \underbrace{\left(F(t, S_t, Z_t) - H(t, X_t^\theta, Z_t)\right)}_{=:\, \Pi(t, S_t, X_t^\theta, Z_t)} \theta(t)\, dt\right] \end{aligned}$$

Note that the function $\Pi(t, S_t, X_t^\theta, Z_t)$ has an economic interpretation: It describes the expected marginal benefit of holding an extra unit of the risky asset, $F(t, S_t, Z_t)$, minus the current purchasing price, $H(t, X_t^\theta, Z_t)$. Following our convention, we write

$$\pi(s, x, z) := \Pi(T, s, x, z) = f(s, z) - h(x, z).$$

We now incorporate the option into the large trader's objective function. Since we do not model the market for options, we treat the number of options held by the large trader as exogenous for the large trader's optimisation problem. Recall that the option, if exercised, pays off an amount $\varphi(P_T)$ at time T, and the price P_T at time T is set according to the rule $P_T = h(X_T^\theta, Z_T)$. Since from the large trader's point of view the pricing rule $h(x,z)$ is exogenously given, we shorten notation by setting $\phi(x,z) := \varphi(h(x,z))$. Hence, if the large trader holds ρ units of such an option, his expected terminal wealth becomes

$$E^{s,x,z}\left[\int_0^T \Pi(t, S_t, X_t^\theta, Z_t)\theta(t)\,dt + \rho I \phi(X_T^\theta, Z_T)\right]. \tag{4.10}$$

Here, I denotes the exercise policy, which we assume to be $= 1$ if the option is exercised, and $= 0$ if not. The exercise policy will be determined endogenously later, see Section 4.3.6.

4.3.3 Admissible Strategies

The choice variables for the large trader are the trading strategy $\theta(t)$ and the exercise policy I. Note that due to our assumption of absolute continuity, the demand process X_t is automatically *predictable*. This is in line with the intuition that the decision on what assets to hold over any period of time must be based upon the information available at the *beginning* of that period. It is also important for technical reasons, see for example Harrison and Pliska (1981) for a detailed elaboration on this issue. On the other hand, since the decision whether to exercise the option or not is made at time T, it can be contingent on all information available to the large trader at that time.

Formally, an admissible policy is a pair (θ, I), where I is an $\mathcal{F}_T$-measurable random variable taking values in $\{0,1\}$, and $\theta(t)$ is a process adapted to $(\mathcal{F}_t)_{t\geq 0}$ such that

$$V^{\theta,I}(t,s,x,z) := E_t^{s,x,z}\left[\int_t^T \Pi(\tau, S_\tau, X_\tau^\theta, Z_\tau)\theta(\tau)\,d\tau + \rho I \phi(X_T^\theta, Z_T)\right]$$

is well-defined for every $s, x, z \in \mathbb{R}$ and $t \in [0, T]$. We can interpret $V^{\theta,I}(t, s, x, z)$ as the *gain* that is made starting at $S_t = s$, $X_t = x$ and $Z_t = z$ at time t and following trading strategy θ and exercise policy I thereafter. The large trader's problem is thus to find an admissible policy (θ, I) which maximises $V^{\theta,I}(t, s, x, z)$. As usual in the theory of stochastic control, we define the *value function* of the problem as

$$V^*(t, s, x, z) := \sup_{(\theta,I)} V^{\theta,I}(t, s, x, z). \tag{4.11}$$

We are now ready to state the main result of this section.

Theorem 4.3.1 (Optimal Strategy) *There exists a non-degenerate solution to the large trader's problem if and only if the pricing rule satisfies the no-arbitrage condition* $H(t, x, z) = E_t^z[\ h(x, Z_T)\ ]$. *In this case, there exist functions* $X^*(t, s, z)$ *and* $I^*(s, z)$ *such that the optimum is attained by exercise policy* $I^* = I^*(S_T, Z_T)$, *and any trading strategy of the form*

$$\theta(t) := \alpha(t)\left(X^*(t, S_t, Z_t) - X_t^\theta\right), \tag{4.12}$$

where $\alpha(t)$ *is any deterministic function such that* $X_t^\theta \longrightarrow X^*(t, S_t, Z_t)$ *as* $t \longrightarrow T$.

The following Sections 4.3.4 ... 4.3.7 provide an intuitive derivation of this result, while a rigorous proof is given in Appendix 4.A. It makes use of standard results from the theory of controlled diffusion processes. For reference, see Fleming and Soner (1993) or Krylov (1980).

4.3.4 Bellman Equation

The key to finding the large trader's optimal policy is the Bellman equation, a partial differential equation which characterises the value function of the problem. We will give an intuitive derivation of this equation, and refer to the standard literature on stochastic control theory for the rigorous proofs. We will see that the boundary condition of the Bellman equation can be used to pinpoint the optimal policy.

To develop an intuitive motivation for the Bellman equation, consider first any arbitrary strategy $\theta(t)$. By optimality of V^* we see that

$$\begin{aligned} & V^*(t,s,x,z) \\ \geq \;& E_t^{s,x,z}\left[\int_t^{t+\delta} \Pi(\tau,S_\tau,X_\tau^\theta,Z_\tau)\theta(\tau)\,d\tau\right] + E_t^{s,x,z}\left[\,V^*(t+\delta,S_{t+\delta},X_{t+\delta}^\theta,Z_{t+\delta})\,\right]\end{aligned}$$

with equality at the optimum. If the value function is sufficiently smooth, we can divide both sides of this inequality by δ and pass to the limit as $\delta \longrightarrow 0$ to obtain

$$\left(\frac{\partial}{\partial t} + L_t^\theta\right) V^*(t,s,x,z) + \Pi(t,s,x,z)\theta \leq 0. \tag{4.13}$$

Here, L_t^θ is the infinitesimal generator (see for example Karatzas and Shreve (1988) for the precise definition) of the joint stochastic process (S_t, X_t^θ, Z_t), i.e.

$$L_t^\theta := \theta\frac{\partial}{\partial x} + \frac{1}{2}\left(\sigma_Z^2(t)\frac{\partial^2}{\partial z^2} + \sigma_S^2(t)\frac{\partial^2}{\partial s^2}\right). \tag{4.14}$$

Since θ was arbitrary, inequality (4.13) holds for all strategies. Intuitively, an optimal strategy should attain equality, so that we should expect the value function $V^*(t,s,x,z)$ to be a solution V of the following partial differential equation:

$$\sup_\theta\left\{\left(\frac{\partial}{\partial t} + L_t^\theta\right) V + \Pi\theta\right\} \equiv 0. \tag{4.15}$$

This equation is known as the "Bellman Equation" associated with the large trader's optimisation problem. Indeed, it can be shown that if the value function $V^*(t,s,x,z)$ is finite and sufficiently regular, then it must be a solution to the Bellman equation (4.15), at least in a generalised sense.[1] Conversely, if a function $V(t,s,x,z)$ is a solution to (4.15) under suitable boundary conditions, then it can be shown to coincide

[1]This result is known as "Bellman's Principle".

with the value function, i.e. $V(t,s,x,z) \equiv V^*(t,s,x,z)$.[2] See Fleming and Soner (1993) or Krylov (1980) for a rigorous mathematical derivation of these results. Let us now take a closer look at the Bellman equation: Substituting the definition of L_t^θ back into (4.15) we get

$$\sup_\theta \left\{ V_t + \frac{1}{2}\left(\sigma_Z^2 V_{zz} + \sigma_S^2 V_{ss}\right) + (V_x + \Pi)\,\theta \right\} \equiv 0.$$

We see that the above equation is *linear* in θ, i.e. in order for the supremum to be finite the co-efficient of θ must be zero. We can thus split up the Bellman Equation into the following two separate equations:

$$V_t + \frac{1}{2}\left(\sigma_Z^2 V_{zz} + \sigma_S^2 V_{ss}\right) \equiv 0, \tag{4.16}$$

$$V_x + \Pi \equiv 0. \tag{4.17}$$

Equation (4.16) is just the heat equation associated with two-dimensional Brownian motion. The Feynman-Kac formula (see e.g. Karatzas and Shreve (1988)) allows us to write its solution (if it exists) in the following form:

$$V(t,s,x,z) = E_t^{s,z}\left[\, V(T, S_T, x, Z_T)\,\right],$$

Note that the right-hand-side of this equation corresponds to the expected terminal value of V, given that the large trader's position x is kept constant from time t onwards. Equation (4.17) has a more economic interpretation. It states that the marginal gain $V_x(t,s,x,z)$ from holding an extra unit of the risky asset at time t, must equal the purchasing cost, $H(t,x,z)$, *minus* the expected marginal benefit, $F(t,s,z)$. To see the intuition behind this, recall that the function $V(t,s,x,z)$ is supposed to give the maximal gain achievable by starting at $S_t = s$, $X_t = x$, and

[2]Results of this kind are called "Verification Theorems".

$Z_t = z$, and trading optimally hence. Now assume that the equation was violated, say $V_x(t,s,x,z) + \Pi(t,s,x,z) < 0$. The large trader could then short δ units of the risky asset at time t, and trade optimally from the new starting point $X_t = x - \delta$, which would yield a gain of $V(t,s,x-\delta,z) \approx V(t,s,x,z) - \delta V_x(t,s,x,z)$. The approximate loss thus incurred, $\delta V_x(t,s,x,z)$, is by assumption smaller than the expected gain, $-\delta\Pi(t,s,x,z)$. We have thus constructed a strategy which yields a higher gain starting from $X_t = x$ than the original $V(t,s,x,z)$, and hence $V(t,s,x,z)$ cannot have been the maximal achievable gain.

4.3.5 Boundary Condition

We will now try to get a better intuition for the boundary condition for the Bellman equation and see how it can help us pinpoint the optimal strategy. Let $V(t,s,x,z)$ be any solution to (4.16) and (4.17). For a given admissible strategy θ consider the process $V(t,S_t,X_t^\theta,Z_t)$. Itô's Lemma gives:

$$\begin{aligned} dV &= \underbrace{\left\{ V_t + \frac{1}{2}\left(\sigma_Z^2 V_{zz} + \sigma_S^2 V_{ss}\right) \right\}}_{=0 \text{ by } (4.16)} dt + V_x\theta\, dt + \sigma_Z V_z\, dW_t^Z + \sigma_S V_s\, dW_t^S \\ &= -\Pi\theta\, dt + \sigma_Z V_z\, dW_t^Z + \sigma_S V_s\, dW_t^S, \end{aligned}$$

where we have omitted the arguments (t,S_t,X_t^θ,Z_t) of V and its derivatives. Upon taking expectations, the integrals with respect to the Wiener processes vanish,

$$\begin{aligned} &E^{s,x,z}\left[\, V(T,S_T,X_T^\theta,Z_T)\,\right] - V(0,s,x,z) \\ &= -E^{s,x,z}\left[\int_0^T \Pi(t,S,X_t^\theta,Z_t)\theta(t)\,dt \right] \\ &= -V^{\theta,I}(0,s,x,z) + E^{s,x,z}\left[\, \rho I \phi(X_T^\theta,Z_T)\,\right]. \end{aligned}$$

Now suppose the value function $V^*(t,s,x,z)$ exists and is sufficiently smooth. We then know that it solves the Bellman Equations (4.16) and (4.17), so that we can substitute $V^*(t,s,x,z)$ for $V(t,s,x,z)$ in the above equation to obtain

$$\begin{aligned} & V^*(0,s,x,z) \\ = \quad & V^{\theta,I}(0,s,x,z) + E^{s,x,z}\Big[\underbrace{V^*(T,S_T,X_T^\theta,Z_T) - \rho I \phi(X_T^\theta, Z_T)}_{=:\, q^*(S_T,X_T^\theta,Z_T;I)} \Big]. \end{aligned} \tag{4.18}$$

This equation has the following interpretation; the quantity $E^{s,x,z}[\ q^*(S_T,X_T^\theta,Z_T;I)\ ]$ measures the difference between the value $V^{\theta,I}(0,s,x,z)$ of any given policy pair (θ,I) and the maximal achievable value $V^*(0,s,x,z)$. The function $q^*(S_T,X_T^\theta,Z_T;I)$ can thus be identified as the *loss* incurred by the large trader by deviating from the optimal strategy. In particular we deduce that $E^{s,x,z}[\ q^*(S_T,X_T^\theta,Z_T;I)\ ]$ must be non-negative, and zero at the optimum. We will reverse this argument to construct the optimal policy and value function. More specifically, we will construct a function $q^*(s,x,z;I)$ and "optimality conditions" $x^*(s,z)$ and $I^*(s,z)$ in such a way that $q^*(s,x,z;I) \geq 0$ and $= 0$ if and only if $x = x^*(s,z)$ and $I = I^*(s,z)$, see Section 4.3.6 below. Based on this we use (4.16) and the Feynman-Kac representation to construct a candidate $V^*(t,s,x,z)$ for the value function which satisfies (4.18), see Section 4.3.7 below. From this and the above arguments we can conclude that an exercise policy $I^*(S_T,Z_T)$ and a trading strategy which ensures that $X_T^\theta = x^*(S_T,Z_T)$ will attain the optimum, and that $V^*(t,s,x,z)$ is indeed the value function of the large trader's problem. The formal proofs for all these statements can be found in Appendix 4.A. We interpret $x^*(s,z)$ as the optimal "target position" towards which the large trader should drive his position X_t^θ in the market as $t \longrightarrow T$, conditional on $S_T = s$ and $Z_T = z$. This is similar to the arguments presented in Back (1992) and Back and Pedersen (1996).

Let us begin our analysis by noting that we can deduce the shape of $q^*(s,x,z;I)$ from the Bellman equations: Evaluating (4.17) at $t = T$ implies:

$$\frac{\partial}{\partial x} q^*(s,x,z;I) = \underbrace{\frac{\partial}{\partial x} V^*(T,s,x,z)}_{= -\pi(s,x,z)} - \rho I \phi'(x,z). \tag{4.19}$$

Given any functions $x^*(s,z)$ and $I^*(s,z)$, it easy to check that

$$q^*(s,x,z;I) = \int_x^{x^*(s,z)} \pi(s,\xi,z)\, d\xi + \rho I^*(s,z)\phi(x^*(s,z),z) - \rho I \phi(x,z) \tag{4.20}$$

satisfies (4.19) with "boundary condition" $q^*(s, x^*(s,z), z; I^*(s,z)) = 0$. We can give an intuitive interpretation of (4.20). Recall that we think of $x^*(s,z)$ as an optimal "target position" towards which the large trader should drive X_t^θ as $t \longrightarrow T$. Recall also that $q^*(s,x,z;I)$ measures the "loss" incurred by the large trader by deviating from the optimal policy. Any marginal decrease in x away from $x^*(s,z)$ for example yields an expected marginal loss of $f(s,z)$ due to the reduced position in the risky asset, contrasted by a marginal gain of $h(x,z)$ due to lower cost of acquiring this position. The total *marginal* loss from deviation in x is hence $\pi(s,x,z)$. Iterating this argument motivates the integral term in (4.20). Similarly, implementing exercise rule I at x instead of $I^*(x,z)$ at $x^*(s,z)$ results in losing $\rho I^*(s,z)\phi(x^*(s,z),z)$ while gaining $\rho I \phi(x,z)$, which motivates the remaining terms in (4.20).

4.3.6 Optimal Policy

By construction, we have $q^*(s, x^*(s,z), z; I^*(s,z)) = 0$. Our remaining task is hence to find $x^*(s,z)$ and $I^*(s,z)$ in such a way that the mapping $(x, I) \mapsto q^*(s,x,z;I)$ has a global minimum at $x = x^*(s,z)$ and $I = I^*(s,z)$. We will proceed in two stages. First we keep the exercise decision fixed and derive optimal positions $x_0^*(s,z)$

and $x_1^*(s,z)$ conditional on no exercise respectively exercise. We then determine the optimal exercise policy $I^*(s,z)$ endogenously.

NEVER EXERCISE:

Let us consider first the case in which the large trader *never* exercises the option, i.e. $I^* \equiv 0$. We define the conditional loss function as

$$q_0(s,x,z;I,x') := \int_x^{x'} \pi(s,\xi,z)\,d\xi - \rho I \phi(x,z).$$

Note that we made the dependency on the target position x' explicit since this is the variable we need to determine. Intuitively, for any given x, the optimal x' should be the one for which the loss from deviating is maximal. Therefore we set

$$x_0^*(s,z) := \arg\max_{x'} q_0(s,x,z;I,x'). \tag{4.21}$$

The necessary first-order condition for this problem is

$$f(s,z) = h(x_0^*(s,z),z). \tag{4.22}$$

Note that, since we assume that $h(x,z)$ is monotonically increasing in x, the mapping $x' \mapsto q_0(s,x,z;I,x')$ will be concave, so that the first order condition (4.22) is also sufficient to uniquely characterise the maximum. Note also that $x_0^*(s,z)$ indeed does not depend on x or I, so that our notation is justified. We now substitute the maximum back into the conditional loss function and define

$$q_0^*(s,x,z;I) := q_0(s,x,z;I,x_0^*(s,z)).$$

Finally, since $h(x,z)$ is monotonically increasing in x, differentiating the first-order condition (4.22) shows that $x_0^*(s,z)$ is monotonically increasing in s.

ALWAYS EXERCISE:

Analogously, we now consider the case in which the large trader *always* exercises the option, i.e. $I^* \equiv 1$. We define the conditional loss function as

$$q_1(s,x,z;I,x') := \int_x^{x'} \pi(s,\xi,z)\,d\xi + \rho\phi(x',z) - \rho I\phi(x,z).$$

Again, x' is the variable we need to determine. As before, we set

$$x_1^*(s,z) := \arg\max_{x'} q_1(s,x,z;I,x'). \tag{4.23}$$

The necessary first-order condition for this problem is

$$f(s,z) = h(x_1^*(s,z),z) - \rho\phi_x(x_1^*(s,z),z). \tag{4.24}$$

If $\phi(x,z)$ is weakly concave in x, the mapping $x' \mapsto q_1(s,x,z;I,x')$ will be concave, so that the first order condition (4.24) is also sufficient to uniquely characterise the maximum. Note also that $x_1^*(s,z)$ indeed does not depend on x or I, so that our notation is justified. We now substitute the maximum back into the conditional loss function and define

$$q_1^*(s,x,z;I) := q_1(s,x,z;I,x_1^*(s,z)).$$

An argument similar to the one used in the preceding paragraph shows that monotonicity of $h(x,z)$ and weak concavity of $\phi(x,z)$ imply that $x_1^*(s,z)$ is also increasing in s. Moreover, a direct computation shows that $x_1^*(s,z) > x_0^*(s,z)$.

ENDOGENOUS EXERCISE:

Continuing our intuitive argument of maximising the loss at the optimum, we should seek to determine the exercise policy $I^*(s,z)$ in such a way that the option is exercised if and only if $q_1^*(s,x,z;I) \geq q_0^*(s,x,z;I)$, or, equivalently, if and only if

$$\int_{x_1^*(s,z)}^{x_0^*(s,z)} \pi(s,\xi,z)\,d\xi \leq \rho\phi(x_1^*(s,z),z). \tag{4.25}$$

Note that this condition shows that the optimal exercise policy depends only on s and z, and not on x or I. The economic interpretation of (4.25) is as follows. Recall that $x_0^*(s,z)$ is the optimal target position conditional on no exercise. Hence the left-hand side of (4.25) is the loss incurred by trading towards $x_1^*(s,z)$ instead of $x_0^*(s,z)$. The right-hand side however represents the gain made from the option's pay-off. Hence (4.25) states that the large trader should exercise the option if and only if the loss incurred by deviating from the optimal position is outweighed by the option's pay-off.

Using the first-order conditions for $x_0^*(s,z)$ and $x_1^*(s,z)$, we show in the appendix (Proposition 4.A.2) that for every z there exists a unique cut-off point $\hat{s}(z)$, defined by the equality

$$q_1^*(\hat{s},x,z;I) = q_0^*(\hat{s},x,z;I), \tag{4.26}$$

such that $q_1^*(s,x,z;I) \geq q_0^*(s,x,z;I)$ if and only if $s \geq \hat{s}(z)$. Note also that by (4.25) the cut-off point does not depend on x or I. We can hence formalise the optimal exercise policy in the following way;

$$I^*(s,z) = \begin{cases} 1 & \text{if } s \geq \hat{s}(z) \\ 0 & \text{otherwise} \end{cases}. \tag{4.27}$$

By construction, $x_0^*(s,z)$ and $x_1^*(s,z)$ are the optimal target positions conditional on no exercise respectively exercise. Intuitively, the unconditional target $x^*(s,z)$ should hence equal $x_1^*(s,z)$ if exercising the option is optimal, and $x_0^*(s,z)$ if not. Therefore we define

$$x^*(s,z) := I^*(s,z)x_1^*(s,z) + (1 - I^*(s,z))x_0^*(s,z).$$

A similar argument regarding the loss function motivates setting

$$q^*(s,x,z;I) := I^*(s,z)q_1^*(s,x,z;I) + (1 - I^*(s,z))q_0^*(s,x,z;I).$$

It is easy to check that $q^*(s,x,z;I)$ is indeed of the form (4.20). In the appendix (Proposition 4.A.3) we show that for any s and z, the mapping $(x,I) \mapsto q^*(s,x,z;I)$ indeed has a global minimum at $x = x^*(s,z)$ and $I = I^*(s,z)$, which qualifies $q^*(s,x,z;I)$ as our loss function.

We argued in Section 4.3.5 that any strategy θ which ensures that $X_t^\theta \longrightarrow x^*(S_T, Z_T)$ as $t \longrightarrow T$ is a candidate for the optimum. It is hence natural to construct θ in such a way that X^θ behaves like a Brownian Bridge with random target. More precisely, set $X^*(t,s,z) := E_t^{s,z}[\, x^*(S_T, Z_T)\,]$ and define

$$\theta(t) := \alpha(t) \cdot \left(X^*(t, S_t, Z_t) - X_t^\theta\right),$$

where $\alpha(t)$ is some adjustment speed factor with $\alpha(t) \longrightarrow \infty$ as $t \longrightarrow T$ fast enough to make X_t^θ converge to $x^*(S_T, Z_T)$. We show in Appendix 4.A that this strategy is indeed optimal.

REMARK: From the above we see that we can write the overall first-order condition for the large trader's terminal position in the following form:

$$h\left(x^*(S_T, Z_T), Z_T\right) = f(S_T, Z_T) + \rho I^*(S_T, Z_T)\phi_x\left(x^*(S_T, Z_T), Z_T\right). \tag{4.28}$$

Since for an optimal strategy, we have $X_T^\theta = x^*(S_T, Z_T)$, the above equation implies that the terminal price of the asset, $h(X_T^\theta, Z_T)$, equals its expected true value, $f(S_T, Z_T)$, *plus* a mark-up $\rho I^*(S_T, Z_T)\phi'(X_T^\theta, Z_T)$. Thus, unlike in the models of Kyle (1985) or Back and Pedersen (1996), the existence of the option implies that the large trader has an incentive to drive prices *away* from the true expected value of the asset.

4.3.7 Value Function

We will now construct a candidate for the value function $V^*(t,s,x,z)$ based on the analysis in the preceding sections. Guided by (4.18) we conjecture that the boundary condition for the value function be given by

$$v^*(s,x,z) := q^*(s,x,z;I) + \rho I\phi(x,z). \tag{4.29}$$

Obviously, $v^*(s,x,z)$ is smooth in x and z, and smooth in s for $s \in \mathbb{R} \setminus \{\hat{s}(z)\}$. Moreover, a straight-forward calculation, using the first-order conditions for $x_0^*(s,z)$ and $x_1^*(s,z)$, shows

$$\lim_{s\downarrow \hat{s}(z)} v^*(s,x,z) - \lim_{s\uparrow \hat{s}(z)} v^*(s,x,z) = q_1^*(\hat{s}(z),x,z;I) - q_0^*(\hat{s}(z),x,z;I) = 0,$$

so that $v^*(s,x,z)$ is continuous also at $s = \hat{s}(z)$. In view of (4.16), we define a candidate for the value function via the Feynman-Kac representation formula:

$$V^*(t,s,x,z) := E_t^{s,z}\,[\,v^*(S_T,x,Z_T)\,] \tag{4.30}$$

We will show in the appendix (Proposition 4.A.4) that under some additional assumptions the function thus defined indeed solves the Bellman Equation (4.15). From the analysis in the preceding sections it should then be clear that $V^*(t,s,x,z)$ *is* the value function, and that the strategy defined above *is* an optimal strategy. Although the formal proof is deferred to the appendix, one remark is due here: A direct computation shows that

$$\frac{\partial}{\partial y}V^*(t,s,x,z) = E_t^z\,[\,h(x,Z_T)\,] = -\underbrace{E_t^{s,z}\,[\,f(S_T,Z_T)\,]}_{=F(t,s,z)}.$$

Recall that in order to satisfy the Bellman Equation, $V^*(t,s,x,z)$ must in particular solve (4.17). From the above we see that this is possible only if $H(t,x,z) = E_t^z[\,h(x,Z_T)\,]$. We have thus found a condition that the pricing rule must satisfy in order for the large trader's problem to have a non-degenerate solution.

4.4 Option Valuation

In this section we will make use of the results of the previous sections to derive the value of the option to the large trader. It will turn out that in contrast to classic derivatives pricing theory, the unit price of the option depends on the quantity held, i.e. the price schedule is non-linear in quantity.

4.4.1 Pricing the Option

The value of the option to the large trader is obviously given by the expected gain that he or she derives from holding it. Denote by V_ρ^* the value function associated with the large trader's problem for a given option position ρ. Recall that from (4.18) we have

$$V_\rho^* = V^{\theta,I} + E\left[\, q_\rho^*(S_T, X_T^\theta, Z_T; I) \,\right],$$

where q_ρ^* is the "loss-function" for the given ρ, and (θ, I) is *any* policy pair. Taking in particular θ to be the optimal trading strategy in the absence of the option, and $I \equiv 0$, we find that the value to the large trader of ρ units of the option which pays off $\varphi(P_T)$ if exercised is given by

$$\begin{aligned}(V_\rho^* - V_0^*) &= E\Big[\, q_\rho^*(S_T, x_0^*(S_T, Z_T), Z_T, I) \,\Big] \\ &= E\Big[\, I_\rho^*(S_T, Z_T) \underbrace{\left\{ \int_{x_0^*(S_T,Z_T)}^{x_1^*(S_T,Z_T)} \pi(S_T, \xi, Z_T)\, d\xi + \rho\phi(x_1^*(S_T, Z_T), Z_T) \right\}}_{=:\, \psi_\rho\,(f(S_T, Z_T))} \,\Big]\end{aligned}$$

Note that by the large trader's first-order conditions, the term in the brackets indeed only depends on $f(S_T, Z_T)$, so that our notation is justified. By definition of the exercise policy, $I_\rho^*(S_T, Z_T) = 1$ if and only if $\psi_\rho(f(S_T, Z_T)) \geq 0$. Thus,

$$(V_\rho^* - V_0^*) = E\left[\, \psi_\rho^+\,(f(S_T, Z_T)) \,\right],$$

which in particular shows that the value of the option to the large trader is positive. The above expectation can be seen to be the Black-Scholes price of a derivative which pays off an amount $\psi_\rho(f(S_T, Z_T))$ if exercised. Indeed, consider the ficticious price process $\hat{P}_t := F(t, S_t, Z_t)$. Note that $\hat{P}_t$ represents the best estimate of the *true* value of the asset, given the large trader's information at time t. Using Feynman-Kac and Itô we find

$$d\hat{P}_t = \sigma_S(t)F_s(t, S_t, Z_t)\, dW_t^S + \sigma_Z(t)F_z(t, S_t, Z_t)\, dW_t^Z.$$

In other words, the process $\hat{P}_t$ is a martingale under the original measure. We summarise these results in the following

Proposition 4.4.1 *The value to the large trader of ρ units of the option is given by*

$$E\left[\, \psi_\rho^+ \left(f(S_T, Z_T)\right) \,\right],$$

which can be identified as the Black-Scholes price of an option with pay-off pattern $\psi_\rho(\hat{P}_T)$, written on the (expected) true value $\hat{P}_t = F(t, S_t, Z_t)$ of the underlying asset.

4.A Appendix: Complements to Section 4.3

Assumption 4.A.1 *Throughout this appendix, we make the following assumptions:*

1. *The function $h(x, z)$ is twice continuously differentiable in x, non-negative and strictly increasing in x, i.e. $h(x, z) \geq 0$ and $h_x(x, z) > 0$ for all x, z.*

2. *The function $\phi(x, z)$ is twice continuously differentiable in x, strictly increasing and weakly concave in x, i.e. $\phi_x(x, z) > 0$ and $\phi_{xx}(x, z) \leq 0$ for all x, z.*

Note that (2) is in particular satisfied if in addition to (1), $h(x,z)$ is weakly concave and $\varphi(p)$ is strictly increasing and weakly concave in p. The latter is in particular true for the European call option pay-off $\varphi(p) = (p - K)$.

Proposition 4.A.2 (Cut-Off Point) *Let $q_0^*(s,x,z;I)$ and $q_1^*(s,x,z;I)$ be defined as in Section 4.3.6. Then there exists a unique cut-off point $\hat{s}(z)$, independent of x and I, such that*

$$q_1^*(s,x,z;I) \geq q_0^*(s,x,z;I) \qquad \text{if and only if} \qquad s \geq \hat{s}(z)$$

PROOF: From the definitions of $q_0^*(s,x,z;I)$ and $q_1^*(s,x,z;I)$ and the first-order conditions (4.22) and (4.24) for $x_0^*(s,z)$ and $x_1^*(s,z)$ we find

$$\frac{\partial}{\partial s}\left(q_1^*(s,x,z;I) - q_0^*(s,x,z;I)\right) = \left(x_1^*(s,z) - x_0^*(s,z)\right) f_s(s,z) > 0,$$

since $f'(s) > 0$ by assumption, and $x_1^*(s,z) > y_0^*(s,z)$ as was shown in Section 4.3.6.

□

Proposition 4.A.3 (Loss Function) *Let the function $q^*(s,x,z;I)$ be defined as in Section 4.3.6. Then, for any fixed s and z, the mapping $(x,I) \mapsto q^*(s,x,z;I)$ has a global minimum at $x = x^*(s,z)$ and $I = I^*(s,z)$.*

PROOF: We consider the two cases, $I^*(s,z) = 0$ and $I^*(s,z) = 1$, separately.

Case (1): $I^*(s,z) = 0$. Here, $x^*(s,z) = x_0^*(s,z)$ and $q^*(s,x,z;I) = q_0^*(s,x,z;I)$. We have

$$\frac{\partial}{\partial x} q_0^*(s,x,z;I) = -\pi(s,x,z) - \rho I \phi_x(x,z) \tag{4.31}$$

$$\frac{\partial^2}{\partial x^2} q_0^*(s,x,z;I) = h_x(x,z) - \rho I \phi_{xx}(x,z) \tag{4.32}$$

Hence, since $\phi_{xx} \leq 0$ by assumption, $q_0^*(s,x,z;I)$ is strictly convex in x by (4.32), so that for each I there exists a unique global minimum. From (4.31) we see that the minimum is attained at $x = x_0^*(s,z)$ if $I = 0$, and at $x = x_1^*(s,z)$ if $I = 1$. We need to show that the overall minimum is at $I = 0$, i.e. that

$$q_0^*(s, x_1^*(s,z), z; 1) \geq q_0^*(s, x_0^*(s,z), z; 0) = 0.$$

Since $I^*(s,z) = 0$, we know already by definition that $q_0^*(s,x,z;I) \geq q_1^*(s,x,z;I)$ for all x and I. Choosing $x = x_1^*(s,z)$ and $I = 1$, this implies

$$q_0^*(s, x_1^*(s,z), z; 1) \geq q_1^*(s, x_1^*(s,z), z; 1) = 0 = q_0^*(s, x_0^*(s,z), z; 0),$$

which is what we had to show.

Case (2): $I^*(s,z) = 1$. Here, $x^*(s,z) = x_1^*(s,z)$ and $q^*(s,x,z;I) = q_1^*(s,x,z;I)$. We have

$$\frac{\partial}{\partial x} q_1^*(s,x,z;I) = -\pi(s,x,z) - \rho I \phi_x(x,z) \tag{4.33}$$

$$\frac{\partial^2}{\partial x^2} q_1^*(s,x,z;I) = h_x(x,z) - \rho I \phi_{xx}(x,z) \tag{4.34}$$

Hence, since $\phi_{xx} \leq 0$ by assumption, $q_1^*(s,x,z;I)$ is strictly convex in x by (4.34), so that for each I there exists a unique global minimum. From (4.33) we see that the minimum is attained at $x = x_0^*(s,z)$ if $I = 0$, and at $x = x_1^*(s,z)$ if $I = 1$. We need to show that the overall minimum is at $I = 1$, i.e. that

$$0 = q_1^*(s, x_1^*(s,z), z; 1) \leq q_1^*(s, x_0^*(s,z), z; 0).$$

Since $I^*(s,z) = 1$, we know already by definition that $q_0^*(s,x,z;I) \leq q_1^*(s,x,z;I)$ for all x and I. Choosing $x = x_0^*(s,z)$ and $I = 0$, this implies

$$q_1^*(s, x_1^*(s,z), z; 1) = 0 = q_0^*(s, x_0^*(s,z), z; 0) \leq q_1^*(s, x_0^*(s,z), z; 0),$$

which is what we had to show.

□

Proposition 4.A.4 (Value Function) *Let $V^*(t,s,x,z)$ be defined as in Section 4.3.7. If the market maker's pricing rule satisfies $H(t,x,z) = E_t^z[\ h(x,Z_T)\ ]$, then $V^*(t,s,x,z)$ satisfies the Bellman Equation (4.15) in the classical sense, with boundary condition*

$$V^*(T,s,x,z) = v^*(s,x,z).$$

Proof: We have seen in Section 4.3.7 that the function $v^*(s,x,z)$ is continuous. Results from the theory of linear parabolic differential equations ensure that there exists a classical solution to (4.16) with boundary condition $V^*(T,s,x,z) = v^*(s,x,z)$. From the Feynman-Kac representation theorem (see for example Karatzas and Shreve (1988, Thm 5.7.6, page 366)) we know that the solution can be represented as

$$V^*(t,s,x,z) := E_t^{s,z}\,[\ v^*(S_T,x,Z_T)\ ]\,.$$

A direct computation then shows that

$$\frac{\partial}{\partial x}V^*(t,s,x,z) = E_t^z\,[\ h(x,Z_T)\ ]\,. = -\underbrace{E_t^s\,[\ f(S_T)\ ]}_{=\,F(t,s)}$$

Hence, if $H(t,x,z) = E_t^z[\ h(x,Z_T^0)\ ]$, the right-hand side of the above equation equals $\Pi(t,s,x,z)$, so that $V^*(t,s,x,z)$ also satisfies (4.17). Following the arguments made in section 4.3.4, this implies that $V^*(t,s,x,z)$ is in fact a solution of the Bellman Equation (4.15), which was to be shown.

□

Proof of Theorem 4.3.1: Define the function $V^*(t,s,x,z)$ as in Proposition 4.A.4. Under our assumption regarding the market maker's pricing rule, we have seen in

Proposition 4.A.4 that $V^*(t,s,x,z)$ solves (4.16) and (4.17). Let (θ, I) be an arbitrary admissible policy. Repeating the arguments made in Section 4.3.5, we apply Itô's Lemma to the process $V^*(t, S_t, X_t^\theta, Z_t)$ and take expectations to obtain

$$V^*(0,s,x,z) = V^{\theta,I}(0,s,x,z) + E^{s,x,z}\left[\, q^*(S_T, X_T^\theta, Z_T; I)\,\right].$$

From Proposition 4.A.3 we know that $q^*(S_t, X_T^\theta, Z_T; I) \geq 0$ with equality if $X_T^\theta = x^*(S_T, Z_T)$ and $I = I^*(S_T, Z_T)$. But this is true by construction for the policy pair (θ^*, I^*). Hence we conclude that

$$V^{\theta,I}(t,s,x,z) \leq V^*(t,s,x,z),$$

with equality if $\theta = \theta^*$ and $I = I^*$. This shows that the policy pair (θ^*, I^*) indeed maximises $V^{\theta,I}(t,s,x,z)$, which is what we had to show. We also see that the function $V^*(t,s,x,z)$ constructed in Section 4.3.7 is indeed the value function associated with the informed trader's problem, as defined in (4.11).

□

Chapter 5

Asymmetric Information and the "Smile" Pattern

5.1 Introduction

In recent years, a fast growing literature has arisen which studies the effects of asymmetric information on asset prices. However, comparatively little has been said about the impact of asymmetric information on the pricing of derivative securities. The present paper investigates the manner in which an option, held by a large, informed trader, changes the nature of equilibrium prices for the underlying asset, and how this feeds back into option prices. As analytical framework, we choose an extension of the continuous-time Kyle (1985) model as laid out in Back (1992) and later extended by Back (1993) and Back and Pedersen (1996). Here, a large trader, who receives a private signal about the fundamental value of the traded asset, interacts with pure noise traders. Prices are set by a risk-neutral competitive market maker.

As in the asymmetric information model of Glosten and Milgrom (1985) the prices set by the market maker naturally satisfy a zero expected profit condition. In this type of model, the informed trader typically drives prices towards the expected value of the asset, thus successively revealing the private information to the market maker. We extend the Back (1992) model by introducing an over-the-counter (OTC) market on which derivative securities can be traded prior to trading in the underlying asset. There is no liquid market for options once trading in the underlying market has commenced. In fact we can show that allowing options to be traded while there is asymmetric information regarding the underlying asset would cause market breakdown. Hence, closing the options market as soon as the underlying market opens can be justified *endogenously* in our model. Unlike Back (1993), we do not introduce additional noise in the options market. The focus of our attention is the endogenous effect that options held by the informed trader have on the underlying market equilibrium and how this feeds back into option prices. These effects would be washed out by additional noise in the options market.

Unsurprisingly, it turns out that the presence of the option creates an incentive for the informed trader to manipulate prices of the underlying asset in order to increase the option's pay-off. Uninformed traders might hence face a price mark-up over the expected true value of the asset. In other words, the informed trader creates a price bubble at the expense of uninformed traders. The size of this bubble is determined by the trade-off faced by the informed trader: On one hand, the informed trader gains from upward price manipulation due to the increased option pay-off. On the other hand, the informed trader incurs a loss from manipulation because he has to buy the underlying asset at prices higher than its expected value. It is the potential existence of such bubbles which constitutes the main difference between our model and the model by Back (1993), in which this kind of manipulation is precluded by the existence of a liquid options market with additional noise.

We derive necessary and sufficient conditions for the existence of equilibria. It turns out that the only feasible equilibrium on both markets requires that the market maker does not believe the informed trader to have an incentive to manipulate prices. This belief can only be rational if equilibrium option prices are "synchronised" with prices on the underlying market. While this synchronicity has to be *assumed* in Jarrow (1994) and Frey (1996), it is an *endogenous consequence* in our model. This has two important implications.

First, it turns out that equilibrium option prices are non-linear in quantity. This is in line with the findings in Jarrow (1994) and Frey (1996), where option prices are determined by the cost of replication, and thus non-linearity arises as a consequence of finite elasticity of the underlying market. It is also consistent with empirical observations made in the OTC market for options, in particular for large quantities.

Second, for a special case, the equilibrium price process for the underlying asset can be shown to follow a generalised geometric Brownian Motion. This allows us to compare our findings with the classic Black-Scholes analysis. In particular, we are able to derive an explicit expression for implied volatility. We show that the implied volatility pattern generated by equilibrium prices in our model displays the famous "skew smile pattern", which has been observed in most financial markets ever since the October 1987 crash. There have been a number of other studies that explain the smile pattern of implied volatility. Most of these obtain the smile by exogenously altering the volatility structure of the underlying price process. Examples include Stochastic Volatility or ARCH models. The reasoning here is quite different. Here, no assumptions are needed regarding the underlying price volatility, the smile pattern rather occurs endogenously as a consequence of the market structure. To the best of our knowledge the only other paper in which the smile has been obtained endogenously is the paper by Platen and Schweizer (1998). Their analysis is based on a modified version of the feedback model developed in Chapter 2 and published previously in Frey and Stremme (1994) and Frey and Stremme (1997). However, in order to explain the smile pattern, Platen and Schweizer have to assume an upward-sloping demand curve for the underlying asset, which implies that the equilibrium they obtain is highly unstable. In other words, Platen and Schweizer have to *assume* an unstable equilibrium in order to explain the smile, while in our model the smile arises as a consequence of *preventing* instability.

The emergence of the skew smile pattern in our model is a result of the attempt to prevent market breakdowns. Empirically, the skewness is a feature which has only been observed *after* the October 1987 stock market crash. The explanation offered by our analysis could be that after the crash, market participants and regulators implemented measures aimed at preventing similar events in the future. This might have brought option prices more in line with the markets for the underlying asset. According to our findings, such an increased level of price synchronicity would indeed result in precisely the implied volatility pattern that has been observed since the crash.

The remainder of this chapter is organised as follows. In Section 5.2, we introduce the model and the mathematical framework for our analysis. The next section gives a brief account of the informed trader's optimisation problem, following the arguments developed in Chapter 4. Section 5.4 analyses the structure of equilibria on both markets, and the implications for the pricing of the option. The resulting price schedule for options and the implied volatility pattern are studied in detail in Section 5.5. Section 5.6 finally concludes.

5.2 The Model

The model is closely related to the one developed by Back (1992) and Back and Pedersen (1996). Their model in turn is an extended continuous-time version of the classic Kyle (1985) model. A single risky asset and a riskless bond are traded continuously over the period from time zero to some final date T. We take the riskless bond as numéraire, thus making interest rates implicit in our model. After time T, the risky asset is liquidated and yields a (random) pay-off V_T. While the distribution of V_T is common knowledge, the actual value is unknown ex ante. There are two agents in the market, a representative *noise trader* whose demand for the risky asset is exogenously given, and a single, large, *informed trader*, who receives a signal about the true value of the asset, V_T. At any time $t \in [0, T]$, we denote the noise trader's demand for the asset by Z_t, and that of the informed trader by X_t. Let $Y_t := X_t + Z_t$ denote the aggregate demand at time t. Trades at time t take place at price P_t, which is set by a competitive market maker. As in Back (1992) and Back and Pedersen (1996) we assume that the market maker only observes aggregate orders Y_t and not individual orders X_t or Z_t.

We extend on the model presented in Back (1992) and Back and Pedersen (1996) by introducing an over-the-counter options market. Prior to trading in the underlying asset, agents can purchase options which, if exercised, pay off an amount $\varphi(P_T)$ at time T. Note that we assume the option's pay-off to depend on the terminal *price* of the asset, P_T, and not its true value, V_T. A simple example would be a cash-settled European call option with strike K, in which case $\varphi(P_T) = (P_T - K)$. Options are issued by the market maker and traded over the counter; there is no liquid options market. Moreover, we assume that options cannot be traded once the market for the underlying asset has opened. We will see later that this restriction is necessary for the existence of equilibria: the closure of the options market is an endogenous consequence of the presence of asymmetric information, see Section 5.4 for details. Note here that we deliberately chose to focus on OTC derivatives, since the results we obtain are radically different from those that pertain in the presence of a liquid options market, as studied for example in Back (1993). Since the options are sold over the counter, the market maker can distinguish individual orders. We maintain however that the market maker does *not* know whether an individual buyer is informed or just a liquidity trader. The price at which the market maker is willing to sell a quantity ρ of the option is therefore independent of the buyer's identity. We denote this price by $\rho H_0(\rho)$, so that the price for *one unit* of the option contract is $H_0(\rho)$.

5.2.1 Informed Trader

As in Back and Pedersen (1996), the informed trader is assumed to receive a continuous signal S_t, successively revealing information about the terminal asset value, V_T. At any point in time t, the expected terminal value of the risky asset, conditional on the current signal S_t, is given by a function $F(t, S_t)$. Based upon the information available at time t, the informed trader forms his or her demand X_t for the risky asset. As in Back and Pedersen (1996) we restrict the informed trader to choosing trading strategies $\theta(t)$ such that the demand X_t is given by $dX_t = \theta(t)\,dt$. Note in particular that this assumption forces the demand process X_t to be *absolutely continuous*. Although this assumption may seem unnecessarily restrictive, it can be shown that even if the informed trader is allowed to choose from a more general set of strategies, the optimal trading strategy will turn out to be absolutely continuous; see Back (1992).

The main difference to the Back and Pedersen (1996) model is the following; In addition to trading in the risky asset, the informed trader may decide to hold a quantity ρ of the option contract. In addition to the trading strategy $\theta(t)$, the informed trader hence has to choose an exercise policy I for the option, which we assume to take on the value 1 if the option is exercised, and 0 if not. The informed trader's objective is thus to choose a quantity of options ρ, a trading strategy θ, and an exercise policy for the option I such as to maximise expected terminal wealth. We will solve the informed trader's optimisation problem in two stages via a backward induction argument: For the second stage (which we will address first), we take the quantity ρ of options held by the informed trader as given and solve for the optimal strategy θ and exercise policy I. This will enable us to quantify the *benefit* that the informed trader derives from holding a certain number of options, so that we can then determine the optimal option position to solve the first stage of the problem. For the second stage, we will make use of the results derived in Chapter 4, see Section 5.3 below. The first stage is dealt with in Section 5.4.2.

5.2.2 Market Maker

The price P_t at which the asset trades at time t is set by the market maker. As in Back (1992) and Back and Pedersen (1996) we assume that the market maker can only observe aggregate orders, Y_t, and not the individual orders X_t or Z_t, or the informed traders signal S_t. Therefore we restrict the market maker to setting prices according to a pricing rule of the form $P_t = H(t, Y_t)$. Also as in Back (1992) and Back and Pedersen (1996) we assume that the market maker acts perfectly competitive, i.e. that prices are set equal to the expected terminal value of the asset, conditional on the information available to the market maker. In addition to setting prices on the market for the risky asset, the market maker also maintains the over-the-counter market for option contracts. That is, the market maker has to choose a pricing schedule $H_0(\rho)$ for the option contract. We extend the requirement of competitive pricing to the options market, i.e. we assume that the price $H_0(\rho)$ reflects the expected value of ρ units of the option. However, it is worth stressing that this assumption can be justified endogenously: it can be shown to be *necessary* for the existence of equilibria, see Section 5.4.

5.2.3 Mathematical Setup and Notation

We will formulate the model within the framework developed in Chapter 4. Note first that the informed trader can invert the market maker's pricing rule $H(t, Y_t)$ to infer the aggregate order flow Y_t and thus the noise trader's demand Z_t from quoted prices P_t. From the informed trader's point of view, S_t and Y_t thus form a sufficient statistic for the current state of the economy. We can hence formulate the model entirely in terms of the two processes S_t and Y_t. For expositional simplicity, we assume that both the noise trader's demand process, Z_t, as well as the informed trader's signal, S_t, are generalised Brownian Motions. It is worth noting however that our results can easily be extended to more general diffusion processes. Formally, we assume that S_t solves the stochastic differential equation

$$dS_t = \sigma_S(t)\, dW_t^S, \tag{5.1}$$

where W_t^S is a standard Brownian motion and $\sigma_S(t)$ is a deterministic function. A trading strategy is an adapted process $\theta(t)$ such that there exists a unique solution to the equation

$$dY_t^\theta = \theta(t)\, dt + \sigma_Z(t)\, dW_t^Z, \tag{5.2}$$

where W_t^Z is a standard Brownian motion, independent of W_t^S, and $\sigma_Z(t)$ is a deterministic function. The solution Y_t^θ to (5.2) is the aggregate order flow that results if the informed trader follows trading strategy $\theta(t)$. The noise trader's demand is reflected in the diffusion term $\sigma_Z(t)dW_t^Z =: dZ_t$ in (5.2). As a special case, denote by Y_t^0 the aggregate demand process corresponding to the "do-nothing" strategy $\theta(t) \equiv 0$. Obviously, the dynamics of this process are simply $dY_t^0 = \sigma_Z(t)\, dW_t^Z$, i.e. identical to the dynamics of the noise trader's demand process, Z_t.

Let $P_t^{s,y}$ be a weak solution to (5.1) and (5.2) conditional on $S_t = s$ and $Y_t^\theta = y$, defined on some suitable measurable space $(\Omega, \mathcal{F})$. We will omit the superscripts s or y whenever there is no ambiguity, and we also write $P^{s,y}$ for $P_0^{s,y}$. Denote by $(\mathcal{F}_t)_{t\geq 0}$ the filtration on $(\Omega, \mathcal{F})$ generated by S_t and Y_t, augmented to satisfy the "usual conditions", see Karatzas and Shreve (1988, Section 1.2) for details. Finally, for notational convenience we introduce the following convention: whenever we are given a function which depends on the state variables *and* time t, we will use lowercase letters to denote the value of this function at time T. For example, we write $f(s) := F(T, s)$ and $h(y) := H(T, y)$.

5.2.4 Equilibrium

The choice variables for the informed trader in our model are, the quantity of option contracts held, ρ, the trading strategy in the underlying asset, $\theta(t)$, and the exercise policy for the option, I. We refer to the triple (ρ, θ, I) as the informed trader's *policy choice*. The market maker on the other hand sets pricing schedules $H_0(\rho)$ for the option contract and $H(t, Y_t)$ for the underlying asset. We are now ready to define the concept of equilibrium in our model:

Definition 5.2.1 (Equilibrium) *An equilibrium is a tupel* $(H_0, H; \rho, \theta, I)$ *such that*

(i) Given the market maker's pricing schedule (H_0, H)*, the informed trader's policy choice* (ρ, θ, I) *is* optimal, *i.e. maximises the informed trader's expected terminal wealth.*

(ii) Given the informed trader's policy choice (ρ, θ, I)*, the market maker's pricing schedule* (H_0, H) *is* rational, *i.e. reflects the assets' expected value, conditional on the market maker's information.*

5.3 The Informed Trader's Problem

In this section, we derive the solution to the second stage of the informed trader's problem, i.e. we take the number ρ of options held as given and solve for the optimal trading strategy $\theta(t)$ and exercise policy I. Note that from the informed trader's perspective, the market maker's pricing rule $H(t, Y_t)$ can be considered exogenously given, so that the informed trader's optimisation problem is a special case of the problem considered in Chapter 4. Therefore, we will give here only a brief outline of the main result, Theorem 5.3.1. A detailed exposition can be found in Chapter 4.

5.3.1 Objective Function

The informed trader's objective is to maximise expected terminal wealth. Following the intuition developed in Chapter 4, terminal wealth in the absence of the option is given by

$$B_T + V_T \cdot X_T = B_0 + V_T \cdot X_0 + \int_0^T \left(V_T - P_t^-\right) \theta(t)\, dt. \tag{5.3}$$

Since B_0 and X_0 are exogenously given and thus irrelevant for the optimisation problem, we will assume $B_0 = X_0 = 0$ and focus on the last term in the above expression. Given any strategy $\theta(t)$, the dynamics of the aggregate demand process, Y_t^θ, are given by (5.2). The market maker sets prices according to the rule $P_t = H(t, Y_t^\theta)$. Note here that since Y_t^θ is by definition continuous, so is the price process P_t. Recall also that given a final value S_T of the signalling process, the expected value of V_T is given by $f(S_T)$. Using the law of iterated expectations, we can hence write the expected terminal wealth of the informed trader for given starting values $S_0 = s$, $Y_0 = y$ and given strategy $\theta(t)$ as

$$E^{s,y}\left[\int_0^T (V_T - P_t)\,\theta(t)\,dt\right] = E^{s,y}\left[\int_0^T \underbrace{\left(F(t,S_t) - H(t,Y_t^\theta)\right)}_{=:\,\Pi(t,S_t,Y_t^\theta)}\theta(t)\,dt\right].$$

Note that the function $\Pi(t, S_t, Y_t^\theta)$ has an economic interpretation: It describes the expected marginal benefit of holding an extra unit of the risky asset, $F(t, S_t)$, minus the current purchasing price, $H(t, Y_t^\theta)$. Following our convention, we write

$$\pi(s,y) := \Pi(T,s,y) = f(s) - h(y).$$

We now incorporate the option into the informed trader's objective function. Since there is no market for options after time 0, we can treat the number of options held by the informed trader as exogenous for the informed trader's optimisation problem. We will analyse the options market in more detail in Section 5.5. Recall that the option, if exercised, pays off an amount $\varphi(P_T)$ at time T. The price P_T at time T is assumed to be set according to the rule $P_T = h(Y_T^\theta)$. Since from the informed trader's point of view the pricing rule $h(y)$ is exogenously given, we shorten notation by setting $\phi(y) := \varphi(h(y))$. Hence, if the informed trader holds ρ units of such an option, expected terminal wealth becomes

$$E^{s,y}\left[\int_0^T \Pi(t,S_t,Y_t^\theta)\theta(t)\,dt + \rho I\phi(Y_T^\theta)\right]. \tag{5.4}$$

Here, I denotes the exercise policy, which we assume to be $= 1$ if the option is exercised, and $= 0$ if not. The exercise policy will be determined endogenously later, see Section 5.3.3.

5.3.2 Admissible Strategies

The choice variables for the informed trader are the trading strategy $\theta(t)$ and the exercise policy I. Note that due to our assumption of absolute continuity, the demand process X_t is automatically *predictable*. This is in line with the intuition that the decision on what assets to hold over any period of time must be based upon the information available at the *beginning* of that period. It is also important for technical reasons, see for example Harrison and Pliska (1981) for a detailed elaboration on this issue. On the other hand, since the decision whether to exercise the option or not is made at time T, it can be contingent on all information available to the informed trader at that time.

Formally, an admissible policy is a pair (θ, I), where I is an $\mathcal{F}_T$-measurable random variable taking values in $\{0,1\}$, and $\theta(t)$ is a process adapted to $(\mathcal{F}_t)_{t\geq 0}$ such that

$$V^{\theta,I}(t,s,y) := E_t^{s,y}\left[\int_t^T \Pi(\tau, S_\tau, Y_\tau^\theta)\theta(\tau)\,d\tau + \rho I \phi(Y_T^\theta)\right]$$

is well-defined for every $s, y \in \mathbb{R}$ and $t \in [0,T]$. We can interpret $V^{\theta,I}(t,s,y)$ as the *gain* that is made starting at $S_t = s$ and $Y_t^\theta = y$ at time t and following trading strategy θ and exercise policy I thereafter. The large trader's problem is thus to find an admissible policy (θ, I) which maximises $V^{\theta,I}(t,s,y)$. As usual in the theory of stochastic control, we define the *value function* of the problem as

$$V^*(t,s,y) := \sup_{(\theta,I)} V^{\theta,I}(t,s,y). \tag{5.5}$$

We are now ready to state the main result of this section.

Theorem 5.3.1 (Optimal Strategy) *There exists a non-degenerate solution to the informed trader's problem if and only if the market maker's pricing rule satisfies $H(t,y) = E_t^y[\, h(Y_T^0)\,]$. In this case, there exist functions $Y^*(t,s)$ and $I^*(s)$ such that the optimum is attained by exercise policy $I^* = I^*(S_T)$, and any trading strategy of the form*

$$\theta(t) := \alpha(t)\left(Y^*(t,S_t) - Y_t^\theta\right), \tag{5.6}$$

where $\alpha(t)$ is any deterministic function with $\alpha(t) \longrightarrow \infty$ as $t \longrightarrow T$ fast enough to force the solution Y_t^θ of (5.2) to converge to $Y^*(t,S_t)$ as $t \longrightarrow T$.

We give a brief outline of the proof of this theorem in the following section. A detailed exposition can be found in Section 4.3. For general reference on the theory of controlled diffusion processes see Fleming and Soner (1993) or Krylov (1980).

5.3.3 Proof of Theorem 5.3.1

The proof of Theorem 5.3.1 is essentially an application of Theorem 4.3.1 in Chapter 4. As discussed there, the value function $V^*(t,s,y)$ can be characterised as the solution to the *Bellman equation* subject to appropriate boundary conditions. In the case considered here, the Bellman equation takes the form

$$\sup_\theta \left\{ V_t + \frac{1}{2}\left(\sigma_Z^2 V_{yy} + \sigma_S^2 V_{ss}\right) + (V_y + \Pi)\,\theta \right\} \equiv 0. \tag{5.7}$$

Since this equation is *linear* in θ, in order for the supremum to be finite the coefficient of θ must be zero. We can thus split up the Bellman Equation into the following two separate equations:

$$V_t + \frac{1}{2}\left(\sigma_Z^2 V_{yy} + \sigma_S^2 V_{ss}\right) \equiv 0, \tag{5.8}$$

$$V_y + \Pi \equiv 0. \tag{5.9}$$

Equation (5.8) is just the heat equation associated with two-dimensional Brownian motion. The Feynman-Kac representation theorem (see Karatzas and Shreve (1988, Theorem 4.4.2)) allows us to write its solution (if it exists) in the following form:

$$V(t,s,y) = E_t^{s,y}\left[\, V(T, S_T, Y_T^0)\,\right],$$

where Y_t^0 is a solution of (5.2) corresponding to $\theta \equiv 0$, i.e. it follows the simplified dynamics $dY_t^0 = \sigma_Z(t)\, dW_t^Z$. Equation (5.9) has a more economic interpretation. It states that the marginal gain $V_y(t,s,y)$ from holding an extra unit of the risky asset at time t, must equal the purchasing cost, $H(t,y)$, *minus* the expected marginal benefit, $F(t,s)$. For a detailed elaboration on the intuition behind this, see Section 4.3.4.

Optimal Strategy

We will now intuitively derive the appropriate boundary condition for the Bellman equation and see how it can help us pinpoint the optimal strategy. Suppose the value function $V^*(t,s,y)$ exists and is sufficiently smooth. We know then that is must solve the Bellman equations (5.8) and (5.9). For any given admissible strategy $\theta(t)$ consider now the process $V^*(t, S_t, Y_t^\theta)$. Following the arguments outlined in Chapter 4, an application of Itô's Lemma in conjunction with equations (5.8) and (5.9) allows us to write

$$V^*(0,s,y) = V^{\theta,I}(0,s,y) + E^{s,y}\Big[\, \underbrace{V^*(T, S_T, Y_T^\theta) - \rho I \phi(Y_T^\theta)}_{=:\, q^*(S_T, Y_T^\theta; I)}\,\Big]. \tag{5.10}$$

Here, $V^{\theta,I}(0,s,y)$ is the value function for the arbitrarily chosen policy (θ, I). The above equation has the following interpretation; the quantity $E^{s,y}[\ q^*(S_T, Y_T^\theta; I)\ ]$ measures the difference between the value $V^{\theta,I}(0,s,y)$ of any given policy pair (θ, I) and the maximal achievable value $V^*(0,s,y)$. The function $q^*(S_T, Y_T^\theta; I)$ can thus be identified as the *loss* incurred by the informed trader by deviating from the optimal strategy. In particular we deduce that $E^{s,y}[\ q^*(S_T, Y_T^\theta; I)\ ]$ must be non-negative, and zero at the optimum. We will reverse this argument to construct the optimal policy and value function. More specifically, we will construct a function $q^*(s,y;I)$ and "optimality conditions" $y^*(s)$ and $I^*(s)$ in such a way that $q^*(s,y;I) \geq 0$ and $= 0$ if and only if $y = y^*(s)$ and $I = I^*(s)$. Based on this we use (5.8) and the Feynman-Kac representation to construct a candidate $V^*(t,s,y)$ for the value function which satisfies (5.7). From this and the above arguments we can conclude that an exercise policy $I^*(S_T)$ and a trading strategy which ensures that $Y_T^\theta = y^*(S_T)$ will attain the optimum, and that $V^*(t,s,y)$ is indeed the value function of the informed trader's problem. The formal proofs for all these statements can be found in Chapter 4. We interpret $y^*(s)$ as the optimal "target position" towards which the informed trader should drive the aggregate position Y_t^θ in the market as $t \longrightarrow T$, conditional on $S_T = s$. This is similar to the arguments presented in Back (1992) and Back and Pedersen (1996).

As shown in Chapter 4, evaluating the Bellman equation (5.9) at $t = T$ implies that the loss function $q^*(s,y;I)$ must be of the form

$$q^*(s,y;I) = \int_y^{y^*(s)} \pi(s,\eta)\, d\eta + \rho I^*(s)\phi(y^*(s)) - \rho I \phi(y), \tag{5.11}$$

for suitably chosen $y^*(s)$ and $I^*(s)$. The intuitive interpretation of this equation is the same as the one discussed in Chapter 4: For any marginal deviation in y away from the optimal $y^*(s)$, the informed trader incurs a marginal loss of $f(s)$ due to the reduced position in the risky asset, while gaining a marginal $h(y)$ due to the lower cost of acquiring that position.

By construction, the loss function defined in (5.11) satisfies the desired "boundary conditions" $q^*(s, y^*(s); I^*(s)) = 0$. Our remaining task is hence to determine $y^*(s)$ and $I^*(s)$ in such a way that for each s the mapping $(y, I) \mapsto q^*(s, y; I)$ has a global minimum at $y = y^*(s)$ and $I = I^*(s)$. Following the arguments developed in Chapter 4, we first keep the option exercise decision fixed and construct optimal targets $y_1^*(s)$ and $y_0^*(s)$ conditional on the option being exercised or not, respectively. We can then establish the optimal exercise decision $I^*(s)$ endogenously. As shown in Chapter 4, the conditional target positions $y_0^*(s)$ and $y_1^*(s)$ can be characterised via the first-order conditions

$$f(s) = h(y_0^*(s)), \quad \text{and} \tag{5.12}$$

$$f(s) = h(y_1^*(s)) - \rho\phi'(y_1^*(s)). \tag{5.13}$$

It can be shown that both $y_0^*(s)$ and $y_1^*(s)$ are *increasing* functions of the signal s, and that $y_0^*(s) < y_1^*(s)$. As shown in Chapter 4, it is then optimal for the informed trader to exercise the option if and only if

$$\int_{y_1^*(s)}^{y_0^*(s)} \pi(s, \eta)\, d\eta \leq \rho\phi(y_1^*(s)). \tag{5.14}$$

Note that this condition shows that the optimal exercise policy depends only on s and not on y or I. The economic interpretation of (5.14) is as follows. Recall that $y_0^*(s)$ is the optimal target position conditional on no exercise. Hence the left-hand side of (5.14) is the loss incurred by trading towards $y_1^*(s)$ instead of $y_0^*(s)$. The right-hand side however represents the gain made from the option's pay-off. Hence (5.14) states that the informed trader should exercise the option if and only if the loss incurred by deviating from the optimal position is outweighed by the option's pay-off.

It is straight-forward to show that there exists a unique cut-off point $\hat{s}$ such that (5.14) is satisfied if and only if $s \geq \hat{s}$. We can hence define the optimal exercise policy $I^*(s)$ to be $= 1$ if and only if $s \geq \hat{s}$, and $= 0$ otherwise. Intuitively, the unconditional target $y^*(s)$ should then equal $y_1^*(s)$ whenever exercising the option is optimal, and $y_0^*(s)$ otherwise. Therefore we define

$$y^*(s) := I^*(s)y_1^*(s) + (1 - I^*(s))y_0^*(s).$$

The loss function $q^*(s, y; I)$ defined via (5.11) can then indeed be shown to have a global minimum at $y = y^*(s)$ and $I = I^*(s)$, see Proposition 4.A.3. From the arguments outlined in the beginning of this section it should now be intuitively clear that any strategy $\theta(t)$ which ensures that $Y_t^\theta \longrightarrow y^*(S_T)$ as $t \longrightarrow T$ is a candidate for the optimum. We hence set $Y^*(t, s) := E_t^s[\, y^*(S_T)\,]$ and define

$$\theta(t) := \alpha(t) \left(Y^*(t, S_t) - Y_t^\theta\right),$$

where $\alpha(t)$ is a deterministic function with $\alpha(t) \longrightarrow \infty$ as $t \longrightarrow T$ fast enough to make Y_t^θ converge to $y^*(S_T)$ as $t \longrightarrow T$. A rigorous proof for the optimality of $\theta(t)$ is given in the appendix to Chapter 4.

REMARK: We can write the overall first-order condition for the informed trader's terminal position in the following form:

$$h\left(y^*(S_T)\right) = f(S_T) + \rho I^*(S_T)\phi'\left(y^*(S_T)\right). \tag{5.15}$$

Since for an optimal strategy, we have $Y_T^\theta = y^*(S_T)$, the above equation implies that the terminal price of the asset, $h(Y_T^\theta)$, equals its expected true value, $f(S_T)$, *plus* a mark-up $\rho I^*(S_T)\phi'(Y_T^\theta)$. Thus, unlike in the models of Kyle (1985), Back (1992) or Back and Pedersen (1996), the existence of the option implies that the informed trader has an incentive to drive prices *away* from the true expected value of the asset.

VALUE FUNCTION

Finally, we can now construct a candidate for the value function $V^*(t,s,y)$ based on the analysis in the preceding sections. Guided by (5.10) we conjecture that the boundary condition for the value function be given by

$$v^*(s,y) := q^*(s,y;I) + \rho I \phi(y). \tag{5.16}$$

Obviously, $v^*(s,y)$ is smooth in y, and smooth in s for $s \in \mathbb{R} \setminus \{\hat{s}\}$. Moreover, a straight-forward calculation, using the first-order conditions for $y_0^*(s)$ and $y_1^*(s)$, shows that $v^*(s,y)$ is continuous also at $s = \hat{s}$. In view of (5.8), we define a candidate for the value function via the Feynman-Kac representation formula:

$$V^*(t,s,y) := E_t^{s,y}\left[\, v^*(S_T, Y_T^0)\,\right]. \tag{5.17}$$

It can be shown that the function thus defined indeed solves the Bellman Equation (5.7), see Proposition 4.A.4. From the analysis in the preceding sections it should then be clear that $V^*(t,s,y)$ *is* the value function, and that the strategy defined above *is* an optimal strategy.

REMARK: A direct computation shows that

$$\frac{\partial}{\partial y} V^*(t,s,y) = \underbrace{E_t^s\left[\, f(S_T)\,\right]}_{=\,F(t,s)} - E_t^y\left[\, h(Y_T^0)\,\right].$$

Recall that in order to satisfy the Bellman Equation, $V^*(t,s,y)$ must in particular solve (5.9). From the above we see that this is possible only if $H(t,y) = E_t^y[\, h(Y_T^0)\,]$. We have thus found a condition that the market maker's pricing rule must satisfy in order for the informed trader's problem to have a non-degenerate solution.

5.4 Equilibrium

In this section we characterise equilibria on both the market for the underlying asset as well as the OTC options market. Again we proceed in two stages. We first analyse equilibria on the market for the underlying asset, taking the informed trader's option position as given. We then solve for the overall equilibrium on both markets, endogenising the informed trader's option position.

Recall that equilibrium requires the market maker's pricing rule to be *rational*, i.e. to reflect the expected true value of the asset, conditional on the available information. For fixed ρ, denote the by $\nu_t(\rho; ds)$ the distribution of S_t, conditional on ρ and the observed order flow up to time t. If we denote by $\mu(d\rho)$ the market maker's beliefs about ρ, then the beliefs regarding the joint distribution of ρ and S_t can be written as $\Lambda_t(d\rho, ds) = \mu(d\rho) \otimes \nu_t(\rho; ds)$. Rationality of the market maker's pricing rule for the underlying asset can then be expressed as

$$H(t, Y_t) = \int_{\mathbb{R}^2} F(t, s)\, \mu(d\rho) \otimes \nu_t(\rho, ds). \tag{5.18}$$

To characterise the dynamics of Λ_t, we will make use of the theory of linear filtering. For general reference, see Kallianpur and Karandikar (1985).

5.4.1 Equilibrium on the Underlying Market

As benchmark, we first consider the two extreme cases in which the market maker either is not aware of the existence of options at all, or possesses perfect information regarding the informed trader's option position. Note that although these extreme cases do not exactly reflect the structure of our model, they nonetheless provide valuable insights and will help us later determine the joint equilibrium on both markets.

Not surprisingly, in the former case an equilibrium exists in which prices are set in the same way as in Back and Pedersen (1996). However, the latter perfect information case only permits an equilibrium in the degenerate situation where no options are held by the informed trader. The main focus of this section, however, will be the case in which the market maker is aware of the existence of options, but is uncertain about the number held by the informed trader.

Throughout this section we will take the informed trader's position ρ in option contracts as given and assume that it is fixed over the entire time interval $[0, T]$. We will endogenise the informed trader's position in Section 5.4.2.

INCOMPLETE INFORMATION

Consider the case in which the market maker is unaware of the existence of OTC derivatives. Formally, we can express this as $\mu = \delta_{\{0\}}$. Due to this misperception of the true market structure, systematic errors might arise in the expectations formed. As a consequence, we shall see that situations might arise in which the informed trader realises increased profits at the market maker's expense. This happens because in some situations it becomes profitable for the informed trader to create a bubble, pushing the options held (further) into the money. The market maker will not react to this mechanism as he is unaware of the existence of the options. We summarise the result in the following theorem.

Theorem 5.4.1 (Incomplete Information Equilibrium) *Define*

$$\Sigma(t) = \int_t^T (\sigma_Z^2(s) - \sigma_S^2(s))ds.$$

By rescaling S_t (and adjusting $F(t, S_t)$ accordingly), we can assume $\Sigma(t) > 0$ for all t. Then, in the case of incomplete information, an equilibrium on the underlying *market is given by the triple* (H, θ, I)*, where* $I = I^*(S_T)$*,*

$$\begin{aligned} H(t,y) &= E_t^y\left[\, f(Y_T^0)\,\right], \quad \textit{and} \\ \theta(t) &= \alpha(t)\left(Y^*(t,S_t) - Y_t^\theta\right). \end{aligned}$$

Here, $\alpha(t) = \sigma_Z^2(t)/\Sigma(t)$*, and* $Y^*(t,S_t)$ *and* $I^*(S_T)$ *are defined as in Section 5.3.3. In particular, the terminal pricing rule is given by* $h \equiv f$*.*

PROOF: Optimality of the informed trader's strategy follows from Theorem 5.3.1. The required convergence of Y_t^θ to $y^*(S_T)$ as $t \longrightarrow T$ for the chosen $\alpha(t)$ follows from the Law of Iterated Logarithms, see Back and Pedersen (1996, Lemma 2) for details. The proof of rationality of the pricing rule follows the lines of the proof in Back and Pedersen (1996). Details are given in Lemma 5.A.1 and Corollary 5.A.2 in the appendix.

□

Note that, even though the market maker's pricing rule is the same as the one obtained in Back and Pedersen (1996), the informed trader's incentive to manipulate markets is reflected in the term $Y^*(t,S_t)$ in the optimal trading strategy. Recall from Section 5.3.3, that the informed trader's first-order conditions imply

$$h\left(Y_T^\theta\right) = f(S_T) + \rho I^*(S_T)\phi'\left(Y_T^\theta\right).$$

That is, the informed trader creates a price bubble which manifests itself in the mark-up $\rho I^*(S_T)\phi'(Y_T^\theta)$ over the true expected value of the asset, $f(S_T)$.

The qualitative conclusion from the analysis in this subsection is interesting. If in the presence of asymmetric information derivative markets are not properly aligned with the underlying market, informed traders will exploit this miss-alignment, thus creating price bubbles in the market for the underlying asset. These bubble arises because the informational content on the derivatives market is not transferred to the market for the underlying asset. This lack of informational synchronicity between the underlying and the options market makes up the difference between the analysis in this section and the equilibrium analysis in Back (1993). There, the market maker can also observe the order flow on the options market, which makes it impossible for the informed trader to profitably manipulate markets. In light of this, one might think that letting the market maker audit the order flows on attached OTC markets should lead to more rational price setting. This conclusion is however not necessarily true. To see why this is so, and to understand better the results of Back (1993), we will need the analysis from the following subsection.

COMPLETE AND PERFECT INFORMATION

We now consider the situation where the market maker is fully aware of the existence of OTC contracts and knows the exact quantity held by the informed trader; formally $\mu = \delta_{\{\rho\}}$. We maintain however the assumption that the market maker cannot observe the informed trader's private signal S_t. In this situation equilibria can be characterised as follows.

Theorem 5.4.2 (Perfect Information Equilibrium) *In the case of complete and perfect information an equilibrium exists if and only if $\rho = 0$, i.e. if the informed trader holds no options at all. In this case, the equilibrium reduces to the one described in Theorem 5.4.1 with $Y^*(t, S_t) = S_t$, that is*

$$\begin{aligned} H(t,y) &= E_t^y\left[\, f(Y_T^0)\,\right], \quad \textit{and} \\ \theta(t) &= \alpha(t)\left(S_t - Y_t^\theta\right). \end{aligned}$$

In particular, the terminal pricing rule is given by $h \equiv f$.

Note that this equilibrium is, not surprisingly, identical to the one obtained in Back

(1992) and Back and Pedersen (1996).

PROOF: We may assume that $h'(y) > 0$. Note from (5.15) that the necessary first order condition for the informed trader's problem can be written as

$$f(S_T) = h(Y_T) - \rho I^*(S_T)\phi'(Y_T). \tag{5.19}$$

Together with the rationality requirement (5.18) this implies

$$\begin{aligned} h(Y_T) &= \int_{\mathbb{R}} f(s)\, \nu_T(\rho; ds) \\ &= h(Y_T) - \rho \int_{\mathbb{R}} I^*(s)\phi'(Y_T)\, \nu_T(\rho; ds). \end{aligned}$$

Consequently we can conclude that

$$\rho \int_{\mathbb{R}} I^*(s)\phi'(Y_T)\, \nu_T(\rho; ds) = 0$$

is a necessary condition for existence of equilibria. Since $h'(y) > 0$ and thus $\phi'(y) > 0$, and since $I^*(s) = 1$ at least for some values of s whenever $\rho \neq 0$, we can finally conclude that $\rho = 0$ is necessary for the existence of equilibria. The sufficiency follows easily from the analysis in the preceding section.

□

The intuition behind the non-existence of equilibrium lies in the the informed trader's incentive to manipulate prices and the fact that in this scenario the market maker anticipates this. From (5.19) we see that it is optimal for the informed trader is to mark up prices relative to the expected true value by an amount $\rho I^*(S_T)\phi'(Y_T^\theta)$. The market maker, anticipating this, would have to adjust prices to account for the price mark-up, only to see the informed trader increase demand to push prices even further up. Hence, no rational pricing rule can exist in this situation.

The results derived in this section are dramatically different from those obtained in Back (1993). If the order flow on the options market is fully observable, no non-degenerate equilibrium can exist in our model, while in Back (1993) this observability in fact increases price accuracy by providing the market maker with additional information regarding the underlying asset. This *qualitative* difference arises because the equilibrium in Back (1993) is upheld by the assumption of additional *exogenous* noise in form of stochastic volatility. Without this additional noise term the results in Back (1993) reduce to the non-existence result established here.

In summary we conclude that situations where market makers can identify trades on the derivatives market, but where prices on the two markets are not aligned, might cause market failure. A stable market equilibrium for the underlying asset can only be sustained if the prices for OTC derivatives are set in such a way that an informed trader would never find them worth holding. This requirement of price synchronicity necessary to prevent manipulative bubbles has previously been addressed in discrete-time models by Jarrow (1992) and Jarrow (1994), and in the related continuous-time generalisation by Frey (1996). However, while Jarrow and Frey simply *assume* price synchronicity, it is an *endogenous* consequence of the market microstructure of our model.

5.4.2 Full Equilibrium

We now return to the scenario in which the market maker is aware of the existence of the options market, but does not know the exact quantity ρ of options held by the informed trader. This is in line with the structure of our model, since we assume that the market maker maintains the OTC options market but cannot distinguish the identity of individual buyers of options. Following the arguments of the previous section, we find that rationality of the market maker's pricing rule together with the informed trader's first-order conditions imply

$$\begin{aligned} h(Y_T) &= \int_{\mathbb{R}^2} f(s)\,\mu(d\rho) \otimes \nu_T(\rho; ds) \\ &= h(Y_T) - \int_{\mathbb{R}} \rho \left(\int_{\mathbb{R}} I^*(s)\phi'(Y_T)\,\nu_T(\rho; ds) \right) \mu(d\rho). \end{aligned}$$

Following the arguments of the preceding section we can conclude that this is only possible if $\rho = 0$ almost surely with respect to μ. In other words, given the informed trader's optimal strategy, rational pricing on the underlying market is only possible if the market maker believes that the informed trader holds no options. The intuition behind this result is analogous to the perfect information case: Whenever the informed trader does hold an option, there will be an incentive to manipulate markets and create a price bubble. Hence, if the market maker has reason to believe that options are held by the informed trader, rationality would require prices to adjust to counteract this manipulation incentive. This in turn would induce the informed trader to increase demand even further. Hence, no rational pricing rule could exist.

As a consequence of this, we can conclude that the only feasible equilibrium is one in which it is *rational* for the market maker to believe that no options are held by the informed trader. This however is only possible if option prices are set in such a way that the informed trader is indifferent between holding or not holding options. In other words, the price of the option must equal the expected additional benefit that the informed trader could derive from holding it. We have thus established the main result of this section:

Theorem 5.4.3 (Full Equilibrium) *The full equilibrium on both markets is given by the quintupel $(H_0, H; \rho, \theta, I)$, where $\rho \equiv 0$ and $I \equiv 0$, and*

(1) The option pricing schedule, $H_0(\rho)$, is set such that the informed trader is indifferent *between holding or not holding options.*

(2) The market maker's pricing rule $H(t, Y_t)$ and the informed trader's trading strategy $\theta(t)$ are set as in Theorem 5.4.2.

Moreover, there exists no equilibrium in which the informed trader finds it optiomal to hold options.

PROOF: It is clear from Theorem 5.4.1 that (H, θ, I) is an equlibrium on the underlying market, if the market maker believes that the informed trader holds no options. This belief on the other hand is rational, since option prices are set in such a way that holding options does not provide any benefit for the informed trader. This in turn implies optimality of the informed trader's option position.

$\square$

The interpretation of this result is very much the same as in the perfect information case. The presence of asymmetric information imposes a requirement of synchronicity between prices on the options market and the market for the underlying asset. Failure to guarantee this synchronicity leads to market breakdown. Hence, equilibrium in our model endogenously determines a price schedule for OTC derivatives which differs drastically from classic option pricing theory. While in the latter, the price of an option reflects its value from the point of view of a price-taker, in our study option prices account for the value of the manipulation opportunity they provide. In the next section, we will analyse the equilibrium option prices predicted by our model in more detail and show that they give rise to the "smile pattern" of implied volatility.

5.5 Option Pricing and the Smile Pattern

In this section we will make explicit the implications that the equilibrium conditions derived in the preceding sections have on the pricing of options. From Section 5.4 we know that the only feasible equilibrium requires option prices to be aligned with the underlying market in such a way that the informed trader cannot derive any additional benefit from holding options. In other words, the equilibrium price $\rho H_0(\rho)$ of a quantity ρ of option contracts must equal the value that these contracts would have to the informed trader. In contrast to classic option pricing theory, the resulting pricing schedule turns out to be non-linear in quantity, i.e. the price per unit depends on the amount demanded. Moreover, in a special case in which equilibrium prices of the underlying asset are given by a classic Black-Scholes model, we are able to explicitly characterise the "implied volatility" and show that it displays the famous "smile pattern".

5.5.1 Arbitrage Option Price

We begin with establishing as a benchmark the arbitrage price of the option in our model. Recall from Section 5.4 that the equilibrium pricing rule for the underlying asset is given by

$$H(t,y) = E_t^y[\, f(Y_T^0)\,].$$

Assuming that the function $H(t,y)$ is sufficiently regular, the Feynman-Kac representation theorem, (see Karatzas and Shreve (1988, Theorem 5.7.6)), implies that $H(t,y)$ solves the partial differential equation $H_t + \frac{1}{2}\sigma_Z^2 H_{yy} = 0$. Hence, by Itô's Lemma the price process $P_t^\theta := H(t, Y_t^\theta)$ for any given strategy $\theta(t)$ satisfies

$$dP_t^\theta = H_y(t, Y_t^\theta)\, dY_t^\theta.$$

Inverting $H(t,y)$ we can write $Y_t^\theta = G(t, P_t^\theta)$, so that

$$dP_t^\theta = \underbrace{H_y(t, G(t, P_t^\theta))}_{=:\,\lambda(t, P_t^\theta)} \left\{\theta(t)\, dt + \sigma_Z(t)\, dW_t^Z\right\}. \tag{5.20}$$

This formulation implies that the market thus described is in fact complete, (see Harrison and Pliska (1981)), which ensures the existence of a unique equivalent martingale measure (EMM). In analogy with the above, we define the process P_t^0 by $P_t^0 := H(t, Y_t^0)$, i.e. the price process corresponding to the "do nothing" strategy. Obviously,

$$dP_t^0 = \sigma_Z(t)\lambda(t, P_t^0)\, dW_t^Z.$$

From (5.20) it is hence clear that the process P^θ under the EMM is distributionally equivalent with the process P^0 under the original measure. From standard derivatives pricing theory we know that the arbitrage-price of an option can be expressed as its expected terminal pay-off under the EMM. The above arguments imply that this is equivalent to the expected pay-off under the original measure with P^θ replaced by P^0. We have thus shown:

Proposition 5.5.1 (Arbitrage Price) *In equilibrium, the arbitrage price of an option which promises to pay an amount $\varphi(P_T^\theta)$ if exercised, is given by*

$$E\left[\,\varphi^+(P_T^0)\,\right] = E\left[\,\varphi^+(f(Y_T^0))\,\right]. \tag{5.21}$$

In particular, this price does not depend on the informed trader's strategy.

The arbitrage price for the option constitutes the benchmark with which we compare the equilibrium option price in our model. Note that the theory of derivatives pricing by arbitrage is based, amongst others, on the assumption of perfectly elastic markets for the underlying asset. This assumption is clearly not satisfied in the model considered here. Hence we expect equilibrium option prices to differ from arbitrage prices. This difference gives rise to the "smile pattern" of implied volatilities.

5.5.2 Equilibrium Option Price

Recall that the equilibrium price of the option is equal to the expected gain the informed trader could derive from holding it. Denote by V_ρ^* the value function associated with the informed trader's problem for a given quantity of options held, ρ. Following the argument developed in Chapter 4, we use (5.10) to write

$$V_\rho^* = V_\rho^{\theta,I} + E\left[\, q_\rho^*(S_T, Y_T^\theta; I)\,\right].$$

Here, q_ρ^* is the "loss-function" for the given ρ (see Section 5.3.3), and (θ, I) is *any* admissible policy pair. Taking in particular $\theta(t)$ to be the optimal trading strategy in the absence of the option, and $I \equiv 0$, we find that the value to the informed trader of ρ units of the option is given by

$$\begin{aligned}
(V_\rho^* - V_0^*) &= E\Big[\, q_\rho^*(S_T, y_0^*(S_T); I)\,\Big] \\
&= E\Bigg[\, I_\rho^*(S_T) \underbrace{\left\{ \int_{y_0^*(S_T)}^{y_1^*(S_T)} \pi(S_T, \eta)\, d\eta + \rho\phi(y_1^*(S_T)) \right\}}_{=:\, \psi_\rho\,(f(S_T))} \,\Bigg].
\end{aligned}$$

Note that by the informed trader's first-order conditions, the term in the brackets indeed only depends on $f(S_T)$, so that our notation is justified. Note also that by definition of the exercise policy, we have that $I_\rho^*(S_T) = 1$ if and only if $\psi_\rho(f(S_T)) \geq 0$. We can thus write

$$(V_\rho^* - V_0^*) = E\left[\, \psi_\rho^+\,(f(S_T))\,\right], \tag{5.22}$$

which in particular shows that the value of the option to the informed trader is non-negative. The above expectation can be seen to be the arbitrage price of a derivative which pays off an amount $\psi_\rho(f(S_T))$ if exercised. Indeed, consider the fictitious price process $\hat{P}_t := F(t, S_t)$. Note that $\hat{P}_t$ represents the best estimate of the *true* value of the asset, given the informed trader's signal. Using Feynman-Kac and Itô we find

$$d\hat{P}_t = \sigma_S(t) F_s(t, S_t)\, dW_t^S.$$

In other words, the process $\hat{P}_t$ is a martingale under the original measure, so that (5.22) is indeed the arbitrage price of the derivative that pays off $\psi_\rho(\hat{P}_T)$ if exercised. We summarise our findings in the following

Proposition 5.5.2 (Equilibrium Price) *The equilibrium price of the option is given by*

$$H_0(\rho) = \frac{1}{\rho} E\left[\, \psi_\rho^+ \left(f(S_T)\right) \,\right]. \tag{5.23}$$

This expression can be interpreted as the arbitrage price of an option which pays off an amount $\psi_\rho(\hat{P}_T)$ if exercised, written on the (expected) true value of the underlying asset, $\hat{P}_t = F(t, S_t)$.

Note in particular that the option price $H_0(\rho)$ may depend on the quantity demanded, ρ. This reflects the fact that the equilibrium price of the option, unlike its arbitrage price, is indeed affected by the informed trader's actions.

5.5.3 Implied Volatility and the Smile Pattern

Throughout this section, we restrict our analysis to the special case in which the option in question is a standard European call option, i.e. $\varphi(P_T) = (P_T - K)$. To make the dependence on the strike price explicit, we denote the equilibrium price of such an option by $H_0^C(\rho; K)$. Also, in order to be able to compare our results with those of classic Black-Scholes theory, we consider the special case where $f \equiv \exp$.

In this case, the fictitious price process $\hat{P}_t$ follows a generalised geometric Brownian Motion:

$$\frac{d\hat{P}_t}{\hat{P}_t} = \sigma_S(t)\, dW_t^S. \tag{5.24}$$

Note that this is just the standard Black-Scholes model with zero interest rates. Recall that in equilibrium, the market maker's pricing rule is set to $h \equiv f \equiv \exp$. It is now straight-forward to see that the informed trader's first-order conditions (5.12) and (5.13) in this case take on the following form:

$$y_0^*(s) = s \qquad \text{and} \qquad y_1^*(s) = s - \log(1-\rho).$$

In order to calculate the *unit price* of the option, we need to evaluate the term $\psi_\rho(\hat{P}_T)/\rho$. A direct computation yields

$$\begin{aligned} \frac{1}{\rho}\psi_\rho(\hat{P}_T) &= \frac{1}{\rho}\Big\{-\hat{P}_T\log(1-\rho) - \rho K\Big\} \\ &= \underbrace{\Big(-\frac{\log(1-\rho)}{\rho}\Big)}_{=:\,\kappa(\rho)}\Big\{\hat{P}_T - \underbrace{\Big(-\frac{\rho}{\log(1-\rho)}\Big)}_{=\,1/\kappa(\rho)} K\Big\} = \kappa(\rho)\{\hat{P}_T - K/\kappa(\rho)\} \end{aligned}$$

That is, the pay-off of the fictitious derivative which represents the equilibrium option price in our model can be expressed as a multiple of a simple European call option with a modified strike price $K/\kappa(\rho)$. We have thus established the main result of this section:

Theorem 5.5.3 *In the case where the informed trader's expectations of the true value of the asset follow a geometric Brownian Motion, i.e. $f \equiv \exp$, the equilibrium price $H_0^C(\rho; K)$ of a European call option with strike price K is given by*

$$H_0^C(\rho; K) = \kappa(\rho) C^{BS}(\Sigma_0^S, \hat{P}_0, K/\kappa(\rho)) = C^{BS}(\Sigma_0^S, \kappa(\rho)\hat{P}_0, K). \tag{5.25}$$

Here $C^{BS}(\Sigma, P, K)$ denotes the Black-Scholes price of a European call option with strike price K, given initial price P and price volatility Σ. Note that we need to use the Black-Scholes pricing theory for time-dependent volatility, with

$$\Sigma_t^S := \frac{1}{T-t}\int_t^T \sigma_S^2(\tau)\, d\tau.$$

The fact that we can identify the equilibrium price of the option as the Black-Scholes price of an option with modified strike price allows us to establish the usual comparative statics results:

Corollary 5.5.4 (Comparative Statics) *The equilibrium price $H_0^C(\rho; K)$ of the European call option is* decreasing *in the strike price K,* increasing *in the signal volatility σ_S, and* increasing *in the initial level of the signal, S_0. Moreover, it is* increasing *in the number of options demanded, ρ.*

All these results follow immediately from the corresponding properties of standard Black-Scholes option prices. Only the last, the dependence of the option price on quantity, is new compared to standard derivatives pricing theory. It is this property which is responsible for the emergence of the "smile pattern" of implied volatility in our model. By virtue of the above theorem, it is clear that implied volatility can be defined as the solution $\Sigma_0^{\mathrm{imp}}(\rho; K)$ of

$$C^{BS}(\Sigma_0^S, \kappa(\rho)\hat{P}_0, K) = C^{BS}(\Sigma_0^{\mathrm{imp}}(\rho; K), \hat{P}_0, K) \tag{5.26}$$

This explicit expression allows us to examine the shape of the "smile pattern" analytically. It can be shown that for most reasonable parameter values, implied volatility is a decreasing function of the option's strike price, cf. Appendix 5.B. These analytical results are confirmed by numerical computations. Calculations for different choices of the model parameters show an implied volatility curve which is monotonically decreasing and convex as a function of the option's strike price. Figure 5.1 shows an example. Both the theoretical as well as the numerical results are consistent with the "skew smile pattern" which has been observed in most financial markets since the October 1987 crash.

As a final point, let us mention that our story of the smile being a consequence of more aligned derivative prices is consistent with recent empirical findings of Christensen and Prabala (1998). They find that after the October 87 crash implied volatilities have become far better predictors of realised volatilities and hence that option pricing schedules have been more consistent and stable after the crash.

5.6 Conclusions

We have studied the manner in which an option, held by a large, informed trader, changes the nature of equilibrium prices for the derivative and the underlying asset. We have shown that in the presence of such an option, the informed trader has an incentive to manipulate prices to push the option into the money, thus creating a price mark-up over the expected true value of the underlying asset. The cost of this manipulative bubble would have to be borne by uninformed traders. As a consequence, it is impossible for the market maker to sustain a stable market equilibrium if he has reason to believe that options are held by the informed trader. The implication of this is that in equilibrium, option prices have to be in line with the market for the underlying asset. We have demonstrated that this synchronicity requirement gives rise to a pricing schedule for options which is non-linear in quantity. Moreover, we were able to derive an explicit expression for implied volatility, and show that it displays the famous "skew smile pattern" which is one of the most prominent empirical puzzles in the theory of financial markets.

There are several conclusions that can be drawn from our analysis. To begin with, our results show that in the presence of asymmetric information, an OTC market for derivative securities which is insufficiently aligned with the market for the underlying can cause market failure. Moreover, bringing the OTC derivatives market in line is only possible as long as private information has not yet arrived. As soon as private information is present, the market maker is unable to set derivatives prices in such way that informed traders can be deterred from manipulating prices and thus cause market equilibrium to break down. Although this problem could be circumvented by introducing additional noise into the options market as in Back (1993), we believe our results to be important and interesting both from a theoretical as well as a practical point of view, in particular in view of the current debate concerning the regulation of derivatives markets. One extreme implication of our results is that derivatives markets would have to be closed down whenever there is asymmetric information regarding the underlying asset. This would be necessary to maintain stability in the underlying market.

We have shown how the smile pattern of implied volatility arises endogenously as a consequence of maintaining market stability. This is in contrast with the majority of the existing literature, where the smile is a result of additional exogenous noise or, as is the case in Platen and Schweizer (1998), of an equilibrium which is inherently unstable. The shape of the smile pattern predicted by our model displays the same skewness that has been observed in most financial markets since the October 1987 crash. This provides an interesting new angle on the interpretation of this phenomenon. Our model explains the smile as a consequence of the attempt to ensure market stability. This is consistent with market participants and regulators, having been alerted by the 1987 crash, have implemented measures to reduce market manipulation and increase market stability.

5.A Appendix: Complements to Section 5.4

Lemma 5.A.1 (Filtering Problem) *Suppose the terminal pricing rule is given by $h \equiv f$, and that the informed trader follows the strategy defined in Theorem 5.4.1. In the case where no options are held, i.e. $\rho \equiv 0$, the distribution $\nu_t(0, \cdot)$ of the informed trader's signal, S_t, conditional on the market maker's information set, is normal with mean Y_t^θ and variance*

$$\Sigma(t) = \int_t^T \left(\sigma_Z^2(s) - \sigma_S^2(s)\right) \, ds.$$

In particular, at the terminal date T, the market maker can infer the signal S_T with certainty, since $S_T \equiv Y_T^\theta$ almost surely.

PROOF: The proof is identical to the one in Back and Pedersen (1996). For $\rho \equiv 0$ and $h \equiv f$, the first-order conditions (5.12) and (5.13) for the informed trader's problem simplify to $y^*(S_T) \equiv S_T$, so that the informed trader's strategy takes the form

$$\theta(t) = \alpha(t) \left(S_t - Y_t^\theta\right).$$

Let $\xi_t := S_t - Y_t$; clearly,

$$\begin{aligned} d\xi_t &= -\alpha(t)\xi_t \, dt - \sigma_Z(t) \, dW_t^Z + \sigma_S(t) \, dW_t^S, \\ dY_t^\theta &= \alpha(t)\xi_t \, dt + \sigma_Z(t) \, dW_t^Z. \end{aligned}$$

These can be interpreted as the transition and measurement equations of a linear continuous-time filtering problem. By linearity, the conditional distribution of ξ_t

given Y_t^θ is normal. Denote its mean by $\hat{\xi}_t$ and its variance by $\hat{\Sigma}_t$. The corresponding Kushner equation (cf. Kallianpur and Karandikar (1985)) implies the following system of equations for $\hat{\xi}_t$ and $\hat{\Sigma}_t$:

$$\begin{aligned} d\hat{\xi}_t &= -\alpha(t)\hat{\xi}_t\, dt + \frac{1}{\sigma_Z(t)}\left(\alpha(t)\hat{\Sigma}_t - \sigma_Z^2(t)\right) dU_t, \\ \frac{\partial}{\partial t}\hat{V}_t &= -2\alpha(t)\hat{\Sigma}_t + \sigma_Z^2(t) + \sigma_S^2(t) - \frac{1}{\sigma_Z(t)}\left(\alpha(t)\hat{\Sigma}_t - \sigma_Z^2(t)\right)^2. \end{aligned}$$

Here, U_t is the innovation process, which is a standard Wiener process with respect to the filtration generated by the market maker's information. It is easy to see that this system is solved by $\hat{\xi}_t \equiv 0$ and $\hat{\Sigma}_t \equiv \Sigma(t)$, see Back and Pedersen (1996, Lemma 3).

□

Corollary 5.A.2 (Rationality) *Under the assumptions of Lemma 5.A.1, the pricing rule defined in Theorem 5.4.1 is* rational, *i.e. prices reflect expected terminal asset values conditional on the market maker's information.*

PROOF: From Lemma 5.A.1 we know that S_t conditional on Y_t^θ is normal with mean Y_t^θ and variance $\Sigma(t)$. Hence, S_T conditional on Y_t^θ is also normal with mean Y_t^θ and variance

$$\Sigma(t) + \int_t^T \sigma_S^2(s)\, ds = \int_t^T \sigma_Z^2(s)\, ds,$$

which is the variance of Y_T^0 conditional on Y_t^θ. Hence, the expected terminal asset value, conditional of $Y_t^\theta = y$, is

$$E_t^y\left[\, f(Y_T^0)\,\right],$$

which was to be shown.

□

5.B Appendix: Complements to Section 5.5

Proposition 5.B.1 (Smile Pattern) *For K sufficiently small,*

$$\frac{\partial \Sigma_0^{imp}}{\partial K}(\rho; K) < 0$$

PROOF: From the Black-Scholes option pricing formula it is immediate to see that

$$P \frac{\partial C^{BS}}{\partial P}(\Sigma, P, K) = -K \frac{\partial C^{BS}}{\partial K}(\Sigma, P, K).$$

Differentiating (5.26) with respect to K, we find

$$\begin{aligned}
&\frac{\partial \Sigma_0^{\text{imp}}}{\partial K}(\rho; K) \frac{\partial C^{BS}}{\partial \Sigma}(\Sigma_0^{\text{imp}}, \hat{P}_0, K) \\
&= \frac{\partial C^{BS}}{\partial K}(\Sigma_0^S, \kappa(\rho)\hat{P}_0, K) - \frac{\partial C^{BS}}{\partial K}(\Sigma_0^{\text{imp}}, \hat{P}_0, K) \\
&= \frac{P}{K}\left(-\kappa(\rho)\frac{\partial C^{BS}}{\partial P}(\Sigma_0^S, \kappa(\rho)\hat{P}_0, K) + \frac{\partial C^{BS}}{\partial P}(\Sigma_0^{\text{imp}}, \hat{P}_0, K)\right)
\end{aligned}$$

It is well-known that a European call option's "vega", i.e. the derivative of its Black-Scholes price with respect to volatility, is positive. Hence,

$$\text{sign}\left(\frac{\partial \Sigma_0^{\text{imp}}}{\partial K}(\rho; K)\right) = \text{sign}\left(-\kappa(\rho)\frac{\partial C^{BS}}{\partial P}(\Sigma_0^S, \kappa(\rho)\hat{P}_0, K) + \frac{\partial C^{BS}}{\partial P}(\Sigma_0^{\text{imp}}, \hat{P}_0, K)\right).$$

Finally, it is straight-forward to see that

$$\frac{\partial C^{BS}}{\partial P} \longrightarrow 1 \qquad \text{for } K \longrightarrow 0.$$

Hence, since $\kappa(\rho) > 1$, the term in the brackets eventually becomes negative for K small enough. This proves the desired result.

□

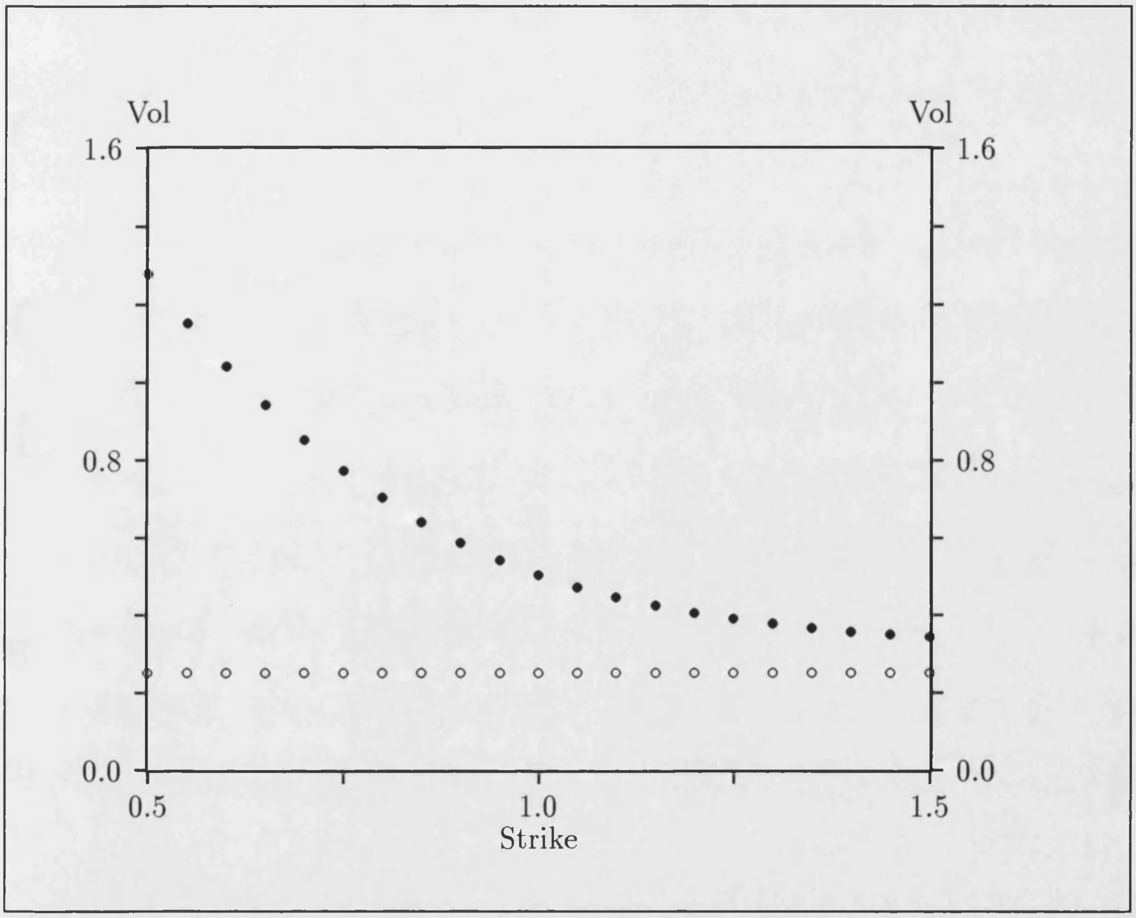

Figure 5.1: The Smile Pattern

This graph shows the implied volatility (•) compared to the signal volatility (○) as a function of the option's strike price. The informed trader is assumed to hold an option on 25% of the market.

Chapter 6

Options as Exchange Rate Policy Instruments

6.1 Introduction

Traditionally, central bank interventions in the foreign exchange market are implemented almost exclusively by spot transactions. Historically, the intensity of intervention increased steadily after the collapse of the Bretton Woods treaty. In the 80's, while the US adopted a "laissez-faire" approach, European central banks had to rely on intervention to keep exchange rates within the bounds of the Exchange Rate Mechanism. Europen intervention culminated in 1985, when the (G5) countries[1] decided to "depreciate the dollar in an orderly fashion", and even more in the late 80's, when the (G6) agreed in the Louvre Accord of February 1987 to "stabilise

[1]US, Japan, Germany, France, UK

exchange rates around their exisiting levels".

The mechanics of traditional central bank intervention are comparatively straightforward. Central banks would simply buy or sell foreign currency in the spot market, thereby affecting money supply. In a "sterilised" intervention, this transaction is neutralised by an offsetting domestic transaction. The survey by Edison (1993) studies the objectives and effectiveness of central bank intervention in the foreign exchange markets. Edison concludes that most central banks' main objective is to "smooth nominal [and real] exchange rates and to achieve a target level of nominal exchange rates." While the evidence with regards to sterilised interventions is inconclusive, unsterilised intervention has proven effective in controlling exchange rates. See also Almekinders (1995).

The use of *options* as instruments of exchange rate policy was first proposed by Taylor (1995). He suggests that central banks should buy *put* options, written on the domestic currency. When the domestic currency depreciates as a result of a speculative attack, Taylor argues, the option is deep in the money allowing the central bank to buy foreign currency at deflated rates, which can then be sold in the spot market to support the domestic currency. Breuer (1997) identifies five main drawbacks of the approach suggested by Taylor, the most important of which is related to what is called the destabilising effect of "delta hedging": The options that the central bank buys are issued by market makers, who incur substantial risk from this transaction. To manage this risk, market makers would typically adopt trading strategies designed to replicate the option's pay-off in order to eliminate the potential liability arising from their short option position. This kind of activity is usually referred to as "delta hedging". In the scenario discussed here, the hedging strategy would require market makers to buy in rising markets and to sell when markets fall. Thus, the implementation of such strategies creates a positive feedback effect, amplifying exchange rate fluctuations. A detailed theoretical analysis of the destabilising effects of dynamic hedging is provided in Chapter 2. Empirical evidence supporting the importance of such effects has been reported, amongst others, by Malz (1995). He finds that even in large, extremely liquid foreign exchange markets, hedging activities can have a significant impact on exchange rates. Similar results have been reported by the Group of Ten (1993).

In view of this, both Wiseman (1996) and Breuer (1997) independently proposed an alternative way of utilising options as exchange rate policy instruments. Rather then *buying* put options, they argue, the central bank should *write* call options. As a result, market participants (other than the central bank itself) would end up with a net long position in option contracts. Strategies designed to hedge the risk inherent in such a position would exhibit negative elasticity, requiring the hedger to buy in falling markets and to sell when markets rise. Thus, Wiseman and Breuer argue, by issuing options the central bank could induce market participants to *stabilise* exchange rates by delta hedging. The results of Chapter 2 confirm that *if* market participants content themselves with hedging their long position in currency options, a stabilising effect could indeed be achieved. However, the mechanism proposed by Wiseman (1996) and Breuer (1997) makes very explicit use of the fact that exchange rates are affected by the hedgers' trading activities. In other words, market manipulation is not an unwanted side-effect but rather the very tool by which the mechanism is implemented. This raises an important question: While it is true that the option induces extremely risk-averse market participants to hedge their positions thus stabilising rates, the situation might change drastically if instead the option is held by speculators who, being risk-neutral, have no incentive to hedge risk but to maximise expected profits.

In this chapter, we consider the situation in which a speculator, having bought the option issued by the central bank, instead of simply hedging her position strategically exploits the leveraged position in the market provided by the option. As theoretical framework, we chose the continuous-time model developed in Chapter 4, which in turn is based on the Back (1992) model. In contrast to the latter however, we assume prices to be given by Walrasian equilibrium rather than set by a market maker. The risk-neutral speculator interacts with "information traders", who base their demand for foreign exchange on the flow of fundamental "news". We assume that the central bank's objective is to stabilise exchange rates and to keep them within a given "target zone". While it is debatable whether this objective is in line with maximising social welfare, the discussion of this question is well beyond the scope of this chapter. We rather content ourselves with noting that this type of objective on behalf of central banks is supported by the empirical evidence, as outlined for example in Edison (1993). In order to achieve their objective, the central bank can intervene using spot transactions or by issuing options. In addition to trading in the spot market, the speculator may purchase the options issued by the central bank.

It turns out that in the absence of options, the speculator has no incentive to manipulate markets. In fact, in this case it is optimal not to trade at all in the spot foreign exchange market. The reason for this is that any potential gain from manipulation is neutralised by the price pressure that results when the speculator "unwinds" her position to realise these gains. In other words, in the absence of options it is only the arrival of "bad news" and the reaction to this by information traders that might push exchange rates outside the target zone. The violation of the target zone hence occurs as a consequence of rational behaviour based on the fundamental value of the domestic currency, rather than as the result of a speculative attack. An option issued by the central bank however creates an alternative way for the speculator to realise the gains from market manipulation. Upon arrival of "bad news", when information traders cause the domestic currency to depreciate, the option creates an incentive for the speculator to squeeze exchange rates even further. Scenarios can arise in which without speculation, the central bank's foreign currency reserves would have been sufficient to support the domestic currency, while they cannot sustain the additional pressure arising from the speculator's manipulation. In other words, instead of *protecting* against speculative attacks, options issued by the central bank in fact *create* an additional vehicle for such attacks.

The remainder of this chapter is organised as follows. In Section 6.2, we introduce the model and the mathematical framework for our analysis. The following section gives a brief account of the speculator's optimisation problem, following closely the arguments developed in Chapter 4. Section 6.4 analyses the structure of equilibria in our model. Two distinct cases are considered; first the case in which the speculator simply hedges her option position, and second the case in which she acts strategically. Section 6.5 concludes.

6.2 The Model

We formulate the model in the general framework developed in Chapter 4, which in turn is based on the model introduced in Kyle (1985) and Back (1992). However, while in the latter prices are set by a competitive market maker, we consider here a Walrasian auction mechanism to determine equilibrium prices.

TRADED ASSETS:

Two assets are traded continuously over the period from time zero to some final date T. The first asset is a riskless domestic bond or money market account. We use the bond as numéraire, normalising its price to one thus making interest rates implicit in our model. The second asset, which we shall refer to as "foreign currency", is risky. We denote the foreign currency's price at time t by P_t, and refer to P_t sometimes also as the *exchange rate*. Note that this is the inverse of the notion of exchange rate as it is used in the U.K., where it specifies the amount of foreign currency that can be bought with one unit of domestic currency.

AGENTS:

There are three types of agents in the economy, *information traders*, the *central bank*, and a *rational speculator*. Without at this point being specific about the precise nature of information traders, we assume that their demand for the foreign currency at any time t, given a quoted exchange rate p, is given by the demand function

$$p \mapsto D(t, Z_t, p). \tag{6.1}$$

Here, Z_t summarises the information available to information traders, which determines their perception of the "true" value of the exchange rate. We can interpret Z_t as a flow of macroeconomic "news" affecting the fundamental value of the domestic versus the foreign currency. We will provide more specific examples of such demand functions in later sections. For the moment, we content ourselves with assuming that $D(t, z, p)$ is differentiable and

$$\frac{\partial}{\partial z} D > 0, \qquad \frac{\partial}{\partial p} D < 0. \tag{6.2}$$

In other words, a decline in Z_t means "bad news" for the foreign currency and "good news" for the domestic currency, causing information traders to short some of their foreign currency holdings.

The central bank on the other hand, implements an exchange rate policy aimed at stabilising exchange rates. In order to achieve this, the central bank may trade on the foreign exchange market. More specifically, the central bank's demand at time t for the foreign currency, given a quoted rate p, is given by the policy function

$$p \mapsto C(t, Z_t, p). \tag{6.3}$$

Note in particular that we allow the central bank's exchange rate policy to depend on Z_t, i.e. we assume (quite reasonably) that macroeconomic news are observable by the central bank. Since the central bank seeks to stabilise exchange rates, it is reasonable to assume that it sells foreign currency when its price is too high, and buys it when the price is low. More specifically, we assume

$$\frac{\partial}{\partial z} C \leq 0, \qquad \frac{\partial}{\partial p} C \leq 0. \tag{6.4}$$

Note in particular that, since the central bank's objective is exchange rate stabilisation, it might be forced to act *against* fundamental macroeconomic news. This is reflected in the non-positive relation between the information process Z_t and the central bank's intervention function. In order to study speculative attacks on the domestic currency, we also assume that the central bank's reserves of foreign currency are limited. Formally, we express this by assuming that $C(t, z, p)$ is *bounded from below*, $C(t, z, p) \geq -\underline{C}$ for some constant $\underline{C} > 0$. We will provide more explicit specifications of the central bank's policy function later.

The third player in the market, the rational speculator, is assumed to be a risk-neutral, perfectly informed profit maximiser. The speculator in this model plays the role of the "large trader" in Chapter 4. At any time t, the speculator observes Z_t and determines his or her demand X_t for the foreign currency. For reasons of expositional simplicity, we restrict the speculator to using trading strategies $\theta(t)$ so that the demand X_t is given by $dX_t = \theta(t)\,dt$. Note in particular that this assumption forces the demand process X_t to be *absolutely continuous*. Although this assumption might seem unnecessarily restrictive, it can be shown that even if the speculator is allowed to choose from a larger strategy set, the optimal strategy will turn out to be absolutely continuous; see Back (1992).

In order to be able to study the effectiveness of options used as exchange rate policy instrument, we furthermore allow the central bank to issue options, written on the foreign currency. We assume that at the time when the central bank issues these options, it cannot observe whether the buyer is an information trader or a speculator. We assume that the option, if exercised, pays off an amount $\varphi(P_T)$ at time T. A simple example would be a European call option with strike price K, in which case $\varphi(P_T) = (P_T - K)$. Options are sold only at date zero, and cannot be traded thereafter. In addition to the trading strategy $\theta(t)$, the speculator hence has to choose an exercise policy I for the option, which we assume to be $= 1$ if the option is exercised and $= 0$ if not. As in Chapter 4, the speculator's objective is then to choose a trading strategy $\theta(t)$ and an exercise policy I such as to maximise expected terminal wealth. Note that both the optimal strategy as well as the exercise policy will depend on the number of options the speculator holds (which could be zero).

EQUILIBRIUM:

In the context of foreign exchange markets, it is appropriate to assume that the total supply of foreign currency is zero, reflecting the fact that all that matters is the *net* balance of foreign currency held by domestic traders versus domestic currency held by traders in the foreign country. The *equilibrium exchange rate* at any time t is then given as the solution P_t of the *market clearing equation*

$$D(t, Z_t, P_t) + C(t, Z_t, P_t) + X_t \equiv 0. \tag{6.5}$$

In order to bring the speculator's profit maximisation problem into the formal framework of Chapter 4, we define the *pricing rule* $H(t, z, x)$ associated with the equilibrium equation (6.5) implicitly via

$$D(t, z, H(t, x, z)) + C(t, z, H(t, x, z)) + x \equiv 0. \tag{6.6}$$

By differentiating this equation, we find that the hypotheses (4.1) imposed on the pricing rule in Chapter 4, are satisfied if

$$\frac{\partial}{\partial z} D + \frac{\partial}{\partial z} C \geq 0.$$

As we did in Chapters 4 and 5, we denote by corresponding lowercase letters any function evaluated at the terminal date T, for example, we write $h(x, z) = H(T, x, z)$. The terminal value to the speculator of the foreign currency is, just as in Section 4.2.2, the price at which the speculator can "unwind" his or her position in the market at time T. Denote this price by P_T^0. Obviously, $P_T^0 = h(0, Z_T)$. Since there is no informational asymmetry in this model, the speculator's expectation at time T of the terminal value of the asset is hence also given by P_T^0. Note that this notion of "unwind price" is in contrast with the majority of the literature. In most studies of similar problems of optimal trading, the terminal value of the asset is simply given by its current market price. In other words, this approach neglects the price pressure arising from the attempt to actually *realise* potential gains from trade. We therefore believe our specification to be more realistic. Note also that this specification actually *reduces* the speculator's incentive to manipulate markets, a fact which allows us to place more emphasis on the effect of options as central bank policy instrument. More precisely, in the absence of the option, any potential gain from market manipulation is consumed by the depreciation caused by the supply pressure that arises when the speculator "unloads" her position in the market. The option however allows the speculator to realise the gains from market manipulation without affecting the underlying exchange rates.

Finally, we define the "fundamental value" of the foreign currency as the equilibrium price that would obtain in the absence of both the speculator as well as the central bank. Denote this value at time t by P_t^F. Obviously, $P_t^F = H^F(t, Z_t)$, where $H^F(t, z)$ is the solution of the equation

$$D(t, z, H^F(t, z)) \equiv 0.$$

We introduce this value only as a benchmark, since it reflects the information traders' perception of what the exchange rate *should be* in the absence of speculation or intervention, given their information. We will use this benchmark in Section 6.4 to demonstrate how in our model situations are likely to occure in which,

(a) the fundamental exchange rate is well within the central bank's target zone, so that in the absence of speculation no intervention would be necessary,

(b) even when the speculator is active, there is no incentive for manipulation, so that still no intervention is needed, but

(c) if the speculator holds an option, the manipulation incentive thus created is strong enough for the exchange rate to be pushed outside the target zone even though the central bank exhausts all its foreign currency reserves in the attempt to support the domestic currency.

6.2.1 Mathematical Setup and Notation

As in Chapter 4, we will assume for expositional simplicity that the fundamental news process, Z_t, is given by a generalised Brownian motion. More specifically, we assume that Z_t is a solution to the stochastic differential equation

$$dZ_t = \sigma_Z(t)\, dW_t^Z, \tag{6.7}$$

where W_t^Z is a standard Brownian motion, and $\sigma_Z(t)$ is a deterministic function. Note however that the results of this chapter can easily be extended to more general diffusion processes. A trading strategy for the speculator is a process $\theta(t)$, adapted to the filtration generated by Z_t, such that the corresponding demand process

$$X_t^\theta := X_0^\theta + \int_0^t \theta(s)\, ds \tag{6.8}$$

is well-defined. Note that we assume the speculator to possess perfect information regarding the structure of the economy, in particular the information trader's demand function $D(t,z,p)$ and the central bank's policy function $C(t,z,p)$. Thus, the news process Z_t is a sufficient state variable for the speculator's optimisation problem.

Let $P_t^{x,z}$ be a weak solution to (6.7) and (6.8) conditional on $Z_t = z$ and $X_t^\theta = x$, defined on some suitable measurable space $(\Omega, \mathcal{F})$. We will omit the superscripts x or z whenever there is no ambiguity, and we also write $P^{x,z}$ for $P_0^{x,z}$. Denote by $(\mathcal{F}_t)_{t\geq 0}$ the filtration on $(\Omega, \mathcal{F})$ generated by Z_t, augmented to satisfy the "usual conditions", see Karatzas and Shreve (1988, Section 1.2) for details.

6.3 The Speculator's Problem

In this section, we will outline the solution of the speculator's optimisation problem. From the speculator's point of view, the equilibrium pricing rule $H(t, X_t, Z_t)$ can be considered exogenously given, so that the optimisation problem to be solved is a special case of the problem considered in Chapter 4. Therefore, we give here only a brief outline of the main result, Theorem 6.3.1. A detailed exposition can be found in Chapter 4.

6.3.1 Objective Function

The speculator's objective is to maximise expected terminal wealth. Following the intuition developed in Chapter 4, terminal wealth in the absence of the option is given by

$$B_T + V_T \cdot X_T = B_0 + V_T \cdot X_0 + \int_0^T \left(V_T - P_t^-\right) \theta(t)\, dt. \tag{6.9}$$

Since B_0 and X_0 are exogenously given and thus irrelevant for the optimisation problem, we will assume $B_0 = X_0 = 0$ and focus on the last term in the above expression. Given any strategy $\theta(t)$, the dynamics of the speculator's demand process, X_t^θ, are given by (6.8). Equilibrium prices are given by the pricing rule $P_t = H(t, X_t^\theta, Z_t)$. Recall that from the speculator's point of view, the terminal value of the foreign currency is given by the "unwind value" $P_T^0 = h(0, Z_T)$. In the context of Chapter 4, we can formalise this as $f(z) = h(0, z)$. Consequently, the *expected* terminal value at any earlier point in time $t < T$ is given by $F(t, z) = E_t^z[\ f(Z_T)\ ]$. Using the law of iterated expectations, we can write the speculator's expected terminal wealth for given strategy $\theta(t)$, conditional on $X_0 = x$, $Z_0 = z$, as

$$E^{x,z}\left[\int_0^T \left(P_T^0 - P_t\right) \theta(t)\, dt\right] = E^{x,z}\left[\int_0^T \underbrace{\left(F(t, Z_t) - H(t, X_t^\theta, Z_t)\right)}_{=:\, \Pi(t, X_t^\theta, Z_t)} \theta(t)\, dt\right]$$

Note that the function $\Pi(t, X_t^\theta, Z_t)$ has an economic interpretation: It describes the expected marginal benefit of holding an extra unit of the risky asset, $F(t, Z_t)$, minus the current purchasing price, $H(t, X_t^\theta, Z_t)$. Following our convention, we write

$$\pi(x, z) := \Pi(T, x, z) = f(z) - h(x, z).$$

We now incorporate the option into the speculator's objective function. For this part of the analysis, we take the number of options held by the speculator as exogenously given. Recall that the option, if exercised, pays off an amount $\varphi(P_T)$ at time T. Since the price P_T at time T is given by $P_T = h(X_T^\theta, Z_T)$, we shorten notation by setting $\phi(x,z) := \varphi(h(x,z))$. Hence, if the speculator holds ρ units of such an option, expected terminal wealth becomes

$$E^{x,z}\left[\int_0^T \Pi(t, X_t^\theta, Z_t)\theta(t)\,dt + \rho I \phi(X_T^\theta, Z_T)\right]. \tag{6.10}$$

Here, I denotes the exercise policy, which we assume to be $= 1$ if the option is exercised, and $= 0$ if not.

6.3.2 Admissible Strategies

The choice variables for the speculator are the trading strategy $\theta(t)$ and the exercise policy I. Note that due to our assumption of absolute continuity, the demand process X_t is automatically *predictable*. This is in line with the intuition that the decision on what assets to hold over any period of time must be based upon the information available at the *beginning* of that period. It is also important for technical reasons, see for example Harrison and Pliska (1981) for a detailed elaboration on this issue. On the other hand, since the decision whether to exercise the option or not is made at time T, it can be contingent on all information available to the speculator at that time.

Formally, an admissible policy is a pair (θ, I), where I is an $\mathcal{F}_T$-measurable random variable taking values in $\{0,1\}$, and $\theta(t)$ is a process adapted to $(\mathcal{F}_t)_{t\geq 0}$ such that

$$V^{\theta,I}(t,x,z) := E_t^{x,z}\left[\int_t^T \Pi(\tau, X_\tau^\theta, Z_\tau)\theta(\tau)\,d\tau + \rho I \phi(X_T^\theta, Z_T)\right]$$

is well-defined for every $x, z \in \mathbb{R}$ and $t \in [0,T]$. We can interpret $V^{\theta,I}(t,x,z)$ as the *gain* that is made starting at $Z_t = z$ and $X_t^\theta = x$ at time t and following trading strategy θ and exercise policy I thereafter. The speculator's problem is thus to find an admissible policy (θ, I) which maximises $V^{\theta,I}(t,x,z)$. As usual in the theory of stochastic control, we define the *value function* of the problem as

$$V^*(t,x,z) := \sup_{(\theta,I)} V^{\theta,I}(t,x,z). \tag{6.11}$$

We are now ready to state the main result of this section.

Theorem 6.3.1 (Optimal Strategy) *There exists a non-degenerate solution to the speculator's problem if and only if the equilibrium pricing rule satisfies* $H(t,x,z) = E_t^z[\ h(x,Z_T)\ ]$. *In this case, there exist functions* $X^*(t,z)$ *and* $I^*(z)$ *such that the optimum is attained by exercise policy* $I^* = I^*(Z_T)$, *and any trading strategy of the form*

$$\theta(t) := \alpha(t)\left(X^*(t,Z_t) - X_t^\theta\right), \tag{6.12}$$

where $\alpha(t)$ *is any deterministic function with* $\alpha(t) \longrightarrow \infty$ *as* $t \longrightarrow T$ *fast enough to force the demand process* X_t^θ *to converge to* $X^*(t,Z_t)$ *as* $t \longrightarrow T$.

Note that the necessary condition, $H(t,x,z) = E_t^z[\ h(x,Z_T)\ ]$ can be interpreted as a no-arbitrage condition, since it requires the equilibrium exchange rate process in the absence of the speculator to be a martingale. In the examples considered below, this condition is satisfied by construction.

We give a brief outline of the proof of this theorem in the following section. A detailed exposition can be found in Section 4.3. For general reference on the theory of controlled diffusion processes see Fleming and Soner (1993) or Krylov (1980).

6.3.3 Proof of Theorem 6.3.1

The proof of Theorem 6.3.1 is essentially an application of Theorem 4.3.1 in Chapter 4. As discussed there, the value function $V^*(t,x,z)$ can be characterised as the solution to the *Bellman equation* subject to appropriate boundary conditions. In the case considered here, the Bellman equation takes the form

$$\sup_\theta \left\{ V_t + \frac{1}{2}\sigma_Z^2 V_{zz} + (V_x + \Pi)\,\theta \right\} \equiv 0. \tag{6.13}$$

Since this equation is *linear* in θ, in order for the supremum to be finite the coefficient of θ must be zero. We can thus split up the Bellman Equation into the following two separate equations:

$$V_t + \frac{1}{2}\sigma_Z^2 V_{zz} \equiv 0, \tag{6.14}$$

$$V_x + \Pi \equiv 0. \tag{6.15}$$

Equation (6.14) is just the heat equation associated with two-dimensional Brownian motion. The Feynman-Kac representation theorem (see Karatzas and Shreve (1988, Theorem 4.4.2)) allows us to write its solution (if it exists) in the following form:

$$V(t,x,z) = E_t^z\,[\,V(T,x,Z_T)\,]\,.$$

Equation (6.15) has a more economic interpretation. It states that the marginal gain $V_y(t,x,z)$ from holding an extra unit of the risky asset at time t, must equal the purchasing cost, $H(t,x,z)$, *minus* the expected marginal benefit, $F(t,z)$. For a detailed elaboration on the intuition behind this, see Section 4.3.4.

OPTIMAL STRATEGY

Following the arguments developed in Section 4.3.5, the boundary condition for the Bellman equation will help us pinpoint the optimal strategy. Suppose the value function $V^*(t,x,z)$ exists and is sufficiently smooth. We know then that is must solve the Bellman equations (6.14) and (6.15). For any given admissible strategy $\theta(t)$, Itô's Lemma applied to the process $V^*(t, X_t^\theta, Z_t)$, together with (6.14) and (6.15) allows us to write

$$V^*(0,x,z) = V^{\theta,I}(0,x,z) + E^z\Big[\underbrace{V^*(T, X_T^\theta, Z_T) - \rho I \phi(X_T^\theta, Z_T)}_{=:\, q^*(X_T^\theta, Z_T; I)}\Big]. \tag{6.16}$$

The above equation has the following interpretation; the quantity $E^z[\, q^*(X_T^\theta, Z_T; I)\,]$ measures the difference between the value $V^{\theta,I}(0,x,z)$ of any given policy pair (θ, I) and the maximal achievable value $V^*(0,x,z)$. The function $q^*(X_T^\theta, Z_T; I)$ can thus be identified as the *loss* incurred by the speculator by deviating from the optimal strategy. In particular we deduce that $E^z[\, q^*(X_T^\theta, Z_T; I)\,]$ must be non-negative, and zero at the optimum. We will reverse this argument to construct the optimal policy and value function. More specifically, we will construct a function $q^*(x,z;I)$ and "optimality conditions" $x^*(z)$ and $I^*(z)$ in such a way that $q^*(x,z;I) \geq 0$ and $= 0$ if and only if $x = x^*(z)$ and $I = I^*(z)$. Based on this we use (6.14) and the Feynman-Kac representation to construct a candidate $V^*(t,x,z)$ for the value function which satisfies (6.13). From this and the above arguments we can conclude that an exercise policy $I^*(Z_T)$ and a trading strategy which ensures that $X_T^\theta = x^*(Z_T)$ will attain the optimum, and that $V^*(t,s,y)$ is indeed the value function of the speculator's problem. The formal proofs for all these statements can be found in Chapter 4. We interpret $x^*(s)$ as the optimal "target position" towards which the speculator should trade as $t \longrightarrow T$, conditional on $Z_T = z$. This is similar to the arguments presented in Back (1992) and Back and Pedersen (1996).

As shown in Chapter 4, evaluating the Bellman equation (6.15) at $t = T$ implies that the loss function $q^*(x, z; I)$ must be of the form

$$q^*(x, z; I) = \int_x^{x^*(z)} \pi(\xi, z)\, d\xi + \rho I^*(z)\phi(x^*(z), z) - \rho I\phi(x, z), \tag{6.17}$$

for suitably chosen $x^*(z)$ and $I^*(z)$. An intuitive interpretation of this equation is given in Chapter 4. By construction, the loss function defined in (6.17) satisfies the desired "boundary conditions" $q^*(x^*(z), z; I^*(z)) = 0$. Our remaining task is hence to determine $x^*(z)$ and $I^*(z)$ in such a way that for each z the mapping $(x, I) \mapsto q^*(x, z; I)$ has a global minimum at $x = x^*(z)$ and $I = I^*(z)$. Following the arguments developed in Chapter 4, we first keep the option exercise decision fixed and construct optimal targets $x_1^*(z)$ and $x_0^*(z)$ conditional on the option being exercised or not, respectively. We can then establish the optimal exercise decision $I^*(z)$ endogenously. As shown in Chapter 4, the conditional target positions $x_0^*(z)$ and $x_1^*(z)$ can be characterised via the first-order conditions

$$f(z) = h(x_0^*(z), z), \quad \text{and} \tag{6.18}$$

$$f(z) = h(x_1^*(z), z) - \rho\phi_x(x_1^*(z), z). \tag{6.19}$$

As shown in Chapter 4, it is then optimal for the speculator to exercise the option if and only if

$$\int_{x_1^*(z)}^{x_0^*(z)} \pi(\xi, z)\, d\xi \leq \rho\phi(x_1^*(z), z). \tag{6.20}$$

The economic interpretation of (6.20) is as follows. Recall that $x_0^*(z)$ is the optimal target position conditional on no exercise. Hence the left-hand side of (6.20) is the loss incurred by trading towards $x_1^*(z)$ instead of $x_0^*(z)$. The right-hand side however represents the gain made from the option's pay-off. Hence (6.20) states that the speculator should exercise the option if and only if the loss incurred by deviating from the optimal position is outweighed by the option's pay-off. Since by (6.20), optimality of exercise depends only on z, we can define the optimal exercise indicator $I^*(z)$ in the obvious way. Intuitively, the unconditional target $x^*(z)$ should then equal $x_1^*(z)$ whenever exercising the option is optimal, and $x_0^*(z)$ otherwise. Therefore we define

$$x^*(z) := I^*(z)x_1^*(z) + (1 - I^*(z))x_0^*(z).$$

The loss function $q^*(x, z; I)$ defined via (6.17) can then indeed be shown to have a global minimum at $x = x^*(z)$ and $I = I^*(z)$, see Proposition 4.A.3. From the arguments outlined in the beginning of this section it should now be intuitively clear that any strategy $\theta(t)$ which ensures that $X_t^\theta \longrightarrow x^*(Z_T)$ as $t \longrightarrow T$ is a candidate for the optimum. We hence set $X^*(t, z) := E_t^z[\, x^*(Z_T)\,]$ and define

$$\theta(t) := \alpha(t)\left(X^*(t, Z_t) - X_t^\theta\right),$$

where $\alpha(t)$ is a deterministic function with $\alpha(t) \longrightarrow \infty$ as $t \longrightarrow T$ fast enough to make X_t^θ converge to $x^*(Z_T)$ as $t \longrightarrow T$. A rigorous proof for the optimality of $\theta(t)$ is given in the appendix to Chapter 4.

6.4 Equilibrium

In this section, we will analyse the structure of the equilibria that arise in two quite different scenarios. First, we consider the case in which the speculator does not exploit her market power but instead simply *hedges* the option position, as assumed in Wiseman (1996) and Breuer (1997). Second, we study the case in which the speculator takes advantage of the fact that the option provides an opportunity to benefit from market manipulation. This case will be the main focus of our attention. Unsurprisingly, in the former case the presence of the option indeed does dampen the fluctuations caused by fundamental news, thus achieving the desired stabilisation of exchange rates. In the latter, speculative case however, the effect is reversed:

While in the absence of the option, the speculator has no incentive to manipulate markets and may even act in the central bank's favour, the introduction of the option provides a way to benefit from manipulation and thus creates an *additional* vehicle for speculative attacks on the domestic currency.

6.4.1 Hedging Equilibrium

As a benchmark, we consider the case in which the speculator does not fully exploit her market power but simply implements a strategy designed to hedge the long option position. Note that, since the speculator's trades affect underlying equilibrium exchange rates, solving the hedging problem is far more involved then in the classic Black-Scholes model. According to the classic theory, the strategy needed to replicate an option's pay-off can be characterised as the solution to a *linear* partial differential equation. However, if the underlying market is imperfectly elastic, the activity of option hedging or replication creates a feedback effect which affects the underlying equilibrium price process. See Chapter 2 for a detailed elaboration on this issue.

The hedging problem in finitely elastic markets was solved by Jarrow (1994) in discrete time, and later by Frey (1996) in continuous time. Frey considers a reduced-form equilibrium specification, in which prices are given by a "reaction function", depending on the large trader's position and some exogenous stochastic factor. The pricing rule derived in (6.6) is a special case of such a reaction function, where the role of the stochastic factor is played by the news process Z_t. Frey shows that, if the option's pay-off function is sufficiently regular, there exists a unique replicating strategy for the large trader. This strategy is characterised by a non-linear partial differential equation, which implies in particular that the cost of replication is not additive, i.e. replicating two options does not necessarily cost twice as much as replicating one.

However, the one feature of replicating strategies which is essential for the analysis here remains valid even in this more complex situation: Strategies designed to *replicate* convex pay-off patterns exhibit *positive* price elasticity, they require buying in a rising market and selling when prices fall. See Frey (1996) for details. Conversely, a strategy designed to *hedge* the risk of an option position exhibits *negative* price elasticity. As a consequence, if the speculator hedges her position, the resulting equilibrium exchange rate, P_t, is less volatile than its fundamental level, P_t^F. In other words, the speculator's hedging activity indeed serves to stabilise exchange rates. In particular, when fundamental news are bad, and information traders would sell domestic currency thus pushing rates outside the central bank's target zone, the speculator's hedging strategy requires *selling* foreign currency thus supporting the domestic currency.

6.4.2 Speculative Equilibrium

We now turn to the case which is the main focus of our attention. We assume that rather than simply hedging the option's pay-off, the speculator acts strategically such as to maximise expected terminal wealth. We have seen in Section 6.3 that the speculator's optimal strategy can be characterised in terms of a "target position" $x^*(Z_T)$ towards which she should trade, conditional on the news process Z_T. The target position is characterised by the first-order condition

$$f(Z_T) = h(x^*(Z_T), Z_T) - \rho \underbrace{I^*(Z_T)\phi_x(x^*(Z_T), Z_T)}_{=: \Gamma_T}. \tag{6.21}$$

In equilibrium, we have $X_T = x^*(Z_T)$, so that the first term on the right-hand side of the above equation is simply the equilibrium exchange rate, $h(x^*(Z_T), Z_T) = P_T$. Recall that the value of the speculator's foreign currency position is determined by the price at which this position can be "unwound" in the market. Formally, this can be written as $f(Z_T) = P_T^0$, so that the first-order condition takes the form

$$P_T = P_T^0 + \rho\Gamma_T. \tag{6.22}$$

The supply pressure resulting from unwinding a large position will cause exchange rates to depreciate, thus neutralizing any potential gain the speculator could have derived from manipulating exchange rates. Thus, in the absence of the option ($\rho = 0$), the speculator has no incentive to trade at all. More precisely, if $\rho = 0$, condition (6.22) simply requires $P_T = P_T^0$, which implies $x^*(Z_T) \equiv 0$. Hence we can interpret the term Γ_T as the additional price pressure resulting from the speculator's incentive to manipulate exchange rates when she holds an option.

Proposition 6.4.1 (Equilibrium) *In equilibrium, exchange rates satisfy*

$$P_T = P_T^0 + \rho\Gamma_T.$$

In particular, in the absence of options ($\rho = 0$), the speculator does not trade at all.

Note that, since $\Gamma_T > 0$, the above equilibrium equation implies that $P_T^0 \geq P_T^F$. In other words, the equilibrium exchange rate in the absence of speculation will always be below the fundamental level. However, if the speculator holds options ($\rho > 0$), situations are likely to occur in which

$$P_T^0 < P_T^F < P_T^0 + \rho\Gamma_T.$$

In other words, while in the absence of the option, exchange rates would not exceed their fundamental level, P_T^F, the existence of the option creates an incentive for the speculator to push exchange rates beyond P_T^F. The intuition behind this result is as follows. From the analysis in Section 6.3 we know that the speculator's incentive to manipulate exchange rates is stronger the more the option is likely to end up in the money. In other words, the speculator responds to the arrival of bad news regarding the domestic currency by manipulating *against* it, i.e. by increasing her position in the foreign currency. This effect is *entirely* due to the existence of the option! Rather than *protecting* the domestic currency from speculative attacks, options issued by the central banks in fact *create* an additional vehicle for such an attack.

6.4.3 An Example with Linear Demand

In this section we will analyse a specific example in which we assume for simplicity that the information traders' demand function is *linear*. More precisely, recall that the "fundamental" exchange rate is given by $P_t^F = H^F(t, Z_t)$. We assume that information traders buy foreign currency if the exchange rate is below its fundamental level and sell if it is above. Formally, we assume

$$D(t,z,p) = \frac{1}{\kappa}\left(H^F(t,z) - p\right),$$

which is obviously compatible with the definition of $H^F(t,z)$. To give an explicit example, we may assume that $P_t^F = H^F(t, Z_t)$ is simply an exponential martingale,

$$H^F(t, Z_t) = \mathcal{E}_t(Z_t) := \exp\left(Z_t - \frac{1}{2}\int_0^t \sigma_Z^2(s)ds\right).$$

Moreover, we assume the central bank's exchange rate policy to be of a very simple form. Suppose the central bank's objective simply is to ensure that the exchange rate at the terminal date, P_T, does not exceed some *intervention level* P_T^I. To simplify matters even further, we assume that at the terminal date T, the central bank does nothing if the exchange rate is in the target zone (below P_T^I), but sells all foreign currency reserves if it is not. In other words, the central bank's policy function at time T can be written as

$$c(z,p) = \begin{cases} 0 & \text{if } p < P^I \\ -\underline{C} & \text{if } p \geq P_T^I \end{cases} .$$

Strictly speaking, this example violates the differentiability assumptions made in Section 6.2. However, the arguments developed below can be obtained by approximating the central bank's policy function by sufficiently smooth functions and passing to the limit. Figure 6.1 shows the pricing rule at the terminal date arising from this specification. It is straight-forward to see that the speculator's first-order conditions in

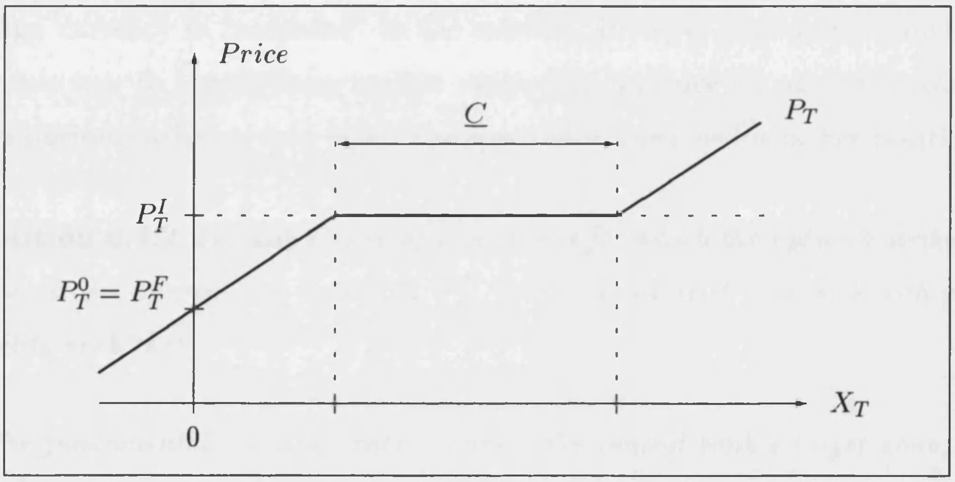

Figure 6.1: Pricing Rule

this case are given by

$$\begin{aligned} P_T^0 &= h(x_0^*, Z_T) \qquad \text{for } I = 0, \\ P_T^0 &= h(x_1^*, Z_T) - \rho\kappa \qquad \text{for } I = 1. \end{aligned}$$

From this we can deduce that $x_0^* = 0$, and

$$\begin{aligned} &\text{if } P_T^0 + \rho\kappa < P_T^I, &&\text{then } x_1^* = \rho, \text{and} \\ &\text{if } P_T^0 + \rho\kappa \geq P_T^I, &&\text{then } x_1^* = \underline{C} + \rho. \end{aligned}$$

In other words, if there is no option, then the speculator will have no incentive to trade, and the terminal exchange rate will be $h(0, Z_T)$. However, if the speculator *does* hold an option, it might be optimal to hold foreign currency. The intuition behind this is simple. Without the option, the speculator cannot realise any gain from manipulating the exchange rate, since any potential gain will be neutralised by the price decline arising from the supply pressure when the speculator's position in foreign currency is "unloaded" in the market. However, the option provides an alternative way to benefit from market manipulation, since its pay-off is settled at the equilibrium exchange rate *before* the speculator unwinds his or her position.

Proposition 6.4.2 *For any choice of parameters for which the option's strike price, K, is below the intervention threshold, P_T^I, there always exist outcomes with positive probability such that*

(a) the fundamental exchange rate is within the central bank's target zone, $P_T^F < P_T^I$, so that in the absence of speculation, intervention is not necessary,

(b) even when the speculator is active, without the option ($\rho = 0$) there is no incentive for manipultation, so that the equilibrium exchange rate equals the fundamental exchange rate, $P_T = P_T^F < P_T^I$, and no intervention is necessary, but

(c) if the speculator does hold an option ($\rho > 0$), the equilibrium exchange rate will be pushed outside the target zone, $P_T > P_T^I$, even though the central bank exhausts all its foreign currency reserves to support the domestic currency.

PROOF: We focus on those values of Z_T for which $P_T^0 = P_T^F < P_T^I$, but $P_T^0 + \kappa\rho > P_T^I$, i.e.

$$h(0, Z_T) < P_T^I < h(0, Z_T) + \kappa\rho.$$

In other words, the fundamental exchange rate is within the target zone, but the marginal benefit from the option carries the exchange rate outside this region. An example of such a situation is depicted in Figure 6.2. Obviously, since $\kappa\rho > 0$, the set of such Z_T is non-empty. In this case, the optimal trading targets for the speculator are $x_0^* = 0$ and $x_1^* = \underline{C} + \rho$. Note that if the speculator trades towards $X_T = x_1^*$, the equilibrium exchange rate would be $P_T^0 + \kappa\rho > P_T^I$. That is, although the central bank in this case sells off all its foreign currency reserves, it does not succeed in keeping the exchange rate within the target zone.

Recall from Section 6.3.3 that x_1^* is the speculator's optimal trading target only in the case when it is optimal to exercise the option. It hence remains to show that there exists a set of Z_T with positive probability for which it is optimal to exercise. From Section 6.3.3 we also know that it is optimal to exercise the option if and only if

$$\int_{x_1^*}^{x_0^*} \left(P_T^0 - h(\xi, Z_T)\right) d\xi \leq \rho \left(h(x_1^*, Z_T) - K\right). \tag{6.23}$$

It is immediate to see that the integral on the left-hand side of inequality (6.23) corresponds to the shaded area in Figure 6.2. A straight-forward calculation gives

$$\int_{x_1^*}^{x_0^*} \left(P_T^0 - h(\xi, Z_T)\right) d\xi = \underline{C} \left(P_T^I - P_T^0\right) + \frac{1}{2}\kappa\rho^2.$$

The right-hand side of (6.23) is obviously equal to $\rho(P_T^0 - K) + \kappa\rho^2$. We can hence conclude that sufficient for optimality of exercise is

$$P_T^0 > \frac{\rho K + \underline{C} P_T^I}{\rho + \underline{C}}. \tag{6.24}$$

Note that the right-hand side of (6.24) is simply a weighted average of K and P_T^I. Since by assumption, $K < P_T^I$, it is obviously always possible to find a range of Z_T for which

$$P_T^0 + \kappa\rho > P_T^I > P_T^0 > \frac{\rho K + \underline{C} P_T^I}{\rho + \underline{C}}.$$

Since by construction, the distribution of P_T^0 has full support on $\mathbb{R}_+$, it follows that the set of Z_T for which the above inequality is satisfied has positive probability.

$\square$

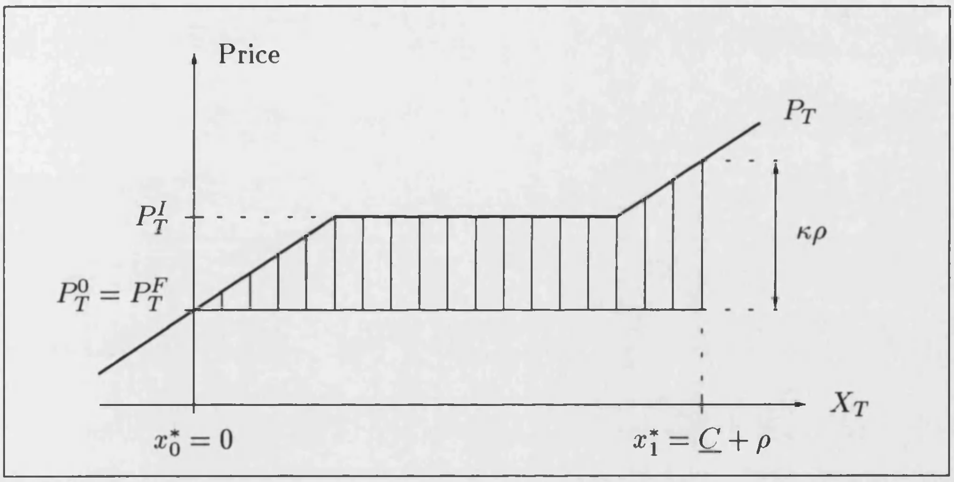

Figure 6.2: Optimal Exercise Decision

6.5 Conclusion

We have studied the theoretical implications of the proposal by Wiseman (1996) and Breuer (1997), according to which central banks should issue options in order to stabilise foreign exchange rates. We have shown that the desired stabilisation effect relies on the assumption that the buyers of such options content themselves

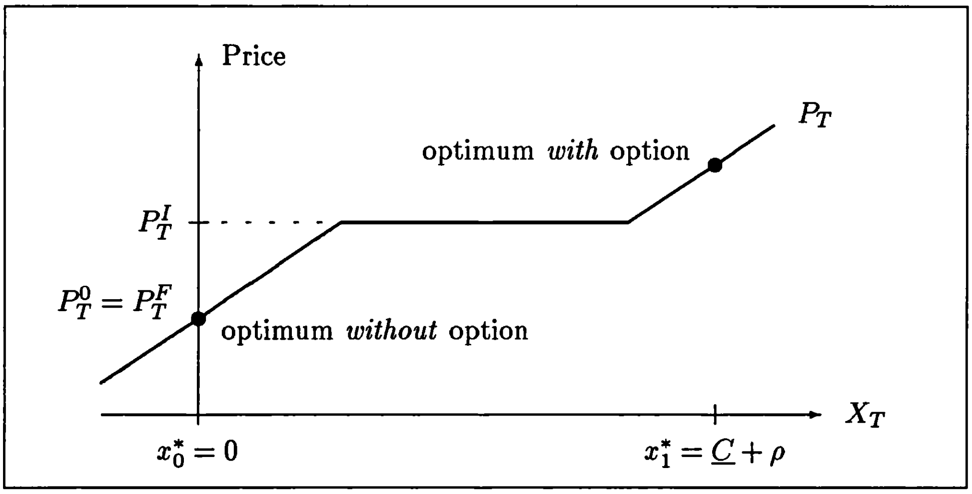

Figure 6.3: Equilibrium

with hedging their positions, rather than exploit strategically the levaraged position in the market provided by the option. Our results confirm that, if this assumption is true, the implementation of the corresponding hedging strategies would indeed reduce exchange rate volatility.

More importantly however, we have shown that if the assumption fails to hold, i.e. if speculators strategically exploit their market power, the existence of such options might in fact achieve the opposite of the desired stabilisation effect. In the absence of options, there is no incentive for speculators to manipulate exchange rates, and fluctuations merely reflect changes in the economic fundamentals. By issuing options, the central bank in fact *creates* an additional incentive for market manipulation. Upon arrival of "bad news", when rational traders adjust their positions thus causing the domestic currency to depreciate, the option creates an incentive for the speculator to squeeze exchange rates even further. Scenarios can arise in which without speculation, the central bank's foreign currency reserves would have been sufficient to support the domestic currency, while they cannot sustain the additional pressure arising from the speculator's manipulation. In other words, instead of *protecting* against speculative attacks, options issued by the central bank in fact create an additional vehicle for such an attack.

While the mechanism proposed by Wiseman (1996) and Breuer (1997) can be shown to achieve the desired exchange rate stabilisation in *some* circumstances, it possesses an even greater potential for destabilisation in others. We therefore believe that despite its potential benefits, the risk inherent in the proposed meachnism is intolerably high.

An obvious question ensuing from the discussion in this chapter is whether there are alternative ways in which options could be used as exchange rate policy instruments without creating an unnecessary risk exposure. More specifically, future research should aim at determining *endogenously* the optimal strategy for the central bank, given their objective.

Chapter 7

Conclusions

In this thesis, we have studied the manner in which different types of trading behaviour affects the nature of equilibrium prices when markets are imperfectly elastic.

We have seen that the additional demand generated by dynamic hedging strategies creates feedback effects which alter the volatility structure of the underlying price process, thus destabilising prices. These results are in line with the empirical evidence. Due to this effect, simple Black-Scholes strategies are no longer sufficient to perfectly replicate an option's pay-offs. However, we were able to demonstrate that such strategies can still be used to "super-replicate" the option's pay-off, even if their implementation causes a violation of the very assumptions they are derived from.

Further, we have investigated the manner in which technical trading affects equilibrium exchange rates. Technical trading leads to feedback effects similar to those inherent in the implementation of dynamic hedging strategies. As a consequence, technical trading causes exchange rate fluctuations of a magnitude far beyond the level justified by changes in the economic fundamentals. In other words, technical trading leads to the emergence of irrational price bubbles. Our results demonstrate that, while being ex-ante irrational, their very existence can make technical trading rules ex-post profitable. In other words, technical trading can be seen as a kind of "self-fulfilling prophecy". These results are in line with the empirical evidence.

We have analysed the interplay between options markets and the markets for the underlying asset in the presence of asymmetric information. An option, held by a large, informed trader, creates an incentive for this trader to manipulate prices in order to push the option into the money. In other words, the option induces the informed trader to create a manipulative price bubble on the underlying market at the expense of uninformed traders. As a consequence, equilibria can only exist if the potential gain from manipulation is incorporated into the option's price. The resulting equilibrium option price schedule gives rise to the "smile pattern" of implied volatility. The particular shape of the smile pattern predicted by our model is in line with the empirical observations made *after* the 1987 market crash. In view of our results, the post-crash shape of the smile could be interpreted as an indication that increased awareness amongst market participants and regulators has lead to a better alignment of option prices with the markets for the underlying.

Finally, we have analysed the implications of imperfect elasticity for the use of options as central bank exchange rate policy instrument. Following up on recent suggestions for central banks to *write* call options in order to induce stabilising "delta hedging", we demonstrated the potential danger inherent in such a strategy. While an option issued by the central bank would indeed induce highly risk-averse market participants to reduce exchange rate volatility by hedging their option position, this is certainly not true if the option is instead used as a speculative instrument. A speculator, having bought the option from the central bank, would have a strong incentive to manipulate exchange rates in order to push the option deeper into the money. Thus, rather then *protecting* against speculative attacks on the domestic currency, the option instead *creates* an additional vehicle for such an attack.

Bibliography

ALLEN, H., and M. TAYLOR (1992): "The Use of Technical Analysis in the Foreign Exchange Market," *Journal of International Money and Finance*, 11(3), 304–314.

ALMEKINDERS, G. (1995): *Foreign Exchange Intervention.* Edwar Elgar Publishing Company, Brookfield, Vermont.

BACK, K. (1992): "Insider Trading in Continuous Time," *Review of Financial Studies*, 5(3), 387–409.

BACK, K. (1993): "Asymmetric Information and Options," *Review of Financial Studies*, 6(3), 435–472.

BACK, K., and H. PEDERSEN (1996): "Long-Lived Information and Intraday Patterns," OLIN School of Business Working Paper, Washington University at St Louis.

BALL, C., and A. ROMA (1994): "Stochastic Volatility Option Pricing," *Journal of Financial and Quantitative Analysis*, 29(4), 589-607.

BLACK, F., and M. SCHOLES (1973): "The Pricing of Options and other Corporate Liabilities," *Journal of Political Economy*, 81(3), 637-654.

BRENNAN, M., and E. SCHWARTZ (1989): "Portfolio Insurance and Financial Market Equilibrium," *Journal of Business*, 62(4), 455–472.

BREUER, P. (1997): "Central Bank Participation in Currency Options Markets," Working Paper, Brown University.

BROCK, W., J. LAKONISHOK, and B. LEBARON (1992): "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns," *Journal of Finance*, 47(4), 1731–1764.

CHRISTENSEN, B. J., and N. PRABALA (1998): "On the Dynamics and Information Content of Implied Volatility: A Bivariate Time Series Perspective," *Journal of Financial Economics*, Forthcomming.

COX, J., S. ROSS, and M. RUBINSTEIN (1979): "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7(3), 229–263.

CURCIO, R., and C. GOODHART (1992): "When Support/Resistance Levels are Broken, Can Profits be Made? Evidence from the Foreign Exchange Market," Discussion Paper 142, Financial Markets Group, London School of Economics.

DAY, R., and W. HUANG (1990): "Bulls, Bears, and Market Sheep," *Journal of Economic Behaviour and Organization*, 14, 299–329.

DE LONG, J., A. SHLEIFER, L. SUMMERS, and R. WALDMANN (1990b): "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98(4), 703–738.

DE LONG, J., A. SHLEIFER, L. SUMMERS, and R. WALDMANN (1990a): "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance*, 45(2), 381–395.

DUFFIE, D., and P. PROTTER (1992): "From Discrete to Continuous Time Finance: Weak Convergence of the Financial Gain Process," *Mathematical Finance*, 2(1), 1–15.

EDISON, H. (1993): "The Effectiveness of Central Bank Intervention: A Survey of the Literature after 1982," Special Papers in International Economics, Princeton University.

EL KAROUI, N., M. JEANBLANC-PICQUÉ, and S. SHREVE (1998): "Robustness of the Black and Scholes Formula," *Mathematical Finance*, 8(2), 93–126.

FLEMING, W., and H. SONER (1993): *Controlled Markov Processes and Viscosity Solutions.* Springer-Verlag, New York.

FÖLLMER, H., and M. SCHWEIZER (1990): "Hedging of Contingent-Claims under Incomplete Information," in *Applied Stochastic Analysis*, pp. 205–223. Gordon & Breach, London.

FÖLLMER, H., and M. SCHWEIZER (1993): "A Microeconomic Approach to Diffusion Models for Stock Prices," *Mathematical Finance*, 3(1), 1–23.

FREY, R. (1996): "Perfect Option Hedging for a Large Trader," *Finance and Stochastics*, 2, 115–142.

FREY, R., and A. STREMME (1994): "Portfolio Insurance and Volatility," Discussion Paper 188, Financial Markets Group, London School of Economics.

FREY, R., and A. STREMME (1997): "Market Volatility and Feedback Effects from Dynamic Hedging," *Mathematical Finance*, 7(4), 351–374.

GENNOTTE, G., and H. LELAND (1990): "Market Liquidity, Hedging and Crashes," *American Economic Review*, 80(5), 999–1021.

GLOSTEN, L., and P. MILGROM (1985): "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.

GOODHART, C. (1988): "The Foreign Exchange Market: A Random Walk with a Dragging Anchor," *Economica*, 55, 437–460.

GOURIEROUX, C., J. LAFFONT, and A. MONFORT (1982): "Rational Expectations in Dynamic Linear Models: Analysis of the Solutions," *Econometrica*, 50, 409–425.

GROSSMAN, S. (1988): "An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies," *Journal of Business*, 61(3), 275–298.

GROSSMAN, S., and J. STIGLITZ (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70(3), 393–408.

GROUP OF TEN (1993): "International Capital Movement and Foreign Exchange Markets: A Report to the Ministers and Governors by the Group of Deputies," Rome.

HARRISON, J., and S. PLISKA (1981): "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and Applications*, 11, 215–260.

HE, H. (1990): "Convergence from Discrete to Continuous Time Contingent Claim Prices," *Review of Financial Studies*, 3(4), 523–546.

HULL, J., and A. WHITE (1987): "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42(2), 281–300.

JACOD, J., and A. SHIRYAEV (1987): *Limit Theorems for Stochastic Processes.* Springer Verlag, Berlin.

JARROW, R. (1992): "Market Manipulation, Bubbles, Corners and Short Squeezes," *Journal of Financial and Quantitative Analysis*, 27(3), 311–336.

JARROW, R. (1994): "Derivative Securities Markets, Market Manipulation and Option Pricing Theory," *Journal of Financial and Quantitative Analysis*, 29(2), 241–261.

KALLIANPUR, G., and R. KARANDIKAR (1985): "White Noise Calculus and Nonlinear Filtering Theory," *Annals of Probability*, 13(4), 1033–1107.

KARATZAS, I., and S. SHREVE (1988): *Brownian Motion and Stochastic Calculus.* Springer Verlag, Berlin.

KRYLOV, N. (1980): *Controlled Diffusion Processes.* Springer-Verlag, New York.

KYLE, A. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53(6), 1315–1335.

LELAND, H. (1980): "Who Should Buy Portfolio Insurance?," *Journal of Finance*, 35(2), 581–594.

LEVICH, R., and L. THOMAS (1993): "The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach," *Journal of International Money and Finance*, 12(5), 451–474.

LUX, T. (1995): "Herd Behaviour, Bubbles and Crashes," *Economic Journal*, 105, 881–896.

MALZ, A. (1995): "Currency Options and Exchange Rates: A Case Study of the U.S. Dollar in March 1995," mimeo, Federal Reserve Bank of New York.

MÉMIN, J., and L. SLOMINSKI (1991): "Condition UT et Stabilité en Loi des Solutions d'Equations Différentielles Stochastiques," in *Séminaire de Probabilité XXV*, vol. 1485 of *Lecture Notes in Mathematics*, pp. 162–177. Springer Verlag, Berlin.

NEELY, C. (1997): "Technical Analysis in the Foreign Exchange Market: A Layman's Guide," *The Federal Reserve Bank of St. Louis Review*, September/October, 23–38.

PLATEN, E., and M. SCHWEIZER (1994): "On Smile and Skewness," Statistics Research Report 027-94, Australian National University, Canberra.

PLATEN, E., and M. SCHWEIZER (1998): "On Feedback Effects from Hedging Drivatives," *Mathematical Finance*, 8(1), 67–84.

SCHEUTZOW, M. (1984): "Qualitative Behaviour of Stochastic Delay Equations with a Bounded Memory," *Stochastics*, 12, 41–80.

SCHÖNBUCHER, P. (1993): "Option Pricing and Hedging in Finitely Liquid Markets," Master's thesis, Exeter College, Oxford.

SWEENEY, R. (1986): "Beating the Foreign Exchange Market," *Journal of Finance*, 41(1), 163–182.

TAYLOR, C. (1995): "Options and Currency Interventions," Working Paper, Centre for the Study of Financial Innovation, London.

WISEMAN, J. (1996): "Mechanisms for Central Banking," Special Paper 84, Financial Markets Group, London School of Economics.